Durham E-Theses

Exploring the Initial Mass Function by Stochastically Lighting Up Galaxies

GREG ASHWORTH

How to cite:

ASHWORTH, GREG (2018) Exploring the Initial Mass Function by Stochastically Lighting Up Galaxies. Doctoral thesis, Durham University.

Use policy

Exploring the Initial Mass Function
by
Stochastically Lighting Up Galaxies

Greg Ashworth

Abstract

In this thesis, the Initial Mass Function (IMF) is studied using the Stochastically Lighting Up Galaxies software suite (SLUG), a package of tools including a stochastic Stellar Population Synthesis (SPS) code and associated analysis packages, including a novel Bayesian inference framework.

Following an introduction to some core concepts, new developments of the SLUG code are described. These include a variable IMF capability which is then applied to broad-band photometry taken from the Legacy ExtraGalactic Ultraviolet Survey (LEGUS), a *Hubble Space Telescope* treasury programme. The physical parameters of star clusters in galaxy NGC 628 are inferred using SLUG's Bayesian inference tools. We find that the posterior probability distributions of the high-mass slope of the IMF are very broad, and we quantify a degeneracy between the IMF and the cluster mass. The inclusion of additional photometric data ($H\alpha$) is found to provide some improvement. However, using mock cluster models we found that only through constraining the mass of the cluster through photometrically-independent means is it possible to accurately recover the IMF slope.

An additional source of information is the UV spectrum, which is dominated by the massive stars whose populations are affected by the high-mass region of the IMF. To be able to exploit this region of the spectrum effectively using equivalent width measurements, the resolution of SLUG's UV spectral synthesiser requires improvement. To this end, the implementation of a high-resolution UV synthesiser is described, and then put to use in a theoretical study of the IMF using mock observations generated with SLUG. The constraining power of UV spectral features when combined with broad-band photometry is quantified, resulting in significant improvement in IMF slope recovery.

Finally the results and limitations of the studies are discussed, and recommendations are made for future studies and improvements.

Exploring the Initial Mass Function by Stochastically Lighting Up Galaxies

by Greg Ashworth

A thesis submitted to Durham University

in accordance with the regulations for

admittance to the Degree of Doctor of Philosophy.

Institute for Computational Cosmology

& Centre for Extragalactic Astronomy

Department of Physics

Durham University

June 2018

CONTENTS

LIST OF FIGURES

LIST OF TABLES

Declaration

The work described in this thesis was undertaken while the author was a research student under the supervision of Dr. Michele Fumagalli and Dr. John Lucey in the Department of Physics at Durham University, during the author's PhD course (October 2014-June 2018). This work has not been submitted for any other degree at Durham University or any other university.

Chapters 3 and 4 have been submitted for publication in the form of papers to Monthly Notices of the Royal Astronomical Society (MNRAS):

- *Exploring the IMF of star clusters: a joint SLUG and LEGUS effort*
  G. Ashworth, M. Fumagalli, M. R. Krumholz, A. Adamo, D. Calzetti, R. Chandar, M. Cignoni, D. Dale, B. G. Elmegreen, J.S. Gallagher III, D. A. Gouliermis, K. Grasha, E. K. Grebel, K. E. Johnson, J. Lee, M. Tosi, A. Wofford, *MNRAS* 2017, Volume 469, Issue 2, pp. 2464-2480, DOI:10.1093/MNRAS/STX935

- *Theoretical predictions for IMF diagnostics in UV spectroscopy of star clusters*
  G. Ashworth, M. Fumagalli, A. Adamo, M. R. Krumholz, *MNRAS* 2018, (Submitted)

All figures in this thesis have been produced by the author, and all analysis contained herein was performed by the author. During this work the author contributed code to the open source Stochastically Lighting Up Galaxies stellar population synthesis software, described in Chapter 2. This took the form of a new variable Initial Mass Function capability and a new high resolution UV spectral synthesiser, containing an equivalent width calculation capability. The author's contributions are detailed at https://bitbucket.org/krumholz/slug2.

The copyright of this thesis rests with the author. No quotation from it should be published without prior written consent and information derived from it should be acknowledged.

"Reality is frequently inaccurate."

- Douglas Adams, The Restaurant At The End Of The Universe (1980)

Acknowledgements

I would like to begin by thanking my supervisor Michele Fumagalli, for his endless reserves of knowledge, advice, and patience. In addition, I must thank Mark Krumholz for his helpful advice and suggestions on dealing with SLUG, Angela Adamo for her detailed comments on my work, my secondary supervisor John Lucey, and the COSMA team who have helped make all of my simulation work possible. I would also like to thank the LEGUS collaboration for making me feel welcome, and for helping me feel like part of something important.

Thanks also go to the many friends and officemates who I have had the pleasure of meeting, from the original 302 team of James, Emma, and Ady, through 101's Ra'ad, Raj, and Daniel, to 317's Steph and Ruari. There are many others to list, but limited room in this paragraph, so particular thanks go to Helen, Piotr, Hannah, Andrew, Tim, and Frank. I'd also like to thank John for keeping me busy with the world of miniature wargaming when I needed it the most.

Most importantly, thanks go to Mum (who will forever be missed), Dad, the family, and finally Kate. Without them, I would not have made it this far.

I hope this thesis scores well in the only categories that matter (as made famous by Robot Wars): Style, Control, Damage, and Aggression.

CHAPTER 1

INTRODUCTION

1.1 The stellar Initial Mass Function

The mass of a star is the property that is most dominant in defining how that star will evolve over the course of its existence (Scalo, 1986). Understanding how the masses of stars are distributed at the time of their formation is therefore the key to understanding many astrophysical phenomena. This distribution goes on to determine the stellar content of star clusters, and therefore the chemical makeup of these populations. This in turn goes on to determine the light and stellar evolution of not only these clusters but also their host galaxies (Chabrier, 2005). The form of this distribution, and indeed whether it is constant in space or time (or both), is a long standing unsolved problem in astrophysics, which underpins the vast majority of galaxy evolution theoretical and observational work. This is complicated by the fact that there is no direct method of determining this distribution from observations (Chabrier, 2003).

This initial distribution of stellar masses, or the stellar Initial Mass Function (IMF), was first codified in 1955 by Edwin Salpeter (Salpeter, 1955). The IMF presented in that seminal paper was derived from the present day luminosity function (distribution of stars by their luminosity), itself obtained from direct star number counts. This first IMF was a single power law of the form:

$$\frac{dN}{dm} \propto m^{\alpha},$$ (1.1)

where N is the number of stars, m is the stellar mass, and α is the power law index. In Salpeter's original paper, this index is calculated to be $\alpha = -2.35$. The IMF is also often seen in its logarithmic form:

$$\frac{dN}{d\log m} \propto m^{-\Gamma}, \tag{1.2}$$

where $\Gamma = 1.35$ as given in Salpeter (1955). During this thesis, we will use the linear convention.

Since Salpeter's seminal paper, the shape of the IMF has been under constant study and refinement, and many (similar) functional forms for the IMF have been proposed. It is now generally understood that the IMF does not constitute a constant power law across all stellar masses, and that there is some flattening as the stellar mass moves downwards towards the Hydrogen burning limit. This flattening of the IMF was first proposed in the the work of Miller & Scalo (1979) who suggested that the IMF took on a log-normal form for low mass stars.

At the present time, two of the most commonly used forms in the literature are the broken power-law form (Kroupa et al., 1993; Kroupa, 2001) commonly referred to as the "Kroupa" IMF, which splits the IMF into multiple segments, each with a power law slope. For example, in Kroupa (2001) this is given as:

$$\xi(m) \propto m^{-\alpha_i}, \tag{1.3}$$

where $\xi(m)$ is the IMF (with $\xi(m)\mathrm{d}m$ representing the number of stars in the interval m to $m + \mathrm{d}m$). The slope α is defined as follows:

- $\alpha_0 = 0.3 \pm 0.7,\ 0.01 \le m/\mathrm{M}_\odot < 0.08$,

- $\alpha_1 = 1.3 \pm 0.5,\ 0.08 \le m/\mathrm{M}_\odot < 0.50$,

- $\alpha_2 = 2.3 \pm 0.3,\ 0.50 \le m/\mathrm{M}_\odot < 1.00$,

- $\alpha_3 = 2.3 \pm 0.7,\ 1.00 \le m/\mathrm{M}_\odot$.

The other commonly used form of the IMF is that of Chabrier (Chabrier, 2003), where the IMF takes on a log-normal form for lower stellar masses, based on the earlier work of Miller & Scalo (1979) and Adams & Fatuzzo (1996). This is suplemented by a power law slope as in the Salpeter form past a certain cutoff mass. Maschberger (2013) gives the Chabrier IMF as:

- $P_{\text{chab}} = A k_1 \frac{1}{m} e^{-\frac{1}{2}\left(\frac{\log m - \log 0.079}{0.69}\right)^2}, \, m < 1\text{M}_\odot$

- $P_{\text{chab}} = A k_2 m^{-2.3}, \, m > 1\text{M}_\odot$

where P_{chab} is the IMF, $k_1 = 0.158$, $k_2 = 0.0443$, and A is a global normalisation constant.

Both of these forms have a similar high-mass end slope to the original Salpeter IMF, with each having differing forms at the low-mass end. Where these forms of the IMF change from the scale-free "Salpeter-like" regime to a more shallow structure at low masses sets a characteristic mass-scale, usually around $1\,\text{M}_\odot$.

Whichever form it takes, the IMF has generally been considered to be universal, constant across different environments and across time. However, this assumption is not made through any theoretical understanding, but rather through there being no fully convincing evidence that the IMF varies (Massey et al., 1995; Bastian et al., 2010). However, as Kroupa (2001) states, there are fundamental reasons why the IMF should, in fact, vary. For example, the IMF should become top-heavy in certain conditions. In star-forming environments with low metallicity gas, cooling is less efficient, and therefore larger stellar masses form from this warmer gas. Indeed, Larson (1998) suggests that the IMF should vary with time, with high-temperature primordial gas clouds leading to much more massive stars being formed due to the corresponding increase in Jeans mass. On the other hand, Adams & Fatuzzo (1996) argue that stars regulate their own masses through winds and outflows, and therefore determine the form of the IMF. They postulate that the initial conditions in a molecular cloud can be "transformed" to produce the resultant IMF for that cloud, with the IMF approaching a log-normal form as the number of parameters that describe the molecular cloud approaches infinity through the central limit theorem.

There are further questions, such as whether star formation is a stochastic process or whether it is a purely deterministic one. We must also question whether the form of the IMF is correlated with the star forming region's properties. One of the more widely known examples of work which does link environment with the IMF's form in a deterministic way is the Integrated Galactic Initial Mass Function (IGIMF) theory described by Kroupa & Weidner (2003), which was later expanded upon in further work (e.g. Weidner & Kroupa, 2005; Weidner et al., 2011). In this model, the maximum stellar mass that can be formed in a star cluster is intrinsically linked with the mass of the cluster itself. This therefore sets an upper limit for

the IMF of the star cluster. The IMF for a composite stellar population (the galaxy) is then a combination of the IMFs of the simple stellar populations (and field stars) that make it up. In this IGIMF model, the IMF is then sampled based on a 'sorted' sampling method (Weidner & Kroupa, 2006), where the stars in the population are assembled from the lowest mass to the highest, thereby causing the cutoff of that population's IMF.

There are many studies where a variable IMF has been invoked to explain observations. However, there are also many cases where this evidence for IMF variability has been subsequently shown to be possible to replicate with a non-varying IMF. There are also studies where there is no convincing argument either way. Attempting to quantify any variability in the IMF is a difficult process. For example, Bastian et al. (2010) details several challenges to measuring the IMF successfully in the Milky Way. The first of these challenges is that many studies of the IMF (even those wishing to search for variability in the IMF) make the assumption that the IMF is constant in either space or time, or even both of these. The spatial assumption comes in when using star counts over large (or even spatially separated) volumes. This requires the assumption that the IMF does not vary in space. The assumption that the IMF is constant in time is invoked when calculating the IMF from a measured Present Day Mass Function (PDMF), as correcting for the losses of high mass stars requires that the IMF is assumed to not vary in time. This reconstruction also requires knowledge of the star formation history of our Galaxy, knowledge that is also required to understand the low-mass end of the IMF if the IMF varies with time. This is due to the fact that the present day mass function will be missing any massive stars that have reached the end of their lives, and indeed there could be several generations of these. Binary systems also complicate IMF determination through the additional luminosity added by unresolved companion stars. Correcting for this effect is difficult due to sensitivity to the fraction of binary stars and the ratio of masses between the stars and their binary companions. Also required is knowledge of the mass dependence of the fraction of binary systems and the distribution of mass ratios across the sample. There are also statistical considerations when studying the IMF. Many IMF studies determine the IMF by fitting a functional form to a dataset and compare this fit to those of other studies. Bastian et al. (2010) claims that disagreement between these parameters is usually seen as evidence for a different IMF, which is complicated further by the fact that many

studies do not provide uncertainties on their fits or rigorous justification for their choice of functional form. In addition, the range of masses sampled during a study along with how the data is binned can both affect the IMF determination, with samples drawn from the same IMF resulting in different fits depending on these two choices.

Nevertheless several recent studies have supported the idea of a variable IMF, a few examples of which we will summarise here. Recently, the work of van Dokkum & Conroy (2010) detailed the detection of an abundance of low-mass stars in early-type elliptical galaxies, suggesting a bottom-heavy IMF. This evidence came from measurements of the strength of distinctive absorption lines (namely the NaI doublet and Wing-Ford molecular FeH band). This was followed up with similar detections by Conroy & van Dokkum (2012a) where observations of absorption line features (including the IMF-sensitive NaI, CaII, and FeH features) were compared to stellar population synthesis (SPS, see Section 1.2) models to derive an IMF that became increasingly bottom-heavy as the velocity dispersion of the parent system increased. This leads to an IMF compatible with that of the Milky Way for low velocity dispersions, but results in an IMF with an extreme low-mass slope (higher than that of the Salpeter IMF) for high velocity dispersion systems. In a later study, Conroy et al. (2013) once again found a bottom-heavy IMF in early type galaxies using a combination of spectral lines and SPS models. However, in this study the result was backed up by simple dynamical modelling, which agreed well with the spectral analysis, resulting in an IMF that becomes more bottom-heavy as velocity dispersion of the parent galaxy increases. Dynamical measurements were also used by Cappellari et al. (2012) to infer a systematic variation in the IMF with the mass-to-light ratios of early type galaxies. These are just a handful of the studies which suggest variation in the IMF for elliptical galaxies. Other examples include Auger et al. (2010); Smith et al. (2012); Spiniello et al. (2012); La Barbera et al. (2013) and more recently Conroy et al. (2017) and Li et al. (2017), amongst others. However, the Smith (2014) study, comparing the work of Conroy & van Dokkum (2012b) with that of Cappellari et al. (2013) suggested that the results of these spectroscopic and dynamical studies of the IMF in early-type galaxies were not always in agreement when considering the same sample of galaxies. They conclude that either one or both of the methods do not take into account all possible limiting factors, or that variation in the IMF's shape may be beyond what

current spectroscopic techniques are able to infer. More specifically, the Conroy et al. (2013) study using a spectroscopic method finds correlations only with the metallicity, whereas the Cappellari et al. (2012) study uses dynamical methods and finds correlation only with dynamics. This may lead us to conclude that there may be residual problems with both methods. The results of these studies therefore remain inconclusive, although the more recent work of Lyubenova et al. (2016) did not find such a disagreement to be apparent on a galaxy-by-galaxy basis when using consistent datasets.

Moving away from elliptical galaxies, correlations between the Hα-FUV ratio and the star formation rate in dwarf galaxies have been observed (Lee et al., 2009, 2016). These correlations could be explained by variations in the high-mass end of the IMF. Meurer et al. (2009) make similar comparisons between the Hα-to-FUV luminosity ratio and Hα and R-band surface brightness in starburst galaxies, once again finding that IMF variation could explain their observations. Gunawardhana et al. (2011) find indications of IMF variation using comparisons between Hα equivalent width and galaxy colours.

The aforementioned IGIMF model has been put forward to explain observations such as these. For example, in dwarf galaxies, the mass dependent cutoff of the IMF would suppress the formation of massive stars in dwarf galaxies, due to the dependence of the cluster mass on the star formation rate of the system. This suppression of massive star formation would result in the lower Hα-FUV ratios seen in, for example, the Lee et al. (2009) and Meurer et al. (2009) studies. However, Boselli et al. (2009) conclude that a universal IMF of a Salpeter-like slope is sufficient to explain the observations, given the wide distribution in Hα-FUV ratio. Rather than immediately invoking a variable IMF, the authors instead argue that inclination (resulting in inaccurate dust extinction), presence of an active galactic nucleus, or bursty star formation histories may also be possible explanations for the observed behaviour of the Hα-FUV ratio. The bursty star formation history angle was expanded upon by Fumagalli et al. (2011), where stochastic SPS modelling (see Section 1.3) was used in conjunction with the Lee et al. (2009), Meurer et al. (2009), and Boselli et al. (2009) observations. The authors of that study concluded that there was no need to invoke the IGIMF as an explanation for the observed Hα-FUV ratio behaviour. They found instead that stochastic sampling of a Kroupa IMF across star clusters was sufficient to explain the observed Hα luminosities, and that

realisations generated using a truncated IGIMF underpredicted these Hα luminosities. This study also showed that the correlations between the star formation rate, cluster mass function, and IMF all arise naturally from stochastic sampling, rendering some of the modifications to the IMF introduced as part of the IGIMF theory unnecessary.

Other works have explored the relationship between Hα and cluster mass, such as the studies of Calzetti et al. (2010) and Andrews et al. (2013), finding no clear mass dependence on the high-mass end slope of the IMF. Corbelli et al. (2009) also find that a randomly sampled IMF is consistent with their observations of the ratio between Hα luminosity and bolometric luminosity in M33. The truncation of the IMF proposed by the IGIMF theory is likewise challenged by Andrews et al. (2014).

The observation of lonely massive stars would support the idea of a stochastically sampled IMF rather than the mass-dependent behaviour of the IGIMF, and indeed several studies have been published detailing such observations (examples include: Bressert et al., 2012; Oey et al., 2013; Lamb et al., 2016; de Wit et al., 2004, 2005; Lamb et al., 2010). However, these observations are not without controversy, with some of these results being disputed by Gvaramadze et al. (2012) through a search for bow shocks near massive OB stars that were thought to be isolated, finding that most are instead likely to be runaways. Instead they conclude that there is no significant evidence that supports the idea that massive stars may form in isolation. In addition, Stephens et al. (2017) present observations of compact clusters of pre-main-sequence stars around what were once considered isolated massive stars, in opposition to the idea of stochastic IMF sampling. In fact, one does not need the massive stars to be totally isolated for the system to be incompatible with the IGIMF, which requires that the stellar mass distribution be smooth up to the maximum value for a given cluster mass. An absence of, for example, B stars around an O star, with the next most massive being an A type star, would also be incompatible with the IGIMF. From even this brief summary, we can see that there is no general consensus on the variability of the IMF, with almost every study in support of one viewpoint responded to from the other side. There is disagreement on everything from whether the IMF varies or not, to whether star formation is a stochastic process or not. Cerviño (2013) on the other hand takes the view that both the stochastic and deterministic methods of modelling a stellar population are in fact equivalent, and the strengths of each method can be used to counter the

weakness of the other when studying the inherently stochastic nature of stellar populations.

It is probable that new tools and techniques will be required if we are to truly understand whether the IMF is universal or otherwise. This work attempts to move towards this through the development of novel ways to infer the form of the IMF from observations. Further information on stochasticity in stellar population modelling will be given in Section 1.3.1.

1.2 Stellar population synthesis

Stellar population synthesis is a useful tool that helps us answer a fundamental question in astrophysics: how do we know what the properties of a stellar population are if we cannot do more than observe its light?

The concept of stellar populations has been present in the literature for some time, with Cerviño (2013) citing Baade (1944) as originating this concept. However, it was not until Renzini (1981) that the concept of the Simple Stellar Population (or SSP) was truly defined. These SSPs are groups of stars for whom there was a single star formation event, with all stars formed at the same time. For such a coeval group of stars, a cluster for example, we can construct a prediction for its light output by effectively integrating over all of the stars that make it up. Beginning with an IMF $\left(\frac{\mathrm{d}n}{\mathrm{d}m}\right)$, which dictates the distribution of stars by mass, and the necessary information to provide a specific luminosity for a star of a given mass and age, $L(m,t)$, we can make a prediction for the specific luminosity, L_v of the population:

$$L_v = M_* \int_0^\infty L_v(m,t) \frac{\mathrm{d}n}{\mathrm{d}m} \mathrm{d}m, \tag{1.4}$$

where M_* is the mass of the SSP. We can use the same technique for objects like galaxies (or Composite Stellar Populations, CSPs), which are composed of stars with a distribution of ages, which itself is described by a star formation history (SFH). Working from the same principles as the SSP (an IMF and a set of stellar luminosities), we now include the star formation history of the population, $\dot{M}_*(t)$, which is the rate of star formation as a function of time t. This results in the specific

luminosity prediction for the CSP:

$$L_\nu = \int_0^\infty \dot{M}_*(t)\,\mathrm{d}t \int_0^\infty L_\nu(m,t)\frac{\mathrm{d}n}{\mathrm{d}m}\mathrm{d}m. \qquad (1.5)$$

Most SPS codes perform essentially these operations, and many have been released over the years. One of the more popular examples of this sort of deterministic SPS code is STARBURST99 (Leitherer et al., 1999).

As can be seen in Equation 1.4, to produce a spectrum for a stellar population modelled with these methods, we need to know the luminosity of the stars that make up the population. To do this we need to know the masses of the stars that make up the stellar population, along with their spectra, at a given time.

First we will discuss how we calculate the properties of the stars. SPS codes use libraries of stellar tracks, which describe how stars with particular properties evolve in time. In the case of STARBURST99, the tracks are given for particular stellar masses, with each track detailing how the stars' parameters (temperature, for example) evolve with time (Leitherer et al., 1999). In many (mostly older) SPS codes, stars are binned such that each group evolves along a particular track (Charlot & Bruzual, 1991). This results in an isochrone, with values at each mass bin. To produce an output spectrum for the SSP, we need to know what stars at each of the points along this isochrone look like. For this reason, SPS codes also provide atmosphere models. These atmosphere models, in the case of STARBURST99 for example, give flux values for stars as a function of surface gravity and effective temperature across the spectrum. Such libraries of atmospheres may either be constructed empirically from observations, or they may be purely theoretical. For example, STARBURST99 contains empirical libraries of atmospheres for O and B type stars created from various observations (Robert et al., 1993; de Mello et al., 2000; Leitherer et al., 2001; Robert et al., 2003), while also containing theoretical spectra for these stellar types (Leitherer et al., 2010). Having assigned a spectrum to each point along the isochrone (either through interpolation or nearest neighbour selection), we can then combine these atmospheres in a sum weighted by the distribution of masses across the SSP (the IMF) which takes the place of the integral in Equation 1.4.

However, this method is not ideal as it is limited by the number of stellar tracks available, and can result in discontinuous behaviour in the calculated light

output of the population as stars move between the bins as they evolve in time. To overcome this undesired behaviour, Charlot & Bruzual (1991) describe a method whereby the track data is interpolated to give isochrones that are continuous in stellar mass. The number of stars of each mass are then determined from the IMF either deterministically or stochastically (see Section 1.3.1).

These deterministic methods however have the effect of including "fractional" high-mass stars in the population if the IMF is assumed to be fully sampled and the population is small.

Tracks and atmospheres are usually available at a range of metallicities, either selected discretely or with the ability to interpolate between the provided values. From a given SSP spectrum, photometry may be created by integrating across a given filter with a given response function. In many SPS codes, the resulting light output may also undergo extinction by dust, using a given extinction curve, and may also include nebular emission. To produce a CSP spectrum (as in the case of a galaxy), one simply needs to combine the appropriate SSPs as in Equation 1.5, weighted by a cluster mass function, summing up to the galaxy's target mass.

1.3 Stochastically Lighting Up Galaxies

1.3.1 Stochastic sampling in SPS

The majority of SPS codes available work as described in 1.2, using a deterministic methods to produce fixed observables for a given input. However, this is not the only available approach. The alternative is to take a stochastic approach to modelling stellar populations. This approach is useful for several reasons. Firstly, the traditional deterministic method of SPS predicts only the mean light output of a given set of parameters defining a stellar population. However, it is possible that there is a large amount of scatter around this mean and that individual stellar populations from a large family of realisations may have observably different luminosities in a given band. This scatter is most evident for small stellar populations, and for these systems the output of a deterministic SPS model may be a poor fit due to the assumption in the deterministic case that the IMF is well sampled (da Silva et al., 2012). Undersampling of the star formation history (SFH) in low star formation rate (SFR) systems can also lead to failure in deterministic models.

At the time of SLUG's initial release, no single general purpose tool existed which allowed exploration of stochastic star formation from scales ranging from a single cluster up to an entire galaxy (da Silva et al., 2012). However, it was not the first SPS model to make use of stochasticity. Other notable stochastic SPS codes include MASSCLEAN (Popescu & Hanson, 2009), which is also capable of producing images in addition to photometry, and the work of Anders et al. (2013) who use the GALEV code. However, the inclusion of stochastic effects when modelling stellar populations goes back many years, with studies such as that presented by Barbaro & Bertelli (1977) using Monte-Carlo techniques to account for fluctuations in colour for systems with few massive stars.

The most common form of stochasticity introduced into SPS models is the stochastic sampling of the IMF. For example, take a simple stellar population. Instead of integrating over an IMF as we would in the deterministic method, we choose to draw individual stars from a probability distribution function (PDF) that has the same form as the IMF. This results in a distribution of output values for a given set if input parameters, rather than the single value obtained from the deterministic method. There are many ways to handle drawing stars from the IMF and each affect the resulting cluster in different ways. We will discuss these sampling methods in more detail in Section 1.4.2.

The traditional deterministic approach to SPS results in photometric measurements which are effectively a mean value, given a large enough set of populations. We do not therefore see the (potentially large) scatter present for individual realisations of a given set of cluster parameters. These effects are particularly prominent in low mass clusters. As an extreme example, consider a star cluster containing a single star, with a mass of $50\,M_\odot$. This star will not have the same luminosity as a group of 50 stars of $1\,M_\odot$, although both configurations are possible if the distribution of stellar masses is a probability distribution. This *sampling stochasticity* is not captured by the deterministic method of SPS as we do not have information on the shape of the luminosity distribution, only its mean. A consequence of this sampling stochasticity is an undersampling of the high-mass end of the IMF (as low mass clusters may be under-luminous compared to a fully sampled IMF). Likewise, in some cases, a cluster may appear to have a top heavy IMF if several high mass stars form within it. Both cases are possible with a stochastic assembly method, depending on the sampling method used.

In fact, it is possible to estimate the mass limit for stellar populations below which stochastic effects become important. Although the following formalism was initially developed for a Salpeter IMF with no upper limit (Elmegreen, 2000), da Silva et al. (2012) find that a Kroupa IMF with an upper limit does not change the result significantly. The total population mass (M) required to expect a single star above mass m is:

$$M \sim 3 \times 10^3 \left(\frac{m}{100 \mathrm{M}_\odot} \right)^{1.35}. \tag{1.6}$$

Using this relation, to expect the formation of a single star at a mass of $120\,\mathrm{M}_\odot$ (the limit of the tracks we use in SLUG), the stellar population must have a mass of $\sim 10^4 \mathrm{M}_\odot$. In order to be able to ignore stochastic effects (such as apparent IMF truncation), a coeval stellar population must have a mass much greater than this limiting value (which we dub M_{trunc}). Even so, additional scatter can also occur in coeval populations whose masses do exceed this limiting value. The number of stars in luminous, short-lived evolutionary phases of can be greatly affected by the sampling of the appropriate mass range.

As mentioned briefly earlier, undersampling of the SFH may also be a cause for concern when using deterministic models. In the case of a composite stellar population like a galaxy, not only must all the clusters that make it up exceed the limiting mass given in Equation 1.6, the SFH must be well sampled as well. For low SFR composite populations, these conditions may not hold. In much the same way as we defined the limiting mass in Equation 1.6 for stochastic effects from the IMF, we are also able to set a limit on the SFR (da Silva et al., 2012), which is defined as follows:

$$\mathrm{d}t = \frac{M_{\mathrm{trunc}}}{\mathrm{SFR}} \ll 10^6 \ \mathrm{yr}. \tag{1.7}$$

This condition is only met for CSPs forming stars at a rate that is consistently $\gg 10^{-2}\,\mathrm{M}_\odot\,\mathrm{yr}^{-1}$. However, the fact that we have discrete units of star formation in the form of clusters amplifies this stochasticity (and we replace the limiting mass M_{trunc} in Equation 1.7 with the mass of the clusters being drawn from the CMF). This discretisation of the star formation history results in a burst of star formation every time a new cluster is formed, and the gaps between the formation events increase the probability that an observation of a galaxy occurs when many of the more massive stars have already evolved off the main sequence (Fumagalli et al.,

2011).

1.4 What is SLUG?

1.4.1 Introduction

Until the release of SLUG, there was no general purpose tool available to study the photometry of galaxies and star clusters (da Silva et al., 2012), whilst simultaneously handling both the stochastic effects of an undersampled IMF, and the stochastic effects introduced by undersampling of the SFH. Since SLUG's initial release, many additional features have been added which make it a unique and powerful tool for studying stochastic effects in stellar populations.

1.4.2 Stochastically constructing stellar populations

The core building blocks of any SLUG simulation are probability distributions. These PDFs are used to represent the IMF, the Cluster Mass Function (CMF), the Cluster Lifetime Function (CLF), and the Star Formation History (SFH). The PDFs in SLUG can be of arbitrary shape and are constructed of segments. Each segment can have one of several functional forms (such as a power law or a normal distribution), with each functional form having a characteristic set of parameters that define it.

To construct an SSP, SLUG begins with an IMF. This IMF is defined by a PDF of arbitrary shape (as described above), although common IMFs are provided with the software (Kroupa, Salpeter, Chabrier etc...). At the beginning of a realisation, SLUG populates the SSP by drawing stars from the PDF that represents the IMF until the SSP reaches its target mass.

However, as SLUG builds star clusters by drawing stars from the IMF until the cluster is completed, it is probable that we will not hit the requested cluster mass exactly. Therefore, we need some form of sampling condition to define the behaviour of the final draw. In SLUG, several sampling conditions are implemented. The default option, and the one that is used throughout this work, is the "stop nearest" condition. Under this condition, if the final star drawn for a cluster takes that cluster above it's target mass, the star is only kept if the difference between the resulting cluster mass and the target mass is smaller than the difference

between the target mass and the mass of the cluster without this final star. Other sampling options are available, such as the "stop before" (where the last star is always discarded), "stop after" (where the last star is always kept), and finally "sorted sampling" which was first introduced by Weidner & Kroupa (2006). Sorted sampling involves arranging stars by increasing mass such that the cluster is filled to its target mass. This has the additional effect of truncating the IMF, as required by the IGIMF model.

To assemble a CSP, SLUG is able to draw SSP target masses from a cluster mass function, also defined by a PDF. This results in a distribution of cluster masses across a galaxy simulation. In addition, SLUG implements a "clustering fraction", which dictates how many stars form in clusters and how many stars form in the field. These field stars are drawn directly from the IMF.

In the SSP simulation (known as "cluster mode"), all stars are formed in a single starburst event, and the population is evolved in time. In a CSP simulation ("galaxy mode"), a star formation history (which takes the form of another probability distribution) can be requested for the population. Stars in clusters share a formation time drawn from the star formation history (as in a "cluster mode" simulation), whereas field star formation times are individually dictated by the the star formation history such that at a given time, the masses of stars formed respect the clustering fraction. In addition, SLUG can simulate the effects of cluster disruption. At formation, each star cluster may be assigned a maximum lifetime from the Cluster Lifetime Function (CLF), which is itself a PDF. When a cluster reaches the end of its life, it ceases to be tracked as a coherent stellar population, although the stars within it remain part of the galaxy and contribute to its integrated properties in the same way as stars formed in the field.

During this study we make use of SLUG's "cluster mode" only, although the modifications described in Chapter 2 are available to the galaxy mode as well.

1.4.3   Lighting up galaxies

Having constructed a stellar population, simple or composite, we can now "light it up" using SLUG to calculate the light output of this population. Using a similar method to that described in Section 1.2 for deterministic models, SLUG calculates the properties of each star in the realisation by interpolating within a library of

stellar evolutionary tracks. Having obtained the properties of the stars present, atmospheres can then be assigned to each star using the provided atmosphere libraries, depending on the properties of each star (effective temperature and surface gravity). The combined spectrum for the stellar population is then calculated from these atmospheres as luminosity per unit wavelength L_λ. SLUG can also include the effects of nebular emission. This is implemented by assuming that all the stars and star clusters with an age less than 10^6 yr have their ionising photons absorbed by a surrounding HII region. This region is assumed to be in photoionisation equilibrium, and is assumed to consist entirely of fully-ionised hydrogen and singly-ionised helium. This gives us:

$$\phi Q\left(\mathrm{H}^0\right) = \alpha_B\left(T\right) n_\mathrm{e} n_\mathrm{H} V,\tag{1.8}$$

where ϕ is the nebular fraction (which is set to 0.73 by default), which represents the fraction of ionising photons that are absorbed by hydrogen within the aperture, and $Q\left(\mathrm{H}^0\right)$ is the hydrogen ionising luminosity (a quantity that SLUG can output for a given simulation, which can be used to calculate Hα luminosity for example). This quantity is given by:

$$Q\left(\mathrm{H}^0\right) = \int_{h\nu=I(\mathrm{H}^0)}^{\infty} \frac{L_\nu}{h\nu}\mathrm{d}\nu,\tag{1.9}$$

where h is Planck's constant, $I\left(\mathrm{H}^0\right)$ is the ionising potential of neutral Hydrogen, ν is the frequency, and $L_\nu = \lambda^2 L_\lambda/c$ is the luminosity per unit frequency (Krumholz et al., 2015a). $\alpha_B\left(T\right)$ is the case-B recombination coefficient. This is calculated at a temperature which may either be defined by the user or set by the age of the stellar population. n_e is the number density of electrons, n_H is the number density of Hydrogen, and V is the volume.

We can combine these to write the nebular luminosity as:

$$L_{\lambda,\mathrm{neb}} = \gamma_\mathrm{neb} n_\mathrm{e} n_\mathrm{H} V = \gamma_\mathrm{neb}\phi\frac{Q\left(\mathrm{H}^0\right)}{\alpha_B\left(T\right)},\tag{1.10}$$

where γ_neb is the wavelength-dependent nebular emission coefficient which is computed by SLUG, taking into account free-free and bound-free emission (H$^+$ and He$^+$), two-photon emission from neutral Hydrogen in the 2s state, and hydrogen recombination lines. Non-hydrogen lines are approximated using tabulated values.

The process is described in more detail in Appendix B of Krumholz et al. (2015a), and is an approximation of the nebular emission. More detailed modelling can be performed by linking SLUG with the CLOUDY radiative transfer code (Ferland et al., 2013), although that is beyond the scope of this work.

The effects of extinction can also be calculated, using an input extinction curve and either a randomised or fixed A_V (V-band extinction). If a fixed value of A_V is chosen, then the wavelength dependent extinction for all the stellar populations generated is calculated using this fixed value and the input extinction curve. If this value is randomised, it is drawn from a probability distribution such that a unique value is generated for each realisation (be it for a number of realisations of individual clusters, or across the clusters in a single realisation of a galaxy).

Having produced a spectrum for the stellar population, SLUG can then use this to calculate photometry. SLUG is distributed with a library of photometric filters and their associated response functions. To produce photometry is therefore just a matter of integrating the output spectra over any number of these photometric filters. In "galaxy mode", light outputs are available both on a cluster-by-cluster basis and integrated over the entire galaxy.

SLUG is also capable of calculating the yields of various isotopes from core collapse supernovae and Asymptotic Giant Branch (AGB) stars, although this capability is not exploited during this work.

1.4.4 Solving the inverse problem with BAYESPHOT

We have discussed how SLUG can be used to construct stellar populations and calculate their light output. However, the SLUG package contains more than just a stochastic SPS code, and is able to also reconstruct the physical parameters of a stellar population given its light output. First introduced by da Silva et al. (2014) as a tool to infer the star formation rate of a galaxy from integrated light photometry, SLUG's Bayesian inference package BAYESPHOT was generalised in Krumholz et al. (2015a) to be a multipurpose method to infer the physical properties of stellar populations from their photometry. In effect, this tool allows us to solve the reverse of the problem handled by the SLUG SPS code. Rather than answering the question of what a cluster looks like given its parameters, instead we can work out what parameters a cluster has based on what it looks like. Due to the stochastic nature

of SLUG model clusters, traditional fitting methods (for example χ^2 methods) would be unsuitable. The stochastic effects introduced by our modelling means that there is no longer a direct one-to-one mapping between a set of model parameters and an output set of observables, and therefore interpreting an observation of a given cluster in the stochastic regime using fitting methods that do assume a one-to-one mapping is an unsuitable course of action. Instead, it is desirable to explore the probability distributions of the parameters that could produce a given set of observables. It is this which BAYESPHOT allows us to do.

At its most basic, BAYESPHOT operates by returning posterior PDFs for a requested set of physical parameters, given a set of input values (photometry, for example) obtained from either observations or other SLUG simulations. The method and its implementation is described in detail by da Silva et al. (2014) and Krumholz et al. (2015a), although the main points will be summarised here.

A stellar population can be described by a vector of N physical properties, $\vec{x}$, where it is normal to set $N = 3$. These three properties are usually cluster mass, cluster age, and extinction when considering star clusters, as in Krumholz et al. (2015b). However, as we will discuss in Chapter 3, other physical properties may be included in the analysis, and for CSPs the SFR takes the place of the cluster mass. In addition to these physical properties, the stellar population also has a vector of M photometric properties $\vec{y}$ (magnitudes across M photometric bands for example, although it could be any observable). This vector of observed light properties, $\vec{y}_{\text{obs}} = \left(y_{\text{obs},1}, y_{\text{obs},2}, ..., y_{\text{obs},M}\right)$, has an associated vector of errors $\vec{\sigma}_{\vec{y}} = \left(\sigma_{y_1}, \sigma_{y_2}, ..., \sigma_{y_M}\right)$, which we assume to be Gaussian. The output of this method is the posterior probability distribution of the stellar populations physical properties, which is therefore given by $p\left(\vec{x}|\vec{y}_{\text{obs}}; \vec{\sigma}_{\vec{y}}\right)$.

In BAYESPHOT, a kernel density estimation algorithm is employed to calculate this posterior PDF (da Silva et al., 2014; Krumholz et al., 2015a). The method begins with the joint PDF of the physical and (true) observable properties of the stellar populations we wish to analyse, given by $p\left(\vec{x}, \vec{y}\right)$. For this joint PDF, the posterior PDF that we wish to calculate is given by:

$$p\left(\vec{x}|\vec{y}\right) \propto \frac{p\left(\vec{x}, \vec{y}\right)}{p\left(\vec{y}\right)},$$ (1.11)

where we define $p(\vec{y})$ as:

$$p(\vec{y}) \propto \int p(\vec{x}, \vec{y})\, \mathrm{d}x, \tag{1.12}$$

which represents the probability distribution for the (true, not observed) observable properties, $\vec{y}$. This comes from Bayes' theorem, which states:

$$p(\vec{x}|\vec{y}) = \frac{p(\vec{y}|\vec{x})\, p(\vec{x})}{p(\vec{y})}, \tag{1.13}$$

which can then be transformed into Equation 1.11 using the probability chain rule to transform from likelihood multiplied by the prior to the joint probability distribution as follows: $p(\vec{x}, \vec{y}) = p(\vec{y}|\vec{x})\, p(\vec{x})$.

The posterior probability distribution for the true values of the observable properties is given by the prior probability distribution for these values multiplied by the likelihood function associated with the measurements of these observables:

$$p(\vec{y}|\vec{y}_{\mathrm{obs}}) \propto p(\vec{y}) G\left(\vec{y} - \vec{y}_{\mathrm{obs}}, \vec{\sigma}_{\vec{y}}\right), \tag{1.14}$$

where the likelihood is a Gaussian and the prior probability distribution $p(\vec{y})$ is given by Equation 1.12. The Gaussian likelihood has the form:

$$G\left(\vec{y}, \vec{\sigma}_{\vec{y}}\right) \propto \exp\left[-\left(\frac{y_1^2}{2\sigma_{y_1}^2} + \frac{y_2^2}{2\sigma_{y_2}^2} + ... + \frac{y_M^2}{2\sigma_{y_M}^2}\right)\right]. \tag{1.15}$$

Convolving Equations 1.11 and 1.14 gives us the posterior PDF for the physical parameters:

$$\begin{aligned} p(\vec{x}|\vec{y}_{\mathrm{obs}}) &\propto \int p(\vec{x}|\vec{y})\, p(\vec{y}|\vec{y}_{\mathrm{obs}})\, \mathrm{d}\vec{y} \\ &\propto \int p(\vec{x}, \vec{y}) G\left(\vec{y} - \vec{y}_{\mathrm{obs}}, \vec{\sigma}_{\vec{y}}\right) \mathrm{d}\vec{y} \end{aligned} \tag{1.16}$$

This has reduced the problem of computing $p(\vec{x}|\vec{y}_{\mathrm{obs}})$ to that of calculating $p(\vec{x}, \vec{y})$ which represents the joint PDF of the physical and photometric properties of the stellar populations of interest. To calculate this joint PDF, SLUG employs the kernel density estimation (KDE) technique as follows:

$$p(\vec{x}, \vec{y}) \propto \sum_i w_i K\left(\vec{x} - \vec{x}_i, \vec{y} - \vec{y}_i; \vec{h}\right). \tag{1.17}$$

To use this technique, BAYESPHOT requires a large library of SLUG model clusters, whose parameters are represented in Equation 1.17 by $\vec{x}_i$ (physical) and $\vec{y}_i$ (observable). The kernel function is represented by K, and w_i is the weight assigned to each library cluster. The bandwidth parameter is given by $\vec{h} = (h_{x_1}, h_{x_2}, ..., h_{x_N}, h_{y_1}, h_{y_2}, ..., h_{y_M})$. The weights are a combination of the sample weighting of the library and the prior:

$$w_i = \frac{p_{\text{prior}}(\vec{x}_i)}{p_{\text{lib}}(\vec{x}_i)}, \tag{1.18}$$

where p_{prior} represents the prior and p_{lib} represents the distribution of parameters that the library clusters are drawn from during library construction. This is often defined such as to produce more clusters in regions where stochasticity is most important (for example, at low cluster mass) to improve computational efficiency.

SLUG uses a Gaussian kernel function for reasons of computational efficiency, which modifies Equation 1.16 as follows:

$$p(\vec{x}|\vec{y}_{\text{obs}}) \propto \sum_i w_i G\left(\vec{x} - \vec{x}_i, \vec{h}_x\right) \int G\left(\vec{y} - \vec{y}_{\text{obs}}, \vec{\sigma}_y\right) G\left(\vec{y} - \vec{y}_i, \vec{h}_y\right) d\vec{y}, \tag{1.19}$$

where $\vec{h}_x$ and $\vec{h}_y$ represent the bandwidth in the physical and observable dimensions of $\vec{h}$ respectively. Using the result that the convolution of two Gaussians gives another Gaussian, we can rewrite Equation 1.19 as:

$$\begin{aligned}
p(\vec{x}|\vec{y}_{\text{obs}}) &\propto \sum_i w_i G\left(\vec{x} - \vec{x}_i, \vec{h}_x\right) G\left(\vec{y}_{\text{obs}} - \vec{y}_i, \sqrt{\vec{\sigma}_y^2 + \vec{h}_y^2}\right) \\
&\equiv \sum_i w_i G\left((\vec{x} - \vec{x}_i, \vec{y}_{\text{obs}} - \vec{y}_i), \vec{h}'\right),
\end{aligned} \tag{1.20}$$

where we have introduced the modified bandwidth:

$$h' = \left(\vec{h}_x, \sqrt{\vec{\sigma}_y^2 + \vec{h}_y^2}\right). \tag{1.21}$$

It is also possible to write down the marginal probabilities for each physical parameter in the same form. For example, for physical parameter x_1 (which could be mass, age, extinction, or indeed IMF slope), the marginal probability is given by:

$$p(x_1|\vec{y}_{\text{obs}}) \propto \int p(\vec{x}|\vec{y}_{\text{obs}})\,\mathrm{d}x_2\mathrm{d}x_3...\mathrm{d}x_N$$
$$\propto \sum_i w_i G\left(x_1 - x_{1,i}, h_1\right) G\left(\vec{y}_{\text{obs}} - \vec{y}_i, \sqrt{\vec{\sigma}_y^2 + \vec{h}_y^2}\right) \tag{1.22}$$

The expression also generalises to the the joint marginal probabilities of the physical variables.

Using KDE techniques allows the use of fast computational methods within SLUG. The first step is to calculate the weights for the SLUG model library. These weights and their corresponding sample points are then stored in a k-dimensional tree structure (the KD tree). The bandwidth selected for the analysis must be appropriate for the library of SLUG models, and is usually tuned by hand such that there are atleast ten models within a single kernel size. However, SLUG does have the ability to automatically estimate a bandwidth for a given library. This automatic bandwidth estimation works by taking a random sample of 5000 clusters from the library (or the whole library if it numbers less than this), and sets the bandwidth to the 90th percentile of the tenth-nearest-neighbour distance across that sample.

With this bandwidth in hand, the KD tree allows very efficient evaluation of the joint or marginal probability distributions as it avoids considering regions of the library that are not relevant for the input observations. SLUG also includes the ability to use Markov Chain Monte Carlo (MCMC) techniques, although it is usually faster to use the KDE method when generating PDFs for only one or two variables at once.

1.4.5 A brief history of SLUG

The SLUG package has been used in many studies since its first release. Here we will summarise a small selection of this work. We begin with two examples produced by members of the SLUG development team.

The first application of SLUG was described in Fumagalli et al. (2011), where SLUG's stochastic capabilities were exploited to explain the systematic variations detected in the ratio of Hα to FUV luminosities ($L_{\text{H}\alpha}/L_{\text{FUV}}$) across different galaxy properties (Meurer et al., 2009; Lee et al., 2009; Boselli et al., 2009). Several studies, for example Meurer et al. (2009) and Weidner et al. (2011) have claimed that these

observations imply a non-universal IMF. The former suggests either a changing upper-mass limit or changing upper-end slope may be responsible for the variations, whereas the latter focuses solely on the IGIMF theory as being a predictor of these results. Fumagalli et al. (2011) use these observations, along with SLUG models, to test whether a standard, but undersampled, IMF could explain these observed $L_{\mathrm{H}\alpha}/L_{\mathrm{FUV}}$ variations. Although for bright FUV luminosities, galaxies produced a $L_{\mathrm{H}\alpha}/L_{\mathrm{FUV}}$ close to that expected for a canonical and fully-sampled Kroupa IMF (and also the IGIMF), at lower luminosities and lower SFR, the data appears to follow the fully-sampled IGIMF prediction rather than that which is expected of a fully-sampled Kroupa IMF. This behaviour is what was taken to be suggestive of a varying IMF following the IGIMF theory. However, stochastic effects are missing from this picture, and when the stochastic effects of both an undersampled IMF and an undersampled SFH are included using SLUG, they found that model galaxies created using the IGIMF are in fact inconsistent with the observations, whereas the models drawn from the canonical Kroupa IMF overlap with the observed sample in $L_{\mathrm{H}\alpha}/L_{\mathrm{FUV}} - L_{\mathrm{FUV}}$ space, and that IGIMF models produce consistently lower than observed $L_{\mathrm{H}\alpha}$. Their final conclusion is that the aforementioned observations are best explained by a standard Kroupa IMF (once stochastic effects are included), especially at low luminosities, and that the discrepancy between the IGIMF and the observations is likely to be larger than any errors for faint Hα luminosities. However, due to sparse data at very low FUV luminosities, variation of the IMF in this regime was not completely ruled out in this study.

