

Durham E-Theses

A Relational-Representational Theory of Perceptual Intentionality: Immanent Content, Emergent Representation

GAURAV SATEESH KUDTARKAR

How to cite:

KUDTARKAR, GAURAV SATEESH (2025) A Relational-Representational Theory of Perceptual Intentionality: Immanent Content, Emergent Representation. Doctoral thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/16330/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Gaurav Sateesh Kudtarkar

A Relational-Representational Theory of Perceptual Intentionality: Immanent Content, Emergent Representation

Abstract

Representationalism about perception is the representational theory of perceptual intentionality. It interprets perceptual states as perceptual *representations* and accounts for their intentionality or external directedness in terms of content, which is the essential representational property. There are two broad types of representationalist theories – reductive-naturalistic representationalism and nonreductive representationalism. Whereas most reductive-naturalistic theories are pitched to solve the content determinacy problem, most nonreductive views are posited to solve the problem of non-veridical perception. In the first part of my thesis, I argue that the central problem of perception is a specific type of content determinacy problem called the ‘distality problem’: how is content as of distal entities rather than proximal intermediaries of the causal chain from the entities to subjects?

In the second part of my thesis, I critically evaluate reductive-naturalistic and nonreductive representationalist theories. Regarding the former, I argue that a crucial ingredient of their reductive project – psychological mechanisms – is not necessary for representation. Thus, their reductive enterprise fails, and they cannot account for any content determinacy problem, including the distality problem. Regarding nonreductive representationalism, I argue that, while their performance on solving the problem of non-veridical perception is mixed, they cannot account for the external directedness and distality of content.

In the third and final part, I develop a novel version of representationalism, which I call *Pluri-relational Immanent Emergent Representationalism (PRIMER)*. In developing this view, I draw upon the views of Aristotle and the Indian classical philosophy of Nyāya regarding universals and perception. A major feature of PRIMER is its claim that content is partly constituted by the co-instantiation in the subject of the same concrete universals that are also instantiated in the target; but the manner of instantiation is different. Content is, therefore, an immanent property of the subject. Further, I argue that content is a strongly emergent property because it possesses novel action-oriented causal powers. Finally, I argue that immanent and emergent content can account for, among other things, the intentional directedness of perceptual representations, the distality of content and the problem of non-veridical perception.

**A Relational-Representational Theory of
Perceptual Intentionality:
Immanent Content, Emergent Representation**

**Thesis Submitted for the Degree of
Doctor of Philosophy**

**By
Gaurav Sateesh Kudtarkar**



**Department of Philosophy
University of Durham
2025**

Contents

Statement of Copyright	i
Acknowledgements	ii
0. Introduction	1
0.1. Delimiting the Boundaries	1
0.2. Representationalism	2
0.3. Research Problem/Question	4
0.4. Solutions	5
0.5. Significance	7
0.6. Methodology	8
0.7. Structure of the Thesis	10
1. The Representational Theory of Intentionality	16
1.1. Intentionality	16
1.1.1. Traditional Theories of Intentionality	17
1.1.2. Contemporary Theories of Intentionality	18
1.2. Notions of Representation	18
1.2.1. Minimal Representation	19
1.2.2. Thin Representation	21
1.2.3. Vehicular Representation	21
1.2.4. Syntactically Structured Representation	24
1.3. Representationalism	25
1.3.1. Arguments for Representationalism	26
1.3.2. An Argument for the Vehicular Notion of Representation	26
1.3.3. Types of Representationalism	27
1.4. Metaphysical Primer	29
1.5. Conclusion	36
2. Representationalism About Perception: Key Terms and Debates	37
2.1. Representationalism about Perception	37
2.1.1. Representationalism and Intentionalism	39
2.1.2. Representationalism and the Phenomenal Intentionality Theory	40
2.2. Representationalism – Analysantia and Theoretical Components	41
2.2.1. Analysantia of the Perceptual Representational Process	41
2.2.2. Theoretical Components of Representationalism	46
2.3. Content	48
2.3.1. Informational/Presentational Content	49
2.3.2. Semantic Content	54

2.3.3. Relationship Between Presentational and Semantic Content	56
2.3.4. Ontological Nature of Content	57
2.3.5. Constitution of Content – Particularism vs Generalism	60
2.3.6. Individuation of Content – Externalism vs Internalism	66
2.3.7. Conceptual vs Non-conceptual Content	68
2.4. The Content Problems	69
2.4.1. Metaphysical Explanation of the Content Problems	70
2.4.2. The Distality Problem and Other Content Determinacy Problems	71
2.4.3. The Primary Problem of Perception	75
2.5. Conclusion	77
3. Reductive-Naturalistic Representationalism: Naturalistic Relations and Functions	79
3.1. Reductive-Naturalistic Representationalism	79
3.1.1. Primacy of Theory of Application	79
3.1.2. Tracking Relation	80
3.1.3. Ontological Reduction of the Tracking Relation/Content	81
3.1.4. Demarcating the Landscape	82
3.2. Naturalistic Relations	84
3.2.1. Natural Signification	84
3.2.2. The Informational Relations – Causation and Correlation	85
3.2.3. Structural Correspondence	86
3.2.4. Difference Between Signification and Representation	88
3.2.5. Insufficiency of Relations	90
3.3. Necessity of Relations	91
3.3.1. An Argument for Content Externalism	91
3.3.2. Argument For Relationality from Action-Oriented Perception	92
3.4. Teleological Functions	103
3.4.1. The Selected Effects (SE) Theory of Functions	103
3.4.2. Information-Carrying Functions – Dretske and Neander	104
3.4.3. Consumer Functions – Millikan	105
3.4.4. Task Functions – Shea	108
3.4.5. Drawbacks of the SE Notion of Functions	108
3.4.6. The Causal Role Notion of Functions	110
3.4.7. The Fitness-Contribution Notion of Functions	110
3.5. Necessity of Teleological Functions	112
3.5.1. Normativity	112
3.5.2. Teleology	115

3.6. Conclusion	117
4. Reductive-Naturalistic Representationalism: Representational Application	118
4.1. Insufficiency of Relations and Functions	118
4.2. The Distality Principle – Neander	121
4.2.1. Problems for Neander’s Account	123
4.3. Constancy Mechanisms – Dretske and Schulte	124
4.3.1. The Constancy-for-Representation Condition	124
4.3.2. Evaluation of the CFR Condition	128
4.4. Use-of-Relations Condition – Millikan and Shea	130
4.4.1. Millikan’s Notion of ‘Use’	130
4.4.2. Shea’s Notion of ‘Use’	132
4.4.3. Problems With the Use Condition	133
4.5. Tracking Is Not Representing	135
4.6. Conclusion	138
5. Against Reductive-Naturalistic Representationalism and Content Particularism	139
5.1. General Problems with Reductive-Naturalistic Representationalism	139
5.1.1. Primacy of Theory of Application	139
5.1.2. Representations Are Merely Explanatory Posits	140
5.1.3. No Causal Efficacy to Content	142
5.2. Against Content Particularism	145
5.2.1. Schellenberg’s Capacitism	146
5.2.2. Schellenberg’s Argument for Content Particularism	147
5.3. Conclusion	156
6. Nonreductive Representationalism	157
6.1. Nonreductive Representationalism	157
6.1.1. Motivations for Nonreductive Representationalism	157
6.1.2. Basic Commitments	159
6.2. Con-relational Representationalism	161
6.2.1. Pautz’s Internalist-Nonreductive Representationalism	162
6.2.2. Con-relational Phenomenal Intentionality Theory (PIT) – Chalmers:	165
6.2.3. Objections to Con-relational Representationalism	168
6.3. Non-relational Representationalism	172
6.3.1. Crane’s Non-relational Aspectual Representationalism	172
6.4. Common Problems with Con-relational and Non-relational Representationalism	178
6.4.1. Essential External Directedness	178
6.4.2. The Distality Problem	180

6.5. Conclusion	181
7. Content Universalism: Immanent Content	182
7.1. Against Standard Content Generalism	182
7.2. Concrete Universals	183
7.2.1. Aristotle on Universals	184
7.2.2. Armstrong	190
7.2.3. The Nyāya-Vaiśeṣika System of Philosophy	193
7.2.4. Constructing a Theory of Universals	199
7.3. Content Universalism	204
7.3.1. Assumptions	204
7.3.2. Constitution of Content	205
7.3.3. Content As the Immanent Permeated Universal	208
7.3.4. Content as Powerful Quality	211
7.3.5. Content-Representation Pairing	212
7.3.6. Content Universalism (Final Version)	213
7.4. Conclusion	213
8. Content Universalism: Emergent Content	214
8.1. Irreducibility of Content	214
8.2. Ontological Emergence	216
8.2.1. Weak Emergence	217
8.2.2. Strong Emergence	218
8.2.3. Strong Emergence of Content	219
8.2.4. Downward Causation and the Causal Closure of the Physical	220
8.2.5. Other Views on the Strong Emergence of Intentionality	224
8.3. Phenomenological Particularity of Content	225
8.3.1. Attention	226
8.3.2. Attention and Particularity	228
8.3.3. Rensink's Triadic Architectural Theory of Attention-Modulated Perception	230
8.4. Conclusion	232
9. PRIME Representationalism	233
9.1. Representational Application/Directedness	233
9.1.1. Reference Magnetism	235
9.2. PRIMER: Pluri-Relational Immanent Emergent Representationalism	236
9.3. Ontological Kind of Relations in PRIMER	238
9.3.1. Defining Internal and External Relations	238
9.3.2. Externality/Irreducibility of Relations in PRIMER	240

9.4. Applications of PRIMER	244
9.4.1. Solution to the Distality Problem	244
9.4.2. Solution to the Problem of Non-Veridical Perception	246
9.4.3. Intentionalism	251
9.4.4. Metaphysics of Subjectivity	253
10. Conclusion	256
10.1. Metaphysical Highlights of PRIMER	256
10.2. PRIMER's Niche	257
10.3. Advantages of PRIMER	257
Bibliography	260

Statement of Copyright

The copyright of this thesis (including any appendices or supplementary materials to this thesis) rests with the author, unless otherwise stated.

© Gaurav Kudtarkar 2025

This copy has been provided under licence to the University to share in accordance with the University's Open Access Policy, and is done so under the following terms:

- This copy can be downloaded for personal non-commercial research or study, without prior permission or charge.
- Any quotation from the work (e.g. for the purpose of citation, criticism or review) should be insubstantial, should not harm the rights owner's interests, and must be accompanied with an acknowledgement of the author, the work and the awarding institution.
- The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

Acknowledgements

My foremost gratitude goes to my supervisors Professors Philip Goff and Matthew Tugby. Their perceptive and piercing feedback was instrumental in streamlining my ideas into what I hope is a coherent view. I will especially remember our joint meetings where we literally triangulated ideas. They were encouraging and kind throughout, making my PhD journey a memorable one.

I would also like to thank Prof. David Papineau at King's College London who looked at an early draft of chapters 3 and 4. His thoughts on the chapter were invaluable, and his advice on general research etiquette will be forever cherished. Thanks are due also to the participants of the BSPS 2024 conference in July 2024 at the University of York, whose comments on a paper helped polish the mentioned chapters.

A huge chunk of my gratitude pie goes to the Department of Philosophy. First and foremost, many thanks to the participants of the regular Metaphysics Reading Group, which sharpened my inclination towards metaphysics and encouraged me to attempt a robust integration of metaphysics and philosophy of perception. Special thanks to Matthew Tugby, James Miller and John Heil, in this regard. Secondly, my thanks go to the academic staff who gave me excellent feedback on various portions of the thesis on which I gave presentations and talks through the years. I am also grateful to the professional and support services staff at the department for making my PhD journey administratively smooth.

Next up, all my dear friends. Special thanks to Mansimer Singh Sethi whose random suggestion one evening in 2017 led me into the alleys of philosophy. Thanks also to my friends in Durham's Philosophy Department who were my social anchor in an otherwise solitary voyage.

No amount of thanks would suffice for my family, especially my mother and sisters. You were my emotional safety net from across two oceans! And Poonam: you were unfailing in your belief in me, professionally and personally; thanks for being there through this roller coaster of a ride!!

I am grateful to the AHRC-Northern Bridge Consortium for funding my research and conference visits.

0. Introduction

“There is in the living act of perception always something that glimmers and twinkles and will not be caught, and for which reflection comes too late.

No one knows this as well as the philosopher”.

– William James (1902/2002, 353)

In the above epigraph, William James, delivering the Gifford Lecture on natural religion in 1901-02 (later to become *The Varieties of Religious Experience*) at the dusk of his career, hits the bull’s eye when he notes the mystery inherent in that most basic of mental capacities – perception. My PhD thesis is among several attempts to chip away at this mystery. And I plan to do this by focussing on the intentionality of perception.

Intentionality is that feature of perceptual states in virtue of which they are directed towards external environmental entities, where the term ‘entities’ is to be understood in a broad sense to include objects, property-instances, and facts (i.e., obtaining state of affairs). The aim of my thesis, titled *A Relational-Representational Theory of Perceptual Intentionality: Immanent Content, Emergent Representation*, is to develop and argue for a novel metaphysical account of perceptual intentionality. Before I introduce my view, let me set the stage.

0.1. Delimiting the Boundaries

First, I will be focusing entirely on sensory exteroception – the perception of the external environment through the sense organs. Sensory exteroception includes visual, auditory, tactual, olfactory and gustatory perception. I will not be dealing with interoception (perception of internal states such as hunger), nociception (perception of bodily damage), proprioception (perception of body position) and kinaesthesia (perception of body movement).

Secondly, a starting assumption of my thesis is that I interpret intentionality in terms of representation – this is the representational theory of intentionality. According to this interpretation, perceptual states are representations of environmental targets, and their directedness just is their representing those targets in virtue of having the essential representational property of content. Content is a term of art in the philosophy and cognitive science literature. For now, content can be understood as the information that a perceptual

representation carries about a target by means of presenting the target in a certain way or under a certain aspect. The theory which accounts for *perceptual* intentionality in terms of representation is known as *representationalism*. Representationalism must be distinguished from a theory about phenomenal consciousness that also sometimes goes by the same name – ‘representationalism’ – but sometimes under the label ‘intentionalism’. Thus, the natural home of my thesis is representationalism, the representational theory of perceptual intentionality.

A third stage-setting point about my thesis is that I will be entirely concerned with certain debates *within* the representationalist camp. My focus will be on how representationalist theories address certain problems of perception, chief among which is the distality problem (more on this in §0.3). I will not be arguing for representationalism over non-representationalist theories of perception, the major ones among which are naïve realism, adverbialism and sense data theory (Fish 2021; Pautz 2021). Nor will I compare representationalist solutions to the problems of perception with non-representationalist solutions.

0.2. Representationalism

Representationalism is the representational theory of perceptual intentionality. A complete representationalist theory consists of two sub-theories: a theory of content and a theory of application – i.e., application of representations to targets.

A theory of content is concerned with accounting for how content is constituted (rather than merely caused). Most theories posit content-constituting relations – to concrete and particular environmental targets, or to abstract and general properties (i.e., universals), or to abstract propositions. The theories which posit content-constituting relations to environmental targets are known as *relational* representationalist theories. Some theories argue that content is an intrinsic property (an adverbial modification or an aspect) of the representation or the subject tokening the representation; these theories are the *non-relational* representationalist theories. Since there is no specific term for theories that posit content-constituting relations to abstract and general entities, I will call them *con-relational* (for content) representationalist theories.

A theory of application explains the application of representations to their targets in virtue of having content. In other words, application is intentional directedness. Representational directedness is often interpreted as a relation to targets. However, some argue that in case of hallucinations, there arguably are no targets. Therefore, the representational directedness

cannot be a relation, because relations imply the existence of their relata. I will, in my thesis, argue that there are targets in *every* case of perception – veridical or non-veridical (hallucinations included). Therefore, I will interpret representational directedness/application as a relation, which I will call the *Representational External Directedness (RED) relation* – to distinguish it from the content-constituting relations that are components of the theory of content.

Representationalist theories can be classified in several ways. A major classification that I will be following in my thesis is: reductive and nonreductive representationalism. Reductive representationalism is also known as reductive-*naturalistic* representationalism since it attempts to ontologically reduce representational content (and, hence, representation) to naturalistic entities. Naturalistic entities are those that are posited by the natural sciences, especially physics. Thus, ontological naturalism is interpreted in this context as synonymous with the thesis of physicalism (Papineau 2023). Nonreductive representationalism, in contrast, denies the possibility of such a reduction.

Reductive representationalism tends to be relational in the matter of content-constitution. For example, Dretske's (1986) causal-informational theory and Schellenberg's (2018) theory of capacitism. In contrast, its nonreductive counterpart tends to be either con-relational – eg., Pautz's (2021) internalist nonreductive representationalism – or non-relational (eg., Crane (2013)). Moreover, they differ in the problems they grapple with. Reductive representationalism is mostly concerned with the content-determinacy problem (Shea 2018, 6; Schulte 2023, 19) – what makes it the case that a representation R has the content $\langle C \rangle$ rather than some other content $\langle C^* \rangle$. There are several content determinacy problems, which I will discuss in chapter 2. However, the most important problem that reductive representationalism focuses on is the disjunction problem – why does representation R have the content $\langle C \rangle$ rather than $\langle C \vee C^* \vee C^\# \vee \dots \rangle$. Nonreductive representationalism, in contrast, is mostly concerned with the problem of misperception or non-veridical perception – the problem of explaining illusions and hallucinations in a way that is compatible with the direct realist account of perception, according to which perception involves direct access to a mind-independent world (Crane and French 2021).

0.3. Research Problem/Question

The core research problem around which my thesis will revolve is the *distality problem* of perception, which is a specific kind of content-determinacy problem. This is the problem of explaining *distal* content. That is, the problem is of explaining how perceptual content is *as of* suitably distal targets, given two conditions: (i) the perceptual system (brain and the central nervous system) is immediately causally sensitive to only incoming nerve stimulation; and (ii) the causal relation between a suitably distal target and the sense organs is mediated by other proximal state of affairs – for example, in the case of visual perception, the light reflecting from the object, the registrations of the sense organs. In other words, the problem is of explaining how the brain – caged in the claustrophobic cranium – can represent a distal world outside.

It should be noted that the distality of content is relative to the sensory modalities. Distal *visual* content is, plausibly, information about the surface properties of objects. However, distal *auditory* content is, arguably, information about the sound waves generated by vibrating objects rather than any property of the objects themselves (Nudds 2009). Similarly distal tactual content is often information about vibrations on the skin of the perceiving organism. This relativity is captured by the term ‘suitably’ distal.

The distality problem, to my mind, is the central problem of perception. I will argue for this in chapter 2. In this regard, I follow the psychologist Egon Brunswik who wrote that the “basic fact of perception is distal focusing” (Brunswik 1956, 61). However, the distality problem has not been given its due. Instead, much philosophical attention is directed at the problem of misperception or non-veridical perception. Although important in its own right, I will argue that any solution to the problem of misperception ought to be based on a solution to the distality problem. This is because even non-veridical perceptions are suitably distal – hallucinations, for example, are always *as of* distal objects and properties. Thus, a solution to the problem of misperception presupposes a solution to the distality problem.

Related to the problem of non-veridical perceptions is the problem of phenomenology of perceptual states. Most perceptual states are characterized by, apart from intentionality, phenomenal consciousness or “phenomenal character” (Fish 2021, 2) – that property in virtue of which there is *something it is like for the subject* to token those perceptual states (Nagel 1974). In other words, some perceptual states are *conscious* perceptual states or perceptual *experiences*. The problem of phenomenology is to explain how perceptual experiences have phenomenal character. It should be noted that not all perceptual states are experiences – some

are pre-conscious or unconscious. Pre-conscious perceptual states are plausibly those that are tokened in the early stages of the perceptual system (eg., early vision), whereas unconscious perceptions include cases such as blindsight.

The centrality attributed to the problems of phenomenology and non-veridical perception is evidenced in the traditional classification of theories of perception, which is based on how they account for the twin problems. Thus, we have the sense data theory, adverbialism/intrinsic qualia theory, naïve realism and representationalism. By focussing on the distality problem, my thesis will, therefore, bring this central problem to relief in the philosophy of perception, and also represent a fresh vantage point from which to theorize about perception.

0.4. Solutions

Although the distality problem might not be the main problem that reductive and nonreductive representationalism tackle, the theories apparently have the resources to address the problem. Reductive-naturalistic representationalism reductively explains the distality of content in terms of three ingredients – (i) content-constituting naturalistic relations, (ii) representational functions, and (iii) psychological mechanisms and principles that ensure the application of representations to distal targets – for example, Neander (2013) posits the distality principle; Burge (2010) and Schulte (2018) posit constancy mechanisms. I will critically discuss these ingredients in chapters 3 and 4, and argue that, while (i) and (ii) are necessary for distal content, (iii) is not. Consequently, their reductive solution fails.

Nonreductive representationalism does not posit content-constituting relations to the environment. Thus, it can explain the distality of content only by repurposing other resources at their disposal. For instance, con-relational representationalism might explain distality in terms of acquaintance with abstract universals (Chalmers 2006, Pautz 2021). Or both con- and non-relational representationalism might invoke the two-dimensional view of content, according to which there is a second stage of content – externalist or wide content – that accounts for distality (Chalmers 2010). I will critically evaluate these attempted solutions in chapter 6 and argue that they ultimately fail.

Instead, I will strike a middle path in my thesis, and will argue for and develop a *nonreductive* and *relational* version of representationalism. I will call my novel version *Pluri-Relational*

Immanent Emergent Representationalism (PRIMER). PRIMER has two theoretical components:

1. *Content Universalism*: This is PRIMER's theory of content, which I will develop in chapters 7 and 8. In constructing this theory, I will draw on the views of Aristotle and the classical Indian philosophical school of *Nyāya* on concrete universals and their role in perception. Content universalism is the thesis that perceptual content is partly constituted by the co-instantiation in the subject and target of the same concrete universals, but the manner of instantiation is different. Whereas universals are instantiated in the *qualification* manner in targets, they are instantiated in the *permeation* manner in subjects. In other words, whereas universals qualify targets, they permeate in subjects. Content is identical to the subjectively immanent (i.e., permeated) universal and is, therefore, a property of the subject. Moreover, I will argue in chapter 8 that content is a *strongly emergent* property of the subject. Therefore, content is irreducible to its constituting base.
2. *Representational Magnetism*: This is PRIMER's theory of representational application/directedness, which I will develop in chapter 9. I will argue that because the same universals are instantiated in different manners in the target and the subject, this grounds the representational relation between them. I will call this relation the Representational External Directedness (RED) relation. Further, I will argue that the RED relation is a weakly emergent entity.

PRIMER invokes relations in both its theories of content and application. The content-constituting relations are the naturalistic relations to the environment and the instantiation relation. Further, the intentional directedness of representations is interpreted in terms of the RED relation. Therefore, it is a pluri-relational theory.

PRIMER accounts for the distality of content in terms of the co-instantiation of the same universals in the subject and the target. More precisely, content is distal because the universals which are qualification-instantiated in distal targets are also permeation-instantiated in subjects. The universals of the proximal causal chain do not permeate in the subject and therefore do not play a role in content-constitution. In this regard, it scores over reductive-naturalistic and other nonreductive representationalist theories.

0.5. Significance

Apart from providing a solution to the distality problem, PRIMER has the following other advantages, which I will discuss in chapter 9. First, it provides an intentionalist account of the phenomenal character of perceptual states. Intentionalism holds that phenomenal character is identical to (strong intentionalism) or supervenes on (weak intentionalism) representational content. Although any theory of phenomenal character is compatible with representationalism, intentionalism is the most commonly espoused account. The main reason for this is that it allows for a unified account of both intentionality and phenomenal character in terms of content.

According to PRIMER, phenomenal character is identical to a certain kind of subjectively permeated universal – the conjunctive universal (Armstrong 1997, 31), which permeates the subject as phenomenal content. Thus, phenomenal character is identical to phenomenal content. PRIMER, therefore, espouses a form of strong intentionalism. PRIMER's intentionalism has an advantage over most other accounts of intentionalism. Most representationalists consider content to be an abstract entity such as a proposition or an abstract mode of presentation. They may be called 'abstractivist representationalists'. The abstractivist representationalists who are also intentionalists therefore explain concrete phenomenal character in terms of abstract content. Kriegel (2011) and Papineau (2021) object to this by noting the fundamental incompatibility of natures of the explanandum (concrete) and explanans (abstract). I call this the content-(phenomenal) character mismatch problem.

PRIMER avoids the content-character mismatch problem because content is a concrete property of the subject. This is because content is the subjectively immanent universal, given that other constitutive conditions have been met. Thus, concrete content is suitable to explain concrete phenomenal character. Concretist representationalism, which considers content as a concrete property, is not very common in the representationalist literature. However, some examples are Schellenberg's reductive and relational representationalist view of Capacitism (Schellenberg 2018) and Crane's nonreductive and non-relational representationalist theory (Crane 2009). My view would occupy a middle ground – it is nonreductive and relational.

Another advantage of PRIMER is its account of hallucinations. I will argue that the represented target is a spatiotemporal region rather than a particular object. Thus, there is a target even in the case of hallucinations. Now, most hallucinations are partial hallucinations, where some aspects of the target are represented veridically, and some hallucinated. The hallucinatory content can be explained by the permeation of deviant universals in the subject which are not

qualification-instantiated in the target. As I will explain in detail in chapter 9, deviant universals are those that do not qualify the target but, nevertheless, permeate in the subject. Thus, PRIMER offers a uniform relational account of content-constitution and representational external directedness for veridical and non-veridical perceptual experiences.

0.6. Methodology

The foregoing discussion summarized the substantive claims and applications of PRIMER. In this section and the next, I lay out the administrative part, starting with the methods and methodology.

0.6.1. Metaphysical analysis:

The method I adopt in my thesis is metaphysical analysis. The *Stanford Encyclopedia of Philosophy* entry on ‘Analysis’ (Beaney and Raysmith 2024) mentions that there are four conceptions of analysis: (i) *reinterpretation* – of the analysandum in order to facilitate analysis; (ii) *regression* – to more fundamental principles, on the basis of which to engage in progression to the analysandum; (iii) *decomposition* – of the analysandum into simpler components (the analysans), to be followed by synthesis; and (iv) *connection* – of the analysandum with other entities in a holistic network. The authors note that in practice analysis involves more than one of these four conceptions.

The metaphysical analysis of perceptual intentionality that I engage in involves the first three conceptions of analysis. First, following a respectable tradition, I *reinterpret* intentionality as representation. Then, I engage in *regression and decomposition* by identifying the more fundamental elements that are the constituents of content and, therefore, representation. These elements are the individually necessary and jointly sufficient conditions for content. Finally, I undertake a *progressive and synthetic task*, and argue that content is not only constituted by, but also strongly emergent from, its constituting base. The notion of constitution that I will use throughout the thesis is non-reductive (Baker 2007) and is compatible with strong emergence. It involves identifying the necessary and sufficient conditions of content. I will discuss this conception in chapter 1.

The metaphysical analysis of perceptual representation is in service of a metaphysical explanation of it. Metaphysical explanations are synchronic non-causal explanations (Brenner, et al. 2021, §1). By ‘synchronic’, it is meant that the explanandum and the explanans (either

essence or ground or constituter) exist at the same time. As I will discuss in chapter 1, there are broadly four types of metaphysical explanations: grounding, essentialist, reductive and constitutive. My aim will be to give a nonreductive constitutive explanation of content and representation.

0.6.3. Methodology

I will deploy the method of metaphysical analysis in consonance with two broad methodological principles. The first is the principle of ontological parsimony or Occam's Razor – "Entities are not to be multiplied beyond necessity" (Baker 2022, §2). There are two kinds of ontological parsimony – qualitative parsimony, which says that entities of different *kinds* must not be needlessly posited; and quantitative parsimony, which says that even entities of the same kind must not be needlessly spawned. Many philosophers, prominently Lewis (1973), argue that only qualitative parsimony ought to be a constraint on metaphysical theorizing. However, as with everything else in philosophy, there have been defences of quantitative parsimony as well (eg., Nolan 1997, Sendlak 2018). Therefore, for the sake of uniformity, I will follow both qualitative and quantitative ontological parsimony.

A second methodological principle is a weak form of methodological naturalism. Methodological naturalism – as opposed to ontological naturalism – is encapsulated by Quine's remark that philosophy must be "continuous with science" (Quine, 1969). Thus, it is the strong claim that "...philosophy and science are both concerned to establish synthetic knowledge about the natural world, and moreover to achieve this by *a posteriori* investigation" (Papineau 2023, §2.1). However, as Kornblith (2016) notes, there are weaker readings of methodological naturalism. One such reading gives *a priori* reasoning a role in philosophy along with space for *a posteriori* evidence. Another reading is one that considers "...the empirical literature as presenting certain constraints on philosophical theories without going so far as to see the solution to philosophical questions as a direct product of scientific work" (Kornblith 2016, 150).

I will adopt the weaker form of methodological naturalism. As I see it, it paves the way for an empirically informed, and to some extent constrained, metaphysics. The empirical sciences that I will advert to most are perceptual psychology, cognitive science and biology.

0.7. Structure of the Thesis

Having hiked up the slope of the substantive and methodological preliminaries, we are now in a position to have a vantage view of the thesis' chapters, in preparation of the trek we will soon embark on.

Chapter 1 – “The Representational Theory of Intentionality” – is an introduction to the theory that interprets the intentionality of mental states in terms of representation. I introduce the various notions of representation at work in different literatures (philosophy of perception, cognitive psychology), which range from a minimal notion that is concerned merely with semantic evaluation conditions of putative intentional states to a structurally robust notion that posits a content-bearing representational vehicle (a physical state), which participates in algorithmic computations.

I then introduce representationalism, which is the representational theory of *perceptual* intentionality. As regards the notion of representation used, I note that reductive-naturalistic representationalism mostly works with the vehicular notion of representation, whereas nonreductive representationalism prefers the minimal or the thin notion. I then argue that the vehicular notion is superior because adopting it aids in better dialogue with perceptual psychology, which always works with the vehicular notion. I end chapter 1 with an inventory of my metaphysical toolkit – the concepts that I will employ and refer to in my thesis; for example, ontological dependence, reduction, and constitution, among others.

In **Chapter 2 – “Representationalism About Perception”** – I go into the details of representationalism. I begin by noting the relationships between representationalism, on the one hand, and intentionalism and phenomenal intentionality theory (PIT), on the other. Intentionalism, unlike representationalism, is a theory about phenomenal character rather than intentionality. On the other hand, PIT is a theory of intentionality, but unlike representationalism, it grounds intentionality in phenomenal character. Next, I delineate the major elements or analysantia of the perceptual representational process: vehicle, target, representational content, and content-constituting relations. I argue that the target of representation is a spatiotemporal region of the environment, rather than distinct objects. This is followed by a delimitation of the theoretical parts of representationalism – a theory of content and a theory of application. A theory of content is of primary importance because it accounts for how content, the essential representational property, is constituted and determined. A theory

of application explains how representations are directed at their targets in virtue of having content.

Then, I discuss topics and debates revolving around representational content. I distinguish between two kinds of content, which I call presentational and semantic content. While the former is non-propositional, the latter is propositional. I then explicate the major debates regarding content – abstractivism vs concretism, particularism vs generalism, externalism vs internalism. At this stage of the thesis, I argue that the primary kind of content is presentational content, which is concrete, non-propositional and non-conceptual. Semantic (propositional) content is determined by presentational content. I also introduce the content-character mismatch problem for abstractivist representationalism that also endorses intentionalism – the problem that abstract content is not suitable to account for concrete phenomenal character.

Finally, I list the various content problems – which includes the problem of misrepresentation or non-veridical perception – and argue that the distality problem is the primary problem of content. The distality problem is about explaining how content consists of distal objects and properties that are attributed to distal targets, rather than proximal properties attributed to the proximal intermediaries of the causal chain between the target and subject.

Next, there follows a set of chapters (chapters 3, 4 and 5) which contains a critical discussion of Reductive-Naturalistic representationalism. Naturalistic representationalism explains the theory of content (i.e., constitution of content) in terms of their theory of application. It posits the tracking relation as the representational external directedness relation, and argues that content is a relational property. It then attempts a reduction of the tracking relation (and, consequently, content) to three naturalistic elements – naturalistic relations, biological (teleological) functions and representational application formulae (RAFs). These elements are deemed to be individually necessary and jointly sufficient for representation.

In Chapter 3 – “Reductive-Naturalistic Representationalism: Relations and Functions” –

I focus on the first two conditions – naturalistic relations and functions. Specifically, I argue that they are individually necessary for content. Naturalistic relations to environmental targets partly constitute content. My argument for this is based on the premise that the essence or nature of perception is to guide actions in order to ensure fitness and survival. This implies that perceptual content needs to be essentially action-oriented, consisting of action-properties or affordances that are attributed to environmental targets. This further implies that the subject needs to be naturalistically related to targets (spatiotemporal regions) in which subjects act.

Biological functions are a second necessary condition for content because they ground the normativity of content, which is what determines semantic content (i.e., evaluation/satisfaction conditions). They also underwrite teleological explanations of representations – they explain why representations exist. I argue that the appropriate notion of function for representationalism is the *fitness-contribution* notion.

In **Chapter 4 – “Reductive-Naturalistic Representationalism: Representational Application”** – I critically evaluate the prominent proposals for the third necessary naturalistic element – Representational Application Formulae (RAFs). RAFs are mechanisms and principles that enable the perceptual system to apply a representation to its distal target, and consequently to track it. The proposals I evaluate are the distality principle by Neander (2017), constancy mechanisms by Dretske (1986) and Schulte (2018), and the use-of-relations condition by Millikan (1984, 2004) and Shea (2018). I argue that none of these are even necessary for representation. Thus, the reductive project of naturalistic representationalism fails. Although two of the three elements of their package reductive solution are necessary, the third is not. Consequently, it is not sufficient to reductively explain representation.

In the first half of **Chapter 5 – “Against Reductive-Naturalistic Representationalism and Content Particularism”** – I give general arguments against naturalistic representationalism. I argue that by prioritizing the theory of application over the theory of content, they reverse the metaphysical order that a genuinely representationalist theory ought to follow. I also argue that it cannot adequately account for the causal efficacy of content. In the second half of the chapter, I focus on a thesis about content-constitution that naturalistic representationalism adopts called ‘content particularism’. It is the thesis that perceptual content is constituted by particular entities (particular objects or property-instances); thus, perceptual content is singular or particular content. I focus on Schellenberg’s (2018) argument for content particularism and argue that content particularism is false. I base my arguments on empirical studies about object-tracking and change blindness.

Next follows **Chapter 6 – “Nonreductive Representationalism”**. Here, I critically discuss two varieties of nonreductive representationalism. While neither endorses content-constituting relations to the environment, the difference between them is as follows. One version accepts content-constituting relations to abstract entities (such as universals) or relations to abstract contents themselves. I call this version *con-relational* representationalism. Proponents of this view are Chalmers (2006) and Pautz (2017, 2021). The other version is thoroughly non-

relational and argues that content is a non-relational property of the subject, and the exemplar here is Crane (2006, 2013). Both versions endorse content generalism, which is inverse to content particularism. Content generalism holds that perceptual content is constituted by general entities such as abstract universals. Thus, perceptual content is general content; for example, abstract modes of presentation and intensions (Chalmers 2004)

I concur with nonreductive representationalism on their espousal of generalism. However, I argue that neither of the two versions can satisfactorily account for the essential external directedness of perceptual representations and the distality of content. The resources available to them – acquaintance relation to abstract entities in the case of con-relational representationalism, and intrinsic psychological facts in the case of non-relational representationalism – are just not suitable to explain essential external directedness and distality, which are external spatial notions. Con-relational representationalism moreover faces the content-(phenomenal) character mismatch problem.

The next three chapters – 7, 8 and 9 – form the constructive part of my thesis, where I develop and argue for a novel nonreductive and relational theory of representationalism – Pluri-Relational Immanent Emergent Representationalism (PRIMER). Chapters 7 and 8 develop PRIMER's theory of content, which I call Content Universalism, which is version of content generalism. According to content universalism, content is partly constituted by the co-instantiation in the subject of the *same* concrete universals that are instantiated in the target, given that other content-constituting conditions – such as naturalistic relations and functions – obtain. However, the manner of instantiation of universals in the subject and target is different – whereas universals are qualification-instantiated in targets, they are permeation-instantiated in subjects. To put it another way, the same universals *qualify* the target and *permeate* in the subject.

In **Chapter 7 – “Content Universalism: Immanent Content”** – I construct a theory of concrete universals drawing on the insights of Aristotle, Armstrong and the classical Indian philosophical school of Nyāya. The main features of this theory are: (i) Intrinsically concrete universals: From Nyāya, I will adopt the view that universals are intrinsically concrete (spatiotemporal) independent of their immanence in particulars; (ii) Universalist Monism: From Armstrong, I will adopt the view that there is only one category of universals – non-substantial universals – that can non-mereologically form complex universals (Armstrong 1997); (iii) Instantiation Monism: From Aristotle and Nyāya, I will adopt the view that there is

one instantiation relation, but which relates its relata (universal and particular) in different manners. While universals are qualification-instantiated in targets qua objects, they are permeation-instantiated in subjects qua subjects.

Based on the above theory of universals, I will argue that perceptual content is identical to the subjectively immanent (i.e., permeated) universals, given three content-constituting conditions – naturalistic relations between the subject and the target, fitness-contribution biological functions of the subject, and the co-instantiation of universals in both target and subject in different manners. This idea is a precisification of Aristotle's view that perception involves the reception of the form (i.e., universal) of the perceived object without its matter.

In **Chapter 8 – “Content Universalism: Emergent Content”** – I argue for the nonreductive aspect of my novel version of representationalism. I specifically argue that content is a strongly emergent property of the subject, and the emergence base are the content-constituting conditions listed above in addition to the subject. Content exhibits the twin marks of strong emergence – it is coterporally dependent on its constituting base, and it has novel action-oriented causal powers. It should be noted that I employ Baker's (2007) interpretation of constitution, which is compatible with strong emergence. According to this interpretation, constitution is a generative relation between the constituted and the constituter, where the constituter is both necessary and sufficient for the constituted, and the constituted entity has distinct causal powers vis-à-vis the constituter.

Given that content is the subjectively immanent universal, it is a general content. I, therefore, explain how the phenomenological particularity of content can be accounted for. I argue that the psychological mechanism of *attention* supplies content its particularity. I discuss Rensink's (2000) triadic architecture theory of perception, which brings together discussion of general content and attention-induced particularity.

Chapter 9 – “PRIME Representationalism” – contains a discussion of PRIMER's theory of application and its solutions to various problems of content. I develop a theory of representational application/directedness, which I call ‘Representational Magnetism’. According to this, representations are directed at targets because the same concrete universals are instantiated (in different manners) in both the target and the subject. In the subject, the universals permeate as the content, whereas in the target, they qualify the target as property-instances. I call the directedness relation the Representational External Directedness (RED) relation.

I further argue that the relations that PRIMER posits, especially the RED relation are external relations – they are irreducible to the intrinsic properties of their relata. Particularly, I argue that the RED relation is weakly emergent. It has distinct causal efficacy compared to its dependence base, which includes the content and target property-instances. However, despite the distinct efficacy, its causal powers are plausibly not novel, but only a subset of the causal powers of content. Therefore, it is weakly emergent.

Then, I demonstrate the applications of PRIMER in addressing various problems and issues. First, I argue that it can straightforwardly solve the distality problem. This is because the universals that permeate in the subject (to constitute content) are the universals that originally qualify *distal* targets. Thus, content is also distal. The universals of the proximal intermediaries do not permeate in the subject and, therefore, are not part of content. Next, I argue that PRIMER can account for non-veridical perceptions as well, especially hallucinations. Since I argue, in chapter 2, that perceptual targets are spatiotemporal regions, I can maintain that there are targets even in the case of hallucinations. The non-veridicality is at the level of content. This is accounted for by the permeation in the subject of deviant universals – universals that do not qualify the target – along with normal universals that do qualify the target. Deviant universals are plausibly qualification-instantiated in other regions. This gives the result that content of hallucinations consists of both veridical and non-veridical aspects. Thus, PRIMER's account of hallucinations is an illusionist account (Ali 2018, Masrour 2020).

I then discuss PRIMER's compatibility with a strong intentionalist account of phenomenal character, according to which phenomenal character is identical to a certain kind of content – the content constituted by the permeation of conjunctive universals. I finally end my chapter and the thesis by suggesting a future area of research – the metaphysics of subjectivity. The specific research question would be the following: since permeation-instantiation is applicable only to subjects, what is it about subjects or subjectivity that grounds this unique manner of instantiation?

In sum, *Pluri-Relational Immanent Emergent Representationalism* (PRIMER) offers a novel theory of perceptual representation that has the potential to address the main problems of perception, account for phenomenal character, and branch into other area of metaphysics and philosophy of mind.

1. The Representational Theory of Intentionality

In this chapter, I introduce the representational theory of intentionality. In §1.1, I briefly explain intentionality – the directedness of mental states – and intentional objects – that towards which mental states are directed. I then note the different ways in which traditional and contemporary theories of intentionality are classified. Contemporary theories are classified into representational and anti-representational theories. In §1.2, I discuss the various notions of representation in philosophy and cognitive science ranging from a minimal notion to a structurally robust notion of representational vehicles – physical states of the cognitive system. §1.3 then introduces the representational theory of perceptual intentionality – representationalism, briefly noting the main arguments for it and the ways of internal classification. I advocate for the employment of the vehicular notion of representation in representationalism. Finally, in §1.4, I provide a primer of the metaphysical concepts that are applied in discussions about the metaphysics of representationalism.

1.1. Intentionality

Mental states are always about or directed at something. Intentionality is the property of directedness or aboutness of mental states towards entities known as intentional objects. The term intentional object is a functional term in the sense that an intentional object is whatever occupies the role of ‘that towards which the intentional mental state is directed’. Thus, intentional objects could be concrete mind-independent objects, properties, mental entities such as concepts, or abstract entities such as numbers and propositions. Kriegel defines intentional objects as “any objects which happen to be such that some intentional act is directed at them” (Kriegel 2008, 80). Similarly, Crane gives the following definition: “...the *object* of an intentional state is what the state is about, or what it concerns, or what it is directed upon” (Crane 2009, 475).

Some philosophers have argued that all and only mental states are intentional. That is, intentionality is both a necessary and sufficient condition for a state to be ‘mental’. This is known as Brentano’s thesis, which goes by the slogan “intentionality is the mark of the mental” (Crane 1998a). Perceptual states, being the paradigmatic mental states, are uncontroversially

characterized by intentionality. Even those mental states that are not about any mind-independent concrete object are characterized by directedness towards intentional objects. For example, when hallucinating a unicorn or thinking of Pegasus, there is no concrete (i.e., spatiotemporal) entity in the environment that the mental states are directed towards. Yet, there is an intentional object because the hallucination and the thought are about the unicorn and Pegasus, respectively¹.

1.1.1. Traditional Theories of Intentionality

The main motivation for the traditional theories of intentionality was not only to account for the directedness of the intentional state, but also to address the problem of “non-existent intentional objects” (Crane 2013, 94). For example, hallucinations are about things which have no concrete mind-independent existence. Kriegel (2008) calls such objects as *merely* intentional objects. They are merely intentional because they depend on the mental acts that are directed upon them for their existence. Accordingly, Kriegel writes: “x is a merely intentional object of act y iff (i) x is an intentional object and (ii) the following counterfactual holds: if there was no y, then there would not be x” (Kriegel 2008, 80).

Traditionally, theories of intentionality have been categorized along two axes – structure and relationality. On the basis of the structure of intentionality, there are two broad categories. First, dyadic intentionality theories – where intentionality involves two elements: the intentional state/act and the intentional object. Second, triadic intentionality theories – intentionality as involving three elements: the intentional state/act, the intentional object and the intentional *content*. Briefly, whereas intentional object is whatever the act is directed at, the content is the way in which the act is directed.

The relational basis of categorization focuses on explicating the nature of the directedness of mental states towards their objects. Relational theories interpret directedness in terms of a relation to intentional objects. Non-relational theories interpret the directedness in non-relational terms – either intrinsic properties of the subject or facts about the subject.

The two axes cut across each other, giving us a four-way categorization of intentionality theories as below:

¹ There have been many accounts of such non-existent intentional objects. For instance, Meinong’s liberal view of objects, according to which they are objects which do not have being (Marek 2024, §5.3.2). Another view is the adverbial view, according to which intentional objects are modifications of the mental state (Kriegel 2008).

	Relational	Non-relational
Dyadic (act, object)	• Early Brentano, before his reistic turn in 1903-1905 (Jacquette 2004, 100; Werner 2017, 133).	• Later Brentano (Crane 2017, 46; Kriegel 2018, 56).
Triadic (act, content and object)	• Twardowski (Twardowski 1894/1977; Jacquette 2004, 112) • Husserl – for veridical perceptions and existent objects (Smith and Smith 1995, 21; Zahavi 2017, 85)	• Husserl – for hallucinations and non-existent objects (Smith and Smith 1995, 21).

1.1.2. Contemporary Theories of Intentionality

Contemporary theories of intentionality are classified based on the issue of representationality. Accordingly, there are two categories, representational and non-representational theories of intentionality. Representational theory of intentionality is the “intentionality as representation” view (Searle 1983, 4), which interprets intentionality in terms of representation. That is, to be intentionally directed towards something (or to be about something) is to *represent* that thing. There are two further types of representational theories – relational (for example, reductive naturalist theories) and non-relational (non-reductive theories)².

Non-representational theories of intentionality account for intentional directedness in terms of a non-representational intrinsic properties (for example, adverbialism) or non-representational relations (for example, naïve realism).

1.2. Notions of Representation

The Representational Theory of Intentionality (RTI) analyses intentionality in terms of mental representations. A mental representation, to state the obvious, is a representation that is mental. I will go with the common assumption that it is uncontroversial what it is to be mental – to be mental is to be physically realized in the central nervous system (brain and spinal cord) and to

² More on relational reductive representationalist views in chapter 3 to 5. More on non-relational non-reductive representationalist views in chapter 6.

exhibit intentionality³. So, the weight of the above question is on what it is to be a *representation*.

It can be instructive to adapt Ryder's (2020) insight of viewing representation-hood as a continuum along the axis of structural complexity. Accordingly, there are the following notions of representation-hood in order of increasing structural complexity:

1.2.1. Minimal Representation

At the minimal end of the continuum, a mental representation is any mental state "that possesses semantic properties" (Ryder 2020, 234). These semantic properties are evaluation conditions. Evaluation conditions prescribe conditions that must obtain in the world in order for the representation to be evaluated on the touchstone of an appropriate norm. If it is a belief representation then the evaluation conditions are truth conditions, and the belief is evaluated as true or false. If it is perceptual representation, the relevant conditions are accuracy or veridical conditions. The semantic property of a representation is commonly known as its *content*⁴.

The semantic property (content) is often interpreted as a proposition, which can be shared by several different 'attitudes' such as belief, desire, hope, etc. Thus, mental representations according to the minimal notion are propositional attitudes – attitudes towards propositions. For example, a representational state 'it is raining' is said to be related to the proposition <it is raining>. And this proposition (not the state) can be paired with the attitudes of belief (belief that it is raining), hope (hope that it is raining), fear (fear that it is raining), etc.

According to the minimal conception, there is no commitment to any physical state that can be identified with the representation. Rather, on the minimal and purely semantic notion, "...mental representation requires only the true application of predicates that [are] both psychological and semantic" (Ryder 2020, 235). Ryder claims that even Dennett with his intentional stance approach (Dennett 1987) can be said to be a representational realist as per the minimal conception.

³ However, in fact, it is a matter of controversy what it is to be mental. Some philosophers argue that the mental or cognitive is a functional notion and can include entities beyond the boundaries of the subject. For example, such as Clark and Chalmers (1998) argue for the extended mind hypothesis.

⁴ More on content in chapter 2 (§2.2). There, I will argue that content has two aspects – a semantic aspect and an informational or presentational aspect. While semantic content is usually an abstract proposition, presentational content could be a proposition or a concrete property. I will argue for concrete presentational content.

1.2.1.1. Personal/Subpersonal explanations

A minimal mental representation, and the other notions that will be discussed below, can be better understood in the context of the personal/subpersonal distinction in psychology. The personal-subpersonal distinction is interpreted in several, often inconsistent, ways. As Drayson (2012) points out, it was first introduced by Dennett to mark a distinction between two levels of psychological *explanation* (Dennett 1969, 93), rather than psychological states. Given a certain explanandum – a cognitive capacity such as perception or external behaviour – a personal level explanation ascribes psychological predicates (the explanans) to the subject or person, whereas a subpersonal level explanation ascribes psychological predicates to a component of the person’s cognitive system.

The subpersonal level of explanation is preceded by a *functional analysis* of the cognitive system (Cummins 1975). This involves an analysis or decomposition of the system into components, which have their separate functions, and these components together contribute to produce the cognitive phenomenon or capacity of interest. The ascription of psychological predicates to *any* level of analysis, including the highest or coarse-grained level, would count as a sub-personal psychological explanation. In sum, “[f]unctional analysis in psychology involves a particular kind of decomposition: the decomposition of the person into *subpersons* to whom we ascribe the sorts of psychological predicates that can explain the personal-level capacities” (Drayson 2012, 5). To be clear, as Dennett emphasizes, only adjusted or attenuated psychological predicates are ascribed to functional components (Dennett 2007, 78). Thus, it steers clear of the homunculus fallacy – “attributing the *whole* mind to a proper part of the system” (Ibid., 88).

Drayson argues that the personal/subpersonal explanatory distinction does not automatically entail a metaphysical distinction of personal and subpersonal *states*. That move depends on further metaphysical assumptions, which will be discussed in the case of more structurally robust notions of representation.

The minimal conception of mental representation stays at the *personal* level of psychological *explanation* without any metaphysical commitment, and at this level psychological predicates such as beliefs and desires are ascribed to the subject. To be in a representational state, therefore, is for the subject “...to instantiate a certain relation to a proposition, rather than to possess an internal state” (Drayson 2012, 9). A metaphysically neutral *sub-personal* level of psychological explanation would be where psychological predicates are used only as *heuristics*

to help identify the underlying physiological and neurobiological causes of behaviour (Ibid., 9).

1.2.2. Thin Representation

A “thin” representation is an identifiable physically realized state of the subject that has semantic properties (Ryder 2020, 235). We get to the thin notion of representation when metaphysical commitments are added to the *personal* level of psychological explanation. Thus, the thin representation is a state that is ascribed to the subject as a whole, rather than any component of the subject’s cognitive system. In having this representation, the subject stands in a relation to a propositional content. The relation to contents is known as propositional attitudes – the belief attitude, the perception attitude, etc. (Byrne 2005, 245) or content relation (Chalmers 2006, 51).

1.2.3. Vehicular Representation

A thicker notion of representation is that it is a physically realized state of the functionally analyzed cognitive system that carries information about some entity (usually an external state of affairs or object) and is causally involved in mental mechanisms and processing in virtue of its physical properties. Such representations are also known as representational *vehicles*, and the information they carry is the representational content. Thus, content is not merely a semantic property, but also an informational property. The represented object or state of affairs is known as the “target” (Cummins 1996, 6).

A phrase that is often associated with the vehicular notion of representation is ‘standing-in’. The representational vehicle is said to stand in for a target state of affairs, in virtue of carrying information about it. This feature is highlighted in several definitions of what a representation is. Friedenberget al, for instance, write tersely that “a representation is something that stands in for something else” (Friedenberget al. 2022, 3). Godfrey-Smith adds a little more detail when he writes, “a representation is one thing that is taken to stand for another, in a way relevant to the control of behaviour... the paradigm case [is] when a person decides to control their behaviour towards one domain, Y, by attending to the state of something else, X” (Godfrey-Smith 2006, 45).

Representational vehicles are states of the functionally analyzed system, and so are *subpersonal* states. The metaphysical assumption that allows us to move from subpersonal explanations to subpersonal vehicles is the *computational* understanding of cognitive capacities. Since

computation requires transitions over identifiable states, the components identified in functional analysis are co-opted as computable states – that is, as representational vehicles.

The idea of computation was introduced by Turing (1936) in the context of a mathematical function, which is basically a mapping from an input domain to an output range (Colombo and Piccinini 2023, 14). “To compute a mathematical function $f: I \rightarrow O$ is to transform an input i belonging to set I into an output o belonging to set O by following an algorithm” (Ibid., 14). An algorithm is a pre-defined set of rules or a “method for calculating the value [i.e., output] of a function” (Crane 2016, 61). In the context of mental computation, the inputs are representational vehicles, and the algorithm is causal relations and transitions among the input vehicles. Given this, computation is a kind of “...content-preserving causal relation among symbols [i.e., vehicles]” (Fodor 1997, 833). By content-preserving is meant the preserving of the same semantic value (truth/accuracy) that the representations had before entering the causal-computational process. Computation, in short, is “causation that preserves semantic values” (Ibid., 832).

As noted above, the vehicular notion of representation is subpersonal only insofar as it is posited in the context of functional analysis of the system. It is not subpersonal in the sense of being unconscious. Drayson emphasizes that the distinction between personal and subpersonal states *does not* map onto the conscious/unconscious states distinction. In fact, “[t]he subpersonal level of explanation can posit conscious states, accessible to introspection” (Drayson 2012, 14). Instead, the distinction between conscious and unconscious states maps on Stich’s (1978) distinction between doxastic and subdoxastic states.

Regarding subdoxastic states, Stich writes that “...subjects have no access to them...[For example], we cannot report, nor are we consciously aware of any information about binocular disparity. Awareness stops at the level of apparent comparative depth” (Stich 1978, 505). In contrast, doxastic states such as beliefs and perceptual experiences “...form a consciously accessible, inferentially integrated cognitive subsystem” (Ibid., 508). There are several ways in which a mental state can be unconscious or subdoxastic; consequently, there are several types of unconscious states (Berlin 2011). Block calls this the Anna Karenina theory of the unconscious: “all conscious states are alike; each unconscious state is unconscious in its own way” (Block 2011, 34). But many of the unconscious perceptual states are representations, nonetheless.

The following table provides examples of the different kinds of representations due to the crosscutting of the personal/subpersonal and the doxastic/subdoxastic distinction:

	Subpersonal states	Personal states
Subdoxastic/ unconscious states	Low-level perceptual representations. For example, low-level vision, which “refers to those processes concerned with separating out the various physical factors in the scene that give rise to the pattern of luminance intensities in the image” (Rensink 2000, 20). Low-level vision includes, for example, the first two of Marr’s three stages of vision – the primal sketch, the 2½ dimensional (2½-D) sketch, and the 3-D model representation (Marr 1982, 37).	Freudian and Jungian unconscious; Subliminal perception; Blindsight
Doxastic/ conscious states	High-level perceptual representations posited by perceptual psychology	Perceptual experiences, beliefs

The vehicular notion of representation is predominant in classical cognitive science. Cognitive science since the 1950s has branched into its different sub-disciplines – classical, connectionist, embodied. However, Dawson writes that “cognitive scientists are united in sharing a key foundational assumption that cognition is information processing” (Dawson 2013, 4). Classical cognitive science is distinguished from the other disciplines by its claim that the information processing is done by means of computation over representations.

An important characteristic of representational vehicles in cognitive science is their systematicity. A representation is a part of a system where transformations in the specific details of a representation maps on to different targets. This is to avoid counting as representations cases where something stands in for just one target rigidly. Thus, Haugeland writes that something is a representation in virtue of “a general representational scheme such that...a variety of possible contents can be represented by a corresponding variety of possible representations” (Haugeland 1991, 62).

1.2.4. Syntactically Structured Representation

The thickest notion of representation is one according to which the representation (i.e., representational vehicle) it is syntactically structured. The analogy is to a linguistic sentence that is structured into words and according to a grammar. Representations are similarly posited to be constituents of a “language of thought” (Fodor 1975). Further, the syntactically structured notion adds that the same representation can be associated with different attitudes such as belief, desire, hope, etc. Thus, the syntactically structured notion “...is of an identifiable internal state [the representation] with semantic properties (as on the thin notion), with the additional requirement that this state can mix-and-match with the attitudes” (Ryder 2020, 239). The attitudes are still ‘propositional attitudes’ since the representations are associated with propositions (which are the content/semantic properties of the representations).

Thus, the syntactically structured representation combines the personal and sub-personal levels. It involves a propositional attitude (personal level state) towards a computable structured vehicle (subpersonal level state). As Fodor puts it, “having a particular propositional attitude is being in some computational relation to an internal representation” (Fodor 1975, 198). The vehicular notion, in contrast, does not explicitly require that the vehicles be pairable to any attitude; the idea is that the functional role (i.e., computation) of the vehicles is sufficient to constitute the attitudes, which are subsumed under representational content. For instance, Shea – a proponent of the vehicular notion – writes that

“I prefer to use ‘content’ for the full representational import of a state, so as to include a specification of its mode of representing [i.e., attitude]...The functional role of a putative representation fixes this full specification...It’s not always the case that the same attitude-neutral vehicle can be reused in different functional roles so as to give it different modes of representing [i.e., attitudes]” (Shea 2018, 178).

The syntactically structured representation is the main currency of Fodor’s Representational Theory of the Mind (RTM). As Fodor writes,

“RTM claims that mental representations are related to propositional attitudes as follows: For each event that consists of a creature’s having a propositional attitude with the content P (each such event as Jones’s believing at time t that P), there is a corresponding event that consists of the creature being related, in a characteristic way, to a token mental representation that has the content P” (Fodor 1997, 830).

As the above quote clarifies, the same content is ascribed to both the propositional attitude such as belief (personal level) and the mental representation (subpersonal level). Therefore, “RTM is a theory about the *relationship* of the personal to the subpersonal levels, specifically, that propositional attitudes are computational relations to subpersonal mental representations” (Von Eckardt 2012, 30).

Other than the above distinguishing feature, RTM shares the claims of vehicular representation. These mainly include the thesis that “thinking is computation” (Fodor 1997, 831) – that is, the mental processes that representations participate in are *computational* processes; and the thesis that content is information (Ibid., 833).

1.3. Representationalism

In this section, I will briefly introduce the representational theory of *perceptual* intentionality, which is better known by the term *representationalism*⁵. It is the view that perceptual states are essentially (i.e., by nature) representational – they are directed towards perceived entities in virtue of having content.

The opposing view – anti-representationalism – holds that perceptual states are not essentially representational. The anti-representational camp is very diverse and includes views such as Naïve Realism (Martin 2004, Logue 2012) and enactive cognition (Varela et al. 1991, Noe 2004). Naïve realism accounts for perceptual intentionality by positing a perceptual relation between subjects and perceived environmental objects; this relation usually goes by the name of “acquaintance” (Fish 2021, 102). According to enactivism, perception is partly constituted by a subject’s externally directed action. Consequently, it accounts for intentionality in terms of the perceiving agent’s actions.

A distinguishing feature of *perceptual* intentionality is that the intentional directedness of perceptual states is *essentially* external. That is, it is part of the essence of perceptual states that they are about objects and states of affairs that are (or appear to be) externally located and occupying spacetime. Pautz calls this the “essential external directedness” (Pautz 2021, 74).

⁵ I use ‘representationalism’ as a thesis about intentionality and as distinct from ‘intentionalism’ which is a thesis about phenomenology. More on this in chapter 2.

1.3.1. Arguments for Representationalism

The aim of my thesis is not to evaluate which of the two camps – representationalism or non-representationalism – has sounder arguments in its favour. I will assume a representationalist understanding of perception and argue for a novel version of it (in chapters 7 to 9). In lieu of detailed arguments, I will merely present the main motivations for representationalism.

There are several arguments in the literature in favour of representationalism – some substantive and some methodological. The substantive arguments emphasize that representationalism better explains non-veridical perceptions such as hallucinations (Nanay 2013, 34; Brogaard 2014, 2; Schellenberg 2018, 104). This is because representations essentially have accuracy conditions (semantic content) which enables the evaluation of perceptions as accurate or inaccurate. Other reasons are that representations best explain the fineness of grain of perceptual experiences (Schellenberg 2018, 104) and the twin features of perceptual experiences – “internal dependence” and “essential external directedness” (Pautz 2021, 103-104).

The methodological arguments emphasize that interpreting perceptual states as representations aids in better dialogue between philosophy, on the one hand, and cognitive science and perceptual psychology, on the other (Brogaard 2014, 2). This is because perceptual psychology predominantly interprets perceptual states (along with other mental states) as vehicular representations. Similarly, representations are a staple in classical (i.e., computational) cognitive science and connectionist cognitive science (Dawson 2013). They are also widely posited by most views espoused under the 4E cognition research programme – Embodied, Extended, Embedded and Enactive cognition (Newen et al. 2018). Within representationalism, reductive naturalistic representationalism is most closely allied with cognitive science and perceptual psychology.

1.3.2. An Argument for the Vehicular Notion of Representation

The primary explanandum of cognitive science and perceptual psychology is the behaviour of cognitive agents, understood broadly – externally directed bodily behaviour as well as internal behaviour such as production of perceptual and other psychological states. Classical computational cognitive science posits representations as the main explanans⁶ (Thagard 2005, 19). Representations, for their part, can explain behaviour because they are about their targets,

⁶ The other explanans is the algorithm which guides computations over representations.

at which behaviour is directed. Thus, the explanatory route from representations to behaviour can be traversed only by explaining how representations are intentionally directed at their targets in virtue of having content.

However, cognitive science is usually silent or deflationary about content (Chomsky 1995, 52-53; Egan 2014), focussing only on explanations of behaviour. This is where philosophical theories of content step in. They give an account of how content is constituted or determined, which can then feed into cognitive science's primary project. Thus, Shea writes about the basic desideratum of a theory of content as follows: "An account of how representational content is constituted in a class of systems should allow us to show why recognizing the representational properties of such systems enables better explanations of behaviour than would be available otherwise" (Shea 2018, 29). Theorizing about mental content, then, apart from being a substantive philosophical enterprise is also an example of a happy partnership between science and philosophy.

Now the notion of representation used in perceptual psychology (and cognitive science, in general) is mostly the vehicular notion of perceptual representations. Thus, naturalist representationalists such as Shea (2017) and Nanay (2013) have espoused a vehicular notion of perceptual representation. Having a common understanding of perceptual representation bodes well for both philosophy of perception and perceptual psychology since it allows for easy cross-permeation of ideas. Therefore, in my thesis, I will adopt the vehicular notion of perceptual representation.

To recall the discussion in §1.2.3, the vehicular notion is subpersonal, but only in the sense that representational vehicles are tokened in a functionally analysed perceptual system. That does not preclude them from being involved in conscious perception – that is in perceptual experiences. Thus, I will use perceptual *experiences* as the paradigm examples of perceptual states in my discussions, and will speak of their representational vehicles and content.

1.3.3. Types of Representationalism

There are several ways of classifying representationalist theories, some of which are as follows:

- (i) Relational vs Nonrelational representationalism (Schellenberg 2018): This classification is based on whether content-constituting relations to the environment are accepted.
- (ii) Particularist vs Generalist representationalism (Schellenberg 2018): Based on whether content is constituted by environmental particulars.

- (iii) Externalist vs Internalist representationalism (Pautz 2021), or wide vs narrow representationalism (Chalmers 2004): Based on whether content is individuated by relations to the environment.
- (iv) Essential vs Contingent representationalism (Papineau 2021), or intentionalist vs non-intentionalist representationalism: Based on whether representationalism is paired with intentionalism, the thesis that phenomenal character is explained by content.
- (v) Reductive vs Nonreductive representationalism (Pautz 2021; Chalmers 2004): Based on whether content and representation can be reduced to physical or naturalistically acceptable entities.

In my discussions and evaluation of current representationalist theories, I will follow the reductive vs nonreductive classification. This is for two reasons. First, it follows the biggest schism in intentionality (and representationalism) research in contemporary philosophy of mind (Kriegel 2013; Papineau 2021, 32). Kriegel calls the two opposing research programmes the “*Naturalist-Externalist Research Programme*” (NERP) and the “*Phenomenal Intentionality Research Programme*” (PIRP) (Kriegel 2013, 1). The NERP attempts to naturalize intentionality by reducing it to naturalistically acceptable non-representational entities⁷. The PIRP, in contrast, offers a non-reductive account of intentionality by grounding it in phenomenal character. The main posit of the PIRP is the phenomenal intentionality theory (PIT), which (roughly) claims that the intentionality of a mental state is constituted by or grounded in the phenomenal character of that state (Bourget and Mendelovici 2019, §2.1). However, the PIRP/PIT is just one way of implementing non-reductive representationalism. Therefore, the two camps may be labelled, in the broadest terms, *Reductive-Naturalistic Representationalism* and *Nonreductive Representationalism*. Nonreductive Representationalism holds that although content and representation may ontologically depend on other entities (naturalistic or otherwise), they do not reduce to them.

It is interesting to note that reductive-naturalistic representationalism mostly works with the vehicular notion of representation (§1.2.3). To recall, according to this notion, representations are physically identified states of a functionally analyzed perceptual/cognitive system. Thus, a representational vehicle is subpersonal level representation. In contrast, nonreductive representationalism mostly works with the minimal or thin notion of representation, according to which representations are states of the entire subject. Thus, it is personal level representation.

⁷ Ontological naturalism is equivalent to physicalism.

The second reason for focussing on the reductive/nonreductive divide is that it brings the causal efficacy of content to the forefront. Reductive representationalism does not give distinct causal efficacy to content, but rather to the entities in the naturalistic reduction base. In contrast, nonreductive representationalism holds that content has distinct causal efficacy vis-à-vis entities that it may ontologically depend on. In other words, content is an emergent entity/property. Content could either be a weakly emergent property (compatible with ontological naturalism/physicalism) or strongly emergent property (incompatible with ontological naturalism/physicalism). In chapters 7 to 9, I will develop a novel version of nonreductive representationalism, and will argue that content is a strongly emergent property. However, against the common tide of nonreductive theories, I will favour a subpersonal level vehicular notion of representation, because (as noted above in §1.3.2) it can aid better dialogue between perceptual psychology and philosophy of perception.

1.4. Metaphysical Primer

“Metaphysics is a serpent that has itself by the tail”

– D. M. Armstrong (1997, 28).

Representationalism and associated theses – such as content externalism and particularism⁸ – are formulated using metaphysical concepts such as constitution, dependence, etc. Therefore, it will be helpful to have a primer of the major concepts that I will be invoking in my discussions.

1.4.1. Metaphysical Explanation

My aim in this thesis is to give a metaphysical explanation of content by means of metaphysical analysis. Metaphysical explanations are to be contrasted with causal and teleological explanations which are *diachronic* explanations – the explanans and the explanandum do not co-exist spatiotemporally. Metaphysical explanations, in contrast, are synchronic non-causal explanations (Brenner, et al. 2021, §1) – the explanandum and the explanans (essence or ground or constituter) co-exist.

Broadly, there are four types of metaphysical explanations: grounding, essentialist, reductive and constitutive. A *grounding* explanation explains the existence of an entity in terms of

⁸ More on these in chapter 2.

grounds – more fundamental entities which non-causally “generate” (Rubenstein 2024) or “give rise” (Brenner, et al. 2021, §2) to it. An *essentialist* explanation, in contrast, explains the *nature* of an entity: what it is for an entity to be the entity it is. The explanans employed is the essence of the entity. A third type of metaphysical explanation is the *constitutive* metaphysical explanation, which explains the existence and/or nature of an entity by citing both necessary and sufficient conditions. Finally, there is the *reductive* explanation where the explanandum is reduced to certain explanantia. The following subsections contain more detailed discussions of essence, grounding, constitution and reduction.

1.4.2. Essence

Most contemporary metaphysicians agree that all entities – whether properties, substances, relations, or abstract entities – have essences or natures. However, there are two competing ways of understanding the essence of an entity – the existential-modal account and the real-definitional account (Roca-Royes 2011, 66).

According to the existential-modal account, the essence of an entity is that property which the entity has of metaphysical necessity. In other words, “P is an essential property of *a* if and only if *a* could not exist without being P” (Roca-Royes 2011, 66). Essence is understood in terms of modality here. The real-definitional account, in contrast, interprets essence as the definition of an entity which tells what that entity is. Lowe, a prominent advocate of the account, puts it thus: “...the essence of something, X, is *what X is*, or *what it is to be X*. In another locution, X’s essence is the very *identity* of X” (Lowe 2018, 16). Lowe emphasizes that essence must not be interpreted as an entity itself. The real-definitional account is a non-modal account in the sense that it explains modality in terms of essence, rather than the other way around. The real-definitional account, in other words, endorses the essentialist theory of modality (Leech 2021, 887). In this thesis, I will be using essence and its cognate essential in the real-definitional sense.

Lowe further distinguishes between general essences and individual essences. Given that an entity X belongs to a kind K, “...X’s *general* essence is *what it is to be a K*, while X’s *individual* essence is *what it is to be the individual of kind K that X is*, as opposed to any other individual of that kind” (Lowe 2018, 16). For example, Tigger and Hobbes, the tigers, have the same general essence of what it is to be a tiger, but different individual essences – the individual essence of Tigger is what it is to be Tigger rather than Hobbes or any other tiger. Even if two qualitatively identical objects have the same general essence, due to which they belong to the

same kind, they will have different individual essences, in virtue of which they are distinct individuals or particulars. In other words, the individual essence of an entity is its *particularity*.

1.4.3. Particularity

Particularity, since it is just the individual essence of an entity, is not an entity itself. Thus, it is not a property of an object. Armstrong (1997) is an advocate of this ontology-lite approach towards particularity. He writes: “...the particularity of particulars is not to be accounted for in terms of universals, or, indeed, in terms of properties taken as particulars. It is an irreducible feature of particulars, distinct from their properties” (Armstrong 1997, 107).

Interestingly, Armstrong argues for the non-property view of particularity on the basis of perception. He starts with the observation that in the case of non-veridical perceptions such as hallucinations there is a perceptual content “*that* something is the case” (Armstrong 1997, 96). However, there are no particular state of affairs that are the intentional object of the perception. Particulars are always concrete – that is, spatiotemporal. Therefore, “...it can be claimed that the particularity of particulars, the fact that they are not exhausted by their properties and relations, is part of the content of perception” (Ibid., 96).

