

Durham E-Theses

Effects of Gamma-Ray Spectral Variability on Fundamental Physics Measurements

JAMIE GRAHAM

How to cite:

GRAHAM, JAMIE (2025) Effects of Gamma-Ray Spectral Variability on Fundamental Physics Measurements. Doctoral thesis, Durham University.

Use policy

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

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/16114/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

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

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

Effects of Gamma-Ray Spectral Variability on Fundamental Physics Measurements

Jamie Andrew Graham

A thesis presented for the degree of
Doctor of Philosophy



Department of Physics

Durham University

2019

Contents

1	Introduction	1
1.1	Active Galactic Nuclei	2
1.1.1	Overview	2
1.1.2	History	2
1.1.3	AGN Unification	5
1.1.4	AGN Classification	7
1.1.5	AGN Across the Spectrum	8
1.1.6	Observational Properties of AGN	11
1.2	Astronomical Gamma-Rays	12
1.2.1	Overview	12
1.2.2	Absorption Cross Sections	13
1.2.3	<i>Fermi</i> Large Area Telescope	14
1.2.4	<i>Fermi</i> -LAT Instrument Response	17
1.2.5	Types of IRF	19
1.2.6	Imaging Atmospheric Cherenkov Telescopes	21
2	Astrophysical Modelling	25
2.1	Modelling	25
2.1.1	Overview	25
2.1.2	Synchrotron Emission	26
2.1.3	Physical Processes	26

CONTENTS

2.1.4	Synchrotron Self-Compton Models	29
2.1.5	Hadronic vs Leptonic Emission	31
2.1.6	Phenomenological Models	32
2.2	Statistics	33
2.2.1	Overview	33
2.2.2	Astronomical Modelling	34
2.2.3	Multivariate Normal and χ^2 Distributions	34
2.2.4	Maximum Likelihood and Wilks's Theorem	37
2.2.5	Wilks's Theorem in Practice	39
3	Spectral Breaks in Centaurus A	45
3.1	Introduction	45
3.2	<i>Fermi</i> -LAT Data Selection, Reduction & Initial Analysis	47
3.3	Centaurus A's Core Spectrum	50
3.4	Model Comparison	55
3.5	Variability	62
3.6	Systematic Uncertainty in IRFs	65
3.7	Physical Interpretations	77
3.8	Conclusions	84
4	Spectral Variability in IC 310	87
4.1	Introduction	87
4.2	Data Reduction & Preliminary Analysis	89
4.3	Spectral Modelling	92
4.4	Constraining the Influence of NGC 1275	100
4.4.1	Extended emission from NGC 1275	101
4.4.2	Localisation of NGC 1275	103
4.5	Temporal Evolution	104
4.5.1	Choice of IRFs	110

4.6	Discussion	113
4.7	Conclusions	117
5	Spectral Variability in Blazars	119
5.1	Introduction	119
5.2	Data Analysis	122
5.3	Integrated Spectra	126
5.4	AGN Variability	128
5.5	A Statistical Basis for Spectral Variability with <i>Fermi</i> -LAT .	135
5.6	Spectral Variability of AGN	137
5.7	Characteristics of Spectral Variability	141
5.7.1	Harder When Brighter?	143
5.7.2	Discussion	145
5.8	Conclusions	150
6	Spectral Variability and Measurement of the Extragalactic Background Light	153
6.1	Introduction	153
6.2	EBL Models	155
6.3	Data Selection	158
6.3.1	<i>Fermi</i> -LAT Analysis	158
6.3.2	Redshift Data	162
6.4	Analysis Methods	163
6.4.1	EBL Fitting	163
6.4.2	Systematic Fitting Uncertainty	165
6.4.3	Joint Likelihood Fit	166
6.5	Method Validation	170
6.5.1	Sources of Systematic Uncertainty	170
6.5.2	Variability Analysis	171

6.6	Results	173
6.7	Discussion	177
6.8	Conclusions	179
7	Conclusions	181
7.1	Acknowledgements	188
A	Catalogue of Spectral Variability	189
A.1	Catalogue Data Selection	189
B	Error on EBL Attenuated Flux	201
C	Phenomenological Spectral Models	203

List of Tables

3.1	A summary of the spectral fits using several models. PL: Power-Law, BPL: Broken Power-Law, LP: Log Parabola. The definitions of each model can be found in Equations 3.1, 3.2 and 3.5 respectively.	
^a	This is compared to that of the best-fit power-law model.	
^b	The broken power-law model does not satisfy Wilks's theorem as it is not nested with a power-law model. The large discrepancy between normalisations is due to the scale energy varying between models, whose fractional difference is then raised to a power of ~ 2.5 .	61

3.2	Table containing the start and end times of the temporal binning used for the lightcurve in Section 3.5	63
3.3	Summary of the broken power law fits with various bracketing IRFs for $\epsilon(E) = 0.1$, $k = 0.13$ and $E_0 = 420$ MeV in the case of the power-law.	73
3.4	Summary of the power law fits with various bracketing IRFs for $\epsilon(E) = 0.1$, $k = 0.13$ and $E_0 = 420$ MeV in the case of the power-law.	76
4.1	A summary of the data selection performed for the preliminary analysis on IC 310.	89
4.2	Summary of the new sources with $TS > 25$ found in our ROI. We have followed the naming convention of the <i>Fermi</i> catalogues for the source column with “PS” identifying the source as newly discovered. ^a PS J0355.3+3910 was no longer a significant source when an analysis centred directly on this source was performed. ^b While not in the 3FGL, this source was identified as 4C+50.11 in Carpenter et al. (2014) and is present in the 4FGL.	90
4.3	A summary of the sky models tested in Section 4.3. ^a The relative number of degrees of freedom (k) in each model as compared to model A. ^b These values are approximate, and are discussed further in Section 4.3. ^c As the $\log(\mathcal{L})$ values returned from the <i>Fermi</i> tools is only relatively correct due to the suppression of model-independent terms, we show only the change in AIC relative to model A.	95
4.4	The parameter values obtained for the initial log parabola and power-law fits to IC 310 within the integrated dataset. .	97

4.5	Detailed information on the binnings assessed in Sect. 4.5.	109
4.6	The significance of this result when varying the IRFs used in the analysis. The event types as defined by the <i>Fermi</i> -LAT collaboration are used in a composite likelihood analysis for PSF and EDISP classes, whereas the combined event class is used for the FB analysis. $TS_{\text{Soft } 1}$ refers to the third time bin (MJD: (55402.66, 55762.66)) and $TS_{\text{Soft } 2}$ refers to the seventh time bin (MJD: (56842.66, 57202.66)).	110
5.1	A catalogue of catalogues. Relevant metadata of the released LAT collaboration catalogues.	121
5.2	A summary of the data selection parameters used within the <i>Fermi</i> science tools in Chapter 5.	123
5.3	Galaxy classification designations used in the 3FGL catalogue as defined by Acero et al. (2015) (Section 4.1, Table 6). Only those designations used in this analysis have been extracted to present here.	124
5.4	A list of the sources removed from the catalogue selection due to their proximity.	126
5.5	Summary of linear fits to PSD slopes in Figures 5.3 and 5.4.	134
A.1	The AGN selected for use in Section 5.2. Definitions as described in text.	189