Following the release of the second version of SLUG (Krumholz et al., 2015a), the newly rewritten code was applied to the study of broad-band photometry from the Legacy Extra-Galactic UV Survey (LEGUS; Calzetti et al., 2015). Harnessing the Bayesian inference capabilities of the SLUG software suite, the properties of the unresolved star clusters in two LEGUS galaxies were explored. Taking the broad-band data from the LEGUS cluster catalogues of NGC 628 and NGC 7793, which was made from *Hubble Space Telescope* photometry in the NUV, U, B, V, and I bands, Krumholz et al. (2015b) were able to calculate full posterior probability distributions for the cluster masses, ages, and extinctions. This Bayesian technique was found to be particularly effective in low mass clusters ($M < 10^{3.5}\,\mathrm{M}_\odot$) due to undersampling of the IMF, and for clusters where there are degeneracies between the mass, age, and extinction. In both cases, a single "best fit" value would

be unsuitable, with the full posterior probability distribution for each parameter being a more reliable result than the single value that would be obtained from a more traditional and deterministic SPS code (in this study SLUG is compared to YGGDRASIL; Zackrisson et al., 2011).

In their study, Krumholz et al. (2015b) demonstrated that the probability distributions determined for the physical parameters are often broad, and in some cases multi-peaked. For individual clusters, there was found to be little dependence of these posterior distributions on the choice of evolutionary tracks, IMF, metallicity, nebular emission treatment, and extinction curve. However, the distributions were sensitive to the choice of priors due to degeneracies which lead to multiple combinations of the physical parameters being viable. On the other hand, over the whole cluster population, the recovered properties were found to be stable with respect to the choice of these parameters, although the age distributions are somewhat dependent on the choice of prior at all ages. At young ages, there was found to be a much increased dependence on prior, track choice, metallicity, and the nebular emission treatment.

Others have used SLUG for a wide range of applications, and we summarise a handful of examples here. Recently, Paalvast & Brinchmann (2017) used SLUG to show that stochastic sampling of the IMF can lead to variations in optical emission lines, and that therefore metallicity can be overestimated in galaxies with low star formation, and that care must be taken when selecting metallicity calibrators in such systems. Rubin et al. (2018) use SLUG to estimate the number and luminosity of young star clusters in the background galaxies of their sample, as part of their study of the circumgalactic medium. Andrews et al. (2013) use SLUG as part of a study of the initial mass function in dwarf starburst galaxy NGC 4214. They find, using SLUG, that an IGIMF-like truncation of the upper end of the IMF does not match well with the data and that the young clusters in NGC 4214 are consistent with a stochastically sampled IMF with a fixed upper mass limit at the canonical value. Building on this, Andrews et al. (2014) study the massive star content of M 83, once again finding that the young clusters present in the galaxy cannot be well explained with a truncated IMF. In the study presented in Sparre et al. (2017), SLUG is used in conjunction with the Feedback in Realistic Environments (FIRE) simulations (Hopkins et al., 2014) to constrain the burstiness of galaxy star formation histories.

1.5   The objectives of this work

This work presents a study of variation of the IMF in the high-mass regime, and aims to predict whether any variation of the IMF would be able to be constrained using various observational methods. Using the stochastic SPS code SLUG, along with photometry and spectra of star clusters (both real and simulated), we will explore the possbility of constraining the slope of the upper end of the IMF in unresolved stellar populations from UV and optical observations. As described above, the shape and universality of the IMF is an unsolved problem, and understanding it is key in our understanding of how stellar populations evolve. Although there are many methods for attempting to infer the IMF (such as line ratios as described earlier), the approach we take allows us to infer the IMF from a stellar population while simultaneously inferring the other parameters which describe the cluster (mass, age, extinction) from a single set of observables.

To study changes in the IMF, the SLUG code had to undergo additional extension and development if we wished to employ our novel Bayesian technique. The first of these extensions was the implementation of a variable PDF capability, which allows the variation of the parameters that describe a PDF (the IMF in this case) across a set of realisations. The second extension involved the development of a high resolution spectral synthesiser, along with the ability to calculate the equivalent widths of arbitrary spectral features, for the UV region of the spectrum. These new features are detailed in Chapter 2, and allow us to tackle the problem of constraining the form of the IMF using a novel method.

In the first application of the new variable PDF capability, *Hubble Space Telescope* observations from the Legacy ExtraGalactic UV Survey (LEGUS; Calzetti et al., 2015) comprised of broad-band photometry of clusters in the nearby galaxy NGC 628 is analysed using SLUG. Bayesian inference techniques are employed in an attempt to constrain the physical parameters that describe these stellar populations. Using our novel variable IMF capability, we are able to go further than previous studies that have been performed using SLUG (Krumholz et al., 2015b), as the IMF slope can now be included in this set of physical parameters for the first time. This study is detailed in Chapter 3.

To investigate further improvement in the likelihood of constraining the IMF, spectral lines in the UV region were employed in a theoretical study. The addition

of equivalent width information to the analysis is discussed in Chapter 4, with the inclusion of the new high-resolution spectral synthesiser detailed in Chapter 2. The constraining power of UV spectral features is then quantified through the use of mock observations generated from SLUG models.

Finally, discussion and conclusions are presented in Chapter 5, with suggestions of possible future developments and directions of study.

CHAPTER 2

AN EVOLVING SLUG

2.1 Introduction

In this chapter, we will discuss the development of the SLUG code that has been
performed as part of this work. Following the brief summary of some of the
work performed with SLUG to date, along with a summary of how the code works
presented earlier in Section 1.4, we now move on to discussing the implementation
of a variable IMF capability in Section 2.2, where the method is discussed and test
examples are presented. Following this, in Section 2.3 we detail the development
of a high resolution spectral synthesiser, along with associated equivalent width
calculation capabilities. Once again, a selection of testing data is presented. This
provides a more detailed description of the development performed by the author
on new capabilities for the SLUG software package, which are then used in the
studies described in Chapters 3 (Variable IMF) and 4 (UV spectral features).

2.2 Implementing a variable IMF

If we are to search for evidence of a variable IMF through inference from broad-band
photometry, we need a tool capable of predicting the IMF for these stellar popula-
tions. Previously, SLUG was unable to perform such a task. Indeed, the parameters
that define the IMF were static, changeable only between runs of the code and fixed
for the duration of that run. In other words, one would have to explore param-

eter space manually, creating new input IMF files for each IMF shape one wished consider. This prevents the user from running large libraries of SLUG simulations with a continuous distribution of clusters with respect to the parameters that define the shape of the IMF. Such a library is necessary if one wishes to investigate IMF variations using the Bayesian inference techniques available in the SLUG software. To be able to produce posterior PDFs of the parameters that define the IMF in much the same way as for the parameters that define the other physical aspects of a stellar population such as its mass or age, SLUG must be able to continuously vary the parameters of the IMF over a set of realisations.

As described in Section 1.3, the IMF is defined as a probability distribution of several segments, with each segment of the distribution defined by whatever parameters may be appropriate to its functional form. At its simplest, consider a single segment IMF where that single segment is a power-law. This power-law is defined by a single parameter, its slope. If we set this slope to be -2.35, we recover a Salpeter IMF (Salpeter, 1955). However, as described above, if we wish to infer the slope of an IMF from a stellar population where the slope is unknown, we need a library of SLUG models with a range of values for this parameter. It is the ability to vary these parameters that we will describe here.

We therefore introduce the concept of a "nested PDF". This concept allows one or more of the parameters that define a probability distribution to themselves be drawn from further probability distributions. This means that at the start of a realisation, we draw new values for any of the PDF's parameters that are set to this new variable mode from whichever further PDFs have been defined for those parameters. The realisation will then be constructed as normal. The process is then repeated for the next realisation.

The first step is to modify the parser which handles the PDF configuration files to allow the user to set a flag to denote variability. This flag can be set for any of the parameters that define a particular segment of a PDF. For example, the index of a power law, or the mean of a normal distribution. Any number of the parameters that define a segment may vary. Likewise, any number of segments that make up a PDF may have variable parameters. A segment parameter that is defined as variable then points to a secondary PDF definition file which defines the distribution from which new values for this parameter will be drawn.

Once the parser is able to deal with these new flags, the next step is to construct

27
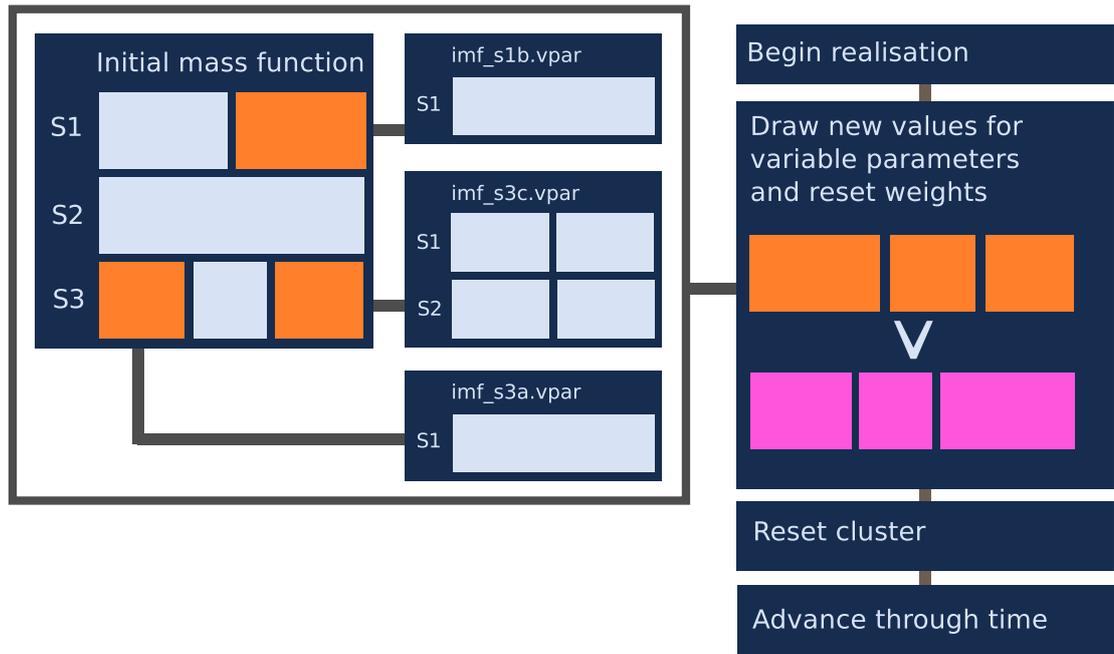
Figure 2.1: A realisation involving a three segment variable IMF, with variable parameters marked in orange. Following through with how SLUG runs, we begin at "begin realisation." At the start of the realisation we draw new values for the variable parameters and reset the weights on each segment of the PDF containing these variable parameters. The construction of this PDF is shown on the left. In this example, segment one ('S1') has its second parameter drawn from a single component PDF. Segment three ('S3') on the other hand has its first parameter drawn from a single parameter single segment PDF, and its third parameter drawn from a two segment PDF where each segment has two parameters. The new PDF is displayed in magenta. The code then resets the star cluster and continues with the realisation. This process is then repeated at the next realisation.

the framework necessary to handle these variable parameters within the SLUG code itself. This requires the creation of machinery to handle the drawing of new values for these segment parameters. The new values are drawn from PDFs defined for each of the segment parameters that have been defined to vary in the configuration files.

This implementation of variable PDFs is very general and in principle can be applied to any PDF utilised in the SLUG code during a simulation. However, at present, we implement this capability only for the IMF.

At the beginning of each realisation, be it a cluster or galaxy simulation, a new value is drawn for each of the segment parameters that are defined as variable. The weights on each segment are then recalculated and the PDF is normalised if required with the weights corresponding to the integral under the segment for the purposes of normalisation as follows:

$$w_i = n_i \int_{x_{i,a}}^{x_{i,b}} f_i(x) \mathrm{d}x,$$ (2.1)

where w_i is the weight on segment i, n is the normalisation, $x_{i,a}$ and $x_{i,b}$ are the segment boundaries and $f_i(x)$ is the function that defines the segment. The expectation values are then recalculated, leaving the IMF ready for use in that realisation. This process is then repeated for each realisation. In cluster mode, this means that each cluster is generated with a unique IMF. In galaxy mode, the clusters in each galaxy share an IMF and each galaxy realisation is created with newly drawn values for its IMF parameters. Figure 2.1 demonstrates the structure of a cluster simulation in SLUG with this new variable PDF capability. The values of any variable parameters are stored in the cluster and integrated property files which are output by SLUG.

Here we provide examples of the capabilities of the new variable PDF mode, which also function as test cases which validate the new mode's capabilities. The first example is shown in Figure 2.2. For this test, 10^4 realisations of a cluster were simulated in SLUG with a Salpeter-like (Salpeter, 1955) variable IMF (a single segment power-law PDF). The power-law index of the IMF (α) was set to vary, with its value set to be drawn from a normal distribution centred on the canonical value of -2.35, with a dispersion of 0.2. As we see in the Figure, the implementation works well, with the resulting distribution of IMF slopes falling very close to the
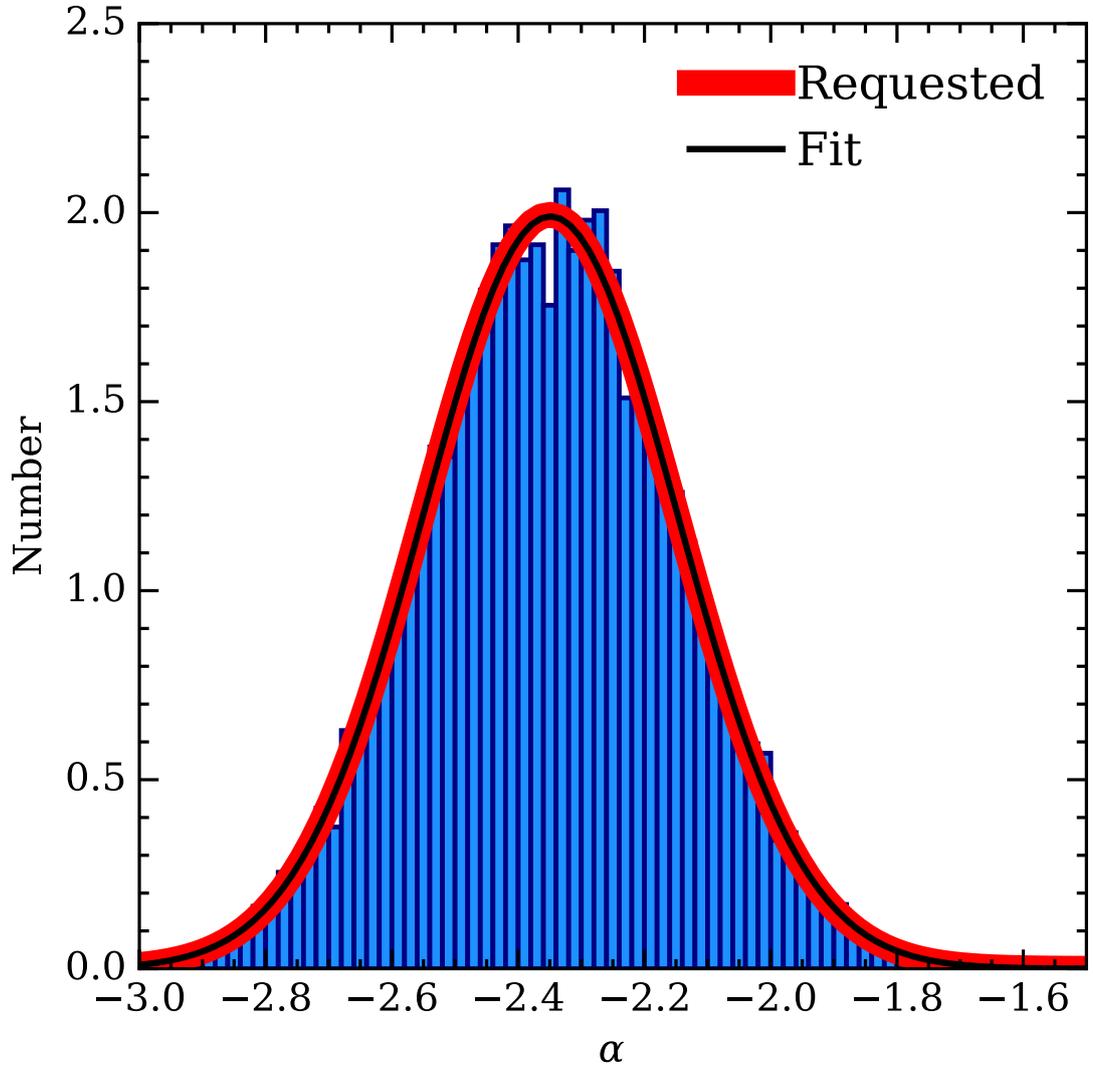
Figure 2.2: Values of the power law index α of a variable Salpeter-like IMF (blue histogram), obtained from 10^4 cluster realisations generated using the newly-developed variable mode in SLUG. The requested normal distribution for α has a mean of -2.35 and a dispersion of 0.2, and is plotted in red. The best fit normal distribution is displayed in black, and closely agrees with the input PDF. This figure originally appears in Ashworth et al. (2017).
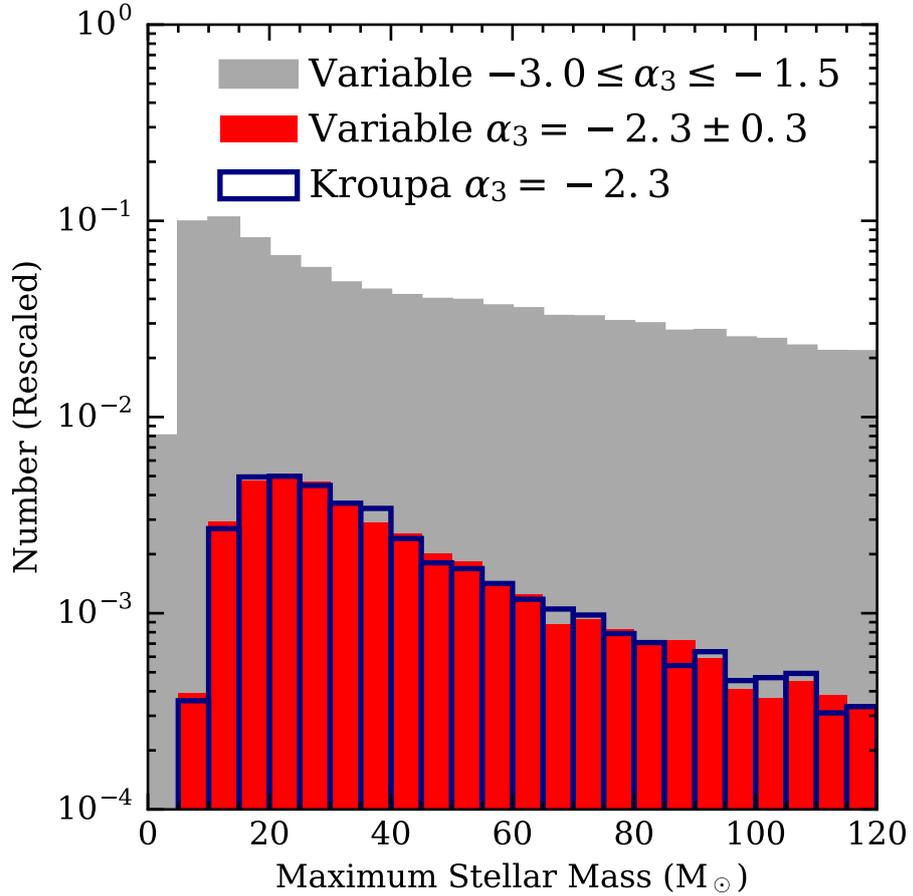
Figure 2.3: Comparison of the maximum stellar mass for two sets of realisations of a $500\,M_\odot$ cluster at $10^6\,$yr. In the first set (blue histogram), the stars are drawn from a canonical Kroupa IMF in 5000 realisations. In the second set, we generate 10^5 realisations using the variable IMF mode, with the upper-end slope α_3 drawn from a flat distribution between -3.0 and -1.5. The subset of realisations with a value of α_3 within ±0.3 of the Kroupa value is shown in red, while the full range is shown in grey. The histograms are rescaled such that there are an equal number of clusters in the bins for the Kroupa and constrained variable cases, such that their size is relative to the complete variable set. This figure originally appears in Ashworth et al. (2017).

requested normal distribution.

Next, we demonstrate the use of the variable PDFs for the case of an IMF with multiple segments, such as the commonly used Kroupa IMF. As described above, due to the flexible nature of the variable mode's implementation, we can choose to vary any number of the parameters that define the slope. In this particular example, we vary only the high-mass slope (α_3) of a three segment Kroupa-like IMF. The values for this slope are drawn from a flat distribution with limits at -3 and -1.5 for each of the 10^5 realisations of a $500\,\mathrm{M_\odot}$ star cluster. The maximum stellar mass in each realisation is plotted in Figure 2.3. As can be seen in the figure, we are able to mimic the behaviour of the fixed-parameter case (dark blue) by extracting a subset of clusters with a slope close to this canonical value from the full variable IMF dataset. We also see, from the shape of the full variable dataset's histogram, that including shallow IMFs with slopes as low as $\alpha_3 = -1.5$ results in a higher probability of drawing massive stars when compared to the steeper Kroupa IMF. This behaviour is what one would expect in this situation, which validates our implementation of the variable IMF method. The method works as expected in both "cluster" mode and "galaxy" mode, with each realisation of the galaxy drawing new values for the IMF slope, which is then used in drawing the stars that make up the clusters and field of that realisation of the galaxy

2.3 Implementation of a high-resolution UV spectrum

The second major modification to SLUG that was undertaken as part of this work was the implementation of a high-resolution UV spectral synthesiser. Until the creation of this extension, the UV spectra that SLUG produces have been limited to a resolution of $\simeq 10\,\mathrm{\mathring{A}}$ by the stellar atmosphere libraries included with the SLUG package. Unfortunately, this resolution is insufficient for equivalent width studies of varying IMFs, as many features would be smaller than a single wavelength element, or even only as big as a few elements at most. In order to be able to use the equivalent widths of UV spectral features in any analysis of the IMF, we therefore need new stellar atmosphere libraries with more finely sampled wavelength grids in the UV region of the spectrum. To this end, we create a new spectral synthesiser

for SLUG, dubbed SB99HRUV. To provide this new synthesiser with the capability of producing high resolution spectra in the UV region, we have incorporated the "IFA" atmospheres of STARBURST99 (Leitherer et al., 1999) into the SLUG atmosphere library. These model atmospheres were first described by Leitherer et al. (2010), and have a resolution of $\simeq 0.4\,$Å and cover the wavelength range between 900 Å and 3000 Å. Therefore, the new synthesiser must use other atmospheres to cover the wavelengths that lie outside this UV region of the spectrum. For the region of the spectrum between 900 Å and 3000 Å the following atmospheres are used:

- Stars classified as Wolf-Rayet stars use the Potsdam PoWR atmospheres (Gräfener et al., 2002; Hamann & Gräfener, 2004; Sander et al., 2012).

- Kurucz (Kurucz, 1992; Lejeune et al., 1997) atmospheres are used for stars with $M < 5\,\mathrm{M}_\odot$. They are also used for stars with $T < 170000$ K in the mass range $5 - 10\,\mathrm{M}_\odot$, and for stars with $T < 22000$ K at masses $> 10\,\mathrm{M}_\odot$.

- For the remaining stars, we use the high resolution IFA atmospheres (Leitherer et al., 2010) if they fall within the boundaries of this library.

- Stars that fall outside of the above atmosphere grids are treated as black bodies.

For the rest of the wavelength range of the spectrum, we use the standard lower resolution atmospheres that are employed by SLUG, mimicking the behaviour of the standard SB99 synthesiser. This selection corresponds to the following:

- Stars classified as Wolf-Rayet stars use the Hillier atmospheres (Hillier & Miller, 1998).

- Stars with temperatures too high or too low for the Kurucz atmospheres are treated as black bodies.

- Stars whose temperatures are above the minimum temperature of the Pauldrach OB atmospheres (Pauldrach et al., 1998), while also having a surface gravity that is both greater than the minimum surface gravity for the Pauldrach atmospheres and less than the maximum surface gravity for that stars temperature in the atmospheres are handled using the Pauldrach OB star atmospheres.

- All other stars use the Kurucz atmospheres.

These high and low resolution regions are then stitched together. Due to the low density of points in some regions of the temperature-surface gravity plane for the high resolution atmospheres, certain parameter combinations lead to large differences in flux at the boundary between the high and low resolution zones. To minimise these discontinuities, we scale the whole of the high resolution zone by a calculated scaling factor. This scaling factor is calculated by integrating under a region of the spectrum (between 2400Å and 2900Å) in both the high and low resolution spectra and then taking the ratio of the two integrals. We recommend that users interested only in the photometry of star clusters use the low resolution spectral synthesiser however, as this method is at best an approximation.

We show an example of this rescaling in Figure 2.4. We plot the spectrum of a $10^4 \, \mathrm{M_\odot}$ cluster at an age of 10^7 yr, generated with a Salpeter IMF at Solar metallicity. The SLUG spectra shown in the figure are the medians of 100 realisations. The spectrum in red is generated using the high resolution atmospheres with no rescaling, and the blue spectrum is generated using the high resolution atmospheres with the rescaling factor described above. In grey is the low resolution SLUG spectrum. Spectra from STARBURST99 are also plotted, with the low resolution spectrum shown in magenta and the high resolution spectrum shown in green. The inflated flux produced when using the high resolution atmospheres is clearly visible in both STARBURST99 and SLUG. The reason for this is that at an age of 10^7 yr, stars of $10 \, \mathrm{M_\odot}$ contribute strongly to the UV region of the spectrum. This region of the atmosphere grid is very sparse, as can be seen in the bottom right of Figure 1 in Leitherer et al. (2010), and also towards the left of Figure 2.5. Due to the nearest neighbour atmosphere selection method, stars in this region usually jump between either model 83 or model 84 of the atmosphere grid (Leitherer 2017, Private Communication), and when combined over the whole cluster, the overall luminosity is overestimated. We see that the rescaled spectrum lies much closer to the low resolution spectrum (with some divergence towards 900Å. As described above, this rescaled spectrum is still at best an approximation.

We note that the work described in Chapter 4 uses the PoWR atmospheres shipped with STARBURST99 as new versions were not yet available at the time this analysis was performed. However, a new revision of the PoWR atmospheres
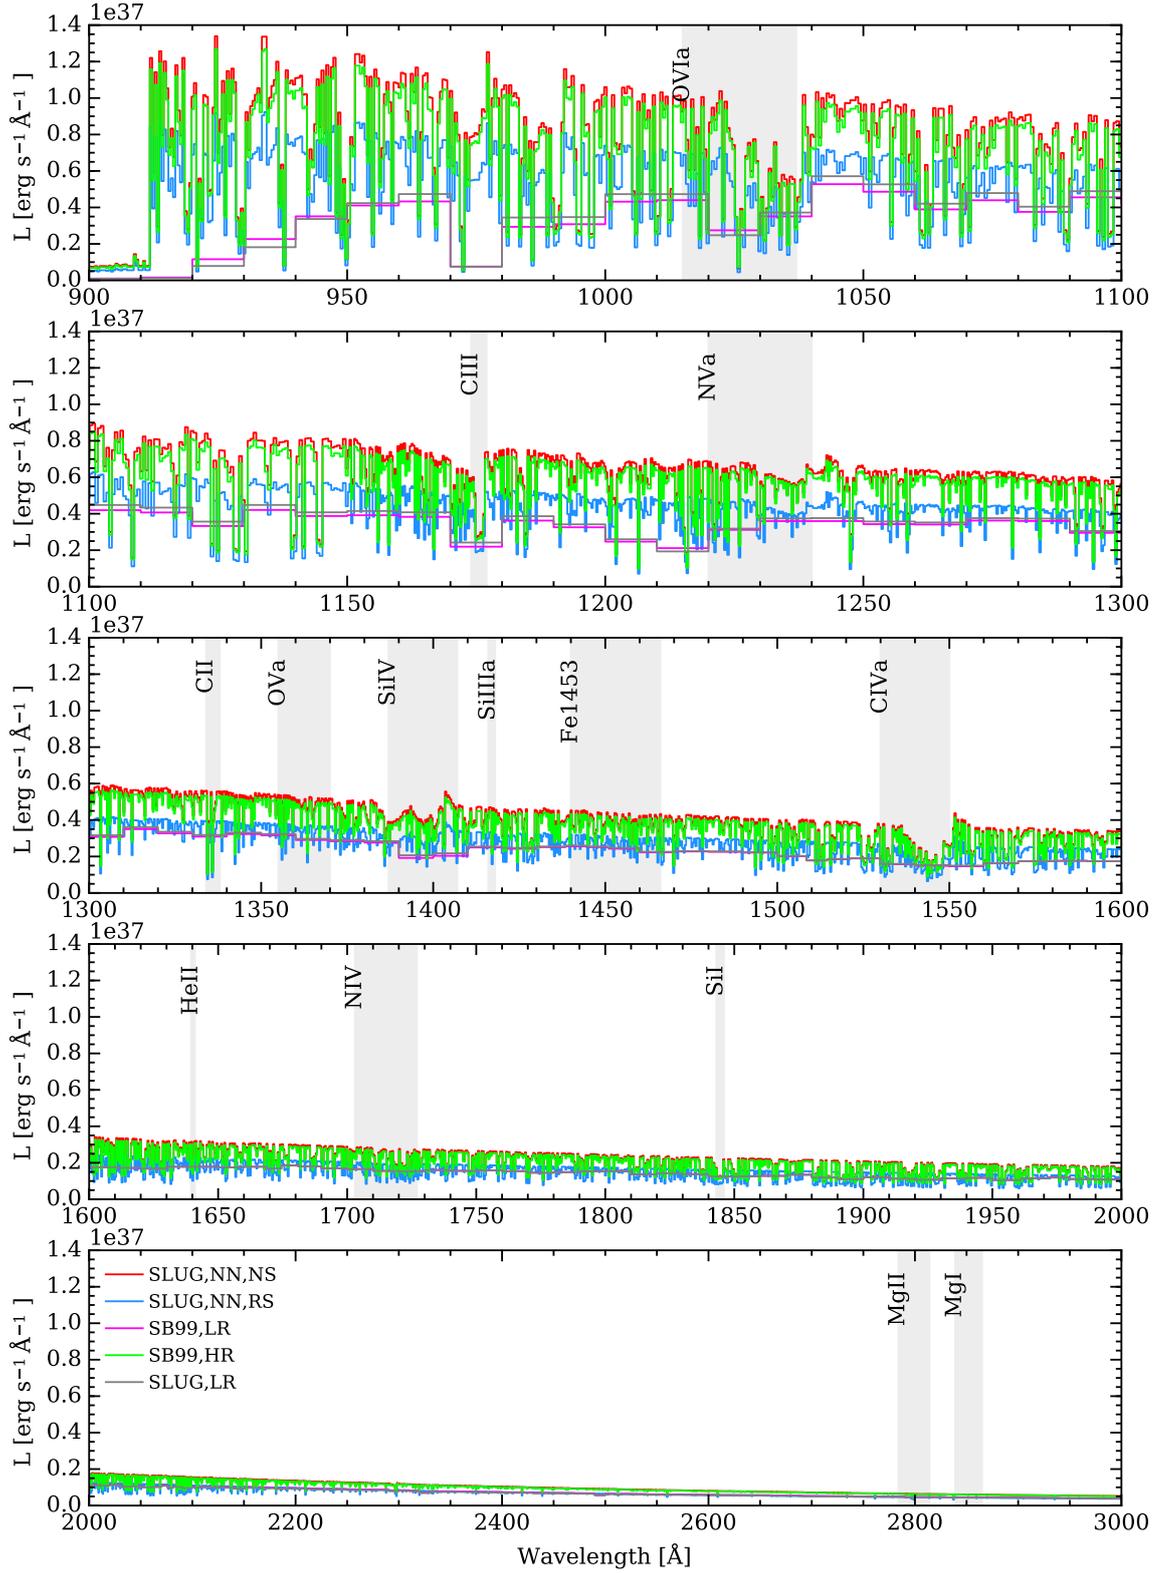
Figure 2.4: Comparison between the spectra of $10^4 \, M_\odot$ star clusters generated with SLUG and STARBURST99. Clusters were generated with SLUG using the nearest neighbour method with rescaling on (blue) and off (red) using the high resolution atmospheres, and with the low resolution atmospheres (grey), with each spectrum being the median of 100 realisations. The STARBURST99 spectra are showin in green (high resolution) and magenta (low resolution). The increased flux in the high resolution case is clearly visible.

has since been made available and these new models are to be reformatted for use with SLUG. These new atmospheres will then be included in the public release [1].

An important decision when implementing the new high resolution UV atmosphere library is the choice of interpolation scheme. Given an OB-type star, how do we choose an atmosphere to represent it? In STARBURST99, the atmospheres are selected using a weighted nearest-neighbour algorithm, with each direction (in the surface gravity - temperature plane) weighted differently. In this case, the surface gravity axis is weighted by a factor of 5 (Leitherer et al., 2010). In SLUG, we implement the same algorithm. However, we also implement a new interpolation scheme which may be toggled with a boolean switch in the spectral synthesiser code. It is this new interpolaton scheme which is selected by default. There are several reasons for providing an alternative interpolation scheme, for example during testing, it was noticed that in some cases minor changes in stellar properties resulted in large changes in light output as the star jumped to a different atmosphere. The weighting for the nearest neighbour finder may be adjusted for both methods.

There are several steps to this new interpolation scheme. We begin with the parameters of the star for which we wish to obtain an atmosphere, namely its temperature and surface gravity. The first step of the interpolation process is to search the atmosphere grid for the two nearest-neighbour atmospheres in the temperature-surface gravity plane. The definition of 'nearest neighbour' is the weighted nearest-neighbour used by STARBURST99, as described above. We then iterate through the next nearest neighbours until we satisfy the conditions required to enclose the star within a triangle of three atmospheres. Having formed this triangle of nearest neighbour atmosphere models, we find the barycentric coordinates of the star of interest within this triangle. These coordinates are then used to weight the atmospheres on each vertex of the triangle. The weighted combination of these atmospheres gives us an interpolated atmosphere model for the star of interest. If the triangle cannot be closed, we default to the weighted nearest-neighbour method used in STARBURST99. An example of the triangle construction for a random selection of stars is shown in Figure 2.5.

To test this new interpolation scheme, we first perturb the parameters of a stellar model present in the library by a small amount, and then compare the calcu-

---

[1] www.slugsps.com
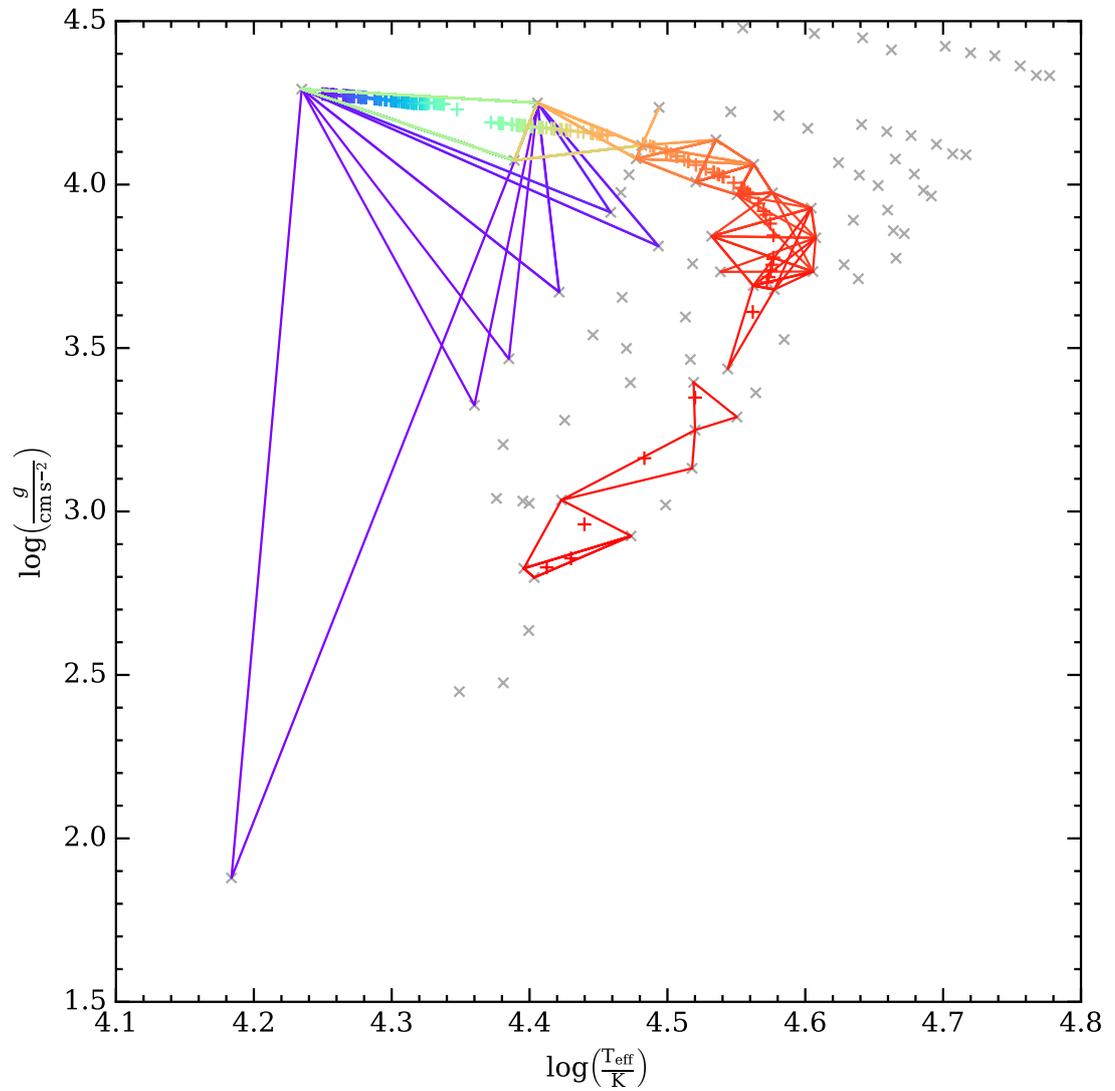
Figure 2.5: A selection of randomly generated stellar parameters positioned throughout the temperature surface gravity plane. The corresponding triangle of interpolation points is marked in the same colour as its marker. The grey crosses mark the location of the atmospheres in the Leitherer et al. (2010) library that is used in this spectral synthesiser.
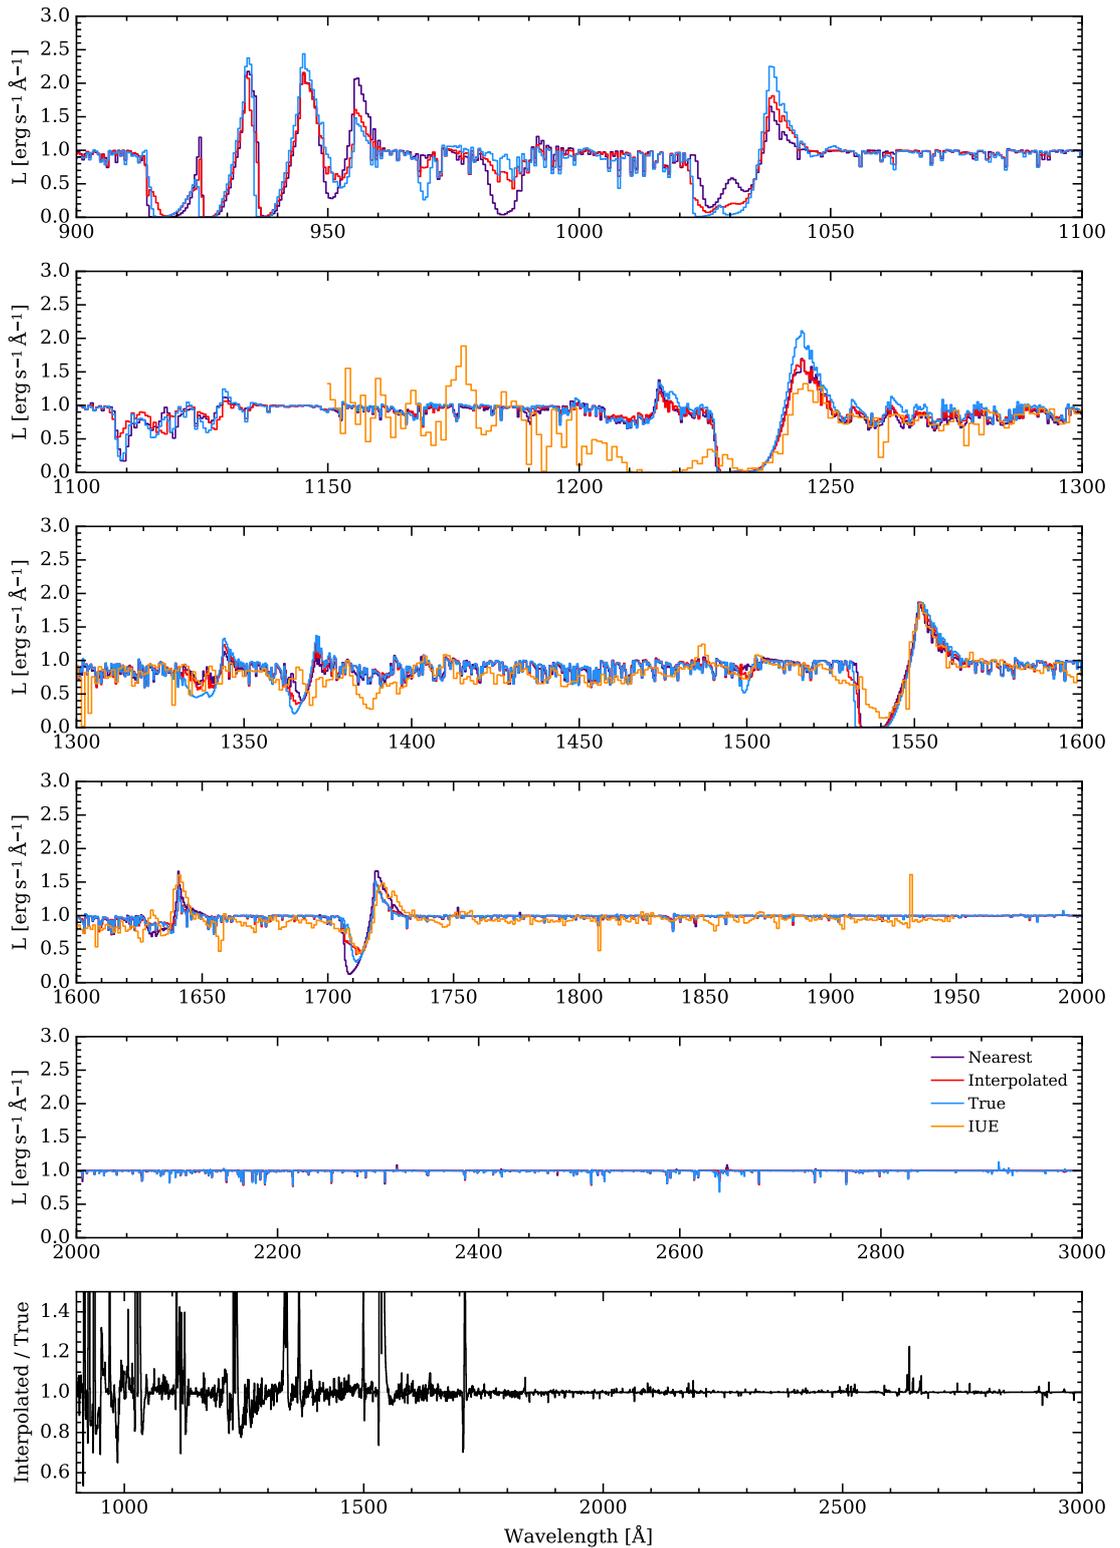
Figure 2.6: Comparison between the spectra for one of the stellar models (model 21) in the IFA atmosphere grid. The true spectrum for that model is plotted in blue. The interpolated spectrum is plotted in red and the nearest neighbour spectrum is plotted in purple, both from an IFA grid with this particular stellar model removed. In orange is shown the observed spectrum for a star of matching type, marked as "IUE" and taken from Walborn et al. (1995a,b). The observed spectrum is the mean of three O3I spectra from this dataset. For more details on companion stars to model spectra, see Leitherer et al. (2010).

38

lated interpolated spectrum to the spectrum for that stellar model. The resulting spectra were found to be very close to the true values. The next test was to remove a star from the atmosphere grid, and choose its parameters as the parameters of our model star. We then compare the spectrum produced by the interpolation with the nearest neighbour atmosphere that would have been selected using the STAR-BURST99 method, the true atmosphere for that star, and an observed atmosphere (as done by Leitherer et al., 2010). These spectra are plotted in Figure 2.6, for Model 21 from the IFA grid. We see that all three atmospheres follow the shape of the observed stellar spectrum (orange), a mean of three O3I stellar spectra taken from Walborn et al. (1995a,b). We also see that in this case the interpolated spectrum (plotted in red) is a closer match to the true atmosphere (blue) than the atmosphere selected using the nearest neighbour method (purple).

In addition to implementing the new spectral synthesiser, we also introduce the ability for SLUG to handle spectral lines, and also the ability to calculate their equivalent width. The user is now able to provide SLUG with a list of line features, and these will then have their equivalent width calculated and stored for each realisation during a run of the code. To define a feature all that is required is choosing a name for that feature (usually the name of the spectral line of interest) along with the wavelength of that line (which is not currently used) and the boundaries of the integration region over which the equivalent width will be calculated.

There are some caveats to this new equivalent width calculation feature. To avoid the uncertainties and complexity introduced by attempting to fit a pseudo-continuum around the line feature, we initially use a truly normalised (or "recti-fied") spectrum (Leitherer et al., 2010) which contains only the O and B type stars covered by the high resolution atmospheres. This is possible because the high reso-lution IFA atmospheres are provided in both line and continuum form, allowing for a perfect normalisation. The downside to this is that the spectra we calculate our equivalent widths from contain *only* the features present in the massive stars covered by these new atmospheres. This also means that once all of the massive stars evolve or die away, this rectified spectrum becomes a featureless continuum. We also do not include any nebular emission in this rectified spectrum.

Once we have the rectified spectrum for the cluster, we can calculate the equiv-alent width of a spectral line. The spectrum is linearly interpolated across the integration region. This interpolated function is then integrated over the region

of interest. Several different integration methods have been implemented, and a comparison of the results is shown in Figure 2.7 for a selection of lines, listed in Table 2.1. In red are the results from the EVAL_INTEG function of the GSL interpolator[2], in purple are the results from the QAGS integrator (also in GSL) and in blue are the results from SLUG's internal INT_TABULATED function. The figure shows that all the methods available produce comparable results. The equivalent width calculator is set to use QAGS by default, with a fall-back to the EVAL_INTEG method if QAGS encounters problems. This integral is represented by I in the following equation:

$$W = \int_{\lambda_A}^{\lambda_B} \left(1 - \frac{F_\lambda}{F_0}\right) d\lambda = (\lambda_B - \lambda_A) - \int_{\lambda_A}^{\lambda_B} \frac{F_\lambda}{F_0} d\lambda = \Delta\lambda - I, \tag{2.2}$$

where W is the equivalent width, λ_A and λ_B are the limits of the integration region (which has width $\Delta\lambda$), and F_λ and F_0 are the total flux and continuum flux respectively.