1.4.4. Ontological Dependence

An entity *F* can be ontologically dependent on another entity *G* for either its existence or individuation. Thus, we have two varieties of ontological dependence – existential dependence and individuation (or identity) dependence. I will be mostly concerned with existential dependence in the context of representationalism and content-constitution.

There are two ways of analysing existential dependence – in terms of modality and essence (Tahko and Lowe 2025). Modal existential dependence can be defined as: “*x* depends for its existence upon *F*s =_{df} Necessarily, *x* exists only if some *F* exists” (Ibid., §2.1), whereas essential existential dependence can be defined as: “*x* depends for its existence upon *F*s =_{df} It is part of the essence of *x* that *x* exists only if some *F* exists” (Ibid., §4.3). It is not crucial for my arguments which type I adopt, and so I will remain neutral towards both.

The intuitive idea is that if *M* existentially depends on *P*, then *P* is a necessary condition for *M*’s existence. That is, *M* could not exist in the absence of *P*. However, *P* need not be sufficient for *M*. For example, water existentially depends on hydrogen and oxygen atoms. However, the atoms are not sufficient for water; rather the sufficient condition is that the atoms bond to form molecules which in turn bond to produce water.

As regards the formal characteristics of ontological dependence, it is mostly argued to be non-reflexive, non-symmetric and transitive (Tahko and Lowe 2025, §2.2).

1.4.5. Grounding

Grounding is a metaphysical generative relation between two entities or classes of entities – the ground and the grounded. If M is grounded in P, then M exists because of or in virtue of P. Moreover, grounding is considered to be explanatory (Bliss and Trogon 2024, §1.1) and has implications on fundamentality (Ibid., §4.3). Thus, if M is grounded in P, then P is more fundamental than M, and P explains why and how M exists. The intuitive idea is behind grounding is that: if M is grounded in P, then necessarily, P is a *sufficient* condition for M's existence (Brenner et al. 2021, §2).

Dependence, in contrast, does not have any implications about fundamentality or explanatoriness. Also, the dependee is a necessary condition for the depender, whereas the ground is a sufficient condition for the grounded. Grounding and dependence also have different formal features. There is widespread agreement (though not consensus) that grounding is irreflexive, asymmetric and transitive (Bliss and Trogon 2024, §2.3).

There is also a distinction between full grounds and partial grounds. The standard definition of partial grounding is that P partially grounds M “just in case [P]... either on its own or together with other facts fully grounds [M]” (Bliss and Trogon 2024, §3.2). In other words, a full ground can generate and explain the grounded by itself, whereas a partial ground can generate and explain the grounded only in conjunction with other facts. Usually, the term ‘grounding’ simpliciter is used to mean partial grounding.

1.4.6. Necessitation/Entailment

Grounding is often associated with the concept of ‘necessitation’. For instance, according to *necessitarianism*, if P fully grounds M, then P necessitates M. And, Ps metaphysically necessitate Ms “...just in case any (metaphysically) possible world in which the former obtains is a world in which the latter obtains” (Bliss and Trogon 2024, §4.1). Paolini Paoletti explains necessitation in terms of implication: P necessitates M just in case “necessarily, the existence...of [P] implies the existence...of [M]” (Paolini Paoletti 2025, 51). The notion of necessity used is metaphysical necessity. For example, a determinate property (scarlet) necessitates a determinable property (red).

Metaphysical necessitation, then, seems to be the same as property entailment, according to which “...property P entails property [M] just in case it is metaphysically necessary that anything that possesses P also possesses [M]” (McLaughlin and Bennett 2023, §3.2). As regards formal features, necessitation or entailment is “reflexive, transitive, and non-symmetric” (Ibid., §3.2).

1.4.7. Supervenience

Supervenience is a relation of “dependent-variation” (McLaughlin 1995, 18). The core idea is that Ms supervene on Ps iff there cannot be any variation in Ms without variation in Ps. Since the emphasis is on any possible variation (“there *cannot* be any variation”), supervenience has an intrinsic modal force (McLaughlin and Bennett 2023, §1). McLaughlin and Bennett argue that supervenience is distinct from both ontological dependence and grounding. For one, they argue that if M supervenes on P, then it neither follows that M depends on P nor that P grounds M (Ibid., §3.5). For example, two effects of a common cause supervene on each other; however, they neither depend on nor are they grounded by each other. For another, supervenience has different formal features as compared to the other two relations – supervenience is reflexive, non-symmetric and transitive.

The supervenience relation is invoked extensively in philosophy of mind, especially in the context of reductive physicalism (Kim 2005). The claim is that supervenience is necessary but not sufficient for the reduction of mental properties to physical properties (McLaughlin and Bennett 2023, §3.3). That is: if M reduces to P, then M supervenes on P; but not vice versa.

1.4.8. Ontological Reduction

The intuitive idea of ontological reduction⁹ of Ms to Ps is captured in the slogans ‘[M]s are nothing over and above [P]s’ or ‘[M]s are nothing more (or other) than [P]s’ (Van Gulick 2001, 2). There are several routes to reduction. The two commonly accepted routes are elimination and identity (Ibid., 4; Paolini Paoletti 2025, 6). Eliminative reduction is when “[M] does not exist and every truth about [M] is actually a truth about [P]” (Paolini Paoletti 2025, 6). For example, demon possession has been eliminatively reduced to schizophrenia with auditory hallucinations (van Gulick 2001, 4). In the case of identity reduction, “[M] is...identical with [P] – but the description given by the term “[P]” more perspicuously grasps the essence of the

⁹ Ontological reduction must be distinguished from theoretical reduction (Nagel 1961) where the statements of a higher-level special science are said to be reducible to the statements of a lower-level science (particularly physics) by means of ‘bridge laws’ which allow the derivation of concepts of the higher-level science from the concepts of the lower-level science.

entity that is both [P] and [M]” (Paolini Paoletti 2025, 6). For example, lightning is identical to atmospheric discharge of static electricity and can be reduced to the latter without elimination. Eliminative reduction is a special type of identity reduction where the reduced description is dropped from the ontological roster altogether.

Apart from the above, there is difference of opinion on what relations between M and P would suffice for M to reduce to P. Kim proposes ‘functionalization’ as the sufficient relation for reduction, and calls his preferred model “*functional reduction*” (Kim 2005, 101). To functionalize a putative reducible entity M is to give a causal-functional definition in terms of the causal roles of the entity: “Having M =_{def.} having some property or other P (in the reduction base domain) such that P performs causal task C” (Ibid., 101). This is followed by identifying “the properties (or mechanisms) in the reduction base that perform the causal task C” (Ibid., 102). Kim calls these lower-level causal performers the “realizers” of the reducible property (Ibid., 102). Given these two steps, we can say that M functionally reduces to P. For example, a gene functionally reduces to DNA molecules, because it is the DNA molecules that code for proteins and are transmitted through reproduction.

In his discussion of reducibility of relations (i.e., the debate about internality vs externality of relations), Paolini Paoletti mentions *determinative reduction*, where determination is the sufficient relation for reduction. Paolini Paoletti defines determination in terms of ontological dependence and necessitation:

“When I shall claim that [the Ps] *determine* [the Ms], I shall invoke *full* dependence *plus* necessitation...More precisely, I shall mean that, when it comes to...[the Ms]’s existence/exemplification/obtaining, [the Ms] fully depend on [the Ps] *and* [the Ps] necessitate [the Ms]” (Paolini Paoletti 2025, 6).

To be sure, Paoletti uses ‘full dependence’ to mean a combination of ontological dependence and grounding (§1.4.4 and §1.4.5)¹⁰. Therefore, in terms of the notions introduced above, determination is a combination of dependence, grounding and metaphysical necessitation/entailment. The intuitive idea is that: if P determines Ms, then it follows that the P is both necessary and sufficient for M, and P metaphysically necessitates M. Given this, M is said to be determinatively reducible to P.

¹⁰ Paolini Paoletti interprets dependence as an explanatory (i.e., grounding) relation. That is, if B depends on A, then not only is A necessary for B’s existence, but it also explains B’s existence (Paolini Paoletti 2025, 5).

1.4.9. Constitution

The term constitution is employed extensively in metaphysics (under the label ‘material constitution’) and the philosophy of mind and perception, including in the context of representationalism. Constitution must be distinguished from causation, on the one hand, and composition, on the other. Causation is commonly understood a diachronic generative relation, where an entity or an event gives rise to another, in the presence of causal conditions. It is diachronic in the sense that the cause and effect do not exist at the same time – the cause ceases to exist after generating the effect.

Constitution and composition, in contrast, are synchronic relations – the constituted/composed and the constituter/composer coterminally exist. Composition is a mereological relation between a whole and its parts. Constitution, in contrast, is a non-mereological relation between an entity and other entities that generate it. Paradigm examples of constitution are a lump of bronze constituting a statue, the extension of an arm from a moving car (in the appropriate circumstances) constituting a signal to turn.

There is no consensus on the nature of the constitution relation. Constitution is sometimes interpreted to be identical to grounding (Dasgupta 2017, 75; Schellenberg 2018, 16). So, ‘P constitutes M’ is synonymous with ‘P grounds M’. Some interpret constitution to be a specific kind of grounding relation (Kirchhoff and Kiverstein 2024, 3). Burge interprets constitution as a combination of ontological dependence and grounding: “...constitutive conditions...are conditions that are necessary, sufficient, or necessary and sufficient to be something of that kind or with that nature, and that are in principle potentially relevant to explaining, understanding, illuminating the kind or nature” (Burge 2010, 58).

Similarly, Baker (2007, 2009) interprets constitution as a generative relation which prescribes both necessary and sufficient conditions for the constituted entity. P constitutes M just in case: P generates or gives rise to M in *M-favourable circumstances*. *M-favourable circumstances* “...are the milieu in which something can have the property of being [an M]” (Baker 2007, 113). For example, in the circumstances of driving and pre-set rules of traffic, an extension of an arm outside the window constitutes a right-turn signal. If P constitutes M, then P is a necessary condition for M, and the combination of P and *M-favourable circumstances* is sufficient (Baker 2007, 162).

Despite the constituter and favourable circumstances being necessary and sufficient for the constituted entity, the latter is not reducible to the former. This is because Baker holds that

“[c]onstituted objects have different causal powers from their lower-level constituters” (Baker 2007, 33). Baker emphasizes the irreducibility of the constituted to the constituters in her slogan that constitution is “unity without identity” (Baker 2007, 166).

Talk of different causal powers launches Bakerian constitution into the orbit of the emergence literature. And indeed, Baker acknowledges the connection between constitution and emergence, noting that constitution could be interpreted as a way of precisifying the generation of emergent properties. Thus, she writes that “...the Constitution View provides for an account of levels of reality. If we add the plausible hypothesis that higher-level primary kinds come into being over time, we get an ontologically robust kind of emergence” (Baker 2007, 237).

In my formulation of my novel theory of representationalism¹¹, I will adopt the Bakerian notion of constitution. The reasons for this are: first, it carves a distinct niche within the conceptual space of metaphysical generative relations (distinct from grounding or dependence, causation or composition). Second, its endorsement of novel causal powers for the constituted entity fits well with my eventual argument that content is an emergent property of the subject.

1.5. Conclusion

In this chapter, I introduced the representational theory of perceptual intentionality – representationalism. I argued for adopting a vehicular notion of representation because it will aid better dialogue between philosophy of perception and the corresponding sciences. I also indicated the specific interpretations of the metaphysical relations and tools that I will employ in developing my novel version of representationalism. In particular, I will use ‘essence’ of an entity as a real definition of the entity; and ‘constitution’ in the Bakerian sense of providing necessary and sufficient conditions without the implication of reducibility.

¹¹ I call my theory Pluri-Relational Immanent Emergent Representationalism (PRIMER). I develop it in chapters 7 to 9.

2. Representationalism About Perception: Key Terms and Debates

In this chapter, I discuss representationalism in detail. In §2.1, I discuss the relationship of representationalism to related views. I distinguish representationalism from intentionalism, followed by noting the similarities between representationalism and phenomenal intentionality theory (PIT). §2.2, then, lays out the elements of the perceptual representation process – vehicle, target and content, among others. It also delineates the theoretical components of a complete representationalist theory. These are the theory of content and the theory of application or directedness.

Next, in §2.3, I explicate topics and debates revolving around the essential representational property – content. In particular, I note the two kinds of content – presentational and semantic – and debates such as particularism vs generalism, internalism vs externalism. Finally, in §2.4, I discuss the various content problems including the content determinacy problems and the problem of misrepresentation or non-veridical perception. I argue that the distality problem – a specific kind of determinacy problem – is the primary problem of perception.

2.1. Representationalism about Perception

Representationalism about perception is the thesis that perception is *essentially* representational. That is, it is the essence (or nature) of perception that perceptual states are content-bearing perceptual representations. Schellenberg puts this point across in terms of constitution: “[a]ccording to representationalism, a perceptual state is constitutively a matter of representing the environment” (Schellenberg 2018, 74). Now, to be a representation is to essentially have content – the informational-cum-semantic property in virtue of which representations stand in for external states and are directed towards them. Therefore, representationalism is also known as the “content view” (Siegel 2010, 4) or the view that endorses the “content thesis” (Schellenberg 2018, 105) – the view/thesis that perceptual states essentially have content.

Brogaard writes that representationalism commits to the “Strong Content View: Experience e has the proposition p as a content iff necessarily, the subject of e bears a propositional attitude

towards p ” (Brogaard 2014, 3). The strong content view is formulated in terms of a propositional attitude, and therefore in terms of the minimal or thin notion of representation discussed in chapter 1 (§1.2), which are personal level notions of representation. However, many reductive-naturalistic representationalist theories adopt a subpersonal vehicular notion of representation, but sometimes combine it with personal level attitude to propositional contents.

The strong content view is contrasted with the “Weak Content View: Experience e has the proposition p as a content iff p is conveyed to the subject of e , and necessarily if e is accurate, then p is true” (Brogaard 2014, 2). The difference between the strong and weak versions, it will be noted, is that in the former there is a commitment to a propositional *attitude* that corresponds to a relation between a subject and a propositional content. Schellenberg calls the weak content view the *association thesis*, whereby perceptual states are merely post-facto associated with propositional contents (i.e., evaluation conditions) without those contents being a constitutive aspect of the perceptual states. As she puts it: “*Association Thesis*: Every experience can be associated with (propositional) content in the sense that sentences can be articulated that describe how the environment seems to the subject, without the content expressed being a proper part of the experience” (Schellenberg 2014, 201).

The weak content view/association thesis could be accepted by any theory of perception, whether representationalist or anti-representationalist (Schellenberg 2014, 201). For instance, Logue (2014) – a proponent of naïve realism – argues that naïve realism is compatible with the weak content view, which she calls “the Mild Content View” (Logue 2014, 222-223). Her argument – the argument from belief generation – is based on the premise that perceptual experiences ground *epistemic appearances*: “a given experience E is associated with a particular *epistemic* appearance that p – the subject is disposed to believe that p solely on the basis of E , given that she is rational and doesn’t suspect that E is misleading” (Ibid., 234). Thus, her argument is:

- “1. Any given experience E is associated with a particular epistemic appearance that p .
2. The best explanation of (1) is that there is a proposition associated with E that captures the way things perceptually appear to the subject in virtue of having E .
3. Hence, for any given experience E , there is a proposition associated with E that captures the way things perceptually appear to the subject in virtue of having E (i.e., the Mild Content View)” (Logue 2014, 235).

2.1.1. Representationalism and Intentionalism

Representationalism is, in effect, the representational theory of *perceptual* intentionality. According to it, the intentional directedness of perceptual states towards their intentional objects is accounted for by their content. Strictly speaking, then, representationalism is neutral regarding how the phenomenal character of experiential perceptual states (i.e., perceptual *experiences*) is determined.

Phenomenal character is that property of some perceptual states in virtue of which there is *something it is like* for the subject to have those states (Nagel 1974). This property is referred to by other terms such as ‘phenomenal consciousness’, and ‘phenomenology’ (Van Gulick 2025, §2.2). Phenomenal character encompasses the subjective experience of perceptual *qualities* such as colours, smells and tastes. It also includes the experience of spatial and temporal *structure* to perceptions. To be sure, not all perceptual states are conscious in this phenomenal sense, whereas all perceptual states have intentionality.

Representationalism may be paired with any account of phenomenal character. For example, Papineau argues that phenomenal character is determined by intrinsic non-representational qualitative properties of perceptual states (Papineau 2021, 83-86). These non-representational properties are known as ‘qualia’ (Block 2004). Tye, in contrast, argues that phenomenal character is determined by the subject’s representation of (or, more precisely, acquaintance with) phenomenal qualities which are externally located in the world (Tye 2009a, 119). He calls the view “property representationalism” (Tye 2021, 71)¹².

However, representationalism is most often coupled with a view known as *intentionalism*. Intentionalism is the thesis that the phenomenal character of perceptual experiences is either identical to or supervenes on representational content. Intentionalism is therefore a representational theory of phenomenal character, as opposed to representationalism, which is a representational theory of intentionality. Papineau calls the pairing of representationalism and intentionalism ‘essential representationalism’: “conscious sensory properties...are one and the same as representational properties” (Papineau 2021, 31)¹³.

There are two versions of intentionalism – strong and weak. Crane (2009) draws the distinction as follows: *strong intentionalism* is the thesis that “the phenomenal character of a mental state

¹² Tye (1995) earlier endorsed a different view about phenomenal character – strong intentionalism.

¹³ Contrasted with essential representationalism is contingent representationalism: “The conscious and representational aspects of experience are only contingently connected” (Papineau 2021, 31).

is identical with its representational content”, whereas *weak intentionalism* holds that “the phenomenal character of a state...supervenes on its representational content” (Crane 2009, 479). A similar distinction is made by Tye (2009b; 2009c) who clarifies the supervenience claim in weak intentionalism as “necessarily any two states that are alike with respect to the relevant representational content are alike phenomenally” (Tye 2009c, 256). That is, weak intentionalism claims that phenomenal character merely covaries with content¹⁴. Thus, Schellenberg explains weak intentionalism as follows: “[t]he content of an experiential state corresponds one-to-one with its phenomenal character in that any changes in content go hand in hand with changes in phenomenal character and vice versa” (Schellenberg 2018, 78).

Intentionalism is sometimes referred to in the literature, confusingly, as ‘representationalism’ (Chalmers 2004, Tye 2009c, Seager and Bourget 2017, Lycan 2023) and sometimes as ‘intentionalism’ (Pautz 2007, Crane 2009, Speaks 2015, Schellenberg 2018). To avoid confusion, I will follow Schellenberg (2018) and use the term ‘intentionalism’ for the thesis about the phenomenal character of perceptual states, and the term ‘representationalism’ for the thesis about the intentionality of perceptual states, irrespective of how phenomenal character is determined. This usage will align with the usage of ‘representationalism’ in cognitive science and perceptual psychology as well, which are concerned predominantly with intentionality.

In my thesis, I will be focussing on representationalism (about perceptual intentionality) rather than intentionalism (about phenomenal character). However, the novel version of representationalism¹⁵ that I will develop in chapters 7 to 9, will have implications on intentionalism. I will briefly note these implications in the concluding chapter 9. I am inclined to favour the pairing of representationalism and intentionalism, mainly because it is the simple and elegant view – representational content explains both the intentionality and phenomenal character of perceptual states.

2.1.2. Representationalism and the Phenomenal Intentionality Theory

Since representationalism is primarily about intentionality, I will include the phenomenal intentionality theory (PIT) under representationalism. As noted in chapter 1 (1.3.3), the PIT is a nonreductive theory about intentionality, and holds that the intentionality of a mental state is

¹⁴ There are other bifurcations of intentionalism – pure vs impure; wide vs narrow; reductive vs non-reductive (Lycan 2023). These are not important to my thesis.

¹⁵ Very briefly, I will argue for a version of representationalism whereby content is constituted by the immanent universals of external targets, and the content then constitutes an irreducible non-naturalistic representational relation to targets.

constituted by or grounded in the phenomenal character of that state (Bourget and Mendelovici 2019, §2.1). In the case of perception, then, perceptual intentionality is grounded in the phenomenal character of perceptual experiences. The intentionality of non-phenomenal states is derived by their being appropriately related to phenomenal intentional states. PIT interprets intentionality usually in terms of representation, and therefore PIT can be interpreted as a version of representationalism. The difference between them is that representationalism mostly (i.e., when paired with intentionalism) accounts for phenomenal character in terms of intentionality (i.e., content), whereas PIT gives an independent account of phenomenal character.

2.2. Representationalism – Analysantia and Theoretical Components

2.2.1. Analysantia of the Perceptual Representational Process

I will be working with the subpersonal vehicular notion of representation, for reasons articulated in chapter 1 (§1.3.2). A distilled description of the vehicular notion of representation is as follows:

Representation: A representation is a physically realized entity – i.e., vehicle – that carries information – i.e., content – about an environmental state of affairs. In virtue of carrying information, it stands in for the environmental state and acquires semantic properties (evaluation conditions). Further, it is causally involved in mental mechanisms and processing in virtue of its formal – i.e., syntactic and structural – features.

The following analysantia or elements of the perceptual representational process can be gleaned from the above description:

2.2.1.1. *The representational vehicle:*

A representation described or picked out in terms of its physical characteristics is called the representational vehicle. Shea notes that “vehicle realism” (Shea 2018, 37) is a basic assumption of many accounts of representation (including his own) in cognitive science. Vehicles are “individual bearers of content picked out in terms of intrinsic processing-relevant non-semantic properties” (Ibid., 39). Content, as will be discussed below, is the informational and semantic property of the vehicles. The processing of mental representations is described in

computational terms as following an algorithm. As Shea writes, “the algorithm must operate over a set of vehicles that can be individuated non-semantically and must follow a series of processing steps that can be specified non-semantically” (Shea 2018, 37). At the same time, the processing needs to be “faithful to” the semantic properties (content) of the representations.

In terms of physical or, more precisely, neurobiological realization of perceptual representations, a representational vehicle is plausibly a distributed set of activated/firing neurons. Thus, Kentridge and Brogaard write: “...representations are generally encoded in a distributed fashion. It is the relative strength (or even relative timing) of responses across multiple cells that encodes information, not the response of individual cells” (Kentridge and Brogaard 2017, 140).

There is an alternative view of representational vehicle that is endorsed by Crane (2014). According to this alternative, the vehicle is the format of representation – eg. linguistic vs imagistic vs auditory – rather than a physically realized entity. For example, an emotion such as pathos can be represented by words, or by a painting, or by music. If vehicle is understood in this sense, then Crane argues that vehicle should be subsumed under phenomenological content, which I will discuss in §2.3. As Crane writes, “...where the phenomenological conception of content is concerned, we should not make the distinction between vehicle and content: it is central to the phenomenology of an experience that what is conveyed to the subject includes its specific vehicle [i.e., format]” (Crane 2014, 252).

2.2.1.2. *The represented target:*

The target is *what* is represented by a representational vehicle. In other words, it is the intentional object – that entity towards which the vehicle is directed. In cognitive science, the target is always a concrete – i.e., spatiotemporal – entity in the environment, for which the representation stands in as a proxy by carrying information about it.

There are different terms used in the literature to refer to the external target of perceptual representation. For instance, Burge uses the term “subject matter” of the representation (Burge 2010, 38). Cohen (2004, 471) and Nanay (2013, 50) call the target “sensory individual”. I prefer to use the term ‘target’, which has been suggested by Neander (2017, 14). This is for two reasons: first, it fits perfectly with the etymology of ‘intentionality’, which is the Latin word *tendere*, where the emphasis is on “mental directedness towards...objects, as if the mind were construed as a mental bow whose arrows could be properly aimed at different targets” (Jacob 2023, §1). Second, intuitively, ‘target’ connotes a concrete mind-independent entity and thus

avoids the confusion that a term like ‘object’ can engender. This is because ‘object’ could refer a physical object or a perceptual object (an object *in* perception)¹⁶.

What is the target of perceptual representation? While it is agreed that the target is a concrete environmental entity, it is debated what kind of entity it is. As Nanay points out, the two main competing views are: ordinary *objects* and spatiotemporal *regions* (Nanay 2013, 51). The view that ordinary objects are the target of representation is the more common of the two. Nanay endorses Matthen’s definition of an object of perception: an ordinary object is a “spatiotemporally confined and continuous entity that can move and take its features with it” (Matthen 2005, 281). This definition has the virtue that it does not rule out non-material objects such as shadows and rainbows as targets.

The alternative view is that spatiotemporal regions in the environment are the targets of perceptual representation. Austen Clark (2000) is a prominent defender of this view. Clark argues that “space-time regions of finite but definite extent” (Clark 2000, 47) are the targets of perception. He continues: “[t]hese regions are physical regions...They must be within ‘signal range’ of the sense organs for physical changes within them to have physical effects on...sense organs” (Ibid., 47).

Many adherents of the projective theory of perceptual experience also hold that spatiotemporal regions are the targets of perceptual projection. For instance, Baldwin writes:

“The basic idea of [the projective] theory is that sense experience incorporates a reference to regions within the subject’s egocentric sensory space – which is just a region of physical space organised from the subject’s ‘point of view’ (or point of hearing, smelling etc.) – and that the sensory quality which identifies the type of an experience is given as ‘projected’ into the region of space referred to in the experience” (Baldwin 1992, 184-185).

Spatiotemporal regions are, of course, not bare regions, but are usually populated by what Armstrong calls “states of affairs” (Armstrong 1997, 1). An atomic state of affairs is an object instantiating properties (i.e., a propertied object) or two objects standing in a spatial or causal relation; a molecular state of affairs is constituted by atomic states of affairs standing in relations to each other.

¹⁶ As will be discussed in more detail in Chapter 8, a perceptual object is different from a perceived object. The *perceived* or *perceptible* object is a mind-independent external object. In contrast, the *perceptual* object is the object as presented; thus, it is an aspect of the content of perceptual representations.

There have been compelling philosophical and empirical reasons to consider spatiotemporal regions as the targets of perceptual representation. First, a major motivation is the consideration of non-visual perception. The intuition that ordinary macro-objects are the targets of perception is influenced by the visuo-centric model of perception. However, perception in other sensory modalities is not necessarily of macro-objects. Auditory perception is a significant example. It is extremely plausible that the targets of auditory representation are not the vibrating objects that generate soundwaves, but the soundwaves themselves. And soundwaves are states of affairs in spatiotemporal regions. For instance, Nudds argues that “...it is plausible that our experience of sounds represents patterns or structures of frequency components instantiated by the sound waves that are detected by the ears” (Nudds 2009, 75).

Second, even for visual perception there is growing recognition in the empirical literature on change detection and attention (Rensink 2000, Rensink 2002) that organisms perceive dynamic scenes rather than objects. Rensink points out that organisms need to perceive not just completed change but dynamic change. Whereas the perception of complete change involves seeing that an object or scene has changed, “the detection of dynamic change refers to the perception of the transformation itself...This suggests that the spatiotemporal continuity of the external entity may be reflected in the spatiotemporal continuity of the internal representation” (Rensink 2002, 249).

Third, the above insight on change detection is corroborated by the evolutionary perspective on perception. Perception has evolved to ensure survival in a dynamic and often dangerous environment. Perception does this by guiding actions in a dynamic environment, which requires perception of the scene in the first instance, and perception of objects secondarily.

Fourth, considering the spatiotemporal region as the target aids in a better account of hallucinations. Suppose one hallucinates a pink elephant in a room. Here the target would be the entire room, which is perceived inaccurately as having a property it actually does not have – pink elephantness. Thus, hallucinations can be interpreted as a kind of illusion, where the target is perceived in way it is actually not¹⁷.

Owing to the force of the above reasons, I favour the view that spatiotemporal regions are the primary targets of perceptual representation.

¹⁷ More on the illusionist theory of hallucination in chapter 9 (§9.3.2).

2.2.1.3. *The representational content:*

Content is the information-cum-semantic property of the representation. It is the information that the representational vehicle carries about the target, and in virtue of which the representation can act as a stand-in for the target. It is also the semantic property which constitutes or determines the evaluation conditions of the representation. These evaluation conditions are mostly considered to be accuracy or truth conditions.

Content is the essential representational and intentional property, since it is of the essence or nature of representations to have content. It is on account of having content that representations are about (are directed towards) their targets.

I will discuss representational content in more detail in §2.3. For now, I want to bring a point of order on the phrasing of representation and content. The phrase “R represents x” is ambiguous as to whether the variable x is the content or the target of the representation R. Often, the phrase is used such that x refers to the content. However, more accurately, a representation represents a target, in virtue of having content. Therefore, in what follows, when using the phrase “R represents x” I will always use “x” to refer to the target. Also, I will use the following bracket symbols in my thesis: <C> for content; {T} for target.

2.2.1.4. *The representational relation and the representation-grounding relation:*

The relation between the representation and the target is the representational relation. This relation holds on account of the representation having content. Von Eckardt explains that the holding of the representational relation depends on other “*ground* relations, such as similarity or causality, in virtue of which...semantic [i.e., representational] relations hold” (Von Eckardt 2012, 32). The ground relations are usually naturalistic relations. Ground relations are also known as content-constituting relations.

Most naturalistic representationalist theories take the representational relation to be the tracking relation between a representation and its target. Since they are also reductive, they argue that the tracking relation is reducible to naturalistic elements such as causation, biological functions and psychological mechanisms¹⁸. In contrast, non-reductive representationalist theories are non-relational to the extent that they do not posit a representational relation in their analysis of representation. Rather, they come in two varieties – one variety focuses only on content-constituting relations, not to environmental targets, but to abstract entities such as universals

¹⁸ More on the naturalistic theories and their ground relations in Chapter 3.

or propositions; the other variety analyses content as a non-relational property of the representation and/or perceiving subject¹⁹.

2.2.1.5. *The consumer of the representation:*

A mental representation is always for a user or consumer of the representation. Since reductive-naturalistic theories are also subpersonal-level theories to the extent that they functionally analyse the perceptual system, they posit that the consumer of a perceptual representation is a component of the system or is the entire system itself. For instance, Millikan (1984, 2004) explicitly makes a consumer sub-system a component of her theory of representation. In fact, the consumer is so central to Millikan's theory that she often labels it "consumer semantics" (Millikan 2013a, 38). Others like Shea (2018), while not explicitly making space for a consumer sub-system in their theories, refer to downstream exploitation or use of representations. Nonreductive theories, because they are mostly personal-level theories, posit the subject himself as the consumer of the representation.

2.2.2. Theoretical Components of Representationalism

Cummins (1996) argues that a complete theory of mental representation involves three sub-theories as components: (i) A theory of representational content; (ii) A theory of target fixation; and (iii) A theory of application of a representation to a target (Cummins 1996, 20).

A theory of representational content must explain "...what it is for something to be representation, and what it is for a given representation to have a particular content" (Cummins 1996, 20). A theory of target fixation must explain the representation's *function* to represent a particular target. This is because something is eligible or fixed as the target of representation based on the function of representation-producers (and consequently, representations). Finally, a theory of application must explain how a represent is *applied* to a target; that is, how a representation is directed towards its target on account of having content. In other words, application is intentional directedness.

Now, most theories of mental representation – especially naturalistic theories, which I will discuss in Chapter 3 and 4 – combine the sub-theories of representational content and target fixation. This is because the fact that a representation has content is partly accounted for by the fact that representations (or representation-producers) have the biological function of

¹⁹ More on non-reductive representationalist theories in Chapter 6.

representing a target. Thus, I will treat them as one theory – the theory of representational content.

Therefore, in effect, a theory of mental representation includes two sub-theories – a theory of content and a theory of representational application/directedness:

2.2.2.1. A theory of content:

This addresses what it is for a representation to have content, and how a representation has content. That is, it explains the constitution of content. Content, to recall, is the information about the target or the way in which the target is presented, and also the semantic property of the representation. A theory of content is also known as *psychosemantics*.

Many theories of content posit relations as partly constitutive of content. I will, hereafter, refer to such relations as *content-constituting relations*. For instance, reductive naturalistic theories posit content-constituting naturalistic relations to external environmental entities. Some nonreductive theories posit acquaintance to mind-independent abstract properties (i.e., Platonic universals).

Some representationalist theories posit relations to abstract contents themselves (interpreted as propositions). For instance, Byrne – a reductive representationalist – writes that in perceptual experiences, subjects stand in the “exing relation” (for experiencing) to propositions (Byrne 2009, 437). Similarly, Chalmers – a nonreductive representationalist – writes that subjects stand in “content relations” to propositional contents (Chalmers 2006, 51). I will follow Chalmers and, hereafter, call such relations *content relations*. Both content-constituting and content relations are the representation-grounding relations as discussed above in §2.2.1.4. Some representationalists – for instance, Crane (2006, 2009) – do not endorse any relations in their analysis of the constitution of content.

2.2.2.2. A theory of application/directedness:

This addresses how the representation is applied to or directed at its target in virtue of having content. In other words, a theory of application demonstrates how content grounds representation. The rough model for this is Frege’s view that sense (i.e., content) determines the referent of a linguistic expression (Frege 1892/1980). This is only a rough model because, for Frege, the same content always determines the same referent. In contrast, a theory of application only explains the directedness towards the target, not the target itself.

Pautz argues that the directedness of perceptual states is an essential feature of perceptual states, which he describes as:

“Essential External Directedness. For many experience-types, a correct definition of them will use certain spatial terms. For instance, a correct definition of what it is to have the orange-experience will somehow include spatial terms such as round and moving to the right. It will have the form: to have the orange-experience is to [... round ... moves to the right]” (Pautz 2021, 74).

Most theories interpret the representational application/directedness in terms of a relation between a representation and its targets. For instance, reductive-naturalistic theories account for application in terms of the tracking relation. I will call the representational directedness relation the *Representational External Directedness (RED) relation*, in order to distinguish it from the relations posited by a theory of content – content-constituting and content relations. The two sub-theories, however, are not equal. The theory of content has primacy over the theory of application. This is because representationalism holds that intentional/representational directedness is in virtue of having content.

2.3. Content

Representational content is the *essential* representational property – what it is to be a representation is defined in terms of having content. It is on account of having content that representations are intentionally directed towards their targets. Thus, Crane writes that “intentional contents are...what constitute something’s being a representation: it is in virtue of the fact that a mental state has an intentional content that it represents what it does” (Crane 1998b, §2).

Content has two aspects or faces – the informational and the semantic. It is the information that the representational vehicle carries about the target, and in virtue of which the vehicle can act as a stand-in for the target. It is also the semantic property which determines the evaluation conditions of the representation. The informational face points outwards since the information carried by the representation is about the target, and the semantic face points inwards since the evaluation conditions apply to the representation.

2.3.1. Informational/Presentational Content

The information carried is often explained in terms of the *way of representing* the target. Alternative descriptions include *mode of presentation* of the target or *how* the target is represented. Crane, for example, writes that

“The idea of representation itself implies representation in a particular way: in representing something in language or in pictures one has to choose some particular way of representing it. The particular way in which the intentional object [i.e., target] is represented is what I call the content of the state” (Crane 2009, 477).

The way of representing the target is the information about the target since it implies capturing certain aspects of the target to the exclusion of others. Therefore, Crane calls this the aspectual side of content: “[r]epresentations can have the same objects but differ in the aspects under which they represent these objects... When a representation represents something under some particular aspect, it inevitably excludes other aspects” (Crane 2013, 97). I will henceforth refer to informational content as ‘presentational’ content.

Similarly, according to Burge, content is the “way of representing” (Burge 2010, 37) or “mode of representation” (Ibid., 37) or “mode of presentation” (Ibid., 385) of the target. Schellenberg, follows suit, and describes content as “*de re* modes of presentation” (Schellenberg 2018, 87). The *de re* qualification is to emphasize that content is constituted by concrete particulars, rather than general entities. As she puts it, “[u]nderstanding perceptual content as constituted by *de re* modes of presentation recognizes that representing a particular is not independent of singling out the particular that is the referent of the sense” (Ibid., 87).

Presentational content is sometimes known as *Fregean content* (Schulte 2023, 5) or *neo-Fregean content* (Shea 2018, 13) because of its similarity to Frege’s notion of sense, as opposed to reference (Frege 1892/1980). According to Frege, the meaning of a linguistic expression has two aspects – reference and sense (analogous to denotation and connotation). Sense is the mode of presentation of the referent: “...there being connected with a sign (name, combination of words, letter), besides that to which the sign refers, which may be called the reference of the sign, also what I should like to call the sense of the sign, wherein the mode of presentation is contained” (Frege 1892/1980, 37). Sense is an abstract (non-spatiotemporal) meaning shareable between speakers and listeners. However, the analogy between presentational content and Frege’s senses is not perfect. For one, not everyone agrees that presentational content is abstract (more on this in §2.3.4).

Presentational content is also sometimes dubbed “cognitive content” (Egan 2014, 124; Schulte 2023, 6) because *how* things are given to the subject is crucial to the cognitive role representations play in the subject’s perceptual system. More specifically, the presentational content plays a role in the performance of the cognitive task, which in the case of perception is perceiving the environment and subsequently acting in the environment.

Given that presentational content involves the way in which the target is represented, that way is given in perception as the perceived object. Thus, Cussins writes that “[t]he right specification of *how* [i.e, mode of presentation] is also a specification of *what* is available in a cognition” (Cussins 2003b, 149). Therefore, Crane also calls the presentational content as “the phenomenological conception of content” (Crane 2014, 247) and explains it as “what is phenomenologically given to the subject” (Ibid., 247) in a perceptual experience.

2.3.1.1. Presentational content in unconscious perception

However, presentational content is not the same as phenomenological content, because there can be presentational content even in unconscious or nonphenomenal perceptual representation. Unconscious perception is perception that is neither phenomenally conscious nor access conscious, as per the distinction made by Block (1995). An example of unconscious perception is blindsight. Patients who have suffered damage or lesions to the primary visual cortex (V1) have blindspots – regions in the visual field where they cannot consciously perceive stimuli such as movement or objects. Yet, in studies where they were asked to guess the type of object or movement, it was found that they could guess with a probability of higher than chance. The best explanation of this finding is that blindsight patients can perceive, but unconsciously. Moreover, “[i]ndividuals with blindsight can even discriminate simple shapes..., react to emotional expression..., and...avoid obstacles while walking...” (Prinz 2015, 376). The case of blindsight implies that there is information carried even in the absence of conscious perception. To recall, for a representation to carry information is for it to have presentational content. Consequently, there is unconscious presentational content in such cases.

There is ample evidence for unconscious perception even in normal non-pathological cases. For instance, Burge (2010) argues “...there is considerable evidence that individuals, not merely subsystems, have unconscious perceptual states. So there is reason to doubt that consciousness is constitutive either to the individual/subsystem distinction or to perception” (Burge 2010, 374). Nonetheless, Burge writes that content of perceptual states is presentational content (Ibid., 37).

Unconscious perceptual representation usually occurs at early stages of the perceptual system – for example, the early visual processing described by Marr (1982). Marr argues that vision involves three stages – the primal sketch, the 2½ dimensional (2½-D) sketch, and the 3-D model representation (Marr 1982, 37).




Figure 2.1: Marr's theory of the stages of visual representation (Gomes 2000)

Each of the stages involve representations. The primal sketch stage involves representing “intensity changes,...local geometrical structure, and...illumination effects like light sources, highlights, and transparency distribution and organization” (Marr 1982, 42). These are the targets of the first stage, which are represented in a 2-D way (Ibid., 37) – this is the presentational content. It is content because there is information carried by the representations at this stage, and information carried just is presentational content.

The next 2½-D sketch stage involves making “explicit the orientation and rough depth of the visible surfaces, and contours of discontinuities...in a viewer-centred coordinate frame” (Marr 1982, 37). But they are represented in 2½ dimensions. Thus, the content of representations at this stage have 2½-D properties. Both the primal and 2½-D sketches are retinocentric, which means that “...the spatial relations represented refer to two-dimensional relations on the viewer's retina, not three-dimensional relations relative to the viewer in the world around him...” (Ibid., 42).

This implies that the contents of these stages are plausibly unconscious since 2-D and 2½-D contents are neither phenomenally experienced nor available for report; only 3-D content is. Therefore, Marr's theory implies that there is a presentational content at lower levels of

perceptual representation which is non-phenomenologically given to the perceptual system. I will call this non-phenomenal content.

2.3.1.2. Phenomenal content

The phenomenological counterpart of nonphenomenal presentational content is phenomenal content. This is the phenomenally conscious presentation of targets. Phenomenal content is the same as phenomenal qualities (such as colours, smells, tastes, and so on) experienced or the phenomenal character of perceptual experiences²⁰. Thus, Kriegel writes:

“Call the sort of representational content that features only...phenomenal properties – *phenomenal content*. I submit the following thesis: For any experience *E*, the phenomenal character of *E* is one and the same as the phenomenal content *E* carries. Phenomenal character is thus a species of representational content” (Kriegel 2002, 180).

Thus, phenomenal content = phenomenal qualities. Phenomenal character/consciousness is the subjective givenness of phenomenal content or qualities. In other words, it is subjective character + qualitative character.

2.3.1.3. Stages of presentational content

As Marr’s theory of visual representation attests, perception is not a one-off affair. It is accretionary and occurs in stages. Correspondingly, there are different stages of presentational content. In Marr’s theory, these are 2D content, 2½-D content and then 3D content. Thus, Pylyshyn (2001) writes that it is an important assumption of visual psychology that

“...our visual representation of a scene is not arrived at in one step, but rather is built up incrementally...Theoretical analyses (e.g. Tsotsos, 1988; Ullman, 1984) have provided good reasons for believing that some relational properties that hold between visual elements, such as the property of being inside or on the same contour, must be acquired serially by scanning a display” (Pylyshyn 2001, 135).

Furthermore, based on empirical studies on attention-modulated perception, it is claimed that there are two stages of perceptual content – pre-attended and post-attended perceptual content (Rensink 2000)²¹. Historically, Husserl made this point with his distinction between noesis and

²⁰ Chalmers (2006) defines phenomenal content in a reverse way: “A phenomenal content of a perceptual experience is a representational content that is determined by the experience’s phenomenal character. More precisely: a representational content *C* of a perceptual experience *E* is a phenomenal content if and only if necessarily, any experience with the phenomenal character of *E* has representational content *C*” (Chalmers 2006, 50-51).

²¹ More on attention-modulated perception in chapter 8.

noema, which are two stages of content (*Cartesian Meditations* §17, 40). Similarly, The Nyāya school of classical Indian philosophy distinguished between indeterminate (*nirvikalpaka*) and determinate (*savikalpaka*) stages of perception (Chadha 2024)²².

2.3.1.4. Content and mode/manner of representation

Crane distinguishes between the intentional (presentational) content and intentional *mode*. By intentional mode, he means the attitudes such as believing, desiring, seeing, hearing, etc (Crane 2009, 477). Chalmers calls this aspect of representations the “manner of representing” (Chalmers 2004, 155). He writes that the manner “...involves a mental characterization of the state of representing...For example, one can represent a content perceptually, and one can represent a content doxastically...At a more fine-grained level, one can represent a content visually or auditorily” (Ibid., 155). Lyons calls it “attitude” (Lyons 1995, 3).

As an example of the distinction between content and mode/manner, consider seeing a Necker cube:




Figure 2.2: Necker cube

The target is the cube. There is one mode/manner of representing, which is visually perceiving or seeing. But there are two ways of visually representing (corresponding to two presentational contents) – one in which the shaded region seems to face forward, and another in which it seems to face backward/downward.

Not all philosophers, however, make a distinction between the intentional content and mode. Most would include the mode within content, since mode is just another *way* – albeit a coarse-grained way – of representing the target. Thus, Shea uses the term ‘content’ “...for the full representational import of a state, so as to include a specification of its mode of representing” (Shea 2018, 178). This inclusionist view was also held by Husserl, who included within content

²² More on the Nyāya school of Indian philosophy in chapter 7.

(noesis and noema) the “thetic character” (Smith and McIntyre 1982, 132), which was his term for mode. I will adopt the inclusionist view going ahead in the thesis.

2.3.2. Semantic Content

The second face of content is the semantic face. Content is the semantic property of the representation, which implies that it is associated with evaluation or satisfaction conditions. More specifically, perceptual representations are associated with accuracy or veridicality conditions – conditions that the world needs to satisfy in order for the perceptual representation to be evaluated as accurate/veridical or inaccurate/non-veridical²³. Thus, Schellenberg writes that “[a]ccuracy conditions need to be distinguished both from [presentational] content and from the truthmaker of that content. The accuracy conditions of perceptual content specify the way the environment of a perceiver would have to be for the content of her perceptual state to be accurate” (Schellenberg 2018, 92-93)²⁴.

For example, Siegel writes that

“...the content of an experience is given by the conditions under which it is accurate. For instance, if an object looks fish-shaped and orange and looks to be at location L, then...the experience is accurate only if there is something fish-shaped and orange at location L, and the contents of the experience include that there is something fish-shaped and orange at location L” (Siegel 2006, 361).

Formulating the above picture in terms of the model of truth bearers/makers/conditions, we can say that the representation is the accuracy bearer, the environmental states of affairs are the accuracy maker, and the content is the accuracy condition. The formulation of accuracy conditions necessarily refers to environmental targets (the accuracy makers). Hence semantic content is also known as *referential content* (Neander 2017, 15; Shea 2018, 13). It should be noted that the accuracy conditions of a mental state include a specification of the mode of presentation of the target. For example, suppose I have the belief ‘George Eliot wrote *Middlemarch*’, and I do not know that George Eliot was the pseudonym of Mary Ann Evans. The truthmaker of this belief is the woman Mary Ann Evans who went by the pseudonym George Eliot. But the truth condition includes the concept <George Eliot> and *not* <Mary Ann

²³ There could be various evaluation conditions for mental representations: truth conditions for beliefs, accuracy conditions for perceptions, fulfillment conditions for desires, and so on.

²⁴ A minor question of detail here is: what is the accuracy-bearer – the presentational content (as Schellenberg says) or the representation? I incline towards the latter. However, taking a stand on this question is not crucial to my thesis.

Evans>. Similarly, Shea holds that “...a correctness condition like *the door is open* and a satisfaction condition like *open the door* are different contents” (Shea 2018, 178).

Further, accuracy conditions are thought to be propositions; hence, semantic content is also known as *propositional content* (Crane 2014, 241-242). Thus, Siegel writes about semantic content that “[u]sually...contents are thought of as some kind of proposition – an abstract object that is the kind of thing that can be true or false” (Siegel 2024, §2). The standard conception of propositions is that they are “...the sharable objects of the attitudes and the primary bearers of truth and falsity” (McGrath and Frank 2024, §1)²⁵. Both these features derive from Frege’s discussion of the sense of a linguistic expression (Frege 1892/1980). To recall, the sense of a linguistic expression is that aspect of meaning which is abstract, shareable and determines the referent. According to Frege, the sense of a sentence is its “thought”, by which he meant the proposition.

Propositions are posited as the meanings/content of sentences and mental states. Focussing on mental states, since propositions are alleged to be the primary bearers of truth/falsity, the implication is that propositions are the primary representations of states of affairs in the world. Mental states are representations only derivatively – in virtue of standing in a relation to propositions – for example, the “content relation” (Chalmers 2006, 51) or the “exing relation” (Byrne 2009, 437). Moreover, since propositions are also entities that are abstract and shareable, most standard conceptions of propositions subscribe to a Platonism about propositions – they exist independently of spacetime and of minds (McGrath and Frank 2024, §7.1). Thus, “[t]he proposition that there are rocks,...<there are rocks>, does not entail the existence of any beings that have or are capable of having mental states” (Ibid., §7.1). Three of the most popular precisifications of Platonic propositions are Russellian/singular propositions, Fregean propositions and propositions as sets of possible worlds (King 2017).

It should be clarified that the presentational/semantic content distinction is not the narrow/wide content distinction. The latter is about the individuation of content – narrow content is entirely individuated by the subject’s intrinsic properties, while wide content is partly individuated by environmental properties and states of affairs (more on this in §2.3.6). While semantic content is usually wide content, presentational content could be either narrow or wide.

²⁵ King (2017) adds two other characteristics to the standard conception of propositions: “[b]elievers in propositions can say that the proposition expressed by a sentence is the piece of information the sentence expresses...Propositions are also alleged to possess modal characteristics like being necessary and possible” (King 2017, §1).

2.3.3. Relationship Between Presentational and Semantic Content

There are two views about the relationship between presentational and semantic content. According to the first, there is no relationship. This is because mental representations (perceptual or non-perceptual) have only one kind of content – semantic content. This view holds that contents are identical to evaluation conditions; or alternatively contents are given by evaluation conditions. Thus, for instance, Siegel claims that “contents are a kind of accuracy condition...[T]he contents of experiences are conditions under which the experience is accurate” (Siegel 2010, 30). Similarly, “...the contents of an experience are given by the conditions under which it is accurate” (Siegel 2024, §2).

Considering content to be only semantic content is the default position in perceptual psychology and reductive-naturalistic representationalist theories (Neander 2017, 15). Many philosophers – such as Byrne (2005) and McDowell (1996) – who argue that perceptual content is always and only propositional also fall into this camp. Byrne, for instance, writes that “...perceiving is very much like a traditional propositional attitude, such as believing or intending...When one has a perceptual experience, one bears the perception relation to a certain proposition *p*” (Byrne 2005, 245).

A second view is that perceptual representation has both kinds of content, and presentational content determines semantic content (Schulte 2023, 7). In other words, the mode of presentation of the target determines the evaluation conditions of the perceptual representation. This view is being adopted in perceptual psychology and cognitive science as well. Thus, Rescorla writes that “...explanation within perceptual psychology must sometimes consider mode of presentation and not simply denotation” (Rescorla 2020, 582).

Burge is a representative of the view that presentational content determines semantic content. He writes that “[r]epresentational content is a perspectival way of representing at the finest explanatory grain relevant both to determining psychological or linguistic kind and to determining veridicality conditions” (Burge 2010, 38). Thus, according to Burge, presentational content has three functions. First, “it is way of representing, or a perspective” (Ibid., 37-38). Second, it helps type-individuate a token mental state “...as being an instance of a certain psychological kind of perceptual state, thought event, belief, or capacity” (Ibid., 39). And third, content determines or constitutes the evaluation conditions (veridicality or accuracy conditions) of representational states.

2.3.4. Ontological Nature of Content

As noted above, semantic content is often interpreted as a proposition. And propositions are interpreted in the Platonic sense as being abstract entities. Therefore, it is widely accepted that semantic content is an abstract (i.e., non-spatiotemporal) proposition, which is the primary bearer of truth/veridicality.

2.3.4.1. Abstract presentational content

The issue with presentational content is more vexed. To recall, presentational content is also known as *Fregean* content due its analogy with Frege's notion of sense. For Frege, sense was an abstract entity. Therefore, many philosophers hold that, by analogy, presentational content is an abstract entity – i.e., proposition – as well. For instance, Burge writes about presentational contents that they "...are abstract entities" (Burge 2010, 42). Similarly, Pautz distinguishes between presentational and semantic content – which he calls "appears-looks" content and "accuracy" content, respectively (Pautz 2009, 484) – and writes about both that they "...are abstract objects which are true or false" (Ibid., 484).

However, considering presentational content as abstract poses a problem for those representationalists who are also intentionalists – essential representationalists, in Papineau's terminology. To recall, intentionalism claims that phenomenal character of perceptual states either supervenes on (weak intentionalism) or is identical to (strong intentionalism) to representational content. The problem stems from the mismatch between the ontological natures of contents and phenomenal character. Therefore, I will call this objection the *content-character mismatch objection*.

2.3.4.2. Content-character mismatch objection

Phenomenal character is the property of a perceptual state in virtue of which there is something it is like for the subject to token the state. It has two aspects: first, qualitative character – the 'something it is like' aspect; for example, phenomenal qualities such as red, vanilla-smell, and so on. Second, subjective character – 'givenness 'for the subject (Kriegel 2009, 1). Phenomenal character, as Kriegel (2011) and Papineau (2021) have argued, is a concrete property. This is because it is instantiated in space-time by a concrete particular – the concrete particular is the representation for Kriegel (2011, 153), and the subject for Papineau (2021, 73-74). Given this, the content-character mismatch objection to intentionalism is that an abstract entity – abstract

content – is not the right sort of entity to explain and account for a concrete fact – the instantiation of phenomenal character.

Kriegel glosses this objection in terms of explanation. He argues that if we allow abstract content to account for concrete phenomenal character, it will violate the principle of “the *explanatory closure of the realm of concreta*” (Kriegel 2011, 141). This principle states that,

“every aspect of the realm of concreta can be fully explained exclusively by appeal to what goes on inside the realm of concreta. On the assumption that concreta are characterized by their spatio-temporal existence..., the principle might be put more precisely as follows: every spatio-temporal fact can be fully explained by spatio-temporal facts” (Kriegel 2011, 141).

Kriegel concedes that instantiated abstracta such as instantiated universals may not fall foul of the principle on a weaker understanding of concreteness. However, many intentionalists – such as Dretske (2003) – accept uninstantiated abstract properties in the case of hallucinations, which would violate the principle. For instance, Dretske writes that “[i]n hallucinating pink rats we are aware of...the properties pink and rat-shaped that something is represented as having...but we are not aware of any object that has these properties...We are aware of pure universals, uninstantiated properties” (Dretske 2003, 73).

Papineau presents the content-character mismatch objection in terms of causation. He first notes that the instantiation of phenomenal character is a concrete “here-and-now” fact (Papineau 2021, 65), where “concrete facts are constituted by some spatio-temporal particular (or particulars) instantiating some first-order property (or relation)” (Ibid., 66). He clarifies that relations to uninstantiated properties – as in Dretske’s hallucination account – do not count as instantiation, or concretisation in any other way, since “if I am mentally related to some property, but without instantiating it, then the resulting relational fact is by no means local, but extended into whatever distal realm the property in question inhabits” (Ibid., 66).

His argument against intentionalism is based on the causal principle: “[O]nly concrete facts can function as causes or effects” (Papineau 2021, 71). His argument is as follows:

“(1) Instantiations of conscious sensory properties constitute concrete facts with causes and effects.

(2) ...[R]epresentational properties constitute abstract facts that cannot feature as causes or effects.

[Therefore]

(3) Conscious sensory properties are not representational properties” (Papineau 2021, 72).

Although Papineau's argument targets strong intentionalism specifically, it generalizes to weak intentionalism as well – since the causal properties of phenomenal facts cannot supervene on non-causal representational facts. The moral of the content-character mismatch objection is therefore that, per intentionalism, if content must account for concrete phenomenal character, then content must be concrete in the first place.

2.3.4.3. Concrete presentational content

Owing to the force of the content-character mismatch objection, some representationalist-intentionalists such as Crane (2014) and Schellenberg (2018) argue that presentational content is a concrete entity – a complex of concretely instantiated properties, for example.

Crane (2014) draws on Frege's notion of an "idea" as opposed to the sense of a linguistic expression to argue for concrete presentational content. In the same seminal paper – "On Sense and Reference" – where he introduces the notion of 'sense', Frege introduces the notion of 'idea':

"The reference and sense of a sign are to be distinguished from the associated idea. If the reference of a sign is an object perceivable by the senses, my idea of it is an internal image...Such an idea is often saturated with feeling; the clarity of its separate parts varies and oscillates...The idea is subjective: one man's idea is not that of another..."

The reference of a proper name is the object itself which we designate by its means; the idea, which we have in that case, is wholly subjective; in between lies the sense, which is indeed no longer subjective like the idea, but is yet not the object itself" (Frege 1892/1980, 38-39).

Thus, while the sense is quasi-objective and abstract, the idea is subjective and concrete. Drawing on Frege's insights, Crane argues that presentational content (which he calls "phenomenological" content) is concrete: "[t]he content in the phenomenological sense is something spatiotemporal, concrete, particular, and specific to its subject. The content in the propositional sense is not" (Crane 2014, 254).

Similarly, Schellenberg argues that perceptual contents are *de re* modes of presentation, which are "...constituted by employing perceptual capacities by means of which we (purport to) single out particulars" (Schellenberg 2018, 85). Since perceptual capacities are concrete capacities, it follows that content is concrete as well. Hence, the qualification of *de re*.

2.3.5. Constitution of Content – Particularism vs Generalism

This subsection discusses an important debate about constitution of content – content particularism vs content generalism. Content particularism holds that the content of perceptual representations is constituted by mind-independent concrete particulars – that is, to targets (Schellenberg 2018, 14). These particulars could be “objects, events, or property-instances in our environment” (Ibid., 15). Particulars are concrete entities that are characterised by their distinct individual essence or particularity (understood in an ontology-lite sense as mentioned in §1.5.2), in addition to a general essence.

Content generalism, in contrast, is the thesis that content is constituted by general entities. These general entities are often interpreted as abstract entities – Platonic universals, property-complexes or propositions²⁶.

2.3.5.1. Content Particularism

Content can be constituted by particulars only if perceptual representations are related to environmental particulars. Thus, there is a relational implication to content particularism. Therefore, Schellenberg formulates content particularism as:

Content Particularism: “A perceptual state M brought about by being perceptually related to the particular α has the property that M ’s content is constituted by α ” (Schellenberg 2018, 58).

Representationalist views that endorse content particularism are also said to be *relational*. As Schellenberg explains, “...the thesis that perception is relational means always that perception is constitutively a matter of a subject being perceptually related to concrete, mind-independent objects, property instances, events, or a combination thereof” (Schellenberg 2018, 75). Representationalist views that posit relations to *contents*, rather than environmental particular entities, are not relational in this sense. The content that is constituted by (relations to) particulars is known as particular content or singular content²⁷.

Schellenberg interprets constitution in terms of grounding (Schellenberg 2018, 16). Grounding, as discussed in chapter 1 (§1.4.4) is a generative explanatory relationship whereby the grounding entity is metaphysically sufficient for the grounded entity. However, Schellenberg

²⁶ However, not all general entities need be abstract. For example, Aristotelian immanent universals are concrete general entities.

²⁷ More precisely, particular content is specific content (Crane 2013, 139). There are two types of specific content: singular and plural content.

notes that any interpretation of constitution is compatible with the above formulation. To reiterate, I will be following the Bakerian interpretation of constitution (chapter 1, §1.4.8) as a synchronic non-causal generative relation whereby the constituter is both metaphysically necessary and sufficient for the constituted, but the constituted is not reducible to the constituter.

Schellenberg clarifies that when she uses the phrase ‘P constitutes C’ she precisely means that ‘P partially constitutes C’²⁸. This is for two reasons. First, most content particularists hold that content is fully constituted by both environmental particulars and by intrinsic factors. Second, and more importantly, most content particularists are not pure particularists²⁹. They hold that content has both particular and general elements (Schellenberg 2018, fn1, 13-14). For example, Schellenberg (2018) argues that perception fundamentally involves the employment of perceptual capacities. Every perceptual experience is a product of the employment of perceptual capacities, and has two levels of presentational content – content type and token content. Content types “are constituted by the perceptual capacities employed [whereas] token contents...are constituted both by the perceptual capacities employed and the particulars (if any) thereby singled out” (Schellenberg 2018, 87-88).

Only token contents are properly particular contents. Content types are not particular because they are constituted by perceptual capacities which are general capacities. Perceptual capacities are general because “...the same perceptual capacity can be employed in many different environments” (Schellenberg 2018, 66). However, Schellenberg argues that they are not fully general either. Rather, content types are “potentially particularized contents” (Ibid., 92). That is, they can become particular contents when there are particulars in the environment that are singled out. She also describes content types as content schema (Ibid., 98), which the token contents actualize. Thus, content types are quasi-general contents³⁰.

Let us return to the formulation of content particularism – the content of perceptual state M, brought about by being perceptually related to particular α , is constituted by α . What the content-constituting perceptual relations are depends on the specific representationalist theory.

²⁸ Schellenberg interprets constitution as grounding (Schellenberg 2018, 16). But she says that her formulation of content particularism/generalism is neutral towards other interpretations of constitution. To reiterate, I will be using Baker’s analysis of constitution as ‘unity without identity’ (§ 1.4.9) in my arguments and formulations.

²⁹ An example of a pure particularist view is naïve realism.

³⁰ Burge also argues that content is structured by both particular (“singular applications”) and general (“perceptual attributives”) constituents (Burge 2010, 380-381).

Schellenberg, for instance, accepts causal relations (Schellenberg 2018, 15) and spatiotemporal relations (Ibid., 32) as perceptual relations. It should be noted that the relation to particulars (targets) is a content-*constituting* relation, i.e., a relation that constitutes content, not merely causes representations to have content. Of course, causation may itself be a constitutive relation – for example in the case of naturalistic representationalism.

An example of particular (singular) content is singular or Russellian propositions. These are “...propositions that are about a particular individual in virtue of having that individual as a direct constituent” (Fitch and Nelson 2024). Such contents are also known as “Russellian contents” (Chalmers 2004, 167). Russellian propositions can also have the particular properties of objects as constituents. For example, if I have a visual experience of a red ball of a certain size, its particular (Russellian) content is: $\langle (b_1), (r_1, s_1) \rangle$. The subscripts indicate the particular object and property instances. There are different precisifications of particular content in order to deal with hallucinations – for example, merely “property-involving” particular contents (Chalmers 2004, 167), or “gappy” propositional contents (Tye 2009, 92) or gappy non-propositional contents (Schellenberg 2018, 89).

A key implication of content particularism is that two perceptual experiences having the same phenomenal character could have different particular contents, if their contents are constituted by qualitatively identical but numerically distinct particulars. To illustrate, suppose a subject visually perceives a red ball b_1 of a particular size at t_1 . As per content particularism, the content of his visual representation is constituted by b_1 with its particular properties. So the content is $\langle (b_1), (r_1, s_1) \rangle$. Now, at t_2 the ball b_1 is replaced by exact duplicate – a qualitatively identical ball with similar properties. The content of her visual representation at t_2 , constituted by b_2 with its particular properties, is $\langle (b_2), (r_2, s_2) \rangle$. Thus, the contents of the subject’s visual representations at t_1 and t_2 are different. However, the phenomenal character of her visual representations at t_1 and t_2 is the same, because the balls b_1 and b_2 are qualitatively identical and therefore phenomenologically indistinguishable. Therefore, as per content particularism, there could be difference in content without a corresponding difference in phenomenal character. Content particularism is, therefore, not straightforwardly compatible with intentionalism about the phenomenal character of perceptual states. Intentionalism, to recall, is the thesis that phenomenal character at least supervenes on, and hence covaries, with content³¹.

³¹ Impure content particularists such as Schellenberg work around this problem by arguing that phenomenal character supervenes only on the general elements of content and not the particular elements (Schellenberg 2018, 88).

Particular content, in virtue of being constituted by particulars, is itself characterized by particularity. Schellenberg distinguishes between two kinds of content particularity – phenomenological and relational particularity. She defines phenomenological particularity as:

“Phenomenological Particularity: A mental state manifests phenomenological particularity if and only if it phenomenally seems to the subject that there is a particular present” (Schellenberg 2018, 17).

Phenomenological particularity does not require the presence of particulars in the external environment to which the perceptual state is related. Phenomenological particularity is only at the level of phenomenology. In contrast, relational particularity is defined as:

“Relational Particularity: A subject’s perceptual state M brought about by being perceptually related to the particular α is characterized by relational particularity if and only if M is [partly] constituted by α ” (Schellenberg 2018, 17).

The idea is that a perceptual state has a non-phenomenological kind of particularity which is grounded in two facts: (i) that there is a relation between the perceptual state and an environmental particular, and (ii) that the perceptual content is constituted by the particular. Therefore, this kind of particularity is more accurately *relational-constitutional particularity* of perceptual states. As Schellenberg notes, it is relational-constitutional particularity that is in debate between content particularists and generalists, since both parties can agree upon phenomenological particularity (Schellenberg 2018, 17).

2.3.5.2. Content Generalism:

Content generalism: A perceptual state M brought about by being perceptually related to the particular α has the property that M ’s content is “constituted only by general elements, and – contra particularism – not even in part by α ” (Schellenberg 2018, 16).

Content generalism implies that although a perceptual representation may be related to particulars (by causal, spatiotemporal or any other relation), the relations do not have any content-*constituting* role. Therefore, content is not constituted by the particulars related to.

Of the several proposals for what are the general contents of perception, most are variations on Fregean propositions. Unlike Russellian propositions, Fregean propositions “...are composed of senses, not individuals, and senses are individuated independently of any individual” (Fitch

and Nelson 2024, §1). In other words, the perceptual content (whether presentational or semantic) is composed of abstract modes of presentations.

One prominent version of Fregean general content is the existentially quantified content or existential content. The existentially quantified content is usually understood as an abstract proposition (Tye 2009, 79). For example, if I have a visual experience of a red ball of a certain size, its existential content is <there is a ball such that it is of a certain size and has redness>. Another type of Fregean general content is an intension, which is “a mapping from scenarios to truth values, where scenarios are maximal epistemic possibilities, or centered possible worlds” (Chalmers 2004, 172). A centered world is “a world marked with a designated individual and a designated time (intuitively, these represent the perspective of the subject who [perceives])” (Chalmers 2010, 372). It should be noted that in the above examples, the general elements constituting content are abstract entities.

The main advantage of content generalism is that it can give a uniform account for indistinguishable veridical and non-veridical perceptual experiences (illusions and hallucinations) (Schellenberg 2018, 16). This account adverts to (a) general contents that are usually abstract entities and hence can be associated with even hallucinations, and (b) endorsing intentionalism, according to which phenomenal character supervenes on content. Thus, the veridical and non-veridical experiences are subjectively indistinguishable because they have the same content, which grounds the sameness of phenomenal character.

A consequence of general contents is that perceptual representations are not essentially relational. Content is not constituted either by external particulars or by relations to those particulars. Instead, perceptual representations involve either (i) constitutive relations to abstract properties or property-complexes (i.e., Platonic universals), or (ii) constitutive relations to general contents, or (iii) no relations at all.

Chalmers (2006) and Pautz (2021) endorse option (i). Chalmers argues that we are primitively related (plausibly by acquaintance) to Platonic universals, which then constitutes “Edenic content” (Chalmers 2006, 70). Similarly, Pautz defines representationalism in terms of subjects standing in the “experientially representing” relation to abstract property-complexes (Pautz 2021, 100). Chalmers (2006) and Pautz (2009) also endorses option (ii), whereby subjects bear constitutive relations to *contents*, rather than targets. These relations could be the propositional attitudes or “content relations” (Chalmers 2006, 51); Pautz labels such a relation the “sensory entertaining” relation (Pautz 2009, 494). The idea is that although content is not constituted by

relations, relations to contents are constitutive of representations in that representations have contents by being related to them. Strictly speaking, these views are non-relational in the sense defined in the section above.

I will call the representational views that endorse options (i) and (ii) – relations to abstract universals and to contents – as ‘*con-relational*’ to distinguish them from the genuinely relational views which accept relations to environmental targets. I will reserve the term ‘*non-relational*’ for those representationalist views that endorse option (iii) – no constitutive relations at all, to contents or to targets. According to non-relational representational views, contents are not constituted by any relations, nor are they entities to which we can bear any relation. Instead, they are properties or aspects of either the representation or the subject. There are two versions of non-relational representationalism (Kriegel 2007), distinguished on the basis of what or who possesses the content property.

First, adverbial or aspectual representationalism (Kriegel 2007, 314). According to it, content is an adverbial modification or aspect of the intentional state. Perceptual representation is interpreted as an activity and content is the way of perceiving, akin to adverbial modification of an action (eg., walking *quickly* vs *slowly*). For example, ‘S perceives a red square’ can be paraphrased as ‘S perceives redly (or red-wise) and squarely (or square-wise)’, where redly and squarely are properties of the perceptual state. “More generally, for any x, representing x does not involve constitutively bearing a representation relation to x; what it involves constitutively is representing x-wise” (Kriegel 2008, 84)³². The adverbial content is also known as the “aspect” of the representation (Mendelovici 2018, 198). Crane is an exemplar of the aspectual representationalism view, and he writes that “...perceptual representation is intrinsic to the perceptual state itself, that it represents the world, even if the state is one which can be analysed into its various components [such as vehicles, contents and targets]” (Crane 2006, 137).

The second non-relational view is *hyphenism* (Kriegel 2007, 314). According to it, content is not a property of the representation but of the subject. For example, ‘S perceives a red square’ can be paraphrased as ‘S is perceives-red-square’ where the predicate ‘perceives-red-square’ denotes a monadic property of the subject.

³² Kriegel (2008) further delineates two versions of adverbialism – inferential and phenomenological adverbialism.

2.3.6. Individuation of Content – Externalism vs Internalism

Particularism and generalism about content must be distinguished from the related, but orthogonal, theses of internalism and externalism about content (Schellenberg 2018, 76). Whereas particularism and generalism are theses about the *constitution* – or existence conditions – of content, internalism and externalism are theses about the *individuation* – individuation and identity conditions – of representational content, and consequently of representational states (Rowlands et al. 2020, §3; Schantz 2004, 14; Schellenberg 2018, 76). According to Lowe, “[i]ndividuation...is an ontological relationship between entities: what ‘individuates’ an object...is whatever it is that makes it the single object that it is – whatever it is that makes it one object, distinct from others, and the very object that it is as opposed to any other thing” (Lowe 2003, 75). In other words, individuation is about what determines the individual essence of contents and, consequently, representations.

However, sometimes, externalism and internalism are formulated and described as theses about the nature (i.e., constitution) of content rather than individuation of content – for example, Burge writes of externalism (or anti-individualism as he prefers to call it) as a thesis about the nature of perceptual representations, and also as a necessary condition for *being* in perceptual representational states (Burge 2010, 61-63). I will interpret externalism and internalism as theses about the individuation of content. There are two opposing views about the individuation of content – content externalism and content internalism.

2.3.6.1. Content Externalism

Content externalism – sometimes labelled “anti-individualism” (Burge 2010, 61) – is the thesis that “...at least some mental content is...individuation dependent³³ on the nature of the environment, understood as a collection of circumstances that reside outside the biological boundaries of the individual.” (Rowlands, et al. 2020, §3). It is sometimes formulated in terms of constitution: content externalism is the thesis that content “...at least in part constitutively depend[s] on features of one’s physical or social environment” (Carter et al. 2014, 64). It should be noted that the role of the environment is only *partly* constitutive, because content also depends on intrinsic properties of the subject.