List of Figures

- 1.1 A sketch of the main features of an active galactic nucleus and its classification depending on the line of sight of the observer. 6
- 1.2 A multiwavelength (and multi-scale) view of Centaurus A. Both the extended lobe structure and core jet emission can be studied across the spectrum. (a) Radio continuum imaging at 1.4 GHz using the ATCA and Parkes 64 m radio telescope (Feain et al., 2011). (b) Microwave imaging combining WMAP K-Band (19–25 GHz, Red; Bennett et al., 2013), Planck030 (20–33 GHz, Green; Planck Collaboration et al., 2014), and WMAP Ka-Band (29–37 GHz, Blue; Bennett et al., 2013). (c) Optical imaging from DSS (Lasker et al., 1990). (d) UV imaging in the GALEX Near UV (1.05–1.69 PHz; Gil de Paz et al., 2007) and GALEX Far UV (1.68–2.23 PHz; Gil de Paz et al., 2007). (e) X-Ray imaging using Chandra ACIS (Red: 0.5–1.0 keV, Green: 1.0–1.5 keV, Blue: 1.5–2.0 keV; Hardcastle et al. 2007). (f) Gamma-Ray imaging using *Fermi*-LAT data with a 90° zenith cut (Red: 100–1000 MeV, Green: 1000-10000 MeV, Blue: >10000 MeV). 10

LIST OF FIGURES

- 1.3 Feynman diagram for pair production interaction in the presence of an atom. In reality, the produced electron pair is almost colinear and must be separated by use of a perpendicular B field. 14
- 1.4 The interaction cross section for photons within matter as a function of energy. Source: www.nuclear-power.net 15
- 1.5 A sketch of the *Fermi*-LAT detector. The incoming dotted line represents a gamma-ray photon of energy sufficient to pair-produce an e^+e^- pair whose tracks shown as two red lines. 16
- 1.6 The orbit of *Fermi* from 2019-05-01 00:00:00–03:00:00 UTC, providing one exposure of the entire sky. The red line represents the projected position of *Fermi* onto the surface of the Earth in this period and the grey polygon represents the area in which the *Fermi*-LAT does not take data due to the SAA. 17
- 1.7 The variation of the 95% containment region of the *Fermi*-LAT PSF with energy. Specifically, the PSF-partitioned IRFs show an improvement from ~ 12 arcmin (PSF3) to just below 3 arcmin (PSF0 as a result of the varying reconstruction quality of the photon events. The EDISP classes show little variation in PSF, meaning that energy reconstruction and PSF reconstruction are not particularly interdependent. FRONT+BACK IRF results are shown for comparison. As the PSF varies depending on orbit, these values were calculated for a year-long exposure in the vicinity of NGC 1275 with weighted averaging over angles of incidence. 19

-
- 1.8 Top: **FRONT** effective area as a function of gamma-ray incidence angle and energy. Bottom: The on-axis effective area with differing type of IRF. The effective area of the PSF and EDISP classes is approximately half that of the **FRONT**- and **BACK**-conversion IRFs as they are partitioned into 4 quartiles as opposed to the physically-motivated IRF categories of front- and back-converting photons within the detector. 20
- 1.9 One week of photon events observed with the *Fermi*-LAT detector coloured according to the EDISP IRF class associated with it. `WP8CTBestEnergyProb` is a parameter resulting from the energy reconstruction algorithms used by *Fermi*-LAT which enumerates the certainty of the reconstruction. 22
- 1.10 A sketch of the electromagnetic cascade resulting from a 1 TeV gamma-ray interacting in the atmosphere, and the resultant Cherenkov radiation as projected onto the ground. The red cascade is taken from simulations using Corsika (Heck et al., 1998) and the image provided by Johannes Knapp available at this URL: <https://www-zeuthen.desy.de/~jknapp/fs/photon-showers.html>. 24
- 2.1 A Feynman diagram showing one of two first-order representations of the Compton scattering process. 28

LIST OF FIGURES

2.2 An instance of an SSC model with its components. Dashed lines represent the synchrotron flux, dotted lines represent the synchrotron self-Compton flux, and the solid line represents the summed flux after attenuation due to the extragalactic background light (EBL). The attenuation was applied using the EBL model of Domínguez et al. (2011), assuming a source at $z = 0.03$. The model used an electron energy spectrum with a broken power-law, a B field of 0.3G, and a jet oriented at 3° to the line of sight in order to provide a plausible model for a *Fermi*-detected BLL. 30

2.3 A sketch of the modelling process. (a) An idealised point source entirely contained within one pixel. (b) The point source convolved with a Gaussian point spread function (independently normalised) indicative of an instrument response. (c) The instrument's response when sampled as a Poisson photon-counting process, resulting in the data measured with an instrument. 35

2.4 An example of overfitting data. (a) A simple quadratic model is used to generate data points with Gaussian errors scaling linearly with $|x|$. (b) The data are fitted with a suitable quadratic model. (c) A sixth-order polynomial is instead used. The data are closer to the model, but the result is clearly absurd. 38

2.5 The reverse cumulative density function for a χ^2 -distributed variable, λ , with degrees of freedom, df , between 1 and 5. Dashed horizontal lines show the probabilities relevant for the sigma values often used as a reference point in this thesis. 40

2.6	Diagram of the number of parameters needed to compare between different models. The diagonal is greyed out due to the self-referential nature of these comparisons and due to symmetry only the top independent values have been listed. An X corresponds to those models which are not nested whatsoever. A number in a cell represents the number of degrees of freedom are required to compare the model in the top row to the model in the first column. An asterisk highlights the issue of parameter degeneracies within the SBPL model where if two power law indices are identical, the break energy becomes entirely degenerate with infinite covariance, breaking the positive-definiteness required for Wilks's theorem to hold.	42
3.1	The TS map of our RoI and the sources contained within. The red contours represent the boundary of the extended template used for the spatial map of the giant radio lobes, and blue crosses represent point sources contained within the RoI, including Cen A in the very centre. Blue circles represent two of the three point sources discovered within this analysis, with the third being outside of the bounds of this image. As noted previously, excess significance (photon counts that are unaccounted for within our model) is present on the edges and within the large-scale structure of the lobes.	51
3.2	The 2D likelihood profile as a function of the core and lobe model normalisations in Cen A. The contours represent $\Delta \log$ likelihood values of 11.829 and 28.744 corresponding to the 3σ and 5σ uncertainty regions respectively. In each bin all other parameters in the model are optimised.	52

LIST OF FIGURES

3.3 The spectral energy distribution (SED) of Centaurus A’s core fitted with a power-law spectrum. The red shaded area represents the 1σ confidence region resulting from the uncertainties on a power-law model fit to this data, with the dashed red line representing the best-fit parameters. The black data points represent the flux predicted in each bin when allowing Cen A’s normalization to vary independently in each energy bin independently. The grey histogram represents the TS of Cen A in each energy bin analysed. 54

3.4 The SED of Centaurus A from Figure 3.3 with the additional broken power-law model shown in blue. 55

3.5 Distribution of the source TS when fitting a broken power-law to a simulated (unbroken) power-law signal with a fixed break energy parameter. As only one degree of freedom has been added, this should follow a χ_1^2 distribution. 58

3.6 A histogram of the β values of 3FGL sources with log parabola spectral models, compared with the value calculated for Cen A whose uncertainty region is shown by the dashed lines and light grey region. The pile-up of sources at $\beta = 1$ is due to the parameter constraints introduced by Acero et al. (2015) to avoid excessive curvature. It should also be noted that values of $\beta=0$ suffer a selection bias as the log parabola reduces to a power-law at $\beta = 0$, meaning these sources do not meet the requirement for a 4σ preference over a power-law. What is most relevant is the lack of any sources with negative curvature, as is found in Cen A. 60