Of course, as SLUG also outputs the full spectrum, it is possible to calculate equivalent widths in post-processing. This also allows the user to include all the stars (including the Wolf-Rayet stars), along with nebular emission. However, this does require the creation of a pseudo-continuum in order to normalise the spectrum before integrating.

With this equivalent width calculator in hand, we then repeat the interpolator testing process described above for the entire atmosphere grid of 86 model stars, of which 50 are positioned such that a triangle may be constructed around them. We then study the residuals between the calculated equivalent widths (using the nearest-neighbour and interpolation methods) and the true value for the missing model star. Across the 34 example lines (see Table 2.1, later used in Chapter 4) we find only one occasion where the interpolation method results in fewer more accurate equivalent widths than the nearest neighbour method (OVI, where the nearest-neighbour method is better 52% of the time). Four examples of these residual comparisons are plotted in Figure 2.8. The red dots mark the models where the interpolation performed best and the blue dots mark the models where the nearest-neighbour method performed best. The triangular red marker indicates the median residual for the interpolation and the blue diamond marker indicates the median

---

[2]http://www.gnu.org/software/gsl/
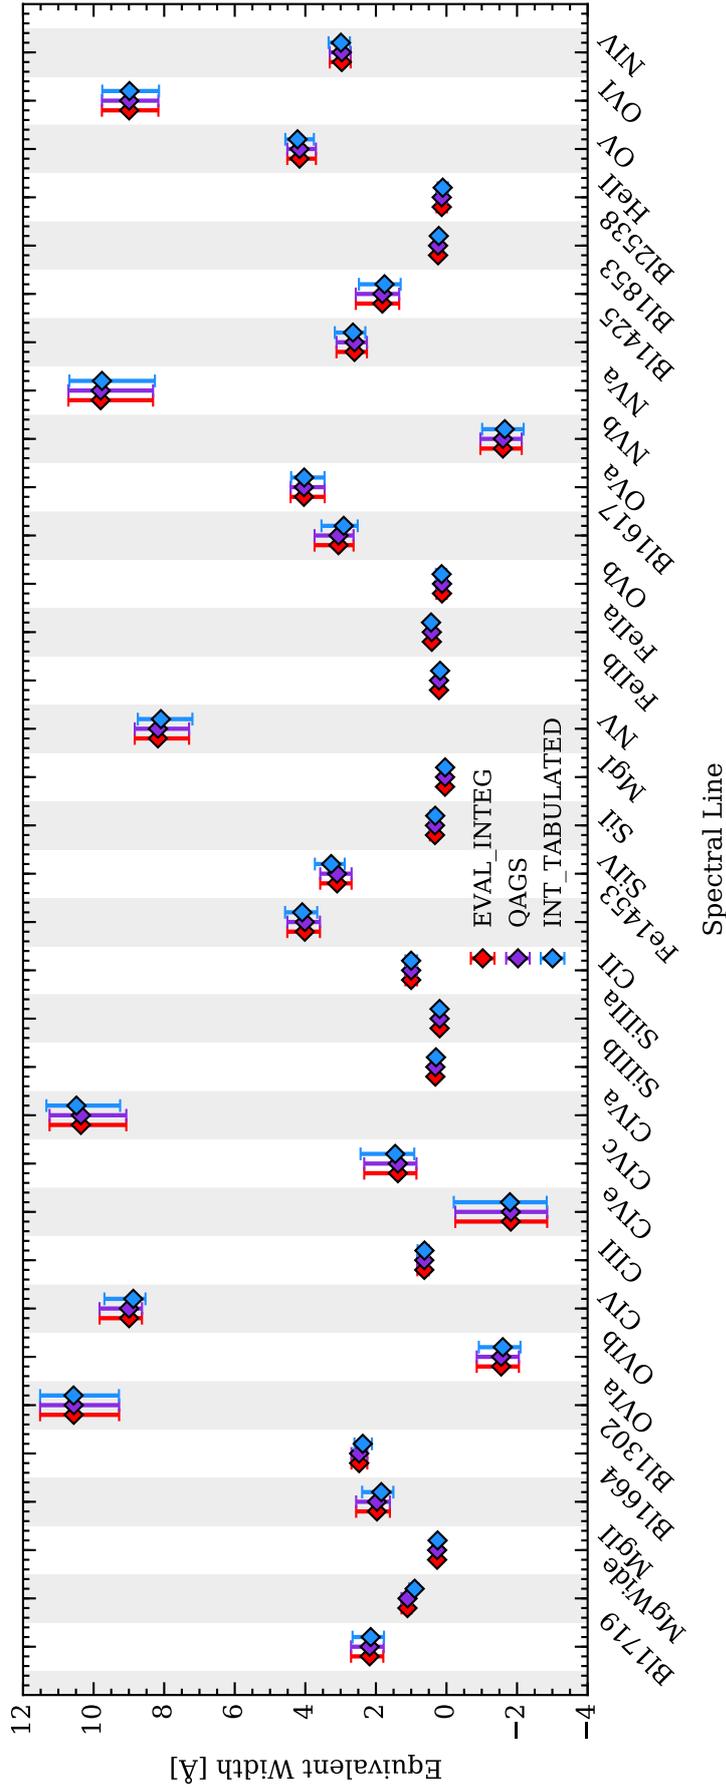
Figure 2.7: Comparison of integration methods available for calculating equivalent widths in SLUG, those being the GSL QAGS method, the EVAL_INTEG method from the GSL interpolator, and SLUG's built-in INT_TABULATED. Results taken from 1000 realisations of a $10^4 M_\odot$ star cluster at 10^6 yr with a Salpeter IMF. The markers show the positions of the median values, with the error bars covering the fifth to ninety-fifth percentile range. All three methods produce comparable results, with the GSL method selected as the default in this work.
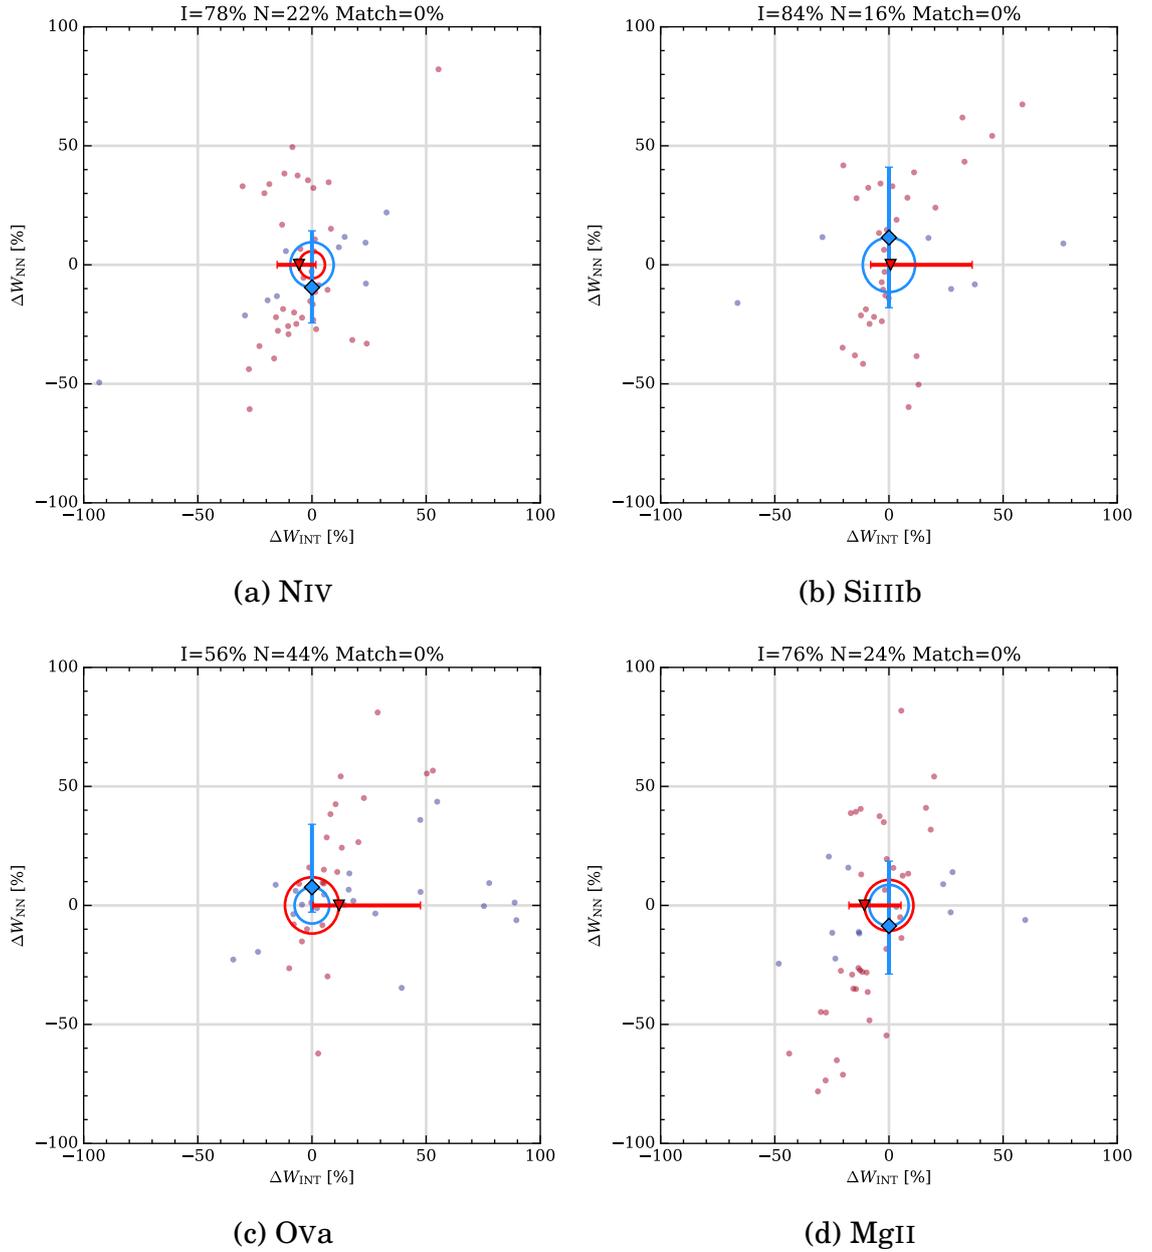
Figure 2.8: Comparison of percentage residuals between equivalent widths calculated from a stellar atmosphere taken from the high-resolution IFA grid, and equivalent widths calculated using the nearest neighbour (NN) and interpolated (INT) methods from the same grid with that atmosphere removed. Markers in red are those where the interpolator performs better, and markers in blue are those atmospheres for which the nearest neighbour method performs better. The blue diamond marker indicates the median residual for the nearest-neighbour method, and the red triangle marks the median residual for the interpolated method. The error bars mark the interquartile range. The percentages above each panel mark the fraction of models for which the interpolator (I) and the nearest-neighbour method (N) were closer to the true value.

Feature Name	Integration Range (Å)	ϕ^{a} (eV)
Ovi	$1015.0 - 1050.0$	138.1189
Ovia	$1015.0 - 1037.0$	–
Ovib	$1037.0 - 1050.0$	–
Ciii	$1174.0 - 1177.0$	47.88778
Siiiib	$1205.0 - 1207.0$	33.49300
Nv	$1220.0 - 1250.0$	97.89013
Nva	$1220.0 - 1240.0$	–
Nvb	$1240.0 - 1250.0$	–
Bl1302 *	$1292.0 - 1312.0$	–
Cii	$1334.0 - 1338.0$	24.3845
Ov	$1355.0 - 1375.0$	113.8989
Ova	$1355.0 - 1370.0$	–
Ovb	$1370.0 - 1375.0$	–
Siiv *	$1387.0 - 1407.0$	45.14179
Bl1425 *	$1415.0 - 1435.0$	–
Siiiia	$1416.0 - 1418.0$	33.49300
Fe1453 *	$1440.0 - 1466.0$	–
Civ	$1527.0 - 1569.0$	64.49351
Civa *	$1530.0 - 1550.0$	–
Civc *	$1540.0 - 1560.0$	–
Cive *	$1550.0 - 1570.0$	–
Bl1617 *	$1604.0 - 1630.0$	–
Heii	$1639.5 - 1641.0$	54.4177650
Bl1664 *	$1651.0 - 1677.0$	–
Niv	$1703.0 - 1727.0$	77.4735
Bl1719 *	$1709.0 - 1729.0$	–
Bl1853 *	$1838.0 - 1868.0$	–
Sii	$1843.0 - 1846.0$	8.15168
Feiia *	$2382.0 - 2422.0$	16.19920
Bl2538 *	$2520.0 - 2556.0$	–
Feiib *	$2596.0 - 2622.0$	16.19920
MgWide *	$2670.0 - 2870.0$	–
Mgii *	$2784.0 - 2814.0$	15.035271
Mgi *	$2839.0 - 2865.0$	7.646236

[a] Ionisation potential

Table 2.1: Table of the line features used in SLUG, where we list the integration regions used for calculating the equivalent width of each feature, along with an ionisation energy if appropriate. The ionisation potentials given were obtained from Kramida et al. (2015). The lines are sorted by increasing leading edge wavelength of their integration region. Spectral features taken from Fanelli et al. (1992) are marked with an asterisk (*). In Civ, the letters a,c,e represent the absorption, centre, and emission respectively (as given in Fanelli et al. (1992)). For P-Cygni features such as Nv, the a and b represent the absorption and emission components respectively. For Siiii the a and b represent two occurrences of the same feature, with a being the stronger in equivalent width. This table is adapted from Chapter 4.

residual for the nearest-neighbour method. The circles are to help make the relative positions of the medians more visible. The error bars mark the interquartile range. In all four panels of this figure, we see that using our new interpolation scheme results in equivalent widths that are closer to the true value for more of the model stars than using the nearest neighbour method does. However, note that in panels (c) and (d), although the interpolator performs better over the sample (being closer for as many as 76% of the test stars in the case of MgII), the median residual for the nearest neighbour case is closer to zero than that of the interpolator case.

In Figure 2.9, the medians and interquartile ranges of the residuals for all of the line features are shown. Although the interpolator is correct more often than the nearest neighbour method, we see that the median residual is less than that of the nearest neighbour method in only one third of the lines tested here, although overall the median residuals are comparable. In both cases we are trying to create an atmosphere where one does not exist, and so neither method will be a perfect representation of the target star. In addition, by removing a star to check how each process performs we are in effect reducing the density of points in that area of parameter space. Therefore, we expect both methods to perform better when used in practice. Overall, we see that both methods are comparable, with the interpolator being more correct more often on a star by star basis, but the median residual of the nearest-neighbour method for a given line being superior more often. We must stress that the analysis presented later in Chapter 4 is self-consistent due to the use of the interpolation scheme for both the SLUG models that are used to generate the mock observation and those SLUG models used to construct the library. However, more detailed comparisons are required if this new capability is to be applied to real spectroscopic data.

In this chapter, and the work presented in Chapter 4, we use rectified spectra generated with O and B stars alone. In STARBURST99 this true continuum is supplemented by pseudo-continua calculated in different ways for different stars, and is further combined with nebular emission. For stars which are treated as black bodies, the continuum is merely the same black body, and so the contribution from these stars is cancelled out in the final rectified spectrum. Stars which are covered by the Kurucz atmospheres are treated in a slightly more complex way. A selection of "continuum" wavelengths is provided, and a continuum spectrum is created by linearly interpolating between the average of the fluxes at these points and the
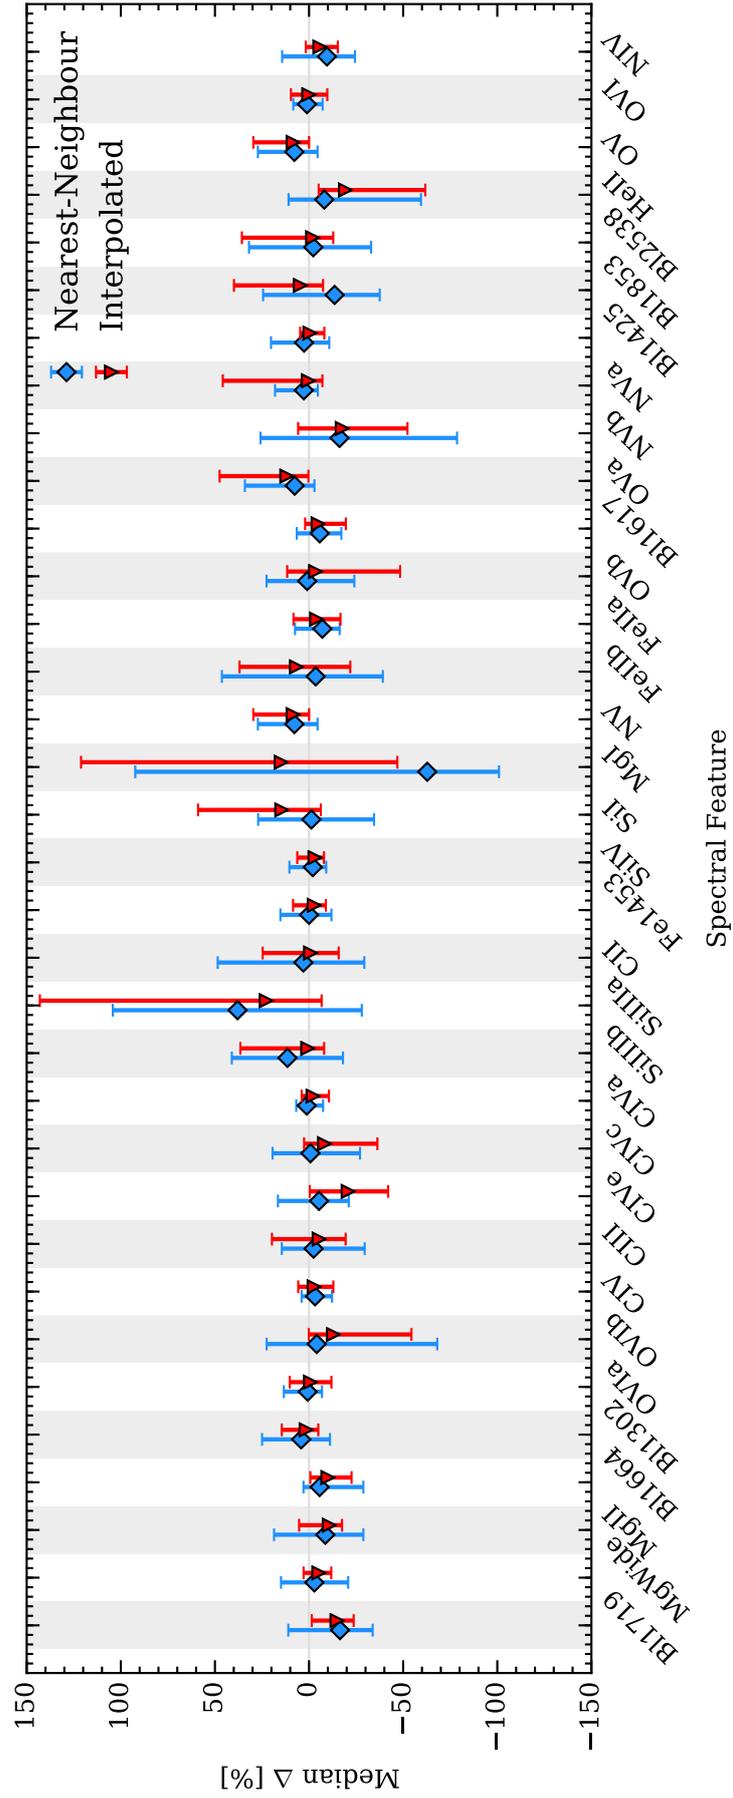
Figure 2.9: Medians and interquartile ranges for the residuals between the calculated equivalent width and true equivalent width for 34 spectral features. Results are plotted for both the nearest-neighbour method (blue diamonds) and interpolator (red triangles).

45

points on each side of these continuum points. Finally, the spectra of Wolf-Rayet stars is normalised by generating a black body spectrum for the WR stars, and then rescaling this by the ratio of the WR spectrum and the Black Body spectrum at a selected point. Following the initial studies presented in Chapter 4, this full-range rectified spectrum was implemented in SLUG (with the caveat that it is an approximation to a normalised spectrum, as in STARBURST99). This extended capability will also be made available in the public release of the code. It is worth mentioning that the initial studies presented in that chapter are self consistent in that both the simulaton library and the mock observations had their equivalent widths calculated from the same rectified spectra, containing only O and B type stars, and with nebular contributions absent. The effect of this on the cluster property constraining power is both positive and negative in different ways. The loss of O and B star features as the cluster ages against an otherwise flat spectrum aids age recovery as the OB-spectrum evolves through time. However, on the other hand, the absence of strongly age-defining features from, for example, Wolf-Rayet stars hampers age determination in another way. To see the effect of changing from an OB-Only to a full-range normalised spectrum, the equivalent widths for a $10^6\,\mathrm{M_\odot}$ cluster at ages of 0.1,1.0, and 10.0 Myr are plotted in Figure 2.10 (using the nearest neighbour atmosphere selection method as to be more comparable with STARBURST99). Most of the lines have comparable equivalent widths, with the difference being under 5% between the two methods. Some line features (between SiIIIb and MgWide in Figure 2.10) do show a more significant difference, however. From these more strongly affected lines, only MgI and MgII are used in the independent line analysis in Chapter 4, and none are present in the set of 5 representative lines. The differences become more striking once we consider certain stages of stellar evolution, for example Wolf-Rayet stars. In Figure 2.11, the spectrum of a $10^6\,\mathrm{M_\odot}$ star cluster at an age of $10^{6.5}$ yr is shown. This cluster was generated using Padova-AGB tracks, the nearest neighbour selection method, and solar metallicity. The IMF chosen was a top heavy Kroupa-like IMF, with an upper-end slope of $\alpha_3 = -1.5$. Both clusters are generated from the same seed, and therefore contain the same randomly drawn stars. It can be clearly seen that a great many UV features are entirely absent in the case of the limited spectrum (blue). Including these features in our analysis would provide further improvements to the age determination for a given cluster. The strength of these features would also be dependent on the number of massive stars,

potentially improving recovery of the IMF as well. The absence of these features, as determined from the equivalent width measurements presented in Figure 2.10, should not have a significant effect on the results presented in the bulk of Chapter 4 as we can see from the ratio of the two shown in the final panel of Figure 2.11 that the most significant discrepancies occur outside of the lines used in that study.

Finally, in Figure 2.12, we perform the same comparison, this time including the same spectra generated using STARBURST99. In this plot, nebular emission is disabled for both the OB-only and all-stars cases in both SLUG and STARBURST99. We see that both SLUG and STARBURST99 produce comparable spectra (with the SLUG spectrum here being the median of 10 realisations). The small differences are likely to be caused by minor stochastic variations, and any differences in SLUG's treatment of the evolutionary tracks. This can be due to SLUG's method of isochrone construction which can result in stars with slightly different properties to those in STARBURST99, a difference that can be exacerbated by the nearest neighnour atmosphere selection method when it comes to generating spectra. In addition, following the creation of these plots, a minor error was discovered by the author in how STARBURST99 normalises the spectra of Wolf-Rayet stars. In the cluster presented here, a maximum error of a few percent was introduced. The authors of STARBURST99 have been notified and this issue will be corrected in the next release of the code.

For more accurate comparison to observations, including the full range of stars in the normalised spectrum would be the best choice, even though the continua generated for most of the stars is an approximation. In addition, the inclusion of additional features which the additional stars produce provides us with additional information with which to determine the physical parameters of the star clusters we observe.

Having calculated equivalent width values as described in Equation 2.2, we must then ask how we may use them in our analysis. Fortunately, BAYESPHOT is a generalised method (see Section 1.4.4) which works between any observable property and any physical property. It is therefore possible for us to feed the newly calculated equivalent width information into BAYESPHOT alongside the broad-band photometry by producing a wrapper script. This script fulfils a similar purpose to SLUG's CLUSTER_SLUG tool, itself a wrapper around BAYESPHOT, but allows the inclusion of equivalent widths with the observational data (or indeed, as the only
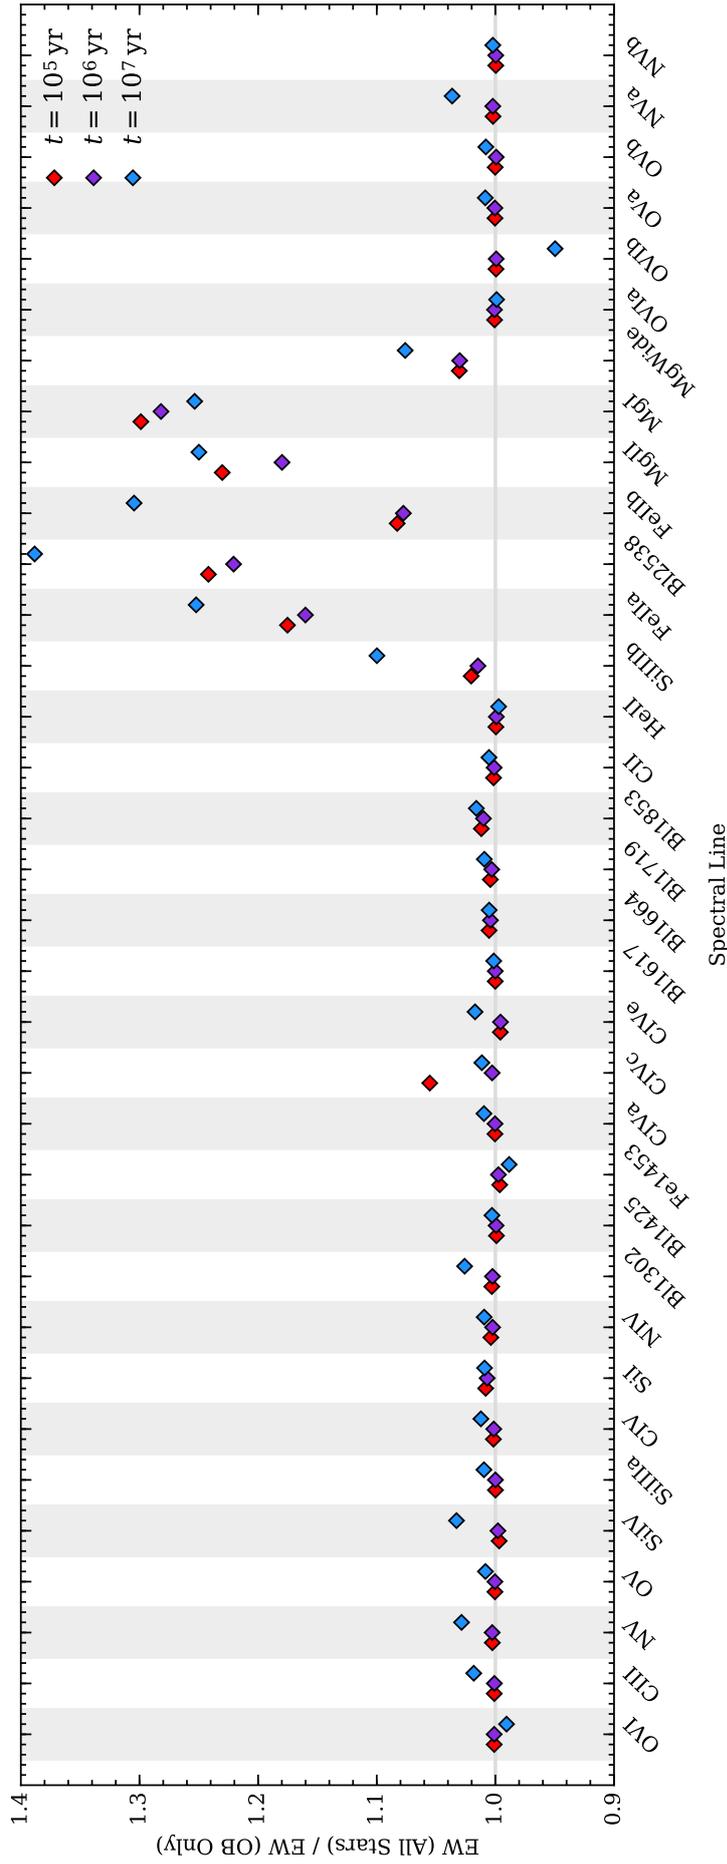
Figure 2.10: Equivalent width ratios for the full list of spectral lines listed in Table 2.1. The lines are calculated from a $10^6 M_\odot$ cluster at 0.1, 1.0, and 10.0 Myr, with a Salpeter IMF. The cluster is at Solar metallicity, and was generated using Padova tracks with AGB. The ratio is between the equivalent width calculated using the spectrum containing O and B stars only, and the equivalent width calculated from a full-range spectrum.
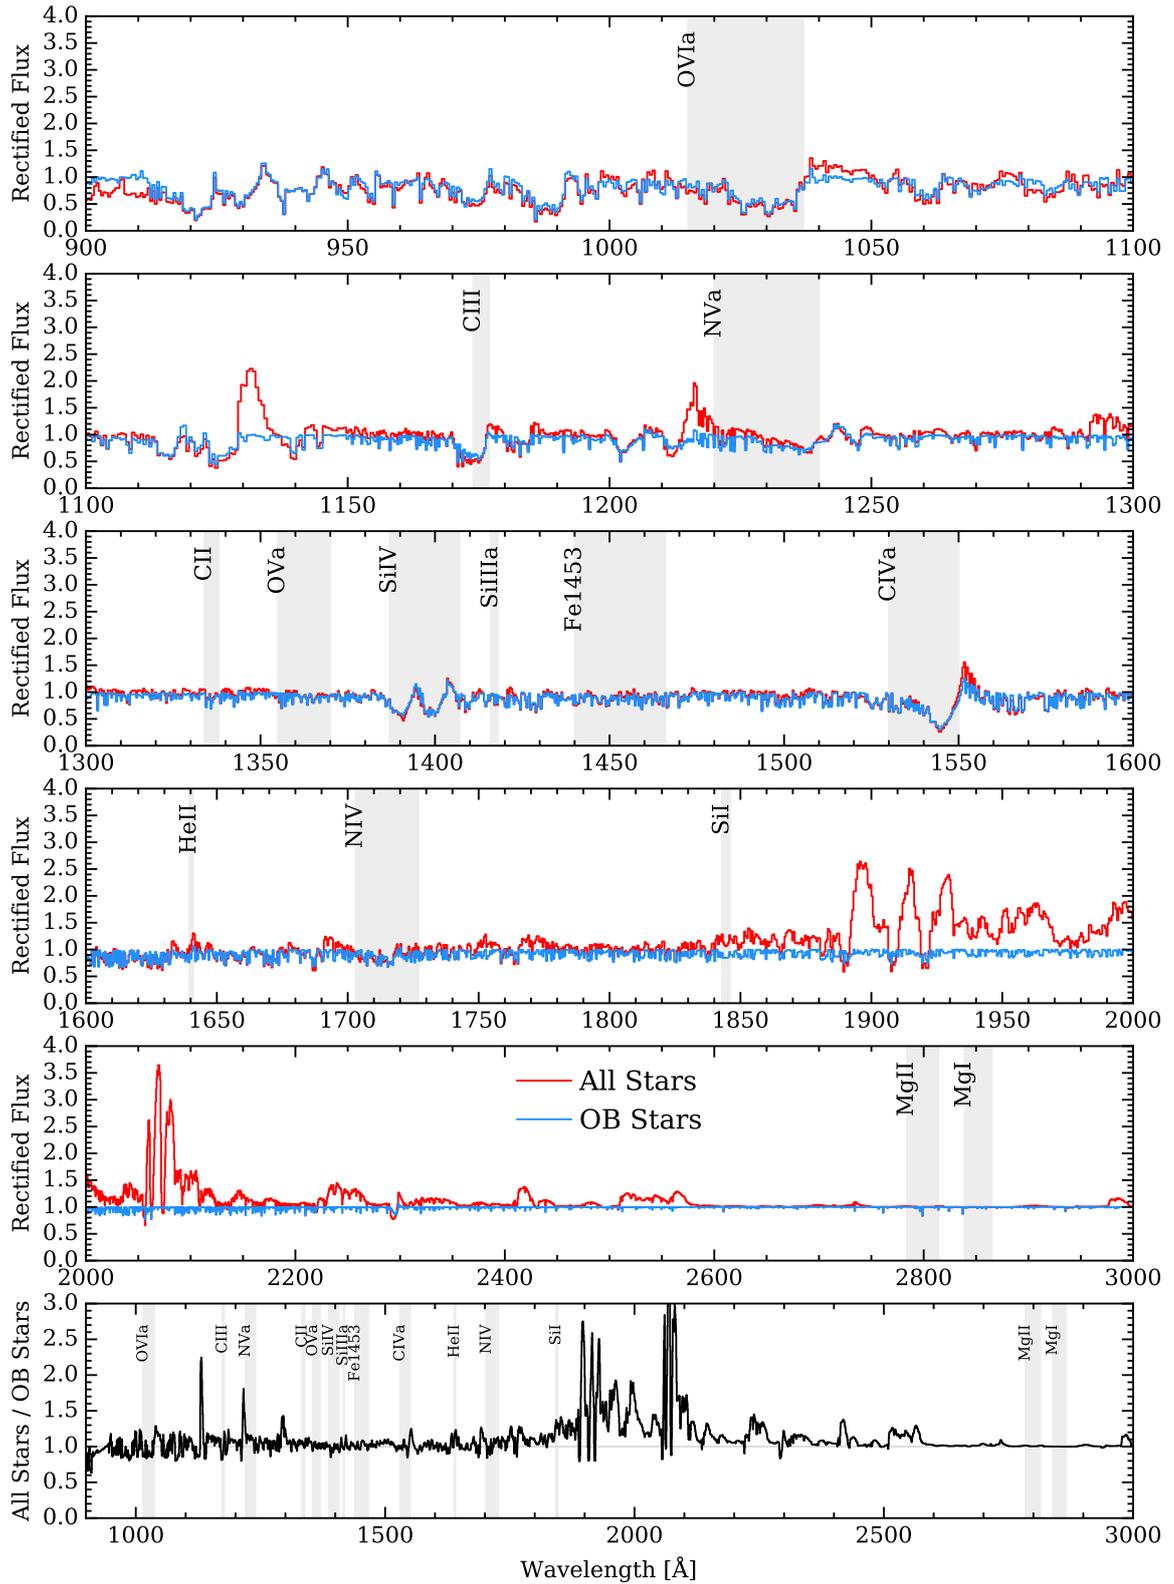
Figure 2.11: Comparison of normalised spectra from a $10^6 \, M_\odot$ cluster at $10^{6.5}$ yr. The spectrum containing only O and B type stars is shown in blue, and the spectrum containing all stars is shown in red. Independent lines from Chapter 4 and their corresponding integration regions are marked in grey. Many of the larger differences are outside the regions covered by our lines of interest however, and including these features would assist in determination of atleast the cluster age.
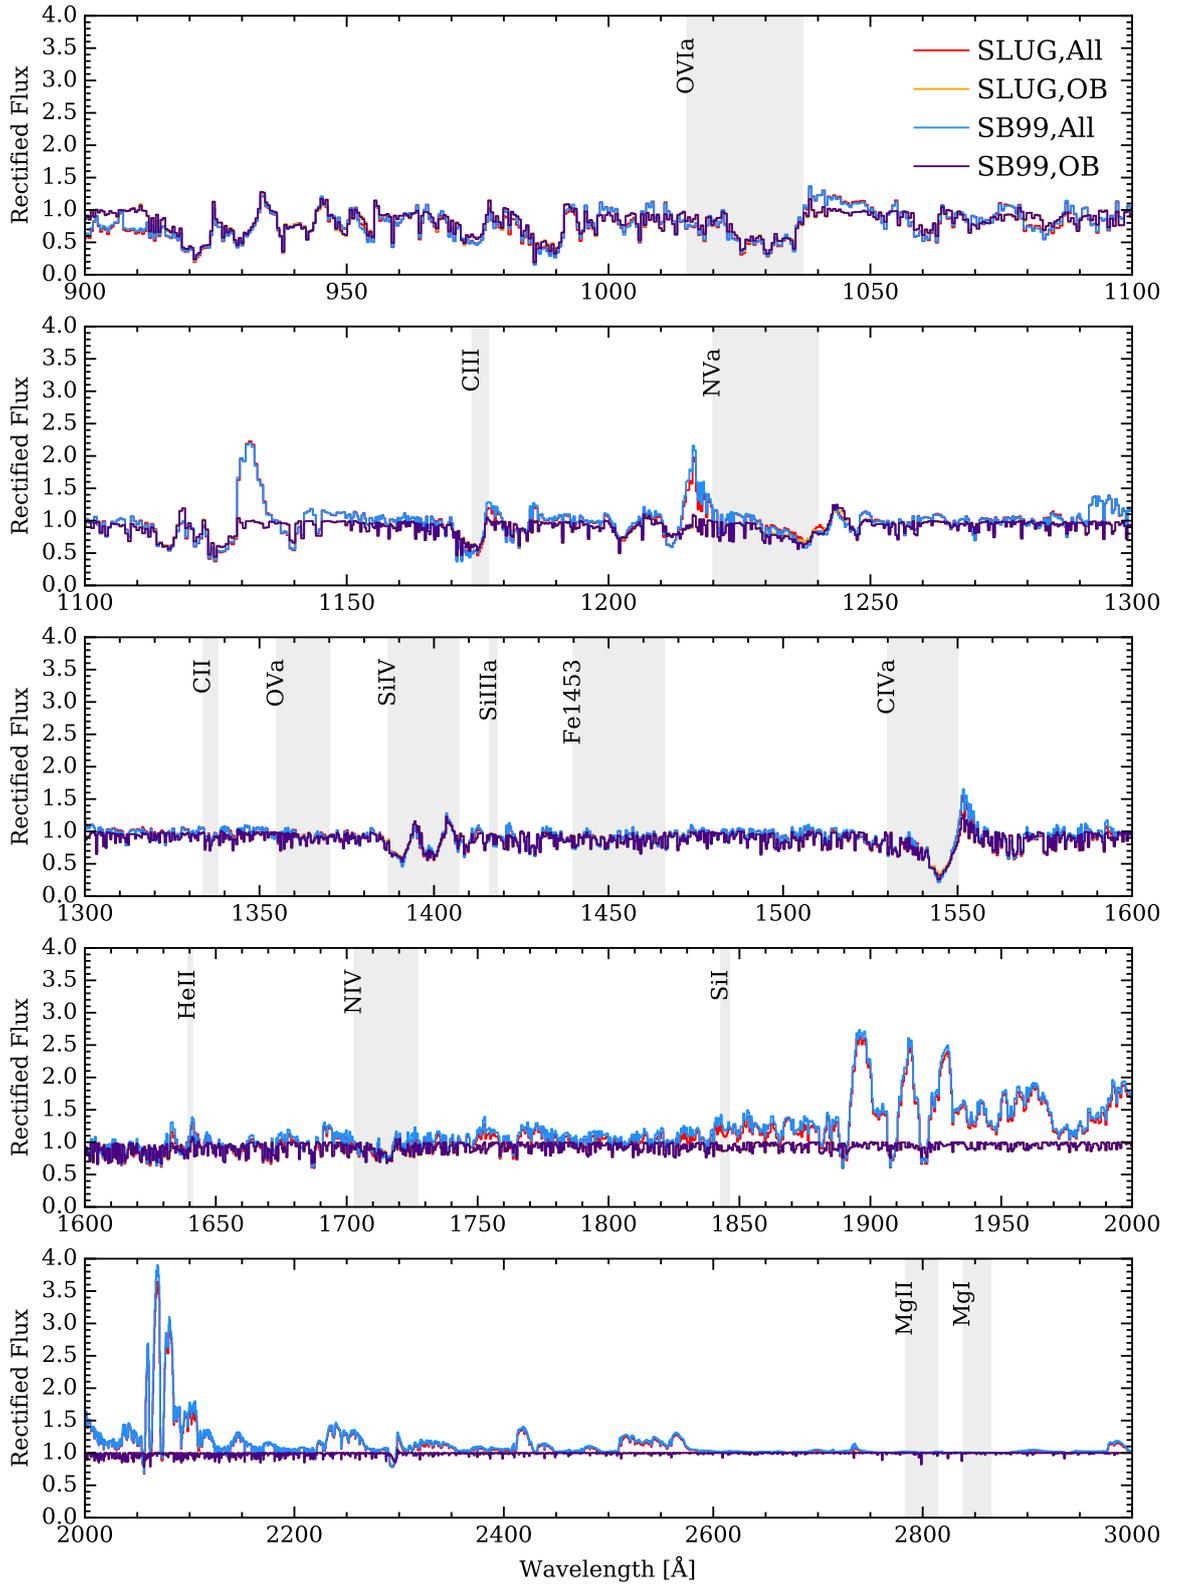
Figure 2.12: Comparison of normalised spectra from a $10^6 \, \mathrm{M}_\odot$ cluster at $10^{6.5}$ yr, generated with a top heavy Kroupa-like IMF with $\alpha_3 = -1.5$. The spectrum containing only O and B type stars is shown in orange (SLUG) and purple (STARBURST99), whereas the spectrum containing all stars is shown in red (SLUG) and blue (STARBURST99). Independent lines from Chapter 4 have their integration regions are marked in grey. The results are comparable between STARBURST99 and SLUG, and many of the larger differences are outside the regions covered by our lines of interest.

observational data).

CHAPTER 3

EXPLORING IMF VARIABILITY IN NGC 628

Preamble

The following chapter is a reproduction of the accepted manuscript for the following paper:

The paper is as it was when accepted, save for minor spelling, citation, and grammar corrections, and the removal of redundant columns from tables such that they fit within the margins of the thesis.

In this paper we describe the development of the new variable IMF mode for SLUG (described more fully in the previous chapter), and then apply these new capabilities in a study of the nearby galaxy NGC 628, building on the previous work carried out by Krumholz et al. (2015b). In that study broad-band photometry of unresolved stellar populations was analysed using SLUG models to predict

the physical parameters of the unresolved star clusters in two galaxies (of which NGC 628 was one, the other being NGC 7793).

Using LEGUS broad-band photometry, in conjunction with large libraries of model star clusters generated using SLUG, we calculate posterior PDFs for the parameters that define the unresolved stellar populations observed in the survey, namely the cluster mass, age, extinction, and in this case the high-mass slope of the IMF. The inclusion of the IMF in the analysis, and therefore the ability to calculate posterior PDFs for the slope parameter, represents a novel feature of this study, as in the previous work Krumholz et al. (2015b) the IMF was a fixed parameter assigned to the catalogue of clusters before analysis. We find very broad posterior PDFs for the IMF, and note a degeneracy between its slope and the mass of the stellar population (which we then quantify). The recovered IMF slope is also found to be sensitive to our choice of priors. In an attempt to break the mass-IMF degeneracy, we include Hα data from the companion Hα-LEGUS survey, finding that this improved our ability to determine the age of a cluster, removing many occurrences of multi-peaked posterior PDFs. However, the improvement in our recovery of the IMF slope parameter was not found to be significant. Using mock observations generated from simulated clusters with SLUG, we then quantify the degeneracies between the physical parameters, and are able show that providing independent constraints on the cluster mass allows for greater accuracy in recovering the slope of the IMF.

Abstract

We present the implementation of a Bayesian formalism within the Stochastically Lighting Up Galaxies (SLUG) stellar population synthesis code, which is designed to investigate variations in the initial mass function (IMF) of star clusters. By comparing observed cluster photometry to large libraries of clusters simulated with a continuously varying IMF, our formalism yields the posterior probability distribution function (PDF) of the cluster mass, age, and extinction, jointly with the parameters describing the IMF. We apply this formalism to a sample of star clusters from the nearby galaxy NGC 628, for which broad-band photometry in five filters is available as part of the Legacy ExtraGalactic UV Survey (LEGUS). After allowing the upper-end slope of the IMF (α_3) to vary, we recover PDFs for the mass,

age, and extinction that are broadly consistent with what is found when assuming an invariant Kroupa IMF. However, the posterior PDF for α_3 is very broad due to a strong degeneracy with the cluster mass, and it is found to be sensitive to the choice of priors, particularly on the cluster mass. We find only a modest improvement in the constraining power of α_3 when adding Hα photometry from the companion Hα-LEGUS survey. Conversely, Hα photometry significantly improves the age determination, reducing the frequency of multi-modal PDFs. With the aid of mock clusters we quantify the degeneracy between physical parameters, showing how constraints on the cluster mass that are independent of photometry can be used to pin down the IMF properties of star clusters.

3.1 Introduction

A core parameter that is common to many astrophysical problems is the stellar initial mass function (IMF), which describes how stellar masses are distributed across a group of stars at the time of their birth. Constraining the form and evolution of the IMF as a function of galaxy properties is of great importance in our theoretical understanding of how star formation proceeds in galaxies and star clusters. Moreover, as the form of the IMF will dictate the makeup of a given star cluster, and hence its mass-to-light ratio, any inference an observer makes about the properties of a given unresolved stellar population via photometry will be strongly dependent on the assumed shape of the IMF (see e.g. Bell et al., 2003).

The IMF is traditionally accepted to be universal, invariant both in time and with respect to the properties of the local environment. Two commonly adopted functional forms that describe the IMF are the broken power law of Kroupa (2001) and the lognormal form of Chabrier (2003) which features a power law tail at the high mass end ($M \geq 1 \, \mathrm{M_\odot}$). Both these functional forms possess a high-mass slope that is consistent with that of the power-law IMF originally presented by Salpeter (1955), with index $\alpha = -2.35$, . This form of the IMF appears to hold in most nearby environments observed to date (Massey et al., 1995; Bastian et al., 2010).

Although theoretical predictions exist for the variation of the IMF with, for example, metallicity and gas temperature (Adams & Fatuzzo, 1996; Bonnell et al., 2006; Krumholz, 2011), at present there is no widely accepted evidence of IMF variation. There is, however, an increasing body of literature supporting the

notion of a non-universal IMF, with variations across different environments. For instance, van Dokkum & Conroy (2010) found evidence suggesting a bottom-heavy IMF in early-type galaxies based on the strength of distinctive absorption lines (NaI doublet and Wing-Ford molecular FeH band), which can be modelled with an overabundance of low mass stars in massive galaxies. Similar detections have been reported in the more recent literature (Conroy & van Dokkum, 2012b), with further evidence in support of a steeper IMF in massive elliptical galaxies coming from dynamical measurements (Cappellari et al., 2012). While both methods consistently imply more bottom-heavy IMFs in massive galaxies, they appear to disagree in detail when applied to individual galaxies rather than to a galaxy population (see Smith, 2014). However, a more recent study by Lyubenova et al. (2016) using consistent data sets has not found this disagreement to be apparent.

At the other end of the galaxy mass function, recent observations have also revealed apparent correlations between the Hα-to-FUV luminosity ratio and the star formation rate (SFR) (Lee et al., 2009, 2016) in dwarf galaxies, which can be attributed to variation in the high-mass end of the IMF. Similar suggestions of variation in the IMF have been made based on comparisons between the Hα-to-FUV luminosity ratio and Hα and R-band surface brightness in starburst galaxies (Meurer et al., 2009). Other indications of a variable IMF have been gleaned using techniques involving the Hα equivalent width and galaxy colours (Gunawardhana et al., 2011), and chemical evolution arguments (Matteucci, 1994; Thomas, 1999).

Attempts to explain these observations include the framework of the Integrated Galactic Initial Mass Function (IGIMF) (Kroupa, 2001; Kroupa & Weidner, 2003; Weidner et al., 2011; Kroupa et al., 2013). In this theory, the most massive stars allowed to form in a cluster (corresponding to the high-mass cutoff of the IMF) are correlated with the cluster mass, which in turn depends on the galaxy SFR. These correlations suppress the formation of massive stars in dwarf galaxies, resulting in a lower Hα-to-FUV ratio, in line with the aforementioned observations.