³³ To recall from the discussion in chapter 1 (§1.4) dependence is of two types – existence and identity or individuation dependence. “The general idea of individuation dependence is that the identity of an entity, as an entity of a...kind *s*, depends upon things (objects, properties, facts, etc.) distinct from it” (Rowlands et al., §3).

The dependence of content individuation on the environment is by way of being *related* to the environment. Thus, content externalism, like particularism, entails relations to concrete particulars in the environment. Content partly individuated by relations to the environment is known as *wide or broad content*.

Content externalism is the default position of naturalistic representationalism. In fact, particularism and externalism form a neat package of theses to endorse. Both involve relations to the environment. While particularism says that content is constituted by relations, externalism says that it is individuated by them.

2.3.6.2. Content Internalism

Content internalism is the thesis that representational content is entirely individuated by factors internal or intrinsic to the perceiver. In other words, “the content of mental states is determined only by the internal, non-relational properties of the subject’s mind or brain” (Schantz 2004, 14). Such content is known as *narrow content*.

It should be noted that there is absolutely no role given to the environment and relations to the environment in constituting narrow content. Thus, Brown writes that the narrow content of a mental state “is a content of that state that is *completely* determined by the individual’s intrinsic properties. An intrinsic property of an individual is a property that does not depend at all on the individual’s environment” (Brown 2022, §1; emphasis added). Similarly, Chalmers writes that narrow contents are “contents of the sort that, if they are represented by a subject, they are represented by any intrinsic duplicate of that subject (regardless of the environment)” (Chalmers 2004, 166). However, content internalists accept *causal* relations between representational vehicles and the environment.

The main motivation for narrow content is to account for the phenomenal character of perceptual experiences. Given (a) the assumption that phenomenal properties are determined entirely by intrinsic properties (Chalmers 2004, 165), and (b) the endorsement of intentionalism (phenomenal properties supervene on content), the argument is that “both phenomenal properties and the representational properties they are equivalent to depend only on a subject’s internal state” (Ibid., 166)

Some content internalists argue for two kinds of content: a narrow content and a wide content. Chalmers (2010) is an example. As noted above, Chalmers interprets Fregean content as an abstract mode of presentation/intension, which is “a mapping from scenarios to truth values,

where scenarios are maximal epistemic possibilities, or centered possible worlds” (Chalmers 2004, 172). Fregean content determines the satisfaction conditions of the perceptual representation, and when combined with the environment yields Russellian content, which is “...a structured Russellian proposition involving objects and properties” (Chalmers 2010, 372). Whereas Fregean content is narrow, Russellian content is wide.

“The Fregean content of an experience, along with the environment, determines its Russellian content. For example,...object *O* is precisely the object that meets the condition laid down by the Fregean content: the condition of being the object that appropriately causes the experience...Of course, if one has a phenomenally identical experience that involves a different object *O'*, the Fregean content will be the same, but the Russellian content will involve *O'* instead of *O*” (Chalmers 2010, 373).

In other words, Fregean intensions are the satisfaction conditions that take us from possibilities (or possible worlds) to truth/accuracy values via Russellian propositions.

2.3.7. Conceptual vs Non-conceptual Content

Another important debate about perceptual content – particularly presentational content – is whether it is conceptual or non-conceptual. Following Heck (2000), it is customary to distinguish between two versions of conceptualism and non-conceptualism – the state view and the content view.

According to the state view, conceptualism and non-conceptualism are theses about mental representations or states, rather than about their content. A conceptual state is such that a subject cannot token it without possessing the concepts needed to specify its content. On the contrary, a non-conceptual state is such that a subject need not possess the concepts needed to specify its content in order to token it. In other words, a conceptual state is concept-dependent, and a non-conceptual state is concept-independent (Bermudez 2009, 459).

The state view usually holds that non-perceptual states (beliefs, desires, and so on) are concept-dependent, whereas perceptual states are concept-independent. Both can have the same type of contents (for example, propositional), but the difference between them is that in the case of perceptual states, the subject need not possess the concepts needed to specify content, whereas in the case of non-perceptual states, the subject needs to possess content-specifying concepts.

According to the content version, conceptualism and non-conceptualism are theses about the content – particularly presentational content – of mental states. Specifically, it entails that

“...there is a fundamental difference between the type of content that perceptual experiences can have and the type of content that beliefs and other propositional attitudes can have” (Bermudez 2009, 459) – whereas propositional attitudes have conceptual content, perceptual representations have non-conceptual content. In other words, perceptual content is not structured conceptually. There have been several proposals for non-conceptual perceptual content. For example, Peacocke (1992) has defended the claim that perceptual states have “positioned scenario” content (Peacocke 1992, 64). Similarly, Cussins (2003) argues that perceptual experiential content is structured as “activity trails” (Cussins 2003, 154).

As Bermudez and Cahen (2024, §3) point out, most non-conceptualists prefer the *state* non-conceptualism view for perceptual states. However, Bermudez argues that content non-conceptualism is more fundamental because it grounds state non-conceptualism. To the question of what grounds the content-independence of perceptual states, “[t]he most obvious answer is that the concept-independence of perceptual states is a function of their distinctive type of content” (Bermudez 2009, 460). Moreover, most of the arguments for non-conceptualism about perceptual states are actually arguments in support of the content view (Ibid., 460). The main argument for non-conceptualism about perceptual content is that perceptual states have a richness of grain that far outstrips our conceptual capacities (Evans 1982). For instance, we can perceptually discriminate vastly more shades of colour than we have concepts for. This indicates that the presentational content of perceptual states does not depend on our conceptual cache.

In my thesis, I will follow Bermudez’s claim that both perceptual states and perceptual content are non-conceptual, but non-conceptualism about content is primary (Bermudez 2009, 461).

2.4. The Content Problems

As mentioned in §2.2.2, the theory of content is the primary component of a complete theory of representation. Since content is the essential representational property, representationalist theories stand or fall based on their theory of content. This is true for theories of mental representation, in general, and for theories of perceptual representation (i.e., representationalism), in particular.

Neander writes that theories of content try to answer two distinct, but interrelated questions or problems. The first asks how representations can have content *at all*. In other words, how do

representations have content as opposed to having no content? Neander calls this the “representational status question” (Neander 2017, 5) since the possession of content determines the status of something as a representation. The second question is the “content determinacy question” (Ibid., 5): how do representations get to have their *particular* and *determinate* content; in other words, how does a representation have content <A> as opposed to an indeterminate content or some other determinate content ? By determinate content is meant content that is non-disjunctive and distal.

The two content problems are interrelated for two reasons: firstly, because content is plausibly always determinate³⁴, answering the status question goes hand-in-hand with answering the determinacy question. Secondly, the answers to both questions usually involve employing the same resources; for instance, naturalistic theories of mental content posit naturalistic relations and biological functions, among other components, to answer both the questions. Von Eckardt writes that the combination of the two problems is the central problem for cognitive science:

“Cognitive scientists explain various aspects of human cognition, in part, by positing representations with content. But what is it about these representations that gives them this content? This is the problem of content determination, arguably, *the* foundational problem for cognitive science – to identify a general “ground” for [mental representations], that is, a set of naturalistic (i.e., nonintentional, non-semantic, non-normative) properties and relations possessed by the representation-bearers...that *determine* their semantic properties and relations” (Von Eckardt 2012, 35).

2.4.1. Metaphysical Explanation of the Content Problems

An important characteristic of answers to the representational status and the content determinacy problems is that they are self-avowedly metaphysical explanations of representation, rather than merely causal or teleological explanations. For instance, Shea writes that “[w]hat we are after is an account of what determines the content of a mental representation, determination in the metaphysical sense (what makes it the case that a representation has the content it does?)” (Shea 2018, 9).

I will call this the basic desideratum of a theory of mental representation, which includes an account of content as well as representational directedness:

³⁴ However, Shea (2018) claims that simpler cognitive systems could have indeterminate contents.

Basic desideratum:

A theory of mental representation must provide a metaphysical explanation of how representations are applied to (or directed at) their targets in virtue of having content.

To put it another way, it is a basic desideratum of a theory of representation to metaphysically explain the representational external directedness (RED) relation between a representation and its target that holds in virtue of having content.

2.4.2. The Distality Problem and Other Content Determinacy Problems

The content determinacy problem is about metaphysically explaining how a representation can have a particular determinate content as against a different determinate content. There are several specific versions of the content determinacy problem – the most important are the distality problem, the specificity problem and the disjunction problem.

The following example, from Neander (2017, 221), illustrates the different versions of the determinacy problem. Suppose a toad visually perceives a cricket. This state of affairs involves a causal chain beginning with light hitting the cricket and ending with the excitation of T5-2 neurons in the optic tectum region of the toad's brain. This neuronal excitation is the representational vehicle, and the target is the cricket.




Figure 2.3: The excitation of T5-2 neurons are the representational vehicles, which horizontally represents the insect target; but they could represent it as any of the vertical candidates.

2.4.2.1. Distality problem:

In the example, the main content determinacy problem is explaining how the representation has distal content <cricket> rather than content that figures any of the proximal intermediaries (retinal firings, neural signals, etc) of the causal chain between the cricket and the representation. This problem is, therefore, known as the “distality problem” of content (Neander 1995, 136; Schulte 2018). As Schulte puts it, “the content of perceptual states is *distal*: perceptions do not represent proximal states of affairs involving e.g. the shape of retinal

images or the wavelength of light, but only distal states, involving e.g. the colour of a flower, the shape of a tree, or the size of a house” (Schulte 2018, 349). The distality problem goes by several other names. Since the distality problem concerns the representation latching on to a target along a horizontal chain of causes, it is also known as the “horizontal” problem of content (Godfrey-Smith 1989, 536). Further, since it is also about the representation stopping at a particular node of the causal chain, the distality problem is also known as the “stopping problem” (Neander 2017, 217).

Distality of content pertains to both presentational and semantic content – in the presentational case, a thing is presented as being distal and having distal properties; and in the semantic case the evaluation conditions advert to distal objects and properties. Since a representation is directed at a distal target by account of having distal content, the derivative problem is of explaining how the representation is directed at a distal target as opposed to a proximal one.

It should be noted that the distal targets of representations need not always be at the beginning of the causal chain, as in the above example. The beginning of the causal chain is the light hitting the cricket, but that is not the target of representation. Another example is auditory perception. Suppose a tuning fork is struck and the vibrations cause a sound wave which results in hearing a sound.




Figure 2.4: The auditory representation – activation in the auditory cortex – plausibly represents properties of the sound wave rather than surface properties of the vibrating tuning fork.

Auditory psychology takes auditory representations to represent properties of the sound wave rather than surface properties of vibrating objects –the tuning fork – although the former is causally dependent on the latter and so is more proximal than the latter. According to Matthew Nudds, for instance,

“it is plausible that our experience of sounds represents patterns or structures of frequency components instantiated by the sound waves that are detected by the ears. It follows that an experience of a sound is veridical just in case it is produced by the pattern or structure of frequency components that would normally produce that experience” (Nudds 2009, 75).

More specifically, it is argued that the auditory representations tokened when we experience loudness and pitch represent the intensity and frequency of the sound wave respectively (Plack 2024).

Therefore, representational targets need to be not merely distal, but “suitably distal” (Neander 2017, 229). “Suitably” distal also indicates that the target need not always be at a distance from the perceiver. For example, in the case of tactual perception, the target is often sensory registrations on the skin. The distality problem, then, is about finding a systematic way of ensuring that representational targets are in the distal “goldilocks zone”, neither overly proximal nor overly distal. The latter aspect of the distality problem – eliminating overly distal content and targets – is also known as the “far-out” distality problem (Schulte 2018, 366).

2.4.2.2. Specificity problem:

In the toad/cricket example, the distal target could be represented in different ways – corresponding to different presentational contents. It is ‘cricket’ to us humans, but the toad who does not (and cannot) possess human concepts (or percepts) could represent it as any of the vertical alternatives in the diagram – item in worm-like motion, cricket, worm or just food – or in any other way. This is the specificity problem. Since the target has, arguably, all of the vertical properties – it is a cricket as well as an item in worm-like motion as well as food – the specificity problem is about determining which properties, at a particular distality level, are part of visual content.

2.4.2.3. Disjunction problem:

The disjunction problem comes in two versions (Ryder 2020b, 255): one version asks, in the context of the above example, why the content is not a disjunction of all the vertically mentioned properties and more such: <item in worm-like motion or cricket or food or...>. The second version of the disjunction problem is the problem of error. It is plausible that the toad’s visual representation is tokened by an origami-craft work of a cricket rather than a real cricket. The error-disjunction problem asks why the content of the visual representation in every case is not < organic cricket or origami cricket or clay blob or... > rather than the actually tokened <organic cricket>. A solution to the error-disjunction problem is crucial for any theory of

representation because if such disjunctive contents are not ruled out, then there would be no misrepresentation. This can be fatal for a theory of representation, since many hold that the capacity to misrepresent is essential to being a representation (Dretske 1986).

The specificity and the disjunction problems are also known as the “vertical” problems of content (Godfrey-Smith 1989, 536), since they concern the choice of a determinate content among other vertical alternatives.

2.4.2.4. Problem of misrepresentation

The problem of error-disjunction is one version of the more general problem of misrepresentation. This is the problem of explaining how and why representations sometimes misrepresent or represent inaccurately. This is, therefore, the problem of explaining non-veridical perceptions. Perceptual illusions and hallucinations are the paradigm cases of non-veridical perceptions (i.e., perceptual misrepresentation).

An illusion is a perceptual state where a target is perceived to have properties that it does not in fact possess. For example, a stick half immersed in water appears to be bent, even though it is not. Although we may know the physical reasons behind the appearance – the refraction of light – the illusion does not disappear. A hallucination, conventionally understood, is a perceptual state where an object (possessing properties) is perceived even though there is no such object in the environment. Although visual hallucinations are the paradigm example, hallucinations in other modalities – especially auditory and haptic hallucinations are more common. Naturalistic representationalism mostly explains misrepresentation as a failure of the perceptual system in performing its biological functions.

The problem of non-veridical perceptions (illusions and hallucinations) is the lot not merely of representationalist theories of perception. Even non-representationalist theories such as naïve realism and adverbialism face it. In its general avatar, the problem is often posed as one of reconciling non-veridical perceptions with the assumption of direct realism. Direct realism is the thesis that we perceive environmental states of affairs directly, without perceiving intermediary entities such as sense data (Crane and French 2021, §1.4). As Crane and French write, the problem of non-veridical perception...

“...is a pervasive and traditional problem about our ordinary conception of perceptual experience. The problem is created by the phenomena of perceptual illusion and hallucination: if these kinds of error are possible, how can perceptual experience be what we ordinarily

understand it to be: something that enables direct perception of the world?” (Crane and French 2021).

This problem is compounded by the phenomenological fact that sometimes non-veridical perceptual experiences can be subjectively indistinguishable from veridical perceptual experiences.

2.4.3. The Primary Problem of Perception

As is evident from the discussion in the above section, there are multiple problems of perception – some are faced by representationalist theories in particular, and some are faced by all theories. But is there a *primary* problem of perception? And if there is, which one is it? Many philosophers consider the theory-neutral problem of reconciling non-veridical perceptions with direct realism as *the* problem of perception. For instance, the *Stanford Encyclopedia of Philosophy* entry on “The Problem of Perception” (Crane and French 2021) is entirely about this problem. Similarly, Smith, in his book *The Problem of Perception* (2002), writes that,

“...when philosophers speak of “the” problem of perception, what they generally have in mind is the question whether we can ever directly perceive the physical world. This issue is a *problem* because of certain arguments [from illusion and hallucination] that have been put forward to the effect that such direct perception is impossible” (Smith 2002, 1).

However, Brunswik – a psychologist – argues that the distality problem is the primary problem:

“The basic fact of perception is distal focusing. No orientation in an organized “world” would be possible without it. Distal focusing is the result of an ecological generalization process on the part of the responding organism. The generalization takes place over the range of the concrete variants in the proximal mediation patterns of the distal variable” (Brunswik 1956, 61).

I follow Brunswik and submit that the *distality problem* is the primary problem of perception. The reasons for this claim are as follows. First, perceptions are always about suitably distal targets. Even non-veridical perceptions are of suitably distal targets. For example, hallucinations are often of distal objects. Thus, the distality problem is common to both veridical and non-veridical perceptions. Any solution to the problem of non-veridical perceptions would need to assume a solution to the distality problem to even take off.

Second, the distality problem is faced by *all* representationalist theories and not merely by causal-informational theories, as is sometimes argued. The traditional understanding is that only causal-informational semantic theories posit relations that are composed of smaller proximal intermediary relations; therefore, only they face the problem of explaining how representations stop at a distal cause rather than a proximal one. Consumer semantic theories – for example, Millikan (2004) – claim that they rely on the structural correspondence relation and the use of representations by consumer sub-systems to constitute content; and these features avoid the distality problem. I will discuss the above relations and the use condition in chapters 3 and 4 in detail. But for now, I want to note that even consumer semantic theories are faced with the distality problem. This is because structural correspondence also can be to proximal states of affairs rather than distal ones. And the ‘use-by-consumers’ condition just pushes the distality problem to the stage of consumers. The claim is that consumers use representations to perform their functions and such use fixes content. But how do the consumers ‘know’ to use the representation such that its content is distal rather than proximal? After all, proximal content is as much suitable for the performance of such functions as is distal content. For example, the tongue snapping system of toads performs its functions even if visual representations have proximal content.

Third, the distality problem is faced by all theories of perception; not merely by representationalist theories. Even an anti-representationalist theory such as naïve realism accepts causal relations between the perceived object and the perceiver (Fish 2021, 126-128), in addition to a *sui generis* perceptual relation known as “acquaintance” (Ibid., 105). Therefore, even naïve realism is faced with the problem of explaining why the perceptual relation is to a distal cause rather than one of the proximal intermediaries.

Fourth, distality of content is a necessary condition for action. In chapter 3 (§3.3.2), I will argue that the evolutionary purpose of perception is action for survival, and this is achieved by perceptual representations having action-oriented content. Action-oriented content is necessarily distal content, since it involves attributing action-affordance properties to distal objects. Thus, a solution to the distality problem is crucial not only to explain perceptual representation, but also the purpose of such representation – action.

Fifth, in the context of the horizontal-vertical content problems discussed above, the horizontal distality problem is more fundamental than the multiple vertical problems. This is for two reasons: (i) To recall, the vertical problems arise at a particular node of the horizontal causal

chain. This means that to even begin answering the vertical problem, the content ascription must stop at a particular node of the chain. In other words, answers to the vertical problem presuppose an answer to the horizontal problem. If a theory fails to answer the horizontal problem, it *a fortiori* will fail to answer the vertical problems.

(ii) A theory of mental representation could arguably live with an unsolved vertical problem, but not with an unsolved horizontal problem. An example is Shea's varitel semantics: it gives an answer to the horizontal problem but is admittedly indeterminate about the vertical specificity problem in the case of simple cognitive systems. Shea argues that a frog's visual representation, for example, is plausibly indeterminate "between representing flying insects and flying nutritious objects. That is a kind of indeterminacy that flows from the limited complexity of the system" (Shea 2018, 157).

Due to the above reasons, I conclude that the distality problem is the primary problem of perception. Also, I will assume that an answer to the distality problem will constitute an answer to the representational status problem of explaining how representations can have content at all. This is because, to reiterate, the solutions to the two problems involve positing the same answers and explanantia. In the rest of the thesis, I will evaluate representationalist theories – both reductive-naturalistic and nonreductive – mainly on the touchstone of their solution to the distality problem; in other words, their account of distal content.

2.5. Conclusion

In this chapter, I expounded and expanded on various strands and debates in representationalism, especially regarding content. I argued that there are two kinds of content presentational and semantic content, and that presentational content is tokened in both phenomenally conscious and unconscious perceptual representations. Further, presentational content determines semantic content (i.e., evaluation conditions). I also noted the content-character mismatch problem for those representationalist-intentionalists who account for concrete phenomenal character in terms of abstract content. Thus, I favour the concretism about content – content is a concrete property of the subject. Further, I argued for the view that targets of representation are spatiotemporal regions rather than distinct objects.

I also listed the various content problems including the problem of misperception or non-veridical perception. I then argued that the distality of content problem is the primary problem,

because even non-veridical perceptions are suitably distal and therefore presupposes a solution to the distality problem. In the rest of the thesis, I will evaluate reductive and nonreductive representationalist theories on the basis of their solution to the distality problem, among other criteria.

3. Reductive-Naturalistic Representationalism: Naturalistic Relations and Functions

In this chapter, I begin my critique of reductive-naturalistic representationalism, focussing on two of the three content-constituting conditions that it posits – naturalistic relations and teleological functions. In §3.1, I describe the basic commitments of reductive-naturalistic representationalism, particularly the tracking relation and the project to ontologically reduce it. In §3.2, I expand on the first content-constituting condition – naturalistic relations to represented targets. Then, in §3.3, I argue that naturalistic relations are necessary for content, based on the assumption that the essence and purpose of perception is to guide actions to ensure fitness and survival. This ultimately requires relations with the environment. §3.4 then moves to the second of the content-constituting elements posited by naturalistic representationalism – teleological functions. I criticize the commonly adopted notion of teleological functions – the selected effects notion – and argue for an alternative notion, fitness-contribution functions. Finally, in §3.5, I highlight why functions are necessary – because they are the source of evaluation conditions (i.e., semantic content) and they underwrite teleological explanations of the representations and representation-producers.

3.1. Reductive-Naturalistic Representationalism

Naturalistic representationalist theories explain the constitution of perceptual content in terms of naturalistic entities – entities that are countenanced by the natural sciences. More specifically, they are ontologically reductive theories and reduce content and representation to naturalistic entities.

3.1.1. Primacy of Theory of Application

To recall from chapter 1, a theory of representation is composed of a theory of content and a theory of application. The former is about how content is constituted and the latter is about how representations are directed at their targets in virtue of having content. Naturalistic theories follow a common strategy of keeping the theory of content subservient to the theory of application. They first explain the application of a representation to its target; that is, the representational external directedness (RED) relation. Then, on the basis of the theory of

application, they explain how representations have content. It should be noted that ‘content’ is used in the naturalization literature mostly in the semantic sense of evaluation or satisfaction conditions (Neander 2017, 15).

Cummins highlights this strategy when he writes that naturalistic representationalist theories are primarily theories of application. He also calls applications of representations as *use* of the representation to represent their targets: “[t]o use a representation is to apply it to a target. Uses, then, are simply applications” (Cummins 1996, 29). Therefore, Cummins labels naturalistic theories as “use theories” (Ibid., 29). Cummins’ usage of the term ‘use’ might elicit confusion since the notion of use in the naturalistic representationalism literature is used more often for a different phenomenon – that of internal causal processing (more on this in chapter 4). Therefore, I will rebrand Cummins’ idea of use theories as *application* theories.

Cummins then argues that content is determined by the application of representations to targets:

“[Application] theories of meaning (representational content) take the meaning of a representation to be determined by what it is or can be applied to. The fundamental idea is simple: in a case of correct [application], content = target... [Application] theories attempt to naturalize content by specifying a sufficient naturalistic condition under which [application] can be supposed to be correct. They then identify the content of [representation] *r* with whatever *r* is applied to when the specified correct [application] condition is satisfied” (Cummins 1996, 29).

Representational application to targets is just the holding of the representational relation between representations and their targets. Naturalistic theories consider this relation to be the *tracking relation*. And the sufficient naturalistic condition that they specify for the tracking relation is a complex condition consisting of three conditions: naturalistic relations, biological functions and representational application formulae.

3.1.2. Tracking Relation

Naturalistic representationalist theories claim that the representational relation that holds between representations and their targets is the tracking relation. That is, representations represent their targets by tracking them. And it is in virtue of the tracking relation that they have content.

What is it for a representation *R* to track a target *T*? The intuitive general idea is that *R* tracks *T* “...if the state in which [*R*] is somehow *depends systemically* upon the state in which [*T*] is”

(Kriegel 2011, 70; emphasis mine). The theory of intentionality that is based on the tracking relation is known as the tracking theory of intentionality (TTI). Cutter and Tye (2011) explain the TTI as follows: “Tokens of a state [R] in an individual x represent that [T] in virtue of the fact that: under optimal conditions, x tokens [R] iff [T], and *because* [T]” (Cutter and Tye 2011, 91; emphasis mine).

The “systemic dependence” of R on T, or the tokening of R “because T” indicate that there is a naturalistic relation underlying the tracking relation between R and T, which – as we will see below – could be causation or correlation or structural correspondence. This is captured in Bourget and Mendelovici’s elucidation of the TTI: “[a]ccording to the tracking theory, intentionality is a matter of detecting, carrying information about, or otherwise correlating with features of the environment” (Bourget and Mendelovici 2014, 214).

To be sure, many naturalistic theories do not explicitly posit the tracking relation. This is especially so with consumer semantics (also known as teleosemantic theories). However, if tracking is understood in a liberal way as above – as systemic dependence – then all naturalistic theories turn out to be tracking theories.

3.1.3. Ontological Reduction of the Tracking Relation/Content

Naturalistic representationalism accounts for distal content by explaining how the tracking relation holds between representations and distal targets. More specifically, they attempt to give a *reductive* metaphysical explanation of the tracking relation in terms of naturalistic entities – entities that are accepted by the natural sciences. Thus, they are *ontologically* reductive theories (Stich 1992, 258; Byrne 2006, 408), and their attempts to reduce representation is also known as the “naturalization of representation/intentionality” (Jacob 2023, §9).

Despite minor differences in detail, all naturalistic theories attempt to reduce the tracking relation (and, consequently, content) to three ingredients or elements: (i) naturalistic relations; (ii) teleological functions; and (iii) representational application formulae (RAFs). The rough general schema for a representation to track (or be about) its target is as follows:

Naturalistic Representationalism: Representation R is about a distal target T (and consequently has distal content C), iff (i) R is naturalistically related to T, (ii) it is the function of the representation-producer to produce R, and (iii) R is applied to T by employing certain representational application formulae.

Naturalistic relations between subjects and targets are the content-constituting relations of the representational process. For example, causation. Teleological functions are mainly biological functions of psychological systems to produce representations, which underwrite evaluations of representations as successful/accurate or not. Finally, representational application formulae³⁵ (RAFs) are the psychological mechanisms – eg. constancy mechanisms – and principles that ensure that representations are about a single distal target.

These three elements are the “necessary and sufficient conditions” (Neander 2017, 155) that the naturalization project aims to give for the tracking relation. More accurately, the claim is that they are *individually necessary* and *collectively sufficient* for the tracking relation and, consequently, mental representation. The reducibility of the tracking relation may be explained either as functional reduction or determinative reduction (chapter 1, §1.4.8). For instance, the tracking relation is amenable to functionalization – i.e., definition in terms of causal role. And the causal role is performed by the three naturalistic elements, which would then be the ‘realizers’ of the tracking relation. Therefore, the tracking relation could be claimed to be functionally reducible to its realizer base. Alternatively, naturalistic representationalism could claim that the tracking relation is determined by the three naturalistic elements – it is dependent on, grounded by and metaphysically necessitated by them. Thus, the tracking relation can be claimed to be determinatively reducible to them.

The ontological reducibility of representation is often paired, and deemed compatible, with explanatory irreducibility. For instance, Burge argues that representations are posited as explanatory primitives in perceptual psychology because “they have genuine explanatory uses and...the explanations that they serve cannot be reduced, without remainder, to other explanations that lack [representational] terms or concepts” (Burge 2010, 63).

3.1.4. Demarcating the Landscape

Conventionally, naturalistic representationalist theories are classified into three groups. The first group is based on which external naturalistic relation between representations and targets is taken to be central in determining contents; accordingly, there are *informational* theories (further subdivided into causal and correlational) and *structural correspondence* theories. The second group involves theories that emphasize internal naturalistic relations – causal or inferential relations among representations within the cognitive system. This group includes

³⁵ I use the term “formulae” because apart from mechanisms, principles and rules are also posited. More on this in chapter 4.

the *conceptual/inferential role* theories of content. The third group is based on answering the distality question on the basis of the teleological functions of representation-producing or -consuming systems; the theories in this group are consequently known as *teleological* theories.

However, these are not mutually exclusive classifications. Schulte and Neander (2022, §3) emphasize that every theory of mental content has the notion of teleological functions built into it, either explicitly or implicitly. Conversely, Millikan writes that teleological theories by themselves cannot explain representation. They are “piggyback theories...[and] must ride on more basic theories of representation perhaps causal theories, or picture theories, or informational theories, or some combination of these” (Millikan 2004, 66). The role of teleological functions is to enable naturalistic relations to zero-in on a particular determinate content. Functions also underwrite evaluation and thus explain misrepresentation (more on this in §3.5.1). For these reasons, the traditional classification might not be very relevant in carving the theoretical landscape at its metaphysical joints. This is corroborated by the recent trend to embrace a plurality of naturalistic relations as content-constituting, instead of privileging one. Shea’s “varitel semantics”³⁶ (2018) is a prominent example.

Setting aside the classification concern, what is important to note is that all naturalistic theories follow the same structure of explanation by employing the three elements mentioned above – naturalistic relations, teleological functions and representational application formulae (RAFs). Therefore, I will not survey the different theories of mental content as conventionally classified, but instead focus on the three elements. This survey will be in service of my argument in the next two chapters that the naturalistic elements are not collectively sufficient for representations to be about their targets. Although naturalistic relations and functions are necessary for representation, RAFs are not; consequently, the entire naturalistic package is not sufficient for metaphysically explain representation. Therefore, naturalistic theories fail in their attempt to reduce intentionality. I will focus on naturalistic relations and functions in the present chapter, and on RAFs in the next chapter.

³⁶ The word ‘varitel’ is a portmanteau composed of *variety* (pluralism about naturalistic relations, among other pluralisms) and *teleology* (a central role to teleological functions).

3.2. Naturalistic Relations

The start line for naturalistic representationalism is relations with environmental targets. As Papineau writes, “...without the external relations we would have nothing worth calling representational content to start with” (Papineau 2021, 47). The key question then is: which are the appropriate naturalistic relations? To identify these, naturalist representationalists look to natural signs (Millikan 2004, 31).

3.2.1. Natural Signification

What is a natural sign? Millikan writes that a “natural sign of a thing is something else from which you can learn of that thing by tracking in thought a connection that exists in nature” (Millikan 2004, 37). A natural sign is also known as an indicator of the thing tracked³⁷. Thus, a natural sign is a state of affairs in the world that stands for another state of affairs. For example, smoke stands for and, therefore, signifies fire, given that there is actually fire. Significantly, natural signs derive their capacity to indicate their signifieds owing to naturalistic relations between them³⁸. Thus, Fred Dretske writes: “[t]he power of signs to...indicate something derives from the way they are related to what they indicate...” (Dretske 1988, 56). The most common candidates for such relations are causation and correlation. That is, natural signs can be caused by or/and be correlated with their signifieds.

Owing to the naturalistic relations between them and their signifieds, natural signs are said to indicate or carry natural information about their signifieds. Natural information is the bare minimum information *that* the signified is the case. It is not the content-information (discussed in chapter 2, §2.2.1), which is a particular way of presenting a target. Moreover, natural signs cannot misinform about their signifieds; they cannot be false or inaccurate. As Neander puts it, “...nothing can (in this sense of the word “information”) carry the information that some state of affairs, P, is the case, unless P is in fact the case” (Neander 2017, 7).

The above point is the most important difference between a natural sign and a representation. Representations can misrepresent and be inaccurate. Natural signs, in contrast, cannot mis-signify or mis-indicate. Thus, Dretske writes, “there can be no *misindication*, only misrepresentation” (Dretske 1988, 56). It should also be noted that the use of signs *as* signs by cognitive agents like humans involves representing those signs. This is why Millikan says that

³⁷ Signification and indication are used interchangeably and synonymously in the literature.

³⁸ It is these relations that are then used as a basis for explaining the representation-target link.

“[t]he notion of a natural sign is at root an epistemic notion” (Millikan 2004, 37). And at the level of such representation there can be misindication. Therefore, red spots need not indicate a disease if someone has realistically painted spots on their face. The point is this: given the fact of the signified, the sign cannot fail to signify it.

3.2.2. The Informational Relations – Causation and Correlation

Causation is the most important relation that connects natural signs to their signifieds. This is the basis on which, for example, smoke is the sign of fire or red spots are the sign of measles. The notion of causation employed in the literature is mostly the Lewisian counterfactual dependence notion. Thus, Neander, who has given the most sophisticated causal-informational theory of content, writes that her theory – informational teleosemantics – relies “...on a non-oomphy notion of causation, such as one defined in a Lewis-style or Woodward-style [interventionist] analysis” (Neander 2017, 154).

Natural signs can also be correlated with, without being caused by, their signifieds. This is usually when both the sign and the signified have a common cause. However, the correlation may hold for any reason, the only stricture being that it should hold for a uniform reason across the domain of signs and the range of signifieds. In other words, as Dretske puts it, the correlations must be nomically grounded and therefore be “counterfactual supporting” (Dretske 1983, 58). Both causation and correlation are informational relations – signs carry natural information about their causes and correlates.

The correlation relation is analysed in the literature in terms of conditional probabilities. More specifically, correlational information is described in terms of raising (or, more generally, changing) the probability of the signified. For instance, Shea writes that “[sign] a being in state F carries correlational information about [signified] b being in state G iff $P(Gb|Fa) \neq P(Gb)$ ” (Shea 2018, 76-77). There is correlation between the sign and the signified not only when the conditional probability of the signified is raised systematically, but also when it is systematically decreased; hence the inequality ($\neq$) symbol.

The theories of mental content that employ causation and correlation as the essential content-determining relations are the informational semantics and the conceptual role semantics theories. Neander’s (2017) informational teleosemantics and Dretske (1988, 2003) take causation to be the main relation, while Shea (2018) prefers correlation. Conceptual role

semantics (CRS) says that causal and inferential roles internal to the cognitive system are determinative of content. Ned Block (1986) is a prominent proponent of CRS.

3.2.3. Structural Correspondence

In addition to the above, another naturalistic relation that is often taken to be the basis of representation is structural correspondence³⁹. This is the relation of correspondence or mapping between the relations or structures that obtain over representational vehicles, on the one hand, and relations obtaining over targets, on the other. Correspondence, in general, is understood in terms of a mathematical function which takes an item from a domain and outputs an item from a range (Shea 2013a, 64). Thus, correspondence is a functional mapping between items in two sets – domain and range – or between the two sets themselves.

In the context of representation, the structural correspondence is between the domain set of vehicles and the range set of targets. It should be noted that the content-constituting correspondence in question is not between individual vehicles and targets themselves; thus, it is not first-order correspondence, although such first-order correspondence does often obtain. The content-constituting correspondence is a second-order one between the *relations* or *structures* of the two sets; hence the term *structural* correspondence. As Shea puts it, a structural correspondence is a “relation-preserving mapping from one set of entities to another” (Shea 2018, 111). To put it differently, structural correspondence is the *inter*-set relation of correspondence between *intra*-set relations or structures. A depiction of structural correspondence is given by Figure 3.1 below (Shea 2018).




Figure 3.1: For the structure/relation H among the represented entities x_i , there is a corresponding structure/relation V among the representational vehicles v_i

³⁹ This relation variously labelled in the literature: structural isomorphism (Cummins 1996), structural similarity (Gładziejewski and Milkowski 2017), structural mapping (Millikan 2004) and structural resemblance (O’Brien and Opie 2004).

The idea behind structural correspondence being the basis of representation is this: given that there is a suitably constrained⁴⁰ structural correspondence between a set R and a set T, the structure or relation in set R represents the corresponding structure in set T, *and* the individual items in set R represent the corresponding items in set T. These individual representations are also termed “structural representations” (Cummins 1996, 110; Shea 2018, 118). In other words, the correspondence relation between the individual representation and target is derived from the primary structural correspondence relation between the sets of representations and targets.

Millikan (1984, 1989, 2004) gives pride of place to structural correspondence as the relation which is essential to determining targets/contents of representations. She refers to it ‘isomorphism’ and ‘semantic mapping’ (Millikan 2004). One of her favourite examples of the content-constituting work of structural correspondence is the dance that honeybees do to communicate to other bees about the presence of nectar at a given location. Here the bee dance is the representation. The direction of the dance represents the direction in which there is nectar, while the number of waggles or time spent wagging represents the distance to it. Focussing on the representation of the distance by the number of waggles, the following is a simplified relation between the two:

Number of waggles (representations)	Distance to nectar (target/content)
1 waggle	400 m
2 waggles	200 m
3 waggles	100 m
4 waggles	50 m
5 waggles	25 m

Here, each waggle represents its corresponding distance because there is a correspondence between the relations among the waggles – a systematic increase of 1 – and the relations between the distances – a systematic relation of ‘half of’.

According to Millikan, a sign is a representation only if it is part of a system or set R of entities_R (the number of waggles in the above example), systematic transformations of which maps on to transformations of another set T of entities_T which are the targets (distance to nectar).

⁴⁰ The constraint is usually that the intra-set relations be natural and be *used* by the cognitive system in certain ways. More on this in chapter 4 (§4.4).

Therefore, she sometimes terms structural correspondence as “semantic mapping” (Millikan 2004, 49)⁴¹.

3.2.4. Difference Between Signification and Representation

Signification only requires the obtaining of naturalistic relations. But this is not sufficient for representation. The major differences between signification and representation are as follows: The main difference, as noted above, is that natural signs cannot be false; they cannot mis-signify. In contrast, the representations can, and often do, misrepresent (Dretske 1986). The common explanation for misrepresentation adverts to malfunctioning of the representational system.

A second difference pertains to *transitivity*. Signification is transitive (Cummins and Poirier 2004, 23). If A is a sign of B, B is a sign of C, and C is a sign of D, then A is a sign of D. Dretske calls this the “xerox principle” (Dretske 1981, 57). In contrast, representation is non-transitive. Accordingly, Cummins and Poirier write: “If r represents a representation r’ of t, r need not be a representation of t...This is because the structure that r shares with r’ need not, in general, be the same as the structure r’ shares with t” (Cummins and Poirier 2004, 23). For example, while the original Mona Lisa painting (r) represents the woman Lisa del Giocondo (t), photographs of the painting (r’) are representations of the painting (r), and not of Lisa del Giocondo.

The difference in the transitivity feature leads to a third difference – whereas signification often involves a chain of smaller signification relations, representation always involves a direct relation between the representation and the target. Usually, the naturalistic signification relation between a sign and its distal signified is mediated by several single proximal sign-to-signified relations transitively linked. Thus, if A is a sign of D, then it is because A is a sign of B, B is a sign of C, and C is a sign of D. B and C are the intermediary links that are “part of the route from A to that more distal affair [D]” (Millikan 2004, 54). Dretske (1981) and Millikan (2004) describe this feature as signs being “nested” or “embedded” within other signs, respectively.

The signification relation is, therefore, a *relation of relations*: it relates the single sign-signified relations, which are its relata, to form the larger signification relation between the sign and a distal signified. Thus, the signification relation is usually a mediated relation. To put it another

⁴¹ The mapping is determined by use of the structure over entities_R by a consumer system in order to perform its functions. Thus, functions and use are the other necessary conditions for Millikan. More on the use conditions in the next chapter.

way, the signification relation is a *chain* and its intermediary relations are its links. Moreover, since there could be any number of relational links, the signification relation is a “multigrade relation” (MacBride 2020): a relation that does not have definite number of relata or degree or adicity or arity. In sum, signification is a mediated and multigrade relation.

The chain-feature of signification is due to the chain-feature of the naturalistic relations that constitute signification. An example, adapted from Millikan (2004), of the signification causal chain-relation is as follows: Neural firings in the brain (A) are a sign of (because, caused by) stimulation of the neurons in the optic nerve (B), which is a sign of retinal stimulation by incoming light (C), which are a sign of the wavelength properties of the light (D), which are a sign of the surface features of the light-reflecting object (E), which are a sign of a human face (F), which ultimately is the sign of Johnny (G).




Figure 3.2: Signification Chain.

This signification chain and Johnny are the truthmakers for the truth of the claim that neural firings (A) are a sign of Johnny (G). Now, A is not only a sign of G, but also of F, E,..., and B. This is because natural signs carry natural information about all the causal/correlational links of the signification chain. Thus, Millikan writes that “...natural signs of distal affairs carry local information about all of the more proximal affairs on the route from them to those distal affairs” (Millikan 2004, 81).

In contrast to signification’s mediated and multigrade relation to a distal sign, representation involves an unmediated, direct and binary/dyadic relation between a representation and a suitably distal target. Now, the representational process, under normal circumstances, is dependent on the naturalistic signification chain. However, a representation differs from a natural sign because it represents a single distal target without at the same time representing the intermediary links of the signification chain. As Millikan puts it, “...it is only intentional

representations that can signify distal affairs without at the same time signifying all the more proximal ones in between” (Millikan 2004, 58).

In the example in Figure 3.2, neural firings (A) *represent* Johnny (G) without at the same time representing the intermediary links of the signification chain. In fact, A can represent G even in the absence of any or all links; it is a different matter that it would be a case of misrepresentation. Thus, representation involves a single unmediated and dyadic relation between the vehicle and the target in contrast to signification’s mediated and multigrade chain of relational links. I call this relation the Representational External Directedness (RED) relation. Naturalistic representationalism calls it the tracking relation.

Correspondingly, representations carry content-information (rather than natural information) about the distal target alone. Now, in cases where the RED relation (for eg. between A and G) is dependent on the signification chain (A-B-C-D-E-F-G), the representation is produced by what Millikan calls its “normal explanation” (Millikan 1984) or “normal mechanism” (Millikan 2004, 69) – mechanisms that were relevant during the natural selection of the representation. Such representations are still natural signs, but of a special kind – they become *intentional* natural signs (Millikan 2004, 73-74). In addition to carrying natural information *that* the target is the case, a representation has the informational-cum-semantic property of content – the representation has a specific way of presenting some aspects of the target (to the exclusion of other aspects) and acquires evaluation conditions⁴².

3.2.5. Insufficiency of Relations

It is widely agreed that naturalistic relations are not sufficient for perceptual content and representation⁴³. The main reason is that naturalistic relations alone cannot solve the *distality* problem. Naturalistic relations between subjects and targets is structured as a chain relation composed of smaller mediated relational links. Thus, representations are more strongly related (by causation, correlation or structural correspondence) to proximal states – whether sensory registrations or proximal environmental states –than to distal states. For example, in the toad

⁴² Cummins and Poirier (2004) mention two other differences between signs and representations: signs are “arbitrary in a way that representations are not” (Cummins and Poirier 2004, 23), and signs are “source dependent in a way that representations are not” (Ibid., 24).

⁴³ Even a theory such as the causal theory of perception (Grice 1961) that does take causation to be necessary and sufficient for perception has to accept non-causal elements. A problem that is faced by the causal theory is the problem of “deviant causal chains” (Macpherson 2009, 504). Arstila and Pihlainen (2009) argue that this problem is often solved by including non-causal conditions such as “some kind of notion of resemblance between a perceived object and the experience caused by it in order to provide conditions for genuine cases of perception” (Arstila and Pihlainen 2009, 404).

and cricket example from Chapter 2, the representation (T5-2 activation) is more strongly caused by and correlated with neural firings than the distal target.




Figure 3.3: Representation of distal target over proximal intermediaries.

Even on a counterfactual account of causation or correlation, the proximal states are more strongly related to the representation: the absence of proximal states is a more reliable inferential cue to the absence of the representation. Thus, naturalistic relations by themselves do not provide a basis for selecting a distal target over a more strongly related proximal one.

For the above reasons, naturalistic relations by themselves are not sufficient to ensure distal content for representations. Causation, or any other naturalistic relation, may at best be a necessary condition for distal content. This point is, of course, almost universally accepted by philosophers in the naturalistic camp. Therefore, they deploy the two other elements in their metaphysical explanatory toolkit – teleological functions and representational application formulae (RAFs).

3.3. Necessity of Relations

Naturalistic theories of representation posit relations as a necessary condition for the constitution (and subsequent reduction) of distal content. But what is the source of this necessity? Burge (2010) argues that the necessity of naturalistic relations for content stems from the adherence to the thesis of content externalism.

3.3.1. An Argument for Content Externalism

As discussed in chapter 2 (§2.2.6), content externalism is a thesis about the individuation of representational content (Rowlands et al. 2020, §3), but is sometimes interpreted as a thesis about the constitution or nature of content (Burge 2010, 61). According to externalism, content

is individuated/constituted by naturalistic relations to environmental entities. So if content externalism is true, then it entails relations as necessary conditions for content.

There have been several thought experiments motivating content externalism, particularly Putnam's much-discussed 'Twin Earth' (Putnam 1975), and Burge's 'Arthritis' (Burge 1979) thought experiments. However, Burge argues that focussing on the veridicality conditions (i.e., semantic content) of perceptual representations provides independent grounds for espousing externalism.

Given that content is the way in which the target is represented, it attributes certain properties to the external target. And in attributing these properties, it sets veridicality conditions for the representation: the representation is veridical/accurate in so far as the environmental target does in fact possess the properties attributed to it. As Burge notes, the relationship between veridicality conditions and the veridicality makers (environmental properties, objects or state of affairs) is systematic and non-accidental (Burge 2010, 80): any change in environmental states is associated with a corresponding change in content and vice versa. Burge argues that this systematicity is accounted for by content-constituting relations to the environment; in other words, by externalism or anti-individualism.

“To explicate the background of systematic connections between the veridicality conditions of perceptual states and physical attributes in the environment that the states are *as of*, one must recognize that the nature of the perceptual states constitutively depends on systematic patterns of causal interaction with attributes in the environment. This conclusion entails perceptual anti-individualism” (Burge 2010, 85-86).

The endorsement of content externalism then entails accepting naturalistic relations as necessary content-constituting conditions.

3.3.2. Argument For Relationality from Action-Oriented Perception

I wish to present a different argument for the necessity of naturalistic relations for content-constitution. This argument is based on the essential nature of perceptual states as being action-oriented. The gist of my argument is this:

1. Perception is essentially for the purpose of action in the environment and on environmental objects. Thus,

Perception only if action: If a subject tokens a perception representation R, then it is for the purpose of action. This can be stated as the slogan “Perception for Action” (Riener and Stefanucci 2024, 170).

2. This purpose is fulfilled by the action-oriented content of perceptual representations. In virtue of having action-oriented content, perceptual representations are essentially action-oriented. Thus,

Action only if action-oriented content: If the purpose of R is action, then R has action-oriented content.

3. Action-oriented content is constituted by relations to environmental objects on which actions are performed. It is a necessary condition for action that perceptual states are related to the environment acted *in* and to the objects acted *upon*. Thus,

Action-oriented content only if constitutive relations to environment: If R has action-oriented content, then R bears content-constituting relations to the environment.

Therefore:

4. Perception only if relations to the environment:

If a subject tokens a perception representation R, then R bears content-constituting relations to the environment. [From 1, 2 and 3].

3.3.2.1. Premise 1: Perception for action

The perception for action slogan speaks of the essential link between perception and action. Perception is essentially (i.e., by nature) for action. It is not a mere causal link. It is uncontroversial that perception and action are causally linked. Perceptual states can act as input in motor mechanisms, which can eventually cause actions; and certain actions such as moving can cause different perceptions in us. An enquiry into an essential link, in contrast, asks whether action is a necessary condition for perception.

Riener and Stefanucci (2024) provide a useful four-fold framework for investigating the link between perception and action, which follows the schema of ‘Perception ____ Action’: Perception AND Action; Perception WITH Action; Perception AS Action; and Perception FOR Action. The first category is the causal link between perception and action mentioned above; so, I will ignore it. The remaining three categories are about the necessary link. They are also included in the *4E Cognition* (Newen et al. 2018) research paradigm in cognitive science, which studies cognition as constituted by some mixture of the 4Es – embodied, embedded,

extended and enactive. The three ‘Perception ____ Action’ framework are not mutually exclusive; thus, a specific view may fall in more than one category.

‘Perception WITH Action’ captures the idea that perception and action are mutually interacting and are together constrained by the environment. This is captured by the idea of the “brain-body-environment coupling” (Newen et al. 2018, Part V). Riener and Stefanucci (2024) cite Gibson’s (1979/2015) ecological approach to perception as a denizen of this category.

‘Perception AS Action’ includes views which hold that perception is a kind of action itself. The enactivist approach to perception (and cognition, in general) is an example (Varela et al. 1991/2016; Gallagher 2023, §6). There are two major strands of enactivism – autopoietic enactivism and sensorimotor enactivism. Autopoietic enactivism (Thompson 2007) argues that cognition is a form of biological self-sustaining activity of the organism. Sensorimotor enactivism (O’Regan and Noe 2001; Noe, 2004) focuses mainly on perception and reconceptualises perception as a form of motor activity. It argues that perception consists in an innate understanding of sensorimotor contingencies, which are dependencies of sensory perceptions on motor changes.

To be sure, enactivism also highlights the importance of the environment in which organisms act. Thus, it belongs to both the ‘Perception AS Action’ and the ‘Perception WITH Action’ frameworks. That is, perception is an action that is constrained by the environment and, in turn, shapes the environment. Both frameworks are usually characterized by a rejection of representationalism. This is most evidently seen in enactivism (Gallagher 2023, 37). Since a discussion of their anti-representationalism is beyond the scope of my dissertation, I will not discuss them further.

Finally, the ‘Perception FOR Action’ framework encompasses views which hold that perception is essentially for action. Drayson (2017) calls this the “action-oriented perception” approach, according to which perception is geared towards, and hence constrained by, possible actions. This framework includes several views which differ in their commitment to the following theses: (a) the *evolutionary purpose* (and hence biological function) of perception is action, (b) perceptual *content* is essentially for action (action-oriented); (c) perceptual *mechanisms* – i.e., internal causal processing of perceptual representations – is for action (Drayson 2017, 288). I believe that (a) and (b) together imply (c) – given that perceptual content is action-oriented for evolutionary reasons, perceptual mechanisms are also action-oriented.

The evolutionary biologist Theodosius Dobzhansky famously wrote that “nothing in biology makes sense except in the light of evolution” (Dobzhansky 1973). Since perception is a biological phenomenon, it follows that nothing about perception makes sense except in the light of evolution. And perception most plausibly evolved to guide actions and thereby ensure the fitness and survival of organisms (more on the fitness-contribution function of perception in § 3.4.7).

The ‘Perception FOR Action’ or action-oriented perception framework includes both representationalist and anti-representationalist views. An example of an anti-representationalist view is Gibson’s view that we perceive affordances in the environment, which are possibilities for action. Representationalist action-oriented perception views have been gaining increasing traction in the last few decades. I will focus on action-oriented content in what follows. I will specifically discuss the views of Cussins (2003), Millikan (1995, 2004) and Nanay (2013).

3.3.2.2. Premise 2: Action only if action-oriented content.

The representationalist action-oriented perception approach holds that the content of at least some perceptual representations is action-oriented. This is primarily the case with presentational content. That is, perceptual representations represent environmental targets not only as having properties, but as affording actions or as being *actionable*. Thus, action-affordances are a way of representing the world.

Action-oriented content has its roots in Husserl. Husserl argued that the contents of intentional perceptual experiences have a *horizon* of possible perceptions, which implicate the action-oriented nature of perception. There are two kinds of horizons – act-horizons and object-horizons. The act-horizon of a perceptual act or state is that aspect of the content which comprises *possible perceptual acts* directed on an object. These are the possible perceptions that would, “...if they occurred, tend to complete my perceptual determination of the object” (Smith and McIntyre 1982, 231). Accordingly, Husserl writes in *Cartesian Meditations*:

“Every subjective process has...an intentional *horizon of reference* to potentialities of consciousness that belong to the process itself. For example,...perception has horizons made up of other possibilities of perception, as perceptions that we *could* have, if we *actively directed* the course of perception otherwise: if, for example, we turned our eyes that way instead of this, or if we were to step forward or to one side, and so forth” (CM §19, 44).

Object-horizon is that aspect of the content which comprises *possible appearances* of the object, corresponding to the possible perceptual acts directed towards the object. Husserl writes

that the object as presented (which is part of the presentational content) appears only under certain aspects. However, the content also contains horizons of the object that fill in the remainder of the object. Husserl gives the example of a die:

“For example: the die leaves open a great variety of things pertaining to the unseen faces; yet it is already “construed” in advance as a die, in particular as coloured, rough, and the like, though each of these determinations always leaves further particulars open. This leaving open, prior to further determinings (which perhaps never take place), is a moment included in the given consciousness itself; it is precisely what makes up the “horizon” (CM §19, 45).

In the contemporary representationalism literature, action-oriented content is prominently implicated in Millikan’s (1995, 2004) pushmi-pullyu representations, Nanay’s (2013) pragmatic representations and Cussins’ (2003) activity trails.

Millikan’s Pushmi-Pullyu Representations:

Millikan builds on the Gibson’s idea that we perceive affordances in the environment and develops a representational approach to understand the perception of affordances. She argues that we perceive affordances by tokening “pushmi-pullyu representations (PPRs)” (Millikan 1995), which are the most basic kinds of perceptual representations⁴⁴. A PPR is simultaneously a descriptive (indicative) and directive (imperative) representation, “...telling both what the case is and what to do about it” (Millikan 2023, 409). However, “...they are not equivalent to the mere conjunction of a pure descriptive representation and a pure directive one but are more primitive than either” (Millikan 1995, 186).

Millikan explains the descriptive-directive tango of a PPR in terms of content – it has both descriptive and directive content. It is the via the directive content of the perceptual PPR that we perceive affordances. As Millikan puts it, the directive aspect of the perceptual PPR “...involves understanding/knowing *how* to let your current perception of your environment guide you such that you and it, together, would achieve a certain effect” (Millikan 2023, 416).

Millikan clarifies that not all perceptions are PPRs, but only those that are intended by the agent for performing an action. For example, if we perceive a ball flying towards us, then our visual perception is a descriptive representation (i.e., has *descriptive* content) representing the properties and motion of the ball. “At the same time, *given that you intend to catch the ball*, it

⁴⁴ Perceptual representations are a prime example of PPRs. Other examples include animal signals such as bee dances (indicate the location of nectar and direct other bees to go), and human linguistic utterances such as “the meeting is adjourned”.

functions as a directive representation, directing the changing positions of your eyes, head, hands and feet in the way needed to catch something that is moving in just that way over there” (Millikan 2023, 417; emphasis mine). Generalizing this point, the directive content of a perceptual representation depends on the biological functions of the perceptual subsystem that consumes the representation (more on functions in §3.3).

Nanay’s pragmatic representations

A similar view that only some perceptual states are action-oriented is held by Nanay (2013). He labels these action-oriented perceptual representations as “pragmatic representations” (Nanay 2013, 4). Pragmatic representations are called *pragmatic* because they have action-oriented content – they represent environmental objects as having “action-properties” – “properties the representation of which is necessary for the performance of an action” (Nanay 2013, 39)⁴⁵, or those that are “relevant to performing the action” (Ibid., 18). Nanay clarifies that contents of pragmatic representations also involve the attribution of non-action properties.

So, what are the action-properties that are attributed to objects? For example, if I need to perform the action of picking up a coffee cup, the action-properties include properties such as size, weight and location that are attributed in an “action-relevant way” (Nanay 2013, 39). An action-relevant way of attributing weight, for instance, is in terms of how the cup would “allow me to exert the appropriate force” (Ibid., 39). Similarly, size as an action-property would be size in terms of my grip range and strength. The attribution of such action-properties is necessary for picking up the cup. In contrast, the representation of size in inches, or of weight in kilos is of no use for the performance of the action.

Elaborating on the nature of action-properties, Nanay writes that they are “relational properties” (Nanay 2013, 40) of the perceived environmental objects, because their attribution depends on relations between the subject and environmental objects. Thus, the action-properties that are attributed to objects are always relative to the subject – they inform the subject as to what actions the targets can afford *to the subject*. Further, Nanay writes that action-properties are typically attributed unconsciously (Nanay 2013, 43). Consciously attributed action-properties, on the other hand, are “thick action-properties” (Ibid., 43). They are the action-properties attributed in perceptual experiences. There are other thick action-properties as well – for example ‘high-level’ or ‘rich’ properties such as edible, climbable, etc. Thick

⁴⁵ They are not sufficient because actual performance requires the conative component as well.

action-properties are properties that are normative in the sense that call out for an action. They have a “demand character”, demanding the performance of an action from us (Ibid., 48). But action-properties in general need not have such a demand character.

Cussins’ activity trails

Millikan and Nanay hold that some perceptual representations are action-oriented. In contrast, Cussins (2003) argues that *all* perceptual representations are action-oriented. Cussins approaches mental representations from the baseline of norms governing them. He argues that perceptual experiences are governed by the norm of *activity guidance* (Cussins 2003, 156), and not truth or accuracy. In contrast, cognitive states such as thoughts are governed by the norm of truth (Ibid., 152).

Cussins argues that perception, which is governed by activity guidance, represents the world as a realm of activity-mediation that is structured by activity trails (Cussins 2003, 154). Activity trails are possibilities of action that are constituted by “affordances and resistances of the environment” (Ibid., 154). Activity trails are aspects of the presentational content of perceptual states and are attributed to the environment. More accurately, they are non-conceptual contents (Ibid., 159) that are available to the subject’s sub-personal substrate of abilities – or the “substrate-domain” (Cussins 1990, 392) – to move in the environment.

Contrasted to perceptual representations are cognitive representations such as thoughts that are governed by the norm of truth. Their presentational content is structured by concepts, and their evaluation conditions (semantic content) are truth conditions. Thoughts represent the world as a “realm of reference” (Cussins 2003, 153) that is structured conceptually into units that can be referred to, and to which a subject can take recourse to in order to assess the mental states as accurate/true or otherwise.

A significant implication of Cussins view about perceptual content is on the semantic content of perceptual representations. Since perceptual representations have activity trails as content and are governed by the norm of activity guidance, their evaluation conditions are not accuracy conditions, but *activity-success conditions*. Cussins writes, accordingly, that

“...our grip on the shape of the normativity comes only through our understanding of the trails of activity. What forms of activity are fitting and so in place, and what forms are “out of place.” Which ways of acting flow well, and which stutter? To understand the norms is to follow the trails of activity” (Cussins 2003, 156)

Moral:

The upshot of the above discussion of Millikan's, Nanay's and Cussins' views is that at least some perceptual representations are action-oriented in virtue of having action-oriented content. I am inclined to side with Cussins in arguing that *all* perceptual representations are essentially actional. This is because, from an evolutionary standpoint, it is plausible that perception evolved to ensure our survival by means of action. Natural selection would not have selected perceptual representations purely for their role in informing us about the world independent of any implications on action. The action need not have been merely the basic actions of partaking, protecting and procreating. It would certainly have included other actions such as social bonding. Any purely informational role that perception seems to play could be explained as a secondary role where there are no pressures to actually act. Yet, even in cases where perceived properties are not *prima facie* action-oriented, the perceptions have dispositional content – as Nanay claims (Nanay 2013, 47) – in the sense that they can guide action when the need and circumstances for the action arise.

Objections:

A potential objection to action-oriented content is that we need not posit a content that consists of action-properties or activity trails, because the *causal* relation between representations and targets can explain the action-oriented nature of representations. That is, merely the fact that representations are caused by environmental objects is sufficient to ground their action-orientedness towards their causes.

This objection could be made from the vantage point of the causal theory of perception (Grice 1961), which holds that a causal relation between an object and the subject is both necessary and sufficient for perception. However, the objection could also come from a naturalistic representationalist standpoint which holds that causation is only a necessary condition for content-constitution.

I have two replies to the above objection. First, the case of hallucinations puts pressure on the view that causal relations can ground the actionality of perceptual representations. Now hallucinations can cause actions as much as veridical perceptions and, therefore, have action-oriented content. The action is, of course, unsuccessful; consequently, the hallucination is an unsuccessful and, consequently, inaccurate perceptual experience. But the hallucinated object does not exist to act as a cause of the representation. Now, some naturalist representationalists

– such as Dretske (2003) and Tye (2014) – claim that a subject represents uninstantiated universals in the case of hallucinations. For instance, Dretske writes that

“[i]n hallucinating pink rats we are aware of something – the properties, pink and rat-shaped that something is represented as having – but we are not aware of any object that has these properties – a pink, rat-shaped, object. We are aware of pure universals, uninstantiated properties. A representation...doesn’t need an object that has the properties the representation represents something to have” (Dretske 2003, 73).

Setting aside other difficulties with abstract uninstantiated properties⁴⁶, the problem is that abstract properties cannot act as causes. Therefore, there can be no causal relations involved in hallucinations, as per naturalist representationalism.

My second reply is this: a causal relation between an environmental target and the representation cannot account for the action-oriented nature of perception because action is necessarily on distal targets, and causation – as was argued in §3.2.5 – cannot account for the distality of content. As I will argue in the next chapter, a causal relation cannot account for distal content even in combination with other naturalistic elements – biological functions and representational application mechanisms. The simplest account of action in the distal environment is that distal content is also action-oriented, presenting environmental targets as possessing action-properties.

To be sure, the view that I am arguing against is the view causal relations, along with other naturalistic conditions, are sufficient to ground actions without action-oriented content having a role in causing actions. I will argue in the next sub-section that causal relations are, in fact, necessary – but for the constitution of action-oriented content. In other words, if causal relations have an action-role in perception, then it is only by (partly) constituting action-oriented content.

A second objection to action-oriented content could come from the content internalist perspective. As discussed in chapter 2 (§2.2.6), content internalism is the view that content is individuated (or constituted) entirely by intrinsic properties of the subject. Therefore, such narrow content – plausibly an abstract Fregean mode of presentation/intension (Chalmers 2004) – does not attribute action-properties to the environment. The assumption is that narrow content is relevant to explaining the phenomenal character of perceptual experiences, rather

⁴⁶ The content-character mismatch problem is a major one. See chapter 2, §2.2.4.2.

than their action orientation. However, on a two-dimensional content view (§2.2.6), narrow Fregean content, in environmental contexts, yields a wide Russellian content that includes objects and their properties as constituents. The claim could be that it is wide content that accounts for the action-role of content.

However, there are two problems with this two-dimensional proposal about action-orientation of content. First, wide content does not entail action-orientation of content. Wide content does involve environmental objects and their properties, but the proposal is silent on the content attributing action-properties to objects. Rather, since wide content is determined by intrinsically determined narrow content, it is very likely that the wide properties of objects are not action-properties. Thus, the burden is on the two-dimensional proposal to explain action-orientation of content. Action-orientation cannot be explained by narrow content, since it is determined by the intrinsic properties of the subject, and action-orientation is essentially externalist.

Second, even if it is assumed that wide content is action-oriented, there is a more general problem that afflicts the two-dimensional view. The key point about the view is that narrow content determines wide content in environmental contexts. This determination is presumably because the narrow content is somehow paired with environmental objects. The problem is that the two-dimensional view faces a *pairing problem*: the problem of explaining *how* narrow content, which is mostly considered to be an abstract mode of presentation, is paired with the environmental objects. This is analogous to the pairing problem that Kim (2011, 50) raises for interactive substance dualism – the problem is that there is no coherent criterion or relation by which immaterial (i.e., abstract and non-spatiotemporal) minds are paired with material (spatiotemporal) bodies. Kim argues that the only plausible pairing relation is the spatiotemporal relation (Kim 2011, 51). But this will not work since minds are *radically* non-spatiotemporal – essentially such that they cannot be spatiotemporal.

Similarly, in the two-dimensional content proposal, the problem is of explaining how abstract non-spatiotemporal content can be paired with an object in an environmental context. In the absence of any pairing relation, the narrow content could be paired with any object. Thus, although the resulting content could arguably guide actions, it would be entirely ad hoc which objects the contents cause to be acted upon.

3.3.2.3. Premise 3: Action-oriented content only if relations to environment.

Next, I argue that action-oriented perceptual content implies that subjects are related to environmental states of affairs⁴⁷. More precisely, these are content-constituting relations – relations that partly constitute the action-properties aspect of content. Thus, relations are a necessary condition for action-oriented content and, consequently, for perceptual representations. This is explicitly emphasized by Nanay who claims that action-properties – which are part of the content of pragmatic representations – are relational properties.

To see how the relational commitment follows from the actionality of representation, we need to focus on the individuation of action-oriented content. Action-oriented content involves attributing action-properties and is, therefore, individuated by reference to environmental states of affairs. This just means that the individuation of content requires relations to the environment. Thus, action-oriented content implies content externalism. In other words, action-oriented content is wide content.

Now, if action-oriented content is *individuation*-dependent on relations to the environment, it is very plausible that it is also *existence*-dependent (or constitution-dependent) on relations to the environment. It would be unlikely that existence and individuation conditions would be of radically different kinds. Therefore, action-oriented content is (partly) constituted by relations to the environment.

3.3.2.4. Summing up:

The argument from action for the relationality of perceptual representations can be summed up as follows:

1. If a subject tokens a perception representation R, then it is for the purpose of action.
2. If the purpose of a R is action, then R has action-oriented content.
3. If R has action-oriented content, then R bears content-constituting relations to the environment.

Therefore:

4. If a subject tokens a perception representation R, then R bears content-constituting relations to the environment. [From 1, 2 and 3].

⁴⁷ To be sure, my argument about relationality and externalism is limited to perceptual representations and content, respectively. I leave it open to whether all mental representation and content is relational and externalist.

The evolutionary purpose of perception is plausibly to guide actions in a distal environment to ensure fitness and survival. The essential action-guidance role of perceptual states is best explained by their having action-oriented content, which is constituted by being related to the external environment. It should be noted that relations to environment only partly constitute the action-properties aspect of content. This is because action-properties (cf. affordances) are always attributed relative to the subject, and hence are also partly constituted by the subject-agent's action capacities.

3.4. Teleological Functions

Many naturalistic theories hold that representations are about their targets partly because it is the biological function of representation-producers to carry information about targets. Some theories emphasize the biological functions of representation-consumers and claim that representations are about their targets partly because they are used in appropriate ways by representation-consumers to fulfil their biological functions, whatever they might be. In either case, the notion of biological functions is a necessary component, along with naturalistic relations, in the metaphysical explanation of the tracking relation between representations and targets (and, consequently, of content).

3.4.1. The Selected Effects (SE) Theory of Functions

The theory of biological functions that is most often employed in theories of mental content is the *selected effects* theory of biological function, according to which the function of a trait or a part (of the larger system) or the entire system is that effect of the trait/part/system that has been selected for in the course of evolution. More specifically, it is that effect of a trait in ancestral organisms which contributed to their reproductive fitness, which in turn has led to the current presence of the trait in descendant organisms. As Millikan writes,

“[a] trait's function is what it actually did – did most recently – that accounts for its current presence in the population, as over against historical alternative traits no longer present. A trait's function, in this sense, is its historical *raison d'être*, though the history involved is, of course, seldom if ever very distant. Call such a function a 'selected function'” (Millikan 1989a, 174).

In short, functions are selected effects (SE). Justin Garson adds that “...the selected effects theory is called an “etiological” theory for this reason, that is, because functions have to do with the causal history of the trait” (Garson 2016, 36).

Applied to the case of mental representations, the idea is that mental representations are the selected effects of representation-producer systems. Significantly, in the case of mental representations, the selection processes involved need not be limited to natural selection, but can also include learning and cultural selection. Shea thus includes “natural selection, learning, and contribution to persistence” as selection forces in his “very modern history” theory of functions (Shea 2018, 62).

SE functions are also labelled “proper functions” (Millikan 1989c, 288) or “normal-proper functions” (Neander 2017, 52) to distinguish them from the generic use of the term ‘function’ which could mean either the causal role of an entity or any of its effects. For example, the nose may function as a spectacle-holder, but that is not its proper function or SE function.

Another important distinction is between a “direct” and a “derived” proper function (Millikan 1989c). According to Millikan, a trait A has the derived proper function F iff “A originated as the product of some prior device that...had performance of F as a [direct] proper function and that...normally causes F to be performed by means of producing an item like A” (Ibid., 288). To illustrate in the case of representations, it is the direct proper function of the representation-producer system S to represent target T; and because it performs the function by means of producing representation R, R has the *derived* proper function of representing T. Thus, both the representation R and the *representing* of T are selected effects of the system S.

Selected effects (SE) functions play a major and explicit role in several theories of mental representation, but with substantial variation. For instance, Dretske (1986, 1997) and Neander (2017) posit information-carrying SE functions; Millikan (1989, 2004) and Papineau (1987, 2017) posit functions of consumer sub-systems which are part of the cognitive system; and Shea (2018) gives prominence to task performance functions of the entire cognitive system. I discuss these variations in the sub-sections that follow.

3.4.2. Information-Carrying Functions – Dretske and Neander

As discussed in §3.2.1, a necessary condition for content is the holding of a signification chain between a putative representation R and its putative target T. Given this, R is a natural sign or indicator of T, apart from being a sign of any proximal intermediary of the signification chain. R carries natural information about T. Now, Dretske argues that R becomes a *representation* (or, as Millikan puts it, an ‘*intentional*’ sign) if it is the *function* of R to be a natural sign or

indicator of T alone; that is, without at the same time having the function of indicating any proximal intermediary of the signification chain (Dretske 1986, 36).

A representation R derives its function to indicate from the direct function of representation-producing system S to indicate or carry natural information⁴⁸ about a distal target by means of producing R. Thus, Dretske write that “a system, S, represents a property, F, if and only if S has the function of indicating (providing [natural] information about) the F of a certain domain of objects. The way it performs its function (when it performs it) is by occupying different states [i.e., the Rs] corresponding to the different determinate values $f_1, f_2, \dots, f_n$, of F” (Dretske 1997, 2).

Similarly, Neander posits “response functions” (Neander 2017, 127) of representation-producers, which are functions to *causally* respond to, and hence carry natural information about, certain distal environmental states of affairs. She writes thus: “[t]o say that sensory-perceptual representation of type R has the function to carry information about stimuli of type C is to say that the system that produces Rs has the function to produce them...in response to C-type stimuli” (Ibid., 146). In other words, the representation-producer has the direct function to produce representations, which have the derived function to carry information, in response to environmental stimuli.