3.7 Lightcurves for energies $E < 2.6$ GeV (top) and $E > 2.6$ GeV (bottom). The blue colourmap represents the change in likelihood from the optimised position in each bin, and the red line represents the overall average flux. 64

3.8 The $\pm 1\sigma$ uncertainty regions as determined by Monte Carlo simulations and on-orbit calibration for 3 different IRF combinations. The point-to-point correlation within this range cannot be shown in this plot, but the correlation is strong in neighbouring energy bins, becoming uncorrelated after half a decade in energy (as this calculation is performed logarithmically). 66

3.9 A sketch of the limiting bracketing IRF functions recommended by Ackermann et al. (2012). Red lines represent tanh alterations and blue lines represent the linear offsets. The black line represents the case assumed by the analysis tools by default. 68

3.10 A comparison between the linear and tanh bracketing IRFs with baseline A_{eff} data provided with the *Fermi* Science Tools. The effective area for the front detector at $\cos(\phi) = 1$ is plotted with various bracketing functions, $B(E)$ for relative uncertainty $\epsilon(E) = 0.1$. **(a)** Linear $B(E)$ for relative uncertainty $\epsilon(E) = 0.1$. **(b)** Tanh bracketing functions as shown in the legend, with a decorrelation energy of 420 MeV. 69

LIST OF FIGURES

3.11 A comparison between the linear and tanh bracketing IRFs. The effective area for the front detector at $\cos(\phi) = 1$ is plotted with various bracketing functions, $B(E)$ for relative uncertainty $\epsilon(E) = 0.1$. Blue: Linear $B(E)$ for relative uncertainty $\epsilon(E) = 0.1$. Red: Tanh bracketing functions as shown in the legend, with a decorrelation energy of 420 MeV. 71

3.12 The $B(E) = \pm 1$ IRFs plotted either side of the naive $B(E) = 0$ fit. 72

3.13 The $B(E) = \pm 1$ IRFs plotted either side of the naive $B(E) = 0$ fit (darker red). The bar chart represents the TS of each binned data point. 73

3.14 A subset of the power-law fits to the bootstrapped linear IRF alterations. The dashed lines represent individual best-fit power laws and are coloured according to the magnitude of $\epsilon(E)$ (more purple corresponds to more negative). 74

3.15 A subset of the power-law fits to the bootstrapped tanh IRF alterations. The dashed lines represent individual best-fit power laws and are coloured according to the magnitude of $\epsilon(E)$ (more purple corresponds to more negative). 75

3.16 Prompt emission from an NFW-distributed WIMP dark matter annihilation to $b\bar{b}$ with a relic cross section. The behaviour before the break is modelled with a simple power-law extrapolated from the broken power-law in Section 3.3. The data above 300 GeV are the results published by Aharonian et al. (2009). 81

3.17 DM fit with a spiky dark matter profile, with a DM mass of 3 TeV annihilating to $t\bar{t}$ results in a best-fit cross section of $1.6 \times 10^{-32} \text{ cm}^2 \text{ s}^{-1}$. The extrapolated power-law jet from the best-fit spectrum below the break energy is shown with the dot-dashed line. The data above 300 GeV are the results published by Aharonian et al. (2009). 82

3.18 MSP population fit with $\rho_{\text{MSP}} \propto r^{-2.4}$ and an extrapolated power-law jet from the best-fit spectrum below the break energy. The data above 300 GeV are the results published by Aharonian et al. (2009). 83

4.1 The spectral energy distribution of IC 310 when fitted with a log parabola model. This fit is performed over the entire 8-year time period and is jointly optimised with the nearby point sources and diffuse background emission templates. . . 91

4.2 The variation of the SED of IC 310 for each model considered in Table 4.3. Each SED was calculated by profiling the normalisation of IC 310 in 10 logarithmically-spaced bins between 100 MeV and 500 GeV. It is readily seen that only the 3 low energy bins are significantly affected by this modelling, the source of our scepticism in Sects. 4.4.1, 4.4.2 and 4.3. Upper limits are shown if the flux density in an energy bin is consistent with 0 within 2σ 96

LIST OF FIGURES

4.3 A TS map of the ROI in a 3° radius around IC 310, showing the relevant point sources for each model assessed in Section 4.2. The statistical significance of PS J0312.8+4121 can be seen as a peak before a point source model is included from model D onwards. Other potential point sources manifesting as smaller peaks can be seen, but we limit ourselves to a 5σ threshold for new point sources. 99

4.4 Left: The *Fermi*-LAT counts map centred upon NGC 1275 with a 1° radius using 0.1° bins. The annulus is coloured according to the TS of a point source with index $\Gamma = 2.45$ at this position in an analysis between 578.5 MeV and 3 GeV using a composite (PSF0 + PSF1 + PSF2 + PSF3) likelihood. The red ellipse represents the localised position of NGC 1275 from Section 4.4.2 and the black ellipse represents the localised position of IC 310 obtained from the 3FGL. Right: The annulus TS values unwrapped and plotted as a function of bearing relative to IC 310. The corresponding colourmap value is plotted underneath. 102

4.5 The SED for each annual bin, where each spectrum shown has a TS of at least 19.33 representing an approximate 4σ cut defined with two degrees of freedom: normalization and spectral index. The shaded areas around each optimised model represent the 1σ flux uncertainty propagated from the uncertainties in spectral model parameters after the likelihood fit. 105

4.6 The 1σ uncertainty contour in the parameter space of a power law for each time bin assessed in Section 4.5. The normalisation (but not spectral shape) of NGC 1275 has been left free to vary for each parameter combination. Bins with optimised parameters outside of this parameter space are excluded. . . . 106

4.7 The likelihood space from Figure 4.6 profiling over normalisation. Solid profiles represent bins in which IC 310 was not detected by VHE telescopes (as determined from Table 4.5 and dashed lines represent VHE detected bins. The black lines represent the sum of the VHE detected and non-VHE detected profiles with the minimum subtracted, and thus the combined maximum. Artefacts are introduced due to the limited region over which the grid scan was performed and also due to the resolution chosen for the grid scan. Bins are chosen for comparison with Figure 4.6. 108

LIST OF FIGURES

- 4.8 The variation in PSF with IRF choice and energy. The data in the top row correspond to data partitioned into the PSF3 IRF class, the data in the middle row into EDISP3, and the data in the bottom row in the summed FB class (all photons would fall into this category). Data in the first column are binned in energy from 100–126 MeV, data in the second column from 1 GeV to 1.26 GeV, and data in the final column from 10–12.6 GeV. The image is zoomed between energy ranges for easier viewing, decreasing from 10°, to 3°, to 1° with increasing energy. These data represent the first annual bin of our lightcurve. The circles plotted show the 95% containment angles of the PSF of the *Fermi*-LAT instrument using the IRFs as listed in the legend. Containment angles were calculated by numerically integrating the PSFs returned by the `gtpsf` tool within the *Fermi* science tools. Asymmetry due to the fisheye effect and inclination angles to the LAT boresight is not considered. The positions of NGC 1275 and IC 310 are marked along the diagonal to show the scales relevant to our analysis. 112
- 4.9 95% containment radii for our region of interest as a function of energy and instrument response function partition. Each line represents an individual IRF class, coloured as described in the legend. Vertical lines represent the 95% containment regions of the PSF3 and FB classes, and the horizontal line represents 0.6°, the separation of IC 310 and NGC 1275. Artefacts are due to the finite resolution used in the `gtpsf` tool used to calculate these values. 114