However, the very same observations can be explained by alternative models that do not resort to deterministic variations of the IMF. For example, stochasticity has been shown to suitably describe the Hα-to-FUV ratio correlation in dwarf galaxies without the need for an explicit correlation between star and cluster masses. Indeed, a combination of the random sampling of clusters from a cluster mass function combined with a random sampling of stars from the IMF (Fumagalli

et al., 2011), appears to satisfactorily reproduce the observed Hα-to-FUV ratio in dwarf galaxies. Similar results have been reported by Calzetti et al. (2010) and Andrews et al. (2013), who have studied the ratio of Hα to cluster mass, finding that the IMF upper end slope has no clear mass dependence (see also Corbelli et al., 2009). Andrews et al. (2014) find that stochasticity is similarly effective in explaining observations of the spiral galaxy M83. Additional challenges to the IGIMF theory are discussed, for instance, in Oey et al. (2013) and Lamb et al. (2016). These authors present observations of what are likely to be solitary O and B stars with no visible evidence of escape from nearby clusters. Such an occurrence would violate the requirement of the IGIMF that the maximum stellar mass correlates with cluster mass. However, previous indications of the possibility of solitary OB stars (de Wit et al., 2004, 2005; Lamb et al., 2010) have been disputed by Gvaramadze et al. (2012), and more recently Stephens et al. (2017) have made arguments in opposition to stochastic IMF sampling as a rebuttal of the IGIMF through observations of compact clusters of pre-main-sequence stars around massive stars previously thought to be isolated.

This ambiguity surrounding the universality of the IMF calls for the development of new techniques and methods that explore the properties of the IMF in different environments. To this end, Bayesian techniques have recently been applied to stellar population synthesis (SPS) models, in an attempt to infer the parameters that describe the IMF. A particularly appealing feature of these Bayesian approaches is the possibility to marginalise over other "nuisance" parameters, and to explore to what extent the IMF is degenerate with respect to other parameters used in SPS models.

For example, Weisz et al. (2015) use Bayesian techniques combined with colour-magnitude diagrams of resolved star clusters to derive a constant IMF high-end slope that is both slightly steeper than the canonical Kroupa value of α_3=-2.3 and constant in Andromeda. Dries et al. (2016) have also recently presented an example of how Bayesian methods can be applied to this problem, successfully recovering the input parameters describing the IMFs in a selection of mock simple stellar populations produced with varying IMFs, provided however that a representative set of stellar templates is available during the analysis.

Following a similar idea, in this work we aim to combine Bayesian techniques with the stochastic SPS code SLUG (Stochastically Lighting Up Galaxies; da Silva

et al., 2012, 2014; Krumholz et al., 2015a) to explore the IMF properties of star clusters. This analysis builds on the method described in Krumholz et al. (2015b), who combined SLUG with Bayesian inference to constrain the mass, age, and extinction of clusters from broad-band photometry. In this work, expanding on their previous analysis, we develop a formalism to also investigate whether the high-mass end of the IMF can be constrained through photometry of individual clusters, while tracking at the same time degeneracies between the other parameters that regulate the mass-to-light ratio of the stellar populations. To achieve this goal, we first need to extend the current capabilities of SLUG, as we describe briefly in Section 3.2, and in Appendix 3.5. Specifically, we enable a new mode in SLUG through which libraries of simulated clusters can be constructed using a continuous distribution of arbitrary IMFs. These libraries can then be compared to actual observations using Bayesian techniques to infer the posterior probability distribution function (PDF) for the parameters that describe the cluster physical properties, including the IMF, within SPS models.

Following the implementation of this extension in the SLUG code, in Section 3.3, we apply this extended Bayesian formalism to photometric data on star clusters from the Legacy ExtraGalactic UV Survey (LEGUS). LEGUS is a Cycle 21 *Hubble Space Telescope* Treasury programme (Calzetti et al., 2015), designed to investigate star formation and its relation with the galactic environment across a representative sample of nearby galaxies within 12 Mpc, on scales ranging from individual stars up to kiloparsec-scale structures. In addition to this application, in Section 3.4 we construct a library of mock star clusters, which we use to explore from a theoretical point of view how accurately the upper-end slope of the IMF (hereafter the α_3 parameter) can be recovered using photometry from star clusters. In doing this, we also quantify the degeneracy between IMF parameters and other physical quantities, such as mass and age. Summary and conclusions follow in Section 3.5, where we discuss future improvements to our method and describe future experiments that could yield tight constraints on the IMF of star clusters. Throughout this work, unless otherwise specified, magnitudes are in the Vega system.

3.2 The SLUG software suite

The core of the SLUG suite is an SPS code that, unlike the majority of traditional SPS codes, such as STARBURST99 (Leitherer et al., 1999), includes the effects of stochastic sampling. Stochastic sampling of the IMF was pioneered in the MASS-CLEAN software of Popescu & Hanson (2009). In SLUG this stochasticity is achieved through generating galaxies by drawing stars and clusters from underlying PDFs which define the IMF and the cluster mass function (CMF) in the simulated galaxy. These populations are then evolved through time to the desired age, accounting for cluster disruption. We refer the reader to previous work for a detailed description of the algorithms adopted in SLUG (e.g. da Silva et al., 2012; Krumholz et al., 2015a). Briefly, in its basic "cluster mode", SLUG constructs star clusters by drawing stars from an IMF until a desired target mass is reached. This sampling method allows SLUG to accurately capture the effects of stochasticity, which are particularly relevant for low mass clusters (e.g. Cerviño et al., 2002; Elmegreen, 2002; Haas & Anders, 2010). Moreover, the ability to populate star clusters (or galaxies) by randomly drawing from the IMF and CMF provides a convenient way to generate large libraries of realistic clusters which are representative of the observed population in nearby galaxies (e.g. Krumholz et al., 2015a).

Along with the core SPS code, the SLUG distribution includes the CLUSTER_SLUG software, which is a package designed for Bayesian analysis of observed cluster photometry (Krumholz et al., 2015a). Specifically, CLUSTER_SLUG uses Bayesian inference to recover the posterior PDFs for physical parameters (e.g. mass, age, extinction) that influence the photometric properties of a given star cluster. To this end, the CLUSTER_SLUG package makes use of a large 'training set' of SLUG models that, through Bayesian analysis combined with kernel density estimation (KDE) techniques (da Silva et al., 2014; Krumholz et al., 2015a), yields the PDFs for the parameters of interest given the observed broad-band photometry of a stellar cluster.

In this work, we extend this technique to the case of varying IMFs. Within the Bayesian framework of CLUSTER_SLUG, the posterior PDFs of parameters describing the functional form of the IMF (for instance, the high end slope α_3 in the Kroupa IMF) can be derived from cluster photometry, provided that SLUG libraries containing a continuous distribution of clusters with respect to the parameter of

interest (e.g. α_3) are available. To this end we have implemented a variable IMF capability in SLUG, with which libraries of simulated clusters can be generated with continuously variable IMF parameters. This modification is discussed in detail in Appendix 3.5.

3.3 Applying SLUG to observed star clusters

Having enabled a variable IMF capability in SLUG, we now combine a large library of simulated star clusters with LEGUS broad-band photometry. This photometry is taken from a star cluster catalogue of the galaxy NGC 628. We use this combination to investigate whether it is possible to constrain the upper-end slope of the IMF with our Bayesian technique.

3.3.1 The NGC 628 photometric catalogues

For this analysis, we make use of the multiwavelength imaging available from the LEGUS project. As part of the LEGUS survey (see Calzetti et al., 2015), local volume galaxies have been imaged with five broad-band filters using the Wide-Field Camera 3 (WFC3), with parallel observations from the Advanced Camera for Surveys (ACS). Of the fifty galaxies in the LEGUS sample, we select NGC 628 for our analysis. The filters used for NGC 628 are as follows: WFC3 F275W, F336W, and F555W, along with ACS F435W and F814W. These may be generally thought of as the NUV, U, B, V, and I bands. In our study, beside broad-band photometry, we also make use of narrow-band Hα data, which have been obtained in the WFC3 UVIS F657N filter as part of the Hα-LEGUS program (Chandar et al. in preparation).

NGC 628 is a nearby face-on grand design spiral galaxy of type SA(s)c at a distance of 9.9 Mpc (Calzetti et al., 2015). Currently the most well studied LEGUS galaxy, it is one of the two galaxies discussed in Krumholz et al. (2015b). Thus, the properties of star clusters in NGC 628 have already been analysed using the SLUG suite for a constant IMF, which provides us with a reference point to which we can compare our results when variations of the IMF are included. Two pointings are available for this galaxy (NGC 628 East and NGC 628 Centre), and here we focus on the the East pointing (hereafter NGC 628E) as done in Krumholz et al. (2015b).

Moreover, Hα data are available for the NGC 628E pointing, which allows us to explore the effects of including photometry that is sensitive to massive stars in our analysis.

We refer the readers to relevant LEGUS papers for a detailed description of the data processing techniques, and particularly on how cluster catalogues are generated (e.g. Adamo et al., 2017; Calzetti et al., 2015; Grasha et al., 2015). Here we provide only a brief description of the most relevant procedures. We make use of clusters that have been detected in at least 4 filters. For all clusters, both an aperture correction and correction for foreground Galactic extinction have been applied following standard procedures. All the clusters have been visually classified, according to the LEGUS definition of class 1-3. Class 1 objects are compact and centrally concentrated clusters, class 2 objects are clusters with elongated surface-brightness profiles, and class 3 objects loosely comprise associations or systems with multiple resolved peaks in surface brightness. We omit class 4 objects from our analysis, as they represent for most part isolated stars, background galaxies, or artefacts.

With this selection, we are left with a cluster catalogue comprised of 259 clusters from NGC 628E, as in Krumholz et al. (2015b) and Adamo et al. (2017). Of these, we further select 241 clusters with a detection in the WFC3 UVIS F657N filter. Of the clusters with positive detections in this filter, 16 are discarded from the following analysis as visual inspection shows that their morphology is indicative of emission not arising from HII regions (Lee et. al. in preparation). Altogether, our selection results in a final catalogue containing 225 clusters in the LEGUS class 1-3.

3.3.2 Building a library of 2×10^8 simulated clusters

The first step towards applying our formalism to the LEGUS cluster catalogue is the generation of a library that provides the training set required for Bayesian analysis in CLUSTER_SLUG. Having enabled the variable IMF mode in SLUG, we now trivially generate an extended library of simulated star clusters, including a continuous range of IMFs with a varying upper-end slope (α_3). Our focus on the massive end of the IMF is justified by the availability of NUV and Hα photometry in LEGUS and Hα-LEGUS.

To construct this reference library with a variable IMF, which we use in all

Name	Tracks	Z	Extinction[a]	ϕ[b]	$\log M$ ($M_\odot$)	$\log T$ (yr)	A_V (mag)	# Realisations	Variable Parameters
pad_020_vkroupa_MW[c]	Padova AGB	0.020	MW	0.5	2–8	5–10.18	0–3	2×10^8	$-3.0 \le \alpha_3 \le -1.5$
pad_020_kroupa_MW[d]	Padova AGB	0.020	MW	0.5	2–8	5–10.18	0–3	1×10^7	–

[a] MW = Milky Way extinction curve.

[b] ϕ is the fraction of ionising photons that produce nebular emission within the aperture, combining the effects of a covering fraction less than unity and some portion of the ionising photons being absorbed directly by dust.

[c] Variable IMF model.

[d] Constant IMF model (Krumholz et al., 2015b).

Table 3.1: Table of the parameters used in generating the two reference libraries of SLUG star clusters used in this work.

subsequent calculations unless stated otherwise, we extend the fiducial library described in Krumholz et al. (2015b)[1] to include a new dimension for the variable upper-end slope α_3. This library, which we dub `pad_020_vkroupa_MW`, contains 2×10^8 clusters distributed in mass as

$$p_M(\log M) \propto \begin{cases} 1, & 2 < \log M < 4 \\ 10^{-(\log M - 4)}, & 4 \leq \log M < 8, \end{cases} \quad (3.1)$$

where the mass M is in units of $M_\odot$. As in Krumholz et al. (2015b), this choice is purely dictated by computational efficiency, as low mass clusters are cheaper to simulate. Furthermore, with this weighting, we generate the largest number of clusters in the mass range where stochasticity is dominant so as to appropriately sample the dispersion in the population arising from random sampling. As described below, however, the choice of weighting does not affect our analysis as we can apply appropriate weighting schemes to produce arbitrary mass distributions. Similarly, we assume an age distribution defined by

$$p_T(\log T) \propto \begin{cases} 1, & 5 < \log T < 8 \\ 10^{-(\log T - 8)}, & 8 \leq \log T < \log T_{\max}, \end{cases} \quad (3.2)$$

which results in the generation of more young clusters, in which stochastic effects associated with massive stars are more prominent. In these equations the time T is in units of years, and $T_{\max}$ is the maximum age for which stellar tracks are available.

The clusters are populated with stars drawn from a variable IMF. Here, we assume a Kroupa-like IMF, defined as a broken power-law comprised of three segments:

$$\alpha_i = \begin{cases} -0.3, & 0.01 \leq M < 0.08, \ i = 1 \\ -1.3, & 0.08 \leq M < 0.5, \ i = 2 \\ -3.0 \leq \alpha_3 \leq -1.5, & 0.5 \leq M < 120, \ i = 3, \end{cases} \quad (3.3)$$

where α_i is the power law index of the i^{th} segment and the stellar masses are in units of $M_\odot$. The value of the third segment slope (α_3) is drawn from a flat

---

[1]The library is publicly available at www.slugsps.com and contains simulated cluster photometry in the selection of filters used in the LEGUS observations of NGC 628

distribution between -3.0 and -1.5 at the start of each realisation.

Photometry is computed using Padova tracks including thermally pulsating AGB stars (Vassiliadis & Wood, 1993; Girardi et al., 2000; Vázquez & Leitherer, 2005), which sets the maximum age $T_{max} = 15$ Gyr. We assume a metallicity $Z = 0.02$, consistent with the solar value, and we apply a uniform extinction to each cluster assuming a Milky Way extinction curve (Cardelli et al., 1989), which we normalise by drawing the visual extinction A_V from a flat distribution in the range $0 < A_V < 3$. In this work we use the default SB99 spectral synthesis mode of SLUG, which emulates the behaviour of STARBURST99 in choosing which atmosphere model is used for a given star. Details of this implementation can be found in Appendix A2 of Krumholz et al. (2015a). While this set of atmospheres is commonly adopted in SPS modelling, we note that our results will naturally depend on this choice.

We further compute the nebular contribution to the broad-band photometry, assuming that 50% of ionising photons are converted into nebular emission (see also Section 3.3.6). The choice of these parameters is dictated by the properties of NGC 628 (see also Section 3.3.6). Krumholz et al. (2015b) demonstrate that the results of CLUSTER_SLUG analysis are mostly insensitive to the choice of tracks, extinction curve, and metallicity, and so for this study we make use of only one choice for each of these parameters. Metallicity, choice of tracks, and ionisation efficiency do become important when considering clusters with ages less than 3 Myr however.

We do not include a treatment of binary stars or stellar rotation in this study. The presence of binaries (especially among massive stars) has been shown to produce appreciable effects on the integrated light of star clusters (Eldridge & Stanway, 2009), resulting in a less luminous and bluer stellar population (Li & Han, 2008). However, as their inclusion requires several additional parameters, binaries are currently not treated in SLUG. Rotation, which likewise affects results of SPS models (Vázquez et al., 2007; Levesque et al., 2012; Leitherer et al., 2014), can be included through the use of the latest Geneva tracks (Ekström et al., 2012) available in SLUG. However, we choose to consider the non-rotating Padova tracks to include the AGB treatment currently not present for the Geneva tracks with rotation. The effects of rotation can be as much as a 50% increase in luminosity of O stars, along with a factor of 2 increase in luminosity between the UV and NIR (Vázquez et al., 2007; Levesque et al., 2012). However, the effects are only

significant if all the stars have initial rotation velocities of $\approx 300\,\mathrm{km\,s^{-1}}$ (Eldridge & Stanway, 2009).

Throughout this analysis, we also make use of the fiducial library of Krumholz et al. (2015b) with a constant Kroupa IMF to compare our results to the case of a non-variable IMF. A summary of the properties of these reference libraries is given in Table 3.1. Our library contains 2×10^8 clusters, which is a factor of 20 larger than the fiducial library of Krumholz et al. (2015b). This choice is dictated by the need to increase the number of clusters to account for the extra free parameter of the IMF slope α_3, while keeping the library size commensurate to the available computational resources. We have performed tests using smaller libraries of 10^8 clusters, finding that we have reached a satisfactory convergence.

3.3.3 Bayesian techniques and choice of priors

With a library of 2×10^8 clusters featuring variable IMFs in hand, we next make use of the Bayesian analysis tool CLUSTER_SLUG to handle the computation of the posterior PDFs of the α_3 parameter. A description of the algorithms adopted in CLUSTER_SLUG has been presented in da Silva et al. (2014) and Krumholz et al. (2015a), and we refer the readers to this work for details. Briefly, the Bayesian analysis tool CLUSTER_SLUG takes a set of absolute magnitudes M_F over a given set of filters, along with the associated errors ΔM_F. By means of a large library of SLUG models, CLUSTER_SLUG exploits kernel density estimation to map the observed photometry into a PDF for the cluster physical parameters. In previous work, the only parameters under consideration were the cluster mass M, age T, extinction A_V. Here, we further extend CLUSTER_SLUG to also output PDFs for IMF parameters that are chosen to be variable, which in our case is the α_3 parameter. During this calculation we set the kernel density estimation bandwidth to $h = 0.1\,\mathrm{dex}$ for the physical variables and to $h = 0.1\,\mathrm{mag}$ for the photometric variables, as in Krumholz et al. (2015b). During testing, a shift to a smaller bandwidth of $h = 0.05$ did not appreciably affect the results of the analysis.

As discussed in the previous section, the simulated clusters in our library are not weighted evenly in parameter space for reasons of computational efficiency and to better sample the range of parameters where stochasticity is most relevant. It therefore becomes necessary to "flatten" the input library, prior to feeding it into

CLUSTER_SLUG. This step is accomplished simply by inverting the weights chosen during the calculation of the library of SLUG clusters. After the library has been flattened, we further need to make a choice of the priors to be used during the Bayesian analysis. Throughout this work, we explore the effects of two different choices of prior. Our first choice is a flat prior, with $p_{\mathrm{prior}} = 1$. Our second choice is a prior in the form of:

$$p_{\mathrm{prior}}(\mathbf{x}) \propto M^{-2}T^{-0.5}, \tag{3.4}$$

for $T > 10^{6.5}\,\mathrm{yr}$. This functional form is physically motivated by the cluster mass distribution, which is observed to follow a power-law with index -2. The time dependence accounts instead for the cluster disruption (see Krumholz et al., 2015b). For young ages of $T \leq 10^{6.5}\,\mathrm{yr}$, where cluster disruption is largely ineffective, we switch to a prior $p_{\mathrm{prior}}(\mathbf{x}) \propto M^{-2}$. As evident from the above equations, we elect to adopt a flat prior for both the IMF high-end slope α_3 and the extinction A_{V} in both cases.

3.3.4 Analysis of NGC 628E with broad-band photometry

Our next step after constructing a library of simulated clusters with variable IMF is to apply our CLUSTER_SLUG analysis to the LEGUS broad-band photometry of NGC 628E. With our extension of the SLUG code that now handles the case of a variable IMF, we use the CLUSTER_SLUG package to calculate the one dimensional (1D) posterior PDFs of all the physical parameters of interest (mass, age, extinction, and α_3) for the 225 clusters we have selected from NGC 628E. Two dimensional PDFs are also computed, which are useful to explore correlations among parameters. At this stage, we apply both flat priors and the physically motivated priors previously presented in Krumholz et al. (2015b) (see Equation 3.4). In this section, we begin our analysis by considering the results obtained with the physically motivated priors, but in Section 3.3.5 we explicitly investigate the effects of the choice of priors.

Before considering the cluster population as a whole, it is instructive to inspect the corner plots for two example clusters, which we choose to bracket the range of behaviour seen in our analysis. As a first example, we show the corner plot for cluster ID 56 in Figure 3.1, where the posterior PDFs are obtained using the physically motivated priors. This is the same cluster that is presented as an example
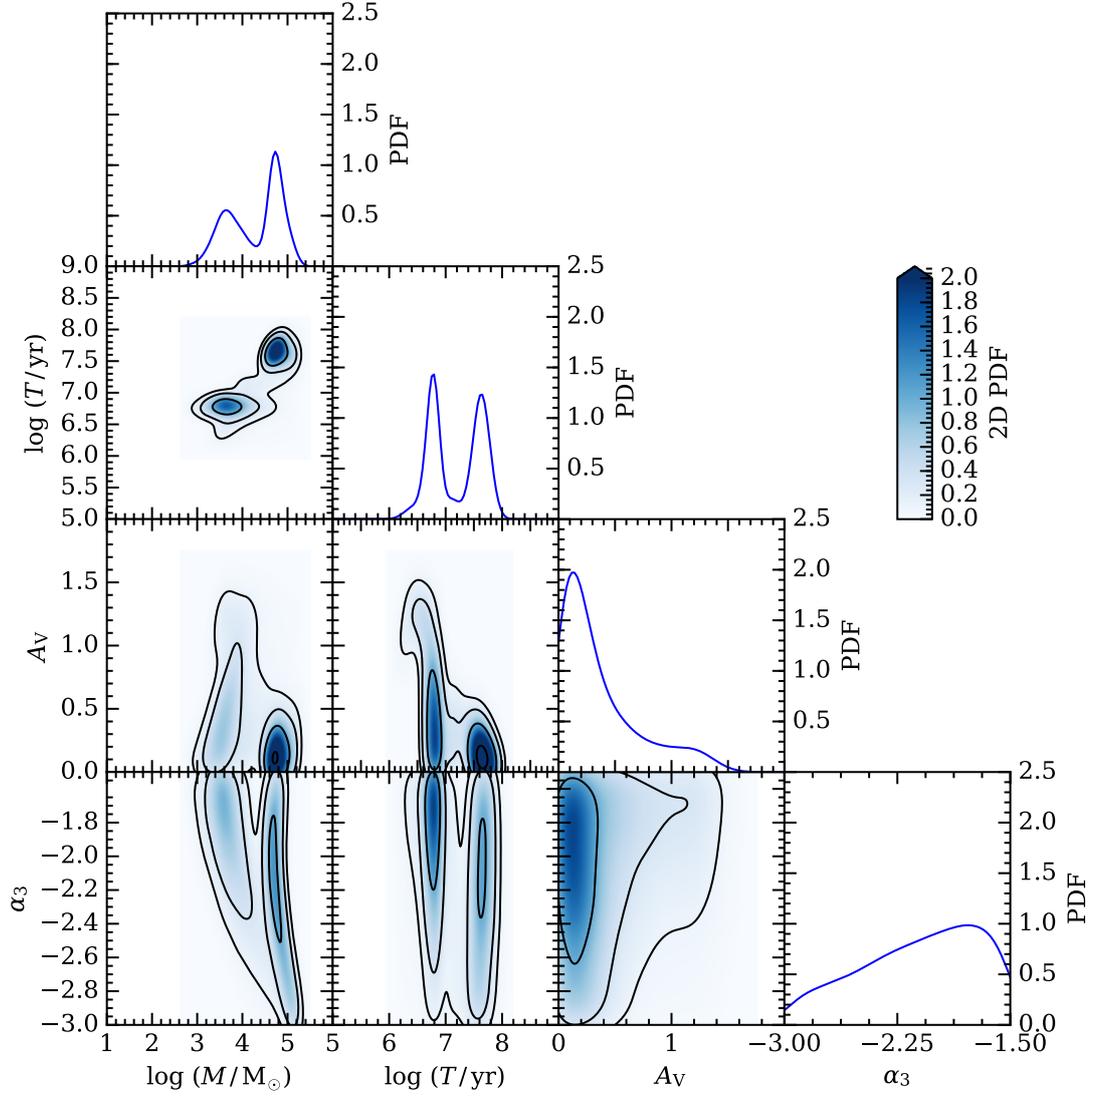
Figure 3.1: Example of a triangle plot for cluster ID 56 in NGC 628E, computed assuming the physically motivated priors from Krumholz et al. (2015b). The 1D posterior PDFs for $\log M$, $\log T$, A_V, and α_3 that are constructed marginalising over all the other parameters are shown in the top panels of each column. The contour plots show instead the joint posterior PDFs, where the intensity reflects the probability density as indicated by the colour bar. In each panel, contours are spaced in intervals of 0.2 unit element. All the PDFs are normalised to have a unit integral.
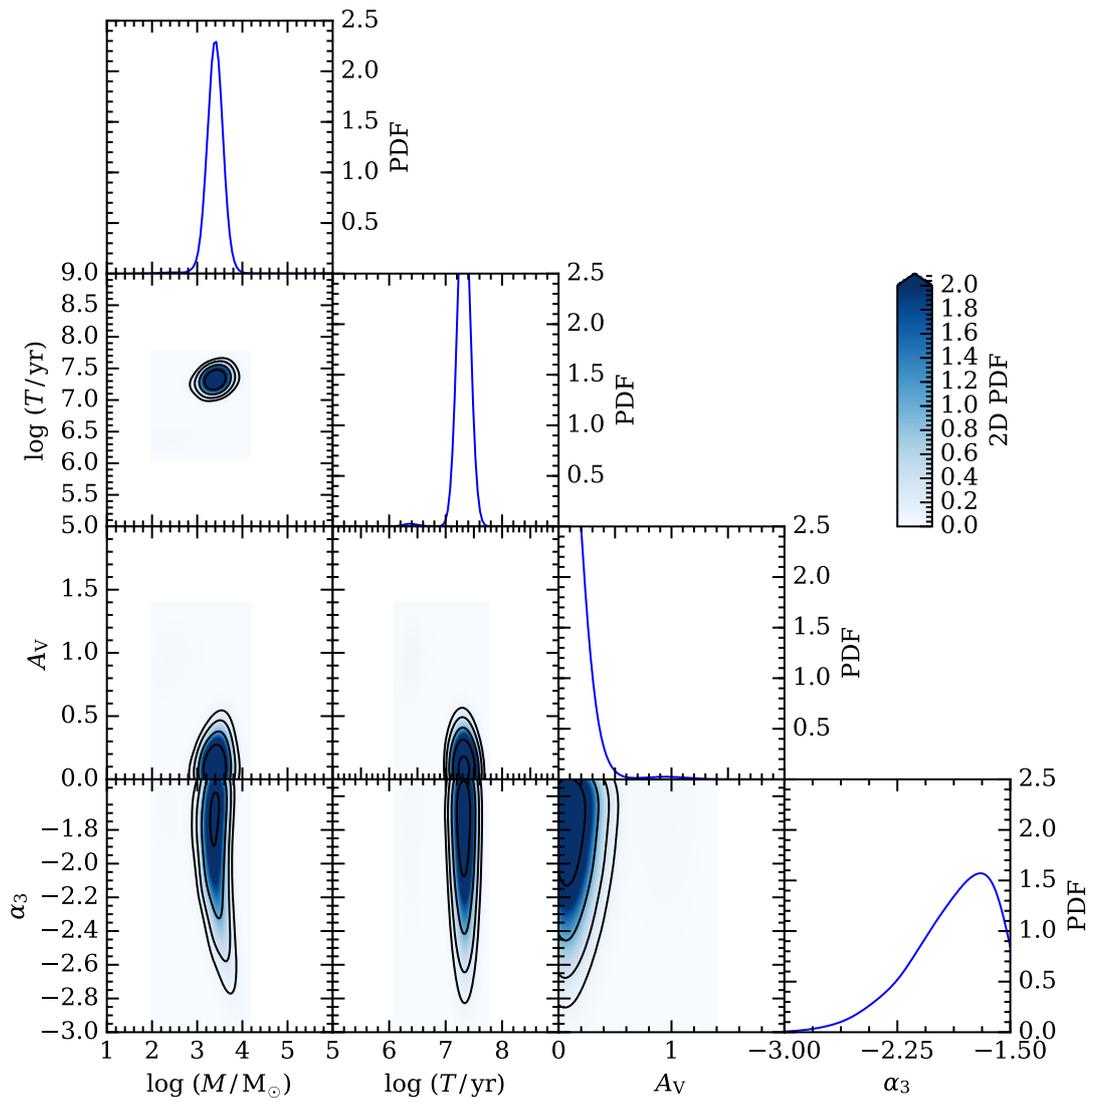
Figure 3.2: As Figure 3.1, but for cluster ID 591.
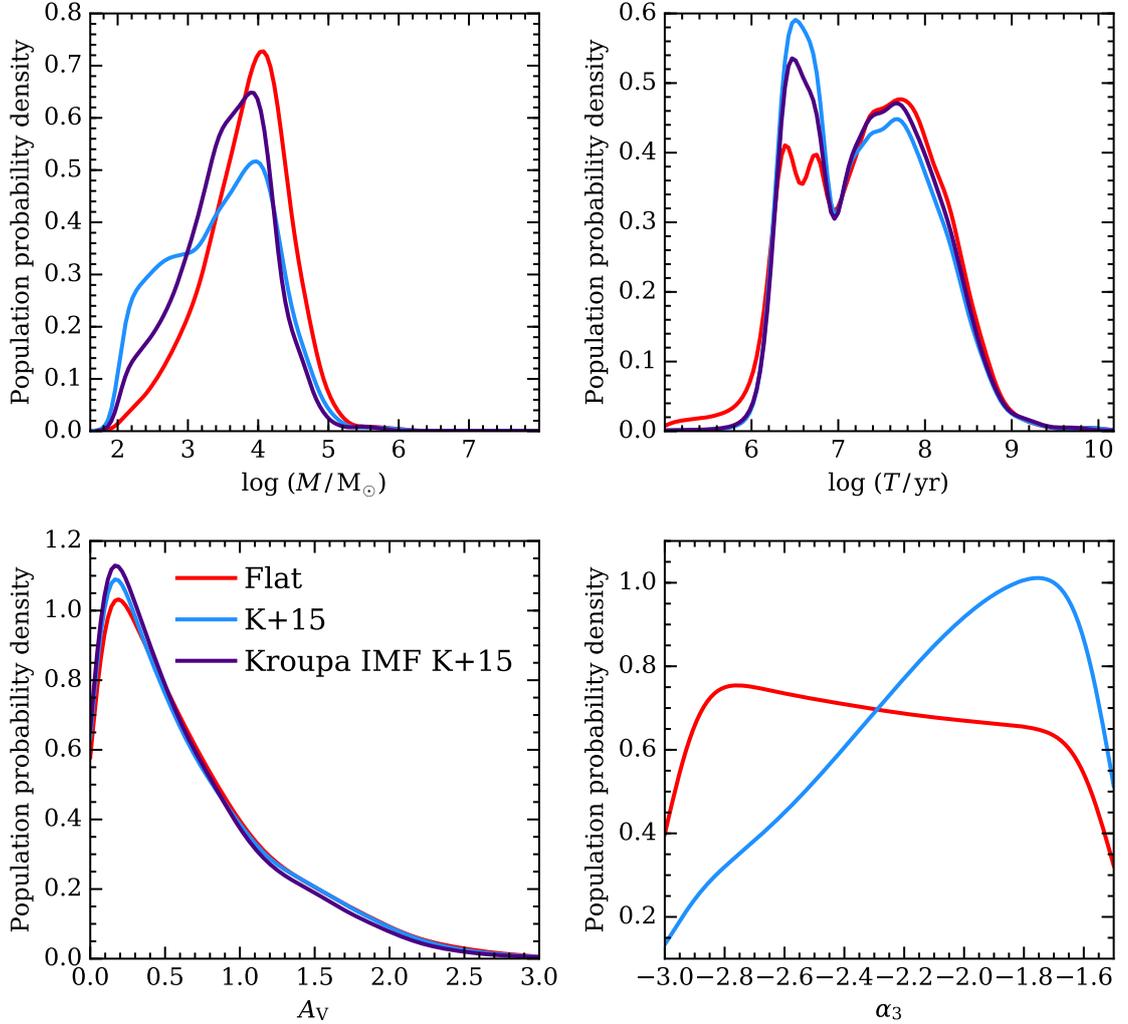
Figure 3.3: The cumulative posterior PDFs for the star clusters in our sample from NGC 628E, computed assuming both flat priors (in red, labelled as Flat) and the physically motivated priors of Krumholz et al. (2015b) (in blue, labelled as K+15). Also included are the cumulative PDFs calculated using the fiducial library of Krumholz et al. (2015b) with a constant Kroupa IMF and physically motivated priors (in purple, labelled as Kroupa IMF K+15). The distributions are comparable overall. There is significant deviation between the two priors in the case of α_3, with the Krumholz et al. (2015b) priors having a significant preference towards more shallow IMF slopes.

in Figure 8 of Krumholz et al. (2015b), and having it as an example here allows us to make a direct comparison with previous work. However, we note that it is not part of the reduced 225 cluster catalogue as it has no observed Hα emission. The shape of the posterior PDFs for mass, age, and extinction computed for a variable IMF slope are quite similar to those recovered by Krumholz et al. (2015b) with a fixed IMF, with the exception of a small excess of probability towards lower masses and ages. This cluster also happens to be representative as to the behaviour of a significant fraction of the clusters in our sample, which exhibit multiple peaks in the posterior PDF for mass and age. Furthermore, and most significantly, the 1D posterior PDF for α_3 is very broad, spanning the full range of possible values with a modest preference for shallow slopes. An inspection of the joint posterior PDFs further reveals that the IMF slope parameter is largely insensitive to age and extinction, and is degenerate with respect to mass.

At the opposite end of the spectrum, we show in Figure 3.2 the case of a cluster (ID 591) for which mass, age, and extinction are well constrained, as is clear from the sharp posterior PDFs of these quantities. Despite the fact that these physical parameters are well constrained, the posterior PDF of α_3 is still quite broad, with most of the probability being contained in the interval $\alpha \approx (-2.3, -1.5)$. As was seen in cluster 56, the corner plot reveals a modest preference of shallow IMF exponents, with $\alpha \gtrsim -2.2$.

The features highlighted in these examples appear to be quite general of the entire sample. This is shown in Figure 3.3, where we display the cumulative posterior PDFs for the entire sample, which we obtain by combining the 1D posterior PDFs for each cluster. It must be noted that we do not analyse the clusters jointly for this, and that the PDFs for the clusters are combined following our analysis.

In addition to the two choices of prior, which we will discuss below, we also include in this figure the cumulative PDFs calculated using the fiducial library of Krumholz et al. (2015b), assuming a constant IMF. A comparison between the posterior PDFs for age and extinction obtained for a constant versus variable IMF (the purple and blue lines respectively) reveals excellent agreement between the two. Besides validating our procedure, the fact that the age and extinction PDFs computed in these two cases are virtually indistinguishable implies that α_3, A_{V}, and T have modest covariance, as suggested by the shape of the joint PDFs in Figure 3.1 and Figure 3.2.

A similar conclusion holds for the posterior PDF of the mass, although differences between the constant and variable IMF cases become more noticeable. Indeed, there is some divergence both at very low and very high masses, particularly with the variable IMF case being skewed on average towards low masses. The origin of this shift, albeit modest, can be attributed to the degeneracy between mass and IMF highlighted above. Indeed, compared to the case of a fixed IMF with $\alpha_3 = -2.3$, the analysis which includes a variable IMF seems to prefer, on average, lower cluster masses, which are populated by stars that are on average more massive. This effect is indeed visible in the cumulative posterior PDF for α_3, which is skewed towards α_3 values that are shallower than the canonical Salpeter value.

Moreover, Figure 3.3 shows that, in line with what is found when examining the PDFs for individual clusters, the cumulative posterior PDF for α_3 is quite broad, implying that only weak constraints on the slope of the IMF can be obtained when comparing broad-band photometry with simulated clusters. As noted above, however, the posterior PDF appears skewed towards a shallow IMF, with $\alpha_3 \gtrsim -2.5$. This effect is modest, and not particularly statistically significant. We can nevertheless inspect the median α_3 recovered for each cluster, and study whether systematic trends can be found as a function of mass and LEGUS class. This analysis is shown in Figure 3.4. When considering the physically motivated priors (squares and diamonds), a noticeable correlation with mass appears. In this figure, we have also separated clusters by their LEGUS class. We see that the majority of the class 3 objects (squares), which include clusters with multiple peaks and what are likely to be associations, are recovered at low masses (see also Adamo et al. (2017)). These objects are characterised by shallow slopes of the upper-end of the IMF. Conversely, class 1 and class 2 clusters (diamonds) are on average more massive, although a fit through the data still indicates a correlation between the median α_3 and the cluster mass.

The correlation between the value of α_3 and the cluster mass can be explained as follows. Due to the degeneracy between mass and IMF, the observed light can be modelled either by a relatively massive cluster with canonical $\alpha_3 = -2.3$, or by a lower mass cluster with a more shallow IMF. The choice of a physical prior of the form $\propto M^{-1}$ skews the mass PDF towards low values (see also Section 3.3.5), making the latter case (lower cluster mass and hence a more shallow IMF)
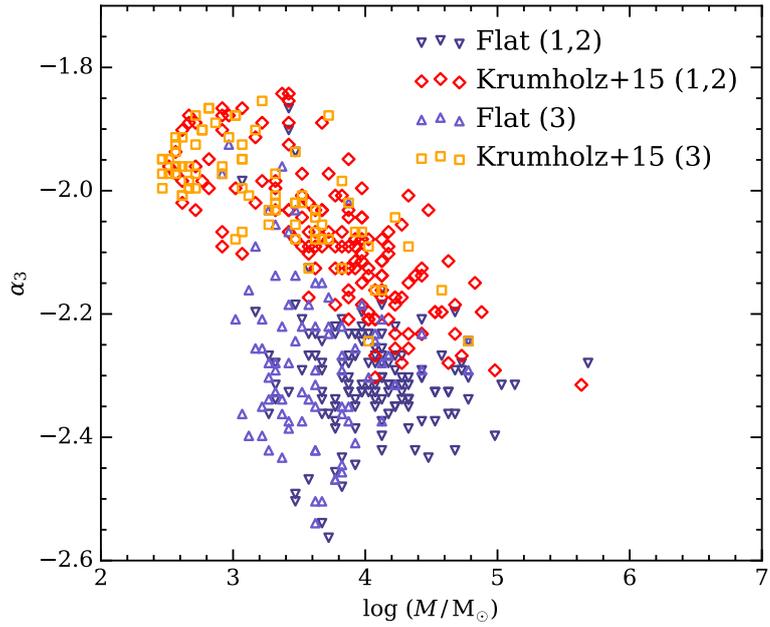
(a) Physically motivated priors
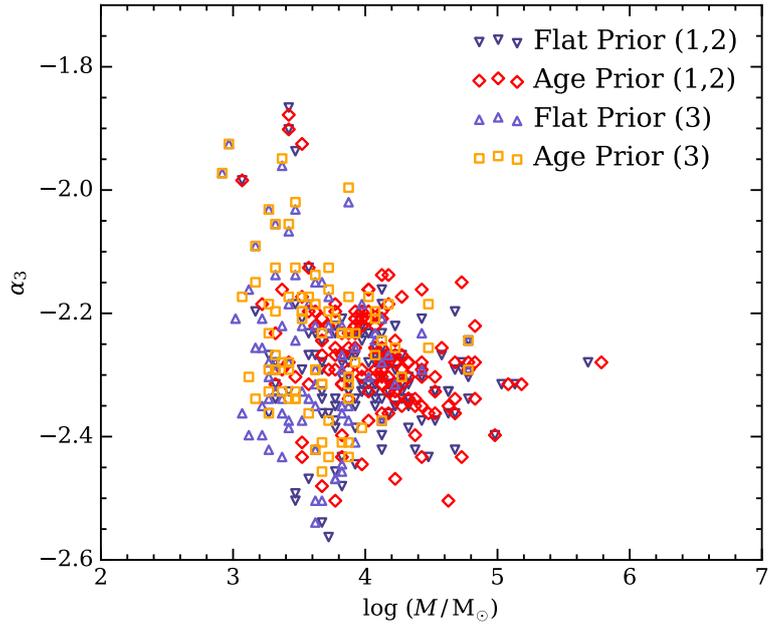
(b) Age prior only

Figure 3.4: Panel (a) shows the correlation between the medians of the posterior PDFs for $\log M$ and α_3, with class 3 objects (light colours) and class 1-2 objects (dark colours) separated for both flat (triangles) and the Krumholz et al. (2015b) priors (quadrilaterals). Note the more apparent correlation between mass and α_3 for the physically motivated priors, with a lower limit of α_3 around the canonical Kroupa value (-2.3). Panel (b) shows the same but with the prior on the mass removed, leaving only a prior on age. This results in a distribution very similar to a purely flat prior.
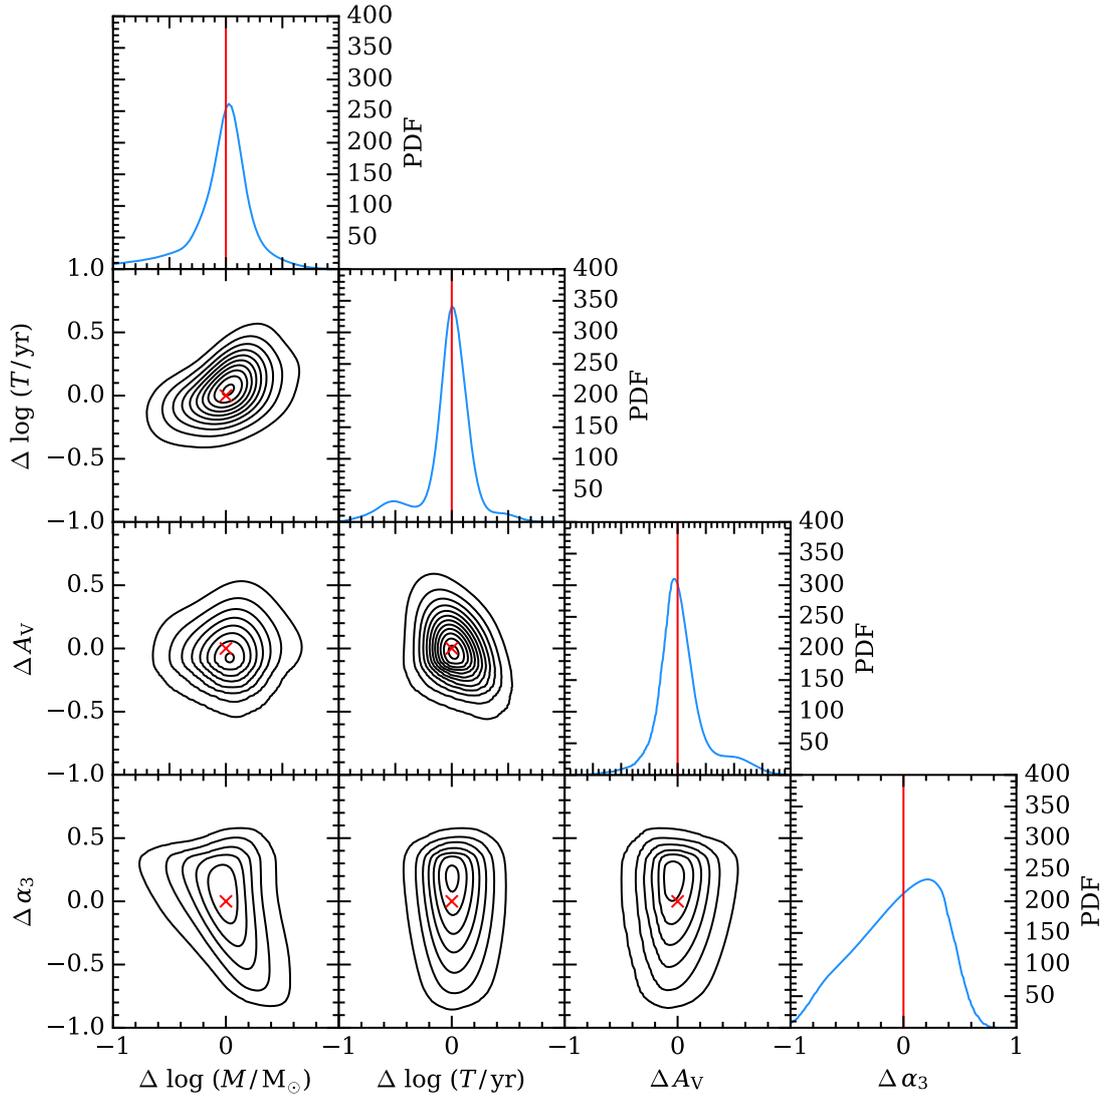
Figure 3.5: Stacks of the one and two-dimensional PDFs (median-centred) for the clusters in NGC 628E. The 2D PDF contours are spaced in steps of 50, beginning at 50. The red crosses and lines mark the zero points.

preferred. Therefore, the combination of imposing a physically motivated prior with a varying upper-end slope of the IMF moves the clusters in the α_3/M plane along a diagonal in the direction of lower mass and shallower slopes compared to the canonical Kroupa IMF. From a physical point of view, correlations between the cluster mass and the IMF slope are expected. For instance, a truncation of the IMF in excess of the effect induced by random sampling is predicted by the IGIMF theory, in which low mass clusters have a bottom-heavy IMF (Kroupa, 2001; Kroupa & Weidner, 2003; Weidner et al., 2011; Kroupa et al., 2013). However, the effect we see in panel (a) of Figure 3.4 is in the opposite direction to the effect predicted by the IGIMF. The fact that the posterior PDFs for individual clusters are very broad and that, as discussed below, the median values are quite sensitive to the choice of prior cautions against far-reaching conclusions of the nature of this correlation. Indeed, at this time we cannot exclude the possibility that the observed trend arises from a mass-dependent correlation at second order that is not explicitly captured in our analysis.

Finally we produce a median-centred stack of the 1D and 2D posterior PDFs for the LEGUS clusters, which is shown in Figure 3.5. Degeneracies between the cluster mass and age, and between the cluster mass and IMF slope (as seen in panel (a) of Figure 3.4) are apparent from the shape of the contours in their respective joint posterior PDFs. The broad shape of the posterior PDFs in α_3 are also visible. The mass-age degeneracy is, at least in part, caused by the need for a higher mass to produce the same luminosity in older clusters.