Significantly, Neander does not limit functions to backward-looking selected effects or outputs of traits – i.e., effects that were selected for in the course of evolution or learning. She also includes the forward-looking responsiveness to potential causes as functions. Thus, production of representations is the function of representation-producer systems not only because such representations – the outputs – were selected for in the past, but also because the representation producer system is causally responsive to environmental targets – the inputs – by means of producing representations. Accordingly, Neander writes: “...functions are selected effects, near enough...Functions are also selected causal roles or selected dispositions, which can involve inputs as well as outputs” (Neander 2017, 130).

3.4.3. Consumer Functions – Millikan

Whereas for Dretske and Neander it is the functions of the representation-producers that are central, according to Millikan (1984, 2004) and Papineau (1987, 2017), it is the function of the

⁴⁸ To recall, natural information is information *that* a target is the case.

representation-*consumers* that is instrumental in fixing the target (and, consequently, the content) of representations. I will focus on Millikan's view in what follows.

Millikan's account of representation is very broad in its application and pertains equally to organisms and cognitive systems. In the organismic case, for example, bee dances are representations produced by foraging bees and consumed/used by nectar-collecting bees. In the case of a cognitive system, both producers and consumers are subsystems that "can overlap in their components and may operate in several capacities even at the same time" (Millikan 2013b, 60). Consumers are "mechanisms or systems for which [representations] supply initial conditions for operation" (Ibid., 61), such that they can in turn perform their biological function – i.e., produce effects that they were selected for producing.

The representation-producer, representation and the representation-consumer comprise a cooperative system, each element with its own biological SE function, which Millikan calls "proper functions" (Millikan 1989c). The representation-producer has the function to produce representations so as to enable the consumer to perform its functions, and both functions have been selected for due to their performance benefitting the ancestors of the cooperative system in the past. The representation itself has the *derived* proper function of enabling the consumer to perform its direct functions.

Apart from positing proper functions, Millikan posits "normal explanations" (Millikan 1984) or "normal mechanisms" (Millikan 2004), and "normal conditions" (Millikan 1984)⁴⁹ for performing the functions. A normal explanation/mechanism is a *causal* explanation of how a particular function "was (typically) historically performed on those (perhaps rare) occasions when it was properly performed" (Millikan 1989b, 284). It was these performances that were operative during the selection of the function across generations. A normal condition, on the other hand, "is a condition, the presence of which must be mentioned in giving a full normal explanation for performance of that function" (Ibid., 285). In other words, normal conditions are causal conditions.

The normal conditions for the performance of the producer's and the consumer's functions are subtly different. The main normal condition for the performance of the producer's function (which is to produce a representation) is the presence of the target (Millikan 1984, 107), whereas the main normal condition for the performance of the consumer's functions is the

⁴⁹ Millikan often capitalizes the 'N' in Normal to highlight the evolutionary characteristic and to distinguish it from statistically normal occurrences.

correspondence between the representation and the target (Millikan 1989b, 287), and this correspondence is itself derived from the structural correspondence between the domain of representations and the range of targets.

For example, in the case of bee dances, the function of the foraging bees (representation-producers) is to produce representations whose function is to represent the target {nectar-in-a-particular-direction-and-distance}. One of the normal conditions for performance of this function is the presence of nectar (the target), and the normal explanation is how the nectar and other normal conditions ultimately cause the production of the dance (the representation). Systematic structural transformations of the dance correspond to differences in the location of target. For instance, a dance at an angle of θ to the vertical represents nectar at the same angle to the direction of the sun; a change of angle to 2θ represents nectar at the angle 2θ to the sun. The function of the collecting bees is to collect nectar, and the main normal condition for this function's performance is that there exists a correspondence between the dance and nectar in the world.

Significantly, it is the *consumer's* function that is key in determining the target of the representation. As Millikan writes, the target of a representation is determined by "what the [representation] needs to correspond to if the consumer is to perform its tasks in its normal way" (Millikan 2004, 79-80). Intuitively, then, the target of a representation is that distal state of affairs that the consumer needs to perform its function, whatever that function might be. The function of representation-producers "is only to produce for their consumers what the consumers need. Their function is only to produce representations that correspond to world affairs by a certain mapping function [i.e., structural correspondence]" (Ibid., 76). In the aforementioned bee dance example, the dances produced by the foraging bees represent {nectar at such and such direction and distance} because the collecting bees use the dances to fulfil their function of collecting nectar, which they can do only if the nectar is actually at such-and-such direction and distance.

Since the consumer's function is central to determining target (and content), Millikan sometimes labels her account "consumer semantics" (Millikan 2013a, 38). Millikan notes that the representation-producer cannot have the *function* to carry natural information, although they often do carry such information when performing their function according to the normal explanation, which involves causal relations to the target and proximal stimuli. Information-carrying is not the function because "accidentally true representations (e.g., representations that

are not causally related to their represented states of affairs, and so do not carry information about them) serve the consumer just as well as information-carrying true representations” (Ryder et al. 2013, 8).

3.4.4. Task Functions – Shea

Shea, for his part, posits *task functions* of the entire cognitive system as central to fixing distal targets. Formally, he defines task functions as follows:

“An output F from a system S is a *task function* of S iff

(a) F is a robust outcome function of S; and

(b) (i) F is a stabilized function of S; or

(ii) S has been intentionally designed to produce F” (Shea 2018, 65).

Since Shea’s focus is on natural representational systems (particularly the mind), he ignores systems with condition (b)(ii) – intentionally designed systems. Thus, the effective notion of function that Shea employs is *stabilized* task function. The processes of stabilization are the processes of selection of the robust outcome F – the function. These processes are evolution by natural selection, learning, and (where S is an organism) persistence of the organism (Shea 2018, 64). Robust outcomes are outputs that S produces in response to a variety of inputs, from different initial conditions, and in response to a wide range of relevant external conditions. Shea understands paradigmatic outputs in a broad sense to include “bodily movements, actions, and consequences of actions” (Ibid., 55).⁵⁰

It is significant to note that the task functions are characterized distally. Shea notes that since task functions, which are the explanandum of a theory of mental representation (Ibid., 34), are distal and externalist, the explanantia – representation and representational content – need to be externalist as well. Therefore, content is externalist because it is partly determined by representational vehicles bearing naturalistic relations to distal targets, which are implicated in distal task functions. This is how functions and relations come together in Shea’s account.

3.4.5. Drawbacks of the SE Notion of Functions

The selected effects (SE) notion is the most widely used notion in biology as well as the naturalistic representationalism literature. This is because it undergirds a naturalistically acceptable teleology of functions – one based on the causal-history of the function and on

⁵⁰ However, he concedes (in footnote 6) that outputs could be internal system effects like the release of hormones (Ibid., 55).

natural selection. However, it has a major drawback – the putative advantage of SE functions providing a teleological explanation of a trait or system is illusory.

As discussed in the section above, SE functions seem to explain the existence of a trait in terms of the natural selection of their effects. The idea is that the trait's presence in ancestral populations of organisms provided them a survival and reproductive advantage. Consequently, natural selection of these fitter organisms led to their descendants having the trait. Thus, "...it might seem that natural selection provides the missing causal link between what [a trait] does in a certain type of organism and its presence in that type of organism...and this might seem to explain the presence of those [traits] in the organisms incorporating them" (Cummins 1975, 750).

However, Cummins (1975) argues that the hope that selection processes can explain why a trait and/or its function exists involves a misunderstanding of evolutionary theory. Whether an organism has a trait (with its function) is determined and explained by its genome and mutations therein, not by natural selection. If the trait is beneficial and ensures reproductive fitness, natural selection only ensures the survival of the fittest organisms. Natural selection, in other words, acts as a filter of *organisms* having a trait; it does not fix the trait or its effects. Thus, Cummins writes:

“We could...think of natural selection as reacting on the set of plans generated by mutation by weeding out the bad plans: natural selection cannot alter a plan, but it can trim the set...But this is not to explain why, e.g., contractile vacuoles occur in certain protozoans, it is to explain why the sort of protozoan incorporating contractile vacuoles occurs” (Cummins 1975, 750-751).

Thus, natural selection only explains the survival of organisms having a trait, not the existence of the trait or its function. In other words, natural selection accounts for the persistence of the function (by means of persistence of the successful organisms); it does not account for its existence. For an effect to be selected to become a function, it must have occurred in the first place. And that occurrence is explained by the genome. Only then can selection forces work on the effect and fix it in future generations of the population. The SE account simply takes the existence of the effect as a given and then works ahead from there. That is why it only explains the persistence of the effect (as a function) in populations, not its existence.

3.4.6. The Causal Role Notion of Functions

The main problem with the SE account is therefore its etiological or causal-historical focus on selection of ancestral organisms. Therefore, Cummins (1975) proposes a non-etiological theory of function which is now known as the “causal role theory of function” (Garson 2016, 82). Cummins begins with the premise that the explanandum of a function ascription is not the presence of a trait (as in the SE theory) but the causal role of the trait. For example, when we say that the function of a representation is to carry information, it is not to explain the presence of the representations, but to explain the causal role representations play in the larger perceptual system. Thus, he writes that “[t]o ascribe a function to something is to ascribe a capacity to it which is singled out by its role in an analysis of some capacity of a containing system” (Cummins 1975, 765). In other words, the function of a trait or organ is its causal contribution to the effects of the larger system of which it is a part.

Cummins’ causal role notion of functions is non-etiological because it looks to the future causal effects of the trait to determine its function rather than its causal history. However, the ascription of function to a system’s component depends on a prior analysis of the system into component parts. And this analysis is relative to the interests and perspectives of those doing the analysis: “functional ascriptions do require relativization to a “functional fact” about a containing system, i.e., to the fact that a certain capacity of a containing system is approximately explained by appeal to a certain functional analysis” (Cummins 1975, 763). Moreover, since a trait can be a part of several systems based on different analyses, it turns out that a trait can several functions (Garson 2016, 86). Due to the above reasons, the causal role notion of function is not ideal as a necessary condition of a metaphysical account of perceptual representation.

3.4.7. The Fitness-Contribution Notion of Functions

The theory of function that I prefer is another non-etiological theory – the fitness-contribution theory of functions (Garson 2016, 67). The fitness-contribution theory holds that “...the function of a trait consists in its typical contribution to the fitness of the organisms that possess it” (Ibid., 67). The fitness-contribution theory of functions gives a teleological explanation of a trait (for eg., representation) in terms of its contribution to fitness. Thus, it has an advantage over SE theory which only explains the persistence of organisms with the trait and not the trait itself.

The notion of fitness in biology traces back to Darwin's *The Origin of Species* (1859)⁵¹. It is the fittest organisms that survive and have reproductive success, which is what constitutes natural selection. Thus, natural selection is based on fitness. There are broadly two notions of fitness – ecological fitness and reproductive fitness. Ecological fitness is interpreted in terms of how well an organism can interact with and navigate its environment to ensure survival and reproduction (Rosenberg and Bouchard 2023, §2). Peacock (2011) argues that there are three faces to ecological fitness: (a) the ability to *compete* for ecological resources; (b) the ability to *cooperate* to share resources; and (c) the ability to *construct* habitats, which implies that “...organisms do not merely hollow out for themselves a small space within a pre-existent environment, but rather, in many cases, alter the whole fabric of their environments” (Peacock 2011, 103). Thus, ecological fitness encompasses all the interactions of an organism with biotic and abiotic factors in the environment.

Reproductive fitness is defined in terms of the offspring an organism is likely to have. For example, “ x is fitter than y in [environment] $E =_{df} x$ has a probabilistic propensity $>.5$ to leave more offspring than y ” (Rosenberg and Bouchard 2023, §4). Most definitions of fitness-contribution functions emphasize reproductive fitness due to its measurability and predictability. However, as Peacock argues, reproductive success ultimately depends on ecological fitness (Peacock 2011, 102).

There are several definitions of fitness-contributing functions in the biological literature⁵². A definition which incorporates the significance of relations to the environment and natural selection is given by Walsh (1996). He defines biological function as follows: “The/a function of a token of type X with respect to selective regime R is to m iff X ’s doing m positively (and significantly) contributes to the average fitness of individuals possessing X with respect to R ” (Walsh 1996, 564). A selective regime is “the total set of abiological and biological (including social, developmental and physiological) factors in the environment of the trait which potentially affect the fitness of individuals with that trait” (Ibid., 564). He calls such a biological function a *relational function*, since it is a function relative to selection regimes.

I have already argued in §3.2.6.2 that the purpose of perception is to guide actions and ensure survival. Consequently, the function of the perceptual representational system is to produce representations to guide actions and ensure survival. This dovetails eminently with the fitness-

⁵¹ The phrase “survival of the fittest” however was coined by Spencer in his *The Principles of Biology* (1864), which Darwin incorporated in the sixth edition of his *TOS*.

⁵² Garson (2016, Chapter 4) provides an overview.

contribution notion of function. The function of the perceptual system is to guide actions to ensure fitness of perceiving organisms and consequently their survival. Therefore, in my arguments for content-constitution in chapter 7, I will posit fitness-contribution functions of perceptual system as a necessary condition for content-constitution.

3.5. Necessity of Teleological Functions

Almost every naturalistic theory of mental representation embraces functions as a necessary component of their account of representation and content (Cummins 1996, 120). The reasons for the necessity of functions are two – (i) they account for the normativity of representation, and consequently explain departures from normativity in the form of misrepresentation); and (ii) they account for the teleology of representations, by undergirding teleological explanations.

3.5.1. Normativity

Functions explain a major characteristic of representations, which distinguishes them from natural signs – their normativity. This is expressed in terms of the evaluation or satisfaction conditions of representations. And this grounds the possibility of error or misrepresentation.

In general, functions – whether biological or otherwise – by their very nature, are associated with standards or norms of fulfilment. In the case of biological and representational functions, Burge clarifies that these are “natural norms” in the sense that the applicability of the norms is “independent of any individual’s setting or acceding to them” (Burge 2010, 311). These natural norms then form the basis of evaluating whether the functions are performed or not. This is the normative nature of functions. Thus, according to Schulte and Neander, the notion of function is normative because “when using this notion, we (and biologists) may speak of systems functioning normally or properly, as well as of abnormal functioning, malfunction, dysfunction, functional impairment, and so on” (Schulte and Neander 2022, §2). Burge clarifies that the normative nature of functions is neither prescriptive in the strong moral sense nor is it purely descriptive as in statistical normality (Burge 2010, 314).

3.5.1.1. Veridicality/accuracy as the perceptual norm

For example, the norm associated with the function to produce beliefs is truth. Thus, beliefs have truth conditions. What is the governing norm for perceptual representational functions? There is differing opinion regarding this. According to most philosophers of perception and

cognitive science, the governing norm of perceptual representations is accuracy or veridicality. Burge, for instance, writes that “[t]he primary natural representational norm that is constitutively associated with perceptual capacity is to perceive things as they are – to form veridical perceptual representation” (Burge 2010, 312). Similarly, Neander claims that accuracy is the norm for the semantic evaluation of perceptual representations (Neander 2017, 18). Accuracy differs from truth in that it admits to degrees.

Veridicality as the norm for representational functions implies that the function of representation-producing systems is not merely to produce representations but to produce accurate or veridical representations. Accordingly, the evaluation conditions of representations are veridical or accuracy conditions, which constitute their semantic content. Producing accurate representations is therefore the result of successful performance of functions, and producing inaccurate representations or misrepresenting is due to malfunctioning or non-performance of representation-producer’s functions.

What are the grounds for claiming that veridicality is the norm of perceptual representational functions? Burge claims that the practice of perceptual psychology provides such grounds. Perceptual psychology makes ineliminable reference to representations whose primary function is to get things in the environment right. In other words, their function is to accurately inform the individual how the environment is. Burge goes so far as to claim that it is *a priori* that veridicality is the norm associated with representational functions: “From the point of view of psychology, whether there are any perceptual states or perceptual systems is an empirical matter. But it is *a priori* that where there are perceptual states or systems, their representational function is to be accurate, or to yield accurate perceptual states” (Burge 2010, 310).

To be sure, veridicality is not the only norm associated with representational functions. For instance, as explained in §3.3.3, Millikan holds that the function of representation producer subsystems is to produce representations for the use of representation consumer subsystems. Thus, we could say that the norm associated with this function is the fulfilment of consumer functions, and the resultant evaluation condition of representations would be fulfilment conditions. Millikan acknowledges that representations can be evaluated as veridical or not, but she emphasizes that their veridicality is a fallout of their primary function of servicing the needs of representation consumers, because the representation consumers will very often require that the representations they use for fulfilling their functions are veridical (Millikan 2013a, 39).

To my mind, the reason for considering accuracy as the only, or the primary, norm for perceptual representations is a tacit acceptance of the assumption that the primary purpose of perception is to inform us or provide us knowledge about the world. This assumption is associated with the “window shade model” of perception (Pautz 2021, 18). The idea is that sense organs act as a window into the world. When they are working, the windows are open, and we (or our brains) have access to how the world actually is. This window shade assumption is common to many theories of perception – naïve realism, representationalism and enactivism. Thus, the further assumption is that perception ought to give us accurate information about how the world actually is. Therefore, perceptual states are evaluated on the touchstone of accuracy, and hence are associated with accuracy conditions.

3.5.1.2. Action-success as the perceptual norm

However, this assumption about the purpose of perception can be, and has been, questioned. For one, “accuracy” is a species-specific construct. What is accurate for one species will not be accurate for another. The visual experience of a human can be evaluated for accuracy only against the normal standard for human visual experiences. A truly accurate representation of the world – one that resembles the world – would be too demanding on cognitive resources. Many philosophers would concede this point and would reply that “...more detail is not necessarily better. We can perceive our environment accurately even if we do not represent every detail in view” (Schellenberg 2018, 94).

Secondly, as discussed in §3.3.2, from an evolutionary perspective, it is more plausible that the primary purpose of perception is to guide our actions in order to facilitate fitness and survival. Providing highly accurate information about the environment is not necessary for this. Rather, what is necessary is information that is adequate for action. And such information may not resemble how the world truly is. In other words, natural selection prioritises fitness over truth⁵³. Accordingly, action-success is the more pertinent perceptual norm, and *action-success conditions* the more suitable semantic content (evaluation conditions). A perceptual representation is successful iff the action that it causes is successful, where the success of the action is analysed in terms of the action fulfilling its biological function.

⁵³ Hoffman (2019) argues for a very strong version of this idea with his “Fitness-Beats-Truth (FBT) theorem” (Hoffman 2019, 61). The FBT theorem states that “evolution by natural selection does not favour true perceptions – it routinely drives them to extinction. Instead, natural selection favours perceptions that hide the truth and guide useful action” (Ibid., xv).

Several philosophers have argued for action-success conditions as the semantic content for perception. For instance, as noted in 3.3.2.2, Cussins argues that the governing norm of perceptual representation is not accuracy but “activity guidance” (Cussins 2003, 156). Governed by this norm, perception represents the world as a realm of activity-mediation that is structured by activity trails (Ibid., 154). Since perceptual representations have activity trails as presentational content and are governed by the norm of activity guidance, their semantic content is not accuracy conditions, but *activity-success conditions*.

Springle and Buccella, similarly, hold that the function of representations is to provide the individual “practical access” to an environmental target (Springle and Buccella 2024, 11). To have practical access is to be positioned to act appropriately towards the target, and an appropriate action is one that, “if it were successfully performed, would non-accidentally solve the practical problem (satisfy the need) it functions to solve” (Ibid., 11). Thus, the norm associated with representations is practical appropriateness (Ibid., 12), and the resultant satisfaction condition is *appropriateness conditions*.

To be sure, despite endorsing action-success conditions as primary, perception can still be evaluated for accuracy, but then accuracy conditions will be “indexed to the action-guiding role of perception” (Schellenberg 2018, 94) – that is, accuracy will be pegged to success and inaccuracy to failure of the actions performed. For example, suppose a man walking in a desert on hot day sees a body of water at a distance. His visual experience has the content <body of water>. He is very thirsty and so rushes to the location. When he reaches it, there are two possibilities: (a) there actually is an oasis. His visual experience can be evaluated as accurate but only because the action – reaching a body of water – caused by the visual experience is successful; (b) there is no body of water. His visual experience can be evaluated as inaccurate (i.e., he saw a mirage), but only because his action is unsuccessful – although he reached the location, it was not a location of water.

3.5.2. Teleology

The normative nature of functions gives them a teleological texture, a purposeful poise. As noted above, functions underwrite normative evaluations of their associated traits and systems. Thus, if the function of the heart is to pump blood and it does not in fact do so, we can judge that the heart is malfunctioning. This judgement is motivated by the thought that the heart exists for the *purpose* of pumping blood, which it is failing to do. In other words, the function is taken to be the telos or purpose of the trait or system. More precisely, functions that are associated

with natural norms – norms that are independent of any individual’s setting or accepting – are “teleonomic” (Neander 2017, 20), since ‘teleological’ applies more broadly to even artifactual functions that are assigned by human creators. However, I will continue using the term ‘teleological’ since it is more widely used in the naturalistic representationalism literature.

Functions are teleological in an explanatory sense as well. The functions of a trait or system give a *teleological* explanation of that trait or system. A teleological explanation of a thing is one which is given in terms of the thing’s effects (or goals or purposes). For example, a teleological explanation of a heart is that it exists to pump blood since that is its selected effects function. The heart – the explanandum – is explained in terms of its pumping effect or purpose – the explanans. Aristotle defines a teleological explanans – the “final cause”, one of his four causes – as “that for the sake of which” the explanandum exists (Leunissen 2010, 11). And of the three kinds of final cause that he introduces in the *Physica*, one is “functions performed by (parts of) natural substances, artifacts and tools” (Ibid., 12). As commentators of Aristotle have noted (Falcon 2023, §2), teleological explanations only explain *why* – not how – something exists or happens. The question of why a thing exists can be satisfactorily and sufficiently answered by citing its function or purpose, whereas the how question is satisfactorily answered by citing either its causes or metaphysical explanantia like ground, essence, etc.

Most accounts of biological functions, including the selected effects (SE) account, ground teleological explanations. Following Mayr (1961), there is a general acceptance in the philosophy of biology of the distinction between *why* and *how* questions, which correspond to why a trait/system exists and how a trait/system works, respectively. SE and other teleological accounts of functions are commonly taken to answer only the why questions⁵⁴. Millikan, for instance, notes that “if you want to know why current species members have [trait] *T* the answer is, very simply, *because T has the function F*, that is, because *T* was selected for because it *did F*” (Millikan 1989a, 174). For example, a reasonable answer to the question “why does the heart exist?” would be a teleological one: “because it pumps blood” or, more accurately, “because hearts in ancestral populations pumped blood”.

The same is true for other teleological accounts of functions. The fitness-contribution and goal-contribution accounts, for instance, explain why a trait has a function in terms of the function’s contribution to the fitness and goals, respectively, of the organism (Garson 2016).

⁵⁴ Neander 2017 is a rare dissenting voice and argues that in answering the why question, SE functions also answer the how question.

In the case of mental representations, then, teleological representational functions explain (i) why the representation-producer or -consumer exists – it exists to produce or consume a representation; and (ii) why the representation exists – it exists to represent a distal target. Since teleological functions are a necessary aspect of the whole naturalistic solution to the distality problem, the teleological explanation that they undergird forms a part of the larger metaphysical explanation of content that is sought by representationalist theories.

3.6. Conclusion

Naturalistic theories posit the tracking relation as the representational relation and attempt to reduce it (and consequently content) to three naturalistic elements: (i) naturalistic relations; (ii) teleological functions; and (iii) Representational Application Formulae (RAFs). These three elements are argued to be individually necessary and jointly sufficient for the tracking relation, and distal content. In this chapter, I discussed the first two content-constituting elements – relations and functions – and concurred with naturalistic representationalism that they are indeed necessary for content.

Specifically, I will follow Shea in being a pluralist about relations. The most important naturalistic relation is causation, but there could be other kinds as well – for example, correlation, spatial relations and structural correspondence. I also briefly argued for the superiority of the fitness-contribution notion of teleological biological functions. I will adopt these two content-constituting elements in developing my own theory of content constitution in chapter 7, which I call ‘Content Universalism’.

4. Reductive-Naturalistic Representationalism: Representational Application

It was noted in the previous chapter that naturalistic representationalism follows the following schema of representation:

Naturalistic Representationalism: Representation R is about a distal target T (and consequently has content C), iff (i) R is naturalistically related to T, (ii) it is the function of the representation-producer to produce R, and (iii) R is applied to T by employing certain representational application formulae (RAFs).

The above three conditions comprise the individually necessary and collectively sufficient reduction base for the tracking relation that obtains between representations and targets. In the previous chapter, I argued that relations and functions are necessary, but not sufficient. In this chapter, I will argue that the third naturalistic condition is not necessary for representation. Thus, the naturalistic reduction project fails.

In §4.1 I will argue that relations and functions, as a pair, are not sufficient for the tracking relation and hence representation. Thus, the entire burden of the reductive project is on RAFs. In §4.2, I consider the RAF proposal by Neander – the distality principle, and note Schulte's argument that it fails to secure distal content. The next RAF proposal I evaluate in §4.3 is the prominent constancy mechanism proposal by Dretske, Burge and Schulte. Based on Green's argument, I conclude that constancy mechanisms are not necessary for representation. Finally, in §4.4, I assess the proposal by Millikan and Shea – the use-of-relations condition. I argue that use – which is just internal causal sensitivity of representations – cannot secure distal representation. Finally, in §4.5, I argue that no RAF, even in principle, can secure distal representation. This is because RAFs are needed to give a metaphysical basis for the tracking relation; but I argue that tracking is not representation.

4.1. Insufficiency of Relations and Functions

As argued in §§3.3 and 3.5, naturalistic relations and teleological representational functions are individually necessary for the distal content. In other words, they are individually necessary

elements of a solution to the distality problem. §3.2.5 also argued that naturalistic relations alone are not sufficient for this solution. If so, would relations and functions together be sufficient for distal content?

The short answer is no. This is because functions alone cannot explain the distality of content to the exclusion of proximal intermediaries. To see why, it should be recalled from the previous chapter (§3.4.3) that Millikan emphasizes the importance of normal explanations or mechanisms for performance of a function, which are causal mechanisms that were operative during the selection of an effect to be a function. Thus, "...for any trait to have had an effect systematic enough to have caused its selection, there will generally be some explanation of how it caused that effect that is also systematic" (Millikan 2004, 69). This implies that not merely the effect, but also the normal causal mechanism is selected for. In the case of mental representations, the normal causal mechanism is the naturalistic signification chain from the target – through the intermediary proximal stimuli – to the representation.

Neander (2017) makes the same point. She claims it is highly plausible and likely that an effect would have been selected to be a function because it was produced by a stable and recurring causal mechanism. Further, it is a general principle that "...if a system was selected for doing one thing *by* doing another then it was selected for doing both" (Neander 2017, 219). For example, if the heart was selected for circulating blood by means of pumping, then it was selected for doing both – circulating and pumping.

The upshot of the above points is that it is not merely a trait's effect that is selected for, but also the evolutionarily normal causal mechanisms for its production. In the case of the representation (the effect/function), accordingly, the normal mechanisms for its tokening will include the naturalistic relations which constitute the signification chain. Now, this implies that if the representation is selected for, so will the signification chain with its transitive relational links. Thus, the pair of relations and functions cannot secure distal content to the exclusion of the proximal intermediaries of the relational chain.

Even with the fitness-contribution notion of functions, distal content cannot be secured by the pair. Suppose it is held that the function of a representation-producer is to produce representations with distal content because it enhances the fitness of organisms. And due to its fitness contribution, it ensured the survival of organisms and was consequently selected. But it could be objected that a conjunctive content with both proximal and distal properties could be more fitness-contributing rather than a content that merely presented the organism with distal

properties. It would be more fitness-contributing because it would enable the organism to avoid misrepresentation. Misrepresentations such as illusions and hallucinations are often due to the tokening of representations with distal content but where the naturalistic relations between subjects and the environment are abnormal. However, if the proximal properties of the naturalistic relations are represented in perceptual representations, then this would preclude the possibility of misrepresentation, thereby enhancing fitness. But this is not the case. Therefore, relations and fitness-contribution functions cannot explain distal content.

Now, most naturalistic representationalists do acknowledge the point that relations and functions cannot secure distal content. However, their argument for this is based on the “breadth-of-application problem” (Burge 2010, 304). This is the problem that if relations and functions are deemed sufficient for representations, then the representational status would apply too broadly. The two conditions are present in several biological cases where, intuitively, it seems that there is no content and hence no representation. Burge explains it as follows:

“Every known explication of representation in terms of functional information, or functional sensory discrimination, or functional correlation, applies too broadly. Such explications apply just as well to discriminative responses of plants, very low-level organisms, and very low-level regulatory processes, as to what is more normally counted representation” (Burge 2010, 303).

A famous example is the case of the anaerobic magnetotactic bacteria (MTB) (Blakemore 1975, Bellini 2009). These are several species of aquatic (mainly freshwater) bacteria that have special organelles called magnetosomes, which act like internal compass needles and cause the bacteria to align with the earth’s magnetic field. Based on this alignment, the MTB move to waters with low oxygen saturation. Their locomotion using the geomagnetic field is known as “magnetotaxis” (Lefèvre and Bazyliński 2013, 498).

In the MTB case, it is clear that naturalistic relations and functions are involved in their magnetotaxis – a causal relation to the geomagnetic field and the biological function of pointing to low-oxygen water zones. The question of debate among philosophers of representation, then, is whether the magnetosomic alignment of the bacteria is a result of *representations* of the magnetic field or not. Most philosophers – eg., Dretske (1986), Burge (2010) and Schulte (2015) – argue that the MTB do *not* represent the geomagnetic field because they lack the relevant representational application mechanisms⁵⁵.

⁵⁵ The relevant representational mechanisms are associative learning mechanisms for Dretske (1986, 33-35), and constancy mechanisms for Burge (2010, 408) and Schulte (2015, 127). More on these in §4.3.

Millikan (1989), however, argues that the bacteria do represent the geomagnetic field since the conditions of her theory are met – a structural correspondence between states of the magnetosomes and states of the geomagnetic field; the biological function of the magnetosomes to point towards oxygen-free water; and the use of the representations to move towards the oxygen-free water (more on the use condition in §4.4).

The breadth of application problem, therefore, motivates most naturalistic representationalists to posit Representational Application Formulae⁵⁶ (RAFs) – principles and mechanisms that ensure that representations are applied or directed to distal targets. Thus, Cummins writes that “the target of a particular tokening of [R] is [T] just in case that tokening of [R] is produced by a mechanism N, and it is the (or a) function of N to produce representations of [T]” (Cummins 1996, 118). The application of representations to distal targets just is the tracking relation, which then constitutes distal contents. To recall from §3.1.1, naturalistic representationalism give primacy to a theory of application – they account for content in terms representational directedness.

Thus, the burden of naturalistic representationalism’s reductive project falls on RAFs. If the RAFs were to turn out as unnecessary for reduction of the tracking relation (and content), then the reductive project would fail. In the rest of the chapter, I will aim to show just that. I will argue that none of the extant RAF proposals seem to be necessary for the tracking relation.

4.2. The Distality Principle – Neander

The first RAF that I consider is Neander’s (2017) distality principle. As discussed in §3.3.1, Neander – in the context of her representationalist theory of informational teleosemantics – argues that a representational vehicle R is a representation when it is produced by a system that has the function of being causally responsive to environmental stimuli and carry information about them. For the representation-producer to be causally responsive to a stimulus implies that the stimulus is the cause of the representation R that the representation-producer produces. Thus, the naturalistic relation that Neander’s theory of informational teleosemantics relies on is causation.

⁵⁶ I call these formulae rather than mechanisms because while some of the techniques posited are mechanisms, others are general principles (as will become clear).

However, naturalistic relations, especially causation, are chain relations having several links. To ensure that the representation is about a suitably distal target rather than any of the proximal intermediaries of the causal chain, Neander plugs in a *distality principle*. She posits the principle as an answer to a question:

“Question: Why does a sensory-perceptual representation R refer to [T] and not to [P] when [P] is a proximal (intermediate) link in a [T]-to-R causal chain?

Answer: In that case, R refers to [T] rather than the more proximal [P] if the system responsible for producing Rs was adapted for responding to [P]s...as a means of responding to [T]s...by producing Rs, but it was not adapted for responding to [T]s as a means to responding to [P]s” (Neander 2017, 222).

Let us return to the toad-cricket example to better grasp the principle:




Figure 4.1.: Causal chain ending in the tokening of representation in a toad’s brain

In the above example, the causal chain extends from light hitting the insect up to the T5-2 excitation, which is the representation. By producing the T5-2 representation, the representation-producer system responds to neural signals as a means of responding to retinal firings as a means of responding to reflected light as a means of responding to a small, elongated moving object. Since it does not respond to small, elongated moving object as a means to respond to anything further, {small, elongated, moving object} is the target of the representation. In other words, the target is the end of the means-end chain of response. Moreover, the target is not any other state of affairs along the vertical line (toad food, cricket, etc.) because the other vertical candidates are not *causally* relevant to token T5-2 excitations. The general implication of Neander’s theory is therefore that surface properties of objects are the prime candidates to be the targets of representations, and consequently figure in their contents (Schulte 2018, 354).

Corresponding to means-end arrangement of the causal response is a hierarchical arrangement of response *functions*. Thus, the toad's representation-producing system has the function of responding to {small, elongated, moving object} by means of having the function of responding to light by means of having the function of responding to neural signals. Schulte (2018) – a fellow informational teleosemanticist – calls the response function which is performed by means of performing other response functions as the *highest* response function, and writes that “[informational teleosemantics] identifies the highest response function of the R-producing system as the content-determining function for R” (Schulte 2018, 357).

4.2.1. Problems for Neander's Account

As noted above, Neander's account has the consequence that surface properties are prime the candidates to be the suitably distal targets for representations. However, Schulte (2018) presents a counterexample to Neander's distality principle, wherein the highest response function of the representation-producer is not to respond to a surface property of an object, but to a more distal property in the causal chain (Schulte 2018, 358-359).

Building on the toad-insect example, Schulte asks us to imagine tiny insects in a toad's natural habitat which seasonally turn red in colour due to an increase in potassium in their blood. Toads catch and eat these insects only when they are red. Now, what is the target of a toad's visual representation when it catches these red insects? According to Neander's informational teleosemantics, it ought to be {tiny red moving object} because the surface properties of being red, tiny and moving are causally efficacious. However, Neander's distality principle gives a different answer. Since the toad's visual system responds to the tiny red moving object as a means to responding to potassium levels in the insects' blood, the highest response function of the visual system is to respond to potassium-richness. Thus, according to Neander's distality principle, the target turns out to be {potassium-rich object}. However, this is an overly distal target. Therefore, Schulte concludes that Neander's solution to the distality principle is unsatisfactory because it falters at the 'far out' distality problem – it cannot rule out overly distal candidates for targets.

Further, Neander's distality principle seems inapplicable to certain non-visual sensory representations like auditory and olfactory representation. As discussed in §1.5.2.1, in the case of auditory representation, auditory scientists and psychologists have argued that auditory representations plausibly represent properties of the sound wave. For instance, when we experience loudness, the auditory nervous system is said to represent the intensity of the sound

wave. As Plack writes, “[i]t seems likely that sound intensity is represented in terms of the firing rates of auditory nerve fibers” (Plack 2024, 127). Thus, the suitably distal target of an auditory representation R in this case is the intensity of the sound wave.

However, according to Neander’s distality principle, R would have the function of responding to intensity as a means of responding to surface vibrations of the object, which cause the sound waves. The function to respond to the surface vibrations would come out as the highest response function, with the result that R has the surface properties of the vibrating object as the target. This result is at odds with the widely accepted view in auditory science.

4.3. Constancy Mechanisms – Dretske and Schulte

The RAF that has been getting increasing traction in the literature is the constancy mechanism condition. Perceptual constancy is stability in the representation of certain target properties of objects in the face of change in: (a) proximal stimulation of the sensory system and (b) perceptual circumstances in the environment. Burge claims that perceptual constancies are capacities for objectification, which is the “...formation of a state with a representational content that is *as of* a subject matter [i.e., target] beyond idiosyncratic, proximal, or subjective features of the individual” (Burge 2010, 397). Schulte presents a distilled version of the textbook definition of constancy: “a perceptual constancy is present whenever we have stability with respect to the representation of object properties in the face of variation in proximal stimulation that is due to changes in perceptual circumstances” (Schulte 2021, 7).

The paradigm examples of perceptual constancies are shape and size constancy. In the case of shape constancy of a coin, for instance, despite variations in the slant of the coin, its distance from perceivers, and the 2D retinal impressions, the 3D shape of the coin is perceived as constant. It is important to note that constancy mechanisms are relevant to generation of stable *unisensory* representations; that is, representations of a single sense modality (visual, auditory, etc).

4.3.1. The Constancy-for-Representation Condition

Many philosophers have argued that the employment of constancy mechanisms is a necessary condition for representation (Burge 2010; Schulte 2018). I will call this the ‘Constancy-for-Representation’ (CFR) condition. To be clear, the CFR condition states that the *exercise* of

constancy mechanisms in the face of proximal changes, rather than their mere presence, is necessary for representation.

4.3.1.1. Dretske's associative learning mechanisms

Constancy mechanisms are a specific application of the associative learning mechanisms posited by Dretske (1986). According to Dretske, associative learning mechanisms are necessary for transforming sensory registrations into representations⁵⁷. So, although he does not refer to constancy mechanisms by that name, his proposal is the same in spirit. To see where associative learning mechanism fit into Dretske's account, suppose there is a mediated causal signification chain as below:




Figure 4.2: Multiple signification chains (Dretske 1986, 33)

There are two different causal chains from a distal object F – which has various properties f_i – to internal state R : causal chain f_1 to R and causal chain f_2 to R . Here, the f_i are properties of the distal object F , the s_i are sensory registrations and the I_i are internal states preceding R . Given this scenario, there are two signification chains: R signifies (i.e., is a natural sign of) I_1 , which in turn signifies s_1 which in turn signifies f_1 which signifies F ; and R signifies (i.e., is a natural sign of) I_2 , which in turn signifies s_2 which in turn signifies f_2 which signifies F . The tokening of R causes certain actions (indicated by the rightward-pointing arrow) such as catching prey.

Suppose further that it is the *function* of the R -producing system to produce R to signify or indicate F . Thus, R has the function of indicating F , and not any proximal intermediary. As discussed in §3.3.2, Dretske holds that having the function to indicate a single distal target is a key necessary condition in making a sign a representation. However, it is not sufficient. This is

⁵⁷ Dretske also favours the use condition, in addition to the associative learning mechanism as a RAF (more on the use condition in §4.4). He writes, for instance, about beliefs: “Beliefs are representational structures that acquire their meaning [i.e., content]...by actually using the information it is their function to carry in steering the system of which they are a part” (Dretske 1988, 81).

because in the above case, R could have the function of indicating F by “...indicating, and having the function of indicating, the occurrence of a certain disjunctive condition – namely, that either f_1 or f_2 (or s_1 or s_2)” (Dretske 1986, 34). In other words, there is nothing to rule out disjunctive proximal stimuli as the target. To eliminate such disjunctive targets, Dretske posits a cognitive mechanism which he calls “associative learning” mechanism (Ibid., 35).

“Learning is a process in which stimuli that indicate the presence of F are, in their turn, indicated by some relevant internal state of the organism (R in this case)” (Dretske 1986, 35). More importantly, learning is the process of tokening a single internal state (R) in order to associate multiple and variable proximal and distal stimuli with a single distal cause (F). The learning mechanism “...not only transforms a variety of different sensory inputs (the s_i) into one output-determining state (R), but is capable of modifying the character of this many-one mapping over time” (Ibid., 35). This ensures that R represents a single distal target F under various circumstances and at different times.

4.3.1.2. Schulte’s Constancy + IMNA proposal

Schulte posits constancy mechanisms as a necessary condition for perceptual representation, especially for solving the distality problem (Schulte 2018) and the breadth-of-application problem (Schulte 2015). I will focus on his treatment of the distality problem. Schulte accepts much of the basics of Neander’s informational teleosemantics, particularly the ideas that (a) causation is the relevant naturalistic relation between a putative representation and its target; (b) representation-producer systems have the response function to produce representations in (causal) response to environmental states of affairs; (c) there is a hierarchy of response functions arranged in terms of means-ends relationship; and (d) causally relevant surface properties are the best candidates for the suitably distal targets of representations.

However, as seen in §4.2.1, Schulte deems that her solution to the distality problem ultimately fails since it does not address the ‘far out’ distality problem. He therefore proposes an alternative solution that involves three conditions. He claims that representation-producers need to employ (1) perceptual constancy mechanisms in producing representation tokens. Given this, representational targets are then (3) the most *immediate* among (2) the *natural* causes of the representations⁵⁸. Given the crucial filtering role played by the naturalness and immediacy conditions, Schulte calls his overall theory of mental content “the IMNA version

⁵⁸ I have numbered the components in the sequential order of their application to states that are candidates to the distal target status.

of informational teleosemantics ('IM' for 'immediacy' and 'NA' for naturalness'))" (Schulte 2018, 366)⁵⁹.

(1) Perceptual constancy mechanisms: While Schulte does not explicitly endorse the strong claim that perceptual constancy mechanisms are necessary for *all* perceptual representation, he does accept that constancy mechanisms are a necessary condition for *determinate* representations⁶⁰. In the toad-insect example from Figure 4.1, the idea is that there can be immense variation in the proximal stimuli and, therefore, the proximal causes of a representation – the wavelengths of light reflected or the degree that retinal impressions make in the toad's eyes. But "...the only external state that qualifies as a normal cause of T5(2) excitation, i.e. as a cause that is always present in normal situations, is the distal state {a SEM [small elongated moving] object is present}" (Schulte 2018, 361).

Schulte (2021) proposes an account of constancy, which he calls the "two-track account (TTA)" of constancy (Schulte 2021, 16). The account starts with the assumption that information about the target is conflated with information about proximal stimuli. The account relies on two proximal variables and sets them off against each other to disentangle information about the distal target and represent it (Ibid., 14-15).

(2) Naturalness of the distal target candidate: Schulte admits that constancy mechanisms alone do not deliver determinate distal targets. That is because, as Dretske (1986) notes in the context of his own RAF proposal of associative learning mechanism, a disjunctive proximal state is an equally stable target candidate as the distal state. In the toad case, for example, retinal impressions $r_1, r_2, \dots, r_n$ are one set of variable proximal stimuli. Now, the disjunctive proximal state $\{r_1 \vee r_2 \vee \dots \vee r_n\}$ is an equally good candidate for the target status as is $\{\text{SEM object}\}$. To rule such states out, Schulte adds a naturalness requirement according to which, properties and states that are natural – in the sense of being non-disjunctive – are better candidates for being distal targets (Schulte 2018, 363-364).

(3) Immediacy of the distal target candidate: The constancy mechanisms and naturalness conditions are still not enough to give us suitably distal targets. To appreciate why, Schulte asks us to consider his counterexample to Neander's account – the potassium-rich red insects (§4.2.1). Both $\{\text{Tiny Red Moving (TRM) object}\}$ and $\{\text{Potassium-rich object}\}$ are equally good

⁵⁹ Garson (2019) criticizes the immediacy and the naturalness conditions.

⁶⁰ That is, Schulte is open to there being representations without constancy mechanisms; but those representations would have indeterminate contents.

candidates for the distal target status after the application of the constancy mechanism and naturalness conditions. Hence, Schulte adds the immediacy requirement. He claims that the most immediate or proximal among the natural causes of a representation is its suitably distal target. Accordingly, in the above example, {TRM object} is the suitably distal target. The immediacy requirement, thus, solves the far-out distality problem.

4.3.2. Evaluation of the CFR Condition

In this section, I will evaluate the Constancy-For-Representation (CFR) condition and eventually argue that constancy mechanisms are not necessary for representation. The idea that perceptual constancies are necessary for representation has been challenged on several fronts. A recent argument by Green (2023) shows representations can be the outputs of psychological mechanisms that are not constancy mechanisms.

As a first step of his argument, he gives a precise account of the working of a constancy mechanism, which he labels the “No Proximal Alternative + (NPA+)” account (Green 2023, 18):

“NO PROXIMAL ALTERNATIVES+ A sensory mechanism *M* is a *constancy mechanism* with respect to a distal dimension **[T]** iff *M* outputs values of a sensory variable **[R]** such that (i) there is a reasonably specific tracking relation between **[R]** and **[T]**, (ii) for any proximal variable **P** that is admissible with respect to *M*, the tracking relation between **[R]** and **[T]** is more specific than the tracking relation between **[R]** and **P**, and (iii) for any sensory variable **X** that serves as input to *M*, the tracking relation between **[R]** and **[T]** is more specific than the tracking relation between **X** and **[T]**” (Green 2023, 18).

In sum, *M* is a constancy mechanism iff its output *R* tracks only one variable with the maximum specificity, which is the distal variable/target. Condition (ii) is significant since it introduces two key notions – admissible variables and tracking specificity.

Green defines the admissible variables, whether proximal or distal, as those variables that have non-derivative causal influence on the mechanism: “...*admissible* proximal variables for a given mechanism *M* include only...those that are “directly” causally relevant to *M*’s operations” (Green 2023, 17). According to this definition, disjunctive proximal variables would be inadmissible. This is because any causal influence of a disjunctive proximal variable on *M*’s operation is parasitic on the causal role of its disjuncts. In the toad example from above, the causal role of the disjunctive proximal state $\{r_1 \vee r_2 \vee \dots \vee r_n\}$ is parasitic upon the individual disjuncts $r_1, r_2, \dots, r_n$. Thus, disjunctive proximal variables are “screened off” with respect to the

workings of a constancy mechanism because any actual causal work is done by the disjuncts (Green 2023, 40). The only admissible variables are the non-disjunctive ones because they have non-derivative causal influence on the working of the mechanism.

The second key notion is tracking specificity. Green understands tracking in terms of systematic correlation between a cause (the tracked) and its effect (the tracker). To understand tracking specificity, suppose an environmental variable T causes a proximal variable P, which then causes a sensory/neural variable R. Then,

“...the tracking relation between [R] and [T] is more specific to the extent that it is *fine-grained* and approximates a *one-to-one mapping*. More precisely, [R] is a more specific tracker of [T] to the extent that: (a) there are many values of both [R] and [T], (b) each value of [T] reliably causes a unique value of [R], and (c) each value of [R] is only reliably caused by a unique value of [T].” (Green 2023, 15).

In contrast, the tracking relation between R and the proximal variable P is less specific because each value of R tracks several values of P (one-to-many mapping). The intuitive idea behind the *specificity* of a tracking relation is that “...a tracking relation is more specific according to how much knowing the value of the cause (tracked) variable reduces uncertainty about the value of the effect (tracking) variable, and vice versa” (Green 2023, 15). In sum, something is a constancy mechanism iff it tracks a distal variable with maximal specificity and there are no admissible proximal alternatives for equally specific tracking. Hence the name “No Proximal Alternative” account.

Having defined a constancy mechanism, Green then argues that outputs of non-constancy mechanisms are also representations. He proposes a thought experiment to motivate what a non-constancy mechanism would look like (Green 2023, 30). He asks us to imagine a creature which has a visual system similar to ours, but with an additional sensory-cum-motor organ to represent distance – a rangefinder. The creature can perceive distance using either visual sense organs or the rangefinder. The rangefinder sends out laser pulses and based on the time taken for the pulses to reflect, it calculates the surface depth or distance.

Let us proceed with the following notations: M is the rangefinder mechanism which produces representations R of the distal target (i.e., distance) T; P is the time taken for the pulses sent by M to be reflected. The specific values of R correspond one-to-one with specific values of P and T despite changes in the colour, shape and texture of an object at a given distance. Green argues that R is a representation because its producer M has the function to track distance, and the

ascription of distal content in terms of veridicality conditions is explanatorily indispensable (Green 2023, 31). Green also notes that theories of representation that rely on teleological functions would also consider R as a representation and “would clearly recommend assigning distal content to [R]” (Ibid., fn 41, 31).

Given this, the question is whether the rangefinder mechanism M is a constancy mechanism. Green argues that it is *not* a constancy mechanism according to the No Proximal Alternative + (NPA+) account. To recall, NPA+ account holds that something is a constancy mechanism if there are no proximal alternatives to a distal variable that is being tracked with maximal specificity. Now, M’s output R tracks the proximal variable P with the same specificity with which it tracks the distal variable T. Since there is a proximal alternative to the distal T, the rangefinder mechanism M is not a constancy mechanism. Thus, he concludes that “...the ascription of distal representational content to [R] could be indispensable without loss of explanatory power, despite the absence of constancy” (Green 2023, 31).

Although he does not show how the rangefinder would fare according to Schulte’s TTA definition of a constancy mechanism, he does note that “[t]he two-track account and NPA+ deliver compatible verdicts for many prototypical constancy mechanisms” (Ibid., 23). Thus, the implication is that even with the TTA definition, the rangefinder would be deemed not a constancy mechanism. The conclusion therefore is that there can be representations which are outputted by non-constancy mechanisms. Therefore, constancy mechanisms are not necessary for representations, and the CFR condition is void.

4.4. Use-of-Relations Condition – Millikan and Shea

The next representational application formula (RAF) that I will evaluate is the condition that the use of naturalistic relations determines the distal target of a representation related to environmental states of affairs. Millikan and Shea are the prominent employers of this formula.

4.4.1. Millikan’s Notion of ‘Use’

It should be recalled from chapter 3 (§3.2.3) that the prominent content-constituting relation for Millikan is structural correspondence (or isomorphism) between the domain set of representations and the range set of targets. Structural correspondence is the inter-set correspondence/mirroring relation between the intra-set relations (transformations) that obtain

within each set. It should also be recalled from §3.3.3 that the selected effects functions (or proper functions) of the representation consumers are central to determining targets and content. Given this background, the target/content of a perceptual representation is fixed by the *use* of the structural correspondence by the representation consumer sub-system for the performance of its proper functions.

What does use of a structural correspondence consist in? Millikan writes that a representation-consumer uses a structural correspondence when it uses the intra-set relation that obtains among representations. This latter use, in turn, consists in the consumer being causally sensitive or responsive to the relation among the representations. To illustrate, let us return to the bee dance example:

Representation set: Number of waggles	Target/content set: Distance to nectar
1 waggle	400 m
2 waggles	200 m
3 waggles	100 m
4 waggles	50 m
5 waggles	25 m

In the above, the collecting bees (consumers) use the structural correspondence between the two sets by means of using the intra-set relation among the different bee dances (representations) of the foraging bees (producers). They use the latter by being causally sensitive to the relation among the dances, where this relation is: ‘systematic increase of one waggle’.

As Millikan puts it, “[t]he relevant differences [and hence relations] between icons that correspond to differences between world affairs are differences that the interpreter mechanisms are causally sensitive to, changing their behaviours with these differences” (Millikan 2013c, 83). Since the intra-set relation obtaining among representations is causally or operatively involved in use, she also calls it an “operative relation” (Ibid., 83). Given this intra-set relation use, the inter-set structural correspondence is itself said to be operatively used: “this correspondence between real relations is operatively involved, or directly implied by what is operatively involved, in the causal mechanism of...use” (Ibid., 83). Thus, the primary use is of

the intra-set relation among representations, while the use of the inter-set structural correspondence is derivative.

Where the naturalistic relation is causation, the use is of the information carried by the representation. Thus, Millikan writes that intentional content is that part of “the *natural* information carried by an intentional representation [that] is...designed for use by the interpreting or consuming part of the cooperative system...” (Millikan 2004, 161). Sometimes, Millikan writes as if it is not the relations among representations, but the individual representations that are used: “not just the origins but the uses of intentional representations are involved in determining their semantic values [i.e., content]. If they are not used for representing any intermediate signs...they do not represent any of these intermediate signs intentionally” (Millikan 2004, 58). But this use can be interpreted as a consequence of the intra-set relation use⁶¹.

4.4.2. Shea’s Notion of ‘Use’

Shea (2018) further expands on Millikan’s notion of use. Shea is a pluralist⁶² about content-constituting naturalistic relations – they could be causation, correlation or structural correspondence. Therefore, he writes about “exploitable relations” (Shea 2018, 36) in general. Following Millikan, Shea holds that relations are exploited or used by the cognitive system when downstream sub-systems are causally sensitive to them. For example, he defines an exploitable structural correspondence as follows:

“An exploitable structural correspondence is a structural correspondence between relation V on vehicles v_m in a system S and relation H on entities x_n in which

- (i) V is a relation that processing in S is systematically sensitive to; and
- (ii) H and x_n are of significance to S” (Shea 2018, 120).

Significance to the system (condition (ii) in the above quote) is glossed in terms of relevance to the system in the performance of its distally characterized task functions. This is similar to

⁶¹ Papineau also privileges use by representation consumers as the main RAF. He writes: “A range of ‘output-oriented’ approaches, including teleosemantics (Millikan 1984; Papineau 1984),...focus on the behavioural effects of representations rather than their causes. They then equate the contents of representations with the conditions required for the responses they prompt to achieve their ends” (Papineau 2021, 48).

⁶² Shea calls his representationalist theory “varitel semantics” - ‘varitel’ being a portmanteau composed of *variety* (pluralism about naturalistic relations) and *teleology* (a central role to teleological functions).

Millikan's requirement that the target is a normal condition for the performance of the representation-producer's functions (chapter 3, §3.4.3).

The use of relations is accomplished by means of use of the vehicles which are the relata on the side of the representational system (environmental states of affairs are the relata on the other side), since the cognitive system has direct access to only the vehicles. This vehicular use is interpreted in computational terms as being dictated by an algorithm that the system implements. An algorithm is "a sequence of operations between representations that leads to an organism performing a function" (Shea 2018, 34). Moreover, the "steps of the algorithm have to map onto causal transitions in the internal processing going on within the system" (Ibid., 35)⁶³.

To sum up Shea's account of content: "content arises from convergence between task functions, internal processes [i.e., use] and exploitable relations: it arises when internal processing over vehicles standing in exploitable relations to the environment implements an algorithm for performing the organism's task functions" (Shea 2018, 198).

4.4.3. Problems With the Use Condition

If use is merely internal causal sensitivity of downstream systems to representational vehicles and naturalistic relations, it does not seem to have the resources to solve the distality problem of explaining *how* the tracking relation between a representational vehicle and a distal target obtains and, consequently, how distal content is constituted. This is for the following reasons.

First, it should be noted that use consists of internal causal relations among representational vehicles. Now, as was discussed in chapter 3 (§3.2.5), external naturalistic relations, including causation, are insufficient to account for distal content. *A fortiori* internal causal relations (i.e., use) are insufficient as well. To be sure, this is well acknowledged, and therefore the condition on use is that it should be for the performance of the functions of consumer subsystems or the entire organism. Thus, the burden of explanation falls on functions.

However, as argued in §3.1, even relations and functions together cannot account for distal content. Functions, especially fitness-contribution functions, contribute to a teleological explanation of the existence of representations and representation-producers. However, they

⁶³ The central role to use – in the sense of internal causal sensitivity – in the constitution of content is the defining feature of conceptual or functional role semantics. Thus, in so far as Millikan's and Shea's theories depend on the use of relations and representations, they share an affinity to conceptual role semantics.

(in partnership with relations) cannot account for a metaphysical explanation, which is synchronic non-causal explanation, of content and representational directedness. And, to recall from chapter 2 (§2.3.1), a metaphysical explanation is the basic desideratum of a representationalist theory.

A second problem for the use condition is the existence of unexploited content. Cummins et al. (2006) define unexploited content as the “information or content carried by or present in a representation that its harbouring system is, for one reason or another, unable to use or exploit” (Cummins et al. 2006, 196). It is uncontroversial that there can be unexploited content in the case of non-mental representations. For example, a person may not be able to extract altitude information from a map, despite being able to extract (horizontal) distance information. This inability is not because there is no altitude content in the map, but is entirely due to the lack of abilities of the map-user. Similarly, Cummins et al. argue that there is unexploited content in the case of mental representations as well, which the brain can exploit only after *learning* to use it. They write:

“...the idea that a brain can learn to exploit previously unexploited structure in its representations is presupposed by all neural network models of learning. Such learning typically consists in adjusting synaptic weights so as to respond properly to input activation patterns. This whole process makes no sense unless it is assumed that...the representational content of input patterns remains the same throughout learning” (Cummins et al. 2006, 197).

They argue that use-based teleosemantic theories like Millikan’s have the consequence that there cannot be any unexploited content. That is because, according to use-based theories, representations can have content only if that content is selected for, and selection forces apply only to those representations that are used or exploited in service of reproductive fitness. However, they claim that “[i]t is certainly possible, and probably common, that the abilities required to exploit various features of representations evolved well after those features appeared in the representations themselves” (Cummins et al. 2006, 198). The upshot of their argument is that content and hence representation precedes use; alternatively, use presupposes that representations have distal content. Therefore, use cannot be necessary for representation.

Thirdly, the use condition seems to presuppose distal content rather than account for it. Shea writes that, “mental representations are physical particulars which interact causally in virtue of non-semantic properties (e.g. their physical form) in ways that are faithful to their semantic properties” (Shea 2018, 31). The causal interaction is the use of the representations, which is

constrained by their content (semantic properties). But the causal interaction of representations can be “faithful to” or “respect” (Ibid., 198) their distal contents only if the causal processing presupposes the distal content and, consequently, the relation to distal targets. For example, a frog’s cognitive system uses a visual representation to have the distal content <food> in order to perform its function of capturing it and feeding. But to use it in this way, the representation should already have that content. The use cannot constitute the content <food>. This shows that content – and the tracking relation to the target – is presupposed by the use, rather than constituted by it. If this is so, then the use condition cannot be necessary for the tracking relation and content.

4.5. Tracking Is Not Representing

To recall, as their solution to the representational status-cum-distality problem of content, naturalistic theories of mental representation posit a metaphysical reductive explanation of the representational tracking relation, which has three elements – relations, functions and RAFs. These are supposed to be individually necessary and collectively sufficient conditions for the tracking relation; more specifically the tracking relation reduces to these three elements. Since content is a relational property that depends on the obtaining of the tracking relation, reducing the tracking relation amounts to reducing content.

In this chapter I have argued that none of the major RAF proposals in the literature are necessary for representation. Therefore, even if the relations and functions conditions are necessary, the whole package is not sufficient for representation; consequently, the naturalistic reductive project fails. Now it may be objected that even if none of the extant RAFs are necessary, it is an empirical question as to which RAF is in fact necessary, and this might be discovered by the natural sciences. When that happens, the tracking relation can be reduced and representation can be given a sufficient and reductive metaphysical explanation.

My reply to this objection is that even if a necessary RAF is discovered and the tracking relation reduced to naturalistic elements, it still cannot account for representation. This is because the tracking relation is not a representational relation in the first place. The representational relation is the relation between a representation and a target that holds in virtue of content. As mentioned in chapter 3 (§3.1.2), there are several interpretations of the tracking relation: Kriegel writes that R tracks T if the state in which R is depends systemically upon the state in

which T is (Kriegel 2011, 70). Cutter and Tye explain tracking as: “Tokens of a state [R] in an individual x represent that [T] in virtue of the fact that: under optimal conditions, x tokens [R] iff [T], and because [T]” (Cutter and Tye 2011, 91). What is common to most interpretations is the systematic dependence condition between the tracker R and the tracked T.

Given this, there are several differences between a tracking relation and a representational relation. First, the tracking relation is irreflexive, whereas the representational relation is non-reflexive. A relation R is reflexive iff, necessarily, an entity stands in the relation with itself. For example, identity is a reflexive relation. A relation is irreflexive iff, necessarily, an entity cannot bear the relation with oneself. For example, being taller than. A *non-reflexive* relation is one where it is not the case that, necessarily, an entity bears the relation with itself. For example, loving: a person may love oneself or not.

The tracking relation is irreflexive because it is necessarily the case that an entity cannot track itself. A tracking relation between a representation and a target is posited precisely because the representations and the target are distinct entities, and the representation stands in for the target. The representational relation might *prima facie* appear irreflexive – the representation and the target in paradigm cases of perceptual representation are, after all, distinct entities. However, this assumption is challenged by self-representational theories of consciousness (Kriegel and Williford 2006), which hold that a mental state is phenomenally conscious iff it represents itself in an appropriate way. For instance, Kriegel writes that a phenomenally conscious state “has subjective character in virtue of suitably representing itself; therefore, it is a phenomenally conscious state (at all) in virtue of suitably representing itself” (Kriegel 2009, 2). This indicates that the representational is non-reflexive at best.

Secondly, the tracking relation is asymmetric, whereas the representational relation is non-symmetric: A relation R is a symmetric relation “iff, necessarily, for any x and any y, Rxy implies Ryx” (Paolini Paoletti 2025, 37). For example, *next to* is a symmetric relation. A non-symmetric relation is one where it is not the case that necessarily, Rxy implies Ryx. For example, *loving* is non-symmetric: if A loves B, then it need not be true that B loves A (although it could be). An asymmetric relation, on the other hand, is one where it is necessarily not the case that Rxy implies Ryx. *Taller than* is an example: if A is taller than B, then necessarily B is not taller than A. Detection is plausibly another asymmetric relation – for instance, if a thermometer detects temperature, then necessarily, the temperature does not detect the thermometer.

The tracking relation, like the detection relation, appears to be asymmetric. If R tracks T, then necessarily T does not track R. For example, if a predator visually tracks a prey, then necessarily the prey does not track back. The symmetry-status of the representational relation is more nuanced. In normal representational cases, where the vehicle and the target are distinct entities, the representational relation, like the tracking relation, is asymmetric: if R represents T, then necessarily it is not the case that T represents R. However, if we allow for the possibility of self-representation, the same entity can be both the vehicle and the target, in which case the representational relation would be symmetric. Thus, in general, the representational relation seems to be merely non-symmetric.

Thirdly, the tracking relation is transitive, whereas the representational relation is non-transitive: A relation R is transitive “iff, necessarily, for any x , any y and any z , Rxy and Ryz together imply Rxz ” (Paolini Paoletti 2025, 38). *Taller than* is an example: if A is taller than B, and B is taller than C, then necessarily, A is taller than C. A non-transitive relation is one where it is not the case that, necessarily, Rxy and Ryz together imply Rxz . *Loving* is an example: if A loves B, and B loves C, it is not necessary that A loves C, although that might be the case. In contrast, an intransitive relation is one where necessarily, it is not the case that, if Rxy and Ryz , then Rxz . *Son of* is an example: If A is the son of B, and B is the son of C, then necessarily A is not the son of C.

Tracking is clearly a transitive relation. If A tracks B, and B tracks C, then necessarily A tracks C. This follows from the systematic dependence understanding of tracking. However, as noted in chapter 3 (§3.2.4), representation is non-transitive. If A represents B and B represents C, then it may be true that A represents C or it may not be true.

Given the above differences, I conclude that the tracking relation is not a representational relation⁶⁴ in the first place. Thus, even if a sufficient reductive base were to be found for the tracking relation, it would have no implication for representation.

⁶⁴ Reference, in contrast, is a representational relation (Burge 2010, 31)

4.6. Conclusion

Over the previous and the present chapter, I surveyed the following content-constituting conditions by naturalistic representationalism:

Relation	Functions	Representational application formulae (RAF)	Proposed by
Causation	Producer response functions	Distality principle	Karen Neander
Causation	Producer response functions	Constancy mechanisms	Peter Schulte
Causation/ Correlation	Indication function	Constancy (learning) mechanisms	Fred Dretske
Structural correspondence	Consumer functions	Use of relations	Ruth Millikan
Plural relations	Task functions	Use of relations	Nicholas Shea

The three conditions are the reductive base for the tracking relation and consequently distal content. In this chapter, I assessed the third condition – the representational application formulae. I argued that none of the prominent RAF proposals are necessary. Even a to-be-discovered RAF would not be necessary, because the tracking relation thus secured is not the representational relation in the first place. Thus, an alternative condition that does not focus on representational application is required. I will argue in chapter 7 that this alternative condition is the co-instantiation (in different manners) of the same concrete universals in both target and subject.

5. Against Reductive-Naturalistic Representationalism and Content Particularism

The aim of this chapter is to highlight a few general weaknesses of reductive-naturalistic representationalism and to argue against one of its key commitments – content particularism, the thesis that perceptual content is constituted by environmental particulars. In §5.1., I explain the general weaknesses: first, the priority to a theory of application to a theory of content; secondly, representations are merely explanatory posits for the purpose of epistemic generalization; and thirdly, there is no, or weak, causal role given to content. In §5.2, I argue against content particularism by critically assessing, and ultimately refuting, Schellenberg's argument in favour it.

5.1. General Problems with Reductive-Naturalistic Representationalism

5.1.1. Primacy of Theory of Application

As noted in chapter 3 (§3.1.1), naturalistic theories make their theory of content subservient to their theory of representational application. However, to recall from Chapter 1 (§1.3), the representational theory of intentionality holds that the intentional directedness of representations towards their targets is on account of or in virtue of content. Thus, according to the RTI, content comes before target application. Content is the essential intentional property of representations, and representational directedness follows from content.

Naturalistic representationalist theories reverse this metaphysical order. They first determine how representations are applied to/about their targets (i.e., the tracking relation), and then explain the constitution of content. Thus, Shea asserts that “[c]ontent is a real property of the system, instantiated in part because of the way the system is embedded in its environment and in part because of the way it is internally configured – a property that is explanatory of the way the system interacts with its environment” (Shea 2013a, 498). In other words, naturalistic theories make content a relational property of the representation or the cognitive system or the subject, which then only constrains internal processing of the vehicles.

Why is it a problem if target is prioritized over content? It is a problem because this is just the ‘weak content view’ or the ‘association thesis’ that was discussed in Chapter 1 (§1.4), which even non-representationalist views such as naïve realism may endorse. To recall, the association thesis states that “[e]very experience can be associated with (propositional) content in the sense that sentences can be articulated that describe how the environment seems to the subject, without the content expressed being a proper part of the experience” (Schellenberg 2014, 201). Since content does not have a role to play in intentional directedness, it is not a genuinely representationalist view.

5.1.2. Representations Are Merely Explanatory Posits

It is pertinent to inquire why naturalistic theories put the cart before the horse. To my mind, the main reason is that representation or intentionality is not the main explanandum of naturalistic theories at all. Rather, the behaviour of organisms is the main explanandum. This behaviour could be the capacity of organisms to perceive or to act in the environment. Representations are the only the explanans – that is, representations are explanatory posits that best explain certain behaviour (perception and action) of certain organisms. If there is any attempt to explain representations (and content), it is only to serve this main goal of explaining behaviour.

For instance, Burge writes that perceptual psychology is mainly concerned with “...an explanation of how individuals perceive” (Burge 2010, 87). And representations only serve as best explanations of this capacity:

“Psychological accounts of the formation of perceptual states invoke representational states that have veridicality conditions...Nontrivial explanatory invocation of states with representational content (hence veridicality conditions) is an *epistemic* guide to demarcating genuine representational states from states on which veridicality conditions are imposed and have no explanatory value” (Burge 2010, 395-396) (emphasis mine).

Similarly, Shea writes of the basic desideratum of a naturalistic theory of mental representation as follows:

“Desideratum

An account of how representational content is constituted in a class of systems should allow us to show why recognizing the representational properties of such systems enables better explanations of behaviour than would be available otherwise” (Shea 2018, 29).

The focus of Shea's 'varietal semantics' theory is "representational explanation" (Shea 2018, 200) of behaviour – representations are explanans and behaviour is the explanandum. And since behaviour is characterized in distal and externalist terms, content is distal and externalist as well. Thus, he writes that "[c]ontents are externalist because the patterns of behaviour to be explained are world-involving: achieving distal effects in the world by reacting to distal objects and properties" (Shea 2018, 198). In many places, Shea writes as if a representational explanation is a matter of generalization and induction over the causal relations that obtain between environmental states of affairs, the internal states in the organism, and eventual behaviour. However, this makes representation (and consequently content) a matter of epistemic simplification of the causal process rather than a metaphysically robust process in its own right. This epistemic role of representations can be noted in Burge's quote above too.

Shea's theory starts with distally characterized task functions or behaviours. These are robust outcomes in the sense that they may be performed by different causal routes – from environmental states through internal states to ultimate behaviour. These causal routes are the "factorized explanations" (Shea 2018, 29) of the task function, because the causal explanation of a behaviour can be divided or factorized into three components: (a) the way environmental states of affairs cause changes to intrinsic physical properties of an organism; (b) the way those input changes cause further internal states in the organism, eventually producing motor movements; and (c) the way those movements cause changes in the external environment. The alternative offered by representational explanations is a generalization over these multiple causal routes:

"[t]he representational explanation does not...simply march in step with the factorized explanation. Causal steps which show up as different in the factorized explanation are unified in the representational explanation. The representational explanation is picking up on patterns in vehicle–world relations that the factorized explanation would miss" (Shea 2018, 202).

In other words, the positing of representations with content is to serve the purpose of generalization: "...vehicles of content will enter into generalizations that 'bridge' across multiple proximal conditions and involve distal states of affairs" (Shea 2018, 202).

In other places, Shea writes that the naturalistic relations obtaining between the subject and the target need to be "unmediated explanatory" relations in order to constitute content (Shea 2018, 83-88; 123-126). Explaining the distinction between mediated and unmediated explanation in the case of the correlation relation, he writes: "a correlation between [representation] R and

condition [P] plays a mediated role in an explanation if its role depends on the explanation adverting to a further correlation between [P] and some further condition [P']; otherwise, it plays an unmediated role" (Ibid., 84).

Shea explains with the example of a frog and its prey (Shea 2018, 152). The task function (explanandum) of the frog is fly capture, in the performance of which representation R (explanans) is tokened by the frog's visual system by carrying correlational information (the relation) about states of affairs in the world. The correlation between R and the distal target {fly at (x, y, z)} provides an unmediated explanation of fly capture. The correlation between R and a more proximal state of affairs {retinal impression} can also explain the fly capture, but only mediatedly by invoking a further correlation between the {retinal impression} and {fly at (x, y, z)}. Therefore, the correlation between R and {fly at (x, y, z)} provides an unmediated explanation of fly capture and hence is the content-constituting relation yielding the content <fly at (x, y, z)>.