-
- 4.10 The analysis from this work plotted with the data presented in Aleksić et al. (2014a). The ‘hard’ spectrum plotted from this work represents the first bin of our analysis (MJD: 54682.66, 55042.66), and the ‘soft’ spectrum represents the third bin (MJD: 55402.66, 55762.66). These data are not simultaneous. 115
- 5.1 The measurement of photon index as measured in this sample compared to the same sample from the more recent 4FGL catalogue. The dashed black line represents a one-to-one correspondence between the 4FGL and this work. 127
- 5.2 The lightcurve of PG 1555+113 from our catalogue analysis as an example. The colour gradients show the $\Delta \log(\mathcal{L})$ from the maximum likelihood value. A change in $\log(\mathcal{L})$ of 0.5 represents the 1σ uncertainty interval, and a change of 2.0 represents the 2σ uncertainty interval. This is true as our scan of the normalisation (a direct mapping to flux) is distributed as $\chi_{n=1}^2$. The red line represents the measured (average) flux value over the 8-year integrated exposure. . . 130
- 5.3 Power spectral density of FSRQ and BLL objects with their associated uncertainty. Uncertainties were calculated using Monte Carlo simulations using the uncertainty in flux measurements, and the 32nd and 68th percentiles of the resultant PSD distributions were taken as the $\pm 1\sigma$ uncertainty region. 132
- 5.4 As Figure 5.3, but showing BLL objects classified by their synchrotron peak frequency. 133

LIST OF FIGURES

5.5 Significance of **(a)** TS_{var} **(b)** $TS_{\text{full var}}$ and **(c)** $TS_{\text{spec var}}$ by optical class and in the case of BLL, the peak synchrotron frequency. The corresponding values are not plotted for FSRQs as they are dominantly LSPs. 138

5.6 Flux variability in the 8-year lightcurves compared to their significance in the seed 3FGL catalogue. 139

5.7 As in Figure 5.6 but showing spectral variability. This plot is however only illustrative as each point will have differing degrees of freedom due to their varied spectral models. . . . 140

5.8 Significance of $TS_{\text{spec var}}$ with time. Blue: The median $TS_{\text{spec var}}$ value of the AGN sample. Grey: The 32nd and 68th percentile ($\approx 1\sigma$) $TS_{\text{spec var}}$ values amongst the AGN sample. Values greater than 8σ are excluded due to the limitations of the `scipy erf-1` routine. 141

5.9 Histogram of spectral index for the monthly lightcurve of 3FGL J1309.3+4304 modelled by a power-law. The black line represents the optimised index in the integrated 8-year dataset, and the grey is the associated 1σ uncertainty. 142

5.10 As with Figure 5.9 except spectral indices are normalised to the 8-year integrated value, and summed for all BLL and FSRQ objects. 144

5.11 Histogram of linear fit slopes to BLL flux-index relationships. A negative slope represents a hardening with increased flux. 146

5.12 An example of one BLL object in the $\log_{10}(F)$ - Γ plane. The uncertainty region on the fit is presented as the 32nd and 68th percentiles of the posterior distribution of the parameters obtained using `LinMix`. 147

6.1	Feynman diagram for pair production to electrons.	156
6.2	The EBL intensity predicted by various models (named in legend) at $z = 0$. We use the the following abbreviations for the published predictions for comparison: INO13 (Inoue et al., 2014), FIN10 (Finke et al., 2010), DOM11 (Domínguez & Ajello, 2015), K&D10(Kneiske & Dole, 2010), FRA08 (Franceschini et al., 2008), GIL12-fiducial (Gilmore et al., 2012). Upper limits from gamma-ray observations and lower-limits from galaxy count predictions are included.	159
6.3	A stacked histogram showing the distribution of redshifts by 3FGL source classification in our sample. The overall height of a bar represents the number of sources in a redshift bin, and the colours are stacked to represent the breakdown within a bin by source type. In general, BL Lac objects dominate at lower redshifts, followed by FSRQs at higher redshifts.	161
6.4	Results of the data challenge set in Section 6.4.2. Simulated AGN power-law spectra parameters are denoted by green squares which effectively represents the “truth” to which Gaussian noise was added. Fitted parameters using the likelihood fitting routine are denoted by black crosses with associated uncertainties.	167
6.5	Parameters measured as in Figure 6.4 except for sources simulated as log parabolae.	168
6.6	Relative uncertainty between fitting SEDs with my MINUIT-based routine and that returned by the <i>Fermi</i> science tools.	168

LIST OF FIGURES

6.7 The TS_{var} calculated in this work plotted against the TS_{var} calculated for a source in the 3FGL. The black dotted line represents a one-to-one relationship between the two values, and the red line represents a linear regression fit ($r^2=0.842$) to the log of the values. The spread is largely due to different data selection processes. 172

6.8 The r^2 value of the linear fit in log space as a function of the number of months of data used in the calculation of the TS_{var} . This is convincing evidence that the correlations are not simply due to similar scaling, but inherent correlation between the sources in the dataset. 173

6.9 EBL normalisation as a function of redshift (where available) for the 6 models considered in this work. As no other work considered the Inoue et al. (2014) model, only our results are presented. Independent measurements are sourced from HESS (H.E.S.S. Collaboration et al., 2013), MAGIC (Ahnen et al., 2016b,a; Mazin et al., 2017), *Fermi*-LAT (Ackermann et al., 2012) and combined (Biteau & Williams, 2015). The grey bar represents the overall measurement excluding spectrally variable sources. 175

6.10 EBL normalisation for our selection of models as described in Figure 6.9, except with sources restricted to only those sources without significant variability, resulting in fewer available redshift bins. 176

Abstract

The *Fermi*-LAT telescope is used to study the effects of spectral variability upon physical measurements made using gamma-ray data. Firstly, two individual AGN are studied, namely Centaurus A and IC 310. In both cases, the spectrum shows a significant hardening above GeV energies. In the case of Cen A, this is shown to be an intrinsic effect, with little variability to be found. On the other hand, in IC 310, this hardening is shown to be an artefact of spectral variability. To assess the prevalence of spectral variability in the population of AGN, first the statistical tools to assign confidence to our measurements are developed. 302 bright BL Lac objects and FSRQs are then analysed to characterise the variability of the AGN population. Over half of the AGN studied were shown to be spectrally variable at the 5σ level, with a far higher fraction variable in flux. This information is then applied to measurements of the extragalactic background light attenuation found in these gamma-ray AGN, revealing a bias towards higher EBL intensity caused by spectrally variable AGN in the sample. We conclude that spectral variability is an effect that must be taken into account in all measurements of fundamental physical parameters from AGN measurements.

Declaration

The work presented in this thesis includes contributions from other authors. Chapter 3 on the topic of Centaurus A dominantly contains this authors contribution to published work (Brown et al., 2017). Other authors, namely A. M. Brown, C. Boehm, T. Lacroix, P. M. Chadwick, and J. Silk all contributed to the original work, but I was responsible for all data analysis including spectra, lightcurves, and interpretation presented here. Dark matter and MSP modelling that is included in the final section of Chapter 3 is the work of T. Lacroix.

Chapters 4 and 5 are entirely my own work.