3.3.5 Effects of the choice of prior

The analysis of clusters from NGC 628E using the CLUSTER_SLUG tools presented in the previous section shows that age and extinction are largely insensitive to the parameters describing the IMF. Moreover, we have shown that the posterior PDF of α_3 is weakly constrained by broad-band photometry. When choosing a physically motivated prior on mass, the recovered PDFs are skewed towards lower masses and shallower IMF slopes compared to the case of a constant Kroupa IMF. As shown above, this effect arises from a degeneracy between mass and α_3, which results in diagonal shifts in the M/α_3 plane. It is therefore expected that the choice of prior on the mass will have a consequential effect on the parameters describing the IMF
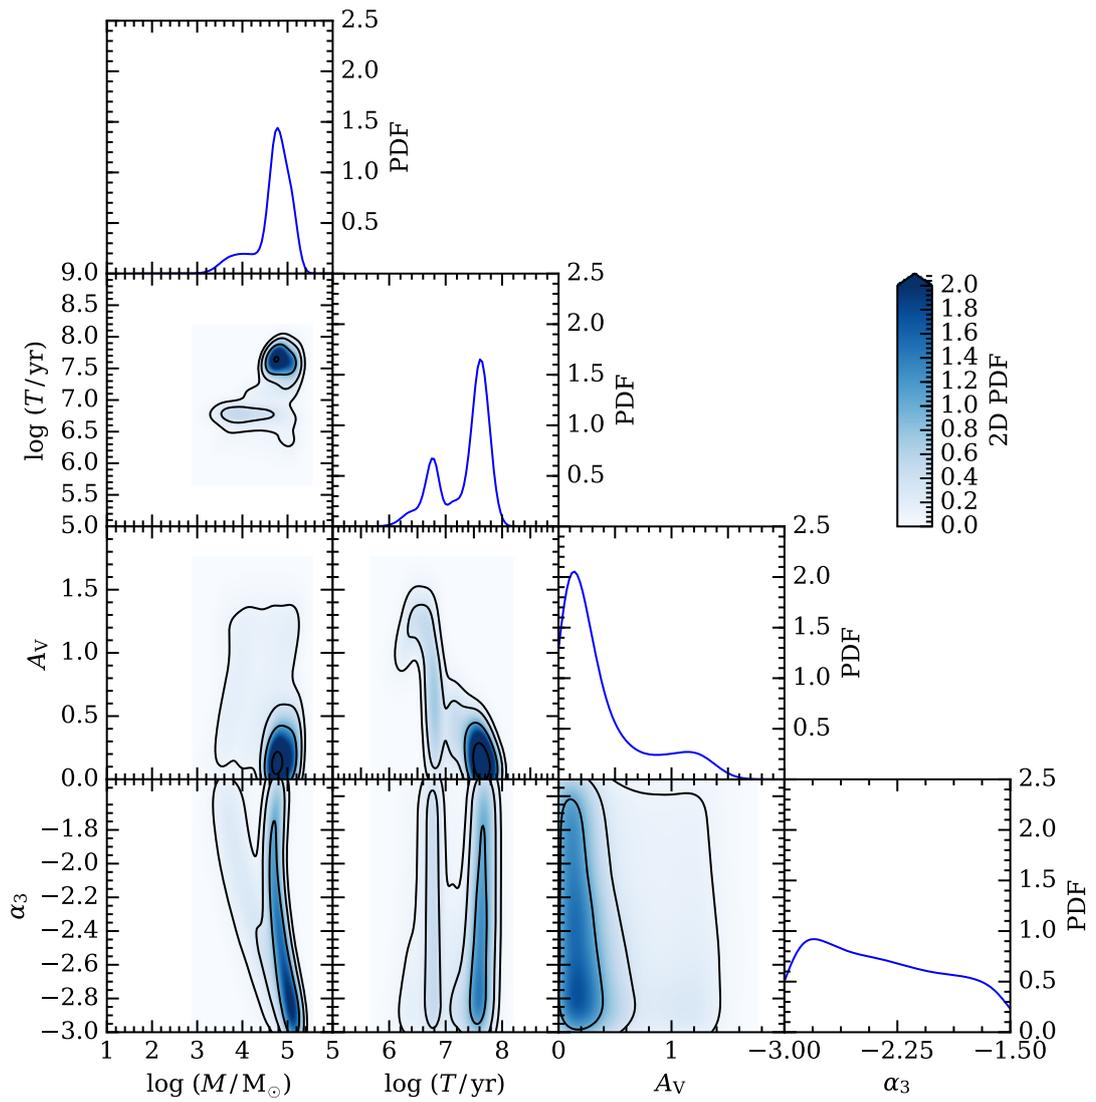
Figure 3.6: As Figure 3.1, but with a flat prior on all physical parameters.
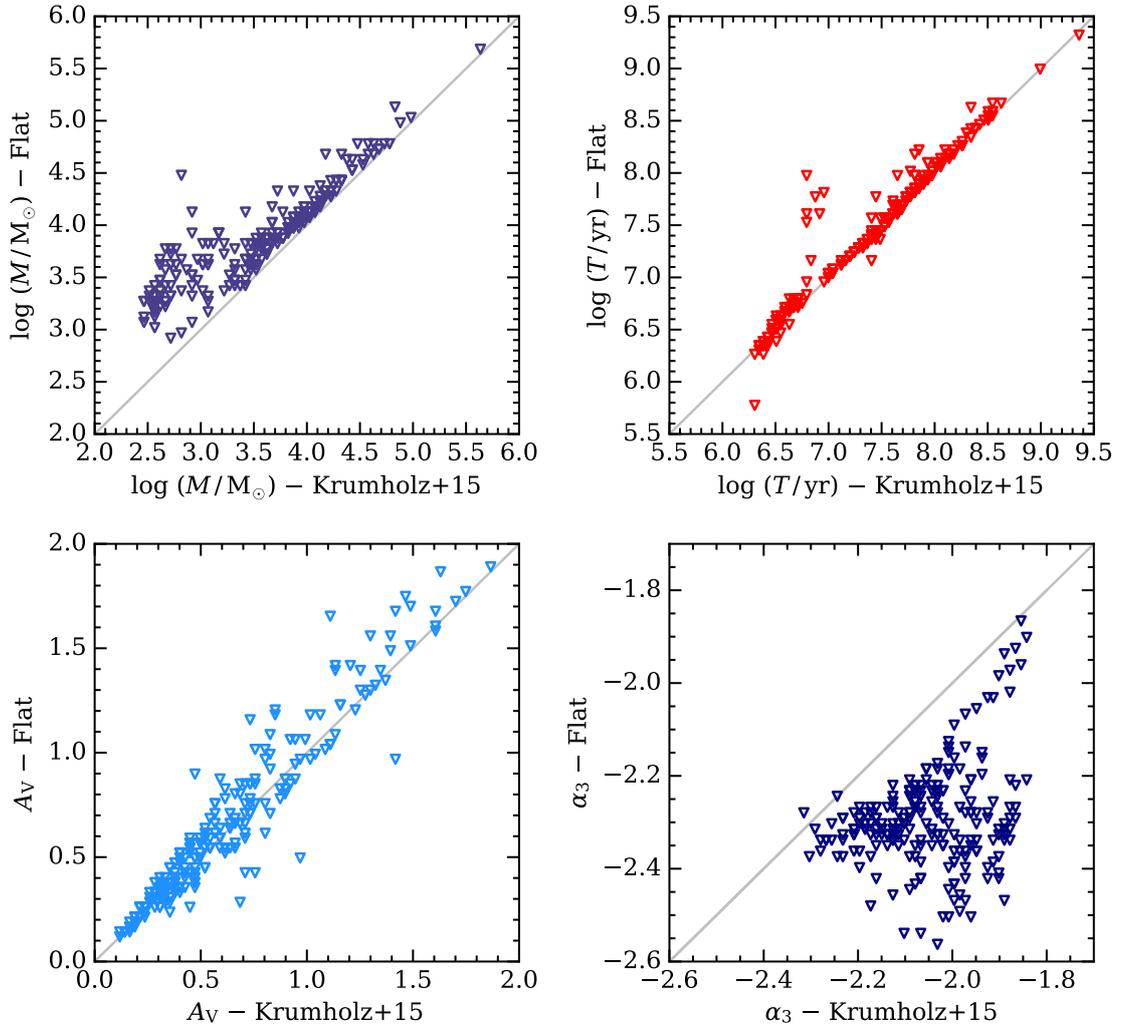
Figure 3.7: Comparison between the median values for each parameter returned by CLUSTER_SLUG for flat priors and the physically motivated priors by Krumholz et al. (2015b). The grey line marks the one-to-one correspondence. While extinction and age are insensitive to the choice of priors, mass and α_3 depend on the assumed prior, owing to a degeneracy between these two parameters.

(and vice versa), as we show in this section.

We start by examining again the corner plot for the example cluster ID 56, this time computed assuming a flat prior on all the parameters, which is shown in Figure 3.6. A comparison with Figure 3.1, in which a physically motivated prior was chosen, highlights how a flat prior shifts the probability from a locus of low mass and young age ($M \sim 10^{3.5}$ $M_\odot$, $T \sim 10^{6.7}$ yr) to a locus of higher mass and age ($M \sim 10^5$ $M_\odot$, $T \sim 10^{7.6}$ yr). As a consequence, in our model based on stochastic sampling, the observed photometry can be realised at higher cluster masses with steeper IMFs ($\alpha_3 \lesssim -2.3$) compared to the case of low cluster mass. In other words, to produce a comparable UV luminosity to that of a higher mass cluster, a low mass cluster needs a larger number of massive stars, something that can be achieved with a more shallow IMF slope.

By repeating the analysis for the full sample but using a flat prior, we see that the behaviour described above is indeed general. Compared to a physically motivated prior, a flat prior induces a shift in the cumulative PDFs (red lines in Figure 3.3) towards higher masses and steeper IMF slopes. Conversely, there is no noticeable difference in the extinction and only a modest change occurs in the age cumulative distribution, with some probability flowing from $T \gtrsim 10^{6.5}$ yr to $T \lesssim 10^{6.5}$ yr. As a consequence of this shift in mass/α_3, the correlation between the cluster mass and the IMF slope disappears when considering a flat prior (triangles in panel (a) of Figure 3.4).

Panel (b) of Figure 3.4 shows the same plot but with the analysis performed using only the age component of the Krumholz et al. (2015b) prior. Here we see a distribution of points that is almost identical to that of a flat prior, suggesting that the mass prior is dominant in its effect on the IMF that we recover. On these plots, we note that the fact the medians are centred around $\alpha_3 \approx -2.3$ is not indicative that the data prefer a Kroupa IMF. Rather, this is a mere reflection of the fact that the posterior PDFs on the IMF slope are very broad and, by construction, are centred around the canonical value of $\alpha_3 \approx -2.3$. Based on this result, it may be tempting to conclude that the choice of a flat prior is preferable over the physically motivated prior of Krumholz et al. (2015b). However, while we do not wish to overinterpret the trend visible Figure 3.4, the choice of our physically motivated prior is preferable given the power-law nature of the cluster mass function. This implies that either our SLUG simulations are not fully describing the properties of clusters

in NGC 628E, or that indeed there is (weak) preference for more shallow IMF slopes in low mass clusters.

Finally, rather than considering the cumulative PDFs, we compare the medians of the posterior PDFs for the physical parameters of individual clusters, as shown in Figure 3.7. Consistent with the previous discussion, the medians of the age and extinction PDFs are insensitive to the choice of prior, while the medians for the mass are systematically offset in the direction of lower masses for a physically motivated prior. This shift is modest (~ 0.2 dex) for clusters above $\sim 10^{3.5}$ M$_\odot$, but it becomes more significant (up to ~ 0.5 dex) at the low mass end. This figure also shows that shifts of a factor of 2 in mass that are modest on a logarithmic scale result in more appreciable variations of a power-law exponent on a linear scale, with $\Delta\alpha_3 \sim 0.2 - 0.3$.

One possible way to avoid the IMF slope having such a dependence on the mass prior would be to analyse all the clusters as a single unit with a common mass function between them. At present, a single cluster is considered more likely to be less massive, as suggested by our choice of prior. However, for an ensemble of clusters which follow the CMF that forms the prior on the mass, we would retrieve a distribution of masses that follows this CMF. In this situation there would not be a preference for every cluster to be less massive, rather for the ensemble of clusters to have a preference for less massive clusters.

3.3.6 Adding Hα photometry to the Bayesian analysis

Preparation of the Hα photometry

As shown in Section 3.3.4, our analysis of star clusters from NGC 628E based on the five broad-band LEGUS filters provides only modest constraints on the IMF slope. We also find that several clusters have broad or multi-peaked posterior PDFs in some of the physical parameters, especially age. As massive OB stars dominate the Hα emission in clusters at younger ages, the inclusion of Hα photometry is expected to improve some of the constraints on the physical parameters of interest. To this end, we include in our analysis data obtained as part of Hα-LEGUS (Chandar et al., in preparation) that are collected through the WFC UVIS F657N filter which, at the redshift of the LEGUS galaxies, encompasses the Hα line.

To derive aperture photometry for Hα data, we follow the method developed by
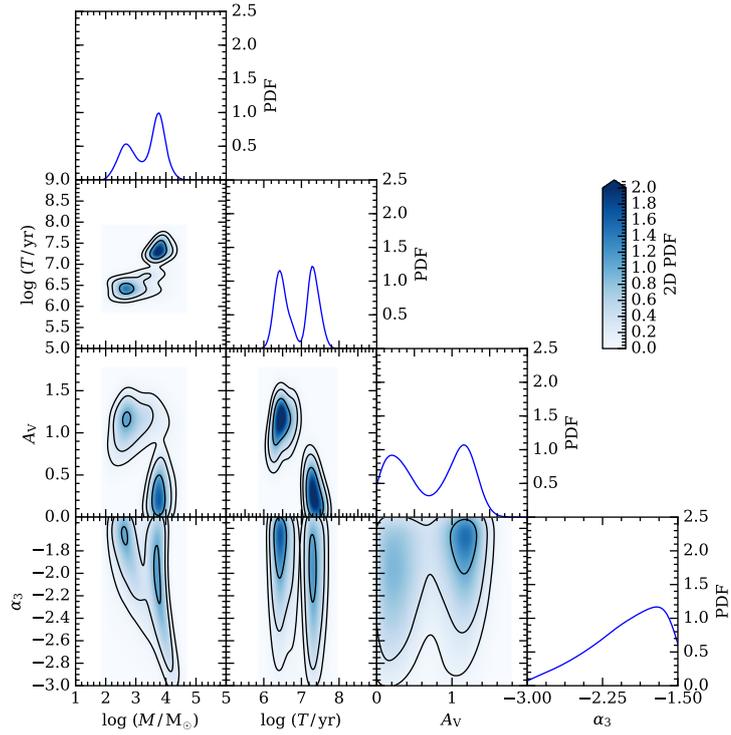
(a) Broad-band filters only
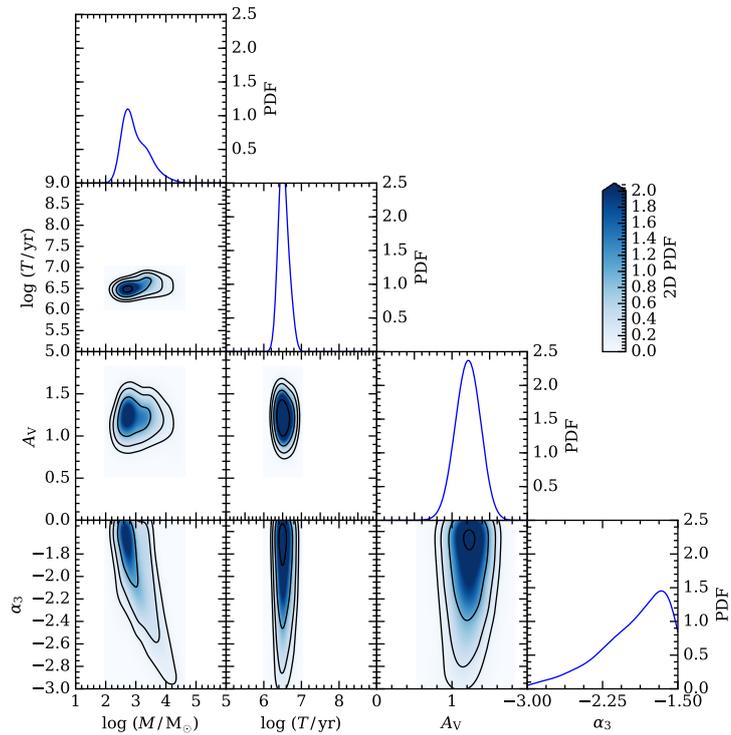
(b) Broad-band + Hα

Figure 3.8: As Figure 3.1, but for cluster ID 292, which has observed Hα emission. Panel (a) is for the standard LEGUS broad-band filters only, whereas panel (b) includes the F657N filter to represent Hα. We see significant improvement in the recovery of the mass, age, and extinction.
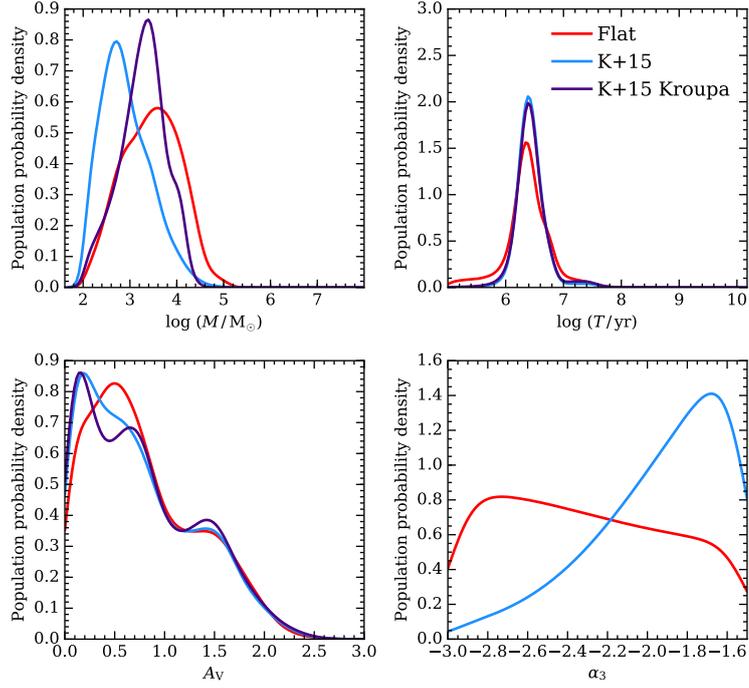
(a) Broad-band filters only
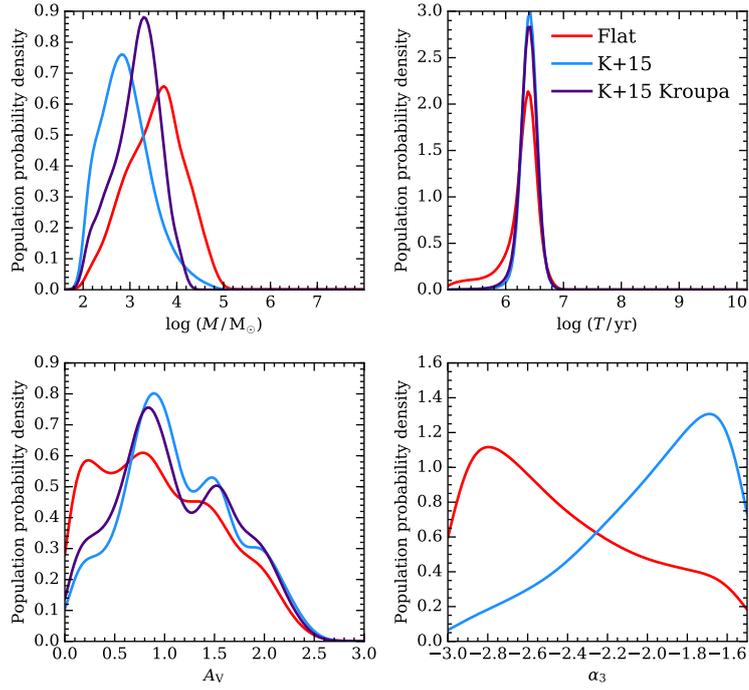
(b) Broad-band + Hα

Figure 3.9: The cumulative posterior PDFs for a subset of young clusters (with median age $T \leq 10^{6.5}$ yr) computed for both choices of priors, and for a constant Kroupa IMF. The upper panels show the PDFs based on broad-band photometry only, while the lower panels show the PDFs computed including the Hα flux. The inclusion of Hα in the analysis results in much sharper PDFs in age, which also lifts some of the degeneracy in extinction. There is also a small shift towards steeper IMF slopes for both priors.

Lee et al. (in preparation), which we briefly summarise here. During the processing of the F657N data, the standard approach for applying aperture corrections to broad-band photometry cannot be trivially applied to the nebular emission, given that the stellar continuum and the nebular emission differ in their spatial extent. To account for this difference and to produce photometry that can be compared to SLUG simulations, we follow a two step procedure.

Firstly, we synthesise an aperture-corrected continuum flux in the F657N filter, $f_{657,c}$. This is achieved by linearly interpolating the aperture-corrected continuum flux measured in the F547M and F814W filters, with a weight tuned to recover the observed aperture-corrected flux in the F657N filter for clusters without Hα emission. This procedure yields a synthesised F657N aperture-corrected continuum flux. Next, we repeat the above procedure, but for fluxes that are not aperture corrected, thus obtaining an interpolated continuum flux in the F657N filter that is not aperture-corrected. This is needed to compute the line emission, as explained next. To derive the net Hα line flux $f_{657,H\alpha}$, we then subtract the synthesised continuum flux (without aperture correction) from the measured flux in the F657N filter.

Combining the two steps, we produce the aperture-corrected total flux (continuum plus line) in the F657N filter as $f_{657,corr} = f_{657,H\alpha} + f_{657,c}$. Thus, at the end of this procedure, we have created a synthetic aperture correction from the weighted average of the aperture corrections on the bracketing filters and applied this to the measured F657N flux. Finally, as for the LEGUS broad-band photometry, we apply a correction for Galactic extinction. This results in a total flux value for the F657N filter which can be used in our analysis and is comparable to the F657N photometry generated by SLUG. This total flux (hereafter referred to as Hα flux for simplicity) is the quantity that we compare directly with the output of SLUG.

The SPS module at the core of the SLUG package predicts the number of ionisation photons produced by stars in each cluster, which we convert into nebular emission following the methods described in Krumholz et al. (2015a). At this stage, we assume that only a fraction ($\phi = 0.5$) of the ionising photons are converted into nebular emission. After convolving the simulated cluster SED with the F657N filter transmission curve, SLUG generates a total flux like the $f_{657,corr}$ we synthesise for the observations as described above. Throughout this work, we do not include Hα

emission from stars themselves (e.g. from Be stars), as they are found to contribute only a few percent compared to the nebular emission.

We note that, due to the extended nature of the Hα emission, there is a possibility that we are missing some fraction of the flux in the F657N filter from extended (but typically low surface brightness) emission. As we will show below, however, we are able to reproduce all the findings of our analysis using a set of mock clusters that, by definition, are not affected by loss of flux on large scales. This implies that any potential loss of flux has a minor effect on our results. Moreover, the choice of a nebular fraction of $\phi = 0.5$ already mitigates this effect. We have also conducted tests to verify that our results are not sensitive to the exact choice of this coefficient. Indeed, we find that repeating our analysis using a library with $\phi = 0.7$ does not have a noticeable effect on the medians of the PDFs we return for the physical parameters.

Results including Hα photometry

To illustrate the effect of including Hα in our analysis, we turn to an example cluster with observed Hα emission (ID 292), which is shown in Figure 3.8. Looking at the differences between panels (a) and (b), we see a significant improvement in the age determination, with the locus at $T \sim 10^{6.5}$ yr being clearly preferred over the older ages. In this particular case, the inclusion of Hα further improves the constraints that can be derived on the mass and extinction, removing the bimodality seen in both. In turn, as the lower mass at $M \sim 10^3$ M$_\odot$ is now preferred, the posterior PDF for α_3 sharpens compared to the case in which no Hα was included.

As we will show in the following section, the improvement in age determination obtained when including Hα in the analysis is common to much of the cluster sample, and is most noticeable in the case of younger clusters. However, the variation in the median of the age PDFs is modest, and we see no appreciable difference in the shape of the cumulative PDFs when considering our entire cluster catalogue. Thus, given that the shapes of the cumulative PDFs for both mass and age are primarily dependent on the scatter among clusters rather than the width of the PDFs for the individual clusters, the shapes of the PDFs shown in Figure 3.3 do not show significant change following the inclusion of Hα. Similarly, we see no appreciable difference in the cumulative PDF for α_3. This is primarily caused by the

broad shape of the α_3 PDFs in the individual clusters, even following the inclusion of Hα.

Finally, we examine the effect of including Hα in the analysis of young star clusters by limiting our choice of clusters to those whose posterior PDF in age has a median of $T \leq 10^{6.5}$ yr. It is for this subset of clusters that the NUV and Hα photometry of the LEGUS and Hα-LEGUS surveys is expected to yield the tightest constraints. The cumulative PDFs for this subset of clusters, computed with both choices of priors, a fixed Kroupa IMF, and with or without Hα, are shown in Figure 3.9. This subset contains 21 clusters for the case of flat priors with no Hα, 24 with Hα, 21 for physical priors with no Hα, 18 with Hα, and 20 for the case of the fixed Kroupa IMF with no Hα with 21 when Hα is included.

Several interesting features can be seen in this figure. Considering first the PDFs derived assuming the physically motivated priors (blue curves), the inclusion of Hα significantly improves the age constraint (compare the uppper and lower panels), thus sharpening the age posterior PDFs. However we see no appreciable shift in the positions of the medians of the PDFs. A more subtle difference is visible when comparing the mass and α_3 PDFs. Conversely, more substantial differences appear when examining the posterior PDFs for the extinction, with the inclusion of Hα favouring extinctions up to ~ 0.5 dex higher. Comparisons with the PDFs computed with a constant Kroupa IMF (purple line) reveal that the differences observed when including Hα are independent of the choice of the IMF.

Moreover, in line with the discussion presented in the previous section, Figure 3.9 shows how critical the mass-α_3 degeneracy is in this type of analysis. At fixed age, a physically motivated prior skews the mass determination towards smaller values, which is compensated by a more shallow slope on the upper-end of the IMF. Conversely, a flat prior (particularly when including Hα) prefers steeper IMF slopes which result in higher masses. In between, the constant Kroupa IMF case with $\alpha_3 = -2.3$ places the preferred masses at intermediate values between those inferred with flat and physically motivated priors. In summary, a key result of this work is that, due to the mass-IMF degeneracy, the choice of prior on the mass strongly affects the resultant shape of the posterior PDF of α_3, preventing robust constraints on the upper end slope of the IMF even when including Hα. Turning this argument around, the assumption of a fixed Kroupa IMF, which is equivalent to imposing a strong prior on α_3, affects the mass determination by up to ~ 0.5 dex for young and
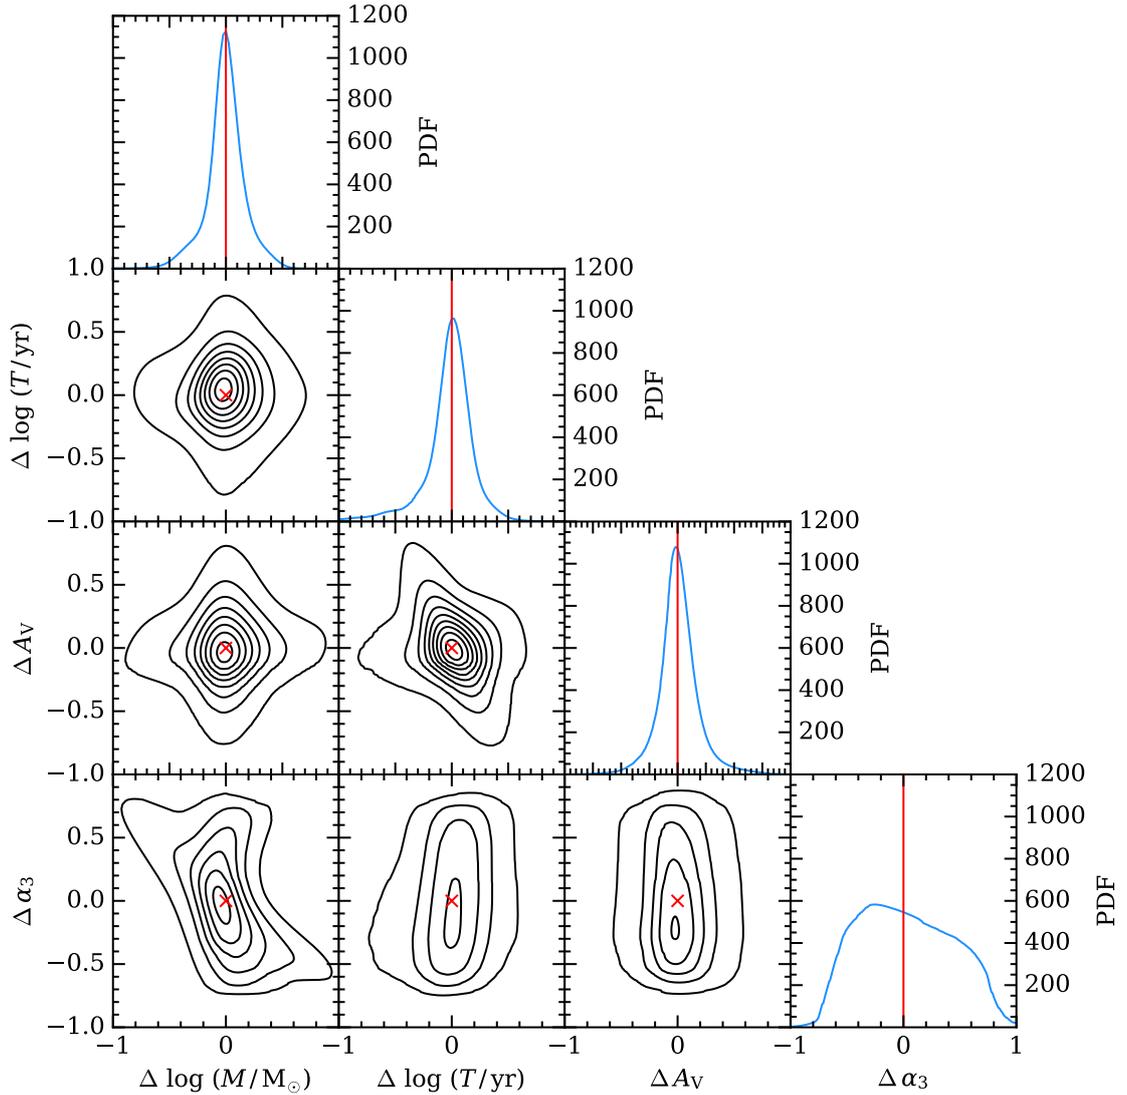
Figure 3.10: Stacks of the one and two-dimensional PDFs (median-centred) for 700 mock clusters simulated with a variable IMF and the choice of parameters listed in Table 3.2. 2D PDF contours are spaced in steps of 100, beginning at 100. The red crosses and lines mark the zero points.

low-mass clusters.

3.4 Analysing the performance of the technique using mock clusters

Having applied our Bayesian formalism to the study of star clusters in NGC 628E, we have found that a degeneracy between mass and the slope of the upper-end of the IMF prevents us from reliably constraining the value of α_3 both when using broad-band photometry alone and when Hα is included. We have also shown how

Name	Tracks	Base IMF	Z	Extinction[a]	ϕ[b]	$\log M$ (M$_\odot$)	$\log T$ (yr)	A_V (mag)	#[e]	α_3
MocksRun1[c]	Padova AGB	Kroupa	0.020	MW	0.5	2–6(1)	6–9(0.5)	0.1,0.5,1,1.5,2	3	−2.8, −1.6(0.2)
MocksRun2[d]	Padova AGB	Kroupa	0.020	MW	0.5	2–6(1)	6–9(0.5)	0.1,0.5,1,1.5,2	3	−2.3

[a] MW = Milky Way extinction curve
[b] ϕ is the fraction of ionising photons that produce nebular emission within the aperture, combining the effects of a covering fraction less than 1 and some portion of the ionising photons being absorbed directly by dust
[c] Variable IMF mocks
[d] Constant IMF mocks
[e] Number of realisations

Table 3.2: Table of the parameters used in SLUG for generating the two grids of mock star clusters used in this work. Within the parentheses we list the adopted step size.
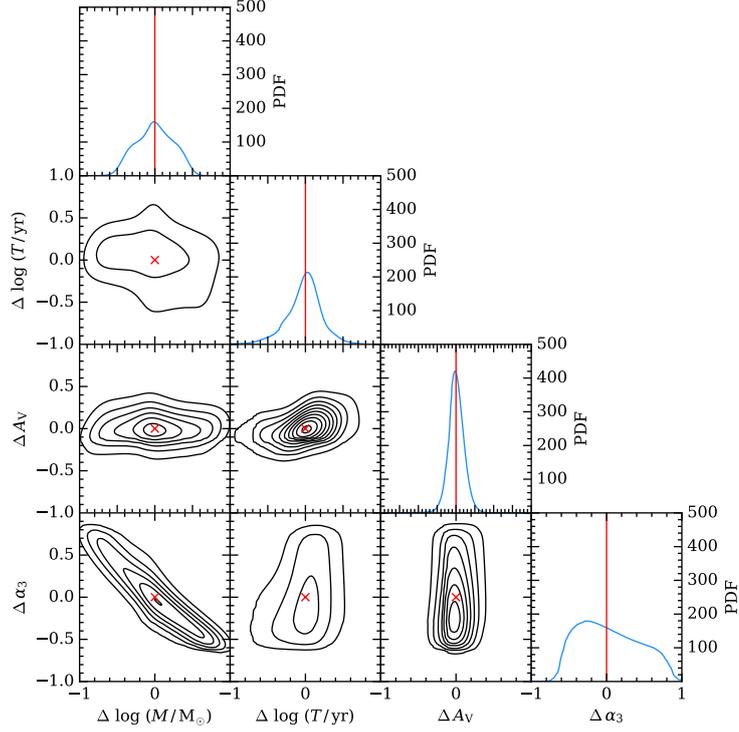
(a) Broad-band filters only
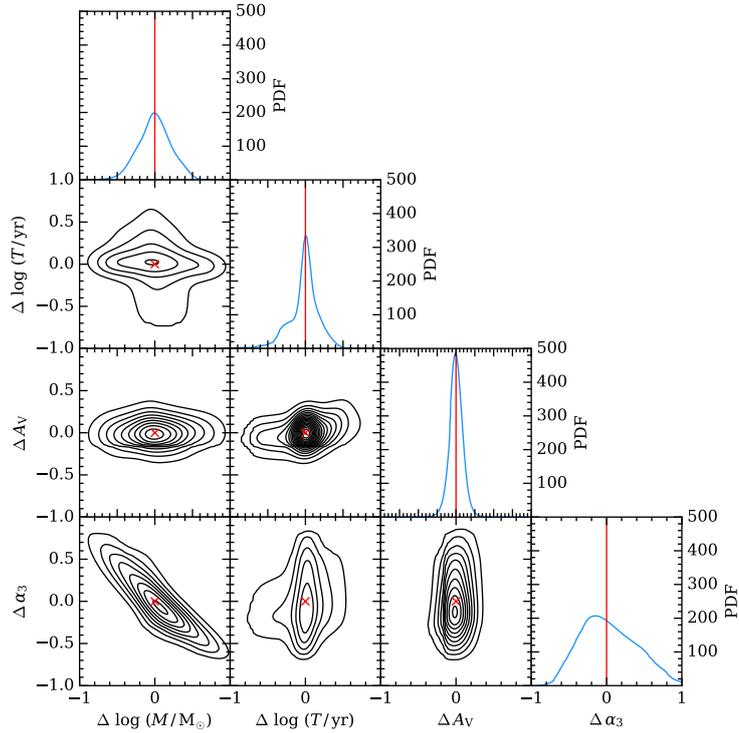
(b) Broad-band + Hα

Figure 3.11: As Figure 3.10, but for a subset of young mock clusters, with age $T \leq 10^{6.5}$ yr. The upper and lower panel show, respectively, the PDFs recovered from broad-band photometry alone, and with the inclusion of Hα. 2D PDF contours are spaced in steps of 50, beginning at 50. The red crosses and lines mark the zero points.

(a) Mass

(b) Age

(c) Extinction

(d) IMF
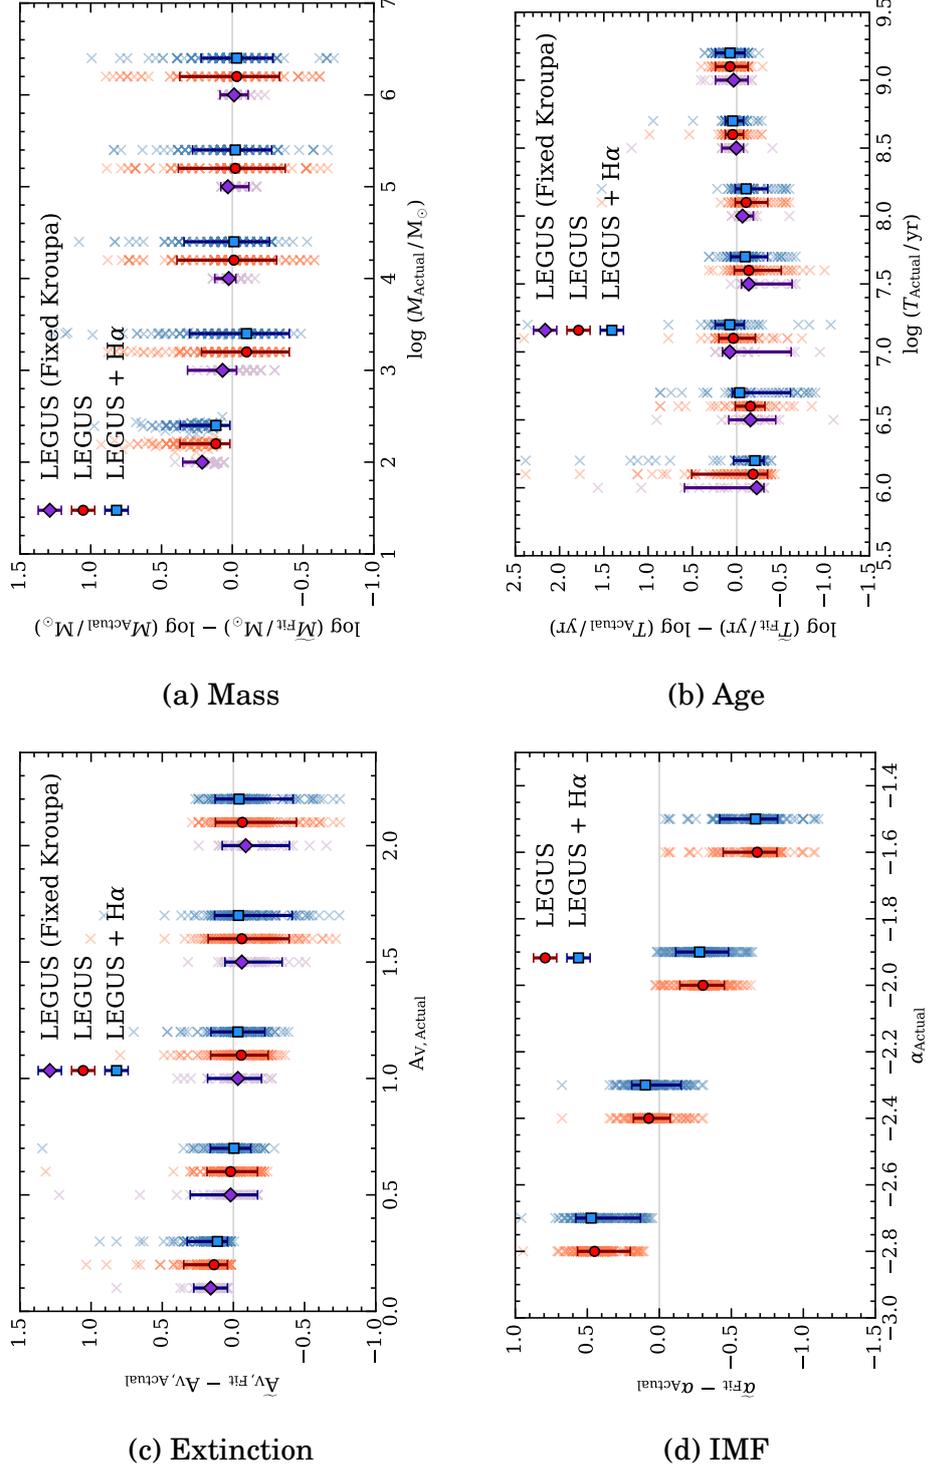
Figure 3.12: Comparisons of the medians of the inferred posterior PDFs for all physical parameters to the true values adopted when simulating the mock clusters. Results for individual clusters are shown with crosses, while the median residuals and the corresponding 10^{th} and 90^{th} percentiles are shown with solid points. We plot analysis including broad-band photometry alone with (red, circles) and without (purple, diamond) a variable IMF. We also show results of the variable IMF case when including Hα (blue, squares). Points are shifted along the x-axis for visualisation purposes, with the leftmost point situated at the true x-axis value.
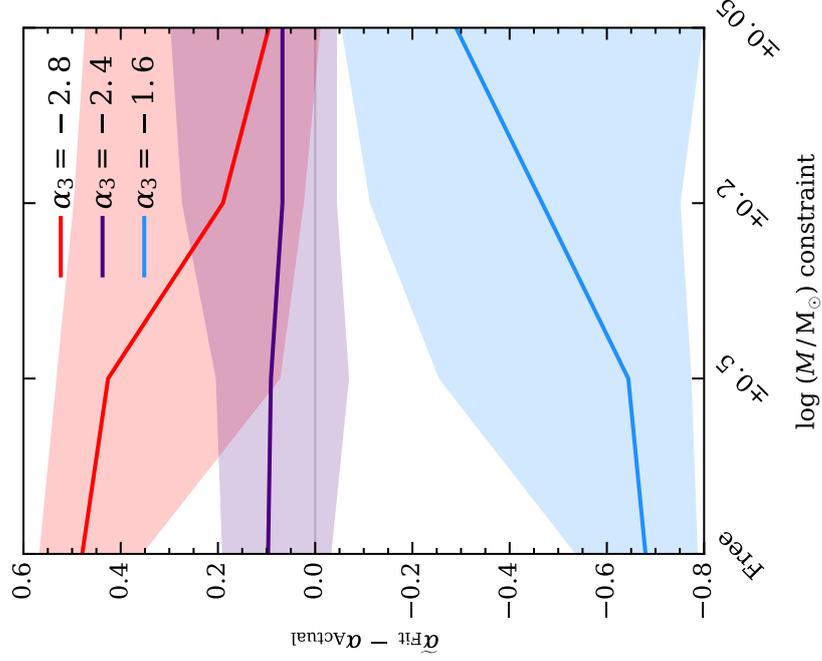
(a) Broad-band filters only
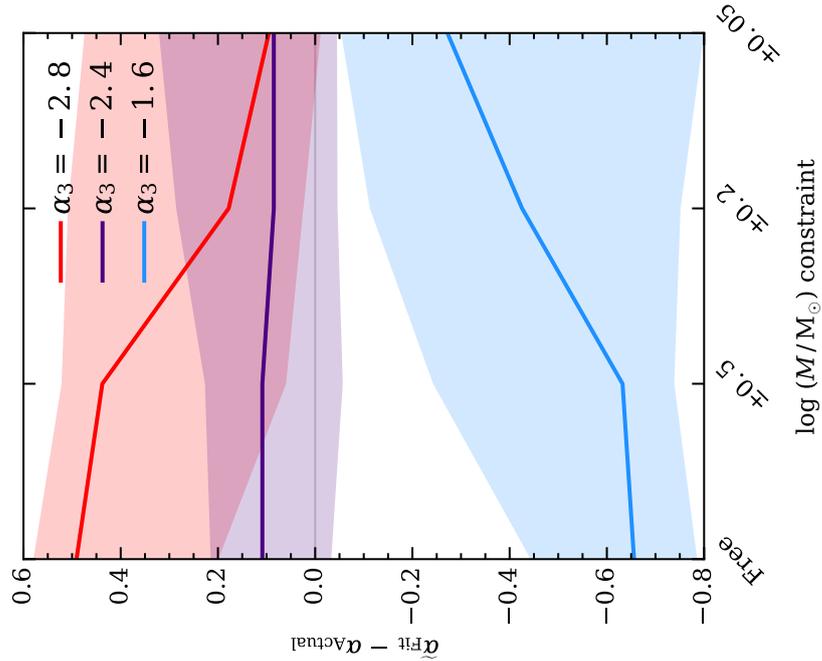
(b) Broad-band + Hα

Figure 3.13: Comparison between the IMF slopes α_3 measured from the inferred posterior PDF (taking the median) and the true value as a function of width of a prior on the mass. Results from three IMF slopes are shown. The filled regions represent the 10th and 90th percentile range. We omit clusters with a target mass of $100\,\mathrm{M}_\odot$ due to the relative variation in the resultant generated masses. Panel (a) was produced using only the 5 broad-band filters, whereas (b) was produced using the additional Hα filter. We see some improvement for the case with no constraints, with marginal improvement for the tighter constraints.
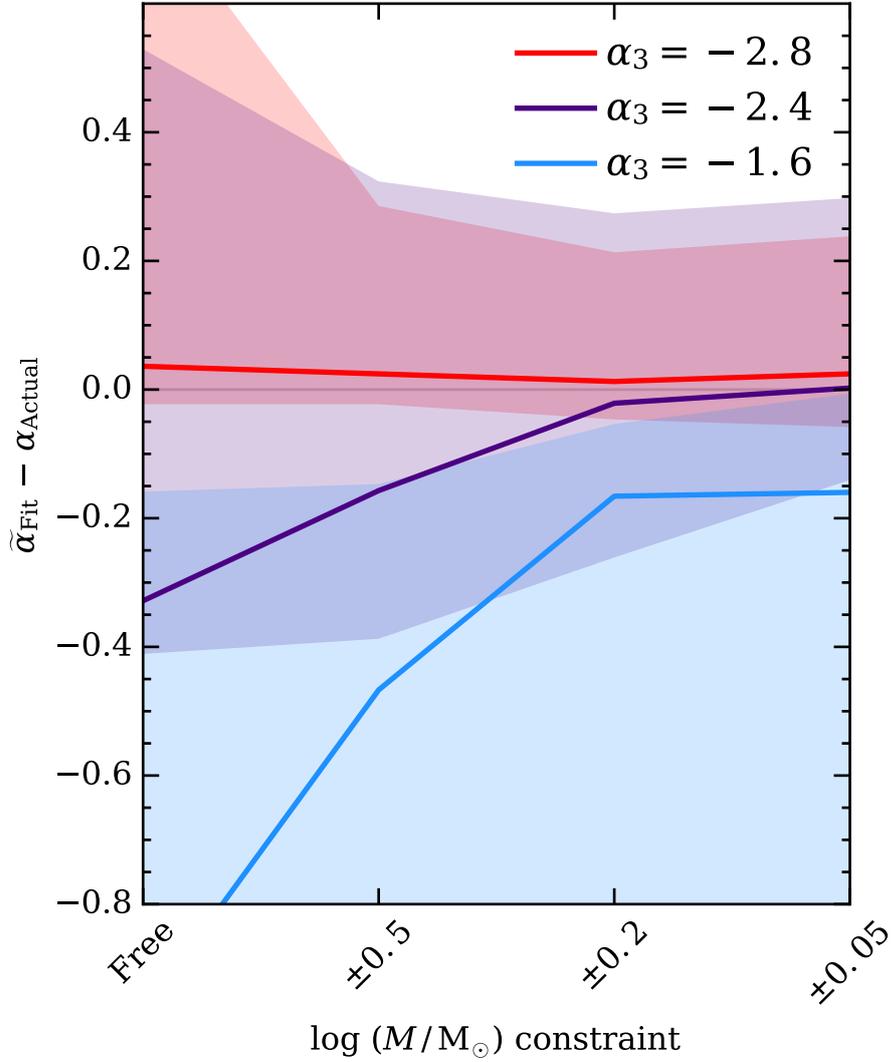
Figure 3.14: Comparison between the IMF slopes α_3 measured from the inferred posterior PDF (taking the mode rather than the median) and the true value as a function of width of a prior on the mass. Results from three IMF slopes are shown. The filled regions represent the 10^{th} and 90^{th} percentile range. We omit clusters with a target mass of $100\,\text{M}_\odot$ due to the relative variation in the resultant generated masses.

the choice of prior on the mass has a knock-on effect on the shape of the posterior PDF for α_3. Equivalently, the choice of a constant IMF may lead to systematic errors in the mass determination of more than 0.5 dex (depending on age).

In this section, we exploit mock clusters to take a more detailed look at our ability to recover the posterior PDFs for physical parameters of clusters using our Bayesian formalism. We also characterise the significance of the mass-α_3 degeneracy, providing forecasts on the ability to constrain the IMF in future experiments when the cluster mass can be determined via techniques that are independent of photometry.