While there is nothing wrong in positing representations as explanans of perception and action, positing them as mere generalizations or explanatory tools undercuts the metaphysical heft of naturalistic theories of representation. It also undermines the representationalist essence of perception. A metaphysical explanation of representation requires considering representation and content as the explanandum. That is not what naturalistic theories do. Moreover, a merely explanatory role to representations undermines the distinct causal efficacy of content. All causal efficacy is rather with the representational vehicles in virtue of their physical syntactic properties.

5.1.3. No Causal Efficacy to Content

To be sure, reductive-naturalistic representationalists disagree about whether content has novel causal efficacy over and above the causal efficacy of syntactic features of representational vehicles (Shea 2018, 206-209). Let us call the two camps – causal irrealists and causal realists about content.

According to the causal irrealist camp, content has no causal efficacy at all, but only an explanatory role. Egan (2004) for instance argues that representations primarily have "mathematical content" (Egan 2004, 122) rather than cognitive content. Mathematical content is just the abstracted mathematical function – i.e., input-output pairing – that describes internal causal processing of vehicles as per an algorithm (Ibid., 127-128). Any cognitive content

ascribed to vehicles is merely an intentional gloss on the computational-causal processing. Moreover, the intentional gloss depends on the pragmatic considerations of researchers who start with an intentionally characterized explanandum such as belief or behaviour (Egan 2004, 125). Cognitive contents are “...a gloss that shows, in a perspicuous way, how the computational/mathematical theory manages to explain the intentionally-described explanandum with which we began and which it is the job of the theory to explain” (Ibid., 128). The implication is that cognitive content has no causal efficacy.

Shea writes that he is “neutral” regarding the causal efficacy of content (Shea 2018, 209). However, Shea identifies more with the causal irrealist camp, as he writes

“My accounts share with the [causal irrealist] camp the view that contents are useful because they allow us to see why an organism’s internal workings are suited to performing certain tasks. Unlike Egan..., I characterize those tasks in non-semantic terms in the first instance ...But...contents for me are partly⁶⁵ a matter of how an organism can perform the computations needed to produce appropriate outputs in appropriate circumstances” (Shea 2018, 207).

Moreover, as discussed in the subsection above, Shea writes that representational explanations are mere generalizations that bridge over factorized causal explanations of task functions. Similarly, Burge argues that representations earn their theoretical keep owing to their non-trivial explanatory role.

Even if any causal powers are attributed to content, the causal irrealist camp reduces them to the causal powers of its reduction base. Thus, content is ‘functionally reduced’ to naturalistic elements via the functionalization route. To recall from chapter 1 (§1.4.8), Kim writes that to functionalize an entity is (i) to define the entity entirely in terms of its causal-functional role in a system; and then (ii) to identify physical realizers of the causal-functional role. Once functionalized, the entity can be reduced to the physical realizers. This is the functional model of reduction, which is the most popular among reductive-naturalistic representationalists.

Coming to the causal realist camp, Dretske (1988) is an exemplar. He argues for causal efficacy for content along with causal efficacy for the representational vehicles *qua* physical states. He follows the “dual explanandum strategy” (Robb et al. 2023, §6.3) of causal relevance to mental states in general. The dual explanandum strategy holds that there are two causal explananda – the bodily movement and the behaviour that the movement constitutes. For eg., the *bodily*

⁶⁵ Shea says ‘partly’ because the other element that constitutes content is the naturalistic relations.

movement of sticking out an arm when driving constitutes the *behaviour* of signalling a right turn (in the context of traffic rules, of course). The claim is then that while physical events (neural states) explain the bodily movement, mental events (beliefs, desires, etc.) explain the behaviour. Thus, there is no causal competition between physical and mental events.

Dretske argues for the following dual explanandum picture of causation:




Figure 5.1: Dretske's model of content-causation (Dretske 1988, 84)

C is the *event* of tokening the representational vehicle *qua* its syntactic features, which causes the event of bodily movement (M). Dretske calls event C the “triggering cause” (Dretske 1988, 42) of bodily movement M. C’s causing M is a fact, and this fact is the behaviour. Alternatively, Dretske describes behaviour as the *process* of C causing M; where “[a] process isn’t a sequence of events which stand in certain causal relations to one another. It is their standing in these relations to one another – one event (or two or more events) producing or bringing about another” (Ibid., 35).

To recall from chapter 3 (§3.3.2), according to Dretske, a necessary condition for a representation to have content is for it to have the function of indicating an external state of affairs. In the above diagram, C has the function to indicate F (and given other conditions – relations and learning mechanisms), it has content. That is why the fact of C indicating F is the fact of C having content. Now, the fact that C indicates F explains the fact that C causes M (which is behaviour). Therefore, Dretske labels the fact that C indicates F (and, consequently, has content) as the “structuring cause” (Dretske 1988, 42) of the *fact* that C causes M (i.e., behaviour). Thus, content-causation is fact causation; more specifically it is the causal explanation of one fact (behaviour) by another (having of content).

The reason content is not a causal *event* is due to its abstract nature:

“Whatever else a meaning [i.e., content] might be, it certainly is not, like an event, a spatio-temporal particular that could cause something to happen. It is, rather, an abstract entity, something more in the nature of a universal property such as redness or triangularity...[I]n exploring the possibility of a causal role for meaning one is exploring the possibility, not of

meaning itself being a cause, but of a *thing's having meaning* being a cause or of the *fact that something has meaning* being a causally relevant fact about the thing” (Dretske 1988, 80).

It should be noted that, strictly speaking, the fact of having content causally *explains* (rather than causes) the fact of C causing M. So, the causal relevance of content is a derivative causal explanatory role, after all.

In contrast, I have argued in chapter 3 (§3.2.6.2) that content has genuine causal efficacy in the sense that it is action-oriented and can cause subjects to act and behave, rather than only produce bodily movements. This implies that content is a concrete property. I will argue for the concrete nature of content in chapter 6, where I will submit that content is the property-instance that results from the instantiation of universals. In chapter 8, I will argue that the action-oriented causal powers of content are novel vis-à-vis its dependence base; in other words, content is a strongly emergent property.

5.2. Against Content Particularism

All naturalistic representationalist theories hold that naturalistic relations to environmental targets are necessary conditions for the constitution of content. In fact, most go further and claim that the relations are borne to particular targets (as opposed to general entities) and that content is partly constituted by the particulars. This is the thesis of content particularism. To recall from chapter 2 (§2.2.5), content particularism is the thesis that “[a] perceptual state *M* brought about by being perceptually related to the particular α has the property that *M*’s content is [partially] constituted by α ” (Schellenberg 2018, 58). Particular content is also known as singular content in the literature. For example, Schellenberg (2018) explicitly endorses both ontological (reductive) naturalism (Schellenberg 2018, 3) and content particularism (Ibid., 66).

Two points must be noted here. First, content particularism claims that content is constituted by environmental particulars by being *related* to them. Thus, in the terms that were introduced in chapter 2 (§2.2.5.1), the particularity of content in question is relational-constitutional particularity and not phenomenological particularity. In the case of phenomenological particularity of content, it only phenomenally *seems* to the subject that there is a particular object. Thus, the particularity is merely part of the presentational content. Relational-constitutional particularity, in contrast, holds that the particularity of content is due to being related to and being constituted by environmental particulars.

Second, the claim of content particularism is not merely that perceptual content is constituted by particulars, but that, more precisely, the *particularity* or thisness (haecceity) of the particular is an aspect of the presentational content. Thus, a perceptual representation, in virtue of having particular content, represents the particularity of the particular, along with its other properties such as colour, shape, size, smell, etc. To be sure, content particularists – such as Schellenberg – understand particularity as not as a distinct property but in the ontology-lite sense of individuating factor (Chapter 1, §1.6.2).

In this section, I will argue, mainly on the basis of empirical considerations in cognitive science and perceptual psychology, that content particularism is false. This implies that content generalism is true for perceptual representations – contents are constituted by or are identical to general entities. I will focus on Schellenberg’s argument for particularism, because it is among the more detailed defences of particularism. However, my argument applies to particularism in general.

5.2.1. Schellenberg’s Capacitism

Content particularism is a central element in Schellenberg’s naturalist representationalist theory of perception, which she calls *capacitism*. Schellenberg starts with the assumption that the main function of perceptual system is to discriminate and single out particulars in the environment (Schellenberg 2018, 67). This function is performed by the employment of perceptual capacities. Thus, capacitism is the view that “perception is constitutively a matter of employing perceptual capacities that function to discriminate and single out particulars” (Ibid., 29). The performance of the function of discriminating particulars also requires the subject to be perceptually related to environmental particulars. Perceptual relations are naturalistic relations – causal (Ibid., 15) and spatiotemporal (Ibid., 32) relations. Given the above conditions – naturalistic relations, functions and capacities – the employment of perceptual capacities yields perceptual representations. In fact, perceptual representation is reducible to the combination of the above conditions, which makes capacitism a reductive naturalist representationalist view (Ibid., 3).

Capacitism endorses content particularism. This is because perceptual content is partly constituted by the employment of perceptual capacities and partly by the particulars that are related to and singled out. Therefore, perceptual content is particular content. More specifically, Schellenberg writes that perceptual content has two levels – content type and token content – both of which are *de re* (i.e., non-propositional) modes of presentation. Content type is

constituted by the employment of perceptual capacities, and token contents is constituted by content types + the environmental particulars related to⁶⁶ (Schellenberg 2018, 87-88). Since particular contents are *de re* modes of presentation, she also calls her version of content particularism as “Fregean particularism” (Ibid., 84).

5.2.2. Schellenberg’s Argument for Content Particularism

Schellenberg calls her argument for content particularism as the singular content argument:

“The Singular Content Argument

- I. If a subject *S* perceives a particular α , then *S* discriminates and singles out α .
- II. If *S* discriminates and singles out α , then *S*’s perceptual state *M* brought about by being perceptually related to α is constituted by discriminating and singling out α .
- III. If *M* is constituted by discriminating and singling out α , then *S* represents α (under a mode of presentation) such that *M* has the property that its content is constituted by α .

From I–III. If *S* perceives α , then *S*’s perceptual state *M* brought about by being perceptually related to α has the property that its content is constituted by α ” (Schellenberg 2018, 66).

The key premise here is premise II since it moves from a statement of *contingent* fact – that perception involves discriminating particulars – to the *constitutive* fact that perceptual states are constituted by discriminating particulars. Therefore, my main criticism will be of premise II. However, I will give my objections to the other premises as well.

5.2.2.1. Premise I: Perceptual function to discriminate particulars.

Premise I is the statement that the function of perception is to discriminate and single out particulars. Schellenberg also claims that perception is *essentially* discriminative. Thus, she writes that “...discriminating and singling out a particular from its surround is a [metaphysically] necessary condition for perceiving the particular” (Schellenberg 2018, 25). To reiterate, Schellenberg includes particular objects, property-instances and events in her use of the term ‘particular’.

Schellenberg juxtaposes the discriminatory view of perception with the attributive view of perception, favouring the former. According to the attributive view, the immediate function of perception is to attribute properties – or “perceptual attributives” – to environmental particulars (Burge 2010, 380). And as Burge explains, perceptual attributives are general entities: “[a]

⁶⁶ In the case of hallucinations, where there are no environmental particulars, the token content is gappy (Schellenberg 2018, 89).

perceptual attributive is an aspect of perceptual representational content that functions to indicate a repeatable type and to group or characterize purported particulars as being of that type” (Ibid., 380). Schellenberg concedes that we do attribute properties in perception, but any such attribution is not part of its essence, but “will be grounded in discrimination” (Schellenberg 2018, 67).

Schellenberg does not provide an explicit argument for stating that discriminating particulars is the function of perception. However, from her writings, the following argument can be constructed:

1. If a subject S tracks the differences between particulars, then S discriminates particulars. This is because “...at the very minimum perceptual discrimination is a matter of tracking the difference between particulars” (Schellenberg 2018, 28).
2. Tracking particulars is an essential feature and function of perception.
Therefore,
3. Discriminating particulars is an essential function of perception.

It should be noted Schellenberg does not use ‘tracking’ in the sense it is used in ‘tracking relation’, which is the naturalistic representational relation. In the latter case, representation R tracks target T iff the state of R depends systematically on the state of T (chapter 3, §3.1.2). In the former case tracking the differences between objects is being able to distinguish between objects and continue to distinguish between them over time. In response to Schellenberg’s argument, I will draw on studies in cognitive science to argue that tracking is not a function of perception at all, let alone an essential function. Rather, it is the function of *attention*, which is a post-perceptual cognitive mechanism.

5.2.2.1.1. Object Tracking

Object tracking has been intensively researched under the Multiple-Object Tracking (MOT) paradigm in cognitive science (Pylyshyn and Storm 1988; Pylyshyn 2001). MOT studies consist of showing participants between 8 and 24 qualitatively identical objects or items, followed by making a subset of these distinct – by briefly flashing them on and off or by other means. Then, all the items are moved around randomly, and the participants are asked to track the previously highlighted subset (Pylyshyn 2001, 142).

Most MOT studies have found that (a) we can simultaneously track four or five independently moving objects (Pylyshyn 2001, 143); and (b) we do not visually track objects by representing their properties or even their locations. In fact,

“...items can be tracked despite the lack of distinctive properties (and, indeed when their properties are changing) and despite constantly changing locations and unpredictable motions. Taken together, these studies suggest that what Marr (1982) referred to as the early vision system...is able to individuate and keep track of about five visual objects and does so without using an encoding of any of their visual properties” (Pylyshyn 2001, 143).

The main point is that the particular properties (and hence the particularity) of the items are not encoded or represented. While the findings of the studies agree with each other, they differ with respect to the mechanism for tracking. I will briefly discuss Pylyshyn’s proposal, which is the most cited.

To track an item, perceptual states need to first individuate it. Pylyshyn proposes a mechanism for individuating the items, which he calls visual indexes. Therefore, he calls his view the “visual indexing theory” (Pylyshyn 2001, 145). The assignment of visual indexes to items presupposes a primitive representation of the items. The visual system first represents external items or targets as “primitive visual objects” (Ibid., 144) that are not individuated. A primitive visual object is a “...cluster of visual features [that is] segregated from the background or picked out as a unit” (Ibid., 145). Pylyshyn further writes that “...primitive visual objects are typically the proximal counterparts of real physical objects” (Ibid., 153). Thus, primitive objects are part of the content of visual representations⁶⁷.

Then, the visual system *selects* a few of these primitive objects and assigns them visual indexes that act as labels. This assigning of indexes is analogous to demonstrative reference via terms like ‘this’ or ‘that’ or, more relevantly, to demonstration by gestures of pointing with fingers. This is why Pylyshyn calls the visual indexes “FINSTs” (*Fingers of Instantiation*) (Pylyshyn 2001, 130). It should be noted that what is indexed (and hence trackable) is not the external targets but the primitive objects that are part of representational content. The primitive visual objects, in virtue of the visual indexes, are individuated and trackable. This process of individuation has the implication that the particularity of perceptual content is not the

⁶⁷ By part, I do not mean component part. Rather, I use ‘part’ to mean an aspect or a phase of the content.

constitutional-relational particularity that Schellenberg claims but is merely phenomenological particularity.

Pylyshyn further argues that the selection of the to-be-tracked primitive objects is a “preconceptual selection” (Pylyshyn 2001, 139). Pylyshyn uses the term ‘conceptual’ in a non-standard way. According to him, to be conceptual is to represent under a particular description or in a particular way – whether or not using concepts as standardly understood. Thus, preconceptual selection is simply non-representational selection. He also describes the selection as a “preconceptual (purely causal) connection” (Ibid., fn1, 128) between the visual system and the primitive objects. Accordingly, he writes that “[t]his sort of causal connection between a perceptual system and an object in a scene is quite different from a representational or intentional or conceptual connection. For one thing there can be no question of the object being *misrepresented* since it is not represented as something” (Ibid., 147). Similarly, he also describes it as a “direct access mechanism” (Ibid., 141) – a “direct connection between...a visual representation and certain token elements in the visual field, a connection that is unmediated by an encoding of properties of the elements in question” (Pylyshyn 2001, 128).

The key point to note is that the assignment of FINSTs (visual indexes) does not happen during perceptual representation. It is an operation performed *on* perceptual content and, therefore, occurs *after* perception. Is there anything more that can be said of this direct-access, non-representational selection that individuates perceptual objects? In the attentional literature in cognitive science and philosophy, it is widely accepted that non-representational selection is the hallmark of *attention*⁶⁸ (Rensink 2013, Scholl 2001). And the consensus is that attention is a post-perceptual cognitive mechanism. Therefore, Pylyshyn claims that attention is a necessary condition for assigning visual indexes and consequently tracking objects. He writes that his FINST-visual indexing theory predicts that

“...indexes would be readily grabbed by any new object that appears in the field of view, so ‘attention’ may be involved in orienting the system to the relevant part of the visual field (i.e. attention may be required to control eye movements and to provide some selection of the inputs to the visual indexing system)” (Pylyshyn 2001, 149).

Thus, it is the post-perceptual mechanisms of attention and indexing – and not perceptual representation – that individuate and therefore particularize the objects that are part of representational content. This is what enables tracking and consequently discrimination of

⁶⁸ I will discuss the role of attention in perception in more detail in chapter 8.

particulars. Therefore, tracking and discrimination is the function of post-perceptual mechanisms⁶⁹.

The lesson from MOT studies – and attentional studies more generally – is that it is not the essential function of perception to track or discriminate particulars. In fact, there is no consensus on what the essential function of perception is. It could be that attributing general properties to particular objects is the primary function of perception (Burge 2010, 380). Given this lack of consensus, as I argued in chapter 3 (§3.3.2), the best that can be said about the essential function of perception – on the back of support from evolutionary biology – is that perception is for action and survival, which involves both attributing general properties to targets and discriminating them.

5.2.2.2. Premise II: Perceptual representations are constituted by discriminating particulars

Premise II states that if S discriminates and singles out α , then S's perceptual state M brought about by being perceptually related to α is constituted by discriminating and singling out α . The premise moves from a contingent fact about discriminating particulars to a constitutive fact about perceptual states (representations) being *constituted* by such discriminating. Schellenberg supports the premise with the following argument:

“...discriminating and singling out a particular from its surround is metaphysically necessary for perceiving the particular. As with any metaphysical necessity claim, we can ask what the source of this necessity is. The source of this necessity is that perceptual states are *constituted* by discriminating and singling out particulars. In other words, the fact that perceptual states are constituted by discriminating and singling out particulars is why discriminating and singling out particulars is necessary for perceiving particulars” (Schellenberg 2018, 26; emphasis mine).

In general terms, Schellenberg is taking a constitution claim to be the source of a modal metaphysical necessity claim. To recall, Schellenberg interprets constitution in terms of grounding (Schellenberg 2018, 16). So, according to Schellenberg, grounding is the source of modality.

In reply, I want to flag this grounding $\rightarrow$ modality relation. As Brenner et al. note, “...it is widely held that the grounds of a fact are modally sufficient for it to obtain, but *need not be modally necessary for it*, while the opposite is true for essence” (Brenner et al. 2021, 2). That is, whereas essence could be the source of modality – this is the essentialist theory of modality

⁶⁹ I will discuss attention in more detail in chapter 8 (§8.3).

which the real definition account of essence endorses (chapter 1, §1.4.1) – ground is not suitable. Similarly, Bliss and Trogdon write that “...grounding isn’t a purely modal notion” (Bliss and Trogdon 2024, 3.1).

Further, as I argued above, Schellenberg’s argument (from tracking) that the essence of perceptual states is to discriminate particulars fails. Thus, even on an essentialist theory of modality, discriminating particulars is not metaphysically necessary feature of perceptual states.

Independent empirical arguments against this premise can be found in change blindness research. Change blindness studies show that perceptual representations are not constituted by discriminating particulars, but by representing the general meaning (known as gist) and spatial features of a target scene. In other words, content consists of general properties.

5.2.2.2.1. Change blindness

Change blindness (Simons and Levin 1997) is the failure or inability of perceivers to detect substantial changes to a scene or state of affairs under certain circumstances. These circumstances under which changes are not detected are when the triggers of change – known as motion transients – are either masked from attention or are so gradual that they do not draw attention. The failure of change detection pertains to “not only detection proper (i.e., the observer reporting on the existence of the change), but also identification (reporting what the change is) and localization (reporting where it is)” (Rensink 2002, 246).

Change blindness has mostly been studied in the case of visual change detection. Most studies of change blindness have been conducted under circumstances where the triggers of change are masked from the attention of perceivers by changing a scene during natural or induced visual disruptions. For example, changes could be made during the saccading movement of the eyes, or during blinks, and so on (Rensink 2002, 252). However, recent studies have focussed on blindness to gradual changes to a scene where there are no visual disruptions. For instance, Simons et al. (2000) found change blindness among participants when they were asked to detect changes in a 12 second-long video where the changes were made at the slow rate of 12 frames per second.

Several explanations of change blindness have been proposed in the psychological literature (Simons 2000). The traditional explanation is the “overwriting” hypothesis (Simons 2000, 8; Frey et al. 2024). According to this, the new representation – that is, representation of the

changed scene – replaces the old representation and thus in a sense writes over the old one. Since the old representation is no longer stored in visual short-term memory (VSTM), there can be no comparison between the old and new representations, and therefore no detection of change. Alternative explanations are the “nothing is compared” hypothesis, the “feature combination” hypothesis, and the “first impressions” hypothesis (Simons 2000, 10). I will discuss the first impressions hypothesis in more detail since, as will become clear, it is well suited to explain change blindness in natural perceptual conditions.

The first impressions hypothesis starts with the assumption that the “primary goal of perception is to understand the meaning and importance of our surroundings” (Simons 2000, 9). This extraction of the meaning or “gist” (Rensink 2000, 36) of the scene is done by the initial representation. In other words, a visual representation primarily represents the gist of the scene. For example, upon entering a room, the visual representation of the scene extracts the gist whether it is an educational gathering, or a religious or a social gathering, etc. The gist “...can provide a stable constraint on the kinds of objects expected, and perhaps even indicates their importance for the task at hand” (Ibid., 36). Clearly, the gist is a general element that is part of the presentational content of a visual representation.

Rensink also argues that visual representations represent the layout or spatial arrangement of the scene. The spatial information extracted is from the allocentric frame of reference – that is, from a fixed world-indexed coordinate system that is independent of the subject’s sense of space (egocentric reference frame). Since the layout is invariant with respect to the subject, “it could be useful for directing eye movements and attentional shifts” (Rensink 2000, 37). In terms of the Gestalt principle of figure-ground distinction (Wagemans et al. 2012, 1194), the gist and the layout are the ‘ground’ information, and so are general elements.

Once the gist and the layout of a scene have been represented, any representation of the specific details is the result of attention, which is a post-perceptual selection of properties from the perceptual content. In other words, representation as of particulars is the result of attention. This is why when attention is masked, changes to details of the scene are not detected, because they are not represented. A change of a specific detail in the scene does not alter its gist or layout; therefore, there is no new perceptual representation of the scene.

The first impressions hypothesis highlights the role of expectations in change detection. Since the observer *expects* the gist of the scene to remain the same and, relatedly, does not expect any specific detail to change, changes in details are not detected (Simons and Mitroff 2001, 203).

Further, the first impressions hypothesis better explains “incidental changes” (Simons 2000, 8), where the observer does not expect any significant changes to the scene (Rensink 2002, 257). It is also a better explanation of gradual changes, because in slow-changing scenarios, observers do not expect the gist of the scene to change drastically. The first impressions hypothesis is, therefore, particularly relevant to explain change blindness in natural perceptual conditions, because incidental gradual changes better approximate natural perceptual conditions.

5.2.2.2.2. Implications on content particularism vs generalism debate

Change blindness studies provide empirical support for content generalism over particularism. According to the first impressions hypothesis, we perceptually represent the gist of the scene rather than the details. Now, the gist of the scene is the general significance of the scene for the perceiver. Thus, it is an abstraction over the details and therefore includes general, rather than particular, entities. The general gist is an aspect of the presentational content of the representation. Thus, perceptual content is general. Therefore, it is more plausible that perceptual representations are constituted by the general gist and spatial layout of the scene rather than by discriminating particulars.

5.2.2.3. Premise III: Constitution of content by particulars

The premise says that if perceptual state M is constituted by discriminating and singling out particular α , then M 's content is constituted by α .

In Premise III, Schellenberg moves from the claim that a perceptual state is constituted by *discriminating* a particular to the claim that the perceptual content is constituted by the *particular*.

Schellenberg's argument for this premise argument begins by noting that discriminating and singling out a particular is the function of perceptual capacities. Although she does not explicitly mention it, it is very plausible that she subscribes to the selected effects (SE) account of functions. This is because she is a naturalistic representationalist, and the SE account of function is the default option for naturalistic representationalism. To recall from chapter 3 (§3.3), according to the SE account, a trait/capacity's function is its effect or output which has been selected – mostly by natural selection, but which could also be selected by cultural selection and learning. Neander expands the scope of the SE account and writes that the function of a trait is not only its selected effect but also its selected causal disposition to produce

the effect in response to causes (Neander 2017, 130). In other words, what is selected is the triad of cause-capacity-effect.

Schellenberg seems to accept Neander's interpretation expansion of the notion of SE function. Accordingly, she argues that "...if singling out a particular has any significance, then the subject's perceptual state is constituted by the particular when she perceives that particular. To think otherwise would be to sever the link between the function of the capacity and its output" (Schellenberg 2018, 26). Thus, the function of a perceptual capacity is not only the perceptual representation (the effect) that is the result of discriminating and singling out particulars, but also the causal disposition to discriminate and single out particulars (the cause). Therefore, if a perceptual representation is constituted by the capacity to single out particulars, then its content is constituted by the capacity and consequently by the cause of the capacity – the particulars singled out.

However, Schellenberg's argument can be resisted by challenging the idea that selected effects function of perceptual states is to discriminate particulars. To my mind, an emphasis on selected functions shows that perceptual capacities have the function of representing and discriminating *general* properties, rather than particular entities. This is because perceptual capacities can be employed by present organisms to discriminate entities that would never have been encountered during the historical selection of their function. If the selected function was to discriminate only particulars that were naturally encountered during the selection process, then perceptual capacities would be incapable of functioning to discriminate novel or non-natural (human-made) entities. Thus, content is constituted by general entities, not particulars.

Moreover, as the discussion on Multi-Object Tracking (MOT) above established, discrimination of particulars, by means of individuating them, is not a perceptual affair, but post-perceptual, plausibly attentional, one. Thus, the particularity of content can be explained in terms of phenomenological particularity rather than relational-constitutional particularity.

5.2.2.4. Summing up:

To sum up my objections to the particularity argument: Premise I can be refuted since tracking and discriminating particulars does not seem to be the function of perceptual states. Rather, research on MOT shows that it is the function of a post-perceptual cognitive mechanism, plausibly attention. Premise II can be refuted since perceptual states are not constituted by discriminating particulars. Rather, research on change blindness indicates that they are constituted by capacities for representing general features of the environment – the gist of the

target scene or likely causes of sensory input. Premise III can be refuted because perceptual states and their content are not constituted by particulars discriminated and singled out. Therefore, content particularism is false – content is not constituted by environmental particulars.

5.3. Conclusion

In this chapter, I put the finishing touches to my argument against reductive-naturalistic representationalism that began in the previous chapter. I argue that the primacy to the theory of application over a theory of content is not genuinely representational. I also highlighted the lack of causal efficacy of content. Finally, I argued against content particularism. Based on my arguments in this chapter, the following are the conclusions that I will adopt in formulating a novel version of nonreductive representationalism in chapters 7 to 9. First, primacy to a theory of content. The representational application/directedness ought to be explained in virtue of content. Second, content generalism – perceptual content is constituted by general entities. I will argue for a specific version of generalism that I will call content universalism – content is partly constituted by co-instantiation of concrete universals, which are general entities by definition. Third, robust causal efficacy to content, which I will argue can be accounted for the strong emergence of content.

6. Nonreductive Representationalism

In this chapter, I survey and assess nonreductive representationalist theories. Almost all such theories are characterized by their rejection of content-*constituting* (rather than content-*causing*) naturalistic relations to the environment. In §6.1, I note the main motivation of nonreductive representationalism, which is a solution to the problem of non-veridical perception. I also note their basic commitments, according to which they may be broadly classified into two groups – con-relational and non-relational representationalism. §6.2., then, gets to the task of assessing con-relational representationalism. I discuss the theories of Pautz and Chalmers and highlight some weaknesses in their views. In §6.3, I assess Crane’s non-relational representational theory. Following this, in §6.4, I note some general weakness of both con- and non-relational versions.

6.1. Nonreductive Representationalism

Non-reductive representationalism focuses on the theory of content – how content is constituted. It argues that content cannot be reduced to physicalist or naturalistic elements, as proponents of reductive-naturalistic representationalism claim.

6.1.1. Motivations for Nonreductive Representationalism

The main motivation for non-reductive representationalism is a solution to the problem of non-veridical perception – that is, the problems of illusion and hallucination. The problem, roughly, is that the occurrence of non-veridical perception puts pressure on the intuitive understanding of perceptual representation, which involves relations to environmental targets. More carefully, the problem of non-veridical perception arises out of the inconsistency of three very plausible theses – Direct Realism, the Phenomenal Principle and the Common Kind Assumption.

Direct Realism (DR) is the thesis that there is a mind-independent concrete reality (realism), and we can directly perceive it. It should be noted that direct realism is compatible with the perception being mediated by naturalistic relations and internal neural processing, because the relations and the processing themselves are not perceived. They are they are processual rather than perceptual intermediaries. Thus, “what direct realism affirms is that perceptual experience of physical reality does not depend on *perceiving* a mediating mental reality” (Genone 2016,

3; emphasis mine). A contrast case is sense data theory which is an indirect realist theory – perceiving external objects depends on perceiving a sense datum.

The Phenomenal Principle (PP), as Robinson puts it, is that “if there sensibly appears to a subject to be something which possesses a particular sensible quality then there is something of which the subject is aware which does possess that sensible quality” (Robinson 1994, 32). In other words, the phenomenal principle prescribes a transition from perceptual experience to an object that is experienced. Therefore, Pautz calls this the “act-object assumption” (Pautz 2021, 27), while clarifying that “...it is left open what those [perceived] objects are (physical objects, mental images, whatever)” (Ibid., 27).

The Common Kind Assumption (CKA) states that “...whatever kind of mental, or more narrowly experiential, event occurs when one perceives, the very same kind of event could occur were one hallucinating” (Martin 2006, 357). In other words, veridical, illusory and hallucinatory experiences belong to the same fundamental kind, which is perceptual experience.

The above three principles form an inconsistent triad. This can be shown in the argument from hallucination:

1. In an hallucination “you have an experience as of an object and its properties but there is no (worldly) object, and there are no (worldly) properties, that you perceive in virtue of having that experience” (Macpherson and Batty 2016, 265).
2. If there appears to a subject to be something which possesses a particular sensible quality, then there is an object of which the subject is aware which does possess that quality (the phenomenal principle).

Therefore,

3. In hallucinations, the object of which the subject is aware is not an everyday worldly object.
4. Veridical, illusory and hallucinatory experiences are of the same kind (the common kind assumption).

Therefore,

5. Even in veridical perception, the objects of which the subject is aware are not everyday worldly objects.

Therefore,

6. Conclusion: We never directly perceive everyday worldly objects (~Direct Realism).

The conclusion of the argument from hallucination is that direct realism, which is plausibly the most common intuition about perception, is false. The above argument is used by sense data theory to argue that the immediate objects of perception are sense data, and any perception of worldly objects is only indirect. However, the conclusion can be resisted by denying either the phenomenal principle or the common kind assumption. Whereas disjunctivism (particularly naïve realistic disjunctivism) opts for a denial of CKA, representationalism denies the phenomenal principle.

A denial of the phenomenal principle implies that perceiving concrete environmental objects is not essential to perception. Consequently, being related to them is not essential to or constitutive of perceptual representations, and therefore, perceptual content.

6.1.2. Basic Commitments

Non-reductive representationalist views do not posit content-constituting relations to environmental targets. As per the terminology introduced in chapter 2 (§2.3.5.2), they are either con-relational or non-relational. Con-relational representationalism accepts either content-constituting relations (for example, acquaintance) between subjects and abstract universals, or content relations to abstract propositional contents. Moreover, these relations are considered to be primitive and irreducible to naturalistic relations. Non-relational representationalism holds that content is a non-relational property or aspect of the representation or subject, and therefore does not accept any relation.

Three points need to be clarified about nonreductive representationalism of both stripes. First, I include phenomenal intentionality theory (PIT) within the ambit of nonreductive representationalism since it accounts for the intentionality of perceptual states in terms of content which is grounded in phenomenal character. Moreover, phenomenal character is in turn not reductively explained in terms of naturalistic elements. Thus, it does not aim to reduce representation to naturalistic elements. Second, perceptual representations do not *constitutively* or *essentially* (i.e., as part of their essence or nature) involve relations to environmental targets. This is compatible with the acceptance of merely *causal* relations to targets. In other words, content may be partly caused by environmental targets but is not constituted by the targets or relations to them.

Third, when it is claimed that content is not constituted by relations to targets or that content is a non-relational property, what is implied is only that representations do not represent in virtue

of bearing relations. This is compatible with representations bearing relations to targets in virtue of having content. In other words, content grounds the representational external directedness (RED) relation (chapter 2, §2.2.2) between representations/subject and the target. Thus, Kriegel – a non-relational representationalist – writes that

“...a non-relational account of intentionality does not imply that representors do not bear relations in virtue of representing. What it implies is rather the converse: that representors do not represent in virtue of bearing relations. Thus it may well be that when a representor represents an existent, it enters a relation to it, and enters that relation in virtue of representing it. This relation we may well call “the representation relation”” (Kriegel 2007, 315).

There is a significant implication that follows from the denial of content-constituting relations to the environment – non-reductive representationalist views reject content particularism, the thesis that perceptual content is (partly) constituted by environmental particulars in virtue of being related to them (chapter 5, §5.2). Consequently, it subscribes to content generalism – content is constituted by general elements (for example, abstract universals or modes of presentation) and, therefore, perceptual content is general content. “There are many different ways of understanding general contents. They can be thought of as *de dicto* modes of presentation, Russellian propositions, or existentially quantified content to name just a few options” (Schellenberg 2018, 77).

Following their rejection of content particularism, nonreductive representationalism holds that the particularity that characterizes perceptual content is *not* relational-constitutional particularity, but only phenomenological particularity (chapter 2, §2.3.5.1) where it only seems to the subject that he is perceiving particular objects with their particular properties. Regarding individuation of content, since they do not accept constitutive relations to the environment, most non-reductive representationalist views espouse internalism about content. Content is individuation dependent entirely on the intrinsic properties of subjects; in other words, perceptual content – more precisely, content of perceptual *experiences* – is narrow content.

In sum, it is common for non-reductive representationalism to espouse either of the following two combinations:

- 1) Con-relationalism + content generalism + phenomenological particularity + content internalism; or
- 2) Non-relationalism + content generalism + phenomenological particularity + content internalism.

6.2. Con-relational Representationalism

Con-relational nonreductive representationalism includes views that endorse a constitutive relation between the subject and (i) abstract universals; or (ii) propositional content; or (iii) both abstract properties and propositional contents. The relation to abstract universals is a content-constituting relation, and constitutes presentational content. The relation to propositional contents is a content relation (Chalmers 2006, 51). To reiterate, these relations are *not* the representational external directedness (RED) relation towards external targets. Thus, with respect to the content relation, Crane cautions that it “...should not be confused with the ‘perception relation’ [i.e., the RED relation] which holds between a subject and a pig when [we] perceive a pig” (Crane 2014, 242).

The relationship between the various relations in con-relational representationalism may be depicted as below:




Figure 6.1: Relations in Con-relational Representationalism

In what follows, I will discuss the con-relational representationalist views offered by Pautz (2017, 2021) and Chalmers (2004, 2006).

6.2.1. Pautz's Internalist-Nonreductive Representationalism

Pautz (2021) is a prominent proponent of con-relational representationalism. He calls his version of the view “internalist-nonreductive representationalism” (Pautz 2021, 170), which he describes as follows:

“Representational view. All sensory-perceptual experiences consist in a basic mental relationship *R* between *subjects* and *ways things might be* such that: *R* plays the cognitive-access role, *R* is existence-neutral, and *R* plays the character-role” (Pautz 2021, 100).

Pautz labels the *R* relation between subjects and the ‘ways things might be’ as the *experiential representation* relation (Ibid., 100). Elsewhere, he calls it the *conscious-of* relation (Pautz 2017, 387). The ‘ways things might be’ in the above description are “*abstract* way things might be” (Pautz 2021, 98). More specifically, they are “...a complex array of [abstract] perceptible properties (shapes, distances, colors, and so on)” (Ibid., 99) which are attributed to external targets. In other words, the subject is related by the experiential representation relation to Platonic universals.

The experiential representation (ER) relation and the abstract universals together constitute the presentational content of a perceptual representation. Thus, the ER relation to abstract universals is a *content-constituting relation*, and not the representational external directedness (RED) relation to concrete targets. The content that is constituted by the ER relation is a “phenomenal content” (Pautz 2009, 494), and the subjects stand in the “sensory entertaining” relation to this phenomenal content. Thus, “...‘the phenomenal content of token-experience *e*’ denote[s] the proposition which one sensorily entertains in having *e*” (Ibid., 494). So, while the ER relation is a content-constituting relation, the sensory entertaining relation is a content relation.

Pautz argues that the ER/conscious-of relation plays several roles. He also puts this in terms of perceptual experience having various *significances* to perceivers (Pautz 2017, 349). First, the ER relation plays the character role. The ER relation, in conjunction with the properties related to, constitutes the phenomenal character of perception states. “[T]he hypothesis is that having an experience with a certain character consists in nothing but experientially representing a certain array of perceptible properties (ways things might be), so that differences in character are always differences in the array of perceptible properties” (Pautz 2021, 100).

Although Pautz does not explicitly address the particularity of content, the character role of the ER relation can also be extended to conferring particularity to content. Pautz calls this the dissimilarity-grounding significance (Pautz 2017, 390) of the ER relation. According to this role, the ER relation can account for the dissimilarities between conscious and unconscious perceptual states. Thus, Pautz writes “the conscious-of [i.e., ER] relation is *unique*: there is no relation R^* that is not a form of consciousness but that is intrinsically similar to the conscious-of relation” (Ibid., 390). For example, when you experientially represent redness, the ER relation to is “necessarily totally unlike your relation to any quality that you are *not* conscious of” (Ibid., 349).

It seems reasonable to conjecture that the ER relation can ground even the dissimilarities between different conscious perceptual states; for example, the dissimilarity between two red experiences. This is because there is a unique ER relation borne to different complexes of properties, or even the same complex. Each fact of ‘ER relation holding between subject and property-complex’ would be a distinct fact, dissimilar from each other. This dissimilarity-grounding role implies that each unique ER relation confers particularity on the property-complex related to, since the property-complex – being composed of abstract universals – is a general entity. And there can be no dissimilarity without first individuating/particularizing the compared entities. Hence, it is the ER relation that accounts for phenomenological particularity.

Secondly, the ER relation plays the cognitive-access role since it grounds the ability of perceivers to have thoughts and beliefs about perceived content. This cognitive significance has two aspects: reasons-grounding significance and determinacy-grounding significance. As regards the reasons-grounding role, Pautz writes that, owing to the ER relation, “having an experience with a certain phenomenology is *sufficient for* having a reason to believe certain things (Pautz 2017, 365). The determinacy-grounding role of the ER relation enables perceivers to have “...beliefs with very *determinate* contents about qualities and other ostensible items” (Ibid., 370).

Finally, the ER relation is existence-neutral since “one can experientially represent a property...so that it *seems* that there exists something with the property, even if in actuality there exists nothing that has that property” (Pautz 2021, 100). Thus, the ER relation can account for hallucinations.

Pautz’s view is internalist in the sense that the experiential representation relation to the abstract perceptible or sensible properties is “...*fully* determined by internal neural states” (Pautz 2021,

172). However, the brain does not generate the sensible properties or objects with sensible properties. “Rather, it merely generates “seeming-states” in which it seems that there are such objects” (Ibid., 173). Summing up his view, Pautz writes: “[y]our brain has an intrinsic capacity to generate richly externally directed experiences in which you stand in an irreducible experience relation to perceptible properties..., even if those properties don’t occur in the brain or indeed anywhere at all (Ibid., 182). Further, Pautz argues that the experiential representation relation is “...irreducible to any physical relationship, even if it depends on the internal physical state of the subject’s brain” (Pautz 2021, 173).

Pautz’s main argument for internalist representationalism is an inference to the best explanation. He argues that it best explains the two significant features of perceptual experience – essential external directedness and internal dependence. To recall from chapter 2 (§2.2), perceptual states are *essentially* directed at states of affairs located in (or seeming to be located in) spacetime. This is the intentional directedness of perceptual states towards an external environment. Internal dependence, on the other hand, is the claim that perceptual experiences as of sensible properties is dependent on neural processing in the brain and central nervous system, more generally.

The combination of the two features gives rise to what Pautz calls the external-internal puzzle about sensible properties:

“External-internal puzzle about sensible properties. Even in totally normal perception, and not just in illusion and hallucination, your experience of sensible properties (pain qualities, smell qualities, color qualities) is shaped by internal processing. Yet you experience sensible properties as “out there”, together with shapes and in various locations, often at a distance from you. How is it that what you seem to experience as “out there” is shaped by internal processing “in here”?” (Pautz 2021, 9).

Pautz argues that reductive-naturalistic representationalism – which he calls “response-independent representationalism” (Pautz 2021, 142) – can easily account for the external directedness of perceptual experiences since they hold that it consists in tracking⁷⁰ (or detecting) external properties (Ibid., 145-146). However, they cannot adequately account for internal dependence. He cites empirical evidence to show that sensible properties are better correlated with internal neural processing than with external physical properties that are

⁷⁰ As discussed in chapters 3 and 4, the tracking relation is reducible to naturalistic relations, functions and representational application formulae.

detected (Ibid., 152-160). Thus, perceptual experience is characterized by “good internal correlation”, but “bad external correlation” (Ibid., 160), which reductive-naturalistic representationalism cannot explain equally well.

In contrast, internalist-nonreductive representationalism can account well for both internal dependence and essential external directedness. Internal dependence is accounted for in terms of dependence on neural processing. In fact, Pautz argues that internalist-nonreductive representationalism supports “across-the-board *experiential internalism*: all aspects of experience are *fully* determined by internal neural states” (Pautz 2021, 172) – that is, experience of not merely phenomenal qualities, such as colour and smell, but also spatial properties. Thus, he says that even a brain in a void could have experiences as of spatial properties and of phenomenal qualities in space. An implication of this internal dependence is that the phenomenal qualities that we experience in perception are not in the world; they are illusions of sorts:

“...internalist-nonreductive representationalism...is uniform in two respects. First, it upholds a uniform internalism about experience: the experience of both colors and shapes is internally generated by the brain without help from the world. Second, it upholds a uniform illusionism: neither colors-as-we-experience-them nor shapes-as-we-experience-them are really in the world” (Pautz 2021, 181).

The external directedness is explained by the experiential representation (ER) relation: “...for you to have visual experiences *just is* for it to seem to you that there are things arranged in space – that is, for you to “experientially represent” that there are things arranged in space” (Pautz 2021, 103). Pautz likens it to the brain projecting the sensible properties out into the world.

“Our brains invented the quality blue and other colors, and enabled us to experientially represent them as occurring in an antecedently colorless world, in order to help us recognize and remember objects. As Cosmides and Tooby (1995, xi) put it, “color is an invention that specialized circuitry computes and then projects onto physically colorless objects”” (Pautz 2021, 171).

6.2.2. Con-relational Phenomenal Intentionality Theory (PIT) – Chalmers:

Many phenomenal intentionality theorists also argue that perceptual representation involves either content-constituting relations to abstract universals or content relations to propositional contents. Chalmers (2006, 2013) presents a version of phenomenal intentionality theory, which

posits both kinds of relations. According to Chalmers, the constitution of representational content happens in two stages:

6.2.2.1. Constitution of phenomenal character:

The phenomenal character of a perceptual state involves being *acquainted* with primitive properties, which are abstract uninstantiated universals. He calls such primitive properties “Edenic properties” (Chalmers 2006, 66). (Note: Chalmers uses ‘properties’ to mean universals, and property-instances to mean instantiations of those universals; I will follow his usage).

Eden is a hypothetical world where primitive or perfect properties are instantiated in objects. These are Edenic properties. The actual world (our world), in contrast, is a non-Edenic world, where the properties instantiated by objects are imperfect properties that only match Edenic properties. To match an Edenic property is to cause a perceptual state that involves acquaintance with the corresponding Edenic property – that is, a perceptual state with phenomenal character. For example, if the Edenic property is perfect redness, then the non-Edenic imperfect property is imperfect redness, which is “plausibly some sort of physical property, such as a surface spectral reflectance” (Chalmers 2006, 72). This surface reflectance property “...matches perfect redness if it normally causes phenomenally red experiences” (Ibid., 72).

Thus, Chalmers writes:

“Phenomenal redness is...the property of representing primitive redness, or perhaps the property of being aware of primitive redness. In Eden, this property is grounded in and perhaps realized by acquaintance with *instances* of primitive redness... Outside Eden, this property is grounded in *acquaintance with the universal* primitive redness: it is in virtue of being acquainted with this universal that non-Edenic subjects are aware of primitive redness” (Chalmers 2013, 350; emphasis mine)

In other words, according to Chalmers, the properties instantiated in the actual world are imperfect universals. But, in perceptual experiences, we are acquainted neither with the imperfect property-instances nor the imperfect universals, but we are acquainted with perfect Edenic universals. It should be noted that Chalmers uses acquaintance, awareness and representation equivalently. To have phenomenal character, therefore, is to be related by acquaintance to primitive Edenic universals.

The acquaintance with Edenic universals constitutes “Edenic content” (Chalmers 2006, 70), which is the presentation of Edenic universals. Edenic content determines phenomenal character. For example, acquaintance with Edenic redness constitutes Edenic content, which determines phenomenal redness (the experience of the phenomenal quality ‘red’). Thus, the major role of Edenic content is to determine phenomenal character. Edenic content may be associated with “perfect veridicality conditions” (Ibid., 69). Perfect veridicality conditions are satisfied only if Edenic universals are instantiated in the world. Thus, in our world, where only imperfect universals are instantiated, Edenic content cannot really be satisfied. However, Edenic content has a regulative role, in that it determines representational content.

6.2.2.2. Constitution of representational content:

Chalmers claims that Edenic content determines ordinary (i.e., non-Edenic) representational content (Chalmers 2006, 72). Representational content itself is of two kinds – narrow Fregean content and wide Russellian content. The Fregean content is a mode of presentation of the imperfect property-instances of objects. It plays the role of an intension that yields Russellian content, which are propositions that include the imperfect property-instances and particular objects as constituents. Both are associated with satisfaction conditions:

“the ordinary Fregean content of a phenomenally red experience will be satisfied (in an environment) iff a relevant object instantiates a property that matches perfect redness (in that environment). This ordinary Fregean content will itself be associated with an ordinary Russellian content: one that is satisfied iff the (actual) object of the experience has P, where P is the property that matches perfect redness in the environment of the original experience” (Chalmers 2006, 73).

The satisfaction conditions (i.e., semantic content) of perceptual experiences are “imperfect veridicality conditions” (Chalmers 2006, 72) because they are satisfied by imperfect properties that only match the perfect Edenic properties. Moreover, Chalmers is a pluralist about perceptual content and argues for perceptual representations having both particular (Russellian) and general (Fregean contents):

“It may be that experiences can be associated with contents of many different sorts by different...*content relations*. For example, there may be one content relation that associates experiences with object-involving contents, and another that associates experiences with existential contents. Each of the different sorts of content, and the corresponding content relations, may have a role to play for different explanatory purposes” (Chalmers 2006, 51).

6.2.3. Objections to Con-relational Representationalism

A major objection to con-relational representationalism pertains to the content-constituting relation that it posits to abstract universals. This is the ‘acquaintance’ relation under different labels. The relation of acquaintance was first posited by Russell (1911), and it is seeing a resurgence in contemporary philosophy of perception and consciousness. Acquaintance can be defined as “...a relation of conscious awareness that is fundamentally distinct from thinking a true thought or forming an accurate judgement, in which the mind has some kind of unmediated confrontation with some portion of reality” (Raleigh 2019, 7).

As noted above, Chalmers (2013) explicitly calls the relation to universals ‘acquaintance’. Although Pautz calls his content-constituting relation ‘experiential representation’, it is clear from the role he gives it that it is acquaintance. In fact, in another paper (Pautz 2017), calls this relation the “conscious-of” relation and speaks of it in the same breath as acquaintance (Pautz 2017, 387-388).

Representationalists across the reductive-naturalistic/nonreductive divide invoke acquaintance to explain the phenomenal character of perceptual experiences. Acquaintance with phenomenal qualities constitutes the presentational content of perceptual representations, which explains their phenomenal character. An exemplar of naturalistic representationalism who invokes acquaintance is Tye (2009a, 2019). Tye describes acquaintance with entities as “encountering...[them] in experience” (Tye 2009a, 101), and writes that,

“The phenomenal character of an experience is out there in the world (or in the body, in the case of bodily sensations)...In being aware of the external qualities, we are aware of phenomenal character. We are confronted by it...[N]ecessarily any change in the external qualities with which we are acquainted in undergoing a given experience generates a change in the phenomenal character of the experience” (Tye 2009a, 119).

The main argument for positing acquaintance with abstract universals is the argument from transparency combined with the need to account for hallucinations. Transparency is the thesis that when we introspect on our perceptual experiences, rather than being aware of properties or features of the experience itself, “[t]he only features of which we are aware and to which we can attend are external features (colors and shapes of surfaces, for example)” (Tye 2021, 33). Tye calls these external properties “sensible property complexes” (Tye 2021, 58). In other words, our experiences are transparent to introspection, which directly gets at external property

complexes. Given the centrality of these property complexes, Tye calls his view “property representationalism” (Ibid., 71).

Similarly, Chalmers writes that transparency is the thesis according to which “(i) our perceptual phenomenal states consist in phenomenal awareness of properties (such as Edenic redness) attributed to the external world, (ii) our core acts of attention are to these properties and (iii) we can also attend to our awareness of these properties, thereby attending to our phenomenal states” (Chalmers 2013, 348)⁷¹.

The further claim that the external qualities we are acquainted with are abstract universals is to give a uniform account of acquaintance for all perceptual experiences. In the case of veridical experiences and illusions, the universals are instantiated (Tye 2021, 66); whereas in the case of hallucinations, where it is supposed that there are no external objects that can possess any properties, the universals are uninstantiated. Thus, Tye writes that “[a]long with (most) other representationalists, I am happy to say that, in the hallucinatory case, the perceiver is conscious of an un-instantiated property” (Tye 2014, 304).

There are several objections to the twin theses of acquaintance with universals and transparency. The main objections to transparency are counterexamples such as blurs. Blurriness is intuitively a feature of a visual experience. When we introspect on a blurry visual experience, we are not aware of any mind-independent blurriness property. Thus, Crane writes that:

“When I say that ‘everything seems blurry’ I don’t mean that it seems as if the things around me are blurry...What I mean is that I am experiencing things in a blurry way. Isn’t this a straightforward case of where one can be ‘directly’ aware of an aspect of one’s experience which is not an aspect of the objects of experience?” (Crane 2006, 130).

Tye’s response to the above counterexample is bit evasive. He writes:

“I agree that when I introspect a blurry experience of an object *O*, I am indeed not turning my attention to blurriness as a feature of *O*. But equally I am not turning my attention to blurriness as a feature of my experience...A so-called ‘blurry experience’ is really an experience that

⁷¹ A related, but stronger, thesis that is popular among a few non-representationalists is “revelation” (Johnston 1992, 223). Revelation is the thesis that introspection and attention not only show that experiences are transparent, but also reveal the nature of perceptual qualities we are acquainted with. As Brewer puts it, “Revelation is the thesis that visual experience whose conscious character is specified by its looking to the subject as though there is an *F* before her is a source of knowledge of what being *F* is, knowledge expressible in this context, and with appropriate attentional direction, by ‘being *F* is being like that’” (Brewer 2019, 279).

represents boundaries of objects indistinctly. When I introspect such an experience, I am simply aware that I am seeing blurrily” (Tye 2021, 42).

Tye’s suggestion seems to be that blurriness is a kind of perceptual judgement (“I am aware *that*”), rather than experience, that we form upon introspection. However, this goes against the phenomenology of blurriness, which presents it as an aspect of perceptual content.

Coming now to the objections to acquaintance with universals. First, the thesis that we are acquainted with uninstantiated universals in hallucinations does not seem to square with the phenomenology of experience. It phenomenologically seems to a hallucinating subject that a concrete object with its properties is spatiotemporally present outside. It does not seem to us that we are perceiving non-spatiotemporal property-complexes. And, acquaintance is a relation of direct ‘encountering’ or ‘confrontation’ with the acquainted. Thus, Duncan writes that the challenge is explaining the perceptual presence of things:

“[P]erceptual representations (even hallucinations) come with a sense of object presence – things really seem to be there, present to the senses, available for attention and demonstration. But it is hard to make sense of this if perceptual representation is just acquaintance with universals, rather than actually instantiated properties” (Duncan 2021, 10).

My second objection is that it is hard to make sense of a conscious awareness relation (which acquaintance, at its core, is) to abstract entities. To be sure, many philosophers claim that we can be related to abstract entities such as propositions and numbers (Liggins 2024, 17). However, such relations are usually considered to be *epistemic relations* such as grasping or believing. Epistemic relations, in contrast to perceptual relations, are relations involved in knowledge. For example, grasping is a relation between a subject and a proposition, which is considered essential to understanding (Strevens 2025). Frege, who was influential in positing the relation of grasping, held that “the reality of thoughts [i.e., propositions] depends on their being grasped, because it is only in so far as they are grasped that they can act on us and have in this way some degree of actuality” (Sacchi 2006, 77)⁷².

Acquaintance, in contrast, is a conscious awareness relation. The phenomenal character of a perceptual state which involves acquaintance is partly dependent on the relation and partly on

⁷² Frege was a primitivist about grasping and considered it unanalyzable (Sacchi 2006, 77). However, there have been several recent theories of grasping. For instance, the inferential theory (Newman 2012) – grasping P consists in the inferences that P can participate in; the phenomenal theory (Bourget 2017) – grasping P consists in having a phenomenal experience with the content P; and the ability theory (Strevens 2025) – grasping P consists in the ability to use knowledge of P in varying circumstances.

the entities that it relates. Now, if the claim is that we can be acquainted with abstract universals, this implies that abstract universals partly determine phenomenal character. But it was argued in chapter 2 (§2.3.4) that abstract content is incapable of explaining concrete phenomenal character – this was the content-character mismatch problem. This problem generalizes to all abstract entities, including abstract universals. They do not seem to be capable of explaining phenomenal character. This puts pressure on the claim that we can be acquainted with such putatively phenomenology-determining abstract universals. In other words, acquaintance is a relation that relates concrete entities.

The intuition that we can be acquainted only with concrete entities is also evidenced in the original usage of ‘acquaintance’ by Russell in the context of “knowledge by acquaintance” (Russell 1910; Russell 1912/1998, chapter 5). Russell defines acquaintance as a direct awareness relation: “we have acquaintance with anything of which we are directly aware, without the intermediary of any process of inference or any knowledge of truths” (Russell 1912/1998, 25). Further, “to say that S has acquaintance with O is essentially the same thing as to say that O is presented to S” (Russell 1911, 108). Russell claims that we are acquainted with several kinds of entities – particular objects sense-data, mental states and thoughts, memories, and abstract Platonic universals (Ibid., 26-28). Most of the entities that we are acquainted with are concrete entities. One entity stands out – abstract universals.

Even though Russell claims that we can be acquainted with abstract universals (Russell 1912/1998, 58), this claim can be challenged. This is because, according to Russell, we are ‘acquainted’ with a universal through a process of abstraction. He expands on this in the case of the universal property ‘whiteness’:

“When we see a white patch, we are acquainted, in the first instance, with the particular patch; but by seeing many white patches, we easily learn to abstract the whiteness which they all have in common, and in learning to do this we are learning to be acquainted with whiteness” (Russell 1912/1998, 58)

The so-called acquaintance with universals is through abstraction, which is mediated by genuine acquaintance with the particular white instances. Therefore, the former does not seem to be genuine acquaintance at all. The burden is on acquaintance theorists, then, to explain how we can have a conscious perceptual relation to abstract entities, especially since it is further held that such a relation can account for concrete phenomenal character.

My third objection to acquaintance whether to universals or concrete entities is that it cannot account for unconscious perception. By definition, acquaintance is a conscious awareness relation. But, as discussed in chapter 2 (§2.3.1), it has been empirically established that there is unconscious (eg. blindsight) and pre-conscious (eg. early-vision) perception. Clearly, acquaintance cannot ground such perception.

Given the above problems with acquaintance and transparency, it seems unlikely that perceptual representation involves either of the two theses. That is, perceptual content is not constituted by transparent acquaintance relation to abstract universals.

6.3. Non-relational Representationalism

The second type of nonreductive representationalism is thoroughly non-relational and holds that content is a modification, property or aspect of the representational state or of the subject. As discussed in chapter 2 (§2.3.5.2), there could be two versions of non-relational representationalism – aspectual and hyphenist. Content is a concrete property and, therefore, non-relational representationalism is a form of concretist representationalism that Pautz mentions as an alternative to his abstractivist representationalism (Pautz 2021, 129).

6.3.1. Crane's Non-relational Aspectual Representationalism

A prominent proponent of non-relational aspectual representationalism is Crane (2006, 2013), in the case of both perceptual states and non-perceptual thoughts. Crane argues for non-relational representationalism within the larger framework of psychologism. He describes psychologism as the thesis according to which mental phenomena, particularly intentionality, must be understood and explained in terms of *irreducible* psychological entities and facts. Psychologism does not attempt an *ontological* reduction of intentionality to other psychological entities; rather it aims merely for an *explanatory reduction* (Crane 2013, 134). This is important because he eventually argues that mental representation is fundamental and irreducible (Ibid., 115).

Psychologism is opposed to anti-psychologism, which explains intentionality/representation in terms of non-psychological entities. These non-psychological entities are of two types: first, semantic entities such as referential and propositional content (accuracy conditions) and external targets (the accuracy makers) (Crane 2014, 4); and second, epistemic or metaphysical

entities, particularly special epistemic and metaphysical relations – such as acquaintance – to abstract universals (Crane 2013, 147-152). Strictly speaking, psychologism about intentionality is neutral between accepting a psychological relation to content – a con-relational account – and embracing a non-relational aspectual view of content. Examples of psychological relations are propositional attitudes such as believing, desiring, etc.; some like Byrne (2009) would even count perceiving as a propositional attitude. Crane favours the latter non-relational view. This is because a non-relational view best accounts for mental states that are about entities that do not exist such as hallucinations, and thoughts about fictitious entities such as Santa Claus. Crane calls such entities “non-existing intentional objects”⁷³ (Crane 2013, 94).

To make sense of non-existent intentional objects, Crane invokes Schiffer’s notion of “pleonastic entities” (Schiffer 2003). Pleonastic entities are entities that can be *inferred* from truths and facts. Schiffer gives the example of fictional entities: “...from the fact that (1) Joyce wrote a novel in which he used ‘Buck Mulligan’ in the pretending way characteristic of fiction, we may infer [that] (2) Joyce created the fictional character Buck Mulligan” (Schiffer 2003, 51). Schiffer calls such inferences “*something-from-nothing transformations*” (Ibid., 52) since they take us from a statement in which no reference is made to an object to a statement in which there is a reference to an object. Pleonastic entities are, therefore, non-existing intentional objects and do not carry ontological weight.

Crane applies this notion to thoughts about non-existing entities and argues that from the truth of statements such as ‘I think of Pegasus, the mythical creature’ we can infer ‘There is *something* of which I think, and that thing is a mythical creature Pegasus’. Such “inferences are valid,...yet the entities we accept by accepting them...are no addition to our ontology” (Crane 2013, 65). The same inference holds for hallucinations: from the truth of ‘I had an hallucination as of a pink elephant’ we can infer ‘There was something which I experienced, and that thing was a pink elephant’.

Having argued that there can be non-existent pleonastic intentional objects, Crane then argues that there cannot be content-constituting relations borne towards such intentional objects. This is because relations presuppose the existence of their relata: if there is no relatum, there is no relation. Since pleonastic intentional objects do not exist in the robust ontological sense, there

⁷³ Among existing entities, Crane includes concrete (i.e., spatio-temporal) mind-independent objects and certain abstract objects such as numbers.

can be no relation borne towards them. Consequently, the con-relational account of intentionality fails for thoughts about non-existent objects and for hallucinations.

Crane also argues for the non-relationality of representation on the basis of the essence or nature of a representation. He holds that the nature of a mental state is exhausted by its mode and content (more on these below), and does not include the intentional object (Crane 2009, 477). Thus, he writes:

“The psychological reality of the act is its reality as a representation; a representation is an intrinsic state (of the person, or the brain, or the soul), and this intrinsic state is something which has a specific, concrete nature of its own. This concrete nature does not consist in some unanalysed relation to a proposition [or external object]. That is a theoretical, external description of the state, not what it is in itself” (Crane 2014, 251).

Therefore, according to psychologism, such intentional facts must then be accounted for in terms of *non-relational* psychological entities and facts. In other words, intentionality must be accounted for in terms of *intrinsic* psychological facts, given the common assumption that “if a condition or property is not relational, then it is intrinsic” (Crane 2006, 135). Although Crane comes to this conclusion in the specific case of intentional states that are about non-existing intentional objects, he intends the conclusion to generalise to all intentional states – those that have existing intentional objects as well⁷⁴. Moreover, Crane intends this conclusion to apply equally to non-perceptual and perceptual representations. Thus, he claims that “...perceptual representation is intrinsic to the perceptual state itself, that it represents the world, even if the state is one which can be analysed into its various components [such as vehicles, contents and targets]” (Ibid., 137). Therefore, representation, in general, must be accounted for in terms of non-relational psychological facts.

So, what are the non-relational psychological facts that account for representation? They are content and intentional mode. Crane defines content as “the way the object of a representation is represented” (Crane 2013, 99). Thus, content captures the idea that a represented object is perceived under a certain aspect and from a point of view (Ibid., 97). In other words, content is presentational content, as described in chapter 2 (§2.3.1). Crane also calls content as the *information* about the object:

“When we form a representation of some object, we ‘open a file’ on that object. We then come to store certain information in the file. But we should not think of the information in the file as

⁷⁴ However, in places Crane seems open to pluralism (Crane 2013, 115).

the [semantic] meaning of the name or other expression which we use to express the thought in question...[T]he information a thinker has may be far richer than the meaning” (Crane 2013, 158-159)

In the case of perceptual *experiences*, Crane also calls this conception of content as the “phenomenological content” (Crane 2014, 247). The phenomenological content is often associated with, but distinct from, semantic content (i.e., evaluation conditions). “In the case of a non-propositional [eg., perceptual] representation, the way the object is represented is one thing, and what would make the representation accurate, or correct, or satisfied is something else, something that can be stated as a proposition” (Crane 2013, 132).

The second representation-grounding fact is intentional mode. Intentional mode corresponds to the type of intentional mental state – beliefs, desires, hopes, perceptions, etc. In the case of perceptual representation, the modes could be defined in a coarse-grained way (eg., perceiving target T as content C) or fine-grained (visually perceiving T as C₁, tactually perceiving T as C₂, etc.).

In sum, Crane argues that mode and content together account for representation. And mode and content are non-relational concrete properties of the subject since we can have hallucinations and think of non-existent intentional objects to which no relations can be borne. As regards the phenomenal character of a perceptual experience, Crane argues that it supervenes on both representational mode and content. Thus, he endorses weak impure intentionalism, which holds that “...the phenomenal character of an experience is determined by its...mode and its content. The phenomenal character of an experience therefore supervenes on its intentional nature” (Crane 2009, 481). It is ‘impure’ because phenomenal character supervenes on content and mode, whereas according to pure intentionalism, phenomenal character supervenes only on content.

6.3.1.1. Phenomenological particularity of content

Content is not merely a generic way of representing, but a particular way of representing a target. Thus, Crane writes about perceptual phenomenological content that “content in the phenomenological sense is something spatiotemporal, concrete, particular, and specific to its subject” (Crane 2014, 254). But, given that content is non-relational, the particularity at issue cannot be relational-constitutional particularity. Therefore, it is phenomenological particularity.

Crane discusses the issue of singularity/particularity extensively in the context of singular *thought*. He states that “...a singular thought is one that *purports* to refer to just one object” (Crane 2013, 140; emphasis mine). From the fact that there can be singular thoughts about non-existing entities such as Vulcan and Pegasus, he concludes that singular thoughts are not dependent on the extra-mental existence of individuals and particulars that are thought about (Ibid., 142-143). Rather, “[s]ingularity is a matter of the psychological or phenomenological role of the thought” (Ibid., 147). More precisely, “...a thought can be singular even if it fails to refer to just one object, so long as it has the cognitive role associated with thoughts that succeed in so referring” (Ibid., 140). Extrapolating this to the case of perceptual states, particularity is a matter of phenomenology, and not relations to perceived particulars.

To account for the singularity of thought, Crane posits two explanantia – the aims of the thinker and internal causal connections of content-bearing representations to other representations. These together also account for the cognitive/phenomenological role of content. First, the aims of the thinker. This comprises “...the thinker is aiming in thought – that is, intending to refer to – just one thing” (Crane 2013, 147). In the case of perception, by analogy, the corresponding explanans would be the function of the perceptual system. This is because the perceptual system “aims” at a target owing to its function of representing it.

The second non-relational explanans is the “dispositional or causal connections [of a representation] to other files or representations” (Crane 2013, 159). Regarding this, Crane hypothesises that:

“...one crucial aspect of these connections is how someone is disposed to treat new information associated with the object of my thought. My file for the moon contains the information natural satellite of the earth. If I came to believe that the earth has more than one satellite, I would not simply add this information to the moon file, but I would ‘open’ another file. My file for the moon (the real actual satellite) would now contain the information that it is not the only satellite of the earth, but I now have another file which also contains this information” (Crane 2013, 159).

I believe the following would be reasonable extrapolation to the case of perceptual representation: Any new information/content at t_1 is a new way of representing the same target as was represented at t_0 , or a way of representing a new target. Now, since a perceptual representation inferentially and causally interacts with other preceding and concurrent representations (perceptual or doxastic), this interaction requires that the content of the perceptual representation be sufficiently particular to enable such interaction. Moreover, the

internal interaction would also determine whether an occurrent content at t_1 is about the same target as represented at t_0 , or is about a different target. Thus, it is internal causal connectivity that determines particularity of content.

6.3.1.2. Objection to Crane's aspectual representationalism

The account of phenomenological particularity given by aspectual representationalism has a couple of weaknesses, especially in the case of perceptual content. These weaknesses pertain to each of the particularizing elements – the aims of the thinker and the causal and functional connections of representations.

The analogue of 'aims of the thinker' in the perceptual case is plausibly the biological functions of the perceptual system. Now, these functions are not under the voluntary control of the perceiver as aiming to think about a particular object is. Moreover, the functions of the perceptual system are general-purpose functions. If the functions were selected for representing particular objects that were encountered during selection, representation-producers would not have been able to produce representations of novel or non-natural (i.e., human-made) targets. But our perceptual system does token representations that can represent novel targets. This indicates that representation-producers and their outputs were selected for their general robustness – performance in various circumstances, even those that have not yet been encountered. This is the reason Shea characterizes functions as "robust outcome functions" (Shea 2018, 55) – outcomes that can be produced in a variety of external circumstances and in a variety of ways. Even under the fitness-contribution notion, functions are general purpose since they serve fitness in variable circumstances.

The second proposal of causal connections seems to presuppose particularity rather than explain it. In Crane's example of the moon, the new file (i.e., representation) is opened because the thought is already singular – about the moon. The opening of the new file does not explain singularity but presupposes it. The same holds for causal connections of perceptual files/representations – they presuppose particularity. For instance, Shea writes that representations "...interact causally in virtue of non-semantic properties (e.g. their physical form) in ways that are faithful to their semantic properties" (Shea 2018, 31). The 'faithfulness' or sensitivity of the internal processing to content implies that the content must already be particular to ground the internal processing.

6.4. Common Problems with Con-relational and Non-relational Representationalism

In chapters 4 and 5, I argued that reductive-naturalistic representationalism ultimately fails in its project of reducing content and representation to naturalistic elements. Thus, non-reductive representationalism is the right way forward to explain representation. However, the extant non-reductive theories have a few weaknesses.

6.4.1. Essential External Directedness

First, nonreductive representationalism – whether con-relational or non-relational – cannot adequately account for the essential external directedness of perceptual states. More precisely, the external directedness of perceptual states is characterized by spatial features, and nonreductive representationalism struggles to account for the spatiality of external directedness.

Con-relational representationalism posits the acquaintance relation – for example, Pautz’s experiential representation (ER) relation – to abstract universals. In the case of veridical perception, these universals are, of course, instantiated in concrete particulars. But in illusions and hallucinations, they are uninstantiated. Now, as noted in §6.2.1, Pautz claims that the ER relation to abstract universals depends only on internal neural states, but accounts for the external directedness of perceptual states (Pautz 2021, 103). However, the problem with this proposal is that the abstract universals that we are related to are outside spacetime, by definition. Thus, a relation to non-spatiotemporal entities does not seem to be suitable to explain the spatial nature of external directedness. Moreover, the role of the ER relation is primarily to account for phenomenology and cognitive capacities (reasons-grounding and thought). If it also accounts for external directedness, then it seems that, for a primitive relation, the ER relation performs a multitude of roles, which seems suspect.