Chapter 6 contains contributions from T. Armstrong with whom this work was presented at the *Fermi* Symposium 2017. The data analysis, statistics, and fitting method are entirely my own. T. Armstrong was responsible for the visualisation of Figures 6.9, and 6.10.

Chapter 1

Introduction

High-energy astrophysics deals with the very smallest and very largest objects in the Universe. The physical processes that allow insight into gargantuan astronomical bodies begin in the microcosm of individual particle interactions. Profound knowledge can be gained from the detection of a single photon or neutrino that has crossed the cosmos to reach our instruments on Earth, making high-energy astrophysics a field at the extremes. At energies of this scale*, the boundaries between traditional astronomy and particle physics begin to blur. While this provides rich research opportunities for both astronomers and particle physicists it also results in a difficult entanglement of phenomena. The clash between astronomy and physics beyond the standard model is a major theme of the research presented in this work, and deconvolving this situation will no doubt continue well into the future. In particular the topics of weakly interacting dark matter models, axion-like particles, and the star formation history of the Universe will be

*For reference, the author considers “high-energy astrophysics” to concern energies in excess of 511 keV, the rest mass energy of the electron, although the boundaries of what is considered to be “high” (or “very high”) has no general consensus and varies between fields and even individuals.

touched upon within the context of the mitigating factors of the variability of emission from extragalactic supermassive black holes. This chapter shall serve as a brief introduction to the relevant astrophysical objects, their associated physics, and the particle physics that can be probed with high-energy observations.

1.1 Active Galactic Nuclei

1.1.1 Overview

Active Galactic Nuclei (AGN) are some of the most extreme objects in the Universe. Highly luminous over many orders of magnitude in frequency, AGN are often violently variable on timescales down to days or even hours. Referring to compact objects found in the heart of galaxies, these individual nuclei can outshine the combined luminosity of the billion or so stars that accompany them. These nuclei are thought to be supermassive black holes, whose emission is powered by the accretion of surrounding matter, and whose resultant relativistic jets are responsible for accelerating particles to enormous energies. The study of AGN has become vital to the explanation of galaxy evolution and astrophysical jets[†], and AGN constitute one of the largest populations of objects among sources detected in the gamma-ray regime.

1.1.2 History

The conventional view of galaxies as simple gravitationally-bound ensembles of stars as held at the start of the 20th century was rather uncontroversial.

[†]Though the contribution of galactic black hole binaries should not be understated (Fender et al., 2004)

This model simply predicted galactic emission to represent the stacked emission of stellar objects resulting in a continuum peaking in the optical–UV. Indeed, the first spectroscopic measurements of spiral galaxies (Fath, 1909) set out to confirm this and found continuous spectra with absorption lines corresponding to solar absorption features in NGC 224[‡] and NGC 3031, but finding “bright” lines in NGC 1068, an object now known to host an AGN. This tantalising glimpse into nuclear line emission from extragalactic objects would be investigated further by Slipher (1917) and Hubble (1926), before Seyfert (1943) revealed 6 galaxies with stellar cores showing high-ionisation emission lines and varying degrees of broadening. Objects with highly luminous cores, but whose host galaxy is easily detectable are now known as Seyfert galaxies in honour of this discovery.

Simultaneously with these spectroscopic measurements, the nascent field of radio astronomy was slowly developing[§] through a series of seminal papers by Karl Jansky (Jansky, 1932, 1933, 1935) which established the astronomical emission consistent with solar emission and diffuse emission towards the Galactic centre. Followed by radio sky surveys (Reber, 1944) and the discovery of multiple discrete, compact radio sources across the sky (Bolton & Stanley, 1948; Ryle & Smith, 1948; Ryle et al., 1950; Smith, 1951), a population of astronomical sources emitting non-thermal radiation had been discovered. Further measurements and improving resolution[¶] allowed the first unambiguous optical identifications of the radio sources Cas A and Cyg A (Baade & Minkowski, 1954), correlating these sources as “galactic-emission nebulosity of a new type” and “an extragalactic affair, two galaxies

[‡]Also known as M31 or the Andromeda Galaxy

[§]Though largely outside of academia at Bell Laboratories

[¶]From 6° by Reber (1944) to 5' by Ryle et al. (1950), leading to resolved extended emission detected by Hanbury Brown et al. (1952)

in collision”[‡] respectively.

Over the next decade surveys improved (especially Edge et al., 1959; Bennett, 1962) allowing the identification of scores more radio sources and in many cases their optical counterparts. Some sources stood out, however. 3C48 (Matthews & Sandage, 1963) and eventually 3C273 (Hazard et al., 1963) were strong radio sources unassociated with galaxies (assumed to be stellar in origin due to the lack of “extragalactic nebulae”) that showed photometric variability but whose broad emission lines corresponded to no known transitions. The mystery of these “radio stars”, now dubbed quasi-stellar radio sources (more commonly QSOs or quasars), was solved upon the realisation (Schmidt, 1963; Oke, 1963; Greenstein, 1963) that the emission lines were that of the redshifted Balmer series given the sources were extragalactic with significant velocities resultant from Hubble expansion**. 3C48 ($z = 0.37$) and 3C273 ($z = 0.16$) were then followed quickly by far more distant objects (e.g. PKS 0237-23 at $z = 2.223$; Arp et al., 1967; Burbidge, 1967; Greenstein & Schmidt, 1967). Notably, the optical characteristics of QSOs — bright, broad emission lines from distant galaxies with a UV excess and violent variability — seemed to suggest a link with the Seyfert galaxies; the question remained as to what energy source could fuel this luminosity?

In the Galactic regime early theoretical estimations by Henyey & Keenan

[‡]Minkowski had been particularly sceptical of this reasoning given Baade’s recent theoretical study of colliding systems (Spitzer & Baade, 1951). A wager was settled when optical spectra found nuclear emission lines thought to be the result of excited gas from a collision. Later work (Jennison & Das Gupta, 1953; Burbidge, 1959) proved the collision theory incorrect, though it is unclear if the wager was ever repaid. Minkowski eventually became a strong supporter of the collision theory (Robinson et al., 1965)

**Or large gravitational redshifts, although this had the downside of being unable to explain the broad emission lines

(1940) and Hard Townes (1947) noted that thermal emission would require unlikely interstellar gas temperatures in excess of 10^5 K — more than an order of magnitude larger than commonly expected — for which Alfvén & Herlofson (1950) suggested the alternative mechanism of synchrotron emission from cosmic ray electrons (generated in supernovae) in the Galactic magnetic field. The extraordinary luminosities (in some cases $> 10^{59}$ erg) seen in QSOs had to be powered by alternative means. The ideas put forward at the time were varied: chain reactions of supernovae (Burbidge, 1961), matter-antimatter interaction (Burbidge & Hoyle, 1957), dense star clusters (Spitzer & Saslaw, 1966), and initial considerations of massive compact objects in which “angular momentum is transferred from the central star to a surrounding disk of gas” and “a magnetic field could be wound toroidally between the central star and a surrounding disk’ [sic]” (Hoyle & Fowler, 1963b,a). This suggestion was prescient, but the connection between a QSO’s energy source and black holes was not fully investigated until Salpeter (1964) and Zel’dovich & Novikov (1964) calculated the energy released from accretion onto a supermassive black hole. The final piece of this puzzle was proved by Lynden-Bell (1969), which argued that such systems could unify the sources responsible for energy production in quasars and Seyferts depending on black hole mass and mass accretion rate.