To achieve these goals, we generate a grid of 700 mock star clusters using SLUG, with masses in the range $\log(M/\text{M}_\odot) = 2 - 6$ in steps of 1 dex, ages in the range $\log(T/\text{yr}) = 6 - 9$ in steps of 0.5 dex, and five choices of extinction, $A_V = 0.1, 0.5, 1.0, 1.5, 2.0$ mag. For the IMF upper-end slope (α_3), we select values ranging from -2.8 to -1.6 in steps of 0.2. These parameters, together with the other SLUG parameters that define the SPS calculation, are listed in Table 3.2.

We also generate a second grid of 175 clusters with a constant Kroupa IMF ($\alpha_3 = -2.3$) spanning the same ranges of mass, age, and extinction. The parameters for this grid are also listed in Table 3.2. Due to the stochastic nature of the SLUG code, clusters simulated with a given choice of parameters do not necessarily have identical properties. For this reason, we generate 3 realisations of each cluster across both of our grids to help prevent the inclusion of rare outliers in our analysis. Using the variable library generated for the analysis of NGC 628E in Section 3.3.2, we apply CLUSTER_SLUG to our grid of mock clusters with a variable IMF. Likewise, we apply CLUSTER_SLUG to the fixed IMF grid using the fiducial library of Krumholz et al. (2015b) as a training set. We use flat priors during this analysis as the clusters are drawn from a flat distribution of physical parameters, without following, for example, a cluster mass function.

At first, we examine the precision with which CLUSTER_SLUG can recover the physical parameters of the clusters. This is achieved by studying the intrinsic width of the posterior PDFs recovered by our Bayesian analysis. A sharp PDF not only means that physical parameters can be recovered to high precision, but also ensures that the global properties of a sample (e.g. via the cumulative PDFs) reflect the intrinsic scatter of the population. Conversely, very broad PDFs reflect not only a low precision in recovering parameters, but also the inability to char-

acterise the underlying scatter of physical parameters within the cluster population. In practice, we stack both the one-dimensional and two-dimensional posterior PDFs returned by CLUSTER_SLUG for all 700 mock clusters with a variable IMF. The PDFs are median-centred before stacking, and to select a representative group of objects, we choose for each cluster the realisation corresponding to the the second-best 'fit' in mass (i.e. second-smallest residual). These stacks are shown in Figure 3.10 and Figure 3.11.

Inspecting the 1D posterior PDFs recovered from broad-band photometry alone (Figure 3.10), we see that the PDFs for the mass, age, and extinction have a well-defined peak. Conversely, the posterior PDF for α_3 is broad and flat, confirming that broad-band photometry alone has weak constraining power on the upper end of the IMF. Additional insight is offered by the analysis of the joint PDFs. The well-known degeneracy in age-extinction is recovered by our analysis, as is the degeneracy between mass and α_3 discussed in the previous sections. Further, the contours on the joint PDFs of age and extinction with α_3 indicate that T and A_V are largely insensitive to variations in the upper end slope of the IMF.

Restricting our analysis to a subset of mock clusters with young ages ($T \leq 10^{6.5}$ yr, Figure 3.11), we see that the performance of our Bayesian technique improves, with the posterior PDFs of all the parameters becoming narrower. Qualitatively, however, the majority of the trends we see in Figure 3.10 remain. By comparing the posterior PDFs computed with and without the inclusion of Hα (right and left panels, respectively), we see that Hα is critical in improving the age determination, breaking the age-extinction degeneracy. The broadness in the age PDF seen in the left panel of Figure 3.11 is reduced greatly following the inclusion of Hα, which is seen in the narrowing of the age PDF in the panel on the right. In turn, a better constraint in age (and extinction) is reflected in a narrower mass PDF, although it remains broad due to the mass-IMF degeneracy. Indeed, with refined PDFs on age and extinction, an evident "banana" shape emerges for the joint PDF for $M - \alpha_3$, which clearly highlights the degeneracy between mass and IMF that has shaped much of the discussion in this work. By means of our Bayesian analysis, we can also accurately quantify this degeneracy (see below), concluding that a variation in $\Delta\alpha_3 \sim \pm 0.5$ maps into a mass variation of $\Delta \log(M/\mathrm{M}_\odot) \sim \pm 1$ dex.

Next, we consider how accurate our Bayesian procedure is in recovering the

underlying physical parameters for each cluster. This is achieved by comparing the medians of the inferred posterior PDFs to the true value used when computing the mocks. As was done for the stacks, in an effort to apply our analysis to a representative sample of clusters, we consider for each set of parameters the realisation that yields the second-best "fit" to the mass. These results are presented in Figure 3.12. Looking at all four physical parameters, we see that the medians of the inferred PDFs are a trustworthy indicator for the true value of the mass, age, and extinction when considering the average population. For the case of constant IMF (purple diamonds), the scatter between individual clusters is modest, of the order of $\sim 0.2 - 0.3$ dex (mag). Conversely, the addition of a variable IMF parameter greatly increases the scatter of residuals on a cluster by cluster basis when compared to the grid of Kroupa IMF clusters. This increase in scatter is especially noticeable in mass, and originates from the $M - \alpha_3$ degeneracy described above. As the mass of the cluster and its IMF slope are highly dependent on each other, the wide scatter in mass is further reflected in the inaccurate values inferred for α_3. Indeed, given the flat and broad posterior PDFs for the IMF slope, the medians of α_3 have no physical meaning and simply reflect the median of the prior on α_3 (in this case a flat distribution between $\alpha_3 = (-3.0, -1.5)$).

Given this result, it is clear that obtaining refined constraints on the upper-end slope of the IMF requires a different approach. It is readily apparent that the ability to apply a prior to the cluster mass that is independent of the broad-band photometry may prove very useful in analysis such as this. Independent constraints on the masses of clusters that may then be used as priors may come, for instance, from dynamical measurements (e.g. Ho & Filippenko, 1996; Gerhard, 2000). The era of 30 m telescopes will make such measurements more feasible.

To illustrate this point in more quantitative terms, we present in Figure 3.13 a study of the accuracy with which α_3 can be recovered as a function of the width of a prior on the mass. Results of three choices of IMF are presented, when leaving the mass unconstrained, and when applying a top-hat prior with width 0.5, 0.2, and 0.05 dex about the target cluster mass, for both broad-band filters only and broad-band filters combined with Hα. This analysis shows that, as expected, the medians of the PDFs on α_3 converge to the true value as the precision with which the mass is constrained increases. Due to the difference between logarithmic intervals in mass and the linear scale in the IMF upper-end exponent, there is a weak

scaling of α_3 with the prior width, with the underlying values being recovered to within $\Delta\alpha_3 \sim \pm 0.2$ when the mass is constrained to better than 10%. However, more realistic constraints on the mass (to within a factor of 2), for instance via dynamical measurements, appear to be useful to rule out extreme variations in the IMF upper-end slope through multiwavelength broad-band photometry, given the assumptions of our model. The inclusion of Hα results in a wider scatter towards zero for the unconstrained case. This advantage becomes much less evident as the mass is constrained however. Finally, in Figure 3.14, the mode rather than the median of the α_3 posterior PDF is used to represent the recovered α_3. Here we see a much faster convergence than we see in Figure 3.13. The scatter also reduces rapidly as we tighten the constraints on the mass, although this improvement is negligible in the $\alpha_3 = -1.6$ case, where the scatter remains very large.

3.5 Summary and conclusions

In this paper, we have developed a Bayesian formalism to study the properties of the IMF in star clusters with multiwavelength broad-band photometry combined with SLUG stochastic stellar population synthesis models. To investigate variation in the IMF, we first extended the capabilities of SLUG, implementing a flexible algorithm by which the PDFs defining physical functions such as the IMF can be set to variable mode. With this variable mode, it is now possible to perform simulations of clusters with an IMF for which specific parameters can be varied continuously across a desired range.

Exploiting this variable mode, which is now part of the main distribution of SLUG, we constructed a library of 20 million star clusters, simulated with a Kroupa-like IMF for which the slope of the high-mass end (α_3) was randomly chosen for each cluster, evenly selecting between -3.0 and -1.5. With this library, we then exploited the CLUSTER_SLUG Bayesian analysis framework described in Krumholz et al. (2015b) to compute the posterior PDF of α_3 for clusters with observed multiwavelength broad-band photometry.

As a proof of concept, we applied this formalism to a catalogue of 225 star clusters from the local galaxy NGC 628, for which broad-band and Hα photometry is available from the LEGUS (Calzetti et al., 2015) and Hα-LEGUS (Chandar et al. in preparation) programmes. After assuming a physically motivated prior for the

mass and age of star clusters, we found that our formalism recovers the PDFs of the core physical parameters of mass, age, and extinction for the entire sample, consistent with those resulting from the analysis of the same clusters with a constant IMF (Krumholz et al., 2015b). Conversely, we found very broad and flat PDFs for α_3, indicating that broad-band photometry alone is unable to tightly constrain the slope of the upper-end of the IMF. This result is primarily driven by a noticeable degeneracy between mass and α_3.

Due to the interlink between mass and IMF slope, we found that the posterior PDF for α_3 is quite sensitive to the choice of priors. Compared to the assumption of a flat prior in mass, a physically motivated choice derived from the power-law nature of the cluster mass function leads to posterior PDFs for the mass that are systematically skewed towards lower cluster masses and with a preference for more shallow IMF slopes, in excess of the canonical Kroupa value of $\alpha_3 = -2.3$. However, the fact that the posterior PDF is very broad, and that it is sensitive to the choice of prior cautions against far-reaching conclusions on the nature of the IMF with present data. Conversely, applying a prior to the cluster ages alone has a negligible effect on our analysis.

Adding Hα fluxes to our analysis does not significantly narrow the posterior PDF for α_3, although it reduces the degeneracy between age and extinction. As a consequence, the posterior PDFs for age and extinction, and to some extent mass, sharpen. This effect is particularly noticeable for the subsets of young star clusters, with ages $T \leq 10^{6.5}$ yr.

We also quantified the accuracy and precision with which our procedure is able to recover the input physical parameters for a grid of mock clusters, simulated across a wide range of physical parameters. We find the posterior PDFs for mass, age, and extinction have, on average, well defined medians. The width of the peaks in the PDFs are particularly narrow when considering a subset of young clusters, and when including Hα in the analysis. We find that we accurately recover the input parameters of our mock clusters, on average, with the medians of the mass, age, and extinction agreeing well with the input parameters. However, we still experience difficulty in recovering our input α_3, finding a very broad posterior PDF across many mock clusters.

Finally, we have presented future avenues for improvement in our ability to constrain the slope of the upper end of the IMF via broad-band photometry. One

possible path, rather than performing the analysis of the photometry on a cluster-by-cluster basis as we do in this study, is to perform a joint analysis of the star clusters with an underlying mass function that is common to all. This ensemble method could lessen the impact of the mass prior on our results. Furthermore, the ability to break the degeneracy between mass and α_3 via priors on the mass that are independent of the photometry (such as dynamical cluster mass measurements) is expected to provide us with the best way forward in our attempts to constrain the IMF slope by photometry. Indeed, we have explicitly shown that, while mass constraints of 10% are required to reliably recover α_3, priors that restrict the allowed mass range to within a factor of two are sufficient to rule out extreme variations in the allowed range of α_3.

In conclusion, while our ability to recover the slope of the upper-end of the IMF via broad-band photometry is limited at present time, there are good prospects to further develop the formalism we have presented here to obtain improved constraints on α_3 with current data, or via future spectroscopic observations in the era of 30 m telescopes.

Acknowledgements

Appendix: Implementing a continuously variable IMF in SLUG

As described in Section 3.2, SLUG is a stochastic stellar population synthesis code which we make use of in our study of IMF variations. This requires a large library of simulated clusters with a wide variety of IMF shapes.

While in its previous version SLUG already handled arbitrary IMF shapes, the parameters specifying the functional form of the IMF were fixed by an input parameter file, which prevented the user from running large simulations with a continuous distribution of clusters with respect to the parameters defining the shape of the IMF. As an example, consider the Kroupa IMF constructed with three power-law segments, having slopes $\alpha_1 = -0.3$, $\alpha_2 = -1.3$, and $\alpha_3 = -2.3$. If we wished to investigate variation in α_3, it would have been necessary to run many libraries of SLUG simulations, each configured with a different input value of α_3. This is clearly not desirable, especially when dealing with simulations of millions of clusters. To circumvent this limitation, we develop a new feature of SLUG which allows users to vary the parameters that control the IMF continuously, in a similar way to what is done with the other physical parameters (e.g. mass, age, and extinction).

To this end, we implement an extension of SLUG's PDF capabilities, introducing nested PDFs. In SLUG, PDFs are generated and handled in a very general way (da Silva et al., 2012, 2014; Krumholz et al., 2015a). Each PDF is described by an arbitrary number of segments, with each segment having a functional form of choice, such as a log-normal distribution or a power-law distribution (a full list of functional forms is given by Krumholz et al., 2015a). In the code, these PDFs are used to represent the IMF, the CMF, Star Formation History (SFH), distribution of extinctions (A_{V}), and the Cluster Lifetime Function (CLF).

Our extension is implemented by enabling a hierarchy of PDFs, by which a "master" PDF can control, at run time, the behaviour of "slave" PDFs. With this new feature, users can define the functional form of the IMF (the "slave" PDF) by means of an arbitrary number of segments as before. However, one or more parameters defining these segments can now be set to a variable mode. The user further specifies the form of one "master" PDF for each variable parameter, from which the numerical values for the variable parameter are drawn during the simulation. Specifically, at run time, SLUG draws new realisations for each of the variable parameters from the "master" PDFs and reinitialises the "slave" PDFs (the IMF in our case) for each cluster accordingly. The end result of this implementation when applied to the IMF case is that we can easily construct libraries of simulations with continuously varying IMFs, using only an input parameter file.

Operationally, the use of this variable feature merely requires users to set the correct flag when defining the PDF segments in the SLUG IMF definition file, and then to create a companion file to define the PDF from which the parameter is drawn. Detailed instructions for the use of this new feature are available in the latest release of the SLUG user manual[2], along with examples. The source code featuring this extension is publicly available from the SLUG repository, and it is now part of the main SLUG release[3].

Here we provide examples of the capabilities of this variable mode, together with tests to validate the new version of the code. As a first example, we consider the simple case of a Salpeter-like IMF (Salpeter, 1955), which is defined by a single power-law segment with slope α. In this run, we set α to variable mode, further providing a Gaussian PDF with mean $\alpha = -2.35$ and dispersion of 0.2 as the "master" PDF. The values of α drawn in 10^4 cluster realisations of a single SLUG run are shown to follow closely the requested normal distribution in Figure 3.15.

Next, we demonstrate the use of the variable PDFs for the case of an IMF with multiple segments. Due to our flexible implementation, we can choose to vary only one of the parameters of this IMF. In this demonstration, we vary only the high-mass slope (α_3) of a Kroupa-like IMF, with its slope generated by evenly selecting between -3 and -1.5 for each cluster realisation. We generate 10^5 realisations of a $500\,\mathrm{M}_\odot$ star cluster, and we plot histograms of the most massive star in each

---

[2]Available at http://slug2.readthedocs.io

[3]Available at http://www.slugsps.com
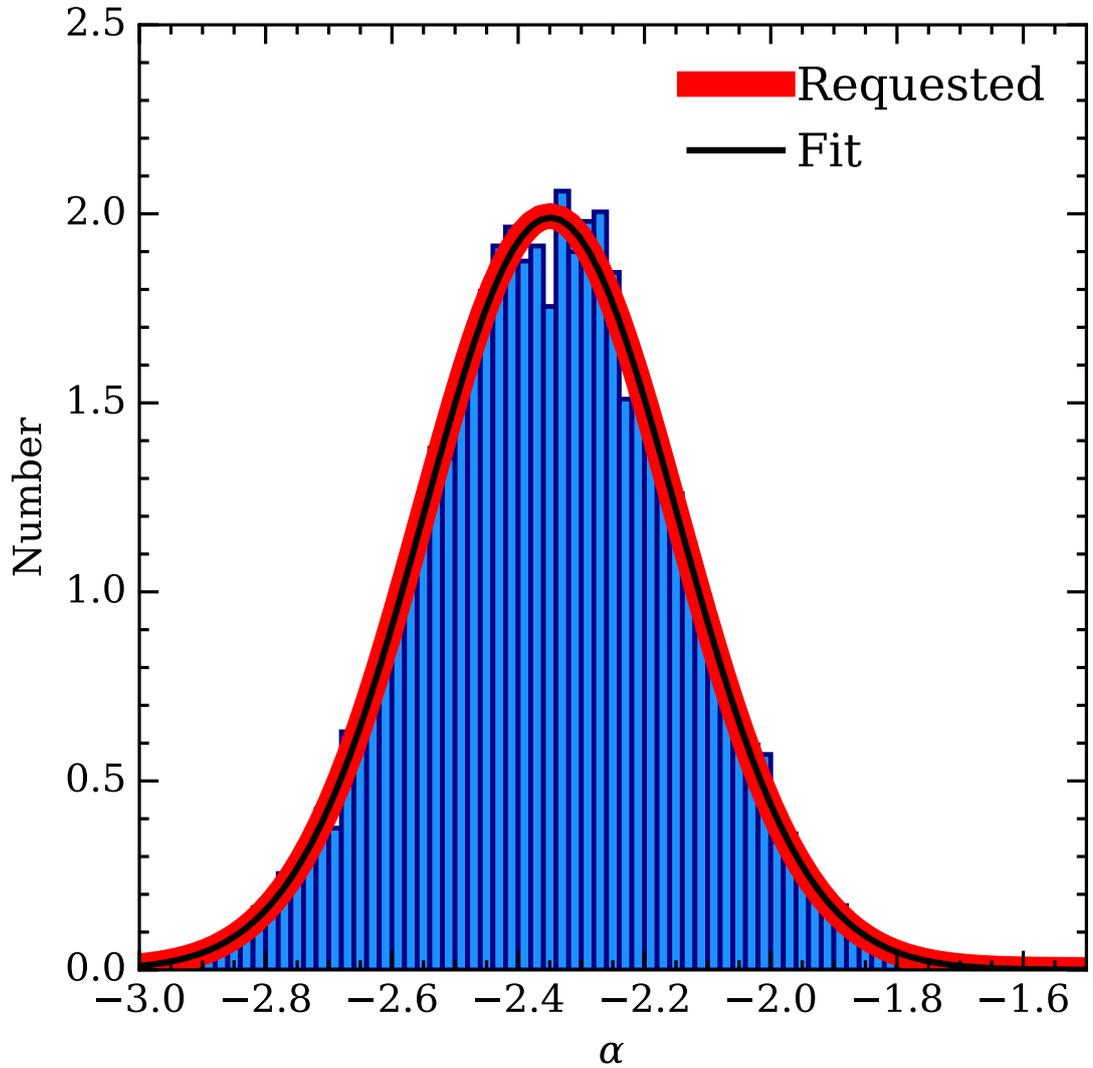
Figure 3.15: Values of the power law index α of a variable Salpeter-like IMF (blue histogram), obtained from 10^4 cluster realisations generated using the newly-developed variable mode in SLUG. The requested normal distribution for α has a mean of -2.35 and a dispersion of 0.2, and is plotted in red. The best fit normal distribution is displayed in black, and closely agrees with the input PDF.
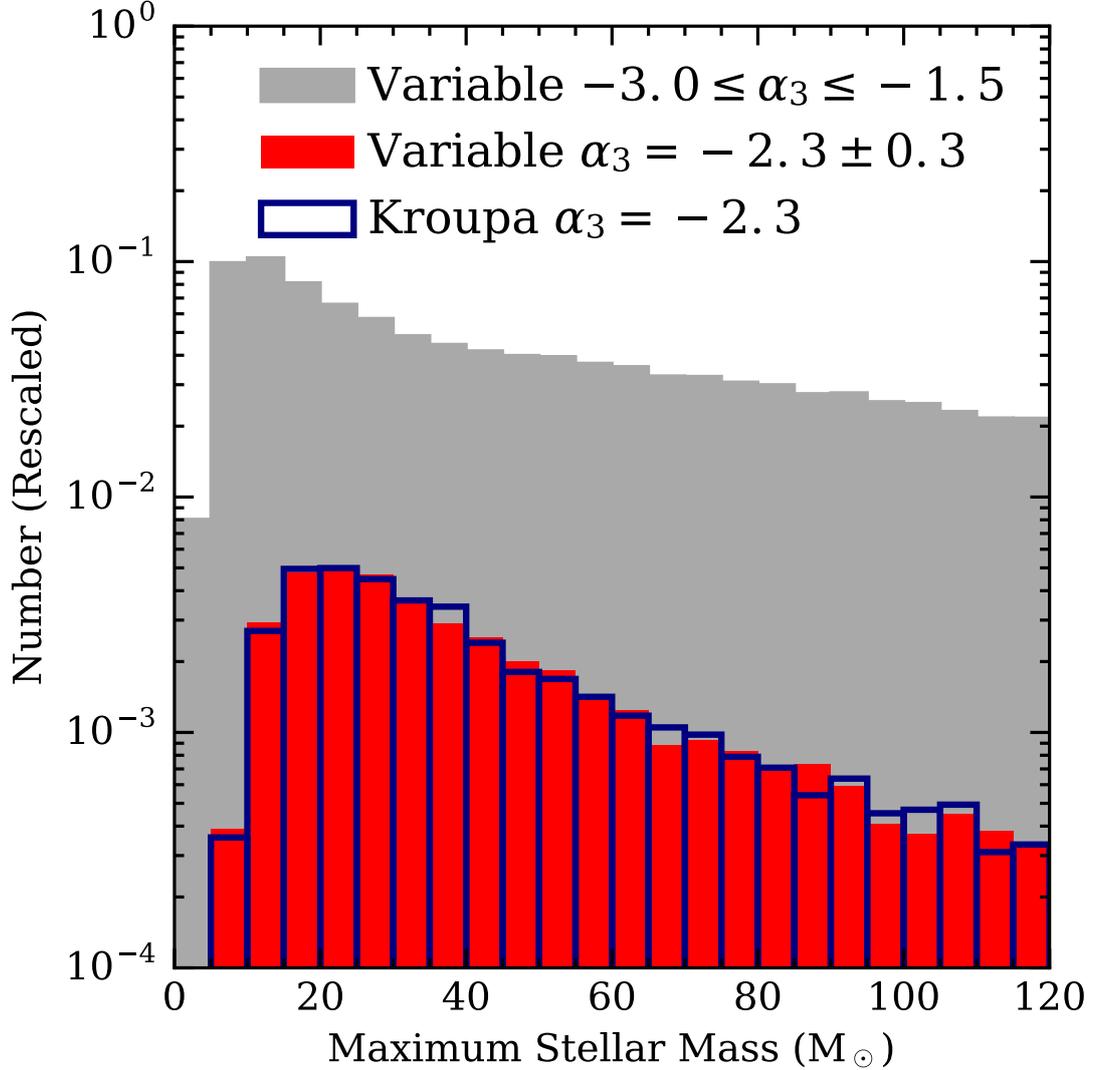
Figure 3.16: Comparison of the maximum stellar mass for two sets of realisations of a $500\,\mathrm{M}_\odot$ cluster at $10^6\,\mathrm{yr}$. In the first set (blue histogram), the stars are drawn from a canonical Kroupa IMF in 5000 realisations. In the second set, we generate 10^5 realisations using the variable IMF mode, with the upper-end slope α_3 drawn from a flat distribution between -3.0 and -1.5. The subset of realisations with a value of α_3 within ± 0.3 of the Kroupa value is shown in red, while the full range is shown in grey. The histograms are rescaled such that there are an equal number of clusters in the bins for the Kroupa and constrained variable cases, such that their size is relative to the complete variable set.

simulated cluster in Figure 3.16. Given enough realisations, we can extract from the resulting library of SLUG models a subset of simulations with $\alpha_3 \approx -2.3$, which we then compare to simulations that have been run in the non-variable mode using a canonical Kroupa IMF. We find excellent agreement when comparing the shapes of the two distributions, which validates our new implementation. Figure 3.16 also shows that, as expected, the inclusion of shallow IMF slopes in the library, with α_3 up to values of -1.5, increases the probability of drawing massive stars compared to a normal Kroupa IMF, skewing the maximum stellar mass distribution to high values.

In summary, our extension to the way SLUG handles PDFs enables users to create large libraries of clusters simulated using a continuous distribution of parameters defining the functional form of the IMF. These libraries can then be used jointly with Bayesian analysis to derive posterior PDFs for the IMF parameters via comparisons with broad-band photometry.

CHAPTER 4

IMF DIAGNOSTICS FROM UV SPECTROSCOPY

Preamble

The following chapter is a reproduction of the submitted manuscript for the following paper:

- *Theoretical predictions for IMF diagnostics in UV spectroscopy of star clusters*
  G. Ashworth, M. Fumagalli, A. Adamo, M. R. Krumholz, *MNRAS* 2018
  (Submitted)

The paper is as it was when submitted to Monthly Notices of the Royal Astronomical Society in 2018, save for minor spelling, grammar, and factual corrections, and the removal of redundant columns from tables and the trimming of figure and table captions such that they fit within the margins of the thesis. In Section 4.6 additional work is presented, with additional analysis for different mock observations covering a narrower range of ages. This additional work will be included in the next revision of the paper.

In this paper, we continue our study of the IMF using the SLUG software suite. Building on the work presented previously in Chapter 3, we move to add more information to our analysis in an attempt to better recover the shape of the IMF. This new information comes in the form of UV spectral features. Young, hot and

massive stars dominate the UV region of the spectrum in a stellar population, and the number of these stars that exist in the population is directly related to the slope of the IMF. However, to include these features in our analysis, SLUG required further development. This development is summarised in this chapter, but is covered in more detail in Chapter 2, in Section 2.3. These modifications included the implementation of the high resolution UV atmosphere library for O and B type stars created by Leitherer et al. (2010), and the development of an equivalent width calculator. These equivalent widths were then able to be passed as additional observational data to BAYESPHOT for use in Bayesian analysis.

Compared to the results presented previously in Chapter 3, we find that including UV line equivalent widths in our analysis greatly improves our ability to constrain the mass, age, and IMF slope parameters for a cluster, compared to using broad-band photometry alone. Indeed, a set of five representative lines was able to reduce the interquartile range of the stacked posterior PDFs calculated for 500 mock observations by $\simeq 32\%$ in the case of the IMF slope parameter. Taking into account the sensitivity limitations that may be present when measuring equivalent widths, we applied cuts to the mock data such that only lines of certain strengths were included in the analysis. Even with a modest sensitivity of $9\,\text{Å}$ in equivalent width, including lines in the analysis was able to reduce the interquartile range of the posterior PDF median residuals by $\simeq 32\%$. Finally, to account for the fact that observations may not be able to probe the full range of ages we present in the analysis, we repeated the study for a narrower range covering $10^{6-7}\,\text{yr}$, finding comparable improvements in IMF slope recovery at $9\,\text{Å}$.

Abstract

We explore the possibility of using UV spectroscopy in combination with broad-band photometry as diagnostic tools for understanding the shape of the Initial Mass Function (IMF) in unresolved stellar populations. Building on our previous work, we extend the Stochastically Lighting Up Galaxies code (SLUG) to include a high-resolution UV spectral synthesizer and equivalent width calculation capabilities. We first gain a qualitative understanding of how UV spectral features behave as the parameters that define a star cluster in SLUG (mass, age, extinction, and IMF slope α_3) are changed. We then exploit Bayesian inference techniques to recover

the α_3 values for clusters simulated with SLUG, using mock observations of these clusters comprised of broad-band photometry and equivalent width measurements of a selection of UV spectral features. We find some improvement when compared to attempts using broad-band photometry alone (with the interquartile range of the α_3 posterior PDF shrinking by $\simeq 32\%$), although we still do not yet fully break the known degeneracy between the cluster mass and α_3. Finally, we make predictions about how effective real observations will be by quantifying our ability to constrain α_3 as a function of limiting equivalent width. We find that observations sensitive to a modest equivalent width of $\simeq 9\,\text{Å}$ are sufficient to improve the recovery of the IMF slope parameter by $\simeq 32\%$ (interquartile range of posterior PDF median residuals), moving to $\simeq 39\%$ when we include all the significant spectral features in the wavelength range $900 - 3000\,\text{Å}$.

4.1 Introduction

Young, bright, massive stars are responsible for much of the ultraviolet (UV) flux of a stellar population, and it is from these stars that many distinctive features in the UV region of the spectrum originate. The number of these stars that exist in a cluster is dependent on the upper end slope of the initial mass function (IMF) (e.g. Salpeter, 1955). The IMF describes the distribution of stars by mass in a stellar population, and is a core ingredient in our understanding of many astrophysical phenomena. As these massive stars contribute strongly to the UV region of the spectrum, this region holds many clues which can aid investigations into the shape of the IMF. Accurately modelling the UV region of a star cluster's spectrum is therefore important if we are to improve our understanding of the IMF from observations of unresolved stellar populations. It is such a task that we undertake over the course of this paper.

Describing the shape of the IMF is an important goal as it is the IMF that defines the makeup of stellar populations. Its shape is generally considered to be universal (Bastian et al., 2010), with most observations suggesting that, at its simplest, it takes the form of a power law of slope ~ -2.3 (Salpeter, 1955), assuming a more shallow slope at low stellar masses (Kroupa, 2001; Chabrier, 2003; Massey et al., 1995; Bastian et al., 2010).

However, the concept of a universal IMF is not without challenge. There exist

many studies which imply variation in the shape of the IMF, and many of these concentrate on the upper end slope. These studies range from theoretical arguments based on temperature and metallicity (Adams & Fatuzzo, 1996; Bonnell et al., 2006; Krumholz, 2011) through to the Integrated Galactic Initial Mass Function (IGIMF) (Kroupa, 2001; Kroupa & Weidner, 2003; Weidner et al., 2011; Kroupa et al., 2013) which is based on the assumption that the most massive star that can form in a star cluster is related to the mass of that cluster (i.e. the upper limit of the IMF). In a recent study, Dib et al. (2017) find suggestions of a variable IMF in the Milky Way through 'lonely' O star fractions, and Schneider et al. (2018) infer a shallow high-mass end IMF slope for 30 Doradus. However, no evidence of IMF variation has been widely accepted in the literature, and a variety of alternative explanations have been proposed to many of these variable IMF indications, from the application of stochastic sampling of the IMF (Fumagalli et al., 2011) to the reported observations of isolated massive OB stars (Oey et al., 2013; Lamb et al., 2016).

In this paper, we continue the work started in Ashworth et al. (2017). In that study, we investigated the possibility of constraining the IMF from broad-band photometry using the Stochastically Light Up Galaxies (SLUG; da Silva et al., 2012, 2014; Krumholz et al., 2015a; Ashworth et al., 2017) stellar population synthesis (SPS) code. Using *Hubble Space Telescope* photometry from the Legacy Extra-Galactic UV Survey (LEGUS; Calzetti et al., 2015), combined with Bayesian analysis exploiting a large library of SLUG models, we demonstrated that accurate recovery of the IMF was non-trivial. Using mock clusters generated with SLUG, we went on to quantify the degeneracy between the mass and the IMF, and showed that using broad-band photometry alone, without accurate knowledge of the cluster mass, we could not put tight constraints on the slope of the IMF.

In this work, we include the UV region of the spectrum in addition to the broad-band photometry used previously. As UV features arise in the spectra of young, hot, and massive stars, including this region of the spectrum in the analysis should allow us to constrain the IMF's high-end slope more closely than was shown to be possible when using broad-band photometry alone. Other studies have made use of this fact, for example, in the recent work of Wofford et al. (2014), the authors use UV spectroscopy to constrain the high-mass end of the IMF (finding it not to be shallower than the Salpeter value).

As in our previous work, we make use of SLUG to perform this study. The paper is structured in the following way. In Section 4.2, we describe new extensions made to the SLUG code in order to produce high-resolution UV spectra. These modifications are then used to construct several libraries of simulated clusters with a wide coverage of the mass-age-extinction-IMF parameter space. In Section 4.3, we continue with a qualitative exploration how the behaviour of the physical parameters that define the stellar populations maps onto UV spectral features and broad-band photometry. Next, in Section 4.4, we produce photometry and equivalent width measurements for a selection of simulated clusters that form a set of mock observations. These observations are then analysed using the Bayesian inference methods used in previous studies (Krumholz et al., 2015b; Ashworth et al., 2017), which we now extend to include spectral features as additional input data. Finally, we make predictions as to what spectral features are of importance when planning observations, and whether observations would be able to recover the high-mass IMF slope of an unresolved stellar population using the methods presented in this paper.

4.2 Implementing high-resolution UV spectra in SLUG

To generate our mock star clusters, we use the SLUG code (da Silva et al., 2012, 2014; Krumholz et al., 2015a; Ashworth et al., 2017). SLUG is a stochastic stellar population synthesis (SPS) code that allows the user to generate star clusters and galaxies by randomly drawing stars from an input IMF, as well as drawing star clusters themselves from a cluster mass function when populating galaxies. SLUG also has the capability to randomly sample the star formation history when populating galaxies. As part of the SLUG suite of software[1], we distribute a Bayesian analysis toolset by the name of BAYESPHOT. The BAYESPHOT tool uses Bayesian inference techniques to calculate posterior probability distribution functions (PDFs) of physical parameters that define the cluster from a set of input observables (generally broad-band photometry; Krumholz et al., 2015a). This is accomplished through the use of kernel density estimation (KDE) techniques, combined with a large library (which we refer to as the 'training set') of SLUG models.

---

[1]Available at www.slugsps.com

4.2.1 Stellar atmosphere grids

Up until now, SLUG has been capable of producing spectra with a resolution of $\sim 10\,\text{Å}$. However, this is not a high enough resolution to allow us to make use of the equivalent widths of absorption line features in our analysis. As it is these features that we wish to use to infer the slope of the high-mass end of the IMF, we have implemented a high resolution spectral capability in the SLUG code. Along with this new capability, the associated machinery required to calculate equivalent widths for arbitrary line features and use them in Bayesian inference has also been implemented.

The highest spectral resolution that SLUG is capable of producing is limited to the resolution of the stellar atmospheres to which it has access. Therefore, in order to produce high resolution spectra, SLUG must have access to high resolution stellar atmosphere models in the wavelength region of interest. To this end, we modified SLUG to handle the high-resolution ($\sim 0.4\,\text{Å}$) atmospheres of Leitherer et al. (2010), which are distributed as part of STARBURST99 (Leitherer et al., 1999). These atmospheres cover a wavelength range of $900 - 3000\,\text{Å}$.

This new spectral synthesiser option, dubbed SB99HRUV in SLUG, covers the UV region of the spectrum ($900 - 3000\,\text{Å}$) using the following atmospheres:

- Stars classified as Wolf-Rayet stars use the Potsdam PoWR atmospheres (Gräfener et al., 2002; Hamann & Gräfener, 2004; Sander et al., 2012).

- We use Kurucz (Kurucz, 1992; Lejeune et al., 1997) atmospheres for all stars with $M < 5\,\text{M}_\odot$, for stars with $T < 170000$ K in the mass range $5 - 10\,\text{M}_\odot$, and for stars with $T < 22000$ K at masses $> 10\,\text{M}_\odot$.

- For all other stars, we use the high resolution IFA atmospheres (Leitherer et al., 2010).

- Stars that fall outside of these atmosphere grids are treated as black bodies, although this is a rare occurrence given our choice of tracks.

For the rest of the wavelength range of the spectrum, we use the standard lower resolution atmospheres (equivalent to selecting the default SB99 spectral synthesizer option, which mimics the atmosphere selection of STARBURST99) that are employed by SLUG. Further information is available in Krumholz et al. (2015a)

and the SLUG user manual. The core differences between this and the high resolution region described above are the use of the Hillier atmospheres (Hillier & Miller, 1998) for Wolf-Rayet stars, and the use of Pauldrach OB atmospheres (Pauldrach et al., 1998) instead of the new IFA atmospheres.

All the atmospheres used during the work performed in this paper are those that are distributed with either SLUG, or in the case of the PoWR atmospheres and the high resolution IFA UV atmospheres, those that are distributed with STARBURST99.

Finally, we stitch the high and low resolution spectra together, replacing the $900-3000\,\text{Å}$ region of the low resolution spectrum with the high resolution spectrum calculated from the high resolution atmospheres.

4.2.2 Selecting atmospheres

When selecting the high-resolution atmospheres, we have implemented both the nearest-neighbour algorithm used by STARBURST99, and a new interpolation scheme whereby the atmosphere for a particular star is calculated through a barycentric interpolation. This can be toggled by simply changing a Boolean value in the high resolution UV spectral synthesiser before compiling SLUG. The nearest-neighbour scheme allows direct comparison with STARBURST99, and is also less computationally intensive.

There are several steps to our new interpolation scheme. We begin with the parameters of the star for which we wish to obtain an atmosphere. First we find the two nearest neighbour atmospheres in the temperature-surface gravity plane that exist within the atmosphere library. We then iterate through the next nearest neighbours until we satisfy the conditions required to enclose the star within a triangle of three atmospheres within this plane. Having formed this triangle of nearest neighbour atmosphere models, we find the barycentric coordinates of the star of interest within this triangle, and then use these to weight the atmospheres on each vertex of the triangle. The weighted combination of these atmospheres gives us an interpolated model for the star. If the triangle cannot be closed, we default to the nearest-neighbour method. We use this interpolated scheme during this paper rather than the nearest neighbour approach due to the sometimes large gaps in the parameter space. The definition of 'nearest neighbour' is the same one used by

STARBURST99 for these atmospheres, with each direction (in the surface gravity - temperature plane) weighted differently (with the surface gravity axis weighted by a factor of 5; Leitherer et al., 2010).

To test the scheme, a test star was selected, and then removed from the atmosphere grid. The spectra of the nearest neighbour star, and the interpolated atmosphere, were then compared to the true atmosphere of the missing star. The interpolated star lay noticeably closer to the true value across most absorption features [2]. Our work is based on mock observations constructed with a consistent choice of interpolation scheme. These mock observations are then analysed using libraries of SLUG models constructed using the same interpolation scheme. Thus, we defer the study of which scheme performs best when applied to observations to future work.

We further note that, due to the low density of points in some regions of the temperature-surface gravity plane, certain combinations of parameters can lead to significant differences in flux between the high resolution spectrum and the low resolution spectrum that we stitch into, as in STARBURST99. This leaves discontinuities at the borders of the high resolution region in some cluster parameter configurations, and affects photometry for filters that cross into this region. To ensure smooth transitions, we integrate a region in both the low resolution and high resolution spectra (covering $2400 - 2900\,\text{Å}$) and scale the high resolution spectrum by the ratio between these integrals, so that we can eliminate any large discontinuities at the points where the two spectra join. It is recommended, however, that users planning to compare SLUG photometry alone to real data use the low resolution spectrum.

4.2.3 Equivalent width calculations

As the high-resolution spectra are provided in both line and continuum form, we are able to normalise the spectrum without the difficulties of attempting to create a pseudo-continuum around each line feature. Therefore, we use only the rectified spectrum (as output by STARBURST99, but for O and B type stars alone) to calculate equivalent widths, which gives us a normalised spectrum consisting of features

---

[2]Even when using the basic nearest-neighbour atmosphere selection scheme, results produced by SLUG will not be directly comparable to STARBURST99 due to the differing track methods that SLUG employs. However, using an older revision of the SLUG code, and using the nearest-neighbour atmosphere selection method, we were able to validate our implementation by reproducing the rectified and non-rectified spectra produced by STARBURST99 for a selection of test clusters.

belonging only to the stars covered by the new high resolution atmospheres. It should be noted that these rectified spectra are not extincted, and contain no nebular emission. We do, however, run SLUG with nebular emission calculations enabled for the work in this paper, and both it and extinction are applied to the photometric filters and the raw non-rectified spectrum. As we are comparing SLUG mock observations to libraries of other SLUG models, this approach is internally consistent for this study.

Finally, the equivalent width is calculated by interpolating linearly between each point in the spectrum (as they are the bin-centres of variable width bins, this is a best approximation), and then integrating under the resulting function across the integration region for a given spectral feature.

Further information on the use of these new features is available in the SLUG documentation[3]. The code is currently publicly available, including the high-resolution IFA atmospheres. The PoWR Wolf-Rayet atmospheres will be added to the distribution when new versions from the PoWR group are made available (W. R. Hamann 2017, Private Communication).

4.3 Grids of high-resolution SLUG models

In order to carry out the work described in this paper, we must generate a large selection of SLUG models. These models will be used for empirical study, to form mock observations, and to form the large library of simulations required by BAYESPHOT. The parameters that define these different models are detailed in Table 4.1.

The 'Library' is the large set of models that makes up the 'training set' we use for our Bayesian analysis. The Library contains 5×10^7 model star clusters, with their masses drawn from a flat distribution between $5 \times 10^2 \, M_\odot$ and $5 \times 10^5 \, M_\odot$, and their ages drawn from a flat distribution between $5 \times 10^4 \, \mathrm{yr}$ and $5 \times 10^7 \, \mathrm{yr}$. The extinction, A_V is drawn from a flat distribution between 0 and 3, and the IMF high-mass slope α_3 is drawn from a flat distribution between -3.2 and -1.3. The actual shape of the distribution is not important during the Bayesian analysis, but a flat distribution makes sense in this context as the clusters we use to generate our input photometry are also generated using flat distributions. Nebular emission

---

[3] http://slug2.readthedocs.io

Name	$M_{*,\mathrm{Max}}$	Z	$\log M$ (M$_\odot$)	$\log t$ (yr)	#[a]	α_3
Library	120	0.020	2.7–5.7(C)	4.7–7.7(C)	5×10^7	$-3.2 - -1.3$ (C)
Mock Observations	120	0.020	3–5(C)	5–7(C)	5×10^2	$-3.0 - -1.5$ (C)
Set 1 : High Cutoff, Solar Z	120	0.020	3,6(D)	5,6,7(D)	10^3	$-3.0, -2.3, -1.5$ (D)
Set 2 : High Cutoff, High Z	120	0.050	3,6 (D)	5,6,7(D)	10^3	-2.3
Set 3 : High Cutoff, Low Z	120	0.004	3,6(D)	5,6,7(D)	10^3	-2.3
Set 4 : Low Cutoff, Solar	30	0.020	3,6(D)	5,6,7(D)	10^3	$-3.0, -2.3, -1.5$ (D)
Set 5 : Low Cutoff, High Z	30	0.050	3,6(D)	5,6,7(D)	10^3	-2.3
Set 6 : Low Cutoff, Low Z	30	0.004	3,6(D)	5,6,7(D)	10^3	-2.3
Squares	120	0.020	3,6(D)	5–7(C)	10^5	$-3.0 - -1.5$ (C)

[a] Number of realisations

Table 4.1: Table of the parameters used in SLUG for generating the mock star clusters used within this work. The base IMF used for all these model sets is the 3 component broken power law of the Kroupa IMF. The break points between the segments of this power law are situated at 0.01, 0.08, 0.5, $M_{*,\mathrm{Max}}$. Z is the metallicity. The (D) represents discrete sampling, with # realisations at each value, and (C) represents continuous sampling, with # realisations drawn from a flat distribution across the given range.

Feature Name	Integration Range (Å)	ϕ^{a} (eV)
Ovɪ	$1015.0 - 1050.0$	138.1189
Ovɪa†	$1015.0 - 1037.0$	–
Ovɪb	$1037.0 - 1050.0$	–
Cɪɪɪ†	$1174.0 - 1177.0$	47.88778
Siɪɪɪb	$1205.0 - 1207.0$	33.49300
Nv	$1220.0 - 1250.0$	97.89013
Nva†	$1220.0 - 1240.0$	–
Nvb	$1240.0 - 1250.0$	–
Bl1302 *	$1292.0 - 1312.0$	–
Cɪɪ†	$1334.0 - 1338.0$	24.3845
Ov	$1355.0 - 1375.0$	113.8989
Ova†	$1355.0 - 1370.0$	–
Ovb	$1370.0 - 1375.0$	–
Siɪv * †	$1387.0 - 1407.0$	45.14179
Bl1425 *	$1415.0 - 1435.0$	–
Siɪɪɪa †	$1416.0 - 1418.0$	33.49300
Fe1453 * †	$1440.0 - 1466.0$	–
Cɪv	$1527.0 - 1569.0$	64.49351
Cɪva * †	$1530.0 - 1550.0$	–
Cɪvc *	$1540.0 - 1560.0$	–
Cɪve *	$1550.0 - 1570.0$	–
Bl1617 *	$1604.0 - 1630.0$	–
Heɪɪ †	$1639.5 - 1641.0$	54.4177650
Bl1664 *	$1651.0 - 1677.0$	–
Nɪv †	$1703.0 - 1727.0$	77.4735
Bl1719 *	$1709.0 - 1729.0$	–
Bl1853 *	$1838.0 - 1868.0$	–
Siɪ †	$1843.0 - 1846.0$	8.15168
Feɪɪa *	$2382.0 - 2422.0$	16.19920
Bl2538 *	$2520.0 - 2556.0$	–
Feɪɪb *	$2596.0 - 2622.0$	16.19920
MgWide *	$2670.0 - 2870.0$	–
Mgɪɪ * †	$2784.0 - 2814.0$	15.035271
Mgɪ * †	$2839.0 - 2865.0$	7.646236

$^{\mathrm{a}}$ Ionisation potential

Table 4.2: Table of the line features used. Integration regions used for calculating the equivalent width of each feature are listed, with an ionisation energy if appropriate. The ionisation potentials given were obtained from Kramida et al. (2015). Lines are sorted by increasing leading edge wavelength of their integration region. Spectral features taken from Fanelli et al. (1992) are marked with an asterisk (*). In Cɪv, the letters 'a','c','e' represent the absorption, centre, and emission respectively as given in Fanelli et al. (1992). For other P-Cygni features such as Nv, the 'a' and 'b' represent the absorption and emission components respectively. For other features such as Siɪɪɪ, the 'a' and 'b' represent two occurrences of the same feature, with 'a' being the stronger in equivalent width. Features marked with a † mark are those used in the full-set analysis as they are sufficiently independent of one another. Sub-regions of larger spectral features are indented.

is calculated using an ionising photon fraction $\phi = 0.73$, the default in SLUG. It must be noted once again, however, that nebular emission (much like extinction) is not included in the rectified spectrum used for calculating equivalent widths from the combinations of high resolution stellar atmospheres. However, much of the increased flux from the nebular emission, and the effects of extinction, would be negated by the normalisation process, although some features would be absent. In addition, as the observations are generated with the same restrictions as the models used in the Bayesian analysis the work is self consistent.

The 'Mock Obs.' models are those whose photometry and equivalent widths form our set of mock observations that we use in our Bayesian analysis. Their masses are drawn from a slightly more narrow set of distributions in mass, age, and IMF than the library, as shown in Table 4.1. This means we have no mock clusters situated right at the edge of the library in any part of parameter space (save for extinction, where we cannot go negative), or even outside of it. This is a desirable property when studying clusters with BAYESPHOT. 'Sets' 1 to 6 are models used for studies of both individual lines and the behaviour of the equivalent width for varying age, mass, α_3, and metallicity. Each set is comprised of 1000 realisations, enough for the median equivalent widths to be acceptably converged. Finally, the 'squares' set is a group of models intended to be used to investigate the changes in equivalent width in the age-α_3 plane. These models are drawn from flat distributions in $\log t$ and α_3.

All of these sets of models have both photometry, along with equivalent widths for the features listed in Table 4.2. A selection of the features are taken from Fanelli et al. (1992). Others were defined by inspection of test model spectra combined with reference to NED[4] and Kramida et al. (2015). The photometric filters are one of the standard five-band filter sets used in the LEGUS project (Calzetti et al., 2015): F275W, F336W, and F555W filters from Wide Field Camera 3 (WFC3 UVIS), and the F435W, and F814W filters from the Advanced Camera for Surveys (ACS). In addition to photometry and equivalent widths, we also generate full spectra for the models referred to as 'Sets' 1 to 6.