Even if it is acceded that the ER relation can account for spatial external directedness, how does it do so? The con-relational representationalist could reply that the ER relation is borne simultaneously to abstract universals and concrete environmental targets. In the latter case, the ER relation would be like the tracking relation that naturalistic representationalism posits. However, this reply would not work because Pautz argues that the ER/conscious-of relation is unique and not identical to the tracking relation or any other physical relation. Accordingly, he writes that

“...there is *no* interesting identification of the form “for x to be conscious of property y is for x to . . . y ”. It is not even identical with some massive disjunction of physical relations...[T]he relation is primitive, just as Russell held that acquaintance is primitive...But rather than holding that consciousness relates to us things (“sense data”) that have properties, the reductive internalist model I favor holds that, in illusion and hallucination at least, consciousness relates us to properties without relating us to things having the properties” (Pautz 2017, 387-388).

Non-relational representationalism faces the problem of accounting for essential external directedness more acutely. This is because content is an intrinsic non-relational property or aspect of either the representation or the subject. So, although it is a concrete spatiotemporal entity, it cannot explain the *external* directedness of perception, since it is an intrinsic property.

In reply, both con- and non-relational representationalism could turn towards the two-dimensional content to account for external directedness. This is Chalmers’ (2006, 2010) view as discussed above. The idea could be that content – whether constituted by acquaintance with abstract universals or non-relationally – is narrow content. It could be a Fregean abstract intension (in the case of con-relational representationalism) or a concrete mode of presentation (non-relational representationalism). Narrow content then determines wide content in environmental contexts – plausibly a Russellian/singular content with propertied objects as constituents. The argument then could be that wide content is capable of accounting for the external directedness of perceptual states.

The problem with the two-dimensional reply is that it does not account for the *essential* external directedness of perceptual states. On this account, perceptual states would be externally directed only contingently or accidentally in the context of environments. It could not explain how a brain in a void, which is not in an environmental context, have externally directed perceptual states. The primary content is narrow content, which is posited to account for phenomenology and cognitive capacities, not external directedness. Wide content is a derivative and contingent content. However, as I argued in chapter 3 (§3.3.2), perceptual states are essentially action-oriented, which implies that they are essentially directed towards the external environment in which they act.

In sum, I conclude that most non-reductive representationalist theories (whether con-relational or non-relational) cannot account for the *essential* external intentional directedness of perceptual states. My diagnosis for this weakness is that they focus entirely on a theory of content to the neglect of a theory of application. This is the exact converse of the problem with

reductive naturalistic representationalist theories, which focus on a theory of application to the detriment of a theory of content.

6.4.2. The Distality Problem

A related second problem facing con- and non-relational representationalism is that they cannot provide an adequate solution to the distality problem. To recall, the distality problem is the problem of explaining how representations can have suitably distal content – information about distal targets to the exclusion of the proximal intermediaries of the naturalistic relational chain between the environment and the subject. It should be noted that, even though non-reductive representationalist views do not posit content-constituting relations to the environment, they do accept causal relations. So, the distality of content is a live problem for non-reductive representationalism as much as for reductive-naturalistic representationalism. To be fair, as mentioned in §6.1.1, the main concern of these theories is the problem of non-veridical perception rather than distality. Nevertheless, the mark of a good theory is its applicability to wider questions.

Now if non-reductive representationalist theories struggle to account for the essential external directedness of content, then they *a fortiori* struggle to account for distality. This is because distality is a stronger form of external directedness – it is external directedness towards distal targets. The acquaintance relation to abstract universals cannot help since the concepts of ‘distality’ and ‘proximality’, which are essentially spatial concepts, do not apply to them.

The two-dimensional content view will not help either. If narrow content is distal, then the burden is on two-dimensionalists to explain how it could be so, since it is determined entirely by intrinsic factors. It could be claimed that, not narrow content but, wide content can account for distality. However, the challenge for this proposal would be to explain how wide content presents distal targets rather than proximal intermediaries of the naturalistic relational (eg., causal) chain. Thus, the distality problem would only be pushed down to the level of wide content.

6.5. Conclusion

In this chapter, I assessed con-relational and non-relational nonreductive representationalist theories. Both endorse content generalism, with which I concur (based on my argument against content particularism in chapter 5). As regards the con-relational version, I argued that the acquaintance relation to abstract universals cannot account for the phenomenal character of perceptual experiences and cannot solve the problem of non-veridical perceptions. As regards Crane's non-relational version, I argued that its account of the phenomenological particularity of content has weaknesses. Both con- and non-relational representationalism, further, have the common weakness that they cannot account for the essential external directedness of perceptual representations and distality, because they do not accept constitutive spatiotemporal relations (such as causation).

7. Content Universalism: Immanent Content

Beginning this chapter and over the next two, I develop and argue for a novel version of nonreductive representationalism that I label *Pluri-Relational Immanent Emergent Representationalism* (PRIMER). In this chapter, I develop PRIMER's theory of content – *Content Universalism*, which is a version of content generalism. Roughly, it claims that content is partly constituted by the instantiation in the subject of target-universals and is inspired by Aristotle's theory of perception. Before getting to it, in §7.1, argue against standard forms of content generalism. Then, in §7.2, I construct a theory of concrete universals, drawing on the views of Aristotle, Armstrong and the classical Indian philosophy of Nyāya. Finally, in §7.3, I put this theory to use and argue for the constitution of content partly by the co-instantiation of concrete universals in both the target and the subject. The other content-constituting conditions are naturalistic relations and teleological biological functions. The content so constituted is identical to subjectively immanent universals.

7.1. Against Standard Content Generalism

In chapter 5 (§5.2), I argued – mainly on the basis of empirical considerations of object tracking and change blindness – that content particularism is false. Consequently, content generalism is true. To recall, content generalism is the thesis that a perceptual state brought about by being perceptually related to a particular α has the property that its content is constituted only by general entities.

Nonreductive representationalist theories – including the con-relational and non-relational theories sampled in the previous chapter – usually endorse content generalism (§6.1.2). Con-relational representationalism holds that the general entities that constitute content are abstract Platonic universals. Consequently, the content is also general and abstract. Now con-relational representationalism also subscribes to intentionalism because content determines phenomenal character. But as argued in chapter 2 (§2.3.4), intentionalist views that start with abstract content face the content-character mismatch problem – phenomenal character is a concrete phenomenon, and cannot be accounted for by, or supervene on, abstract content.

Non-relational representationalism – such as Crane's (2006, 2009) view – holds that general content is a concrete property and therefore does not face the content-character mismatch

problem. But it faces problems in accounting for phenomenological particularity of content. Besides, both con-relational and non-relational nonreductive theories face the common problems of not being to adequately explain the essential external directedness and, more importantly, the distality of content.

Given this situation, in the present chapter and the next two (chapters 8 and 9), I propose and argue for a novel version of nonreductive representationalism that endorses a kind of content generalism, which I will call *content universalism*. Very roughly, content universalism is the thesis that perceptual content is partly constituted by co-instantiation in the subject of the same concrete universals that are instantiated in the represented targets. Content is then identical to the immanent universals in the subject. Supplementing content universalism, which is a theory of content, is the theory of application which I label *representational magnetism* (chapter 9). Together, they constitute my novel version of representationalism – *Pluri-Relational Immanent Emergent Representationalism* (PRIMER). In the present chapter, I will develop and defend the thesis of content universalism. In chapter 8, I will defend the idea that content is a strongly emergent property of the subject. And in chapter 9, I will argue for representational magnetism, according to which content partly grounds the Representational External Directedness Relation (RED) between the representation and the target. Before getting to developing content universalism, I wish to set out/argue for certain views about universals.

7.2. Concrete Universals

Universals are properties that are shareable by multiple particular objects such that each particular instantiates the very same universal. Armstrong, accordingly, describes a universal as being “strictly identical” (1997, 21) across its particular instances, and as a “one running through many” (Armstrong 2004, 140). Therefore, universal are general properties. The relationship or connection between universals and particular objects is known as instantiation⁷⁵. Thus, particular objects are said to instantiate universals. I will call the particular objects the ‘instantiators’ of universals.

Universals are contrasted with tropes, which are particular properties that cannot be shared between entities. For example, suppose two roses are the same shade of red in colour. A

⁷⁵ Instantiation is often used synonymously with another term – ‘exemplification’ (Orilia and Paolini Paoletti 2025, §2). However, Lowe (2006) uses instantiation and exemplification differently.

universal theorist would say that the two roses (or their particular red colours) instantiate one and the same universal – redness. The two roses (and their particular red properties) resemble each other because they share the same universal. This is the classic ‘One Over Many’ argument for universals (Ney 2023, 62). Depending on the specifics of the theory of universals, either the roses or their particular red properties instantiate, and are therefore the instances of, the universal redness. A trope theorist, in contrast, would say that the red properties of the two roses are distinct properties that do not instantiate any universal. Any resemblance between the two is a primitive relation.

There are two broad kinds of universals – Platonic and Aristotelian universals. Platonic universals are those that can exist even when not instantiated by their particular instantiators. When uninstantiated, they exist in abstract form outside space and time. Therefore, Platonic universals are also known as *transcendent universals* (Orilia and Paolini Paoletti 2025, §1.3). Many nonreductive representationalists take content to be constituted by Platonic transcendent universals.

In contrast, Aristotelian universals are always instantiated by particular objects, and cannot exist uninstantiated. Tugby (2022), for instance, writes that Aristotelian universals depend on concrete entities for their existence (Tugby 2022, 29). They are therefore known as *immanent universals* because they are ““immanent” or “*in rebus*” (“in things”)” (Orilia and Paolini Paoletti 2025, §1.3). Since immanent universals are always instantiated in their particulars, which are concrete (spatiotemporal), the universals are themselves concrete. Moreover, since they are universals and multiply shareable, they are also general.

Since I will later argue that concrete universals have a content-constituting role (in §7.3.2), I will discuss the metaphysics of concrete universals in the rest of this section. In particular, I will look at the views of Aristotle, Armstrong (a contemporary Aristotelian) and the classical Indian philosophical system of Nyāya-Vaiśeṣika.

7.2.1. Aristotle on Universals

In his early work *Categories*, Aristotle listed ten categories of being, understood in the broadest sense (i.e, entities and actions): “(1) substance; (2) quantity; (3) quality; (4) relatives; (5) somewhere; (6) sometime; (7) being in a position; (8) having; (9) acting; and (10) being acted upon (*Categories* 1b25–2a4)” (Studtmann 2024, §2.1). These 10 categories can be classified

into four groups, based on two axes –beings that are *said of* other beings, and beings that are *present in* other beings:

	Not Said-Of	Said-Of
Not Present-In	<u>Not Said-Of and Not Present-In:</u> Substantial particulars or particular objects. Aristotle calls them “primary substances”. For example, Socrates (the man), Garfield (the cat).	<u>Said-Of and Not Present-In:</u> Substantial universals such as kinds, genera and species. For example, human-ness or cat-ness. They are always said of particular substances and other substantial universals. For example, cat-ness (genus) is said of persian cat-ness (species); cat-ness and persian cat-ness (genus and species) is said of Garfield.
Present-In	<u>Not Said-Of and Present-In:</u> Non-substantial particulars. For example, the particular colour instance red or a particular shape instance. These non-substantial particulars are also known as <i>modes</i>⁷⁶ (Lowe 2006; Heil 2012). They are always present in particular substances.	<u>Said-Of and Present-In:</u> Non-substantial universals. For example, colouredness or redness. They are always present in particular substances. For example, both colouredness and orange-ness are present in Garfield. Some non-substantial universals are said of other non-substantial universals and of non-substantial particulars. For example, coloured-ness (genus) is said of orange-ness (species); coloured-ness and orange-ness (genus and species) are said of the particular orange-instance (individual).

⁷⁶ Both Lowe (2006) and Heil (2012) distinguish modes from tropes. Tropes, like modes, are property-instances; however, they are mostly posited in the context of a one-category ontology – eg. Williams (1953) – which does not accept substances or universals (Lowe 2006, 96-97).

7.2.1.1. Universalist Dualism:

As the above classification indicates, there are two kinds of universals – substantial universals and non-substantial universals. I will call this position universalist dualism⁷⁷. Corresponding to the two kinds of universals are the two kinds of particulars – objects (substantial particulars) and modes (non-substantial particulars).

Substantial universals are those that are ‘said of’ particular objects (‘primary substances’). They are also known as essential universals since they constitute the essence (or real definition) of a particular. Aristotle calls the substantial universals ‘secondary substances’ due to their role in defining primary substances. Examples of substantial universals are the kinds, species and genera of particulars. *Non-substantial or accidental universals*, on the other hand, are those universals that are ‘present in’ substances. For example, universals of qualities such as colour, smell, shape, etc. So, if the substantial universal of Garfield is ‘cat-ness’, his accidental universal is ‘orangeness’. Thus, there are two ways for a particular substance to instantiate universals – by having them *said of* it, and by having them *in* it.

Aristotle’s notion of universals dovetails into his notion of form in his mature theory of hylomorphism developed in his later works *Physics* and *Metaphysics*. Hylomorphism is Aristotle’s theory of substance, change and causation. According to hylomorphism, every substance or particular object, especially living organisms, is a compound of two metaphysical entities – matter (*hyle*) and form (*morphe*). Aristotle distinguishes between matter and form primarily on the basis of potentiality and actuality. As Shields quotes Aristotle: “‘Matter exists in potentiality, because it may move into a form; and to be sure, when it exists actually, it is in its form’ (*Metaphysics* 1050a15–16)” (Shields 2013, 67). Thus, matter and form can be defined as: “x is matter =_{df} x exists in potentiality; [y] is form =_{df} [y] makes what exists in potentiality exist in actuality” (Ibid., 67). Matter is also described as a principle of “individuation” (Ainsworth 2024, §3) or “numerical diversification” (Loux 2007, 155) or “potential being” (Simpson 2023, 5). Form, in contrast, is the principle of “unity” (Shields 2019, 219) or “actual being” (Simpson 2023, 5).

Aristotle’s notion of form has been interpreted in various ways, particularly in contemporary Neo-Aristotelian literature, as Shields (2019) points out. For instance, forms are interpreted as

⁷⁷ Lowe (2006), a Neo-Aristotelian, also follows Aristotle’s classification and posits two categories of universals – kinds (=substantial universals) and attributes (=non-substantial universals), and two categories of particulars – objects and modes – in his four-category ontology.

structures that are component parts of substances (Koslicki 2008), structures that are merely organizing principles (Fine 1994; Johnston 2006), powers (Koons 2014) and functional roles (Shields 2019). However, an interpretation that harmonizes with Aristotle's earlier work is that *forms are universals* that are always immanent in matter. The combination of form/universal and matter gives rise to substances. This interpretation is championed by Loux (2005, 2007).

The universalist interpretation of forms has the advantage that it gives universals, interpreted as forms, a role in causation and change. This is because form, being the principle of actual being, is deemed by Aristotle as one of four causes – the formal cause⁷⁸ – of an entity. Moreover, the universals = forms interpretation is buttressed by the parallels between Aristotle's independent discussions of universals and forms. For instance, corresponding to the distinction of substantial and accidental universals is the distinction between substantial and accidental forms (Shields 2013, 71; Simpson 2023, 5)⁷⁹. Writing of substantial forms, Loux says that "...forms are universals whose predication of the matter delivers individuals falling under substance kinds" (Loux 2007, 155). Thus, substantial forms play the same role as substantial universals. Accidental forms are those universals that characterize a substance in particular ways – as being of a certain colour, shape, etc.

In sum, the two kinds of universals/forms – substantial and non-substantial – play two roles in Aristotle's metaphysics. First, they account for the resemblance and difference among entities – entities resemble each other insofar as they instantiate the same universals; conversely, they differ from each other insofar as they instantiate different universals. Second, they account for causation and change in entities. To be sure, both matter and universals play a role in causation.

7.2.1.2. Aristotle on instantiation:

Universals, for Aristotle, are always immanent – if they exist, then they are always instantiated in particulars. As mentioned above, Aristotle claimed that there are two ways or manners in which universals are instantiated by particular objects – by *being said of* particulars and *being in* particulars. These two manners of instantiation may be called 'predication-instantiation' and 'inherence-instantiation'.

⁷⁸ The four Aristotelian causes are: "(1) the "that out of which"...or the material cause; (2) the "what it is to be"...or the formal cause; (3) the "that from which the origin of motion or rest comes"...or the efficient cause; and (4) the "that for the sake of which"...or the final cause" (Leunissen 2010, 11-12).

⁷⁹ Of course, there is a major difference between Aristotle's substance-universal theory and hylomorphism. In the former, the universal is instantiated by a substance, while in the latter the form is instantiated by prime matter.

(i) Predication-instantiation: The *being said of* mode of instantiation is when the universal is merely predicated of the particular without being in it. On the one hand, substantial universals (eg. catness) are predicated of objects (Garfield) and other substantial universals (persian catness); on the other hand, non-substantial universals (colouredness) are predicated of modes (orange) and other non-substantial universals (orangeness). For example, Garfield instantiates the universal ‘cat-ness’, because ‘catness’ is said of him.

(ii) Inherence-instantiation: The *being present in* mode of instantiation is explained in terms of a “fundamental inherence relation” (Cohen and Reeves 2025, §2). That is, the universal *inheres in* the particular object. Both non-substantial universals (orangeness) and non-substantial particulars (orange) inhere in substantial particulars (Garfield). Importantly, Aristotle clarifies that although universals inhere in particulars, they are not mereological parts of the latter (Studtmann 2024, §1).

Marmodoro (2021) gives a common interpretation of the two manners of instantiation as *qualification* of a substance. That is, whether a universal is said of or present in a substance, it qualifies the substance. Thus, she writes that,

“...the instantiation of a universal is not the presence of the universal in the substance: e.g., the instantiation of ‘wisdom’ in Socrates is not the presence of ‘wisdom’ in Socrates; it is, rather, Socrates being wise. I thus understand Aristotle’s position as follows: a property’s instantiation in a substance is a type of metaphysical ‘transformation’... It is a transformation of the property, from being an abstract entity,...e.g. as ‘wisdom’, to being the qualification of a particular concrete subject, e.g. a wise Socrates” (Marmodoro 2021, 38).

7.2.1.3. Aristotle’s theory of perception and universals:

Aristotle gave a central role to universals/forms in his theory of perception. According to him, in the process of sensory perception, the sensory faculties receive the form of the sensible or perceived object without its matter (Caston 2012, 328). The change that occurs in the sensory faculty as a result of receiving the form is termed as the ‘*enforming*’ of the sensory faculty with the form of the sensible object. Since forms are best interpreted as universals, this implies that perception involves the reception by the sensory faculty of the universals instantiated in perceived object. To be clear, Aristotle is not saying that we perceive universals. Rather, the idea is that we perceive concrete objects and their modes by means of receiving their respective universals.

Shields encapsulates Aristotle's other observations on perception in the following definition of perception (for some subject *S* and some sensible object *O*):

“*S* perceives *O* if and only if: (i) *S* has the capacity requisite for receiving *O*'s sensible form; (ii) *O* acts upon that capacity by enforming it; and, as a result, (iii) *S*'s relevant capacity becomes isomorphic with that form” (Shields 2020, §6).

The “sensible” form in clause (i) is just the universals of a sensible, i.e., perceived, object. This includes both substantial and non-substantial universals. The isomorphism requirement in clause (iii) holds that the sensory faculty becomes *like* – i.e., similar to – the perceived object. This requirement has been interpreted in two ways – literally and non-literally. The literalist interpretation has it that the sensory faculty or organ becomes exactly the same as the sensible object by instantiating the universals originally instantiated in the object. Thus, “*likeness* amounts to shared-property exemplification” (Shields 2020, Supplement 3). For example, seeing a red object results in the eye literally instantiating redness.

Shields notes that non-literalist interpretations are more popular among Aristotle's commentators. These argue that the likeness between the sensory faculty and perceived object is merely metaphorical and calls for further elucidation. One non-literalist interpretation is the “intentionalist” interpretation (Shields 2020, Supplement 3). According to this, the sensory faculty becomes like its sensible objects “...by coming to *symbolize* them in one way or another. In its simplest form, this approach regards the likeness involved in perception as akin to the likeness obtaining between a house and its blueprint” (Ibid., Supplement 3). The intentionalist approach is similar to representationalism since the enforming of the sensory faculty is similar in role to presentational content – it is in virtue of the enforming that the sensory faculty symbolizes or represents the object. Thus, the sensory faculty is the representational vehicle, the object is the target, and the received form/universal is the content.

The legacy of Aristotle is thriving in contemporary metaphysics, as evidenced in numerous papers, monographs and essay collections⁸⁰. Prominent Neo-Aristotelians are Armstrong, Lowe, Marmodoro, among others. In what follows, I will discuss Armstrong's views on universals and their instantiation.

⁸⁰ For instance, Tahko (ed.) (2011), *Contemporary Aristotelian Metaphysics*. Novotný and Novák (eds.) (2014), *Neo-Aristotelian Perspectives in Metaphysics*; Corkum (2025), *Neo-Aristotelian Metaphysics*.

7.2.2. Armstrong

For Armstrong, the most fundamental ontological category is the state of affairs (Armstrong 1997, 1). “An atomic [i.e., basic] state of affairs exists if and only if a particular has a property, or a relation holds between two or more particulars. These properties and relations are, of course, universals. We will call the particulars, properties and relations the *constituents* of the states of affairs” (Ibid., 20). It should be noted that universals are *non-mereological* constituents of a state of affairs. Also, Armstrong admits monadic and polyadic universals; the latter correspond to relations. Therefore, the structure of atomic states of affairs is aF or Rab .

Armstrong explains universals as unsaturated properties, or properties with a gap. If there is the state of affairs aF (a is F), then the universal is ‘__F’ or “__’s being F ” (Armstrong 1997, 28), rather than merely ‘ F ’. Thus, universals can be interpreted as state of affairs schemas. Armstrong calls them “state-of-affairs *types*”:

“The universal is a gutted state of affairs; it is everything that is left in the state of affairs after the particular particulars involved in the state of affairs have been abstracted away in thought. So it is a state-of-affairs type, the constituent that is common to all states of affairs which contain that universal” (Armstrong 1997, 28-29).

The particular abstracted from its universals (i.e., its properties) is termed “thin particular” (Armstrong 1997, 109). A thick particular in contrast is the propertied particular which is just the state of affairs aF . Thus, the state of affairs aF involves a thin particular a instantiating the universal ‘__F’.

7.2.2.1. Universalist Monism:

In contrast to Aristotle, Armstrong is a *universalist monist* in the sense that he admits only one category of universals – the non-substantial universals. Moreover, Armstrong endorses a sparse theory of universals, according to which the universals are whatever properties natural science posits (Armstrong 1997, 25). Therefore, paradigm examples are the universals corresponding to fundamental physical properties such as mass-hood, electric charge-hood, etc.

However, Armstrong makes space for kinds – such as catness – in his ontology by positing complex universals. Complex universals are of two types – conjunctive universals and structural universals. A “conjunctive universal” (Armstrong 1997, 31), as the name suggests, is a universal formed from the conjunction of two or more simpler universals that are always

co-instantiated in the same particular. For example, the kind (substantial) universal ‘electronhood’ is a conjunction of the universals of mass, spin and charge (Ibid., 67).

Writing about the formation of conjunctive universal F&G, Armstrong writes that “[g]iven that F and G are distinct universals, then F&G can be a universal, provided always that a particular exists at some time which is both F and G” (Armstrong 1997, 31). The criterion by which F&G can be counted as a novel universal is its distinct causal powers, which cannot be reduced to the summation of individual causal powers or F and G taken separately. The universals “...together are ‘synergistic’, bestowing a power that is less or more than the sum of the powers bestowed by one conjunct in abstraction from the other” (Ibid., 32).

The second type of complex universal is the “structural universal” (Armstrong 1997, 32). Structural universals are applicable to mereologically structured objects – that is, to wholes composed of parts. Where F is a structural universal, “[t]o be an F a particular must be made up of just two nonoverlapping parts, one of which instantiates universal G while the other instantiates universal H, with the G part and the H part linked by the (external) relation R” (Ibid., 32). Examples of structural universals include kinds of mereologically structured objects. For example, a methane molecule (CH₄) is composed of four atoms of hydrogen and one atom of carbon bonded together. So, ‘Methanehood’ is a structural universal constituted by the universals of ‘Hydrogenhood’ and ‘Carbonhood’ being instantiated in the respective atoms, and the atoms bonded together by the relational universal of bonding B (Ibid., 34).

7.2.2.2. Role of universals – resemblance and causation:

The primary role of universals is to account for “...the resemblances and differences that we find among particulars” (Armstrong 1997, 25). This contributes to the ‘One Over Many’ argument for universals. A second crucial role for universals is that they confer causal powers on the particulars instantiating them. More specifically, it is a law of nature that confers causal powers on particulars, but a law of nature is itself a higher-order relational universal of necessitation between two lower-order monadic universals.

Suppose there are two universals F and G, and there is a law that states that every instance of F nomically or contingently necessitates (Armstrong 1983/2016, 79) an instance of G. This law is expressed as N(F,G) (where N stands for nomic necessitation). “Although N(F,G) does not obtain of logical [or metaphysical] necessity, if it does obtain then it entails the corresponding Humean or cosmic uniformity: (x) (Fx $\supset$ Gx)” (Ibid., 79). The law N(F,G) is itself a state of

affairs since it involves the N relation between two universals F and G. This view of laws is known as the Dretske-Tooley-Armstrong (DTA) view⁸¹.

Armstrong further argues that the law $N(F,G)$ is also a higher-order universal. Thus, Armstrong writes:

“...if we can also accept that $N(F,G)$ is a universal, instantiated in the positive instances of the law, then, I think, it will be much easier to accept the primitive nature of N. It will be possible to see clearly that if N holds between F and G, then this involves a uniformity at the level of first-order particulars)... $N(F,G)$ is [also] a state of affairs. Rab is a state of affairs involving first-order particulars falling under a first-order universal. $N(F,G)$, by contrast, involves second-order particulars (first-order universals) falling under a second-order universal” (Armstrong 1983/2016, 82).

Causation between aF and bF, for Armstrong, is just lawful necessitation from aF to bF. And this necessitation is accounted for by the law (=higher-order universal) $N(F,G)$ – that is, by the relation N between two universals F and G. Thus, it is the higher-order universals which confer causal powers on particulars. First-order universals are purely categorical or qualitative properties (Armstrong 1996, 16).

7.2.2.3. Armstrong on instantiation:

Following Aristotle, Armstrong subscribes to the *principle of instantiation*. Armstrong puts it as follows: “The Principle of Instantiation for properties asserts that, for each [universal] property, P, there exists (not necessarily now) a particular, x, such that Px” (Armstrong 1978, 76). This means that all universals are immanent in their particulars.

Armstrong does not consider instantiation as a relation. He calls instantiation the “fundamental tie or nexus” between the universal and particular (Armstrong 1989, 109). Lowe, another prominent Neo-Aristotelian, also does not consider instantiation a full-fledged relation, and instead calls it a “*relationship*” (Lowe 2006, 91). In his earlier works, Armstrong took instantiation to be a primitive and unanalysable tie. However, in his later years, he gave an analysis of instantiation in terms of *non-mereological intersection* of universals and particulars.

Developing a proposal first made by Baxter (2001), Armstrong (2004) analyses instantiation as an intersection between a thin particular (a particular abstracted from its properties) and

⁸¹ Because Dretske, Tooley and Armstrong independently proposed the idea in the late 1970s that laws are relations between universals.

universals. He gives the analogy of a cross, which is made up of a vertical and a horizontal piece of wood intersecting each other. Armstrong calls it only a rough analogy because in the case of the cross, the intersection is mereological overlapping. However, the intersection of a thin particular and a universal is non-mereological.

Armstrong claims that the intersection account of instantiation is one way of fleshing out Plato's assertion in *Paramenides* that "particulars...participate in their universals" (Armstrong 2004, 140). He further claims that since the intersection is a sort of non-mereological *unity* between the particular and universal, it is a case of "partial identity" (Ibid., 141). Partial identity is contrasted with the strict identity of universals across its instances⁸². To sum up, instantiation is non-mereological intersection of a universal and a particular that amounts to a partial identity between them.

7.2.2.4. Taking stock:

So far, we have encountered two opposing views about universals – the Platonic view that universals are transcendent and abstract, and sometimes are instantiated in particulars; and the Aristotelian view that universals are always immanent in particulars, and in virtue of their immanence are concrete. We went into detail into two developments of the Aristotelian view – by Aristotle himself and by Armstrong. However, there is space for an intermediate view, one on which universals can exist independently of particulars, yet are concrete even when existing independently. This is a view found in the classical Indian philosophical system of Nyāya-Vaiśeṣika, to which I turn next.

7.2.3. The Nyāya-Vaiśeṣika System of Philosophy

The Nyāya-Vaiśeṣika system of philosophy comprise two closely allied philosophical schools – the Nyāya and the Vaiśeṣika. Vaiśeṣika is the older of the two and provides the metaphysical system common to both. The Nyāya school focuses more on epistemology (including perception) and logic. The Vaiśeṣika is traditionally said to be founded by Kaṇāda (2nd century BCE) who wrote its first treatise the *Vaiśeṣikadarśana*. The origin of the Nyāya school of

⁸² Other cases of partial identity include resemblance of universals (Armstrong 1997, 47), partial identity between a conjunctive universal and its conjuncts (Ibid., 31) and mereological partial identity between wholes and their parts, and between overlapping parts (Ibid., 18).

philosophy is traditionally attributed to Gautama (c. 100 CE), who composed the first Nyāya treatise the *Nyāyasūtra*⁸³.

The Nyāya-Vaiśeṣika system posits 7 fundamental ontological categories (Potter 1977, 49): (1) substance (*dravya*), (2) quality or mode (*guṇa*), (3) action (*karma*), (4) universal (*sāmānya*), (5) individuator or particularity (*viśeṣa*), (6) inherence (*samavāya*), and (7) non-being (*abhāva*). The first five categories are related to each other by the sixth category – fundamental relation of inherence. Inherence plays the same role in Nyāya-Vaiśeṣika metaphysics that instantiation plays in Aristotelian metaphysics, and more (to be discussed below). I will focus on universals and inherence in what follows, since the other categories are not essential to my argument for content universalism.

The canonical Nyāya-Vaiśeṣika definition of universals (*sāmānya*) is: “a universal is something that is (a) eternal, (b) unitary, and (c) located in a plurality of things – substances, qualities, or motions” (Chadha 2014, 288). Potter (1977) encapsulates the other features of universals in the following summary:

“The fully developed Nyāya-Vaiśeṣika view of universals is that they are real, independent, timeless, ubiquitous entities which inhere in individual substances, qualities, and motions and are repeatable, i.e., may inhere in several distinct individuals at once and/or at different times and places” (Potter 1977, 133).

The above definition and summary capture the following features of Nyāya-Vaiśeṣika universals (I’ll henceforth call them ‘Nyāya universals’): They are ‘real’, which means that universals are not conceptual creations, but are real entities. Secondly, Nyāya universals are ‘timeless’ in the sense that they are eternal and indestructible (Sharma 1962, 167). Thirdly, Nyāya universals are ‘unitary’, which means that there is only one universal corresponding to a property or kind. Further, the universal is mereologically simple – it has no parts or constituents.

Fourthly, Nyāya universals are ‘independent’ in the sense that, although often inhering in particulars, they can exist uninstantiated (Potter 1977, 139). Thus, Potter writes that the Nyāya philosopher “should be conceived to be a realist in the scholastic sense; he believes that universals exist *ante rem*, independently of both thought and instantiation” (Ibid., 135). Fifthly,

⁸³ The classical Nyāya school developed for over a millennium benefitting from its intense debates with Buddhist epistemology. It then evolved into the Navya-Nyāya or the Neo- Nyāya school starting with Gaṅgeśa Upādhyāya (c. 1350).

universals are ‘ubiquitous’ in the sense that they are “all-pervasive entities” (Ibid., 139). A universal, being ‘unitary’, pervades space and time, and is everywhere all at once.

The conjunction of the fourth and fifth features has a remarkable implication – Nyāya universals *always* exist in space and time, even when uninstantiated. That is, universals are always concrete, whether instantiated or not. Thus, Vaiśeṣika universals occupy a logical space between Platonic and Aristotelian universals. Like Platonic universals, they can exist even when not instantiated. Unlike Platonic universals and like Aristotelian universals, they are always concrete (spatiotemporally located). However, their concreteness is not dependent on their immanence.

7.2.3.1. Universalist Dualism:

Like the Aristotelian school, the Nyāya school accepts two broad kinds of universals – substantial universals and non-substantial universals. However, even within non-substantial universals, there is a bifurcation into quality universals (eg., redness), and motion or action universals (eg., upwardness) (Sharma 1962, 168). Universals are instantiated in their particulars by the primitive relation of inherence (*samavāya*), which is itself one of the seven categories.

7.2.3.2. Inherence:

The inherence (*samavāya*) relation is a unique feature of the Nyāya-Vaiśeṣika system in the context of Indian philosophy. The Nyāya school employs it for various purposes. Relating universals to particulars is just one of them. The full curriculum vitae of inherence is as follows: (i) every kind of universal inheres in its dedicated instances/particulars – substantial universals (eg., catness) inhere in particular substances (eg. Garfield); quality universals (orangeness) inhere in qualities/modes (orange), and action universals (runningness) inhere in particular actions (running); (ii) modes inhere in substances; (iii) particular actions inhere in substances; (iv) particularity inheres in substances; (v) the whole inheres in its parts; and (vi) the effect (eg. cloth, pot) inheres in its cause (threads, pot halves)⁸⁴.

As the above list testifies, inherence has a role in explaining properties, particularity, composition and material constitution. It is a relation that holds between two substances (effect inheres in cause; whole inheres in parts), between a substance and a quality, and between a

⁸⁴ The cause in which the effect inheres is known as the ‘inherence cause’, and is one of three kinds of causes (more on this later).

substance and an action. In the case of quality and action universals, there is a double inherence. They first inhere in their modes and particular actions, respectively, which in turn inhere in substances.

The canonical definition of inherence is the one given by Praśastapāda (c. 5th century CE) – “Inherence connects things that are inseparably connected (*ayutasiddha*) and stand to each other in the relation of located and locus” (Potter 1977, 302). The inseparability of the related entities is interpreted in terms of spatiotemporal colocation and coincidence (Shastri 1976, 394). The related entities are different, although they appear identical in perception. Thus, Praśastapāda also describes inherence as the relation of identity or unity (*tādātmya*). Moreover, one of the relata – the instantiator entity – is described as the locus or container, and the other – the instantiated – is the located or contained.

The characteristics of inherence are as follows. First, the inherence relation is eternal in the sense that its existence is not dependent on any contingent entities or even universals that it relates (Shastri 1976, 388). In contemporary terms, inherence would be an “external relation” (MacBride 2020, §3) – irreducible to the intrinsic natures of its relata. Whereas one or both of its relata may be destroyed, the relation of inherence itself is not destroyed. This implies that inherence is a fundamental and primitive relation.

Second, inherence is one (Shastri 1976, 389). The feature of oneness has two aspects: (i) It is ‘one’ in the sense that inherence is a simple relation without constituent relations. (ii) It is ‘one’ in the sense that there is just one kind of inherence relation, which relates relata of different categories. Thus, the inherence of universals in substance, or the inherence of modes in substances, or the inherence of whole in parts (both of which are substances) is owing to same inherence relation. However, as listed above, the pairs of entities related by inherence are very different in nature; therefore, it follows that the *manner* of inherence differs in each case. Nyāya philosophers argue that the difference in the manner of inherence is due to the potentiality of the locus relatum, i.e., the instantiator (Shastri 1976, 389). Thus, catness inheres in Garfield and not in any of his particular qualities or in any other object due to the potentiality of Garfield to be a cat. Similarly for the universal ‘orangeness’ – it inheres in orange and not (directly) in Garfield. The invocation of potentiality is strikingly similar to Aristotle’s idea of matter being the principle of potentiality.

Third, inherence does not require another inherence relation for relating its relata. This avoids an infinite regress. Thus, Sridhara (c. 10th century CE) writes “[b]eing itself in the nature of a

relation, [inherence] subsists in an object by itself (without the need of any other relation)” (Shastri 1976, 386). Fourthly, although eternal and all-pervasive inherence is neither a universal which needs instantiation nor is an instantiation of any other universal.

7.2.3.3. Role of universals – resemblance and causation:

The role of universals in Nyāya metaphysics is the same as in Aristotelian metaphysics – to account for (a) resemblance and difference among entities, and (b) causation (Chadha 2014, 289). I will focus on the causation role, since it is the more interesting one.

Nyāya theory of causation posits a causal constellation – a collection of causal conditions, known as *sāmagri* (Potter 1977, 55) – producing an effect. Taking the example of a bronze statue (the effect), the causal constellation involves three causes (Ibid., 56-57): (i) the *inherence cause*, which are the parts of the statue. It is called so, because the effect inheres in the cause – the parts. (ii) the *non-inherence causes*, which are “closely related” to the inherence cause, without themselves inhering in the effect. The non-inherence causes includes properties (i.e., modes) of the lump such as its colour, its mass, etc. These are, of course, instantiations of universals in the inherence cause, but not the effect. (iii) the *instrumental cause*, which is a miscellaneous bunch of all other causes, such as the efficient cause (the sculptor), the tools, the actions required to produce the statue, etc.

The above three causes are together sufficient for the effect. The Nyāya view is similar to the Humean view of constant conjunction, because once the causal constellation obtains, the effect necessarily follows. Thus, Udayana (c. 10th century CE) writes that “[c]ausation just means regular connection between something prior to the effect and the appearance of that effect” (Potter 1977, 560). Significantly, Udayana clarifies that the necessary connection is between the universals involved in the causes: “...the regularity must hold between universals and not merely between individuals. For if it is not so, the very identification of an individual will be rendered impossible” (Ibid., 561). Shastri (1976) elucidates the role of universals in the case of the inherence cause (eg. the bronze lump). “If the [universal]...is not accepted, there will be no single determinant of the [inherence cause] of effects, or that of connection or disconnection [between cause and effect]; and this lack of a single determinant will necessitate the enumeration of many causes” (Shastri 1976, 322). In other words, the universal (bronzeness) of the inherence cause (lump) makes it determinate such that it can cause the effect (statue). If not for the universals, then there would be nothing about the particulars to ground the regular conjunction between the cause and the effect.

7.2.3.4. The Nyāya theory of perception and universals:

Universals play a major role in the Nyāya theory of perception as one of several possible objects of perceptions. Gautama, in the *Nyāyasūtra*, defines perception thus: “Perception is a cognition which arises from the contact of the sense organ and object and is not impregnated by words, is unerring, and well-ascertained” (Chadha 2024, §2.2).

Thus, perception has four individually necessary and jointly sufficient conditions: (1) it is produced from contact between sense organ and an object; (2) it is “non-verbal” (Potter 1977, 162; Matilal 1986, 228); in other words, it is non-conceptual; (3) it is non-erroneous or “non-deviating” (Matilal 1986, 228); and (4) it is well-ascertained or “having the nature of certainty” (Ibid., 228).

The first necessary condition of perception is the sensory organ-object contact. This is the perceptual relation between the sense organ and the perceptible object. As Phillips explains, Nyāya philosophers argued that the perceptual contact relation involves *physical* contact between the sense organs and objects. Further each of the sense organs has a different manner of being in contact with its perceptible objects (Phillips 2019, 55)⁸⁵. For example, in the case of vision, the contact happens at the location of the distal object when “...the visual organ itself expands as rays stretching out of the body as swiftly as light, which carries it, coming into contact with objects at a distance” (Ibid. 55). In the case of other sensory modalities, the contact happens at the site of the organ.

Condition (3) – non-erroneousness – indicates that ‘perception’ is a success term. Illusions and hallucinations are strictly speaking not perceptions. Thus, Nyāya subscribes to disjunctivism (Phillips 2019; Vaidya 2020), which is the theory that veridical perceptions, illusions and hallucinations have nothing in common except the possibility that they could sometimes be phenomenologically indistinguishable (Soteriou 2020, §1). Vācaspati Mīśra (c. 950 CE) interprets conditions (2) and (4) as delineating two kinds of perception – indeterminate or non-conceptual (*nirvikalpaka*) perception, and determinate or conceptual (*savikalpaka*) perception (Chadha 2024, 7). Concepts are interpreted as the products of mental activity; therefore, non-conceptual perception is one that is free from “...the mental activity of synthesizing or judging” (Ibid., 7). It should be noted that these are sequential stages of perception rather than types.

⁸⁵ The sense→object contact is then followed by the *mind*→sense contact and still later by *self*→mind contact.

Thus, “[a]ll perception is determinate, but it is necessarily preceded by an earlier stage when it is indeterminate” (Sharma 1962, 182).

According to Nyāya philosophers, we can perceive entities belonging to all seven categories – not merely particular objects, but also universals and even absences^{86,87}. Regarding the perception of universals, Nyāya philosophy argues that only inhered (i.e., immanent) universals can be perceived. Jayanta Bhatta (c. 9th century CE) further argues that universals are first perceived in non-conceptual or indeterminate perception. According to him, “...at the stage of indeterminate perception, universals and particulars are given but not *as* universals or *as* particulars. It is only at the stage of determinate perception, wherein the mind is involved in the perception of identity and difference, that universals and particulars are cognized as such” (Chadha 2014, 293).

7.2.4. Constructing a Theory of Universals

In the previous three subsections, I have discussed diverse views on universals – from ancient Greece through classical India to the contemporary West; from Aristotle through Nyāya to Armstrong. I will now draw some lessons and construct a provisional theory of universals that I will later apply to develop my view of content universalism. In drawing these lessons, I will apply, as far as possible, the theoretical principle of ontological parsimony or Occam’s Razor – “[e]ntities are not to be multiplied beyond necessity” (Baker 2022, 2). In other words, I will adopt those views about universals which involve the least ontological commitments, from among alternatives that are equally explanatory.

7.2.4.1. Concrete universals:

The Aristotelian (Aristotle and Armstrong) view of universals is that universals exist only if they are immanent in their particulars. The Nyāya view is that universals spatiotemporally exist prior to particulars, but are often immanent. Although this is not the place for a detailed comparison of the merits and demerits of the two views, I have some remarks.

⁸⁶ In fact, the Nyāya view of perception is so liberal as to include even extraordinary perception within its remit. The theory discussed in this chapter is “ordinary perception”. “Extraordinary perception” includes three types – (i) perception by association, (ii) perception of all particulars in virtue of perceiving a universal instantiated in one particular, and (iii) Yogic or “intuitive perception of all objects - past, distant and future - due to some supernormal powers generated in the mind by devout meditation” (Chatterjee 1939, 248).

⁸⁷ Due to extreme diversity of perceptible entities, one kind of perceptual contact relation is not enough to enable every kind of perception. Therefore, Uddyotakara (c. 7th century CE) lists six types of sense organ-object contact (Potter 1977, 162), in virtue of which there are six types of perception. Three of the six types involve perception of immanent universals.

As noted above, Aristotle and Aristotelians endorse the principle of instantiation, which has the implication that universals are ontologically dependent for their existence on their particulars. Thus, Tugby writes that Aristotelians “...accept some form of the so-called principle of instantiation, which says that universals are (generically) existentially dependent on their instances” (Tugby 2022, 29). This gives an ontological priority to the particulars. On the other hand, universals are explanatorily prior to their particulars in the sense that universals explain the property-instances of particulars, their resemblance and their causal powers. This brings the ontological and the explanatory projects of Aristotelianism into tension.

The Nyāya view of universals, on the other hand, holds universals both ontologically and explanatorily prior to particulars. To recall, the Nyāya view is that universals pervade all of space-time, which implies that they are intrinsically concrete. Their concreteness does not depend on their immanence in particulars. In other words, they are *intrinsically concrete* universals as opposed to Aristotelian *immanently concrete* universals.

The idea of intrinsically concrete universals can be better grasped by noting their similarities with the quantum fields that Quantum Field Theory (QFT) posits. According to QFT, the most fundamental physical entities in the world are quantum fields that pervade all of space-time. The properties (eg. electric charge, spin, mass, etc.) and particles (photons, electrons, quarks, etc.) that fundamental physics posits are *excitations* or disturbances of the quantum fields. There are two kinds of quantum fields – *force fields* for three of the four fundamental forces/interactions (electromagnetism, weak nuclear force and strong nuclear force); and *matter fields* for the fundamental particles (Hobson 2017, 60).

Regarding the primacy of fields over particles, the Nobel-winning physicist Frank Wilczek writes:

“Undoubtedly the single most profound fact about Nature that quantum field theory uniquely explains is *the existence of different, yet indistinguishable, copies of elementary particles*. Two electrons anywhere in the Universe, whatever their origin or history, are observed to have exactly the same properties. We understand this as a consequence of the fact that both are excitations of the same underlying ur-stuff, the electron field. The electron field is thus the primary reality” (Wilczek 2006, 335-336).

As the above quote indicates, the resemblance among electrons – their qualitatively identical properties – is accounted for by a single fundamental reality – the field. Given this fact, it is

not too much of a stretch to say that the electron field plays the role of the universal that is commonly instantiated by its excitations – the electrons.

In what follows, I will adopt the intrinsically concrete view of universals because it harmonizes the ontological and metaphysical explanatory priority of universals over particulars, and it is in line with the best theory of natural science about what our universe is like.

7.2.4.2. Universal Monism:

To recall, there are two broad views about kinds of universals – universalist dualism and monism. Dualism (Aristotle and Nyāya) holds that there are both substantial universals (i.e., kinds) and non-substantial universals, whereas monism (Armstrong) holds that there is just one category of universals. Adjudicating between universalist dualism and is not a straightforward affair. *Prima facie*, dualism seems to have more metaphysical primitives than monism – substantial universals as a distinct category in addition to non-substantial ones. However, monism posits non-mereological combination/conjunction of universals, which is admittedly mysterious. Nevertheless, the Occam's Razor principle seems to favour monism. Therefore, I will proceed accepting the view that there is one category of universals which can form complex (conjunctive and structural) universals, as Armstrong argues.

Corresponding to universalist monism is instantiator or particular monism – there is only one kind of particular that instantiates universals, and this is the substance. Thus, I do not accept modes as instantiators. However, when a substance instantiates universals, they become the property-instances of the substance.

7.2.4.3. Role of universals:

Common to all the views discussed above is the duality of roles for universals – they are posited to account for resemblance and difference among particulars, and the causal powers of instantiating substances⁸⁸. In the contemporary literature on the metaphysics of powers, these roles are glossed as two aspects of a property (or two distinct properties themselves) – the *categorical* or *qualitative* aspect, and the *dispositional* or *powerful* aspect, respectively (Allen 2016, 140). These two aspects correspond to the standard way of understanding properties, which is that (a) properties are ways things *are* (Lowe 2006, 90; Heil 2012, 4) – their categorical/qualitative side, and (b) properties ground ways in which things act – their dispositional/powerful side. Regarding the dispositional aspect of properties, Heil writes that

⁸⁸ To be sure, Aristotle gives a causal role to the instantiating substance as well.

“...we appeal to objects’ properties to explain what those objects *do* or *would do*. This suggests a connection between properties and *powers* or *dispositions*” (Heil 2012, 59). For example, by possessing the property of sphericity (assuming there is such a property), a ball is spherical (a quality) and, at the same time, is disposed to roll over smooth surfaces (a power).

Related to these two aspects, there is a thriving debate about what the essence of a property is. There are three options: (a) The essence of a property is exhausted by its qualitativity; that is, a property is purely a categorical property or quality; (b) The essence of a property is exhausted by its powerfulness; that is, a property is a pure disposition or power; or, (c) A property has a mixed essence – it is both a power and a quality.

Option (a) is known as *categoricalism* (Armstrong 1983, 1996). Philosophers who endorse categoricalism explain the powers of a property in terms of laws of nature (eg., Armstrong 1983). Option (b) is known as *dispositional essentialism* (Bird 2007). Dispositional essentialists explain the qualities of a property in terms of its powers. More precisely, the way a thing is explained in terms of the manifestation of the thing’s power to act.

Option (c) is known as the *two-sided view*, according to which, “there is just one kind of property...[which] have two sides to them. They have a qualitative side... But they also have a power side” (Mumford 2007, 84). The two-sided view itself comes in two varieties – the dual aspect view and the powerful qualities view. The dual aspect view (Williams 2019, 113-119) holds that a property has a *composite* essence. In other words, it is dual aspects – a powerful and qualitative aspect. In contrast, the *powerful qualities* view holds that “[a] property’s dispositionality and its qualitativity are not aspects or properties of the property. Rather they inseparably constitute its nature” (Heil 2012, 61). In other words, the powerfulness of a property is identical to its qualitativity; hence this view is also known as the identity view. It is endorsed prominently by Martin (2008) and Heil (2003, 2012)⁸⁹.

Although an argument for it is outside the scope of my thesis, I am inclined to endorse the powerful qualities view, because it can account for both the powerfulness and qualitativity of properties in the simplest way – identity. Thus, I will endorse the claim that universals are powerful qualities. Of course, substances have the power and the quality only when universals

⁸⁹ Related to the essence-of-properties debate is the debate pertaining to how many kinds of properties there are. The options are (i) Monism - there is just one kind of property. Monism is of three types (a) dispositional monism – all properties have dispositional essences; (b) categorical monism – all properties have categorical essences; (c) mixed monism – all properties have two-sided essences; (ii) Dualism – there are properties of two kinds – some properties are pure powers and some are pure qualities.

are instantiated in substances. Universals can, therefore, be interpreted as the *source* of the powerfulness and qualitativity of substances. Further, I will assume that all universals are first-order universals; thus, I do not need to adopt Armstrong's view that it is higher-order universals (i.e., laws) that account for the causal powers.

7.2.4.4. Instantiation Monism:

A significant characteristic of the instantiation or inherence relation that is brought to relief in the above discussions is the plurality in the roles that it plays. To recall, Aristotle argued some universals are predicated of substances, while others inhere in substances; and this difference is accounted for partly by the manner of instantiation – predication-instantiation and inherence-instantiation, respectively. Armstrong interprets the manner of instantiation as an intersection between substances and universals.

Nyāya-Vaiśeṣika is extremely plural about the roles of inherence/instantiation – positing that universals inhere in their particulars (substances or modes or actions), wholes inhere in their parts, effects (constituteds) inhere in their causes (constituters), and particularity inheres in substances. Given the diversity of the pairs linked by inherence, it is plausible to assume that there are different manners of inherence for each pair.

Now, it could be debated whether there are different manners of the same relation here, or there are different relations altogether. Lowe takes the latter route, and posits three different relationships (he does not call them relations, since he considers them primitive) – instantiation, characterization and exemplification (Lowe 2006, 18). However, the demand of ontological parsimony weighs the balance in favour of the 'single relation-different manners' view. This is because there is something common to all the different manners of instantiation – they connect or tie the relata (the instantiated and the instantiator) in a primitive way. I use 'primitive' in the sense of unanalyzable in terms of other relations, and not requiring further relations to relate the relata. Endorsing multiple primitive relations is surely metaphysically profligate. Moreover, the assumption of a single relation with different manners does not compromise on the explanatory benefits of instantiation. Due to the above reasons, I will adopt the view that there is a single instantiation relation which can have different *manners of instantiation*.

A final point about instantiation is that I will consider it to be full-fledged relation, rather than as a *relationship* – as Lowe (2006, 91) considers it – or as a connection or tie – as Armstrong (Armstrong 1989, 109) – interprets it. This puts instantiation in the same broader kind as naturalistic relations. The only difference is that instantiation is a primitive relation – it is

unanalyzable into other relations, and it does not require any further relations to relate its relata. More precisely, instantiation is an external relation that is not reducible to the intrinsic properties of its relata, and is over and above them.

7.3. Content Universalism

7.3.1. Assumptions

In this section, I set out the details of content universalism. The following are my assumptions, from the previous section:

- i. Intrinsically concrete universals: Universals are concrete independent of their instantiation/immanence. However, they are often instantiated.
- ii. Universalist monism: There is one category of universals – non-substantial simple universals. However, these simple universals can form complex conjunctive universals, as Armstrong argues.
- iii. Instantiation monism: There is one kind of instantiation relation. However, there are multiple ways or manners in which universals are instantiated in substances.
- iv. Universals are powerful qualities: They ground the qualities of substances and are the source of their causal powers.

Now for the view proper. On a first approximation, content universalism can be formulated as follows:

Content Universalism (interim version): The content of perceptual representation *R* is partly constituted by being *perceptually related* to a target *T* in the performance of biological functions, and partly constituted by the *co-instantiation* in the subject *S* of the same universals U_i instantiated in *T*.

The details of content-constitution now need to be fleshed out. Before getting into these details, three points need to be clarified. First, by ‘content’ hereafter I will refer exclusively to *presentational content* – the way in which target is presented. Semantic content is determined by presentational content. Second, to reiterate, I consider the target of representation to be a spatiotemporal region of the environment rather than particular objects (as argued for in chapter 2, §2.2.1.2). There will plausibly be several universals instantiated in the target region – simple

and conjunctive – which would then constitute distinct states of affairs and propertied objects in the target region.

Third, I interpret constitution in the Bakerian sense (chapter 1, §1.4.9) as supplying necessary and sufficient conditions for the constituted entity (in this case, content). Part constituters are, therefore, only necessary for content. Content only ontologically depends on its part constituters. Moreover, the constituted entity – content – is not reducible to the constituters, because it has distinct causal powers over and above the constituters⁹⁰.

7.3.2. Constitution of Content

Regarding the process of content-constitution, I submit that it plausibly involves three elements (naturalistic relations, biological functions and co-instantiation of target universals) working in two stages:

1. The holding of (i) naturalistic relations between the subject and the target for the performance of (ii) biological functions.
2. (iii) *Co-instantiation* in the subject of the universals instantiated in the target.

The above three elements are individually necessary and jointly sufficient for content, without content being reducible to them.

7.3.2.1. Content-constituting naturalistic relations:

I argued, in chapter 3 (§3.3), that content-constituting naturalistic relations – such as causation and structural correspondence – are necessary for content. My argument was based on the premise that the essence and purpose of perception is to guide actions in order to ensure fitness and survival. This implies that perceptual content is essentially action-oriented. That is, perceptual content presents environmental targets (i.e., spatiotemporal regions) as possessing action-properties – properties that are relative to the subject's action-capacities, and therefore relevant to the subject's dispositions to act in the environment. Action-oriented content, in turn, implies that it is constituted by relations to environmental targets.

7.3.2.2. Fitness-contributing functions

I also argued in chapter 3 (§3.4.7) that fitness-contribution functions are necessary for content-constitution. This follows from the essence of perception which is to guide actions for fitness

⁹⁰ More on the causal powers of content in chapter 8.

and survival. Perception enables both ecological fitness and reproductive fitness. Ecological fitness is interpreted in terms of how well an organism can interact with and navigate its environment, whereas reproductive fitness is defined in terms of the number of offspring an organism is likely to have.

7.3.2.3. Co-instantiation of universals of target:

In chapters 3 and 4, although I agreed with naturalistic representationalist theories about the necessity of relations and functions, I argued that their third putative necessary element – naturalistic representational application formulae – was not necessary for content and, therefore, representation. Here, I want to posit an alternative final necessary element. This is based on Aristotle's insight that perception involves the reception of the form of the perceived object. I submit that the final necessary element is the instantiation in the perceiving subject of the same universals that are instantiated in the target. I will call this the *co-instantiation* requirement.

Before expanding on the co-instantiation requirement, I want to set aside another proposal for how universals may constitute content. This is the proposal that we are *acquainted* with immanent universals which then constitutes perceptual content. As discussed in chapter 6, this strategy is adopted by many representationalist theories, which posit acquaintance relations with abstract universals. However, I argued there (§6.2.3) that acquaintance cannot be a content-constituting relation. To reiterate briefly, acquaintance is a direct conscious awareness relation, which is of the nature of encountering or confronting the acquainted entity. However, (a) we can have unconscious perception, which acquaintance does not explain; (b) we do not *perceive* universals; rather we perceive concrete environmental targets; and (c) both the relata of the acquaintance relation need to be spatiotemporal; since abstract Platonic universals are non-spatiotemporal, it is hard to say how we can be acquainted with them.

The final content-constituting element, therefore, is the co-instantiation in the subject of the same universals which are instantiated in the target. I will call these universals the 'target-universals'. Content universalism, insofar as it posits co-instantiation of target-universals in the subject, can be considered as a precisification of Aristotle's view that perception involves the reception by the sensory faculty of the universals of the perceived object.

7.3.2.3.1. Permeation manner of instantiation:

Prima facie, the idea of target-universals being instantiated in the subject would seem unintuitive and even wrong. Suppose that redness is a universal that is instantiated in a tomato (the target). The objection to my proposal would be: if perceiving a red tomato involved the subject instantiating redness as well, then this implies that subject or their brain (or a region of it) is red. Surely, this is not the case. To stymie such an objection, I want to invoke – from the discussion in §7.2.4 – the insight that there are different manners in which universals can be instantiated in their particulars. I submit that the manner in which universals are instantiated in the target is different to the manner in which the same universals are instantiated in the subject.

What justifies the difference in the manner of instantiation in the target and the perceiver? I submit that it is their essences or nature. The nature of target *qua* object is completely different to the nature of the perceiver *qua* subject. The perceiver *qua* subject is characterized by subjectivity or the first-person perspective or the property of “for-me-ness” (Zahavi and Kriegel 2016); whereas the target *qua* object is what is perceived. Therefore, I argue that the manner in which a universal is instantiated in an entity *qua* object is different to the manner in which it is instantiated in an entity *qua* subject. It follows that the same universals are instantiated differently in the target and perceiver. Now, to be sure, the perceiver is not only a subject but also an object himself – since he can possess properties such as mass, surface reflectance, etc. In such a case the universals which are instantiated in the perceiver *qua* object and *qua* subject are instantiated in different manners.

What are the different manners of instantiation in the target and the subject? Since I assume that there is only category of universals (universalist monism), I will also assume that there is only one manner of instantiation of universals in an entity *qua* object. I will call this manner, following Marmodoro (2021), *qualification-instantiation*. Therefore, universals are qualification-instantiated in targets (*qua* objects); in short, *universals qualify targets*. Of course, as Aristotle argued, there can be fine-grained manners of qualification-instantiation itself – particularly, predication-instantiation (universals are ‘said of’ targets *qua* objects) and inherence-instantiation (universals are present in targets). However, I will work with only the more general qualification-instantiation.

Now, coming to the manner in which universals are instantiated in *subjects*. The universals that are instantiated in subjects during perception do not qualify them as they do targets – the target-universals are neither predicated of subjects nor do they inhere in subjects. Therefore, I want

to propose a *sui generis* manner in which universals are instantiated in subjects *qua* subjects – the *permeation-instantiation*.

Permeation has connotations of entering, being present in, and spreading. Thus, the target-universals enter and are present *in* the mind of the subject during perceptual representation. Since permeation is instantiation, the universals that permeate the subject are the *properties* of the entire subject, rather than of any specific part (eg. brain) or even of the representational vehicles that the subject tokens. This is the spreading aspect of permeation in action. Permeation, in general, also has connotations of temporary presence in the locus. This applies to perceptual representation as well. Thus, universals are in the mind of the subject but only for the duration of the perceptual episode, unlike in the case of targets, where universals qualify them for as long as they exist. In sum, in perception, *whereas universals qualify targets, they permeate in subjects*. To put it differently, in perception, there is objective (i.e., in target) qualification by universals, on the one hand, and subjective permeation of universals, on the other. In both cases, universals are immanent.

7.3.3. Content As the Immanent Permeated Universal

Taking stock, I have argued that content is constituted by (i) naturalistic relations between subject and target; (ii) fitness-contributing functions of the subject; and (iii) subjective permeation of the same concrete universals that qualify the targets. Given these three individually necessary and jointly sufficient conditions, content can be said to be identical to the permeated universals. To put it differently, content is a property-instance of the subject due to permeation-instantiation of universals.

An example will be helpful at this juncture to illustrate the permeation of concrete universals and resultant tokening of content. Suppose in an environmental target (spatiotemporal region), there are several universals that are qualification-instantiated. There are monadic simple universals corresponding to shape, edges, texture, etc. Let us call these Shapeness, Edgeness, Textureness. Suppose further, as Armstrong suggests, that simple universals can form conjunctions to form complex or conjunctive universals. Complex universals would include universals such as ‘Surface Reflectanceness’ (plausibly a conjunction of simple universals Edgeness $\wedge$ Textureness) and Kinds (a highly conjunctive or structural universal). Suppose still further that there are polyadic universals as well, whose instances are the external relations between the instances of monadic universals. All of these universals together *qualify* the target and carve it up into particular states of affairs and objects.

Along comes a subject who has the function of navigating the environment to ensure fitness and survival, and he is naturalistically related to the target region (by causation or structural correspondence). Now, as noted in chapter 2 (§2.3.1), perception occurs in stages and corresponding to the stages there are different stages of presentational content. Let us suppose that in the initial stages of perception, simple target-universals such as Shapeness and Edgeness permeate the subject. This mix of functions+relations+permeation constitutes non-conscious content of early-stage perceptual representations, where perceptual representations represent edges and shapes, plausibly in 2-D or 2½-D sketch as Marr hypothesised (Marr 1982, 37).

At later stages of perception, complex target-universals such as Surface Reflectanceness₁, Surface Reflectanceness₂, Kind₁ and so on permeate the subject. These qualify the target as the property instances Surface Reflectance₁, Surface Reflectance₂. When Surface Reflectanceness₁, Surface Reflectanceness₂ permeate the subject, they permeate as the property-instances ‘Red’ and ‘Blue’ respectively. However, the subject is not literally red or blue. These property-instances are the content – they are the subjective presentation of (and information about) the corresponding objective property-instances Surface Reflectance₁ and Surface Reflectance₂. Thus, whereas Surface Reflectanceness₁ qualifies the target as the property-instance Surface Reflectance₁, it permeates in the subject as the content <Red>; similarly, Surface Reflectanceness₂ qualifies the target as Surface Reflectance₂, but it permeates in the subject as the content <Blue>. Thus, the same universal, due to different manners of instantiation, is instantiated as different property-instances in the target and subject. The dual instantiation is depicted in the diagram in the next page:




Figure 7.1.: Constitution of content and the Representational External Directedness (RED) relation⁹¹ (U: Universals; PI: Property-instance; C: Content; R: Representation).

A significant implication of content universalism is that the permeation of complex universals is what accounts for conscious content (more on this in chapter 9, §9.4.3). Phenomenal qualities such as ‘red’, ‘low-pitch sound’, ‘the smell of vanilla’, and so on, are not in the world but are the contents of our representations – they are the ways (eg. reddish way) in which the target is represented due to the permeation of target-universals. A content which presents the target in phenomenally characterized ways or, alternatively, a content which is identical to phenomenal qualities is known as “phenomenal content” (Kriegel 2002, 180).

As immanent universals, contents, whether phenomenal or not, are properties of the subject. However, since they are permeated universals, they do not characterize the subject in the same way qualified universals would characterize targets. For example, the phenomenal content <Red> does not make the subject or the subject’s mind literally red. Rather, contents (= permeated universals) are properties only in the sense that they are present in the mind of the subject.

⁹¹ Content-Representation pairing will be explained in §7.3.5, and the RED relation in chapter 9.

Phenomenal contents are representational properties of the subject, because they partly ground the Representational External Directedness (RED) relation between a perceptual representation and the target (more on this in chapter 9). Thus, they must be distinguished from intrinsic non-relational phenomenal properties of the representation or the subject, which are also known as ‘qualia’ (Block 2004). These non-representational properties are posited by adverbialism (Fish 2021, 85) and by non-intentionalist representationalist theories such as Papineau’s (2021)⁹². In contrast, phenomenal content is a representational and a relational property (possessed on account of naturalistic relations and the permeation relation).

7.3.4. Content as Powerful Quality

In sum, the same universal when qualifying targets and permeating in subjects amounts to different property-instances. This difference – for example, Surface Reflectance₁ vs Red – is a difference in the categorical/qualitative side of property-instances. Correspondingly, there is a difference in the causal powers the universal bestows to the instantiators – the target and the subject. For example, Surface Reflectanceness₁ qualifies the target as the property-instance Surface Reflectance₁. This grounds the target’s causal power of reflecting light that has a wavelength between 620 and 700 nanometres. The same universal permeates in the subject as the content <Red>. This has different causal powers, plausibly the power of emotionally affecting the subject so that the subject acts differently in different contexts (Maier et al. 2016) – for example, to avoid danger (when in the presence of blood, an angry opponent, and danger signals) or to obtain food (in the context of ripe fruit, meat, etc.).

Although in the above example, the universal is the same, the difference in causal powers of the property instances (Surface Reflectance₁ vs Red) is accounted for by the difference in the instantiators (target vs subject) and the manners of instantiation (qualification vs permeation). To appreciate the role of the instantiator and the manner of instantiation in bestowal of causal powers, we need only look at Aristotle’s view of universals and instantiation. First, Aristotle held that causal powers were grounded in both the universal/form (principle of actuality) and matter (principle of potentiality). Thus, the instantiator makes a difference to the eventual causal power tokened. Second, the idea that the manner of instantiation makes a difference, although not explicit in Aristotle, is inspired by his view in the following way. To recall, Aristotle distinguished between predication- and inherence-instantiation in substances. For example, suppose the universal ‘Surface Reflectanceness₃’ (or ‘Orangeness’) inheres in

⁹² I will discuss Papineau’s view of phenomenal character in chapter 9, §9.4.3.

Garfield but is predicated of his mode ‘Surface Reflectance₃’. Owing to this difference, there is a difference in causal powers. The mode has the power of reflecting light of a certain wavelength, but Garfield has the power to attract humans. Therefore, I submit that while universals are the source of an instantiator’s causal powers, the causal powers are modified by the nature of the instantiator (subject vs object) and the manner of instantiation (permeation vs qualification).

7.3.5. Content-Representation Pairing

In the above subsection, it was proposed that content is identical to the permeated universal in the subject. Therefore, content is the property of the subject. But perceptual representations, which are states of the subject, are said to *have* content. How can we associate a subject-wide property with a specific perceptual representation that is tokened in the brain of the subject? This may be called the *content-representation pairing* problem. I submit a solution to this problem.

Representational vehicles are the distributed set of neural firings in the subject that are the causal outcome of the naturalistic relations borne to the target. Now the distributed set of neural firings (N) is the property of the subject *s*. It is the qualification-instantiation of a universal (plausibly a conjunctive one) in the subject *qua* object. The state of affairs of the subject having this property can be expressed as *sN*, and it is this state that is the representational vehicle. Similarly, the permeated universal = content (C) is also a property of the subject *qua* subject. This state of affairs can be expressed as *sC*. Now both states *sN* and *sC* are states involving the same subject *s*; and the properties R and C are concurrently colocated in *s*. I submit that this concurrent colocation and the naturalistic relations involved – which is a common denominator between them – are what warrant our pairing of the content C with the representational vehicle *sN*.

Can the relation of concurrent colocation be sufficient to account for content-representation pairing? Yes. This is because it is a spatiotemporal relation. And Kim has argued, in the context of causation, that spatiotemporal relations are “...the only way of generating pairing relations” (Kim 2011, 51). That is, the cause and effect can be paired with each other in virtue of their spatiotemporal relations. Of course, in the case of causation, the spatial relation is not exact colocation and the temporal relation is of succession. Nevertheless, the general point about spatiotemporal relations being ideal pairing relations applies. If spatiotemporal relations are

sufficient to pair cause and effect, then they are also sufficient to pair the colocated states – content and representation.

7.3.6. Content Universalism (Final Version)

To sum up my discussion of content-constitution, content universalism can be formulated in its final version as follows:

Content Universalism (final version): The content of a perceptual representation R is constituted by: (i) naturalistic relations between a subject S and target T for the performance of fitness-contributing biological functions, and (ii) the permeation-instantiation in S of the same universals U_i which are qualification-instantiated in T. Given the above conditions, the content is identical to the permeated universals U_i .

7.4. Conclusion

In this chapter, I developed a theory of content – Content Universalism – as part of my larger representationalist theory PRIMER – Pluri-Relational Immanent Emergent Representationalism. As the final formulation of content universalism suggests, the three individually necessary and jointly sufficient conditions for the constitution of content are: (i) naturalistic relations between a subject and the target; (ii) fitness-contribution biological functions of the subject; and (iii) subjective permeation-instantiation of the same universals that are qualification-instantiated in the target. Although they are jointly sufficient for content, content is not reducible to them. In fact, in the next chapter, I will argue that content is a strongly emergent property with novel causal powers and so cannot be reduced to its constituting conditions.

Furthermore, fitness-contribution functions and the subjective permeation of universals complement each other in the explanatory work they do. While functions give a teleological explanation of why content and representations exist, subjective permeation of universals gives a metaphysical explanation of how content is constituted.

8. Content Universalism: Emergent Content

In this chapter, I argue that not only is content an irreducible property of the subject but also a strongly emergent property. In §8.1, I argue that content is not reducible under either of two plausible models of reduction – functional reduction and determinative reduction. In other words, content is an emergent property. The only question that remains is whether it is weakly emergent or strongly emergent. This leads to a discussion about ontological emergence in §8.2. Ontological emergence is characterized by the twin marks of ontological dependence and distinct causal efficacy of the emergent entity. I argue for the strong emergence of content because the distinct causal efficacy of content is due to its possessing novel causal powers over and above its dependence/emergence base. Finally, in §8.3, I tackle a challenge that content universalism, as a version of content generalism, faces – the challenge of explaining phenomenological particularity of content. I argue that this is accounted for by attention, which is a post-perceptual psychological mechanism.

8.1. Irreducibility of Content

I argued in the previous chapter that, according to content universalism, naturalistic relations, biological functions and subjective permeation-instantiation of universals are individually necessary and jointly sufficient conditions for content. This schema looks similar to naturalistic representationalism's reductive project of reducing content to their preferred individually necessary and jointly sufficient conditions. So, does content universalism imply reducibility of content? I argue that it does not.

To recall from chapter 1 (§1.4.8), there are at least four routes to reduction – elimination, identity, functionalization and determination. Elimination and identity can be straightforwardly eliminated as possible routes to content reduction. That leaves functionalization and determination. Now, naturalistic representationalism can argue for the reducibility of content via either functionalization or determination. But neither of these options is applicable for content under content universalism.