1.1.3 AGN Unification

The prospect of unifying the many classifications of AGN into a single class of objects under different observational and environmental conditions is an obviously satisfactory one, and has been hinted at strongly in the previous section. While I have made this out to be inevitable, this is a *post hoc* argument and the evidence for unification has rarely been unequivocal. The

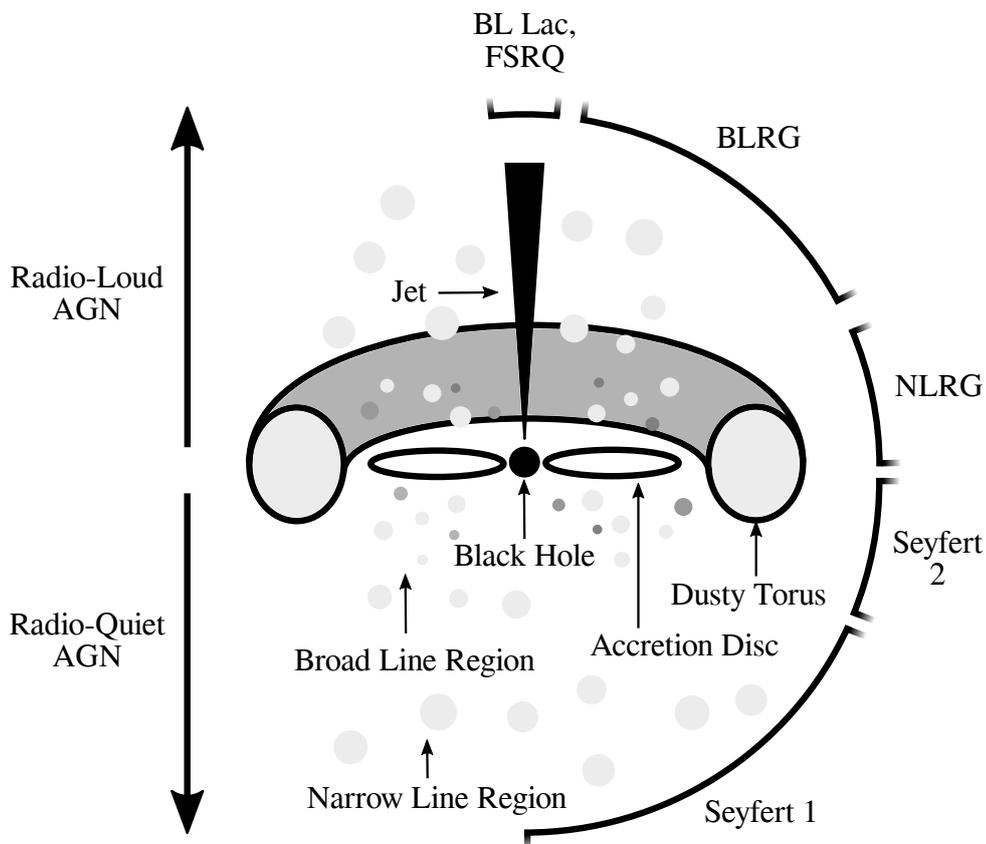


Figure 1.1: A sketch of the main features of an active galactic nucleus and its classification depending on the line of sight of the observer.

first descriptions of AGN unification schemes published by Urry & Padovani (1995) and Antonucci (1993) were based on line of sight angle to the radio jet and radio loud / radio quiet states. This is the dominant view held today and is summarised in Figure 1.1.

This schema posits that active galaxies contain supermassive black holes at their centres with masses of order 10^6 – $10^{9.5}M_{\odot}$ (Kormendy & Richstone, 1995), powered by an accretion disc. The accretion disc is itself very lumi-

nous, with luminosities that can approach the Eddington limit^{††}. In some cases, relativistic jets are formed around the black hole, leading to strongly beamed emission that can appear orders of magnitude brighter to some observers. The gas closely surrounding the nucleus is generally separated into two “regions” – the broad line region (BLR) and the narrow line region (NLR) – distinguished by the width of their nuclear emission lines due to Doppler broadening. In addition, this model requires a torus of dusty clouds surrounding the compact object at a scale of around a parsec (Krolik & Begelman, 1988) to obscure observation of the BLR from certain angles. From these components, the axisymmetric nature of this model leads to the diverse range of AGN classifications as described in Figure 1.1.

Our unified model as presented above is certainly useful and has been largely supported by most observations of AGN, but there are still many open questions as to these classifications. For instance, Tran (2001) showed that only some ($\sim 50\%$) Seyfert 2 objects contain an obscured BLR, and a dependence on luminosity in this relationship (Wu et al., 2011) might suggest that the lack of a BLR is related more to central engine luminosity than obscuration. Indeed many of the open questions listed in the conclusions of Urry & Padovani (1995) are still under debate, but for now it shall remain the “standard model” of an AGN.

1.1.4 AGN Classification

AGN can be broadly classified in a number of ways, depending on their spectral features, radio properties, and morphology. In the optical regime, Seyfert galaxies are divided into Type 1 and Type 2 based on the pres-

^{††}The Eddington limit is the maximum luminosity at which outward radiation pressure balances the inward gravitational pull of gravity; it sets an upper bound on stable accretion onto compact objects.

ence or absence of broad emission lines, while quasars and blazars typically represent more luminous or jet-dominated systems. Another widely used division is between radio-loud and radio-quiet AGN, depending on the ratio of radio to optical luminosity.

For radio-loud AGN, the Fanaroff–Riley classification (Fanaroff & Riley, 1974) is commonly employed to distinguish between two morphological types based on radio maps. FR I sources are generally less luminous and exhibit core-brightened jets that fade with distance from the nucleus, whereas FR II sources show edge-brightened lobes with strong terminal hotspots and are typically more powerful. This distinction also tends to correlate with excitation class: FR I sources are more often associated with low-excitation radio galaxies (LERGs), and FR II with high-excitation radio galaxies (HERGs) and quasars. The FR classification plays an important role in understanding the large-scale structure of radio jets and lobes, which are relevant for interpreting both synchrotron and inverse-Compton emission mechanisms.

1.1.5 AGN Across the Spectrum

One of the most outstanding features of AGN is that they have such extremely high luminosities originating from such a compact region, often outshining their host galaxies. An important addendum is that this is true across an huge range in frequency, from the radio to the gamma-rays in radio-loud AGN. Figure 1.2 shows the view of Centaurus A (Cen A), the closest AGN to Earth^{††}, in several wavebands.

The broadband spectrum of an AGN is due to thermal emission from

^{††}Though there is some disagreement in its distance, estimates ranging from 3 Mpc (Majaess, 2010) to 4.7 Mpc (Crnojević et al., 2016)

the accretion disc peaking in the optical–UV–Soft X-Ray, the obscuring torus dominates in the IR, and non-thermal emission from the jet dominates almost everywhere else.

Throughout this thesis, the term "AGN" is used to refer specifically to those active galactic nuclei that are bright in gamma-ray emission, typically observed by instruments such as the *Fermi*-LAT. Not all AGN fall into this category: gamma-ray bright AGN are predominantly radio-loud, jet-dominated sources such as BL Lacs and flat-spectrum radio quasars (FSRQs), which are collectively known as blazars when viewed close to the jet axis. These comprise the majority of extragalactic gamma-ray sources in the Fermi-LAT catalogues (e.g. Acero et al., 2015; The Fermi-LAT collaboration, 2019).