We note that at ages less than a few Myr, star clusters are embedded in gaseous clouds and are therefore unable to be probed by observations. However, in this theoretical study, the simulations at the age-point of 0.1 Myr are an illustrative

---

[4]The NASA/IPAC Extragalactic Database (NED). is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

extreme low-age case.

4.3.1 Effects of physical parameters on spectral features

To investigate the effects of changing the variable physical parameters on star cluster spectra in a qualitative manner, we run several sets of SLUG models for a variety of IMF slopes, metallicities, cluster masses, and cluster ages. Their parameters are listed in Table 4.1 as Sets 1–6.

As a first example, we look into the effects of changing the value of α_3 on the equivalent widths of the spectral features, and how these equivalent widths then evolve with time. This is done for both low mass ($10^3\,\mathrm{M_\odot}$) clusters and high mass ($10^6\,\mathrm{M_\odot}$) clusters to get a handle on the effects of stochasticity on the equivalent widths we observe. Results are shown in Figure 4.1 for the NIV feature, where the points correspond to the median value from the set of 1000 realisations, and the error bars cover the extent of the 5th to 95th percentile range. We stress that the size of the error bars is not Poisson error, but arises from the stochastic sampling of the IMF (an effect most pronounced in low mass star clusters and clusters with bottom heavy IMFs).

The effects of stochasticity are readily apparent when comparing the low and high mass clusters (panels (a) and (b) to panels (c) and (d) respectively), with the low mass clusters having a wide scatter in equivalent width at all times. This stochasticity is mostly absent for the large $10^6\,\mathrm{M_\odot}$ clusters, with some scatter being apparent for the extremely steep IMF, as expected, due to the small possibility of generating a large star instead of the much more probable large number of small stars.

In the case of the IMF with the $120\,\mathrm{M_\odot}$ cutoff (panels (a) and (c)), we see an upward trend in equivalent width for all three of the α_3 values, although it is somewhat obscured by the scatter in the case of the low mass cluster with a steep IMF. We see much the same behaviour for the truncated IMF in panels (b) and (d), with an overall increase in equivalent width compared to the IMF with the higher mass cutoff. This truncation results in a reduction of extremely hot and massive stars, which in turn leads to less of the Nitrogen being highly ionised, and therefore results in more absorption at young ages for shallower IMF slopes. Already we are able to infer that including lines in our analysis helps to constrain α_3, although this

(a) N IV: $\log\left(\frac{M}{M_\odot}\right) = 3$, IMF cutoff at $120\,M_\odot$

(b) N IV: $\log\left(\frac{M}{M_\odot}\right) = 3$, IMF cutoff at $30\,M_\odot$

(c) N IV: $\log\left(\frac{M}{M_\odot}\right) = 6$, IMF cutoff at $120\,M_\odot$

(d) N IV: $\log\left(\frac{M}{M_\odot}\right) = 6$, IMF cutoff at $30\,M_\odot$
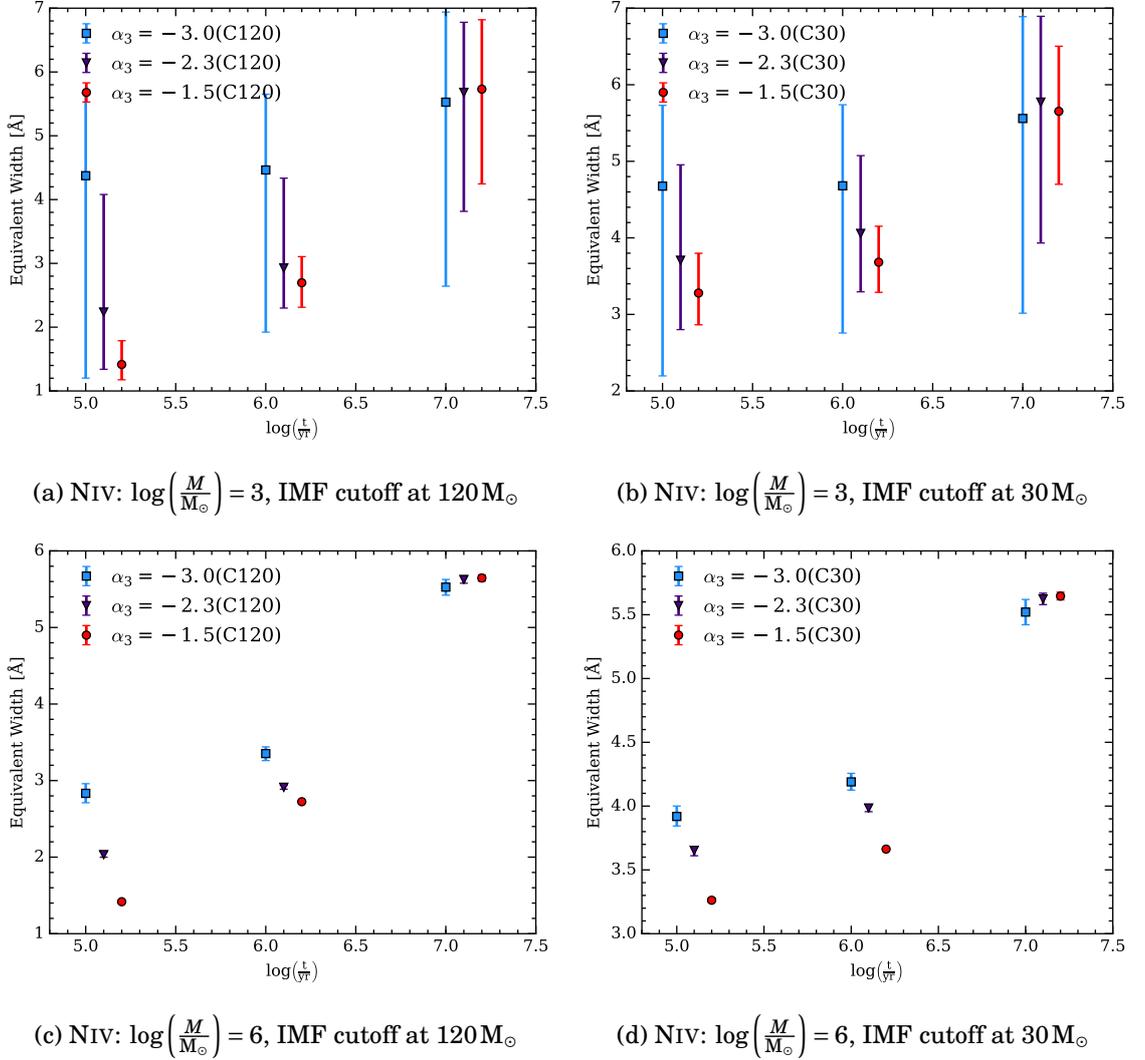
Figure 4.1: Evolution of equivalent width of N IV with cluster age for a range of IMF upper-end slopes. Panels (a) and (b) are clusters with a mass of $10^3\,M_\odot$, whereas panels (c) and (d) are clusters with masses of $10^6\,M_\odot$. Panels (a) and (c) have the IMF cutoff at $120\,M_\odot$, whereas panels (b) and (d) have the IMF cutoff at $30\,M_\odot$. The points correspond to the median value from the set of realisations, and the error bars cover the fifth to ninety-fifth percentile range. Blue squares correspond to $\alpha_3 = -3.0$, purple triangles correspond to $\alpha_3 = -2.3$, and red circles correspond to $\alpha_3 = -1.5$. The leftmost point of each trio is positioned at the correct age value, and the others are shifted slightly to aid readability.
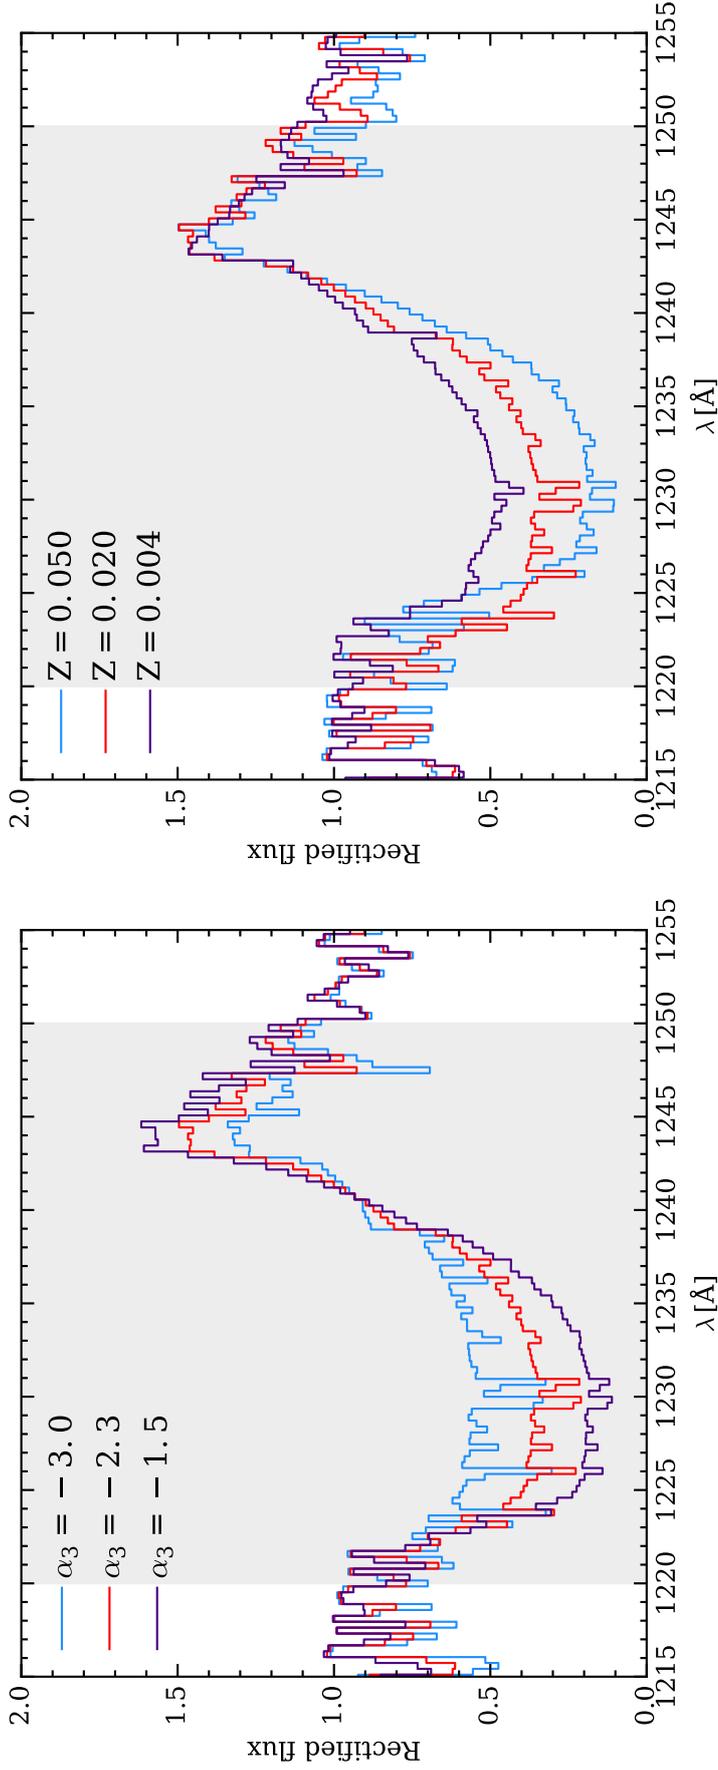
Figure 4.2: Renormalised stacked spectra of 1000 model star clusters of mass $10^6 M_\odot$ at an age of 10^5 yr. In panel (a), each set of 1000 was generated at a different IMF ($\alpha_3 = -3.0, -2.3, -1.5$), at solar metallicity. In panel (b), each set of 1000 was generated using a different metallicity, with $Z = 0.004, 0.020, 0.050$. All models were generated using the canonical Kroupa IMF as a base IMF shape. The spectrum shows the region around the Nv line, with the integration region used for the calculation of the equivalent width outlined in a light grey colour. We see the line increase in strength as the IMF becomes shallower. Note also the degeneracy in the absorption region between the IMF and the metallicity. The emission region is unaffected by changing metallicity.

may be hampered by stochasticity at low masses.

We now move to examine the shape of the spectra themselves, a study which relates to the analysis of stacked spectra. We also look at how metallicity affects the spectrum, and how this compares with the effects caused by variations in the IMF's parameters. A sample of the spectrum around the Nv features is given in Figure 4.2, with panel (a) showing how the spectrum varies at a fixed metallicity (solar), fixed age (10^5 yr), and fixed mass (10^6 $M_\odot$) for changing values of the IMF slope ($\alpha_3 = -3, -2.3, -1.5$). Panel (b) shows the how the spectrum varies at the same fixed age and mass, but with a fixed IMF slope of $\alpha_3 = -2.3$, and metallicities of $Z = 0.004, 0.020, 0.050$. In the changing IMF case, we see that the strength of the absorption line feature increases as the IMF becomes more shallow, due to the increased probability of producing hot, massive stars. We see that changing metallicity has a similar effect on the Nv line that changing the IMF does, and a high-metallicity environment appears to be degenerate with a top heavy IMF in the "Nva region" of the feature. However, the "Nvb integration region" (the emission feature) is not appreciably affected by the changing metallicity, whereas it does differentiate between the three IMF slopes.

Unfortunately, we are at present unable to vary the metallicity continuously in SLUG because we lack a set of atmosphere models that is sufficiently densely sampled to allow for effective interpolation. Should a continuous variation of metallicity be implemented in SLUG, it would be possible to introduce it as a fifth physical variable in the Bayesian analysis. This is a degeneracy which should be explored in future work when the necessary features have been implemented in SLUG. For the rest of this study, however, we maintain a constant solar metallicity.

Figure 4.3 shows how the equivalent width of two spectral features changes through the age-α_3 plane. Each panel contains the same set of 10^6 cluster realisations drawn randomly from the distributions listed in Table 4.1 (the 'Squares' set of models). The equivalent width assigned to each grid element is the median of all the realisations that fall within that element. Panel (a) shows the plane for the CIII feature, and panel (b) shows the plane for the CIVa feature, chosen as they both illustrate well how the equivalent width changes with the two variables, the latter being a strong line and the former being weak. We see that both CIII and CIVa (to a lesser extent) allow us to differentiate the IMF slope at young ages, whereas for older clusters we lose this distinguishing power as the massive stars
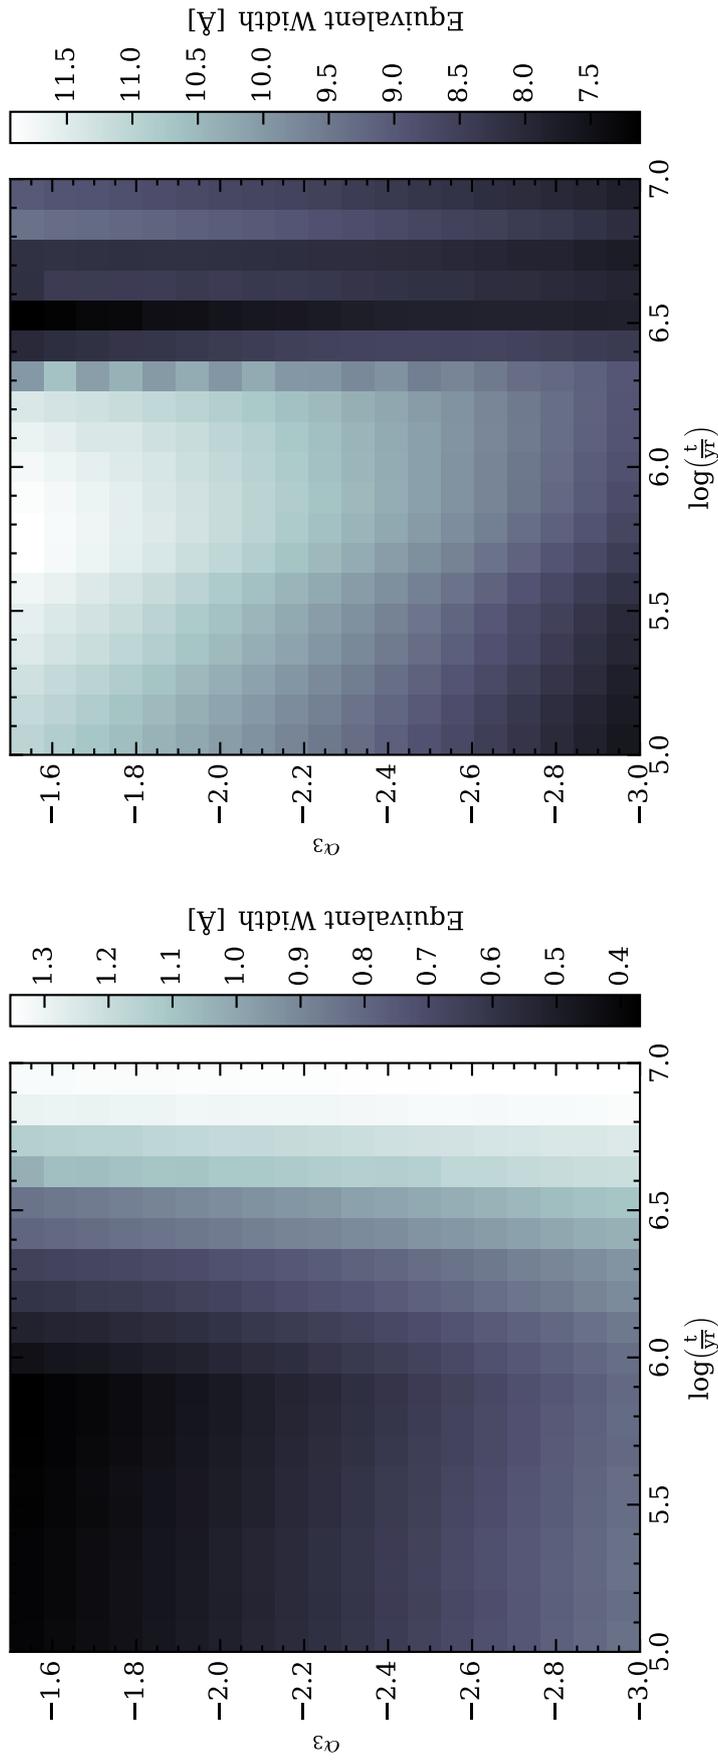
Figure 4.3: Two dimensional plots of equivalent width versus changes in cluster age or IMF slope for two spectral regions (CIII and CIVa), where the equivalent width reported for a grid cell is the median of the equivalent widths for the clusters that reside in that area of parameter space. In both panels we see that the equivalent width of these lines may be used to differentiate the various possible values of α_3 at young ages. In addition, the combination of the two lines provides us with a way to also differentiate the age of the cluster in conjunction with the value of α_3.

that are responsible for the line features evolve outside the main sequence. Note that, unlike the CIV line which increases in strength as the IMF becomes shallow, the CIII line becomes weaker as this happens. This effect occurs because a shallower IMF results in a higher proportion of the cluster's light coming from stars that lack the C III feature due to them being hot enough to ionise carbon past the C^{+2} state responsible for producing the line.

We once again note that observations are limited in how young an age they are able to probe, and would be likely limited to ages of a few Myr or greater. However, even at these ages, we still see useful differences in the equivalent widths in Figure 4.3.

4.4 Analysis of mock observations

Having discussed the qualitative behaviour of the equivalent width as a function of physical parameters, we now move to study more quantitatively how combining the equivalent widths of spectral features with broad-band photometry may be used to constrain the slope of the IMF, α_3.

The mock observations are produced from the photometry and equivalent widths of a selection of SLUG model clusters. These clusters were generated by drawing their physical parameters randomly, as listed in Table 4.1 ("Mock Obs."). It is using this set of mock observations that we attempt to recover the underlying physical properties of the corresponding simulated clusters. For each set of mock observations we extract the photometry and equivalent widths, assigning each a nominal error of 0.05. We must stress that the mock clusters are not taken from the library, and are generated independently. However, as we use the same tracks and atmospheres to generate the mock clusters, the result is somewhat optimistic compared to observations.

The first step is to generate posterior PDFs (1D and 2D) for the physical parameters of each mock cluster that we have generated. An example (cluster 19) is plotted in Figure 4.4, with red markings corresponding to the true values of each parameter. The analysis was performed four times for each cluster. Firstly, using one of the LEGUS filter sets (as used for NGC 628e, see Adamo et al. (2017)), which is comprised of the F275W, F336W, and F555W filters from WFC3 UVIS, and the F435W, and F814W filters from ACS. Next the procedure is repeated using the
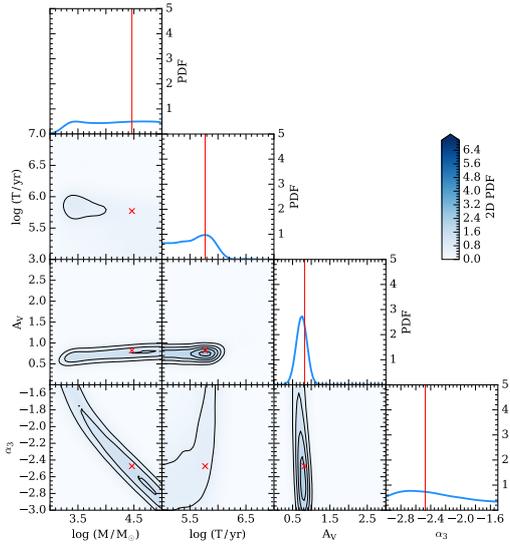
(a) LEGUS Filters
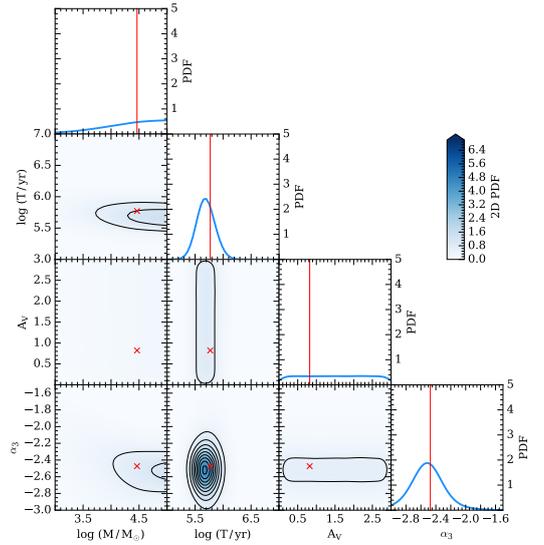
(b) Spectral features: CIII, CIVa, NVa, OVa, OVIb
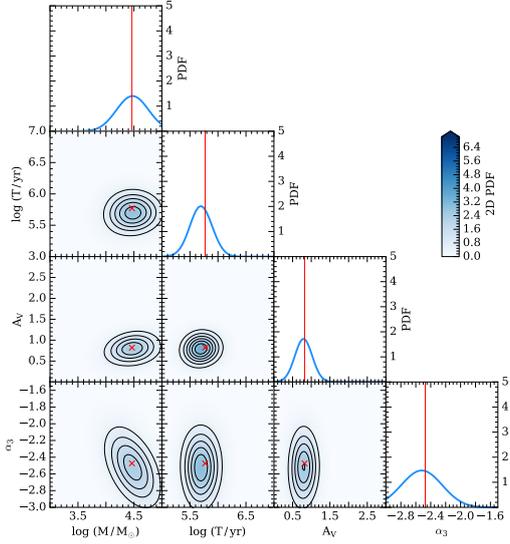
(c) Filters + CIII, CIVa, NVa, OVa, OVIb
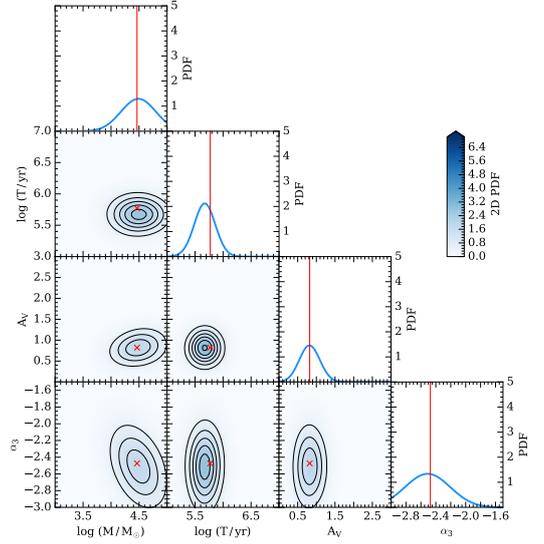
(d) Filters + All independent spectral features

Figure 4.4: Example corner plot for one of the mock observation clusters, cluster 19. The 1D posterior PDFs for $\log M$, $\log T$, A_V and α_3 are shown in the top panels of each column. The joint posterior PDFs are represented by contours (beginning at 0, with increments of 0.5 unit element). They are also represented by a 2D map, where the intensity of the colour reflects the probability density as indicated on the colour bar. All PDFs are normalised to have unit integral. In panel (a) are the results when the analysis is performed using only the LEGUS filters, in panel (b) we show results from using the representative selection of lines only, and in panel (c) we show results from using a combination of the two. In panel (d) we present the result from analysis using all the line features that are independent (see Table 4.2). Note the improvement in the recovery of both the cluster's age and IMF slope when including the spectral lines in the analysis. The red crosses and lines mark the positions of the true values for the parameters that describe this cluster.

following spectral features in place of the photometry: CIII, CIVa, Nva, Ova, Ovib. These features were selected as a representative set of five, with reference to their appearance in the IMF-age plane (c.f. Figure 4.3), which we use to identify those lines most able to distinguish the cluster age or the IMF slope whilst also giving us a range of line strengths and including an emission feature. Next, the analysis is repeated using both the photometry and the equivalent widths of these spectral features together, first using this representative set of five lines, and then finally using all the spectral features listed in Table 4.2 marked with a † symbol. We cannot include all the line features in the table in our analysis at the same time, as not all the features are independent, as some of our features are sub-features of a single feature we also list, are blends containing many features, or are even multiple occurrences of the same line. For example, we list both the full line feature and the separate components of the lines with P-Cygni profile. For the purposes of this analysis, we choose to use only the absorption component when using the "full" independent line list, and indeed the shoulder features may depend greatly on the assumed wind model and therefore may include more uncertainties. We also discard the weaker feature when a line occurs more than once, and we remove the blend features which contain multiple lines.

Panel (a) of Figure 4.4 shows the results of using the broad-band photometry alone. We see quite good recovery of the cluster age and the extinction. However, recovery of the cluster's mass and the value of α_3 is poor, due to the underlying degeneracy between the two. Conversely, in panel (b) we see that lines alone recover the age and the IMF slope much more successfully, at the expense of the mass and the extinction. This is expected as A_V is unconstrained by construction, and the equivalent width (having been calculated from a normalised spectrum) is weakly dependent on the total cluster mass. In panel (c), we see the combination of the two results in good recovery for all the physical parameters. Indeed, the inclusion of the UV lines further refines our determination of the cluster's age, and helps to alleviate the mass-IMF degeneracy that effects analysis performed with broadband photometry alone (Ashworth et al., 2017), although it does not fully break it. Panel (d) is the same as panel (c), but the analysis was performed with all the line features. The results are mostly comparable, with slight shifts towards more accurate mass, age, and IMF, and a spreading out in the extinction direction. We conclude therefore, that the representative set of five lines is sufficient to greatly

Range (Percentiles)	$\Delta \log M$ (M$_\odot$)	$\Delta \log t$ (yr)	ΔA_{V}	$\Delta \alpha_3$
5-95 (MC)	-0.61 (-35%)	-0.63 (-49%)	$+0.18$ ($+31\%$)	-0.34 (-23%)
25-75 (MC)	-0.30 (-42%)	-0.18 (-41%)	$+0.08$ ($+36\%$)	-0.24 (-35%)
5-95 (TC)	-0.84 (-39%)	-0.87 (-55%)	$+0.14$ ($+20\%$)	-0.50 (-26%)
25-75 (TC)	-0.33 (-41%)	-0.22 (-45%)	$+0.07$ ($+27\%$)	-0.24 (-32%)

Table 4.3: Table of the difference in inter-quartile and inter-percentile ranges for the median-centred (MC) and true-centred stacks (TC) shown in Figure 4.6, changing from the case with photometry alone to the case including both photometry and the selected UV features. The numbers in parentheses correspond to the percentage change in the ranges, compared to the original value.

improve the recovery of the cluster's physical parameters.

In Figure 4.5 we plot the residuals (the difference between either the median or mode of the posterior PDFs and the true value of the corresponding parameter) for ten of the mock observation clusters. We see much the same effect as we see in Figure 4.4 in most cases, and there is some improvement overall in the recovery of the IMF. Note also how the medians and modes move much closer together, as the posterior PDFs become narrower. The overall trends seen for Cluster 19 in Figure 4.4 apply more generally across our mock observations.

Finally, we quantify the overall behaviour of the PDFs across our entire input dataset by looking at stacks of the all the posterior PDFs for 500 mock observation clusters, which we show in Figure 4.6. In panels (a) and (b) the posterior PDFs are centred on their medians, whereas those in panels (c) and (d) are centred on the true values of each parameters for each of the clusters in the stack. Panels (a) and (c) are produced using only the five LEGUS filters in the analysis, whereas panels (b) and (d) are produced using a combination of the both the photometry and the equivalent widths of the lines in our representative set of five. The 5th and 95th percentiles are marked with orange dot-dashed lines, and the first and third quartiles are marked with indigo dashed lines. The zero marker, on which the PDFs are centred as they are stacked, is the solid red line (or cross in the 2D PDFs).

As seen in the previous examples, there is significant improvement overall in our recovery of the physical parameters once the spectral features are included. Indeed, we see for the true-value centred stacks that we move from quite wide distributions around the value of interest to quite narrow peaks. This can be seen more clearly by looking at the changes in the interpercentile (fifth to ninety-fifth) and interquartile
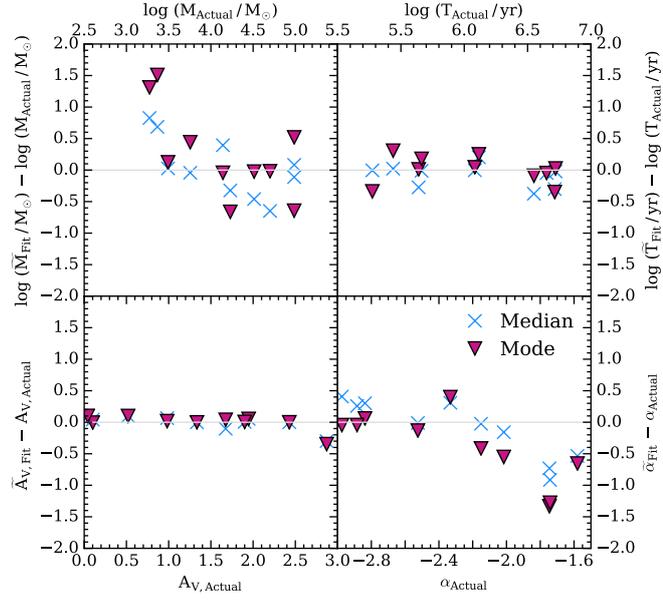
(a) LEGUS filters
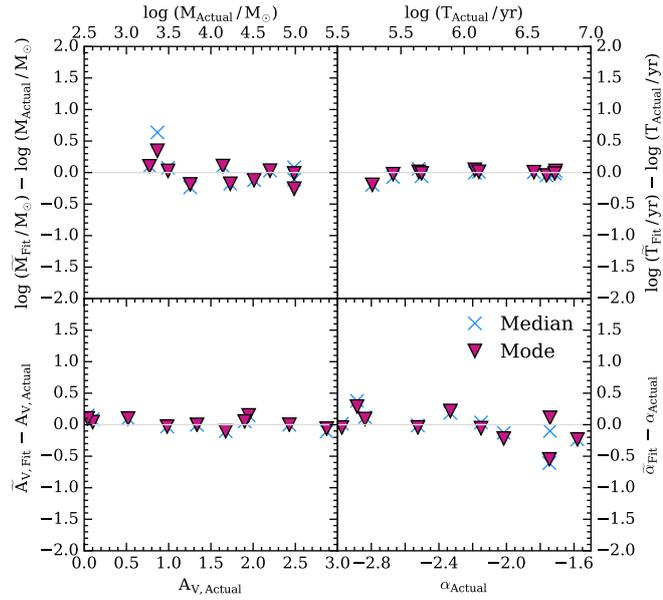
(b) LEGUS filters + lines (CIII, CIVa, NVa, OVa, OVIb)

Figure 4.5: Residuals of the 1D posterior PDFs for the parameters $\log M$, $\log t$, A_{v}, and α_3 for 10 of the clusters generated to provide mock observations. The residuals are calculated from both the medians of the distributions (blue crosses) and the modes of the distributions (purple triangles), when compared to the true values for each of these parameters. We see that the inclusion of the lines into the analysis improves the recovery of all of the parameters, and that in most cases the median and mode move closer together, indicating a narrowing of the distributions as well.

(a) LEGUS Filters, Median centred

(b) Lines + LEGUS filters, Median centred

(c) LEGUS Filters, True centred

(d) Lines + LEGUS filters, True centred
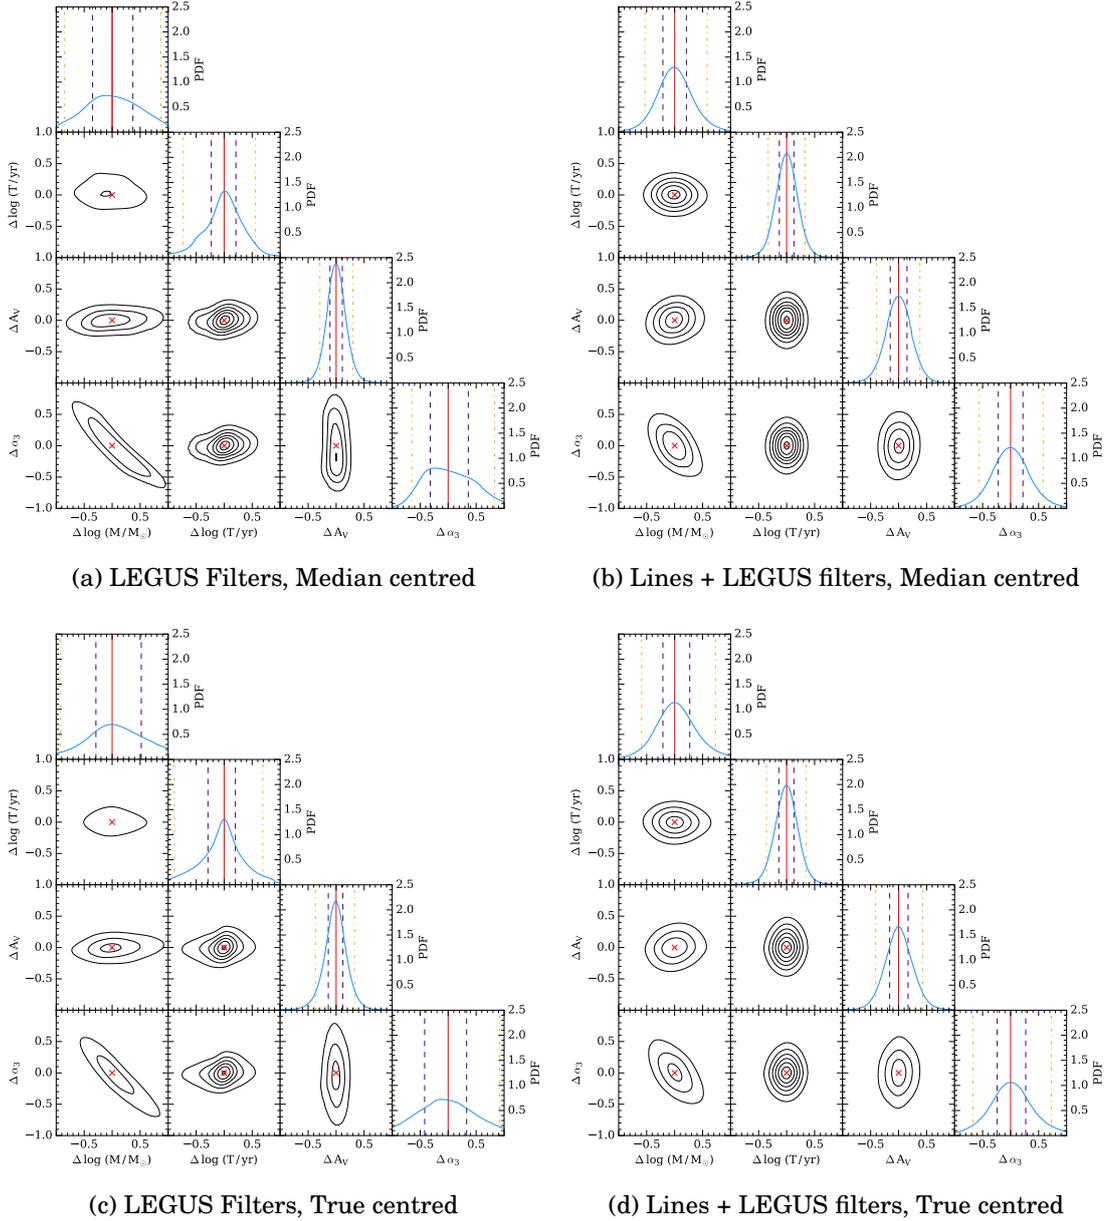
Figure 4.6: Stacks of the one and two-dimensional PDFs for 500 clusters generated as mock observations, as described in Table 4.1. The stacked PDFs are normalised by the number of clusters being stacked, with contours spaced in steps of 0.5, beginning at 0.5. The red crosses and lines mark the zero points. Panels (a) and (c) are the results of analysis using the LEGUS filters only, and are median-centred and centred on the true values respectively. Likewise, panels (b) and (d) show the results of analysis using both the lines (CIII, CIVa, NVa, OVa, OVIb) and the LEGUS filters, once again median-centred and true-value-centred respectively. Note the significant improvement in the recovery of the cluster age when including the lines in the analysis. We also see improvement in the recovery of the cluster mass and IMF, although with a residual $\log M$-α_3 degeneracy in the joint posterior PDF for those parameters. The fifth and ninety-fifth percentiles are marked with orange dot-dashed lines, and the first and third quartiles are marked with indigo dashed lines. The change in the interpercentile ranges (for the true-centred cases, panels (c) and (d)) are given in Table 4.3.

124

(first to third) ranges, as listed in Table 4.3 for both the true-centred and median-centred cases. While the median-centred case (marked MC) is a useful companion to panels (a) and (b) of Figure 4.6, the true-centred rows (marked TC) illustrate the improvement (or worsening) of our recovery of the parameters. In effect, we see a reduction in the overal scatter (as represented by the 5-95 range) across all parameters, save for the extinction, with improvements having a mean of $\sim 25\%$. For the interquartile range, we see a mean reduction in width of $\sim 23\%$, with a narrowing of $\sim 32\%$ for α_3. The broadening of the posterior PDFs for extinction is caused by the lack of extinction information in the rectified spectra we use for calculating the equivalent widths of our spectral features. This results in the almost flat posterior PDFs that we see in Figure 4.4. It is the addition of the lines, which are insensitive to the extinction, to the extinction sensitive photometry which causes the spreading out of our distributions here.

While we are unable to entirely break the mass-IMF degeneracy described by Ashworth et al. (2017) (curved contours in the joint $\log M$-α_3 PDF are still visible even with the equivalent widths included in the analysis), we are able to make significant improvements on constraining the value of not only α_3, but the mass and age of the clusters as well. We also note that some residual scatter is intrinsic due to the stochasticity in our models, and cannot be reduced by adding more clusters. Indeed, during testing we found no significant difference in scatter between stacks of 10, 100, or 500 clusters for a fixed Kroupa IMF case. Our analysis has therefore shown, when introducing additional information in the form of UV spectral lines, that it is possible to tighten the constraints on α_3, and therefore possible to rule out extreme variations in α_3 as well.

4.4.1 Observational constraints

To aid the planning of new observations, in this section we aim to predict what spectral features are required to gain sufficient improvement in recovery of the IMF when compared to the case of using pure broad-band photometry. What is of interest is specifically what information is needed to get the improved recovery of α_3 that we desire. To this end, we take the empirical approach of adding in features by their overall strength (equivalent width) until there is a point of diminishing returns, where the addition of further spectral features does not improve the posterior PDFs

by an appreciable amount.

The most basic method by which we can do this is to limit which spectral features we include in our Bayesian analysis by imposing limits on the equivalent width to which observations may be sensitive. This is a simple approach which allows us, to first approximation, to gauge the sensitivity required by real observations, and thus the wavelength coverage, resolution, and exposure time needed to constrain α_3 in a given sample of clusters.

To this end we select 100 example clusters from the 'Mock Obs.' model set, and repeat the Bayesian analysis presented in Section 4.4. However, instead of using a fixed selection of filters for the entire dataset, we limit the inclusion of spectral features to those which have an equivalent width that passes a particular threshold for each cluster individually. We choose limits of $12, 9, 6, 3$Å, along with the case where we have all the line features included, and the case where we use broadband photometry alone. In the case of the 9Å limit, this gives us between zero and three line features across the input clusters.

In panel (a) of Figure 4.7, we see the residuals of the 1D posterior PDFs for the four physical parameters for this set of input clusters. The markers represent the median residual value for each equivalent width limit, with the error bars covering the interquartile range. The markers in red represent the residuals calculated from the difference between the median of the posterior PDFs and the true values, whereas the purple points represent the residuals calculated using the difference between the mode of the posterior PDFs and the true values.

We note that even at a sensitivity limit of 9Å, we see definite improvement in age, mass, and IMF determination. As sensitivity increases, we see a slow improvement in recovery for these parameters, although the median residuals are close to zero across the range. However, the residuals for the extinction increase. This is caused by the line features having no extinction applied to them, resulting in an almost flat PDF for the A_V parameter as seen in Figure 4.4. As more and more spectral features are included, the posterior PDF for extinction is spread out from the narrow peak we obtain when using filters alone, as more and more noisy PDFs are combined with it. Over this sample of 100 clusters, we find that even with a limit of 9Å we are able to significantly improve the accuracy at which we are able to infer the value of α_3, and indeed that further improvements in sensitivity provide limited improvement over this. In both age and cluster mass, however, further

increases in sensitivity continue to visibly improve results.

The improvement in our ability to constrain the IMF is more visible in panel (b) of Figure 4.7. In this figure we see the evolution of the absolute value of the median of the residuals, as well as the evolution of the interquartile range, as we increase the sensitivity. Looking at the interquartile ranges, we see that our overall ability to constrain the IMF improves significantly with the introduction of equivalent width data, with improvement in α_3 being $\sim 32\%$ (median residual) or $\sim 37\%$ (mode residual) at the 9Å limit. However, further improvements gained by increasing our sensitivity are small, finally coming to $\sim 39\%$ (median residual) or $\sim 61\%$ (mode residual) when we include all the spectral features in the analysis. We see similar behaviour in mass, whereas the improvement in age is more consistent across the range of equivalent width limits. Our constraints on extinction loosen, for the reason explained earlier.

4.5 Summary and conclusions

In this paper we have developed an extension to the SPS code SLUG, whereby the resolution of the output spectra has been greatly increased in the UV region (from $\simeq 10$Å to $\simeq 0.4$Å). Also, a new equivalent width calculation mechanism has been implemented to take advantage of these new higher resolution spectra. These new capabilities were then leveraged to perform a theoretical study of the high mass end of the IMF and its relationship to observable quantities for unresolved stellar populations.

Having performed a qualitative exploration of how the equivalent widths of spectral features are affected by changes in the various physical parameters ($\log M$, $\log t$, A_V, α_3) that describe our model star clusters, we applied the Bayesian analysis capabilities of BAYESPHOT to mock observations generated with SLUG. These observations (both broad-band photometry and the equivalent widths of UV line features) were analysed using a large library of SLUG models which covered a wide parameter space. We found similar difficulty in tightly recovering the value of α_3 as in Ashworth et al. (2017) where degeneracy between the cluster mass and the IMF slope hampered this recovery. Repeating the analysis while including the equivalent widths of five representative spectral features resulted in a significant improvement in the recovery of the physical parameters for the clusters when aver-
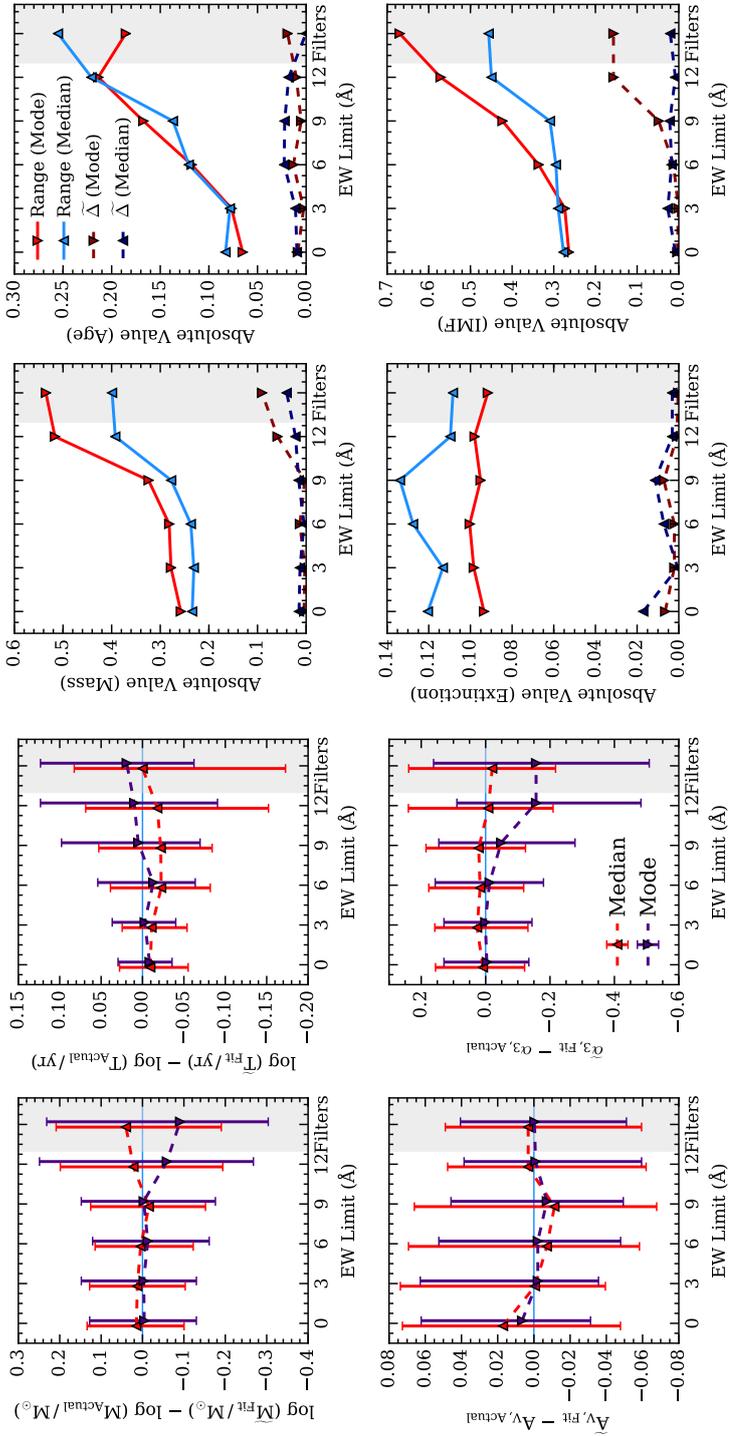
Figure 4.7: The plots in panel (a) show the residuals of the 1D posterior PDFs for the parameters $\log M$, $\log t$, A_v, and α_3 for the first 100 mock observation clusters. The residuals are calculated from both the medians of the distributions (red upward-pointing markers) and the modes of the distributions (purple downward-pointing markers), when compared to the true values for each of these parameters. For each cluster we have set a value of equivalent width below which no spectral features are included in the analysis. The markers correspond to the median value of the residual at each equivalent width limit, and the error bars cover the interquartile range. The markers are offset by ± 0.2 around their limit value for visual clarity. Panel (b) illustrates the behaviour of both the interquartile range, and the absolute value of the median residual for both the case where it is calculated from the modes of the PDFs (reds, downward triangles) and the medians (blues, upward triangles). The markers in panel (b) are not offset, as there are no errorbars to overlap.