As Kim (2005) explains, to functionalize an entity is (i) to define the entity entirely in terms of its causal-functional role in a system; and then (ii) to identify physical realizers of the causal-functional role. In other words, the reducible entity's causal powers, if any, are entirely the

causal powers of the realizing entity. As noted in chapter 5 (§5.1.3), functional reduction is the most popular model of reduction in naturalistic representationalism.

Content (= permeated universal) cannot be functionally reduced to its constituting base – naturalistic relations, fitness-contributing functions and co-instantiation of target universals – for two reasons. First, content cannot be functionalized. As argued in chapter 7 (§7.3.4), content is a powerful quality – it not only bestows causal powers on the subject, but also is a qualitative property that determines the presentational aspect of the representation. In the specific case of phenomenal content (of perceptual experiences), the qualitative aspect is identical to phenomenal character/qualities. Now, Kim, who proposes the functional model of reduction, himself writes that phenomenal qualities cannot be functionalized and are, therefore, irreducible: “...qualia are not functionalizable, and hence physically irreducible. Qualia, therefore, are the “mental residue” that cannot be accommodated within the physical domain. This means that global physicalism is untenable” (Kim 2005, 170)⁹³. Secondly, as argued in chapter 3 (§3.3.2), content is action-oriented, which means it can cause actions in subjects that are directed at environmental targets. Content consists of action-properties for this purpose. The action caused by content is not merely the bodily movements of organisms. Thus, content has causal powers that cannot be identified with the causal powers of any of its constituting conditions. Hence, content cannot be functionalized.

Content cannot even be determinatively reduced. As noted in chapter 1 (§1.4.8), determination is a combination of dependence, grounding and metaphysical necessitation. Content, according to content universalism, is not determined by the three elements because, although it is dependent on and grounded by the latter, it is not metaphysically necessitated by them. This is because of the possibility of illusions and hallucinations. As I will explain in the next chapter (§9.4.2), hallucinations are a result of permeation of universals that do not qualify the target – I call these deviant universals. Thus, even if the three conditions obtain, deviant universals may be permeation-instantiated, which means that content is not metaphysically necessitated by the three conditions.

The failure of either functionalization or determination is ample motivation to claim that content is an irreducible property of the subject. The irreducibility of content still leaves us

⁹³ Kim separates the intentional and the phenomenal sides of perceptual experiences and argues that whereas the intentional aspect can be functionalized and is, therefore, reducible, the phenomenal side cannot be (Kim 2005, 174). Consequently, he argues for an attenuated or “defective” (Ibid., 174) reductive physicalism.

with two options regarding the nature of content – content is either weakly emergent or strongly emergent. In what follows, I will discuss these two options and argue that it is plausible that content is a *strongly* emergent property. Before beginning, it should be noted that I will be using emergence in the sense of *ontological emergence* rather than epistemic emergence. According to ontological emergence, the putative emergent properties have causal powers that are distinct from their emergence base. Epistemic emergence says that a putatively emergent property of a system unpredictably or unexpectedly emerges from lower-level domain, given an adequate knowledge of the properties and the laws governing the lower-level domain.

Accordingly, the distinction between strong and weak emergence that I advert to is a distinction between types of ontological emergence. Some philosophers, however, argue that the strong/weak emergence distinction just maps on to the ontological/epistemological emergence – for example Chalmers (2006), who writes that “a high-level phenomenon is *weakly* emergent with respect to a low-level domain when the high-level phenomenon arises from the low-level domain, but truths concerning that phenomenon are *unexpected* given the principles governing the low-level domain” (Chalmers 2006, 245). However, as Hendry et al. (2019, 2) note, there are arguments to the effect that the strong/weak distinction is present in both ontological and epistemic emergence.

8.2. Ontological Emergence

Ontological emergence of any entity (usually a property) is characterized by two marks: ontological dependence on lower-level entities (the emergence base) and causal autonomy of the emergent property vis-à-vis its base. Thus, “an emergent property...is one that has causal powers that are distinct from the causal powers of the lower level properties on which it [depends]” (Crane 2001, 216). The ontological dependence is usually understood as synchronic (Anjum and Mumford 2017, 99) or cotemporal (Wilson 2021, 1) dependence. This means that once the emergent entity comes into existence, it and the base co-exist. However, some – for example, Humphreys (2008, 2016) and O’Connor and Wong (2005) – argue that the dependence is diachronic in the sense that the base properties either cease to exist or are transformed after the emergence. Corresponding to the synchronic and diachronic dependence, there is a distinction between synchronic and diachronic emergence.

An example of a diachronic emergence is Humphreys’ transformational emergence:

“Transformational emergence occurs when an individual a that is considered to be a fundamental element of a domain D transforms into a different kind of individual a^* , often but not always as a result of interactions with other elements of D , and thereby becomes a member of a different domain D^* . Members of D^* are of a different type from members of D . They possess at least one novel property and are subject to different laws....” (Humphreys 2016, 60).

In the case of content universalism, content is dependent on its constituting elements, such that there is no transformation of the elements nor do they cease to exist. Thus, the dependence is synchronic or cotemporal. Consequently, the account of emergence that applies to content is synchronic emergence. Hereafter, I will refer to synchronic ontological emergence as merely ‘emergence’.

The primary characteristic of emergence is the causal autonomy of emergent entities. This autonomy can be either understood in a weak or a strong way, which gives the two kinds of emergence – weak and strong emergence. Weak emergence is also known as “realization” (Baysan 2019, 79; Wilson 2021, 55).

8.2.1. Weak Emergence

A property is said to weakly emergent from its emergence base if its causal powers are derived or inherited from (Kim 1992, 18) its emergence base. This inheritance is commonly understood in terms of the weakly emergent entity’s causal powers being a subset of the causal powers of its emergence base. Accordingly, Wilson (2021) explains weak emergence as follows: “What it is for token feature [M] to be Weakly...emergent from token feature P on a given occasion is for it to be the case, on that occasion, (i) that [M] cotemporally materially depends on P, and (ii) that [M] has a non-empty proper subset of the token powers had by P (Wilson 2021, 72). The subset account implies that the causal powers of the weakly emergent entity are “*token identical*” (Ibid., 58) – i.e., numerically identical – to the causal powers of the emergence base.

Despite the causal powers of the weakly emergent property being token identical to its emergent base, it is argued that it can have distinct causal efficacy as compared to its base. This is what undergirds the irreducibility of the weak emergent property to its lower-level base. Shoemaker – who calls weak emergence by the name “realization” (Shoemaker 2001, 2007) – argues for the distinct causal efficacy of the weakly emergent property by applying the proportionality criterion. As he puts it,

“[w]here the only causal features of property [P] that play a role in producing an effect are ones that belong to property [M], of which [P] is a...realizer property, there seems a good sense in

which considerations of proportionality favour the instantiation of [M] over the instantiation of [P] as a cause of the effect” (Shoemaker 2001, 81).

That is, Shoemaker argues that since the causal powers of the realized property M are just enough to produce an effect, they are more proportional to the effect than the causal powers of the lower-level realizer property P which can causally produce *more* than the effect.

8.2.2. Strong Emergence

A property is said to be strongly emergent if has at least one novel causal power – that is, at least one causal power over and above the causal powers of its lower-level dependence base. Wilson gives a formal definition of strong emergence as follows:

“What it is for token feature [M] to be Strongly...emergent from token feature P on a given occasion is for it to be the case, on that occasion, (i) that [M] coterminally materially depends on P, and (ii) that [M] has at least one token power not identical with any token power of P” (Wilson 2021, 53).

Related to the causal novelty of emergent entities is another feature – emergent entities are fundamental or basic, despite being dependent on their bases. Thus, Barnes writes that “[a]n entity x is ontologically emergent iff x is fundamental and dependent” (Barnes 2012, 884). Similarly, Baker claims that emergent properties are “ontologically basic” (Baker 2019, 210), while Wilson writes that they are “fundamentally novel” (Wilson 2021, 48). The fundamentality of strongly emergent properties emphasizes their ontological novelty aside from their causal novelty.

There are several ways the fundamentality of strongly emergent properties is understood. Most importantly, it is understood as non-additiveness – a strongly emergent property is ‘over and above’ its base (Barnes 2012, 885; Wilson 2021, 47). That is, it is not merely a summation or arrangement of its base entities. Barnes also interprets fundamentality of strongly emergent properties in terms of truthmaking. The idea is that the emergence base cannot be a sufficient truthmaker for the existence of the emergent property. “You will need the emergent thing itself (and it looks plausible to think that the emergent entity will serve as the truthmaker for its own existence)” (Barnes 2012, 885).

O’Connor and Wong (2005) interpret fundamentality in terms of nonstructurality. A structural property of an object is simply its property of having parts which stand in certain relations – such as the arrangement relation – to one another. In other words, “...there is *nothing more* to

having the structural property than being composed by parts having certain other properties and bearing certain relations to one another – it is ontologically reducible” (O’Connor and Wong 2005, 663). In contrast, strongly emergent properties are “*wholly nonstructural*” (Ibid., 663). They are “...basic properties, token-distinct in character and propensity from any microphysically structured properties of their bearers” (Ibid., 664).

8.2.3. Strong Emergence of Content

Having noted the difference between weak and strong emergence, the pertinent question is whether content is a weakly emergent or a strongly emergent property of the subject. The dependence/emergence base of content includes its constituters – naturalistic relations between the subject and the target, fitness-contributing functions of the subject and co-instantiation of the universals in the target and subject. The emergence base also includes that subject who permeation-instantiates target universals. To adjudicate the weak or strong dependence issue, we need to look at the causal powers of content.

In chapter 3, it was argued that perceptual content is action-oriented. That is, perceptual representations present targets as having action-properties, which are “properties that need to be represented in order for the agent to perform an action” (Nanay 2013, 4). These action-properties are part of the presentational content of representations – they present targets as affording actions. Thus, action-properties are always attributed to targets relative to the subject, and hence are also partly constituted by the subject-agent’s action capacities. The action-oriented nature of content has been supported by several philosophers. For instance, Millikan (1995, 2023) argues that perceptual representations are pushmi-pullyu representations having both descriptive and directive content. The directive content is closely tied to the biological functions of the perceptual subsystem that consumes the representation. Similarly, Cussins (2003) argues that perceptual experience represents the world as a realm of activity trails, which are possibilities of action. Along similar lines, Nanay (2013) argues that perceptual representations are “pragmatic representations” that represent targets as having action-properties.

Content universalism accounts for the action-oriented nature of content. When target universals (those that qualify the target) permeate in the subject, they become content. Now, due to the difference in the manner of instantiation (permeation vs qualification) and the nature of the instantiator (subject vs target), the content (= permeated universal) in the subject does not *resemble* the property-instances (= qualified universal) in the target. Rather, the content

presents the target as having action-properties. For example, the universal ‘Surface Reflectanceness₁’ qualifies the target as the property-instance ‘Surface Reflectance₁’. The same universal permeates in the subject as the content <Red>. Now, <Red> is not only a phenomenal content, but also an action-oriented content. It attributes the action property ‘being red’ to the target. It has the causal power, plausibly, of emotionally affecting the subject so that the subject acts differently in different contexts (Maier et al. 2016) – for example, to avoid danger (when in the presence of blood, an angry opponent, and danger signals) or to obtain food (in the context of ripe fruit, meat, etc.).

From the above considerations, it is clear that content has distinct causal efficacy. My further claim is that the causal powers of content are totally novel vis-à-vis its emergence base. This is because the powers to cause actions in organisms, rather than merely reflexive bodily movements, are not the causal powers of either the universals permeating in the subject or of the naturalistic relations with the target. The novel causal powers can be explained as due to the contribution of the subject’s intrinsic properties qua subject, the subject’s relational properties (due to standing in naturalistic relations), the permeating universals and the permeation-manner of instantiation. Therefore, I submit that content is a strongly emergent property.

Even on the ontological novelty/fundamentality front, content clocks in as strongly emergent. It ticks all the criteria for emergent fundamentality. First, it is ‘over and above’ its base entities, including the universal that permeates subjects. Second, it is a distinct truthmaker for truths about representation. Finally, it is non-structural and simple with respect to the subject, even if the universal that permeates the subject is a conjunctive or structural universal. This is because content is not a property that is defined in terms of the micro- or macro-structure of the subject permeation-instantiating the universal.

Given the above considerations of causal and ontological novelty, I conclude that content is a strongly emergent property. It, therefore, follows that representation and intentionality are strongly emergent phenomena.

8.2.4. Downward Causation and the Causal Closure of the Physical

An implication of the strong emergence of content – and its concomitant novel causal powers – is the phenomenon of downward causation (Crane 2001, 216; Kim 2005, 40). Downward causation consists in emergent entities causally influencing their lower-level base entities. The

paradigm examples are of mental events causing physical actions. For instance, the desire M to have coffee (a mental state emergent from P) causes the motor neurons P* in my brain to fire leading to a causal chain that eventually results in me getting up from my chair.




Figure 8.1: Downward mental-to-physical causation

Many physicalists object to downward causation because it implies the rejection of a principle that is considered by many physicalists as a core dogma of physicalism – the Causal Completeness of the Physical (CCP)⁹⁴. A standard formulation of the CCP principle is: “All physical effects have complete physical causes (‘complete’ in the sense that those causes on their own suffice by physical law to fix the chances of those effects)” (Papineau 1993, 22). A more precise statement of CCP that takes the levels of reality into consideration is: “Every lower-level physical effect has a sufficient purely lower-level physical cause” (Wilson 2021, 41).

Downward causation implies that the sufficient cause of physical effect P* needs to include the higher-level mental event M, apart from physical event P. Thus, it does not have a sufficient lower-level cause, and therefore this violates the CCP principle. In other words, the lower-level domain is not causally complete.

To be sure, weak emergentism – the thesis that accepts weak emergence – could also endorse downward causation. However, every case of downward causation would be considered consistent with the CCP principle, since the causal powers of the weakly emergent event M are a subset of and, therefore, token identical to those of the base event P. However, the downward causation endorsed by strong emergence involves novel causal powers and, therefore, violates the CCP principle. It is this violation of the CCP principle that arguably makes strong emergentism non-physicalist. For instance, van Gulick writes that “radically [i.e., strongly]

⁹⁴ The Causal Completeness principle is also known in the literature as Causal Closure (Gibb 2015). However, Marcus (2005) and Humphreys (2016) argue that they are distinct principles. The distinction is not important to my arguments; so I will ignore it.

emergent powers would pose...a direct challenge to physicalism, since they would threaten the view of the physical world as a closed causal system” (van Gulick 2001, 19).

However, others argue that a rejection of the CCP principle does not entail the rejection of physicalism because the CCP principle is not essential to physicalism (Anjum and Mumford 2017, 106). In fact, philosophers who argue for emergence in the natural sciences such as chemistry (Hendry 2010) and physics⁹⁵ (Gibb et al. 2019, Part 3) embrace the jettisoning of the CCP principle but not their physicalist credentials. I will look at Hendry’s arguments against the CCP principle below.

Hendry (2010, 2021) argues that there are empirical grounds for emergence and the rejection of the CCP in chemistry. A prime candidate for emergence in chemistry is molecular structure. Molecular structure is emergent from “quantum-mechanical systems of nuclei and electrons interacting via [mainly] Coulomb forces” (Hendry 2010, 190). More precisely, molecular structure is the result of the breaking of symmetry present in the emergent base. The asymmetries of the emergent structure imbue it with novel causal powers⁹⁶. For example, “the hydrogen chloride molecule...has an asymmetrical charge distribution that explains its acidic behaviour and its boiling point” (Ibid., 186). In this regard Hendry writes that “...if the acidic behaviour of hydrogen chloride molecule is conferred by its asymmetry, and the asymmetry is not conferred by the molecule’s physical basis according to physical laws, then *prima facie* ontological reduction fails” (Ibid., 187).

Hendry argues that the novel causal powers of molecules would include the powers of downward causation, even in a future amended version of quantum mechanics:

“On any conservative amendment to quantum mechanics, the explanation of why molecules exhibit the lower symmetries they do would appear to be holistic, explaining the molecule’s broken symmetry on the basis of its being a subsystem of a supersystem (molecule plus environment). This supersystem has the power to break the symmetry of the states of *its* subsystems without acquiring that power from its subsystems in any obvious way. That looks like downwards causation” (Hendry 2010, 187).

⁹⁵ For instance, Lancaster (2019) argues for the emergence of particle properties due to interactions in the context of Quantum Field Theory.

⁹⁶ Hendry explains that causal powers of complex systems, in general, are represented in quantum-mechanical Schrodinger equations by Hamiltonian operators. Emergent novel causal powers are represented by configurational or non-resultant Hamiltonians as opposed to resultant Hamiltonians for causal powers of realized systems (Hendry 2010, 185).

This, he argues, is a strong indication that symmetry-breaking attests to the implausibility of the CCP principle.

Hendry also argues that symmetry-breaking invalidates the main argument for the CCP principle – the “argument from fundamental forces” (Papineau 2002, 250), which is itself based on the principle of conservation of energy. The argument from fundamental forces starts with the premise that most natural phenomena are accounted for by a small number of fundamental physical forces, which conserve energy. Any “spontaneous special force” (Ibid., 249), for example a force supplied by a strongly emergent entity, would violate the principle of conservation of energy. Thus, the conclusion is that “...all apparently special forces characteristically reduce to a small stock of basic physical forces which conserve energy” (Ibid., 250). This implies the causal completeness of the physical, since all lower-level physical forces are sufficient and complete for any effect.

Against the argument from fundamental forces, Hendry argues that

“actual quantum-mechanical explanations of chemical structure and bonding seem to presuppose unexplained symmetry-breaking, and so do, in effect, employ configurational [i.e., higher-level] forces. This undercuts any empirical support they could offer to the completeness of physics with regard to chemical systems, and with it, ontological reduction” (Hendry 2010, 187).

Instead of the CCP, Hendry argues for the plausibility of an alternative principle – the *ubiquity of physics*. According to this principle, physical principles and laws are universally applicable, and “physical principles constrain the motions of particular systems though they may not fully determine them” (Hendry 2010, 188). Ubiquity is compatible with strong emergence and downward causation: “The mere applicability of physical principles to chemical bonding requires only ubiquity, and does not rule out downward causation” (Ibid., 188).

The lesson from the above discussion is that not all physicalists accept the CCP principles as essential to physicalism. To be sure, in my argument for content universalism and the strong emergence of content, I am not concerned with whether content universalism is physicalist/naturalist or non-physicalist/non-naturalist. I am more concerned about the irreducibility (strong emergence) of content.

8.2.5. Other Views on the Strong Emergence of Intentionality

In the emergence literature, the most common candidates offered as emergent entities are mental properties, and even among them consciousness (Chalmers 2006b, 246) and libertarian free will (Wilson 2021, 265). Nonetheless, there have been arguments for the emergence of intentionality as well. These other arguments therefore corroborate my own argument for the strong emergence of content.

One example is Crane (2013). He writes: “...I deny that mental representation (intentionality) is reducible to anything else, or definable in other terms, or can be wholly elucidated (for example) in terms of simple causal notions like information or indication....[M]ental representation [is] basic...” (Crane 2013, 115). To recall from chapter 6, Crane explains representation in terms of intentional mode and content. Thus, when he claims that representation is basic, his intention is not to claim that it is primitive and unanalyzable. Rather, the implication is that mental representation is strongly emergent.

Similarly, Baker argues for the emergence of intentionality, but outside the representationalist framework. According to her view, which she calls “Practical Realism” (Baker 1995, 6) intentional states such as beliefs are not brain states, but they are irreducible facts about persons (Ibid., 153). Taking beliefs as the paradigmatic intentional state, Baker writes:

“Whether a person *S* has a particular belief...is determined by what *S* does, says, and thinks, and what *S would* do, say, and think in various circumstances, where “what *S would* do” may itself be specified intentionally. So, whether *S* believes that *p* is true depends on there being relevant counterfactuals true of *S*” (Baker 1995, 154).

Thus, intentional states do not consist in standing in a propositional attitude to a proposition, but in having practical implications which are described by counterfactuals. She further clarifies that believing need not manifest in overt behaviour, but only that it makes some a difference, even if it is to other intentional facts about the person. Therefore, “Practical Realism should be considered not a form of behaviorism but a form of radical relationism” (Baker 1995, 156).

The practical implications of mental states implies that they have causal powers of action in various counterfactual scenarios. Expanding on this point, she writes elsewhere that intentional properties are strongly emergent – they “add to the basic ontology and introduce new causal powers but do not introduce any immateriality into the world” (Baker 2019, 208). In order to

explain how strongly emergent properties, in general, emerge, Baker posits a constitution account, the details of which were given in chapter 1 (§1.4.9). As she puts it, “[t]he fundamental idea of constitution is this: when a thing of one primary kind is in certain circumstances, a thing of another primary kind – a new thing, with new causal powers – comes to exist” (Baker 2007, 32). Baker holds that the constitution-emergence account is not incompatible with ontological naturalism/physicalism, and therefore pitches it as a version of “non-reductive physicalism” (Baker 2009, 120).

8.3. Phenomenological Particularity of Content

Content universalism is a version of content generalism, the thesis that content is constituted by general, as opposed to particular, entities. Every content-constituting entity/condition is a general entity. The universals that qualify the target and permeate in the subject are paradigm general entities – they can be literally shared between particulars. Naturalistic relations are general in the sense that the same kind of relations may obtain between different relata; biological functions are general since they contribute to fitness in a variety of circumstances (within a broad ecological habitat). Finally, the instantiation relation – whether qualification- or permeation-instantiation – is general since it is a primitive relation that relates multiple relata. Thus, immanent emergent content is general content. General content is also known as existentially quantified content (McGinn 1982) or “existential content” (Tye 2009b, 544).

Existential contents are called ‘existential’ due to the notion of existential quantification in symbolic logic. The existential quantifier ‘ $\exists x (Px)$ ’ is interpreted as ‘there is some x , such that x is P ’. Existential contents, therefore, present a target region as populated by objects of a certain size, shape, colour and involved in certain relations with the subject, but without presenting their particularity. As McGinn writes, “...we are to say that a given experience is as of *a* book that is brown, thick, and has the words ‘The Bible’ inscribed upon it; we are not to say, when giving the content of the experience, *which* book it is that is seen” (McGinn 1982, 39).

However, many of our perceptual experiences are as of *particular* objects and their particular properties. What accounts for the particularity of our perceptual experiences? More precisely, what accounts for the particularity of perceptual content?

It should be noted that the particularity that is compatible with content universalism is phenomenological particularity, rather than relational-constitutional particularity (chapter 2, §2.2.5.1). Although content is not constituted by any particular entity, it phenomenologically seems to us that particular entities are represented. In fact, non-reductive representationalism, in general, has this implication. For example, as discussed in chapter 6, Crane argues that “[s]ingularity [i.e., particularity] is a matter of the psychological or phenomenological role of the thought” (Crane 2013, 147). His non-relational representationalist view explains particularity in terms of intrinsic features of the representation itself. These features are the aims of the thinker and the information about the intentional object carried by the representation – both of which account for the psychological/phenomenological role of the thought episode. However, I argued that Crane’s proposals are ultimately inadequate to account for content particularity.

Instead, I submit that phenomenological particularity is plausibly accounted for by *attention*. This suggestion was introduced in chapter 5 (§5.2.2) in the argument against content particularism. Multi-Object Tracking (MOT) studies are best explained by positing a *post*-perceptual mechanism that individuates/particularizes the objects that are constituents of presentational content. This post-perceptual mechanism is attention. The particularizing role of attention is bolstered by independent studies and theories about the nature of attention.

8.3.1. Attention

Attention is a cognitive process of *selection* of features of perceptual states for the purpose of performing a further function – such as filtering, binding, etc. Significantly, the “units of attentional selection” (Scholl 2001, 1) are the features of perceptual content, not features of the external target. In other words, attention does not select from the perceived target but from the way in which the target is represented.

Rensink gives the following explication of attention:

“an attentional process [is] one that is *contingently selective*, with that selectivity controlled via *global considerations* (e.g., tracking a particular person of interest). From this perspective, “attention” is more an adjective than a noun. Any globally controlled process of limited capacity – such as binding visual features, or placing them into vSTM [visual short term memory] – would be “attentional,” because limited capacity implies selectivity of one form or other” (Rensink 2015, 9).

Specific kinds or definitions of attention then depend on delimiting the ‘further function’ of attentional selection. For instance, Kentridge and Brogaard argue that attention is “selection for precision” of perceptual experiences (Kentridge and Brogaard 2017, 141). In contrast, Wu holds that attention is “selection that guides task performance” (Wu 2017, 149). Rensink gives a “function-centred taxonomy” (Rensink 2013, 98) of attention into five kinds – sampling, filtering, binding, holding and indexing. Of particular relevance to particularity are the functions of holding and indexing (more on these below).

8.3.1.1. Attention is neither representational nor perceptual:

A key point to note is that attention is *not a representational process*. Attending to X does not involve representing X in virtue of having an attentional content; it only involves selecting it. Pylyshyn (2001) emphasizes this point in the context of his visual indexing theory, as was noted in Chapter 5 (§5.6.2.1.2). To recall, Pylyshyn argues that attention is a “preconceptual selection” of pre-attentive primitive objects, for the purpose of indexing or individuating. By ‘conceptual’ he means ‘representational’, whether or not the representation is done under concepts as standardly understood. In other words, attention is non-representational selection. Thus, Pylyshyn writes that attentional indexing is a “direct connection between...a visual representation and certain token elements in the visual field, a connection that is unmediated by an encoding of properties of the elements in question” (Pylyshyn 2001, 128).

Also, attention is *not a perceptual process*, although it is closely associated with the latter. For one thing, perception is not necessary for attention. This is because there could be non-perceptual attention as well; for example, we can attend to our thoughts, memories, etc. For another, the question of whether attention is necessary for perception is a hotly debated one, and there are forceful arguments to the effect that attention is not necessary for perception (Wu 2017). More precisely, it is argued that attention is not necessary for phenomenally conscious perception. For instance, Block argues, on the basis of empirical evidence (Block 2007a; 2007b), that phenomenal consciousness “overflows accessibility” (Block 2007a, 487). The overflow consists in the fact that “...the capacity of the phenomenal system is higher than the capacity of the cognitive access system that underlies reportability of phenomenal states” (Block 2007b, 530). A mental state is accessible (or access-conscious) “if, in virtue of one’s having the state, a representation of its content is (1)...poised for use as a premise in reasoning, (2) poised for rational control of action, and (3) poised for rational control of speech” (Block 1995, 231). And it is widely agreed that attention is necessary for access (Campbell 2015, 589-

590). Thus, if phenomenal consciousness overflows access, it also overflows attentional selection. The conclusion of the overflow argument is that “...the cognitive system underlying the reporting of phenomenology is distinct from the system underlying phenomenology” (Block 2007b, 531). The implication is that attentional mechanisms are distinct from perceptual mechanisms.

The implication of the mutual non-necessity of attention and perception is that perception and attention are independent mental processes. However, there is near consensus that attention *modulates* perception (Carrasco et al. 2004, 312). This in turn implies that attention, when it is associated with perception, is a *post*-perceptual mechanism.

8.3.2. Attention and Particularity

Most theories of attention hold that the attention selects from perceptual content that is pre-attentively structured into rudimentary objects. In other words, the *units* of attentional selection are “pre-attentive” objects (Cavanagh et al. 2023, 1649) that are part of the content of perceptual representations. Attention is not directed directly at the external targets of representation.

Pre-attentive objects are rudimentary or primitive objects that are synthesized into perceived objects by attention. Pre-attentive objects are known by several terms in the literature, according to the role they play in different theories of attention-modulated perception – “preattentive object files” according to Wolfe and Bennett (1997, 26), “primitive objects” in Pylyshyn’s (2001, 144) visual indexing theory for visual perception, and “proto-objects” in Rensink’s (2000, 20) triadic architectural theory of perception.

The idea is that once these rudimentary objects are selected by attention, they become the stable and persisting objects that are present to the mind. The post-attended objects are also referred to as “perceptual objects” (Kubovy and Van Valkenburg 2001, 102) – objects *in* perception (Cavanagh et al. 2023, 1649) – to distinguish them from the physical objects that are the targets of perceptual representation (objects *of* perception) and the pre-attentive proto-objects. A defining characteristic of perceptual objects (as opposed to pre-attentive objects) is the Gestalt psychological characteristic of *figure-ground organization* (Wagemans et al. 2012, 1194). Kubovy and Van Valkenburg explain it as follows:

“A *perceptual object* is that which is susceptible to figure-ground segregation...Early [perceptual] processing produces elements that require grouping. Grouping occurs following

the principles described by the Gestalt psychologists...; it produces Gestalts, or perceptual organizations, which are also putative perceptual objects. Attention selects one putative object (or a small set of them) to become figure...and relegates all other information to ground...The putative objects that become figure are perceptual objects, whereas the ground remains undifferentiated information” (Kubovy and Van Valkenburg 2001, 102).

A second important characteristic of perceptual objects is their spatiotemporal *stability* or *coherence* (Rensink 2013, 106) as opposed to the volatility of the pre-attentive objects. A third defining characteristic of perceptual objects is their individuality or *particularity* (Rensink 2000, 23). Of Rensink’s five attentional kinds, attentional holding and attentional indexing are pertinent to particularity.

Attentional *holding* is the selection of pre-attentive objects to create a representation as of a stable object. Attention “holds”, metaphorically speaking, the volatile pre-attentive objects into a stable perceptual object. Attentional holding ensures the spatiotemporal continuity of the perceptual object. “Such continuity is achieved via the idea of *coherent* representation. Here, properties are linked not only across space but also across time, so as to refer to a single persisting object” (Rensink 2013, 106).

Attentional *indexing* or *individuating* (Rensink 2015, 11) is the selection of pre-attentive objects to create not merely coherent perceptual objects, but also individual or *particular* objects (Rensink 2013, 108). That is, attentional indexing is the process that accounts for the phenomenological particularity of perceived objects. Attentional indexing is a necessary condition for the tracking of objects (Pylyshyn 2001; Rensink 2013, 108). As discussed in chapter 5 (§5.2.2), a theory of visual perception that gives a key role to attentional indexing is Pylyshyn’s visual indexing theory (Pylyshyn 2001). To recall, visual indexes or FINSTs (fingers of instantiations) are demonstrative labels assigned by attentional selection to pre-attentive objects (termed “primitive objects” in the theory). It is these visual indexes that account for the individuality or particularity of objectified content, and consequently enable their tracking.

In sum, attention selects from the presentational content and in the process supplies it with particularity. This is corroborated by phenomenological evidence. An experienced driver who takes the same route back home has perceptual experiences of the areas that he drives through without attending to them. His experience does not present him with particular objects. It is only when he attends to an unusual aspect of his perceptual content – say a road accident – that

he has the experience of particularity. Thus, attention has to work on representational content to generate particularity. Particularity is therefore the product of an intentional-attentional tango.

8.3.3. Rensink's Triadic Architectural Theory of Attention-Modulated Perception

A theory of attention-modulated perception that ties together the notions of general content and attention-induced phenomenological particularity is Rensink's (2000) triadic architecture theory. Focussing on the visual perceptual system, Rensink posits a triadic architecture (Rensink 2000, 34) to the visual system – being composed of three independent but interrelated subsystems:




Figure 8.2: Triadic architecture of perception (Adapted from Rensink 2000, 35)

8.3.3.1. Low-level visual system⁹⁷:

This system is the primary representational stage of perception, where the target of representation is the environmental scene, rather than any individual physical object. The scene is represented in the mode of proto-objects, which form the first layer of presentational content

⁹⁷ Rensink clarifies that “low-level” vision “...refers to those processes concerned with separating out the various physical factors in the scene that give rise to the pattern of luminance intensities in the image. These processes are generally considered to be retinotopic, and carried out by a set of processors operating in parallel across the visual field” (Rensink 2000, fn2, 20).

of perceptual representations. Proto-objects are “...relatively complex assemblies of fragments [of properties] that correspond to localized structures in the world” (Rensink 2000, 22).

Proto-objects have two characteristics. First, they are extremely volatile and there is a “constantly regenerating flux” (Rensink 2000, 20) of proto-objects. Second, proto-objects are a “quick and dirty map of scene structure” (Rensink 2015, 6) as they contain information about line segments, edges, shadows, other features, and even three-dimensional orientation. However, they are not stable, are not differentiated into figure over ground, and are not individuated. In other words, proto-objects are the general content of perceptual representations.

8.3.3.2. The setting system:

The setting system involves the extraction of certain kinds of information from the proto-object content level, which can then guide the direction and shifts of attention. Two kinds of information are extracted (Rensink 2000, 36-37): (i) the *gist* or general meaning of the represented scene. For example, upon entering a room, the visual representation of the scene extracts the gist whether it is an educational gathering, or a religious or a social gathering, etc. (ii) the *layout* or spatial arrangement of the scene. The spatial information extracted is from the allocentric frame of reference – that is, from a fixed world-indexed coordinate system that is independent of the subject’s sense of space (egocentric reference frame). Again, both kinds of information are general (rather than particular) information.

The setting system also involves a repository of gist and layout information called *scene schema*. It could “...include an inventory of objects likely to be present in the scene, along with various aspects of layout...” (Rensink 2000, 37). The scene schema is stored in long term memory and influences the expectations of the subject when perceiving a scene.

8.3.3.3. The attentional system:

The third component is the attentional system which involves the direction of attention towards proto-objects to constitute them into stable visual objects. Rensink writes that attention selects only a small number of proto-objects at a time, usually only one. Once selected, attention acts as a metaphorical hand that holds each proto-object to form “...a stable object, with a much higher degree of coherence over space and time” (Rensink 2000, 20). A coherent object is “...integrated over some extent of space and has continuity over some duration of time” (Rensink 2015, 6).

The stable and coherent object is termed a “visual object” (Rensink 2015, 7) – a specific type of the more general ‘perceptual object’. Associated with each visual object is a “...*nexus*, a single structure containing a summary description of the attended object, for example, its size, overall shape, and dominant colour” (Rensink 2000, 23). The nexus of information about the visual object also contains information that is relevant to individuate and particularize the objects attended to. Moreover, it also contains information about the location of the proto-objects. This information aids in the tracking of attended visual objects. Rensink therefore compares the location information to Pylyshyn’s visual index, and the attentional links with proto-objects to Pylyshyn’s FINSTs (fingers of instantiations) (Ibid., 27). Thus, it is the attention-induced nexus that particularizes the stabilized visual objects and enables their tracking.

In sum, Rensink’s theory has it that the primary perceptual content is structured presentationally as proto-objects which is general content. Some parts of it remains general – becoming the gist and the layout – and some parts of it are particularized by attention.

8.4. Conclusion

In this chapter, I fleshed out further details of Content Universalism. I argued that content (= permeated universal) is a strongly emergent property. It has novel causal powers of causing actions in subjects. I also considered the violation of the causal closure principle that strong emergence entails, and discussed arguments to the effect that the principle is not so sound after all. The fact that content has novel causal powers dovetails with my argument in chapter 3 that content is action-oriented because it consists of action-properties that are attributed to targets. Bringing that insight here, the universals that permeate subjects can be said to permeate as action-oriented content. Given that the content so constituted is general content, I then clarified how content can also be characterized by phenomenological particularity. Based on studies on attention in cognitive science, I argued attention selects aspects of content and thereby particularizes it.

To be sure, content universalism is only a theory of content – one half (albeit the primary half) of a representationalist theory. The second half – the theory of representational application (i.e., directedness) – will be outlined in the next chapter.

9. PRIME Representationalism

The aim of the present chapter is to develop a theory of application/directedness – the second half of my theory of Pluri-Relational Immanent Emergent Representationalism (PRIMER) – and to highlight the applications of PRIMER to problems of content. In §9.1, I develop PRIMER's theory of application, which I call 'Representational Magnetism'. It posits a directedness relation to targets that I call the Representational External Directedness (RED) relation. §9.2 brings both content universalism (elaborated in the previous two chapters) and representational magnetism to form PRIMER. I note that PRIMER, although nonreductive, is compatible with naturalism. A notable feature of PRIMER is the many relations that it posits. So, in §9.3, I enquire into the externality vs internality status of these relations, especially the RED relation. I argue that the RED relation is an external relation and therefore irreducible. Finally, in §9.4, I discuss the relevance of PRIMER to solving the distality problem and the problem of non-veridical perception. I also note its compatibility with strong intentionalism – the thesis that phenomenal character is identical to a certain sort of content.

9.1. Representational Application/Directedness

A theory of application explains how representations are directed at their targets in virtue of having content. Thus, a theory of application is based on a theory of content. In the last two chapters, I developed a novel theory of content – content universalism. According to it, content is the permeated universals in the subject, provided that the same universals qualify the target and there are naturalistic relations to the target for the performance of fitness-contributing biological functions.

I will interpret the directedness of representations as a relation between the representation and the target. This interpretation has been challenged – for example, by Crane (2006) – due to the possibility of hallucinations. It is argued that relations require the existence of both relata (representation and target), but there are no represented targets in hallucinations. Therefore, representation does not involve relations – whether at the content-constituting stage or the application stage. However, I have argued in chapter 2 (§2.2.1) that the represented target is a spatiotemporal region, rather than any individual object, and is present even in cases of hallucinations. Thus, in every case of perception, both the representation and the target are

present, which therefore facilitates the interpretation of representational directedness as a relation. I will call this the *Representational External Directedness* (RED) relation.

The key to developing a theory of directedness is to note that the universals being instantiated (in different manners) in the target and the subject are one and the same. They are numerically identical. I submit that it is the numerical identity of the universals permeating in the subject and qualifying the target that grounds the RED relation between the perceptual representations and the target.

The RED relation may be understood on the analogy of electromagnetic attraction between two particles having the same amount but opposite polarities of electric charge. For example, an electron has an electric charge of $-1e$ and a proton has an electric charge of $+1e$. This can be interpreted as both particles instantiating the same property – electric charge – in different manners or polarities. The electric charges generate an electromagnetic field around the particles (Schumm 2004, 51), and due to the opposite polarities of the charges, the two particles attract each other.




Figure 9.1: Analogy of electromagnetism

The electromagnetic attraction happens by the exchange of virtual photons, the quanta or particles of the electromagnetic field. Thus, Schumm writes that “[f]rom the point of view of quantum field theory, the electromagnetic force is a consequence of the exchange of virtual photons – the quanta of the electromagnetic field – between electrically charged objects” (Schumm 2004, 57). If electromagnetic force (attraction and repulsion) is interpreted as a relation between two particles, then the relation of attraction is, in a way, actualized by the exchange of virtual photons.

Perceptual representation can be understood along the lines of electromagnetic attraction. The same universal (cf. electric charge) is instantiated in different manners in the target – qualification (cf. negative charge) – and the subject – permeation (cf. positive charge). This difference in the instantiated property-instances creates an asymmetry which grounds the RED relation (cf. electromagnetic attraction) between the representation and the target. It should be noted that although the RED relation is dependent on the subjectively permeated universal and the objectively qualified universal, its relata are the representation and the target.

To be sure, the analogy of representation with electromagnetic attraction is only metaphorical. But both phenomena contain the core idea that the same universal instantiated differently creates an asymmetry which gives rise to a relation. Given the parallels with electromagnetism, I will call my theory of representational directedness ‘*Representational Magnetism*’. Another reason behind giving it this label is that the term connects the theory of representational directedness with the view of ‘reference magnetism’, which is a hypothesis in metasemantics (Lewis 1983, Sider 2011) and which, as Chalmers (2012) has hinted, may have applications for intentionality.

9.1.1. Reference Magnetism

Reference magnetism is primarily a metasemantic theory in linguistics about how linguistic expressions refer to entities in the world. It holds that the referential power of linguistic expressions is determined partly by their use in a language community and partly by certain metaphysically privileged or “elite properties” (Dunaway 2020, 91) in the world. The intuitive idea is that the metaphysically elite properties exert a sort of magnetic attraction on linguistic expressions.

Referential magnetism was first pitched (without endorsement) by Merrill (1980) and later popularized by Lewis (1983). Lewis argues that the metaphysically elite properties that partly determine reference are *natural properties*: “[r]eference consists in part of what we do in language or thought when we refer, but in part it consists in eligibility of the referent. And this eligibility to be referred to is a matter of natural properties” (Lewis 1983, 371)⁹⁸. Similarly, Dunaway suggests that the elite properties are unified (i.e., non-gerrymandered) properties that play certain theoretical roles. In particular, an elite property is one which “...confers similarity

⁹⁸ Naturalness, for Lewis, is a matter of degree. The *perfectly* natural properties are the sparse and fundamental physical properties. Properties similar to fundamental properties are natural in proportion to their similarity.

on its bearers, contributes to the law-like character of certain generalizations, and makes some properties projectable” (Dunaway 2020, 98). To be projectable is to be usable in epistemological inductions.

Chalmers (2012) has broached the suggestion that natural properties could have a reference-magnetic role to play in the grounding of intentionality as well. Thus, he hints that a complete theory of intentionality is one in which “...intentionality is grounded partly in acquaintance and/or consciousness (or whatever grounds them), partly in inferential role, and partly in norms or naturalness” (Chalmers 2012, 10). However, Chalmers does not develop the reference-magnetic role of natural properties.

Representational magnetism – my theory of directedness – may serve as a specific application of the reference magnetism theory in the case of representation/intentionality. It has the advantage that it implicates properties in *both* the target (property-instances = qualified universals) and the subject (content = permeated universals), instead of merely on the side of the target. This gives a better explanation of the magnetic pull exerted by the target.

9.2. PRIMER: Pluri-Relational Immanent Emergent Representationalism

This completes the development of my representationalist theory. It contains two components: (1) Content Universalism, which is the theory of content; and (2) Representational Magnetism, which is the theory of directedness.

The three main ontological features of my view are:

- (i) Relations: Content Universalism invokes naturalistic relations and the primitive instantiation relation as necessary conditions for content; and Representational Magnetism posits the representational external directedness (RED) relation.
- (ii) Content is subjectively immanent (more specifically, permeated) universals: The universals which are permeation-instantiated – and, therefore, immanent – in the subject is the perceptual content.
- (iii) Emergence: Content is a strongly emergent property. Further, as I will argue in the next section, the RED relation is weakly emergent.

Therefore, I will call my representationalism theory ‘*Pluri-Relational Immanent Emergent Representationalism* (PRIMER). In sum, PRIMER holds that content is an emergent property

that is identical to immanent (permeated) universals, which is constituted by naturalistic relations between a subject and a target, and by the permeation-instantiation of target universals in the subject; the strongly emergent content and property-instances of the target ground the weakly emergent representational directedness relation between a perceptual representation and the target.

PRIMER is a nonreductive theory of perceptual representation. This is because content can be neither functionally nor determinatively reduced to its constituting base; in fact, content is strongly emergent. Despite being nonreductive, it can qualify as a naturalistic theory. In this context, it will be instructive to distinguish, following Armstrong (1997), between two distinct ontological theses – naturalism and physicalism⁹⁹. Naturalism is “...the contention that the world, the totality of entities, is nothing more than the spacetime system” (Armstrong 1997, 5). Defining physicalism, on the other hand, is a more vexed matter. It is commonly agreed that physicalism is the thesis that everything that exists is physical or supervenes on the physical; and ‘physical’ means whatever is posited by the (completed) science of physics (Stoljar 2024). Armstrong defines physicalism as the thesis that

“(1) all fundamental universals, whether properties or relations, are those studied by physics...;
(2) all fundamental laws are connections holding between these fundamental universals and other laws are no more than the fundamental laws operating under specific boundary conditions” (Armstrong 1997, 6).

It is clear that PRIMER is compatible with naturalism as defined by Armstrong. The universals posited by PRIMER are intrinsically concrete (i.e., spatiotemporal), and so are the other elements, including the permeation-instantiation relation. PRIMER’s compatibility with physicalism might seem to be shaky due to the features of strongly emergent content and the permeation manner of instantiation. First, strong emergence has been traditionally considered to be anti-physicalist since it violates the causal closure of the physical (CCP) principle. However, as discussed in chapter 8, (§8.2.4), many philosophers of science accept strong emergence in the natural science and argue for its compatibility with physicalism. Although strong emergence violates the causal closure of the physical principle, the principle itself is arguably not integral to physicalism, as Hendry (2010) argues. Thus, physicalists can embrace strong emergence without qualms. Secondly, the novel manner of instantiation that I posit –

⁹⁹ Naturalism and physicalism are sometimes interpreted as referring to the same view. For example, Papineau (2023) discusses them interchangeably.

permeation-instantiation – could possibly be part of a future completed physics. There is no *a priori* reason to believe that permeation-instantiation is anti-physicalist.

Taking a firm stance on whether PRIMER is compatible with physicalism is not essential to my thesis; so, I will remain neutral towards the question. However, what is clear is that PRIMER is a nonreductive *naturalistic* representationalist theory.

9.3. Ontological Kind of Relations in PRIMER

A significant distinction in the ontology of relations is made between two kinds of relations – external and internal relations. In the context of this distinction, it might be asked: what is the externality/internality status of the relations in PRIMER? Paolini Paoletti (2025) notes that there are several (often incompatible) ways of characterizing the distinction between internal and external relations (Paolini Paoletti 2025, 54). One prominent way of characterising the distinction is in modal terms: Internal relations are those that supervene on (Lewis 1986) or are necessitated by (Armstrong 1997) the intrinsic natures of the relata. Thus, Armstrong writes that “[a]n internal relation is one where the existence of the [relata] entails the existence of the relation” (Armstrong 1997, 87). External relations are relations that are not internal in this sense.

Another prominent way of marking the distinction is in terms of reducibility. Internal relations are held to be reducible to their relata; more precisely, they are reducible to the intrinsic properties of their relata. Thus, Heil writes that internal relations are “no addition to being” (Heil 2021, 8). External relations, in contrast, are not reducible to their relata; they are “something *in addition* to whatever they relate” (Ibid., 8). Paolini Paoletti argues that the different ways of characterizing the internal-external distinction ultimately boil down to the reducibility question. Thus, “what matters for the distinction between internal and external relations is the reducibility of the former” (Paolini Paoletti 2025, 54).

9.3.1. Defining Internal and External Relations

As discussed in chapter 1 (§1.4.8), there are different routes to reduction – elimination, identity, functional reduction and determinative reduction. Paolini Paoletti uses determinative reduction in his discussion of internality of relations, and I will follow him in this. To recall, P determines

M just in case M is ontologically dependent on, is grounded by, and is metaphysically necessitated by P.

A simplified version of Paolini Paoletti's definition of an internal relation of is as follows:

(Simplified Paoletti Internal Relation– SPIR) R is an internal relation iff it is necessarily the case that: (i) the holding of R is *determined* by some other entity, and (ii) the holding of R does not depend on any relation in its transitive dependence chain that is not *determined* by non-relational entities¹⁰⁰.

Condition (i) may be called the *immediate reducibility* requirement. Paolini Paoletti extends the scope of the immediate reduction base by including entities apart from the relata. Thus, he says: “[w]hen I write that “R is determined by...something else”, I mean that it is determined by...something else for its holding between its relata” (Paolini Paoletti 2025, 75). In other words, internal relations may be determined by entities other than the intrinsic properties of their relata. For example, the tracking relation that naturalistic representationalism posits is determined by its relata (the vehicle and environmental target), naturalistic relations and teleological functions.

However, the immediate reducibility of a relation is not sufficient to ensure its ontological status as an internal relation. This is because even if a relation R is determined by other entities (and seems *prima facie* reducible and internal), the determining base could include relations or non-relational entities that, in turn, depend on other relations. In general, since dependence is transitive, there could be a chain of dependence with a relation at any level. Now, if these other relations in the transitive dependence chain are not *determined* by non-relations – that is, if the dependee relations are irreducible – the original relation R would not count as fundamentally internal, since relations have not been excised from our ontology altogether.

Therefore, Paolini Paoletti adds condition (ii) to his definition of internal relations, which effectively says that the relation R must not depend on any irreducible relation at any level of its transitive dependence chain. If there are any relations at any level of the dependence chain, then they must be determined by non-relations, and hence be reducible to non-relations themselves. This will ensure that there is no “irreducible relationality in the universe” (Ibid., 54). Condition (ii) may, therefore, be called the *ultimate reducibility* criterion.

¹⁰⁰ Paolini Paoletti uses technical notions such as “dependence*” (Ibid., 51) in his actual definition, which are not essential to my discussion. The underlying idea, as he emphasizes, is reducibility.

According to Paolini Paoletti, internal relations must satisfy both conditions (i) and (ii) in the definition above. In other words, internal relations are *both* immediately and ultimately reducible to non-relational entities. Examples of internal relations are ‘being taller than’, and ‘being older than’. Conversely, external relations are those that satisfy at most one of the conditions. Depending on how they satisfy the two conditions, Paolini Paoletti lists three types of external relations (Paolini Paoletti 2025, 62-66):

External Relation	Condition (i): Determined by other entities	Condition (ii): Dependence chain has no irreducible relations	Explanation
Type 1	✓	✗	Determined by other entities, but the dependence chain includes irreducible relations.
Type 2	✗	✓	Not determined by other entities, and not dependent on any irreducible relation.
Type 3	✗	✗	Not determined by other entities, but depends on other irreducible relations.

Paolini Paoletti mentions a fourth type of external relation, which he calls *essential relation*. A relation R is an essential relation if it “constitutes the essence of further entities” (Paolini Paoletti 2025, 64). That is, an essential relation is one that constitutes part of the essence of at least one of its relata. The example he cites is that of origin essentialism. For example, it is part of my essence that I am the son of my mother. In other words, my essence is partly constituted by my relation of ‘being a son of’ to my mother. Arguing for the externality of such a relation, Paolini Paoletti writes: “[i]f R is an essential relation, it is not determined by [and, therefore, not reducible to] the entities it is essential to. Therefore, it is *prima facie* external” (Ibid., 124).

9.3.2. Externality/Irreducibility of Relations in PRIMER

In this section, I want to argue for the externality of relations posited in PRIMER, whether as part of the theory of content – Content Universalism – or the theory of application – Representational Magnetism. In particular, I will argue that the following are external: (a) the content-constituting instantiation relation between concrete universals, on the one hand, and target and subject, on the other; (b) content-constituting causation; and (iii) RED relation.

9.3.2.1. Content-constituting relations:

Content-constituting relations fall into two categories: permeation-instantiation, and naturalistic relations. The instantiation relation – whether permeation-instantiation or qualification-instantiation – is a Type-2 external relation, because it is a fundamental and primitive relation. It violates condition (i) of SPIR since it is not determined by any entity. It satisfies (ii) because it is not dependent on any irreducible relation, in virtue of not depending on any entity, relation or non-relation. The permeation-instantiation relation is also plausibly an *essential relation*. This is because it partly constitutes the essence of perceptual content. It is, after all, part of origin-essence of content to be generated by the permeation of universals.

The status of the other content-constituting relation – naturalistic relations – is uncertain. If the naturalistic relation in question is either correlation or structural correspondence, then it is internal, since both correlation and structural correspondence satisfy both conditions of SPIR – they both determined by their relata, and are not dependent on irreducible relations. However, if the naturalistic relation in question is causation, then the question turns on whether causation is external or internal.

Causation is commonly regarded as an external relation (Yates 2016, 156; Paolini Paoletti 2025, 258). Paoloni Paoletti argues that it is either a Type-2 or a Type-3 external relation (Paolini Paoletti 2025, 255). This is because it fails to satisfy condition (i) the SPIR definition – it is not determined by its relata – the cause and effect – and, therefore, cannot be reduced to the intrinsic properties of its relata. Whether it violates condition (ii) (Type-3) or not (Type-2) depends on whether the cause depends or not on any other irreducible relation.

However, many neo-Aristotelians who posit intrinsic causal powers argue that causation is an internal relation (Lowe 2016, 107; Heil 2021, 41-42). They hold that causation is the mutual manifestation of intrinsic causal powers of interacting objects/properties. Thus, Heil writes: “If a property’s identity is bound up with dispositions it would confer on its possessors, and if causal relations are the manifesting of powers, then causal relations would appear to be a species of internal relation” (Heil 2012, 148). However, some dispositional essentialists – such as Anjum and Mumford (2018) – hold that manifestations of causal powers are not necessitated by the powers, but only tended towards. This would mean that causal manifestation is not an internal relation. Although an argument in favour of the externality of causation would be take me too far afield, I will go with the majority opinion that causation is indeed an external relation.

9.3.2.2. The RED relation:

As for the Representational External Directedness (RED) relation, it is plausibly a Type-3 external relation – it violates both conditions of the SPIR definition. Let us start with condition (ii). The RED relation violates condition (ii) since its transitive dependence chain has the irreducible relation of instantiation – it depends on content which itself depends on the permeation-instantiation relation holding between the subject and universals.

The RED relation also violates condition (i) because it is not *determined* by other entities, even though it is *dependent* on its dependence base, which includes content, the qualifying universal in the target, and its relata (the representation and the target). The RED relation is not determined by any of its dependees because, as will be discussed below, the RED relation has distinct causal efficacy vis-à-vis its dependence base; and having distinct causal efficacy is the hallmark of irreducibility and emergence.

The RED relation's causal powers consists in the fact that the holding of the RED relation between the subject and the target can cause changes to the target over an evolutionary significant period of time. This is indicated by studies in sensory ecology, a sub-discipline of biology, which studies "...how animals acquire, process, and use information in their lives, and more recently the role of sensory systems in evolutionary change" (Stevens 2013, 4). This amounts to an attempt "...to understand what they perceive in their environments and how that is going to affect their interactions with that environment" (Burnett 2011, 75).

Studies in sensory ecology indicate that sensory perception not only evolves in response to the properties of represented targets, but also that the targets co-evolve due to the pressures of perception and interaction with perceiving animals. For example, not only does colour vision evolve in response to surface reflectance by objects, but also there is reciprocal co-evolution of surface reflectance properties in plants and animals. Chittka and Menzel (1992), for instance, have argued that flowers developed coloration properties (i.e., surface reflectance properties) in response to the pre-existing capacity of pollinators like bees (which belong to the order Hymenoptera) to see in certain wavelengths. Thus, they write that "...flower colours did not evolve before pollinator colour vision, but it is most likely that Hymenoptera possessed important components of their present colour vision systems prior to the first appearance of a colourful flower" (Chittka and Menzel 1992, 179). Summarizing their findings, Yong poetically writes: "Their [bees'] style of trichromacy evolved hundreds of millions of years

before the first flowers appeared, so the latter must have evolved to suit the former. Flowers evolved colors that ideally tickle insect eyes” (Yong 2022, 115).

Studies like the above indicate that perceptual representation – and, therefore, the RED relation to the targets – has the causal power to bring changes in what properties are instantiated in the targets over an evolutionary pertinent timescale. The causal powers of the RED relation then are analogous to the power of water to erode rocks over a long timescale rather than the power of a hammer to shatter them instantly.

The question then is: is the RED relation a strongly emergent or a weakly emergent property? In the previous chapter, the distinction between weak and strong emergence was marked on the basis of the strength of causal autonomy of an emergent property. A weakly emergent property has causal powers that are a subset of the causal powers of its emergence base. However, since this subset is more proportional to an effect than the emergence base’s causal powers, the weakly emergent property has distinct causal efficacy with respect to its base. A strongly emergent property has novel causal powers not had by its base, and therefore has stronger causal autonomy. It was argued that content is a strongly emergent property, due to its novel action-oriented causal powers.

It is plausible that the RED relation is a weakly emergent (dyadic) property. Its causal power is plausibly inherited from and, therefore, a subset of the causal powers of content, because content is the main dependee entity for the RED relation – the RED relation, after all, holds in virtue of content. Now the causal powers of content are mainly geared towards causing actions in perceiving subjects, actions which are directed at environmental targets. The RED relation therefore inherits a subset of the externally directed causal powers of content. The difference between the causal powers of content and the RED relation is that while content causes organisms to act in on targets, the RED relation causes changes *in* targets. The latter causal power is evidenced by the reciprocal co-evolution of surface reflectance properties in plants and animals, as argued in the study by Chittka and Menzel (1992). Thus, the subset of causal powers possessed by the RED relation are more proportional to the effect (change in targets) than the powers of content. Therefore, we may conclude that the RED relation is weakly emergent, and irreducible to its emergence base.

9.4. Applications of PRIMER

It is often said that the best argument for a theory is its applications. In this section, I will discuss some potential applications of Pluri-Relational Immanent Emergent Representationalism (PRIMER), particularly the thesis of content universalism, to various problems and issues. To be sure, these are suggestions towards solutions, and represent areas of future research and development for PRIMER. I will discuss its applications to the problems of distality and non-veridical perception, and the issue of intentionalism.

9.4.1. Solution to the Distality Problem

As argued in chapter 2, the distality problem is the central problem of perceptual representation. The problem is of explaining how representations can have distal contents – contents that present distal targets to the exclusion of proximal intermediaries of the naturalistic relational chain.

Content universalism presents a straightforward solution to the distality problem. A necessary condition for content is that the same universals are qualification-instantiated and permeation-instantiated in the target and the subject, respectively. Given this, representations have suitably *distal* content because the same universals that qualify suitably *distal* targets permeate in the subject. None of the universals of the proximal chain permeate in the subject and therefore none contribute to the constitution of content.

For example, suppose a subject visually and olfactorily represents a rose:




Figure 9.2: Visual and Olfactory representation

In visual representation, the universals that qualify the region of the rose – say Rosehood (a conjunctive universal), Surface Reflectanceness₁, and Surface Reflectanceness₂ – permeate in the subject as the contents <Rose>, <Red> and <Green> respectively. The universals that qualify the proximal intermediate region which contains the naturalistic relational chain do *not* permeate in the subject. Therefore, the content of the visual representation is as of the distal rose with its colours. In olfactory representation, however, the target is the spatiotemporal region *surrounding* the rose, particularly between the rose and the subject. The universals that qualify the chemical odorant molecules in the ambient region permeate in the subject. Therefore, the content of the olfactory representation is as of rosy smell but at the location of the ambient region. This is why we can smell the odorants of objects even when they are not present in the same room – we represent not the object region, but the spatial region around it. Auditory representation is similar to olfactory representation in that the represented target is the region where the sound waves propagate, rather than the region where the objects that produce the sound waves are located (Nudds 2009).

The distality solution offered by PRIMER has advantages over the solutions offered by other non-reductive representationalist theories. In chapter 6 (§6.4), it was argued that the solutions offered by non-reductive representationalism are inadequate. One possible solution is the postulation of a secondary wide content that is dependent on relations to the environment, apart from the primary narrow content. However, given that the relations are transitively linked naturalistic chain relations, the distality problem raises its head again. Con-relational versions might invoke the abstract universals that subjects are acquainted with to account for distality. However, I argued that, since they are abstract and therefore non-spatiotemporal, the spatial concepts of distality and proximity do not apply to them. PRIMER's solution, in contrast, relies on suitably distal target-universals permeating in the subject. The concepts of 'distality' and 'proximity' can be applied to the target-universals because they are concrete universals.

At this juncture, I want to anticipate a potential objection or, rather, a challenge: What determines which universals permeate in the subject in perceptual representation? For instance, in visual perception, the target is the entire spatiotemporal region before the subject. Given this, what determines that the universals qualifying a specific portion of the target – the rose region rather than the surrounding regions – permeate in the subject.

In reply, I think it is plausible that the biological functions of the perceiving subject determine which target-universals are salient for permeation. The function of visual perception is to act

on the object region – by grasping, attacking, and so on. Therefore, the universals of the object region permeate. The function of olfactory perception, on the other hand, is plausibly to act with respect to the *surrounding* area – by approaching, avoiding, etc. – and so the universals of the ambient region permeate. To be sure, I am not claiming that the permeation of the universals depends on the functions. The permeation and functions are both co-equal necessary conditions for content. Rather, what I am claiming is that which universals permeate is determined by the biological functions.

9.4.2. Solution to the Problem of Non-Veridical Perception

My main concern in the thesis has been the solution to the distality problem. However, I believe that PRIMER, particularly content universalism, can also provide a solution to the important problem of misperception or non-veridical perceptions. A detailed exposition of the solution is beyond the scope of my thesis. So, I will merely hint towards a solution.

The three content-constituting elements work in tandem in two stages: first, a subject is naturalistically related to an environmental target region for the performance of functions, which is followed by, second, the permeation of target-universals in the subject. The content-constituting process is a temporally extended process¹⁰¹. If the first stage occurs in the evolutionarily normal way, the same universals qualifying the target permeate in the subject. This is what constitutes veridical perception.

As for non-veridical perception, my hypothesis is that any *defect* in the first stage results in *deviant* universals permeating in the subject. By deviant universals I mean universals which are not qualification-instantiated in the target. It should be recalled from chapter 7 (§7.2.4) that universals are *intrinsically* (i.e., non-immanently) concrete. So, even if they do not qualify the target, it is possible that they could permeate in the subject. These universals are plausibly qualification-instantiated somewhere in the actual world; however, my hypothesis would work even if they were not instantiated anywhere.

What could a defect in the first stage – where the subject is naturalistically related to a target for the performance of functions – consist in? It could plausibly consist of at least two types of defects. First, deviant naturalistic relations between the subject and target; for example, something as banal as the refraction of light when it passes from a water to air, which causes

¹⁰¹ This does not detract from the synchronic or cotermporal existence of the constituters and content after the content has emerged.

the illusion of a bent stick, or even something as extreme as an evil neuroscientist plugging electrodes into our brains to cause neural stimulation. In the latter case, the deviant causal chain would extend from the scientist through the machine to our brains. Second, the defect could be due to a general malfunction of the perceptual system – that is, a failure of biological functions. Such malfunctioning could be due to subjective factors such as illness or consumption of hallucinatory drugs.

The defective first stage would lead to a defect in the second stage – a deviant universal would permeate the subject which does not qualify the target. The content constituted by the deviant permeated universal would be at odds with the universal qualifying the target. This mismatch between content and target is what accounts for error and misrepresentation.

Now, as argued in chapter 3 (§3.5.1), the primary evaluation conditions of perceptual representations are action-success conditions, and veridicality conditions are determined by them. If the action caused by the permeated action-oriented content is successful, the content and the representation can be evaluated as veridical. In the case of non-veridical perception, the action is often unsuccessful; consequently, they can be evaluated as non-veridical.

It should be noted that the target is a spatiotemporal region rather than any object; so, the same general account applies to both illusion and hallucinations. More specifically, in the case of an illusion, a deviant universal permeates *instead of* a normal universal that qualifies the target. For example, suppose we have the illusion of a bent stick, when in reality it is a straight stick in water. The explanation would be this: There is a deviant causal relation between the spatiotemporal region containing the stick-in-water and the perceiver due to refraction of light when travelling from water to air. Due to this, the target-qualifying universal ‘Straightness’ (assuming such a universal exists) does not permeate in the subject. Instead, a different universal ‘Bentness’ permeates, which then becomes the presentational content of the perception. Thus, we have an illusion.

In contrast, in the case of a hallucination, a deviant universal permeates in the subject *in addition to* normal target-qualifying universals. As Macpherson notes, most hallucinations are partial hallucinations – where some features of a region are represented veridically and some are hallucinated (Macpherson 2013, 8)¹⁰². Masrour calls them “embedded hallucinations”

¹⁰² Ali (2018, 614) notes that there are many related bifurcations of hallucinations: impure vs pure; partial vs total; imperfect vs perfect. Then there is a distinction between psychologists’ and philosophers’ hallucinations (Byrne and Manzotti 2022, 332).

(Masrour 2020, 746), where the hallucinated feature is embedded in a veridically perceived scene. The assumption that targets are spatiotemporal regions is in harmony with this fact. For instance, take the Shakesperean example of Macbeth hallucinating a dagger. Macbeth hallucinates the dagger in the courtyard of his castle as he making his way to the King's chamber. He perceives most of the target region (the courtyard) accurately since he is, after all, navigating through it. The hallucinated dagger is an addition to all of this. The explanation for this hallucination is that in addition to the target-qualifying universals permeating in Macbeth, an additional deviant universal – ‘Knifeness’ – permeates in him as well. This is owing to his psychological turmoil (and probably a mental disorder). The permeated deviant universal becomes the content of his hallucination.

The above account of hallucination diverges from the traditional definition of hallucination, which is: “**Trad Hallucination**: you have an experience as of an object and its properties but there is no (worldly) object, and there are no (worldly) properties, that you perceive in virtue of having that experience” (Macpherson and Batty 2016, 265). Instead, PRIMER's account of hallucination bears an affinity with an illusion as traditionally defined: “**Trad Illusion**: you perceive a (worldly) object but you misperceive one or more of its properties” (Macpherson and Batty 2016, 264). PRIMER, after all, says that in hallucination there is a target which is misperceived as being qualified by *additional* universals, which do not really qualify the target. A traditional illusion, in comparison, also involves perceiving a target in a way it is actually not, but with the difference that an object is perceived to have a property *instead* of the one which it actually has.

PRIMER's illusionist account of hallucinations is not entirely maverick though. Similar accounts have been proffered in the literature. For instance, Ali argues for construing hallucinations as kinds of illusions, and he argues for one such view which he calls “the illusionist theory of hallucinations” (Ali 2018, 607)¹⁰³. He argues that on the illusionist theory, “...hallucinations do involve perceptual contact with the surroundings” (Ibid., 626) and therefore it is relational theory.

Masrour (2020) calls the strategy of explaining hallucinations in terms of illusions *Displacing*. “In Displacing, an aspect of experience that is taken to be presenting a non-existent object is redescribed as misrepresenting a property of a perceived external concrete particular. The

¹⁰³ Ali (2018, 607) cites others who endorse similar illusionist theories of hallucinations – Alston 1999, Watzl 2010 and Raleigh 2014.

redescription turns what we regard as a hallucination into an illusion” (Masrour 2020, 743). For example, suppose John is hallucinating a pink elephant in the corner of a room. The situation must *not* be analyzed as ‘John is veridically perceiving a room *and* hallucinating a pink elephant in the corner of the room’. Rather, on the displacing strategy, John non-veridically perceives the room as having a property that it actually does not have – pink-elephantness. Thus, Masrour writes: “...the object of John’s experience is the room. The aspect of John’s experience that corresponds to the pink elephant is re-described as corresponding to a property that he experiences the room as having” (Ibid., 746).

Another similar theory is Byrne and Manzotti’s (2022) Gerrymandered Object Theory of Hallucination (GOTH). According to this, subjects are always aware of concrete physical objects in hallucinations (Byrne and Manzotti 2022, 345), not mind-dependent objects like sense data or mind-independent abstract objects. In some cases, the concrete objects are ordinary objects – for example, when we hallucinate a living friend. However, sometimes the concrete objects may be spatiotemporally scattered objects that the mind amalgamates and perceives as one object (Ibid., 344). Byrne and Manzotti call such spatiotemporally scattered objects “*gerrymandered objects*”, which are “extraordinary scattered objects, some with radically diverse parts, and some with large spatiotemporal gaps” (Byrne and Manzotti 2022, 341). Examples of mundane gerrymandered objects are a *flock* (of birds), a *constellation* (of stars). Byrne and Manzotti claim that there could be extraordinary gerrymandered objects as well such as a flying pig, a trout-turkey, etc.

The Gerrymandered Object Theory of Hallucination (GOTH) holds that “[h]allucinators are always aware either of ordinary objects or of diverse spatiotemporally scattered objects that appear as one” (Byrne and Manzotti 2022, 345). Suppose John hallucinates a flying pig. According to GOTH, the object of the hallucination is a physical pig (not necessarily any particular pig, such as the Empress of Blandings) and physical wings amalgamated together. Thus, they write:

“Despite looking like a pig, the gerrymandered object is not composed of parts that, were they all appropriately glued together, would form a pig, or indeed anything like a pig...[T]he proposal is that the typical object of a pig-hallucination is an amalgam of pigs that, to the hallucinator, looks like a single pig” (Byrne and Manzotti 2022, 347).

Hallucinations are like illusions since they involve misrepresenting actual but scattered objects (the scattered pig and wings) as a single object (flying pig). However, they differ from illusions

to the extent that the “[gerrymandered] object of a hallucination is not currently affecting the subject’s retinas...” (Byrne and Manzotti 2022, 347).

In sum, PRIMER treats all cases of non-veridical perception (traditional illusions and hallucinations) as illusions. PRIMER’s account gets at the heart of misrepresentation and error, which is that the content and target can come apart. As Cummins writes, “...error lives in the gap between target and content, a gap that exists only if targets and contents can vary independently” (Cummins 1996, 7). PRIMER gives just such an account of how content and target can vary independently, in terms of deviant universals constituting content.

A potential objection that may be levelled against PRIMER’s illusionist account of non-veridical perception is that it is inconsistent with representational magnetism – PRIMER’s sub-theory of directedness. To recall, representational magnetism states that the intentional directedness of representations is grounded by the same universals qualifying the target and permeating in the subject. The different instantiations of the same universals act like opposite electromagnetic charges that attract each other. The objection, then, could be this. In the case of a non-veridical perception, the deviant universal (by definition) only permeates the subject without qualifying the target. Thus, the account is inconsistent with representational magnetism.

In reply, I want to emphasize that in non-veridical perception, the target scene is largely veridically perceived, and the non-veridical feature is perceived either instead of (in the case of illusions) or in addition to (in hallucinations) the target. In other words, the non-veridical properties are embedded into the veridical content of representations. The veridical part of the content is due to target-qualifying universals permeating in the subject. In virtue of this, the representation is directed at the target.

The non-veridical part of content is due to deviant universals permeating in the subject. These deviant universals are plausibly qualification-instantiated in spatiotemporal regions that the subject is not naturalistically related to. Or, the deviant universals may be uninstantiated altogether – after all, the universals that PRIMER posits are intrinsically concrete entities (like quantum fields) and so the subject may be related by permeation-instantiation to them even if they are not qualification-instantiated anywhere. Now, supposing the deviant universals qualify spatiotemporal regions elsewhere, i.e., non-locally. Given that they also permeate in the subject, this would imply – by representational magnetism – that the subject’s representations are also directed at these non-local regions. This is an implication I am ready to accept. In sum,

in a non-veridical perception, the subject is directed both at the target – in virtue of the permeation of target-qualifying universals, and at non-local regions – in virtue of the permeation of deviant universals that qualify non-local regions.