It is also important to note that there exists a substantial population of AGN that are radiatively inefficient. These sources lack many of the standard structural components seen in more luminous AGN, such as the accretion disk, dusty torus, and broad and narrow line regions. This is typically due to low accretion rates and/or advection-dominated accretion flows (ADAFs), and they are often weak or undetectable in high-energy emission. Radiatively inefficient AGN are therefore underrepresented in gamma-ray surveys and are not the focus of this work.

The emission processes involved will be described in more detail in Section 2.1.

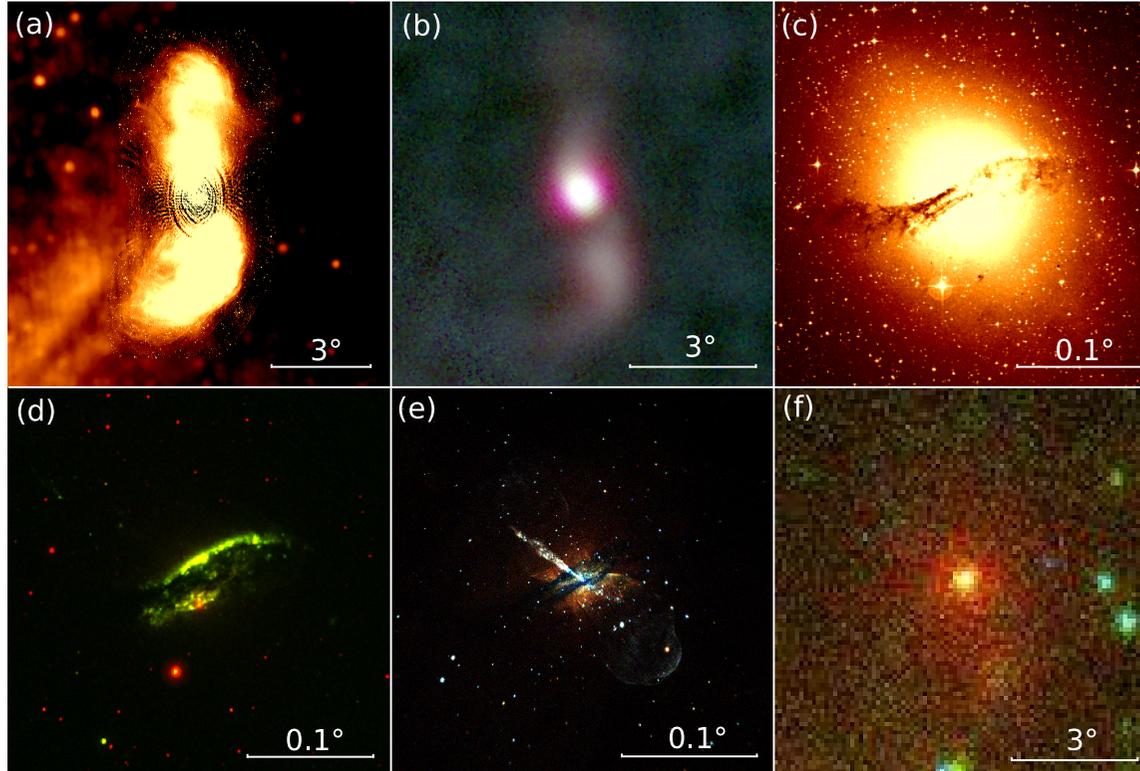


Figure 1.2: A multiwavelength (and multi-scale) view of Centaurus A. Both the extended lobe structure and core jet emission can be studied across the spectrum. (a) Radio continuum imaging at 1.4 GHz using the ATCA and Parkes 64 m radio telescope (Feain et al., 2011). (b) Microwave imaging combining WMAP K-Band (19–25 GHz, Red; Bennett et al., 2013), Planck030 (20–33 GHz, Green; Planck Collaboration et al., 2014), and WMAP Ka-Band (29–37 GHz, Blue; Bennett et al., 2013). (c) Optical imaging from DSS (Lasker et al., 1990). (d) UV imaging in the GALEX Near UV (1.05–1.69 PHz; Gil de Paz et al., 2007) and GALEX Far UV (1.68–2.23 PHz; Gil de Paz et al., 2007). (e) X-Ray imaging using Chandra ACIS (Red: 0.5–1.0 keV, Green: 1.0–1.5 keV, Blue: 1.5–2.0 keV; Hardcastle et al. 2007). (f) Gamma-Ray imaging using *Fermi*-LAT data with a 90° zenith cut (Red: 100–1000 MeV, Green: 1000–10000 MeV, Blue: >10000 MeV).

1.1.6 Observational Properties of AGN

Active galactic nuclei (AGN) are characterised by their high luminosities, often outshining the combined stellar emission of their host galaxies across a broad range of wavelengths. One of the defining properties of AGN is their compactness — observational evidence shows that the emitting region must be very small. This conclusion arises from the observation of rapid variability: in some AGN, the flux can vary on timescales of hours or even minutes (e.g. Ulrich et al., 1997). Such rapid changes imply that the emitting region is confined to a volume no larger than the light travel time over the variability timescale.

Moreover, very long baseline interferometry (VLBI) observations of radio-loud AGN reveal highly collimated jets emerging from regions smaller than a parsec in size (Zensus, 1997), again pointing to a compact central engine.

The extraordinary luminosity and compactness of AGN are difficult to explain without invoking accretion onto a supermassive black hole (SMBH). The gravitational potential energy released by matter falling into a black hole provides a natural and efficient power source. This model is further supported by observations of stellar and gas dynamics in the centres of nearby galaxies, where high-resolution spectroscopy reveals Keplerian motion around an unseen massive object (Kormendy & Richstone, 1995; Ferrarese & Merritt, 2000).

Additionally, the Eddington argument — relating observed luminosity to the maximum possible luminosity of a body in hydrostatic equilibrium under radiation pressure — provides a lower limit on the mass of the central engine, often 10^6 – $10^{10} M_{\odot}$, again implying SMBHs (Soltan, 1982).

1.2 Astronomical Gamma-Rays

1.2.1 Overview

Why does gamma-ray astronomy exist? Given the large-scale projects in astronomy such as the Cherenkov Telescope Array and the *Fermi* Large Area Telescope (hereafter *Fermi*-LAT) this is an obviously rhetorical question, but the answer hasn't always been obvious. There are some big questions to answer^{§§}: Should we actually see anything with a gamma-ray telescope? Can this tell us something about the universe that we couldn't decipher at other wavelengths? How do you plan to detect an astronomical gamma-ray?

Firstly, as we shall see in Section 2.1, the physical processes inherent to the high energies created in the vicinity of supermassive black holes in AGN and in other dense compact objects such as pulsars are easily capable of accelerating cosmic rays to energies of PeV (H.E.S.S. Collaboration et al., 2016). Due to the high energies explored in gamma-ray astronomy, it can help us answer fundamental questions in astronomy (e.g. by probing cosmological effects with the extragalactic background light; Fermi-LAT Collaboration et al. (2018a)) and probe the energies relevant in particle physics (e.g. WIMP dark matter, axion-like particles, and other beyond the standard model physics; Balázs et al., 2017). Many gamma-ray telescopes are reaching or exceeding sensitivities required to detect the expected flux of relic abundance WIMP dark matter models (Balázs et al., 2017).