128

aged over the whole set of 500 mock observations, with the interquartile range of the IMF posterior PDF shrinking by 32%. However, we still see evidence of the mass-IMF degeneracy noted in previous work.

Finally, we provided predictions for the required sensitivity in equivalent width needed to make improvements in recovery of α_3 over using broad-band photometry alone. We find that a modest limit of $9\,\text{Å}$ provides sufficient improvement (with the interquartile range of the median-residuals shrinking by 32% for the IMF), with further improvement as more line features are included in the analysis.

In conclusion, we find that the inclusion of UV spectral features does improve recovery of the physical parameters of unresolved stellar populations in our theoretical study, over that possible with broad-band photometry alone. However, we do not find that it makes it possible to fully break the mass-IMF degeneracy that exists in these systems. Having applied cuts to the included line features, we find observations sensitive to a modest equivalent width are capable of improving the recovery of the IMF high-mass slope α_3 to a significant degree, with a scatter narrow enough to either rule out or indeed distinguish extreme variations in the IMF.

Acknowledgements

This research has made use of the NASA/IPAC Extragalactic Database (NED), which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

For access to the data and codes used in this work, please contact the authors or visit www.slugsps.com.

4.6 Additional analysis: Narrowing the age coverage

As noted in Section 4.3, our studies cover a wider range in ages than would be generally visible to observations, as stars younger than a few Myr are likely to be embedded. In addition, it is possible that the assumption that the stars in the stellar population are coeval may break down at ages as young as 10^5 yr. For example, Kudryavtseva et al. (2012) report age spreads in young star clusters which are atleast as large as this. As the study is purely theoretical, this wide age range was instructive in understanding the effect of the IMF on the output spectrum at its most extreme. However, as observations cannot generally probe to such young ages as 10^5 yr, and are instead limited to a few Myr as a minimum age limit, we now repeat our analysis over a more constrained range of ages of $10^{6-7}\,yr$. The parameters that define the clusters used to generate the new library and mock observations are given in Table 4.4.

As before, we generate 500 sets of mock observations with both photometry (across the same LEGUS filterset) and equivalent widths (for the lines listed in Table 4.2). In Figure 4.8, we see very similar behaviour to that seen in Figure 4.6, with the changes in the interquartile and interpercentile (5-95) ranges given in Table 4.5. The relative improvement in recovery of α_3 when restricting ourselves to clusters in the $1-10\,\mathrm{Myr}$ range is smaller than we saw in Table 4.3, with the stacked PDFs in the true-centred case narrowing their interquartile ranges by 25% compared to 32% when including the younger clusters. In addition, apart from in age (where we have effectively applied a more narrow prior), the PDFs are slightly broader than we saw in Figure 4.6. Even so, we still see that for clusters in the age range likely to be seen by observations, we are able to produce an improvement in IMF recovery of 25% when using the representative set of 5 lines.

Name	$\log M$ ($M_\odot$)	$\log t$ (yr)	#[a]	α_3
Library	2.7–5.7 (C)	5.7–7.7 (C)	5×10^7	$-3.2 - -1.3$ (C)
Mock Observations	3–5 (C)	6–7 (C)	5×10^2	$-3.0 - -1.5$ (C)

[a] Number of realisations

Table 4.4: Table of the parameters used in SLUG for generating the higher-density library and mock observations for the model clusters with a narrower range of ages. The base IMF used for all these model sets is the 3 component broken power law of the Kroupa IMF. The break points between the segments of this power law are situated at 0.01, 0.08, 0.5, $M_{*,\mathrm{Max}}$. Z is the metallicity. The (D) represents discrete sampling, with # realisations at each value, and (C) represents continuous sampling, with # realisations drawn from a flat distribution across the given range. All simulations were run with solar metallicity, and an IMF cutoff of $M_{*,\mathrm{Max}} = 120 M_\odot$

Range (Percentiles)	$\Delta \log M$ ($M_\odot$)	$\Delta \log t$ (yr)	ΔA_V	$\Delta \alpha_3$
5-95 (MC)	-0.45 (-26%)	-0.33 (-38%)	$+0.17$ $(+26\%)$	-0.25 (-16%)
25-75 (MC)	-0.22 (-31%)	-0.08 (-29%)	$+0.06$ $(+24\%)$	-0.20 (-27%)
5-95 (TC)	-0.63 (-29%)	-0.56 (-50%)	$+0.15$ $(+19\%)$	-0.36 (-18%)
25-75 (TC)	-0.25 (-31%)	-0.09 (-29%)	$+0.05$ $(+17\%)$	-0.20 (-25%)

Table 4.5: Table of the difference in inter-quartile and inter-percentile ranges for the median-centred (MC) and true-centred stacks (TC) shown in Figure 4.8, changing from the case with photometry alone to the case including both photometry and the selected UV features. The numbers in parentheses correspond to the percentage change in the ranges, compared to the original value. The models included in this analysis are limited to ages of 10^{6-7} yr, as shown in Table 4.4.

In Figure 4.9 we repeat the sensitivity tests shown previously in Figure 4.7, this time using the models covering a narrower range of ages. The results are similar to what was seen in Figure 4.7, with similar improvements in mass and α_3, and improved recovery of the cluster age. Compared to the interquartile range of the median residuals on α_3 shrinking by 32%, with the restricted range of ages we see the same level of improvement. This improves to 61% when we include all the spectral information.
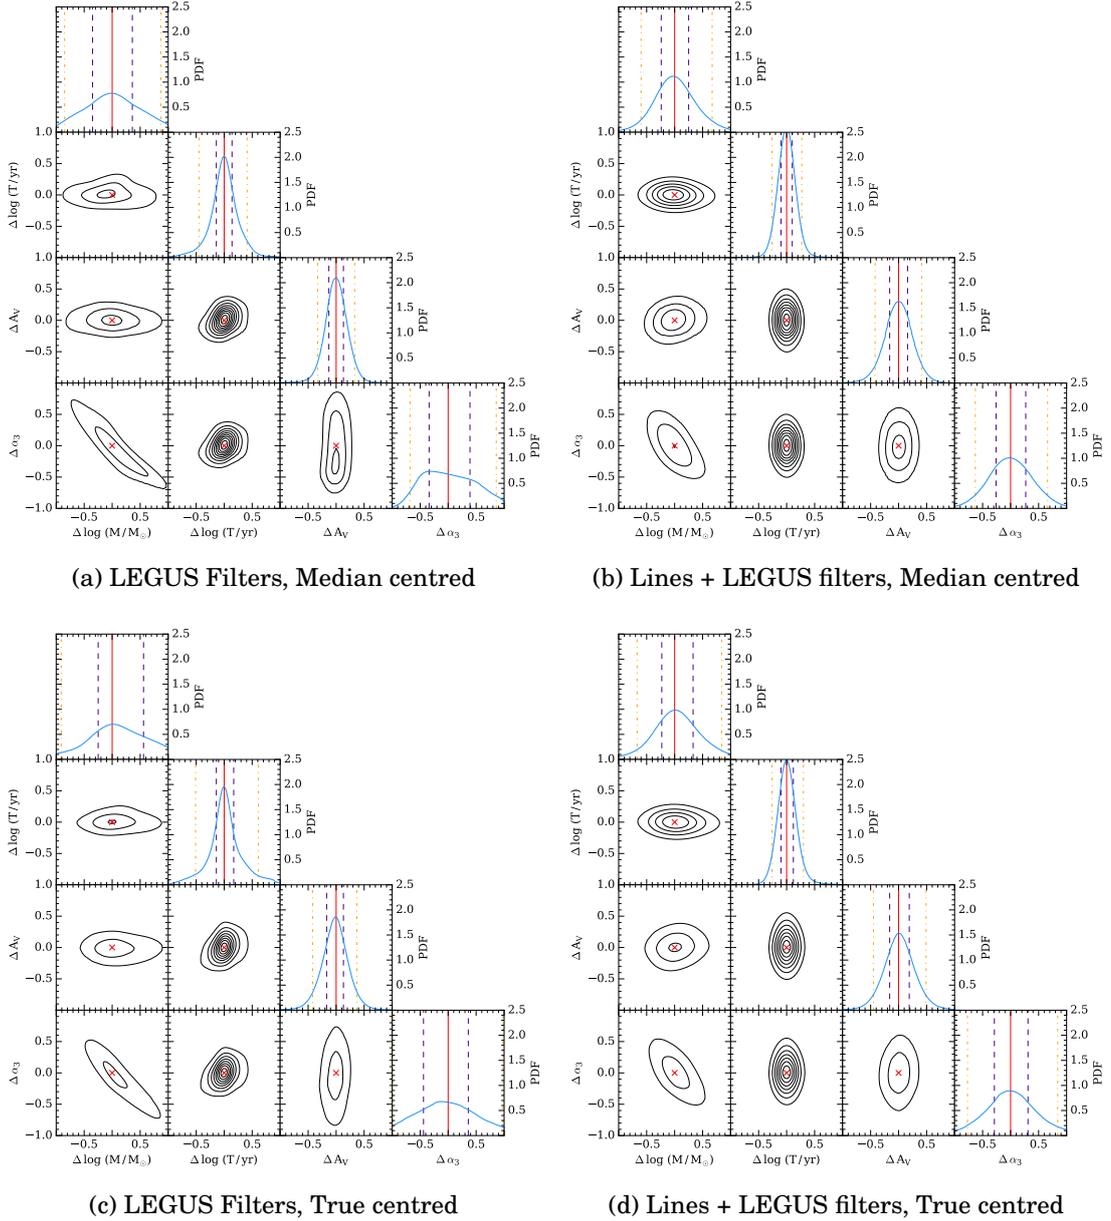
Figure 4.8: Stacks of the one and two-dimensional PDFs for 500 clusters generated as mock observations, as described in Table 4.4. The stacked PDFs are normalised by the number of clusters being stacked, with contours spaced in steps of 0.5, beginning at 0.5. The red crosses and lines mark the zero points. Panels (a) and (c) are the results of analysis using the LEGUS filters only, and are median-centred and centred on the true values respectively. Likewise, panels (b) and (d) show the results of analysis using both the lines (CIII, CIVa, NVa, OVa, OVIb) and the LEGUS filters, once again median-centred and true-value-centred respectively. Note the significant improvement in the recovery of the cluster age when including the lines in the analysis. We also see improvement in the recovery of the cluster mass and IMF, although with a residual $\log M$-α_3 degeneracy in the joint posterior PDF for those parameters. The fifth and ninety-fifth percentiles are marked with orange dot-dashed lines, and the first and third quartiles are marked with indigo dashed lines. The change in the interpercentile ranges (for the true-centred cases, panels (c) and (d)) are given in Table 4.5.

132

Figure 4.9: The plots in panel (a) show the residuals of the 1D posterior PDFs for the parameters $\log M$, $\log t$, A_v, and α_3 for the first 100 narrow age range mock observation clusters. The residuals are calculated from both the medians of the distributions (red upward-pointing markers) and the modes of the distributions (purple downward-pointing markers), when compared to the true values for each of these parameters. For each cluster we have set a value of equivalent width below which no spectral features are included in the analysis. The markers correspond to the median value of the residual at each equivalent width limit, and the error bars cover the interquartile range. The markers are offset by ± 0.2 around their limit value for visual clarity. Panel (b) illustrates the behaviour of both the interquartile range, and the absolute value of the median residual for both the case where it is calculated from the modes of the PDFs (reds, downward triangles) and the medians (blues, upward triangles). The markers in panel (b) are not offset, as there are no errorbars to overlap.

133

CHAPTER 5

DISCUSSION & CONCLUSIONS

5.1 Summary

During this work, variability in the upper-end slope of the IMF has been studied using the stochastic SPS code SLUG, along with its Bayesian analysis tools. What follows is a brief summary of each chapter, beginning with the code development detailed in Chapter 2.

5.1.1 Code Development

In order to undertake the studies that form this thesis, the SLUG code had to be extended to perform tasks it was not previously capable of. These extensions were described in Chapter 2. The main aim of this work was to study variation of the upper end of the IMF, and up until the beginning of this study, SLUG was unable to continuously vary the parameters that describe the IMF. To this end, a variable PDF module has been developed for the SLUG code which allows arbitrary parameters that define the IMF to be drawn randomly from other probability distributions. This new capability allows a library of SLUG simulations to be generated with continuously varying IMF parameters. This is required if one wishes to use SLUG's Bayesian inference tools to constrain, for example, the upper end slope of the IMF. The code performed well in testing and was ready for use in the analysis of LEGUS broad-band photometry, as detailed in Chapter 3.

Following the difficulty in accurately constraining the slope of the IMF that was evident during that study, additional sources of information were sought to improve our constraints on the IMF slope. Including UV spectral features in the analysis was put forward as the next step, as that area of the spectrum has many contributions from the massive stars whose number is affected by the slope of the upper-end of the IMF. However, SLUG's original $\simeq 10\,\text{Å}$ resolution in the UV region was too low to perform equivalent width studies. To this end, the high resolution stellar atmospheres of Leitherer et al. (2010) were chosen to be added to SLUG's included atmospheres, which cover massive stars in the UV. A new high resolution spectral synthesiser was implemented to take advantage of these new atmospheres, and new machinery was implemented to calculate and store equivalent widths of arbitrary integration regions allowing any spectral feature to be included in analysis. In addition to this, an interpolator-based atmosphere selector was developed as an alternative to the nearest neighbour method used by STARBURST99 for the same atmosphere grid. This interpolator had comparable or better performance to the nearest-neighbour method, when recovering the equivalent width for a test set of line features for known stellar atmospheres, producing smaller residuals more often than the nearest-neighbour method, but with its median residual being larger. However, as the test is an excercise in creating data where there is none, this comparable performance is acceptable. This new UV feature capability was then applied to a theoretical study of IMF variation using a number of grids of mock observations generated with SLUG. This study was detailed in Chapter 4. In addition, we explored the effects of including all stars in the normalised spectrum we use to calculate the equivalent widths, finding that it had a minimal effect across most of the lines we consider. However, due to the presence of, for example, age-defining features from Wolf-Rayet stars, it would be useful to use this full spectrum in combination with a more comprehensive line list. This would assist in the recovery of cluster properties (including the IMF), even with the approximate normalisation performed on the full-star spectrum rather than the true normalisation we have for the O and B stars. However, due to the theoretical nature of the study presented in Chapter 4, the use of O and B star features to validate the new high resolution atmospheres is self consistent, and still provides useful results.

5.1.2 Exploring IMF variability using broad-band photometry

In Chapter 3, we detailed a study of the IMF of unresolved stellar populations. Specifically, multiwavelength broad-band photometry from the LEGUS project (Calzetti et al., 2015) was analysed using the Bayesian inferrence tools included in the SLUG software package. *Hubble Space Telescope* photometry across five filters (WFC3 F275W, F336W, and F555W, ACS F435W and F814W) was used, along with synthetic WFC3 F657N measurements constructed using Hα observations from the Hα-LEGUS project. The chosen cluster catalogue was that of NGC 628 (Adamo et al., 2017), which was then trimmed such that the remaining clusters were also observed in Hα, with morphology that was indicitive of Hα emission from HII regions.

The first step was to replicate the analysis performed in Krumholz et al. (2015b), where a large library of 10^7 SLUG models was used in conjunction with SLUG's Bayesian analysis framework in order to constrain the mass, age, and extinction of the unresolved stellar populations in three catalogues, of which our NGC 628 catalogue was one. In that study, the authors found that many of the posterior probability distributions were multi-peaked, and were sometimes quite broad. Before our analysis could take place, as this time we wished to include the slope of the IMF as a constrainable parameter, we required a different library of SLUG models than that available from the previous study. To this end, we generated a library of 2×10^8 new SLUG models, with the IMF high-mass slope α_3 varied continuously across the library.

On using this new library for our analysis, we found much the same behaviour across the cluster catalogue as was found in the previous study by Krumholz et al. (2015b), with many of the clusters exhibiting broad or multi-peaked posterior PDFs. In addition, we observed that the posterior PDFs of the IMF slope were extremely broad, to the point where it was diffiult to discern the slope of the IMF in most cases. A degeneracy between the cluster mass and the IMF slope was also revealed and quantified.

In an attempt to break this degeneracy, Hα photometry in the form of a synthetic WFC3 F657N filter measurement was introduced. As Hα emission is an indicator of the the presence of massive stars, it is therefore linked to the high-mass slope of

the IMF. However, although there was some improvement in the age determination, the mass-IMF degeneracy remained unbroken.

To quantify the degeneracy between the IMF and the cluster mass, we next used a library of mock clusters generated with SLUG to create sets of mock photometry, including the same *Hubble Space Telescope* filters. These were then analysed in the same manner as the real LEGUS photometry. Once again, a degeneracy between the mass and the IMF emerged, although the age and extinction were both well constrained. Tightening the priors on the cluster mass improved recovery of the IMF slope dramatically, with $\Delta\alpha_3 \simeq 0.2$ when the cluster mass was known to an accuracy of 10%.

We concluded that further information was required to enable accurate recovery of the IMF slope when using SLUG, with one potential choice being photometry-independent measurements of the cluster mass.

5.1.3 Including UV line features in IMF studies

Having encountered the mass-IMF degeneracy described in Chapter 3, and the associated difficulties in constraining the IMF using broad-band photometry, we next looked to the UV region of the spectrum. As massive O and B type stars contribute significantly to the UV region of the spectrum, and the number of these stars present in a stellar population is highly dependent on the slope of the IMF, we included spectral features from this spectral region in our Bayesian analysis to improve IMF parameter recovery. However, SLUG was previously incapable of producing spectra in this wavelength range with sufficiently high resolution to enable equivalent width studies of line features. To this end, we implemented the high-resolution spectral libraries described by Leitherer et al. (2010), with the development of this new synthesiser described in detail in Chapter 2 and summarised above. In addition to a new spectral synthesiser, an equivalent width calculator was developed to allow the measurement of the strength of arbitrary spectral features in SLUG simulations.

With these new capabilities in hand, the qualitative behaviour of SLUG model spectra was explored as the cluster-defining parameters were changed, finding that many UV features were affected by changing the IMF slope. In addition, we noted some degeneracies, for example, between changing metallicity and changing IMF

slope in the case of the trough of the Nv feature. This qualitative study was then also used to pick five representative lines which were both effective in discerning the IMF but also covered a range of strengths and behaviours.

Following this qualitative study, 500 SLUG cluster models were simulated with random physical properties (with masses, ages, extinctions, and IMF slopes generated from a range spanning a slightly smaller space than that covered by the library). Their accompanying observable properties were calculated as well, (both broad band photometry across a LEGUS filter set and equivalent widths of many UV features). These model clusters would form a set of mock observations that were then analysed in much the same way as the LEGUS photometry studied in Chapter 3. However, with the new UV spectra capabilities, these mock observations now included equivalent widths for a selection of UV features in addition to broad-band photometry across a LEGUS filter set.

As in Chapter 3, a large library of SLUG models was required in order to perform the Bayesian inference. However, the library produced for the Chapter 3 study was unsuitable for use as it did not contain equivalent width values, which we wish to include in our analysis. A new, tighter library was generated over a parameter space slightly larger than that used to generate the mock observations, allowing us to leverage our knowledge of the properties of our "observed" clusters to produce a more compact libary while maintaining a high density of points.

Using this new library, the Bayesian inference technique was applied to the 500 mock observations. The analysis was performed using just the LEGUS filter set, a representative selection of five UV lines, a large set of indepedent lines, and finally the two line sets each in combination with the filter set to give a total of five sets of posterior PDFs. As in the previous study, a degeneracy between the cluster mass and the IMF slope was apparent, with very broad posterior PDFs present for both parameters across many of the test clusters. Cluster age and extinction were both well recovered however, with the latter more so. In addition, we see that the representative set of five lines (containing a mix of strong, weak, absorption, and emission features) was sufficient to greatly improve recovery of the cluster age and IMF when used alone (at the expense of the mass and extinction, to which the equivalent widths are insensitive), and across all parameters except the extinction when combined with broad-band photometric data. Including the full set of independent lines resulted in further (if minor) improvements in mass, age, and IMF

determination. Having stacked the posterior PDFs (including the joint PDFs), we were then able to measure the interquartile ranges and quantify the improvement (or otherwise) in parameter recovery. Also, using two different centring methods (centring the stacks on either the true value of each parameter for each cluster, or centring on the median value of each PDF) allowed us to explore both the shape of the PDFs themselves along with how close the inference was able to come to the correct answers. Using the set of five representative filters resulted in a 41% reduction in the interquartile range for the mass, 45% in the age, and most importantly 32% for the IMF slope α_3 when considering the true centred stacks. However, the extinction's stacked posterior PDF was broadened, with an increase of 27% in the interquartile range.

Finally, to provide a prediction of the constraining power of equivalent width observations at different sensitives, the first 100 mock observations were selected and reanalysed using the Bayesian inference method. However, unlike the previous study, this time the included line features were limited on a cluster-by-cluster basis to those whose strength exceeded a given equivalent width limit. It was found that even at a modest limit of 9Å it was possible to improve recovery of the IMF slope by 32%. Here this percentage refers to a narrowing of the interquartile range of the residuals between the median of the posteror PDF in IMF and the true value for each cluster. This was further improved to 39% when considering all independent line features.

In conclusion, it was determined that including equivalent width measurements of UV spectral features in Bayesian inference studies of unresolved stellar populations greatly improves our ability to constrain the value of α_3, even at modest sensitivities.

As the age range covered in this study extends to younger ages than would typically be observable, the analysis was repeated using mock clusters in the range 10^{6-7} yr. Using a new and more dense library and new mock clusters, similar behaviour was observed to the wide age study. Recovery of the IMF was similar, with the relative improvement obtained by including line features being 25% compared to 32% when including younger clusters. When restricting included line features by their equivalent width, we once again find similar relative improvements, with the interquartile range of the median residuals shrinking by the same amount (32%) at the 9Å limit as in the case where extremely young

clusters are included in the analysis.

5.2 Discussion of results

The main objective of the work presented in this thesis was to investigate whether it was possible to constrain the form of the IMF from observations using stellar population synthesis techniques. Due to the controversy surrounding the form of the IMF and whether it varies with location or time, rigorous statistical methods should be employed when determining the IMF from a given set of observations. Although Bayesian analysis of stochastic SPS models had been used to recover the physical properties of star clusters in previous studies (Krumholz et al., 2015b), this technique had not been used to constrain the IMF before. By including the slope of the high-mass end of the IMF in our analysis, we were able to further develop these techniques and quantify our ability to recover this slope along with any degeneracies between it and the other physical parameters. This work therefore provides a strong statistical foundation for studying the form of the IMF using observations.

The main conclusion that can be drawn from the study presented in Chapter 3 is that it is difficult to constrain the slope of the IMF for unresolved stellar populations using broad-band photometry alone. The emergence of a degeneracy between the cluster mass and the slope of the IMF was the cause of much of this difficulty. However, the posterior PDF for the IMF was broad enough that the posterior PDFs for the other physical parameters (cluster mass, age, and extinction) were comparable to those produced in the three-parameter study conducted by Krumholz et al. (2015b), with similar multi-modal features across the cluster catalogue.

In Chapter 3, the importance of prior selection noted by Krumholz et al. (2015b) was reinforced. Indeed, for certain multi-peaked posterior PDFs it was seen that the physically motivated priors shifted probability to the opposite peak when compared to flat priors, with the flat priors preferring larger values of α_3 and larger cluster masses. The choice of prior used with SLUG's Bayesian framework can therefore have strong effects on the recovered values of parameters (as we saw with either a top heavy or bottom heavy IMF being preferred depending on the chosen prior), and it is best to use a prior that is motivated by the physical conditions of the environment the data is taken from. While we use physically motivated priors in Chapter 3, in the study presented in Chapter 4 we use flat priors as the mock

observations are taken from clusters drawn from a flat distribution rather than a realistic galactic environment.

The degeneracy of cluster mass with IMF slope is somewhat intuitive, in that a cluster which is bright in a way that could be caused by the presence of many massive stars could contain many massive stars from either a top heavy IMF or from being massive enough to contain such a large number of massive stars just by virtue of its size. In addition, we must also consider the scales of the two parameters. The IMF slope parameter is an exponent, whereas the mass is not. The difficulty in breaking this degeneracy is evident from the study conducted in Chapter 3, where even the inclusion of Hα data that is sensitive to massive stars does not provide much in the way of improvement. The inclusion of equivalent width data in the UV in Chapter 4 does provide some improvement (and in fact dramatically improves recovery of the IMF slope) but the degeneracy is still present in the joint mass-IMF posterior PDF, albeit diminished. Only through an photometry-independent mass measurement of the target cluster could we provide narrow constraints on the IMF slope when using photometry alone.

One possible limitation of the work presented in Chapter 3 is the density of the simulation library. Although the returned PDFs are similar to the those provided by the analysis performed by Krumholz et al. (2015b), it is likely that the simulation library's density is less than ideal. Changing the bandwidth did however not noticably change the results. It would be interesting to study the ideal bandwidth for the simulation library, or computational resources willing, construct an even more dense library. This would be possible by limiting the parameter space covered by the library. The library used in our study was wide in coverage, to match the library used in the Krumholz et al. (2015b) study and to provide a general purpose variable IMF library. In our later study presented in Chapter 4, we choose a library which is much more narrow in coverage. However, the number of points included was also reduced due to the additional memory load of the equivalent width measurements stored for this library.

There has been some previous work on constraining the IMF using (rest frame) UV spectral features (Pettini et al., 2000; Steidel et al., 2004; Quider et al., 2009). However, unlike the work presented in Chapter 4, these studies concern entire galaxies rather than working on a cluster-by-cluster basis. In addition, they assume a continuous star formation rate for the entire galaxy rather than the formation of

star clusters throughout the age of the galaxy. Thirdly, they perform no rigorous statistical analysis of the IMF, instead fitting the CIV and NV lines by eye to model spectra produced by STARBURST99, at two or three slopes and high-mass cutoffs of a Salpeter-like IMF.

The core objective of this work was to discover if it were possible to constrain variations in the high-mass region of the IMF through a combination of observations and stochastic SPS models, and we have gone some way to answering that question. Indeed, although we were unable to closely recover the slope of the IMF through the use of broad-band photometry alone, we were able to quantify a degeneracy between the IMF and the cluster mass. Exploring this degeneracy allowed us to show significant improvement in IMF recovery when independent mass measurements are used as a prior. Then, through the inclusion of UV spectral lines, we were able to significantly improve our constraining power over the IMF slope over using broad-band photometry alone. Although this latter work was a purely theoretical study, it provides the groudwork for future development of this novel technique, and it would be possible to apply this technique to future *Hubble Space Telescope* spectroscopy in tandem with photometry. The statistical techniques presented address some of the concerns presented in Bastian et al. (2010) of a lack of rigour in some IMF analysis, and provide a more complete picture than studies performed using only deterministic stellar population models (especially in the case of lower mass clusters).

5.3   Future work

The SLUG software package has seen significant development over the course of this work, with new features to allow for variation of probability distributions (specifically implemented for the IMF), a new high-resolution UV spectrum capability, and equivalent width calculation tools. However, further improvements can be made. The way the variable IMF capability is implemented (see Chapter 2) is very general. It is therefore possible to extend this variation to other probability distributions used by SLUG, such as the cluster mass function or star formation history. For example, the parameters of the CMF for a given galaxy could be estimated using Bayesian inference if a library with varying CMF parameters could be generated. There has been recent work within the LEGUS collaboration regarding the form of

the CMF. Adamo et al. (2017) report that, much like the IMF, there is much discussion on the shape and universality of the CMF. In their study, they demonstrate that a steeper power law, or a Schechter function, is a better fit for the young star cluster mass function in NGC 628 than the "standard" power law index of −2. In a review of the lifecycle of star clusters, Adamo & Bastian (2015) claim that, although many studies have presented a consistent −2 power law slope for the CMF, these studies do not all consider the same ranges of masses. In addition, they report that several recent studies have shown a turnoff in the CMF which can be reproduced with a Schechter function, as in the LEGUS study. Messa et al. (2018) find a similar steepening at the high-mass end of the CMF in M 51. Adamo & Bastian (2015) argue that the characteristic mass of this Schechter function appears to vary with galactic environment. This variation of parameters provides a timely and appropriate application for a future variable CMF capability in SLUG.

A further extension would be to extend the high-resolution spectral capability from the UV up into the optical region of the spectrum, improving on the current resolution of $\simeq 20\,\text{Å}$. Using high-resolution atmospheres distributed with STARBURST99, it would be possible to create a new spectral synthesiser, either outputting a spectrum with a singular high resolution component in the optical region, or a spectrum with both UV and optical components at high resolution. Obtaining high resolution atmospheres in any region (for example, the UV) for both the line features and continua of the stars covered by the atmosphere library would allow the inclusion of all stellar types in the normalised spectrum (and hence the calculated equivalent widths) output by SLUG. At present, only features from the massive O and B type stars covered by the Leitherer et al. (2010) atmospheres are included in equivalent width calculations (as discussed in Chapter 4 and previously in this chapter). Including contributions from all star types (for example, the features present in the atmospheres of Wolf-Rayet stars) would provide a more complete representation of the spectrum of a given cluster, and a more complete treatment of the equivalent widths. The additional features present in the atmospheres of stars of certain evolutionary phases would provide additional constraining power on the cluster age, and potentially the IMF as well. We see the first inklings of this in the preliminary tests shown at the end of Chapter 2, where many intense UV features are visible at, for example, $10^{6.5}$ yr are clearly visible when including the full range of stars in the normalised spectrum

used for equivalent width calculations.

Extending the high resolution capability of SLUG into the optical regime can provide additional tools for tackling the problem of IMF variation. For example, Hoversten & Glazebrook (2008) (and more recently Gunawardhana et al. (2011)) use a combination of Hα equivalent widths and $g - r$ optical colour to demonstrate a systematic variation of the IMF.

The infrared would also provide us with additional information. For example, Gazak et al. (2013) show how J-H colour can be used to quantify the age of a cluster. This method is shown to be able to determine whether a young, massive star cluster is older or younger than $5.7 \pm 0.8 \, \text{Myr}$, information which could then be used to set priors on the ages of clusters to be analysed using SLUG.

It would also be possible to expand SLUG's inference capabilities into the metallicity domain through the use of continuously varying metallicities across a set of SLUG model realisations. At present this is not possible as SLUG does not contain a set of atmosphere models that are sufficiently densely sampled in metallicity to allow interpolation between them. With these capabilities, we could not only infer the metallicity of a stellar population using the same Bayesian techniques as used for other physical properties, but also explore more fully the degeneracies hinted at in Chapter 4 between the IMF slope and the metallicity of the population.

In summary, for significant improvement in SLUG's capabilities, new atmosphere models models at high resolution across a wide area of the spectrum (UV, optical, IR), with continua and line spectra, densely sampled in metallicity, would be required.

The study conducted in Chapter 4 was purely theoretical, and as such provides an idealised situation for recovering the IMF from spectra. Although, due to SLUG's stochastic nature, it is not simply a trivial problem of finding the exact match in the library, it would be instructive to attempt to recover the IMF for clusters with an independently calculated α_3 value. As discussed in the conclusions of Chapter 4, the real-world performance improvements of the method presented in that study will depend on how well the UV features of the spectrum are modelled. High resolution UV spectral observations of a full resolved stellar population, the central cluster of 30 Doradus (Crowther et al., 2016) for example, would provide suitable input equivalent widths for Bayesian analysis with SLUG, whilst also allowing an independent calculation of the IMF through star counts.

From a more completionist viewpoint, it would also be possible to run all the LEGUS cluster catalogues through SLUG's Bayesian analysis with a variable IMF slope parameter. Although it is likely that results would be much the same as seen for standard fixed-IMF analysis of the LEGUS catalogue (as was the case for most of the clusters in NGC 628), it would be interesting to study the IMF across such a wide selection of different galaxies.

A limitation of the models generated during this work is the absence of binary systems in the star clusters we build. Binary star systems are thought to make up a large proportion of the star systems present in galaxies, with perhaps all O and B stars residing in binary systems, as summarised by Eldridge et al. (2017). Indeed, the presence of binaries affects the evolution of a given stellar system due to events such as mergers and mass transfer between the two stars, and their presence results in a given population of stars appearing more blue than would be predicted by models that do not contain binaries (Eldridge, 2012). The same study also notes that due to the increase in ionising radiation from these binaries, that measurements such as Hα will be affected. They also note that X-ray binaries would provide a very large increase in ionising flux indeed, having stronger effects on both the Hα and UV flux ratios from a stellar population, with the effect particularly large on smaller stellar populations. The effect of mergers and mass transfer would be to push the IMF to a shallower slope, as mass transfer may cause stars to increase in mass with time. Capturing this effect as SLUG evolves a stellar population would therefore assist in constraining the IMF. Continuously varying the binary fraction across a simulation library would also allow this to form an additional parameter when performing a Bayesian inference analysis on a given observation. The implementation of binaries in the SLUG code would be an instructive exercise, and it would be interesting to quantify the effects of including these multi-star systems on the studies presented in this work.

Another limitation is the absence of stellar rotation in our models. This is due to the use of the Padova tacks with thermally pulsating AGB stars which do not include stellar rotation. However, as SLUG is distributed with the Geneva stellar tracks (Ekström et al., 2012), which do include the effects of rotation, it would be possible to study what effects (if any) including rotation has on the IMF slopes that we recover. One must also consider the limits of the IMF we chose. We have already shown how there can be significant differences between clusters generated

with a $30\,\mathrm{M}_\odot$ mass cutoff compared to the canonical $120\,\mathrm{M}_\odot$ for example. There are some observational indications for stars with masses greater than this cutoff (For example, the recent work of Schneider et al., 2018, who find evidence for stars with masses up to $200\,\mathrm{M}_\odot$ in 30 Doradus). However, as SLUG does not have tracks or atmospheres which extend outside the parameter space we explore, it would not be possible to investigate higher mass cutoffs at present. If more comprehensive tracks and atmospheres become available, then this would be a possible avenue for development.

5.4   Concluding remarks

Overall, this study has presented novel methods whereby the IMF of unresolved stellar populations may be better understood. Through the use of Bayesian inference techniques, and a combination of broad-band photometry and UV equivalent width measurements of modest sensitivity, a noticeable improvement in constraining power can be achieved over the use of broad-band photometry alone when using this new framework. With further refinement of the methods used in the theoretical study presented here, these techniques will allow future work to more closely constrain the high-mass end of the IMF using broad-band photometry and UV spectra. If further high resolution atmosphere libraries were to be included, this novel approach could be extended through to other regions of the spectrum as well.

BIBLIOGRAPHY

Adamo A., Bastian N., 2015, preprint, p. arXiv:1511.08212 (arXiv:1511.08212)

Adamo A., et al., 2017, ApJ, 841, 131

Adams F. C., Fatuzzo M., 1996, ApJ, 464, 256

Anders P., Kotulla R., de Grijs R., Wicker J., 2013, ApJ, 778, 138

Andrews J. E., et al., 2013, ApJ, 767, 51

Andrews J. E., et al., 2014, ApJ, 793, 4

Ashworth G., et al., 2017, MNRAS, 469, 2464

Auger M. W., Treu T., Bolton A. S., Gavazzi R., Koopmans L. V. E., Marshall P. J., Moustakas L. A., Burles S., 2010, ApJ, 724, 511

Baade W., 1944, ApJ, 100, 137

Barbaro C., Bertelli C., 1977, A&A, 54, 243

Bastian N., Covey K. R., Meyer M. R., 2010, ARA&A, 48, 339

Bell E. F., McIntosh D. H., Katz N., Weinberg M. D., 2003, ApJS, 149, 289

Bonnell I. A., Clarke C. J., Bate M. R., 2006, MNRAS, 368, 1296

Boselli A., Boissier S., Cortese L., Buat V., Hughes T. M., Gavazzi G., 2009, ApJ, 706, 1527

Bressert E., et al., 2012, A&A, 542, A49

Calzetti D., Chandar R., Lee J. C., Elmegreen B. G., Kennicutt R. C., Whitmore B., 2010, ApJ, 719, L158

Calzetti D., et al., 2015, AJ, 149, 51

Cappellari M., et al., 2012, Nature, 484, 485

Cappellari M., et al., 2013, MNRAS, 432, 1862

Cardelli J. A., Clayton G. C., Mathis J. S., 1989, ApJ, 345, 245

Cerviño M., 2013, "New~Astron.~Rev.", 57, 123

Cerviño M., Valls-Gabaud D., Luridiana V., Mas-Hesse J. M., 2002, A&A, 381, 51

Chabrier G., 2003, PASP, 115, 763

Chabrier G., 2005, in Corbelli E., Palla F., Zinnecker H., eds, Astrophysics and Space Science Library Vol. 327, The Initial Mass Function 50 Years Later. p. 41 (arXiv:astro-ph/0409465), doi:10.1007/978-1-4020-3407-7_5

Charlot S., Bruzual A. G., 1991, ApJ, 367, 126

Conroy C., van Dokkum P. G., 2012a, ApJ, 760, 71

Conroy C., van Dokkum P. G., 2012b, ApJ, 760, 71

Conroy C., Dutton A. A., Graves G. J., Mendel J. T., van Dokkum P. G., 2013, ApJ, 776, L26

Conroy C., van Dokkum P. G., Villaume A., 2017, ApJ, 837

Corbelli E., Verley S., Elmegreen B. G., Giovanardi C., 2009, A&A, 495, 479

Crowther P. A., et al., 2016, MNRAS, 458, 624

Dib S., Schmeja S., Hony S., 2017, MNRAS, 464, 1738

Dries M., Trager S. C., Koopmans L. V. E., 2016, MNRAS,

Ekström S., et al., 2012, A&A, 537, A146

Eldridge J. J., 2012, MNRAS, 422, 794

Eldridge J. J., Stanway E. R., 2009, MNRAS, 400, 1019

Eldridge J. J., Stanway E. R., Xiao L., McClelland L. A. S., Taylor G., Ng M., Greis S. M. L., Bray J. C., 2017, PASA, 34, e058

Elmegreen B. G., 2000, ApJ, 539, 342

Elmegreen B. G., 2002, ApJ, 577, 206

Fanelli M. N., O'Connell R. W., Burstein D., Wu C.-C., 1992, ApJS, 82, 197

Ferland G. J., et al., 2013, RMxAA, 49, 137

Fumagalli M., da Silva R. L., Krumholz M. R., 2011, ApJ, 741, L26

Gazak J. Z., Bastian N., Kudritzki R.-P., Adamo A., Davies B., Plez B., Urbaneja M. A., 2013, MNRAS, 430, L35

Gerhard O., 2000, in Lançon A., Boily C. M., eds, Astronomical Society of the Pacific Conference Series Vol. 211, Massive Stellar Clusters. p. 12 (arXiv:astro-ph/0007258)

Girardi L., Bressan A., Bertelli G., Chiosi C., 2000, A&AS, 141, 371

Gräfener G., Koesterke L., Hamann W.-R., 2002, A&A, 387, 244

Grasha K., et al., 2015, ApJ, 815, 93

Gunawardhana M. L. P., et al., 2011, MNRAS, 415, 1647

Gvaramadze V. V., Weidner C., Kroupa P., Pflamm-Altenburg J., 2012, MNRAS, 424, 3037

Haas M. R., Anders P., 2010, A&A, 512, A79

Hamann W.-R., Gräfener G., 2004, A&A, 427, 697

Hillier D. J., Miller D. L., 1998, ApJ, 496, 407

Ho L. C., Filippenko A. V., 1996, ApJ, 466, L83

Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, MNRAS, 445, 581

Hoversten E. A., Glazebrook K., 2008, ApJ, 675, 163

Kramida A., Yu. Ralchenko Reader J., and NIST ASD Team 2015, NIST Atomic Spectra Database (ver. 5.3), [Online]. Available: http://physics.nist.gov/asd [2017, May 16]. National Institute of Standards and Technology, Gaithersburg, MD.

Kroupa P., 2001, MNRAS, 322, 231

Kroupa P., Weidner C., 2003, ApJ, 598, 1076

Kroupa P., Tout C. A., Gilmore G., 1993, MNRAS, 262, 545

Kroupa P., Weidner C., Pflamm-Altenburg J., Thies I., Dabringhausen J., Marks M., Maschberger T., 2013, The Stellar and Sub-Stellar Initial Mass Function of Simple and Composite Populations. p. 115, doi:10.1007/978-94-007-5612-0_4

Krumholz M. R., 2011, in Treyer M., Wyder T., Neill J., Seibert M., Lee J., eds, Astronomical Society of the Pacific Conference Series Vol. 440, UP2010: Have Observations Revealed a Variable Upper End of the Initial Mass Function?. p. 91

Krumholz M. R., Fumagalli M., da Silva R. L., Rendahl T., Parra J., 2015a, MNRAS, 452, 1447

Krumholz M. R., et al., 2015b, ApJ, 812, 147

Kudryavtseva N., et al., 2012, ApJ, 750, L44

Kurucz R. L., 1992, in Barbuy B., Renzini A., eds, IAU Symposium Vol. 149, The Stellar Populations of Galaxies. p. 225

La Barbera F., Ferreras I., Vazdekis A., de la Rosa I. G., de Carvalho R. R., Trevisan M., Falcón-Barroso J., Ricciardelli E., 2013, MNRAS, 433, 3017

Lamb J. B., Oey M. S., Werk J. K., Ingleby L. D., 2010, ApJ, 725, 1886

Lamb J. B., Oey M. S., Segura-Cox D. M., Graus A. S., Kiminki D. C., Golden-Marx J. B., Parker J. W., 2016, ApJ, 817, 113

Larson R. B., 1998, MNRAS, 301, 569

Lee J. C., et al., 2009, ApJ, 706, 599

Lee J. C., Veilleux S., McDonald M., Hilbert B., 2016, ApJ, 817, 177

Leitherer C., et al., 1999, ApJS, 123, 3

Leitherer C., Leão J. R. S., Heckman T. M., Lennon D. J., Pettini M., Robert C., 2001, ApJ, 550, 724

Leitherer C., Ortiz Otálvaro P. A., Bresolin F., Kudritzki R.-P., Lo Faro B., Pauldrach A. W. A., Pettini M., Rix S. A., 2010, ApJS, 189, 309

Leitherer C., Ekström S., Meynet G., Schaerer D., Agienko K. B., Levesque E. M., 2014, ApJS, 212, 14

Lejeune T., Cuisinier F., Buser R., 1997, A&AS, 125, 229

Levesque E. M., Leitherer C., Ekstrom S., Meynet G., Schaerer D., 2012, ApJ, 751, 67

Li Z., Han Z., 2008, ApJ, 685, 225

Li H., et al., 2017, ApJ, 838

Lyubenova M., et al., 2016, MNRAS, 463, 3220

Maschberger T., 2013, MNRAS, 429, 1725

Massey P., Johnson K. E., Degioia-Eastwood K., 1995, ApJ, 454, 151

Matteucci F., 1994, A&A, 288, 57

Messa M., et al., 2018, MNRAS, 473, 996

Meurer G. R., et al., 2009, ApJ, 695, 765

Miller G. E., Scalo J. M., 1979, ApJS, 41, 513

Oey M. S., Lamb J. B., Kushner C. T., Pellegrini E. W., Graus A. S., 2013, ApJ, 768, 66

Paalvast M., Brinchmann J., 2017, MNRAS, 470, 1612

Pauldrach A. W. A., Lennon M., Hoffmann T. L., Sellmaier F., Kudritzki R.-P., Puls J., 1998, in Howarth I., ed., Astronomical Society of the Pacific Conference Series Vol. 131, Properties of Hot Luminous Stars. p. 258

Pettini M., Steidel C. C., Adelberger K. L., Dickinson M., Giavalisco M., 2000, ApJ, 528, 96

Popescu B., Hanson M. M., 2009, AJ, 138, 1724

Quider A. M., Pettini M., Shapley A. E., Steidel C. C., 2009, MNRAS, 398, 1263

Renzini A., 1981, Annales de Physique, 6, 87

Robert C., Leitherer C., Heckman T. M., 1993, ApJ, 418, 749

Robert C., Pellerin A., Aloisi A., Leitherer C., Hoopes C., Heckman T. M., 2003, The Astrophysical Journal Supplement Series, 144, 21

Rubin K. H. R., Diamond-Stanic A. M., Coil A. L., Crighton N. H. M., Moustakas J., 2018, ApJ, 853, 95

Salpeter E. E., 1955, ApJ, 121, 161

Sander A., Hamann W.-R., Todt H., 2012, A&A, 540, A144

Scalo J. M., 1986, Fund. Cosmic Phys., 11, 1

Schneider F. R. N., et al., 2018, Science, 359, 69

Smith R. J., 2014, MNRAS, 443, L69

Smith R. J., Lucey J. R., Carter D., 2012, MNRAS, 426, 2994

Sparre M., Hayward C. C., Feldmann R., Faucher-Giguère C.-A., Muratov A. L., Kereš D., Hopkins P. F., 2017, MNRAS, 466, 88

Spiniello C., Trager S. C., Koopmans L. V. E., Chen Y. P., 2012, ApJ, 753

Steidel C. C., Shapley A. E., Pettini M., Adelberger K. L., Erb D. K., Reddy N. A., Hunt M. P., 2004, ApJ, 604, 534

Stephens I. W., et al., 2017, ApJ, 834, 94

Thomas D., 1999, MNRAS, 306, 655

Vassiliadis E., Wood P. R., 1993, ApJ, 413, 641

Vázquez G. A., Leitherer C., 2005, ApJ, 621, 695

Vázquez G. A., Leitherer C., Schaerer D., Meynet G., Maeder A., 2007, ApJ, 663, 995

Walborn N. R., Parker J. W., Nichols J. S., 1995b, VizieR Online Data Catalog,

Walborn N. R., Nichols-Bohlin J., Panek R. J., 1995a, VizieR Online Data Catalog,

Weidner C., Kroupa P., 2005, ApJ, 625, 754

Weidner C., Kroupa P., 2006, MNRAS, 365, 1333

Weidner C., Pflamm-Altenburg J., Kroupa P., 2011, in Treyer M., Wyder T., Neill J., Seibert M., Lee J., eds, Astronomical Society of the Pacific Conference Series Vol. 440, UP2010: Have Observations Revealed a Variable Upper End of the Initial Mass Function?. p. 19 (arXiv:1011.1905)

Weisz D. R., et al., 2015, ApJ, 806, 198

Wofford A., Leitherer C., Chandar R., Bouret J.-C., 2014, ApJ, 781, 122

Zackrisson E., Rydberg C.-E., Schaerer D., Östlin G., Tuli M., 2011, ApJ, 740, 13

da Silva R. L., Fumagalli M., Krumholz M., 2012, ApJ, 745, 145

da Silva R. L., Fumagalli M., Krumholz M. R., 2014, MNRAS, 444, 3275

de Mello D. F., Leitherer C., Heckman T. M., 2000, ApJ, 530, 251

de Wit W. J., Testi L., Palla F., Vanzi L., Zinnecker H., 2004, A&A, 425, 937

de Wit W. J., Testi L., Palla F., Zinnecker H., 2005, A&A, 437, 247

van Dokkum P. G., Conroy C., 2010, Nature, 468, 940