9.4.3. Intentionalism

The third and final application of PRIMER that I wish to discuss is its potential account of the phenomenal character of perceptual experiences. The phenomenal character of perceptual experiences is the subjective givenness of phenomenal qualities such as ‘red’, ‘low-pitch sound’, and ‘the smell of vanilla’. Thus, phenomenal character has two aspects – subjective character and qualitative character (Kriegel 2009, 1). Phenomenal character is also known as phenomenal consciousness. As discussed in chapter 2 (§), the account of phenomenal character that is most often combined with representationalism is intentionalism – the thesis that phenomenal character is identical to or supervenes on content.

PRIMER is eminently compatible with intentionalism, especially strong intentionalism. To recall, strong intentionalism holds that phenomenal character is identical to representational content. In this sub-section, I want to suggest a strong intentionalist account under the auspices of PRIMER. According to PRIMER, content is the permeated universals in subjects. PRIMER also accepts Armstrong’s distinction between simple and complex (conjunctive or structural) universals as an assumption. This distinction can ground the distinction between phenomenally non-conscious and conscious perceptual states. Simple universals such as ‘Shapeness’ and ‘Edgeness’, when subjectively permeated, constitute non-phenomenal content. The non-phenomenal content is still presentational content in the sense that it presents a target property, but non-qualitatively (i.e., unconsciously). In contrast, when conjunctive universals permeate in the subject, they constitute phenomenal content, content that is identical to phenomenal character. I use ‘phenomenal content’ in the way defined by Kriegel: “For any experience *E*, the phenomenal character of *E* is one and the same as the phenomenal content *E* carries. Phenomenal character is thus a species of representational content” (Kriegel 2002, 180).

For example, the universal ‘Surface Reflectanceness₁’ is plausibly a conjunctive universal which has other simple universals like ‘Texture₁’, ‘Edgeness₁’,... as conjuncts. ‘Surface Reflectanceness₁’ qualifies a target as the property-instance Surface Reflectance₁, and during perception permeates in the subject as the phenomenal content <Red>. Thus, the phenomenal character of ‘subjective givenness of Red’ is identical to the tokening of the content <Red> in the subject. This implies strong intentionalism. Similarly, a highly conjunctive universal such

as ‘Shapeness^Edgeness^Surface Reflectance-ness^...’ permeates the subjects as the phenomenal content <Object o with properties F, G, H, ...>. There could be a weak intentionalist reading of PRIMER as well, where phenomenal character only supervenes on content. However, this would still leave us with the burden of accounting for the origin of phenomenal character and the nature of its supervenience with inherited universals.

A significant implication of the strong intentionalist account of PRIMER is that phenomenal qualities such as colour and smells are not in the world, as many physicalist representationalists such as Tye (2009, 2021) and Dretske (2003) hold. For instance, Tye (2009) endorses the view that phenomenal character is identical to external phenomenal qualities that a subject represents, and is acquainted with, in perception:

“The phenomenal character of an experience...is out there in the world (or in the body, in the case of bodily sensations)...It is a complex of properties represented by the experience. In being aware of the external qualities, we are aware of phenomenal character...This is why necessarily any change in the external qualities with which we are acquainted in undergoing a given experience generates a change in the phenomenal character of the experience” (Tye 2009, 119).

Tye (2021) labels the above view “property representationalism” (Tye 2021, 58). In contrast, PRIMER holds that phenomenal qualities are properties of the subject, since phenomenal content is a property of the subject. However, they are not the intrinsic non-representational properties that adverbialists (Fish 2021, 85) other naturalist representationalists such as Papineau (2021) posit. For instance, Papineau writes that “...conscious sensory properties are intrinsic qualitative properties of subjects. It is not essential to sensory experiences that they relate subjects to objects or properties beyond themselves” (Papineau 2021, 83)

In contrast to the internalist qualia view above, PRIMER holds that phenomenal content/character is an extrinsic or relational property of the subject – a property which is the permeation-instance of a concrete universal, and it is a representational property – it grounds the external directedness of a representation towards its target. Thus, PRIMER is also in opposition to those nonreductive representationalist views – such as Chalmers’ (2006) and Pautz’s (2021) – which hold that phenomenal character is determined by acquaintance with Platonic abstract universals (chapter 6, §6.2).

The advantage of PRIMER’s strong intentionalist account is that it does not face the content-character mismatch problem. To recall, this problem is faced by those who hold that content is an abstract entity (for eg., a proposition), and abstract entities are incapable of grounding

concrete phenomenal character. However, according to PRIMER, phenomenal content is a concrete property – it is the permeated universal. PRIMER shares this advantage with other concretist intentionalist views such as Crane’s (2009).

9.4.4. Metaphysics of Subjectivity

The final application of PRIMER that I want to consider is its implication on the metaphysics of subjectivity. This implication is highly cursory and points to future research potential. PRIMER posits a special manner of instantiation – the permeation-instantiation – to explain the constitution of content. Permeation-instantiation is applicable only in the case of subjects. Universals permeate in subjects *qua* subjects and, given other content-constituting conditions, the permeated universal is the representational content.

In claiming this, PRIMER is assuming the subjectivity of content. Subjectivity is also known in the literature as the “first-person perspective” or “for-me-ness” or “subjective givenness” (Zahavi and Kriegel 2016) or “subjective character” (Nagel 1974; Nida-Rümelin 2023) of mental states. Subjective character is often associated with phenomenal character/consciousness, with ‘qualitative character’ being the other side of the phenomenal character coin. However, subjectivity is implicated in both phenomenal and non-phenomenal content. Even non-phenomenal perceptual states are from a first-person perspective, although the content is not characterized by phenomenal consciousness. This is because even non-phenomenal states are directed towards their targets, and directedness is always *from* a perspective. Thus, Burge, writing about content in general, says that content is a “perspectival way of representing” (Burge 2010, 38).

Returning to PRIMER, it might be asked: what is it about the subject or subjectivity that grounds permeation-instantiation? There could be three possible answers to this question – one that is compatible with physicalism, another that espouses substance dualism, and a third answer that endorses panpsychism. The physicalist answer could be that subjectivity is ultimately to be explained in purely physical terms. It supervenes on or is grounded in the physical properties of fundamental entities, and when they are complexly arranged into living organisms, subjectivity manifests. The other two answers would comprise non-physicalist answers.

The substance dualist answer is that subjects or selves belong to a distinct ontological category as opposed to material substance. Traditional Cartesian substance dualism holds that subjects

are wholly immaterial substances in the sense that they do not possess any physical property. However, there are other versions of substance dualism as well. Lowe (2008), for instance, espouses ‘Non-Cartesian substance dualism’ (NCSD), which holds that subjects or persons are “*emergent* substances” (Lowe 2008, 118), ontologically dependent yet distinct from bodies, having causal powers that “...*complement* those of our bodies, rather than being either reducible to or in rivalry with them” (Ibid. 118). Furthermore,

“According to NCSD, it is *I* [i.e., person or self], and not my body nor any part of it, who am the bearer of mental properties, just as Descartes maintained. However, unlike Descartes, the advocate of NCSD...may maintain that I possess certain physical properties in virtue of possessing a body that possesses those properties: that, for instance, I have a certain shape and size for this reason” (Lowe 2008, 95).

How could substance dualism be explanatory of the different manners of instantiation in PRIMER? Roughly, the argument could be that each ontological category has a distinct manner of instantiation – qualification for purely material substances, and permeation for subjects *qua* subjects. This proposal would require a lot more detailing though, and is ripe for future research.

A second answer could be from a panpsychist or panprotopsychist metaphysic (Goff et al., 2022). Panpsychism is the view that some kind of mental property is both fundamental and ubiquitous in the universe. Panprotopsychism is the view that precursors of mental properties – which may be called proto-psyche – are fundamental and ubiquitous. Thus, every material entity has the mental (or proto-psyche) property in question. However, not all material aggregates manifest this property. Different theories have different criteria for such manifestation or expression of the mental property. Contemporary panpsychist views hold that phenomenal consciousness (both subjective and qualitative character) is fundamental and ubiquitous. Thus, they are more accurately labelled *panexperientialism*, “the view that conscious experience is fundamental and ubiquitous” (Goff et al. 2022, 2.1). Panpsychism is experiencing a resurgence in contemporary philosophy of mind – see for instance, Brüntrup and Jaskolla 2017¹⁰⁴.

There are several distinctions within the panpsychism literature; an especially relevant one in the context of PRIMER is the distinction between constitutive and non-constitutive or emergent

¹⁰⁴ Even naturalist representationalists such as Tye have recently espoused panpsychist views (Tye 2021, 99).

pan(proto)psychism (Goff et al. 2022; Chalmers 2017). Constitutive panpsychism holds that “facts about human and animal consciousness are not fundamental, but are grounded in more fundamental kinds of consciousness, e.g., facts about micro-level consciousness” (Goff et al. 2022, §2.2). Alternatively, it holds that macro-consciousness (the consciousness of subjects) is constituted or realized by micro-consciousness (Chalmers 2017, 25). Non-constitutive or emergent panpsychism claims that facts about macro-consciousness “are among the fundamental facts” (Goff et al. 2022, §2.2). In other words, macro-consciousness is “strongly emergent from micro-[consciousness] experience and/or from microphysics” (Chalmers 2017, 25).

How could PRIMER contribute to the panpsychism revolution? Roughly, subjectivity – or rather proto-subjectivity – could be argued to be a fundamental and ubiquitous property. If this view is combined with a view of fundamental properties as powerful qualities, the claim could then be that proto-subjectivity manifests as subjectivity in the presence of other reciprocal partner properties that together constitute life. In other words, life is the criterion for manifestation of subjectivity. Moreover, subjectivity could be argued to be strongly emergent from the proto-subjectivity of the subject’s micro-structural parts. Thus, all living organisms are subjects – they are characterized by subjectivity – and subjectivity then grounds the permeation-manner of instantiation. This specific version of panpsychism would be ‘*panprotosubjectivism*’.

To be sure, the above two options are mere suggestions. It would take my thesis too far on a tangent to adjudicate between them. However, PRIMER offers a promising platform to inquire into the metaphysics of subjectivity.

10. Conclusion

In this thesis, my aim was to provide a theory of perceptual intentionality, interpreted in terms of representation. The main research problem guiding me towards this end was the problem of distal content within a representationalist framework. I have labelled my novel representationalist theory Pluri-Relational Immanent Emergent Representationalism (PRIMER). The theoretical components of PRIMER are: (i) Content Universalism – the theory of content; and Representational Magnetism – the theory of application or directedness.

10.1. Metaphysical Highlights of PRIMER

The main metaphysical planks of PRIMER are as follows:

Immanent Content: Content is the essential representational property in virtue of which representations are directed towards their targets. According to PRIMER's theory of content – content universalism – content is constituted by three conditions: naturalistic relations between subject and the target, fitness-contribution functions of the subject, and the permeation-instantiation in the subject of the same universals that are qualification-instantiated in the target. The last content-constituting condition implies that content is identical to permeated (and, therefore, immanent) universals.

Emergent Content: Although the three content-constituting conditions are together sufficient for content, content is not reducible to them. This is because content is a strongly emergent property of the subject. It has novel causal powers of causing actions in subjects. These causal powers harmonize well with the action-oriented nature of content which follows from the essential purpose of perception to guide actions for fitness and survival.

Pluri-relationality: PRIMER posits relations as part of its theory of content and its theory of application. Content universalism posits content-constituting naturalistic relations and the instantiation relation (with its different manners – qualification and permeation). Representational Magnetism posits the Representational External Directedness (RED) relation between the representation and the target that holds in virtue of content and the property-instances of targets. This is because the same universals permeate in a subject as content and qualify a target as property-instances. Moreover, the RED relation is a weakly emergent relation and has distinct causal efficacy vis-à-vis content, although its causal powers are plausibly a subset of the causal powers of content. Since content is strongly emergent and the

RED relation is weakly emergent, PRIMER is a thoroughly nonreductive version of representationalism.

10.2. PRIMER's Niche

PRIMER occupies a distinct niche in the representationalist literature due to the following characteristics. First, it is both relational and nonreductive. To recall, to be relational is to posit content-constituting relations to the environment. In other words, it is to externalist with respect to content individuation. Most relational/external representationalist views are reductive. For example, the naturalistic representationalists discussed in chapter 3 and 4. Conversely, most non-relational/internalist views tend to be nonreductive; for example, the nonreductive views surveyed in chapter 6. To recall, such internalist views include con-relational views which posit to relations to abstract universals or content, and thoroughly non-relational views. PRIMER occupies a middle ground by being relational and nonreductive. Content is partly constituted and individuated by relations to the environment, and owing to its strong emergence, it is irreducible.

Secondly, the theory of universals which content universalism uses is heavily influenced by theory of universals propounded by the classical Indian philosophy of Nyāya. As discussed in chapter 7, the Nyāya theory is that universals are intrinsically concrete entities that pervade spacetime and can exist uninstantiated. Their concreteness does not depend on their being immanent in particulars. Thus, Nyāya universals occupy a middle ground between Platonic and Aristotelian universals. Like Platonic universals, they can exist uninstantiated; unlike Platonic universals, however, they are concrete even when uninstantiated. And like Aristotelian universals, they are concrete; unlike Aristotelian universals, however, they can exist uninstantiated. The universals that content universalism posits are akin to Nyāya universals. They can be also compared to the quantum fields posited by Quantum Field Theory, which, as discussed in chapter 7, play the role of universals.

10.3. Advantages of PRIMER

PRIMER has distinct advantages over rival representationalist theories. Over reductive-naturalistic representationalism it has the following advantages. First, it was argued in chapter

5 that naturalistic representationalism prioritises the theory of application over theory of content – it explains content in terms of representational directedness (i.e., tracking relation) – which does not make it genuinely representational. In contrast, PRIMER explains representational directedness (i.e., RED relation) in terms of content and restores the metaphysical priority of content over representational directedness.

Secondly, as argued in chapter 4, naturalistic representationalism cannot account for the distality of content because the representational application formulae (RAFs) that they posit to account for the tracking relation are not necessary for it. In its place, PRIMER proposes the condition of subjective permeation of the same universals that qualify distal targets, rather than proximal intermediaries. Thus, the distality of content is secured. As noted in chapter 9, universals and their instantiation is naturalistic in the sense of belonging to the spatiotemporal realm. Although PRIMER's account of content is nonreductive, it is naturalistic. Thirdly, PRIMER gives robust causal efficacy to content, since content is a strongly emergent property and has novel causal powers. This is in stark contrast to naturalistic representationalism which – as discussed in chapter 5 – either gives no causal role to content or an attenuated causal explanatory role.

PRIMER has advantages over its other nonreductive counterparts as well. First, it was noted in chapter 6 that correlational representationalism posits the acquaintance relation to abstract universals to account for phenomenal character. But, as a consequence, it faces the content-character mismatch problem – abstract content is unsuitable to account for concrete phenomenal character. In contrast, PRIMER accounts for phenomenal character in terms of content – it claims that phenomenal character is identical to phenomenal content, which is a certain kind of permeated universal. Thus, phenomenal content is concrete and is eminently suitable to be phenomenal character.

Second, con-relational representationalism also accounts for non-veridical perceptions such as hallucinations in terms of acquaintance with abstract universals. However, this has the weakness that acquaintance relation is a conscious direct awareness relation and, so, acquaintance with *abstract* universals must be phenomenologically captured. However, it always seems to us that concrete objects with their properties are phenomenologically present. As an alternative solution, PRIMER accounts for hallucinations by arguing that deviant concrete universals (those that do not qualify the target) permeate in the subject along with

normal universals. Thus, there is a non-veridical aspect to concrete content, which accounts for the hallucinatory experience.

Finally, neither con- and non-relational representationalism can satisfactorily explain the essential external directedness of perceptual representations and the distality of content. This is because the resources they avail – acquaintance with abstract universals and intrinsic psychological facts, respectively – are not essentially spatial. In contrast, because the content-constituting relations of PRIMER (including instantiation relation) are spatial, they can account for both the features.

Bibliography

- Ainsworth, Thomas. 2024. "Form vs. Matter." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman. Fall.
<https://plato.stanford.edu/archives/fall2024/entries/form-matter/>.
- Ali, Rami. 2018. "Does Hallucinating involve Perceiving?" *Philosophical Studies* 175 (3): 601-627.
- Anjum, Rani Lill, and Stephen Mumford. 2017. "Emergence and Demergence." In *Philosophical and scientific perspectives on downward causation*, edited by Michele Paolini Paoletti and Francesco Orilia, 92-109. New York: Routledge.
- Anjum, Rani Lill, and Stephen Mumford. 2018. *What Tends to Be: The Philosophy of Dispositional Modality*. New York: Routledge.
- Armstrong, D. M. 1997. *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Armstrong, D. M. 2004. *How do Particulars Stand to Universals?* Vol. 1, in *Oxford Studies in Metaphysics*, edited by Dean W. Zimmerman, 139-154. Oxford: Oxford University Press.
- Armstrong, D. M. 1978. *Universals and Scientific Realism*. Vol. 1 and 2. New York: Cambridge University Press.
- Armstrong, D. M. 1983/2016. *What is a Law of Nature?* Cambridge: Cambridge University Press.
- Armstrong, D.M. 1989. *Universals: An Opinionated Introduction*. Boulder: Westview Press.
- Arstila, Valtteri, and Kalle Pihlainen. 2009. "The Causal Theory of Perception Revisited." *Erkenntnis* 70 (3): 397-417.
- Baker, Alan. 2022. "Simplicity." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. <https://plato.stanford.edu/archives/sum2022/entries/simplicity/>.
- Baker, Lynne Rudder. 1995. *Explaining Attitudes: A Practical Approach to the Mind*. New York: Cambridge University Press.
- Baker, Lynne Rudder. 2019. "Intentionality and Emergence." In *The Routledge Handbook of Emergence*, edited by Sophie Gibb, Robin Findlay Hendry and Tom Lancaster, 206-214. New York: Routledge.
- Baker, Lynne Rudder. 2009. "Non-reductive Materialism." In *The Oxford Handbook of Philosophy of Mind*, edited by Brian McLaughlin, Ansgar Beckermann and Sven Walter, 109-127. New York: Oxford University Press.
- Baker, Lynne Rudder. 2007. *The Metaphysics of Everyday Life: An Essay in Practical Realism*. New York: Cambridge University Press.
- Baldwin, Thomas. 1992. "The Projective Theory of Sensory Content." In *The Contents of Experience*, edited by Tim Crane, 177-195. New York: Cambridge University Press.
- Barnes, Elizabeth. 2012. "Emergence and Fundamentality." *Mind* 121 (484): 873-901.

- Baysan, Umut. 2019. "Emergence, Function and Realization." In *The Routledge Handbook of Emergence*, edited by Sophie Gibb, Robin Findlay Hendry and Tom Lancaster, 77-86. New York: Routledge.
- Beaney, Michael, and Thomas Raysmith. 2024. "Analysis." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman.
<https://plato.stanford.edu/archives/fall2024/entries/analysis/>.
- Bellini, Salvatore. 2009. "On a Unique Behavior of Freshwater Bacteria." *Chinese Journal of Oceanology and Limnology* 27 (1): 3-5.
- Berlin, Heather A. 2011. "The Neural Basis of the Dynamic Unconscious." *Neuropsychanalysis* 13 (1): 5-31.
- Bermúdez, José Luis. 2009. "The Distinction between Conceptual and Nonconceptual Content." In *The Oxford Handbook of Philosophy of Mind*, edited by Brian McLaughlin, Ansgar Beckermann and Sven Walter, 457-473. New York: Oxford University Press.
- Bermúdez, José Luis, and Arnon Cahen. 2020. "Nonconceptual Mental Content." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Summer.
<https://plato.stanford.edu/archives/sum2020/entries/content-nonconceptual/>.
- Bird, Alexander. 2007. *Nature's Metaphysics: Laws and Properties*. Oxford: Oxford University Press.
- Blakemore, Richard. 1975. "Magnetotactic Bacteria." *Science* 190 (4212): 377-379.
- Bliss, Ricki, and Kelly Trogdon. 2024. "Metaphysical Grounding." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta and Uri Nodelman.
<https://plato.stanford.edu/archives/sum2024/entries/grounding/>.
- Block, Ned. 1986. "Advertisement for a Semantics for Psychology." *Midwest Studies in Philosophy* 10 (1): 615-678.
- Block, Ned. 1995. "On a Confusion About a Function of Consciousness." *Behavioral and Brain Sciences* 18: 227-247.
- Block, Ned. 2007b. "Overflow, Access, and Attention." *Behavioral and Brain Sciences* 30 (5-6): 530-548.
- Block, Ned. 2011. "Perceptual Consciousness Overflows Cognitive Access." *Trends in Cognitive Sciences* 15 (12): 567-575.
- Block, Ned. 2004. "Qualia." In *Oxford Companion to the Mind*, edited by Richard Langton Gregory, 785-789. New York: Oxford University Press.
- Block, Ned. 2011. "The Anna Karenina Theory of the Unconscious." *Neuropsychanalysis* 13 (1): 34-37.
- Bourget, David, and Angela Mendelovici. 2019. "Phenomenal Intentionality." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Fall.
<https://plato.stanford.edu/archives/fall2019/entries/phenomenal-intentionality/>.

- Bourget, David, and Angela Mendelovici. 2014. "Tracking Representationalism." In *Philosophy of Mind: The Key Thinkers*, edited by Andrew Bailey, 209-235. New York: Bloomsbury Academic.
- Brenner, Andrew, Anna-Sofia Maurin, Alexander Skiles, Robin Stenwall, and Naomi Thompson. 2021. "Metaphysical Explanation." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Winter. <https://plato.stanford.edu/archives/win2021/>.
- Brewer, Bill. 2019. "Visual Experience, Revelation, and the Three Rs." In *Acquaintance: New Essays*, edited by Jonathan Knowles and Thomas Raleigh, 277-292.
- Brogaard, Berit. 2014. "Introduction: Does Perception Have Content?" In *Does Perception Have Content?*, edited by Berit Brogaard, 1-35. New York: Oxford University Press.
- Brown, Curtis. 2022. "Narrow Mental Content." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. <https://plato.stanford.edu/archives/sum2022/entries/content-narrow/>.
- Brunswik, Egon. 1956. *Perception and the Representative Design of Psychological Experiments*. Berkeley and Los Angeles: University of California Press.
- Brüntrup, Godehard, and Ludwig Jaskolla. 2017. *Panpsychism: Contemporary Perspectives*. New York: Oxford University Press.
- Burge, Tyler. 2005. "Disjunctivism and Perceptual Psychology." *Philosophical Topics* 33 (1): 1-78.
- Burge, Tyler. 1979. "Individualism and the Mental." *Midwest Studies in Philosophy* 4 (1): 73-121.
- Burge, Tyler. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.
- Burge, Tyler. 2014. "Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints." *Philosophy and Phenomenological Research* 88 (2): 485-501.
- Burnett, S. 2011. "Perceptual Worlds and Sensory Ecology." *Nature Education Knowledge* 3 (10): 75.
- Byrne, Alex. 2009. "Experience and Content." *The Philosophical Quarterly* 59 (236): 429-451.
- Byrne, Alex. 2006. "Intentionality." In *The Philosophy of Science: An Encyclopedia*, edited by Sahotra Sarkar and Jessica Pfeifer, 405-410. New York: Routledge.
- Byrne, Alex. 2005. "Perception and Conceptual Content." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa, 231-250. Malden, MA: Wiley-Blackwell.
- Byrne, Alex, and Riccardo Manzotti. 2022. "Hallucination and Its Objects." *Philosophical Review* 131 (3): 327-359.
- Campbell, John. 2015. "Perceptual Attention." In *The Oxford Handbook of the Philosophy of Perception*, edited by Mohan Matthen, 587-602. New York: Oxford University Press.
- Carrasco, Marisa, Sam Ling, and Sarah Read. 2004. "Attention Alters Appearance." *Nature Neuroscience* 7 (3): 308-313.

- Carter, J. Adam, Jesper Kallestrup, S. Orestis Palermos, and Duncan Pritchard. 2014. "Varieties of Externalism." *Philosophical Issues* 24 (1): 63-109.
- Caston, Victor. 2012. "Aristotle's Psychology." In *A Companion to Ancient Philosophy*, edited by Mary Louise Gill and Pierre Pellegrin, 316-346. Malden, MA: Blackwell.
- Cavanagh, Patrick, Gideon P. Caplovitz, Taissa K. Lytchenko, Marvin R. Maechler, Peter U. Tse, and David L. Sheinberg. 2023. "The Architecture of Object-based Attention." *Psychonomic Bulletin & Review* 30 (5): 1643-1667.
- Chadha, Monima. 2024. "Knowledge of Universals." *British Journal for the History of Philosophy* 1-21.
- Chadha, Monima. 2014. "On Knowing Universals: The Nyāya Way." *Philosophy East and West* 64 (2): 287-302.
- Chadha, Monima. 2024. "Perceptual Experience and Concepts in Classical Indian Philosophy." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman. Summer. <https://plato.stanford.edu/archives/sum2024/entries/perception-india/>.
- Chalmers, David J. 2010. "Consciousness and Its Place in Nature." In *The Character of Consciousness*, by David J. Chalmers, 103-140. New York: Oxford University Press.
- Chalmers, David J. 2017a. "Panpsychism and Panprotopsychism." In *Panpsychism: Contemporary Perspectives*, edited by Godehard Brüntrup and Ludwig Jaskolla, 19-47. New York: Oxford University Press.
- Chalmers, David J. 2006a. "Perception and the Fall from Eden." In *Perceptual Experience*, edited by Tamar Szabo Gendler & John Hawthorne, 49-125. New York: Oxford University Press.
- Chalmers, David J. 2006b. "Strong and Weak Emergence." In *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion*, edited by Philip Clayton and Paul Davies, 244-255. Oxford: Oxford University Press.
- Chalmers, David J. 2004. "The Representational Character of Experience." In *The Future for Philosophy*, edited by Brian Leiter, 153-181. New York: Oxford University Press.
- Chalmers, David J. 2012. "Twentieth Excursus: Reference Magnets and the Grounds of Intentionality." In *Constructing The World: Extended Edition*, by David J. Chalmers, 1-11. Oxford: Oxford University Press. <https://consc.net/books/ctw/extended.html>.
- Chittka, Lars, and Randolph Menzel. 1992. "The Evolutionary Adaptation of Flower colours and the Insect Pollinators' Colour Vision." *Journal of Comparative Physiology A* 171: 171-181.
- Chomsky, Noam. 1995. "Language and Nature." *Mind* 104 (413): 1-61.
- Clark, Andy, and David J. Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7-19.
- Clark, Austen. 2000. *A Theory of Sentience*. New York: Oxford University Press.
- Clark, D. L., and R. A. Clark. 2016. "Neutral Point Testing of Color Vision in the Domestic Cat." *Experimental Eye Research* 153: 23-26.
- Cohen, Jonathan. 2004. "Objects, Places, and Perception." *Philosophical Psychology* 17 (4): 471-495.

- Cohen, S. Marc, and C. D. C. Reeve. 2025. "Aristotle's Metaphysics." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman.
<https://plato.stanford.edu/archives/spr2025/entries/aristotle-metaphysics/>.
- Colombo, Matteo, and Gualtiero Piccinini. 2023. *The Computational Theory of Mind*. Cambridge: Cambridge University Press.
- Corkum, Phil. 2025. *Neo-Aristotelian Metaphysics*. Cambridge: Cambridge University Press.
- Crane, Tim. 2017. "Brentano on Intentionality." In *The Routledge Handbook of Franz Brentano and the Brentano School*, edited by Uriah Kriegel, 41-48. New York: Routledge.
- Crane, Tim. 2009. "Intentionalism." In *The Oxford Handbook of Philosophy of Mind*, edited by Brian McLaughlin, Ansgar Beckermann and Sven Walter, 474-493. New York: Oxford University Press.
- Crane, Tim. 1998b. "Intentionality." *Routledge Encyclopedia of Philosophy*. Taylor and Francis.
<https://www.rep.routledge.com/articles/thematic/intentionality/v-1>.
- Crane, Tim. 1998a. "Intentionality as the Mark of the Mental." In *Contemporary Issues in the Philosophy of Mind*, edited by Anthony O'Hear, 229-251. Cambridge: Cambridge University Press.
- Crane, Tim. 2011. "Is Perception a Propositional Attitude?" In *The Admissible Contents of Experience*, edited by Katherine Hawley and Fiona Macpherson, 83-100. Chichester: Wiley-Blackwell.
- Crane, Tim. 2006. "Is There a Perceptual Relation?" In *Perceptual Experience*, edited by Szabo Tamar Gendler and John Hawthorne, 126-146. New York: Oxford University Press.
- Crane, Tim. 2014. "The Given." In *Aspects of Psychologism*, by Tim Crane, 235-255. Cambridge, MA: Harvard University Press.
- Crane, Tim. 2016. *The Mechanical Mind: A Philosophical Introduction to Minds, Machines and Mental Representation*. New York: Routledge.
- Crane, Tim. 2013. *The Objects of Thought*. New York: Oxford University Press.
- Crane, Tim. 2001. "The Significance of Emergence." In *Physicalism and its Discontents*, edited by Carl Gillett and Barry Loewer, 207-224. Cambridge: Cambridge University Press.
- Crane, Tim, and Craig French. 2021. "The Problem of Perception." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Fall .
<https://plato.stanford.edu/archives/fall2021/entries/perception-problem/>.
- Cummins, Robert. 1975. "Functional Analysis." *Journal of Philosophy* 72 (20): 741-764.
- Cummins, Robert. 1989. *Meaning and Mental Representation*. MIT Press.
- Cummins, Robert. 1996. *Representations, Targets and Attitudes*. Cambridge, MA: MIT Press.
- Cummins, Robert. 1983. *The Nature of Psychological Explanation*. Cambridge, MA: The MIT Press.

- Cummins, Robert, and Pierre Poirier. 2004. "Representation and Indication." In *Representation in Mind: New Approaches to Mental Representation*, edited by Hugh Clapin, Phillip Staines and Peter Slezak, 21-40. Amsterdam: Elsevier.
- Cummins, Robert, Jim Blackmon, David Byrd, Alexa Lee, and Martin Roth. 2006. "Representation and Unexploited Content." In *Teleosemantics: New Philosophical Essays*, edited by Graham Macdonald and David Papineau, 195-207. Oxford: Oxford University Press.
- Cussins, Adrian. 2003. "Experience, Thought and Activity." In *Essays on Nonconceptual Content*, edited by York Gunther, 147-163. Cambridge, MA: The MIT Press.
- Cussins, Adrian. 1990. "The Connectionist Construction of Concepts." In *The Philosophy of Artificial Intelligence*, edited by Margaret A. Boden, 368-440. Oxford: Oxford University Press.
- Cutter, Brian, and Michael Tye. 2011. "Tracking Representationalism and the Painfulness of Pain." *Philosophical Issues* 21 (1): 90-109.
- Dasgupta, Shamik. 2017. "Constitutive Explanation." *Philosophical Issues* 27 (1): 74-97.
- Dawson, Michael R. W. 2013. *Mind, Body, World: Foundations of Cognitive Science*. Edmonton: Athabasca University Press.
- Dennett, Daniel C. 1969. *Content and Consciousness*. Routledge & Kegan Paul.
- Dennett, Daniel C. 2007. "Philosophy as Naive Anthropology: Comment on Bennett and Hacker." In *Neuroscience and Philosophy: Brain, Mind, and Language*, edited by M. Bennett, D. C. Dennett, P. M. S. Hacker and J. R. Searle, 73-95. New York: Columbia University Press.
- Dennett, Daniel C. 1987. *The Intentional Stance*. Cambridge, MA: The MIT Press.
- Dobzhansky, Theodosius. 1973. "Nothing in Biology Makes Sense except in the Light of Evolution." *The American Biology Teacher* 35 (3): 125-129.
- Drayson, Zoe. 2012. "The Uses and Abuses of the Personal/Subpersonal Distinction." *Philosophical Perspectives* 26 (1): 1-18.
- Drayson, Zoe. 2017. "What is Action-Oriented Perception?" In *Logic, Methodology and Philosophy of Science: Proceedings of the 15th International Congress*, edited by Hannes Leitgeb, Ilkka Niiniluoto, Päivi Seppälä and Elliott Sober, 287-300. College Publications.
- Dretske, Fred. 2003. "Experience as Representation." *Philosophical Issues* 13 (1): 67-82.
- Dretske, Fred. 1988. *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: The MIT Press.
- Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Dretske, Fred. 1986. "Misrepresentation." In *Belief: Form, Content and Function*, edited by Radu Bogdan, 17-36. New York: Oxford University Press.
- Dretske, Fred. 1997. *Naturalizing The Mind*. Cambridge, MA: MIT Press.

- Dunaway, Billy. 2020. *Reality and Morality*. New York: Oxford University Press.
- Duncan, Matt. 2021. "Acquaintance." *Philosophy Compass* 16 (3): e12727.
- Egan, Frances. 2014. "How to Think about Mental Content." *Philosophical Studies* 170 (1): 115-135.
- Falcon, Andrea. 2023. "Aristotle on Causality." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman.
<https://plato.stanford.edu/archives/spr2023/entries/aristotle-causality/>.
- Fine, Kit. 1994. "Essence and Modality." *Philosophical Perspectives* 8 (Logic and Language): 1-16.
- Fish, William. 2009. *Perception, Hallucination, and Illusion*. New York: Oxford University Press.
- Fish, William. 2021. *Philosophy of Perception: A Contemporary Introduction*. New York: Routledge.
- Fitch, Greg, and Michael Nelson. 2024. "Singular Propositions." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman.
<https://plato.stanford.edu/archives/spr2024/entries/propositions-singular/>.
- Fodor, J. A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Fodor, J. A. 1975. *The Language of Thought*. New York: Harvard University Press.
- Fodor, J. A. 1997. "The Representational Theory of Mind." *American Behavioral Scientist* 40 (6): 829-841.
- Frege, Gottlob. 1892/1980. "On Sense and Reference." In *Translations from the Philosophical Writings of Gottlob Frege*, edited by Peter Geach and Max Black, 36-56. Oxford: Blackwell.
- Friedenberg, Jay, Gordon Silverman, and Michael J. Spivey. 2022. *Cognitive Science: An Introduction to the Study of the Mind (Fourth Edition)*. Los Angeles: Sage .
- Gallagher, Shaun. 2023. *Embodied and Enactive Approaches to Cognition*. Cambridge: Cambridge University Press.
- Garson, Justin. 2016. *A Critical Overview of Biological Functions*. Cham: Springer.
- Garson, Justin. 2019. "Do Constancy Mechanisms Save Distal Content?" *The Philosophical Quarterly* 69 (275): 409–417.
- Genone, James. 2016. "Recent Work on Naive Realism." *American Philosophical Quarterly* 53 (1): 1-25.
- Gibb, Sophie. 2015. "The Causal Closure Principle." *The Philosophical Quarterly* 65 (261): 626-647.
- Gibson, James J. 1979/2015. *The Ecological Approach to Visual Perception: Classic Edition*. New York: Psychology Press.

- Gładziejewski, Paweł, and Marcin Miłkowski. 2017. "Structural Representations: Causally Relevant and Different from Detectors." *Biology & Philosophy* 32: 337-355.
- Godfrey-Smith, Peter. 2006. "Mental Representation, Naturalism, and Teleosemantics." In *Teleosemantics: New Philosophical Essays*, edited by Graham Macdonald and Papineau. David, 42-68. Oxford: Oxford University Press.
- Godfrey-Smith, Peter. 1989. "Misinformation." *Canadian Journal of Philosophy* 19 (4): 533-550.
- Goff, Philip, William Seager, and Sean Allen-Hermanson. 2022. "Panpsychism." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta.
<https://plato.stanford.edu/archives/sum2022/entries/panpsychism>.
- Gomes, Herman. 2000. *Marr's Theory: From primal sketch to 3-D models*. Accessed July 2025.
https://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/GOMES1/marr.html.
- Green, E. J. 2023. "Perceptual Constancy and Perceptual Representation." *Analytic Philosophy* 00: 1-41.
- Grice, H. P. 1961. "The Causal Theory of Perception." *Proceedings of the Aristotelian Society (Supplementary Volume)* 35 (1): 121-52.
- Haugeland, John. 1991. "Representational Genera." In *Philosophy and Connectionist Theory*, edited by William Ramsey, David E. Rumelhart and Stephen P. Stich, 61-89. Hillsdale, NJ: Laurence Erlbaum.
- Heil, John. 2021. *Relations*. Cambridge: Cambridge University Press.
- Heil, John. 2012. *The Universe as We Find It*. Oxford: Oxford University Press.
- Hendry, Robin Findlay. 2010. "Ontological Reduction and Molecular Structure." *Studies in History and Philosophy of Modern Physics* 41 (2): 183-191.
- Hendry, Robin Findlay. 2021. "Structure, Scale and Emergence." *Studies in History and Philosophy of Science* 85: 44-53.
- Hendry, Robin Findlay, Sophie Gibb, and Tom Lancaster. 2019. "Introduction." In *The Routledge Handbook of Emergence*, edited by Sophie Gibb, Robin Findlay Hendry and Tom Lancaster, 1-19. New York: Routledge.
- Hobson, Art. 2017. *Tales of the Quantum: Understanding Physics; Most Fundamental Theory*. New York: Oxford University Press.
- Hoffman, Donald. 2019. *The Case Against Reality: How Evolution Hid the Truth from Our Eyes*. New York: W. W. Norton & Company.
- Humphreys, Paul. 2016. *Emergence: A Philosophical Account*. New York: Oxford University Press.
- Humphreys, Paul. 2008. "Synchronic and Diachronic Emergence." *Minds and Machines* 18 (4): 431-442.
- Husserl, Edmund. 1931/1960. *Cartesian Meditations: An Introduction to Phenomenology*. Translated by Dorion Cairns. The Hague: M. Nijhoff.

- Jacob, Pierre. 2023. "Intentionality." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta and Uri Nodelman. Spring.
<https://plato.stanford.edu/archives/spr2023/entries/intentionality/>.
- Jacquette, Dale. 2004. "Brentano's Concept of Intentionality." In *The Cambridge Companion to Brentano*, edited by Dale Jacquette, 98-130. New York: Cambridge University Press.
- James, William. 1902/2002. *The Varieties of Religious Experience: A Study in Human Nature*. London: Routledge.
- Johnston, Mark. 1992. "Constitution is Not Identity." *Mind* 101: 89-106.
- Johnston, Mark. 1992. "How to Speak of the Colors." *Philosophical Studies* 68 (3): 221-263.
- Kentridge, Robert W., and Berit Brogaard. 2017. "The Functional Roles of Attention." In *Current Controversies in Philosophy of Perception*, edited by Bence Nanay, 139-147. New York: Routledge.
- Kim, Jaegwon. 1992. "Multiple Realization and the Metaphysics of Reduction." *Philosophy and Phenomenological Research* 52 (1): 1-26.
- Kim, Jaegwon. 2011. *Philosophy of Mind*. 3rd. Boulder, CO: Westview Press.
- Kim, Jaegwon. 2005. *Physicalism, or Something Near Enough*. New Jersey: Princeton University Press.
- Kim, Jaegwon. 1993. "The Non-Reductivist's Troubles with Mental Causation." In *Mental Causation*, edited by John Heil and Alfred Mele, 189-210. Oxford: Oxford University Press,.
- King, Jeffrey C. 2017. "The Metaphysics of Propositions." *Oxford Handbooks Online*.
doi:10.1093/oxfordhb/9780199935314.013.26.
- Kirchhoff, Michael, and Julian Kiverstein. 2024. "Diachronic Constitution." *Manuscrito* 47 (1): 1-33.
- Klumpp, Stefan, and Damien Faivre. 2016. "Magnetotactic Bacteria: Magnetic Navigation on the Microscale." *European Physical Journal Special Topics* 225 (11): 2173-2188.
- Kornblith, Hilary. 2016. "Philosophical Naturalism." In *The Oxford Handbook of Philosophical Methodology*, edited by Herman Cappelen, Tamar Szabó Gendler and John Hawthorne, 147-158. Oxford: Oxford University Press.
- Kriegel, Uriah. 2018. *Brentano's Philosophical System: Mind, Being, Value*. Oxford: Oxford University Press.
- Kriegel, Uriah. 2007. "Intentional Inexistence and Phenomenal Intentionality." *Philosophical Perspectives* 21 (1): 307-340.
- Kriegel, Uriah. 2002. "Phenomenal Content." *Erkenntnis* 57 (2): 175-198.
- Kriegel, Uriah. 2009. *Subjective Consciousness: A Self-Representational Theory*. New York: Oxford University Press.
- Kriegel, Uriah. 2008. "The Dispensability of (Merely) Intentional Objects." *Philosophical Studies* 141 (1): 79-95.

- Kriegel, Uriah. 2013. "The Phenomenal Intentionality Research Program." In *Phenomenal Intentionality*, 1-26. New York: Oxford University Press.
- Kriegel, Uriah. 2011. *The Sources of Intentionality*. Oxford: Oxford University Press.
- Kriegel, Uriah, and Kenneth Williford. 2006. *Self-Representational Approaches to Consciousness*. Cambridge, MA: The MIT Press.
- Kubovy, M., and D. Van Valkenburg. 2001. "Auditory and Visual Objects." *Cognition* 80 (1-2): 97-126.
- Kuphaldt, Tony R. 2025. "Conventional Versus Electron Flow." *All About Circuits*.
<https://www.allaboutcircuits.com/textbook/direct-current/chpt-1/conventional-versus-electron-flow/>.
- Landy, Michael S., Martin S. Banks, and David C. Knill. 2011. "Ideal-Observer Models of Cue Integration." In *Sensory Cue Integration*, edited by Julia Trommershäuser, Konrad Kording and Michael S. Landy, 5-29. Oxford: Oxford University Press.
- Leech, Jessica. 2021. "From Essence to Necessity via Identity." *Mind* 130 (519): 887-908.
- Lefèvre, Christopher T., and Dennis A. Bazylinski. 2013. "Ecology, Diversity, and Evolution of Magnetotactic Bacteria." *Microbiology and Molecular Biology Reviews* 77 (3): 497-526.
- Leunissen, Mariska. 2010. *Explanation and Teleology in Aristotle's Science of Nature*. Cambridge: Cambridge University Press.
- Lewis, David. 1973. *Counterfactuals*. Oxford: Blackwell.
- Lewis, David. 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61 (4): 343-377.
- Liggins, David. 2024. *Abstract Objects*. Cambridge: Cambridge University Press.
- Logue, Heather. 2014. "Experiential Content and Naive Realism: A Reconciliation." In *Does Perception Have Content?*, edited by Berit Brogaard, 220-241. New York: Oxford University Press.
- Logue, Heather. 2012. "Why Naive Realism?" *Proceedings of the Aristotelian Society* 112: 211-237.
- Loux, Michael J. 2005. "Aristotle on Matter, Form, and Ontological Strategy." *Ancient Philosophy* 25 (1): 81-123.
- Loux, Michael J. 2007. "Form in Aristotle's Constituent Ontology." In *Substance and Attribute: Western and Islamic Traditions in Dialogue*, edited by Christian Kanzian & Muhammad Legenhausen, 145-162. Lancaster, LA: Ontos Verlag.
- Lowe, E. J. 2003. "Individuation." In *The Oxford Handbook of Metaphysics*, edited by Michael J. Loux and Dean W. Zimmerman, 75-95. New York: Oxford University Press.
- Lowe, E. J. 2018. "Metaphysics as the Science of Essence." In *Ontology, Modality, and Mind: Themes from the Metaphysics of E. J. Lowe*, edited by Alexander Carruth, Sophie Gibb and John Heil, 14-34. Oxford: Oxford University Press.

- Lowe, E. J. 2008. *Personal Agency: The Metaphysics of Mind and Action*. New York: Oxford University Press.
- Lowe, E. J. 2006. *The Four-Category Ontology: A Metaphysical Foundation for Natural Science*. Oxford: Clarendon Press.
- Lowe, E. J. 2016. "There Are (Probably) No Relations." In *The Metaphysics of Relations*, edited by Anna Marmodoro and David Yates, 100-112. Oxford: Oxford University Press.
- Lycan, William. 2023. "Representational Theories of Consciousness." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta and Uri Nodelman. Winter. <https://plato.stanford.edu/archives/win2023/entries/consciousness-representational/>.
- MacBride, Fraser. 2020. "Relations." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Winter. <https://plato.stanford.edu/archives/win2020/entries/relations/>.
- Macpherson, Fiona. 2009. "Perception, Philosophical Perspectives." In *The Oxford Companion to Consciousness*, edited by Patrick Wilken, Timothy J. Bayne and Axel Cleeremans, 502-508. New York: Oxford University Press.
- Macpherson, Fiona. 2013. "The Philosophy and Psychology of Hallucination: An Introduction." In *Hallucination: Philosophy and Psychology*, edited by Fiona Macpherson & Dimitris Platchias, 1-38. Cambridge, MA: MIT Press.
- Macpherson, Fiona, and Clare Batty. 2016. "Redefining Illusion and Hallucination in Light of New Cases." *Philosophical Issues* 26 (1): 263-296.
- Maier, Markus A., Russell A. Hill, Andrew J. Elliot, and Robert A. Barton. 2015. "Color in Achievement Contexts in Humans." In *Handbook of Color Psychology*, edited by Mark D. Fairchild, and Anna Franklin Andrew J. Elliot, 568-584.
- Marcus, Eric. 2005. "Mental Causation in a Physical World." *Philosophical Studies* 122 (1): 27-50.
- Marek, Johann. 2024. "Alexius Meinong." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta and Uri Nodelman. Spring. <https://plato.stanford.edu/archives/spr2024/entries/meinong/>.
- Marmodoro, Anna. 2021. "Instantiation." *Metaphysics* 4 (1): 32-46.
- Marr, David. 1982/2010. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Cambridge, MA: The MIT Press.
- Martin, C. B. 2008. *The Mind in Nature*. New York: Oxford University Press.
- Martin, M. G. F. 2006. "On Being Alienated." In *Perceptual Experience*, edited by Tamar Szabo Gendler & John Hawthorne, 354-410. New York: Oxford University Press.
- Martin, M. G. F. 2004. "The Limits of Self-Awareness." *Philosophical Studies* 120 (1-3): 37-89.
- Masrour, Farid. 2020. "On the Possibility of Hallucinations." *Mind* 129 (515): 737-768.
- Matilal, Bimal Krishna. 1986. *Perception: An Essay on Classical Indian Theories of Knowledge*. Oxford: Oxford University Press.

- Matthen, Mohan. 2005. *Seeing, Doing, and Knowing: A Philosophical Theory of Sense Perception*. Oxford: Oxford University Press.
- Mayr, Ernst. 1961. "Cause and Effect in Biology: Kinds of causes, predictability, and teleology are viewed by a practicing biologist." *Science* 134 (3489): 1501–1506.
- McDowell, John. 1996. *Mind and World*. Cambridge, MA: Harvard University Press.
- McDowell, John. 1994. "The Content of Perceptual Experience." *The Philosophical Quarterly* 44 (175): 190–205.
- McGinn, Colin. 1982. *The Character of Mind*. New York: Oxford University Press.
- McGrath, Matthew, and Devin Frank. 2024. "Propositions." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman.
<https://plato.stanford.edu/archives/fall2024/entries/propositions/>.
- McLaughlin, Brian P. 1992. "The Rise and Fall of British Emergentism." In *Emergence or Reduction?: Essays on the Prospects of Nonreductive Physicalism*, edited by Ansgar Beckermann, Hans Flohr and Jaegwon Kim, 49–93. New York: De Gruyter.
- McLaughlin, Brian P. 1995. "Varieties of Supervenience." In *Supervenience: New Essays*, edited by Elias E. Savellos and Ümit D. Yalçın, 16–59. New York: Cambridge University Press.
- McLaughlin, Brian, and Karen Bennett. 2023. "Supervenience." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta and Uri Nodelman.
<https://plato.stanford.edu/archives/win2023/entries/supervenience/>.
- Mendelovici, Angela. 2018. *The Phenomenal Basis of Intentionality*. New York: Oxford University Press.
- Merrill, G. H. 1980. "The Model-Theoretic Argument Against Realism." *Philosophy of Science* 47 (1): 69–81.
- Millikan, Ruth Garrett. 1989a. "An Ambiguity in the Notion "Function"." *Biology and Philosophy* 4: 172–176.
- Millikan, Ruth Garrett. 1989b. "Biosemantics." *The Journal of Philosophy* 86 (6): 281–297.
- Millikan, Ruth Garrett. 2023. "Comment on Artiga's "Teleosemantics and Pushmi-Pullyu Representations"." *Erkenntnis* 88 (1): 409–417.
- Millikan, Ruth Garrett. 1989c. "In Defense of Proper Functions." *Philosophy of Science* 56 (2): 288–302.
- Millikan, Ruth Garrett. 1984. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: The MIT Press.
- Millikan, Ruth Garrett. 1996. *Pushmi-pullyu Representations*. Vol. 9, in *Philosophical Perspectives*, edited by James Tomberlin, 185–200. Atascadero, CA: Ridgeview Publishing.
- Millikan, Ruth Garrett. 2013b. "Reply to Godfrey-Smith." In *Millikan and Her Critics*, edited by Dan Ryder, Justine Kingsbury and Kenneth Williford, 59–62. Chichester: Wiley-Blackwell.

- Millikan, Ruth Garrett. 2013a. "Reply to Neander." In *Millikan and Her Critics*, edited by Dan Ryder, Justine Kingsbury and Kenneth Williford, 37-40. Chichester: Wiley-Blackwell.
- Millikan, Ruth Garrett. 2013c. "Reply to Shea." In *Millikan and Her Critics*, edited by Dan Ryder, Justine Kingsbury and Kenneth Williford, 81-86. Chichester: Wiley-Blackwell.
- Millikan, Ruth Garrett. 2004. *Varieties of Meaning*. Cambridge, MA: The MIT Press.
- Mumford, Stephen, and Rani Lill Anjum. 2011. *Getting Causes From Powers*. New York: Oxford University Press.
- Nagel, Ernest. 1961. *The Structure of Science*. New York: Harcourt, Brace and World.
- Nanay, Bence. 2013. *Between Perception and Action*. Oxford: Oxford University Press.
- Neander, Karen. 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics*. Cambridge, MA: The MIT Press.
- Neander, Karen. 1995. "Misrepresenting & Malfunctioning." *Philosophical Studies* 79 (2): 109-141.
- Neander, Karen. 2013. "Toward an Informational Teleosemantics." In *Millikan and Her Critics*, edited by Dan Ryder, Justine Kingsbury and Kenneth Williford, 21-40. West Sussex: Wiley-Blackwell.
- Newen, Albert, Shaun Gallagher, and Leon De Bruin. 2018. "4E Cognition: Historical Roots, Key Concepts, and Central Issues." In *The Oxford Handbook of 4E Cognition*, 3-15. Oxford : Oxford University Press .
- Noë, Alva. 2004. *Action in Perception*. Cambridge, MA: The MIT Press.
- Nolan, Daniel. 1997. "Quantitative Parsimony." *British Journal for the Philosophy of Science* 48 (3): 329-343.
- Novotny, Daniel, and Lukáš Novák, . 2013. *Neo-Aristotelian Perspectives in Metaphysics*. London: Routledge.
- Nudds, Matthew. 2009. "Sounds and Space." In *Sounds and Perception: New Philosophical Essays*, edited by Matthew Nudds and Casey O'Callaghan, 69-96. Oxford: Oxford University Press.
- O'Brien, Gerard, and Jon Opie. 2004. "Notes Toward a Structuralist Theory of Mental Representation." In *Representation in Mind: New Approaches to Mental Representation*, edited by Hugh Clapin, Phillip Staines and Peter Slezak, 1-20. Amsterdam: Elsevier.
- O'Connor, Timothy. 2021. "Emergent Properties." *Stanford Encyclopedia of Philosophy*. Edited by Edward N Zalta. <https://plato.stanford.edu/archives/win2021/entries/properties-emergent/>.
- O'Connor, Timothy. 2018. "For Emergent Individualism." In *The Blackwell Companion to Substance Dualism*, edited by Jonathan J. Loose, Angus J. L. Menuge and J. P. Moreland, 368-376. Oxford: Wiley-Blackwell.
- O'Connor, Timothy, and Hong Yu Wong. 2005. "The Metaphysics of Emergence." *Noûs* 39 (4): 658-678.

- O'Connor, Timothy, and John Ross Churchill. 2010. "Is Non-reductive Physicalism Viable within a Causal Powers Metaphysic?" In *Emergence in Mind*, edited by Graham Macdonald & Cynthia Macdonald, 43-60. New York: Oxford University Press.
- O'Regan, J. Kevin, and Alva Noë. 2001. "A Sensorimotor Account of Vision and Visual Consciousness." *Behavioral and Brain Sciences* 24 (5): 883-917.
- Paoletti, Michele Paolini. 2025. *The Ontology of Relations*. New York: Routledge.
- Papineau, David. 2023. "Naturalism." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman.
<https://plato.stanford.edu/archives/fall2023/entries/naturalism/>.
- Papineau, David. 1993. *Philosophical Naturalism*. Oxford: Blackwell.
- Papineau, David. 1987. *Reality and Representation*. New York: Blackwell.
- Papineau, David. 2017. "Teleosemantics." In *How Biology Shapes Philosophy: New Foundations for Naturalism*, edited by David Livingstone Smith, 95-120. Cambridge: Cambridge University Press.
- Papineau, David. 2021. *The Metaphysics of Sensory Experience*. New York: Oxford University Press.
- Papineau, David. 2002. *Thinking About Consciousness*. Oxford: Oxford University Press .
- Pautz, Adam. 2007. "Intentionalism and Perceptual Presence." *Philosophical Perspectives* 21 (1): 495-541.
- Pautz, Adam. 2021. *Perception*. Oxon: Routledge.
- Pautz, Adam. 2017. "The Significance Argument for the Irreducibility of Consciousness." *Philosophical Perspectives* 31 (1): 349-407.
- Pautz, Adam. 2009. "What are the Contents of Experience?" Edited by Katherine Hawley and Fiona Macpherson. *Philosophical Quarterly* 59 (236): 483-507.
- Peacock, Kent A. 2011. "The Three Faces of Ecological Fitness." *Studies in History and Philosophy of Biological and Biomedical Sciences* 42 (1): 99-105.
- Peacocke, Christopher. 1992. *A Study of Concepts*. Cambridge, MA: MIT Press.
- Phillips, Stephen. 2019. "Perception in Nyāya." In *The Senses and the History of Philosophy*, edited by Brian Glenney and José Filipe Silva, 49-61. New York: Routledge.
- Plack, Christopher J. 2024. *The Sense of Hearing*. Oxon: Routledge.
- Potter, Karl H., ed. 1977. *The Encyclopedia of Indian Philosophies: Indian Metaphysics and Epistemology: The Tradition of Nyāya-Vaiśeṣika up to Gaṅgeśa*. Vol. 2. Princeton: Princeton University Press.
- Prinz, Jesse J. 2015. "Unconscious Perception." In *The Oxford Handbook of the Philosophy of Perception*, edited by Mohan Matthen, 371-389. New York: Oxford University Press.
- Putnam, Hilary. 1975. "The Meaning of 'Meaning'." In *Minnesota Studies in the Philosophy of Science*, Vol. 7, 131-193. Minneapolis: University of Minnesota Press.

- Pylyshyn, Z. W. 2001. "Visual Indexes, Preconceptual Objects, and Situated Vision." *Cognition* 80 (1-2): 127-158.
- Pylyshyn, Z. W., and R. W. Storm. 1988. "Tracking Multiple Independent Targets: Evidence for a Parallel Tracking Mechanism." *Spatial Vision* 3 (3): 179-197.
- Quine, W. V. O. 1969. *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Raleigh, Thomas. 2019. "The Recent Renaissance of Acquaintance." In *Acquaintance: New Essays*, edited by Jonathan Knowles & Thomas Raleigh, 1-30. Oxford: Oxford University Press.
- Rensink, Ronald A. 2015. "Attention and Perception." In *Emerging Trends in the Social and Behavioral Sciences*, edited by Robert A. Scott, Stephen M. Kosslyn and Marlis C. Buchmann, 1-14. Hoboken NJ: Wiley.
- Rensink, Ronald A. 2002. "Change Detection." *Annual Review of Psychology* 53 (1): 245-277.
- Rensink, Ronald A. 2013. "Perception and Attention." In *The Oxford Handbook of Cognitive Psychology*, edited by Daniel Reisberg, 97-116. New York: Oxford University Press.
- Rensink, Ronald A. 2000. "The Dynamic Representation of Scenes." *Visual Cognition* 7 (1-3): 17-42.
- Rescorla, Michael. 2015. "Bayesian Perceptual Psychology." In *The Oxford handbook of philosophy of perception*, edited by Mohan Matthen, 694-716. Oxford: Oxford University Press.
- Rescorla, Michael. 2020. "Perceptual Co-Reference." *Review of Philosophy and Psychology* 11 (3): 569-589.
- Riener, Cedar, and Jeanine Stefanucci. 2024. "Perception and/for/with/as Action." In *The Routledge Handbook of Embodied Cognition*, edited by Lawrence Shapiro and Shannon Spaulding, 169-177. Oxon: Routledge.
- Robb, David, John Heil, and Sophie Gibb. 2023. "Mental Causation." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta and Uri Nodelman.
<https://plato.stanford.edu/archives/spr2023/entries/mental-causation/>.
- Robertson Ishii, Teresa, and Philip Atkins. 2023. "Essential vs. Accidental Properties." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman.
<https://plato.stanford.edu/archives/spr2023/entries/essential-accidental/>.
- Robinson, Howard. 1994. *Perception*. London: Routledge.
- Roca-Royes, Sonia. 2011. "Essential Properties and Individual Essences." *Philosophy Compass* 6 (1): 65-77.
- Rosenberg, Alexander, and Frederic Bouchard. 2021. "Fitness." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Fall.
<https://plato.stanford.edu/archives/fall2021/entries/fitness/>.

- Rowlands, Mark, Joe Lau, and Max Deutsch. 2020. "Externalism About the Mind." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Winter.
<https://plato.stanford.edu/archives/win2020/entries/content-externalism/>.
- Russell, Bertrand. 1911. "Knowledge by Acquaintance and Knowledge by Description." *Proceedings of the Aristotelian Society* 11 (1): 108–128.
- Russell, Bertrand. 1912/1998. *The Problems of Philosophy*. Oxford: Oxford University Press.
- Ryder, Dan. 2020. "Problems of Representation I: Nature and Role." In *The Routledge Companion to Philosophy of Psychology*, edited by Sarah Robins, John Symons and Paco Calvo, 233–250. Oxon: Routledge.
- Ryder, Dan. 2020. "Problems of Representation II: Naturalizing Content." In *The Routledge Companion to Philosophy of Psychology*, edited by Sarah Robins, John Symons and Paco Calvo, 251–279. Oxon: Routledge.
- Ryder, Dan, Kingsbury Justine, and Kenneth Williford. 2013. "Introduction." In *Millikan and Her Critics*, edited by Dan Ryder, Kingsbury Justine and Kenneth Williford, 1–20. Chichester: Wiley-Blackwell.
- Sacchi, Elisabetta. 2006. "Fregean Propositions and their Graspability." *Grazer Philosophische Studien* 72 (1): 73–94.
- Schantz, Richard. 2004. "Introduction." In *The Externalist Challenge*, edited by Richard Schantz, 1–33. Berlin: Walter de Gruyter.
- Schellenberg, Susanna. 2014. "The Relational and Representational Character of Perceptual Experience." In *Does Perception Have Content?*, edited by Berit Brogaard, 199–219. New York: Oxford University Press.
- Schellenberg, Susanna. 2018. *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press.
- Scholl, Brian J. 2001. "Objects and Attention: The State of the Art." *Cognition* 80 (1–2): 1–46.
- Schulte, Peter. 2022. "Constancy Mechanisms and Distal Content: a Reply to Garson." *The Philosophical Quarterly* 72 (1): 229–237.
- Schulte, Peter. 2023. *Mental Content*. Cambridge: Cambridge University Press.
- Schulte, Peter. 2018. "Perceiving the World Outside: How to Solve the Distality Problem for Informational Teleosemantics." *The Philosophical Quarterly* 68 (271): 349–369.
- Schulte, Peter. 2015. "Perceptual Representations: A Teleosemantic Answer to the Breadth-of-Application Problem." *Biology & Philosophy* 30 (1): 119–136.
- Schulte, Peter. 2021. "The Nature of Perceptual Constancies." *Philosophy and Phenomenological Research* 103 (1): 3–20.
- Schulte, Peter, and Karen Neander. 2022. "Teleological Theories of Mental Content." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Summer.
<https://plato.stanford.edu/archives/sum2022/entries/content-teleological/>.

- Schumm, Bruce A. 2004. *Deep Down Things: The Breathtaking Beauty of Particle Physics*. . Baltimore: Johns Hopkins University Press.
- Searle, John R. 2018. "Are there Non-Propositional Intentional States?" In *Non-Propositional Intentionality*, edited by Alex Grzankowski & Michelle Montague, 259-271. Oxford: Oxford University Press.
- Searle, John R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. New York: Cambridge University Press.
- Sendłak, Maciej. 2018. "On Quantitative and Qualitative Parsimony." *Metaphilosophy* 49 (1-2): 153-166.
- Sharma, Chandradhar. 1962. *A Critical Survey of Indian Philosophy*. USA: Barnes & Noble, Inc.
- Shastri, D. N. 1976. *The Philosophy of Nyaya-Vaisesika and its Conflict with the Buddhist Dinnaga School: Critique of Indian Realism*. Varanasi: Bharatiya Vidya Prakashan.
- Shea, Nicholas. 2013b. "Millikan's Isomorphism Requirement." In *Millikan and Her Critics*, edited by Dan Ryder, Kingsbury Justine and Kenneth Williford, 63–80. Chichester: Wiley-Blackwell.
- Shea, Nicholas. 2013a. "Naturalising Representational Content." *Philosophy Compass* 8 (5): 496-509.
- Shea, Nicholas. 2018. *Representation in Cognitive Science*. Oxford: Oxford University Press.
- Shields, Christopher. 2013. *Aristotle*. 2nd. New York: Routledge.
- Shields, Christopher. 2020. "Aristotle's Psychology." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. Winter.
<https://plato.stanford.edu/archives/win2020/entries/aristotle-psychology/>.
- Shields, Christopher. 2019. "Hylomorphic Offices." *Ancient Philosophy Today* 1 (2): 215-236.
- Shoemaker, Sydney. 2007. *Physical Realization*. Oxford: Oxford University Press.
- Shoemaker, Sydney. 2001. "Realization and Mental Causation." In *Physicalism and its Discontents*, edited by Carl Gillett and Barry Loewer, 74-98. New York: Cambridge University Press.
- Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- Siegel, Susanna. 2006. "Subject and Object in the Contents of Visual Experience." *The Philosophical Review* 115 (3): 355-388.
- Siegel, Susanna. 2024. "The Contents of Perception." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta and Uri Nodelman. Fall.
<https://plato.stanford.edu/archives/fall2024/entries/perception-contents/>.
- Siegel, Susanna. 2010. *The Contents of Visual Experience*. New York: Oxford University Press.
- Simons, D. J. 2000. "Current Approaches to Change Blindness." *Visual Cognition* 7 (1-3): 1-15.
- Simons, D. J., and D. T. Levin. 1997. "Change Blindness." *Trends in Cognitive Sciences* 1 (7): 261-267.

- Simons, D. J., and S. R. Mitroff. 2001. "The Role of Expectations in Change Detection and Attentional Capture." In *Vision and Attention*, edited by Michael Jenkin and Laurence Harris, 189-207. New York: Springer.
- Simons, D. J., S. L. Franconeri, and R. L. Reimer. 2000. "Change Blindness in the Absence of a Visual Disruption." *Perception* 29 (10): 1143-1154.
- Simpson, William M. R. 2023. *Hylomorphism*. Cambridge: Cambridge University Press.
- Smith, A.D. 2002. *The Problem of Perception*. Cambridge, MA: Harvard University Press.
- Smith, Barry, and David Woodruff Smith. 1995. "Introduction." In *The Cambridge Companion to Husserl*, edited by Barry Smith and David Woodruff Smith. New York: Cambridge University Press.
- Smith, David Woodruff, and Ronald McIntyre. 1982. *Husserl and Intentionality: A Study of Mind, Meaning, and Language*. Dordrecht: D. Reidel Publishing Company.
- Soteriou, Matthew. 2020. "The Disjunctive Theory of Perception." *Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta.
<https://plato.stanford.edu/archives/sum2020/entries/perception-disjunctive/>.
- Speaks, Jeff. 2015. *The Phenomenal and the Representational*. New York: Oxford University Press.
- Springle, Alison A., and Alessandra Buccella. 2024. "Practical Perceptual Representations: A Contemporary Defense of an Old Idea." *Synthese* 203: 1-18.
- Stevens, Martin. 2013. *Sensory Ecology, Behaviour, and Evolution*. Oxford: Oxford University Press.
- Stich, Stephen. 1978. "Beliefs and Subdoxastic States." *Philosophy of Science* 45 (4): 499-518.
- Stich, Stephen. 1992. "What Is a Theory of Mental Representation?" *Mind* 101 (402): 243-261.
- Stoljar, Daniel. 2024. "Physicalism." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman.
<https://plato.stanford.edu/archives/spr2024/entries/physicalism/>.
- Strevens, Michael. 2025. "Grasp as a Universal Requirement for Understanding." *Philosophical Studies* 1-19.
- Studtmann, Paul. 2024. "Aristotle's Categories." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta & Uri Nodelman. Winter.
<https://plato.stanford.edu/archives/win2024/entries/aristotle-categories/>.
- Tahko, Tuomas E., ed. 2011. *Contemporary Aristotelian Metaphysics*. Cambridge: Cambridge University Press.
- Tahko, Tuomas E., and E. J. Lowe. 2025. "Ontological Dependence." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta and Uri Nodelman.
<https://plato.stanford.edu/archives/sum2025/entries/dependence-ontological/>.
- Thagard, Paul. 2005. *Mind: Introduction to Cognitive Science*. Cambridge, MA: The MIT Press.

- Thompson, Evan. 2007. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, MA: Harvard University Press.
- Tugby, Matthew. 2022. *Putting Properties First: A Platonic Metaphysics for Natural Modality*. Oxford : Oxford University Press.
- Turing, Alan. 1936. "On Computable Numbers, with an Application to the Entscheidungsproblem." *Proceedings of the London Mathematical Society* 42 (1): 230-265.
- Twardowski, Kasimir. 1894/1977. *On the Content and Object of Presentations: A Psychological Investigation*. Translated by R. Grossmann. The Hague: M. Nijhoff.
- Tye, Michael. 2009a. *Consciousness Revisited: Materialism without Phenomenal Concepts*. Cambridge, MA: MIT Press.
- Tye, Michael. 2009c. "Representational Theories of Consciousness." In *The Oxford Handbook of Philosophy of Mind*, edited by Brian McLaughlin, Ansgar Beckermann and Sven Walter, 253-267. New York: Oxford University Press.
- Tye, Michael. 2009b. "The Admissible Contents of Visual Experience." *Philosophical Quarterly* 59 (236): 541-562.
- Tye, Michael. 2021. *Vagueness and the Evolution of Consciousness: Through the Looking Glass*. Oxford: Oxford University Press.
- Tye, Michael. 2014. "What is the Content of a Hallucinatory Experience?" In *Does Perception Have Content?*, edited by Berit Brogaard, 291–330. Oxford: Oxford University Press.
- Vaidya, Anand Jayprakash. 2020. "Multi-Factor Causal Disjunctivism: a Nyāya-Informed Account of Perceptual Disjunctivism." *Sophia* 60 (4): 917-940.
- Van Gulick, Robert. 2001. "Reduction, Emergence and Other Recent Options on the Mind/Body Problem: A Philosophic Overview." *Journal of Consciousness Studies* 8 (9-10): 1-34.
- Varela, Francisco J., Evan Thompson, and Eleanor Rosch. 1991/2016. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: The MIT Press.
- Von Eckardt, Barbara. 2012. "The Representational Theory of Mind." In *The Cambridge Handbook of Cognitive Science*, edited by Keith Frankish and William Ramsey, 29-49. Cambridge: Cambridge University Press.
- Wagemans, J., J. H. Elder, M. Kubovy, S. E. Palmer, M. A. Peterson, M. Singh, and R. von der Heydt. 2012. "A Century of Gestalt Psychology in Visual Perception: I. Perceptual Grouping and Figure-Ground Organization." *Psychological Bulletin* 138 (6): 1172–1217.
- Walsh, D. M. 1996. "Fitness and Function." *British Journal for the Philosophy of Science* 47 (4): 553-574.
- Werner, Sauer. 2017. "Brentano's Reism." In *The Routledge Handbook of Franz Brentano and the Brentano School*, edited by Uriah Kriegel, 133-143. New York: Routledge.
- Wilczek, Frank. 2006. *Fantastic Realities: 49 Mind Journeys and a Trip to Stockholm*. Singapore: World Scientific.

- Williams, D. C. 1953. "The Elements of Being." *The Review of Metaphysics* 7 (1-2): 3-18; 171-92.
- Williams, Neil E. 2019. *The Powers Metaphysic*. Oxford: Oxford University Press.
- Wilson, Jessica M. (forthcoming). n.d. "On the notion of diachronic emergence." In *Rethinking Emergence*, edited by Amanda Bryant and David Yates. Oxford: Oxford University Press.
- Wilson, Jessica M. 2021. *Metaphysical Emergence*. Oxford: Oxford University Press.
- Wu, Wayne. 2017. "Attention and Perception: A Necessary Connection?" In *Current Controversies in Philosophy of Perception*, edited by Bence Nanay, 148-162. New York: Routledge.
- Yates, David. 2016. "Is Powerful Causation an Internal Relation?" In *The Metaphysics of Relations*, edited by Anna Marmodoro & David Yates, 138-156. Oxford: Oxford University Press.
- Yong, Ed. 2022. *An Immense World: How Animal Senses Reveal the Hidden Realms Around Us*. London: The Bodley Head.
- Zahavi, Dan. 2017. *Husserl's Legacy: Phenomenology, Metaphysics, and Transcendental Philosophy*. Oxford: Oxford University Press.