With regard to the detection of gamma-rays, there are certainly a few issues to overcome. Telescope design differs hugely across the electromagnetic spectrum, from the radio to the near-infrared where antennae become unable to satisfactorily detect photons (Hansen, 1981), whereupon the CCD

^{§§}Please note, this is not an exhaustive list

and the photoelectric effect take charge, before standard reflecting optics must give way to grazing incidence mirrors or coded aperture masks at X-ray energies. By the time we reach gamma-ray energies, the lack of an ability to refract or reflect photons requires a substantial redesign, and the penetration depths of charged particles above a few hundred GeV pose even bigger problems for smaller detectors.

In this chapter we will look at the development of the field of gamma-ray astronomy, its origins, and the basic physics behind the emission and detection of gamma-rays. Specifically we explore the inner workings of the *Fermi* Large Area Telescope (hereafter *Fermi*-LAT), its effective area, point spread function (PSF), and other strengths and drawbacks with respect to the analyses we will perform in this work.

1.2.2 Absorption Cross Sections

The photoelectric effect, responsible for the interactions that allow CCDs and CMOS sensors to operate, scales inversely with the cube of photon energy above the highest atomic binding energy (tens to hundreds of keV). The interaction cross section of a pair-production interaction (See Figure 1.3) increases above twice the rest mass energy of the particles created (in our case, usually e^+e^- at 2×511 keV). There is a gap between these two dominant effects in which Compton scattering becomes extremely relevant around the MeV scale, a phenomenon which has been suggested for use within a space-based detector (Moiseev et al., 2015). Above a GeV, pair production becomes the only realistic interaction process to measure gamma-rays, for which the relevant quantities are plotted in Figure 1.4. In the case of the *Fermi*-LAT instrument, materials with high nuclear charge (Z) are used to convert photons to electrons, but at higher energies this is impractical and

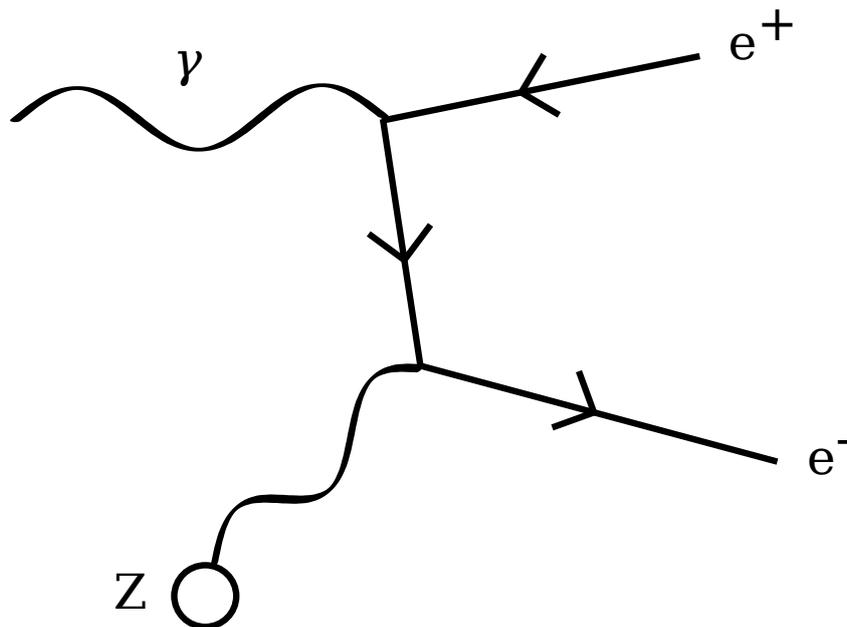


Figure 1.3: Feynman diagram for pair production interaction in the presence of an atom. In reality, the produced electron pair is almost colinear and must be separated by use of a perpendicular B field.

the atmosphere of the Earth is used to make this conversion in the case of imaging atmospheric Cherenkov telescopes (IACTs).

1.2.3 *Fermi* Large Area Telescope

The *Fermi* Large Area Telescope (henceforth *Fermi*-LAT) is a space-based pair-conversion gamma-ray telescope capable of detecting gamma-rays with energies in the approximate range 20 MeV–1 TeV. The LAT has a particularly large effective field of view (FoV) of 2.6 sr ($\sim 20\%$ of the entire sky), an energy resolution of $\sim 4\%$ and an angular resolution of $\sim 0.1^\circ$ at 1 GeV (Michelson, 1996; Ackermann et al., 2012). Functionally, the LAT is composed of a series of Silicon tracker modules interspersed with material designed to convert incoming gamma-rays into e^+e^- through pair produc-

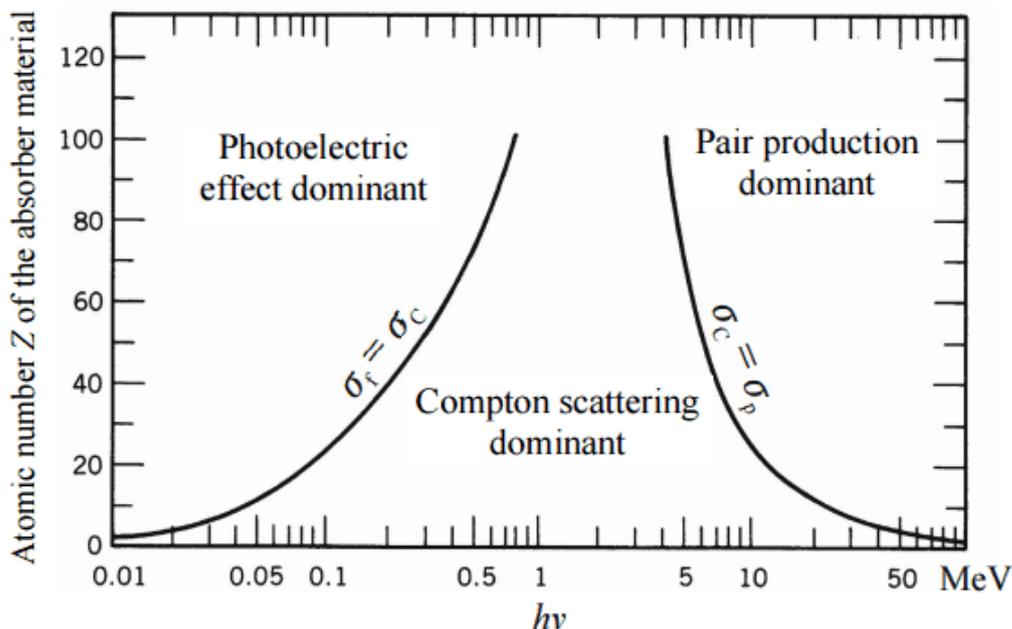


Figure 1.4: The interaction cross section for photons within matter as a function of energy. Source: www.nuclear-power.net

tion followed by a CsI calorimeter, this combination allowing measurement of both the direction and energy of the incoming photon. To allow rejection of incident charged particles, a segmented anti-coincidence shield is also included. A sketch of this instrument can be found in Figure 1.5.

The *Fermi* Gamma-ray Space Telescope¹⁴ has been in orbit and operational since August 2008 and is functional today, far beyond its original estimated lifetime of 5 years (Michelson, 1996). The large FoV and rapid low-Earth orbit allow for all-sky coverage every 3 hours (every 2 orbits), excepting some periods of downtime while in proximity to the South Atlantic Anomaly (SAA) as shown in Figure 1.6. All-sky imaging on such short timescales is fundamentally important to allow long timescale monitoring of gamma-ray sources from periodic galactic pulsars and binaries to flaring

¹⁴The LAT is joined by a gamma-ray burst detecting module, the GBM

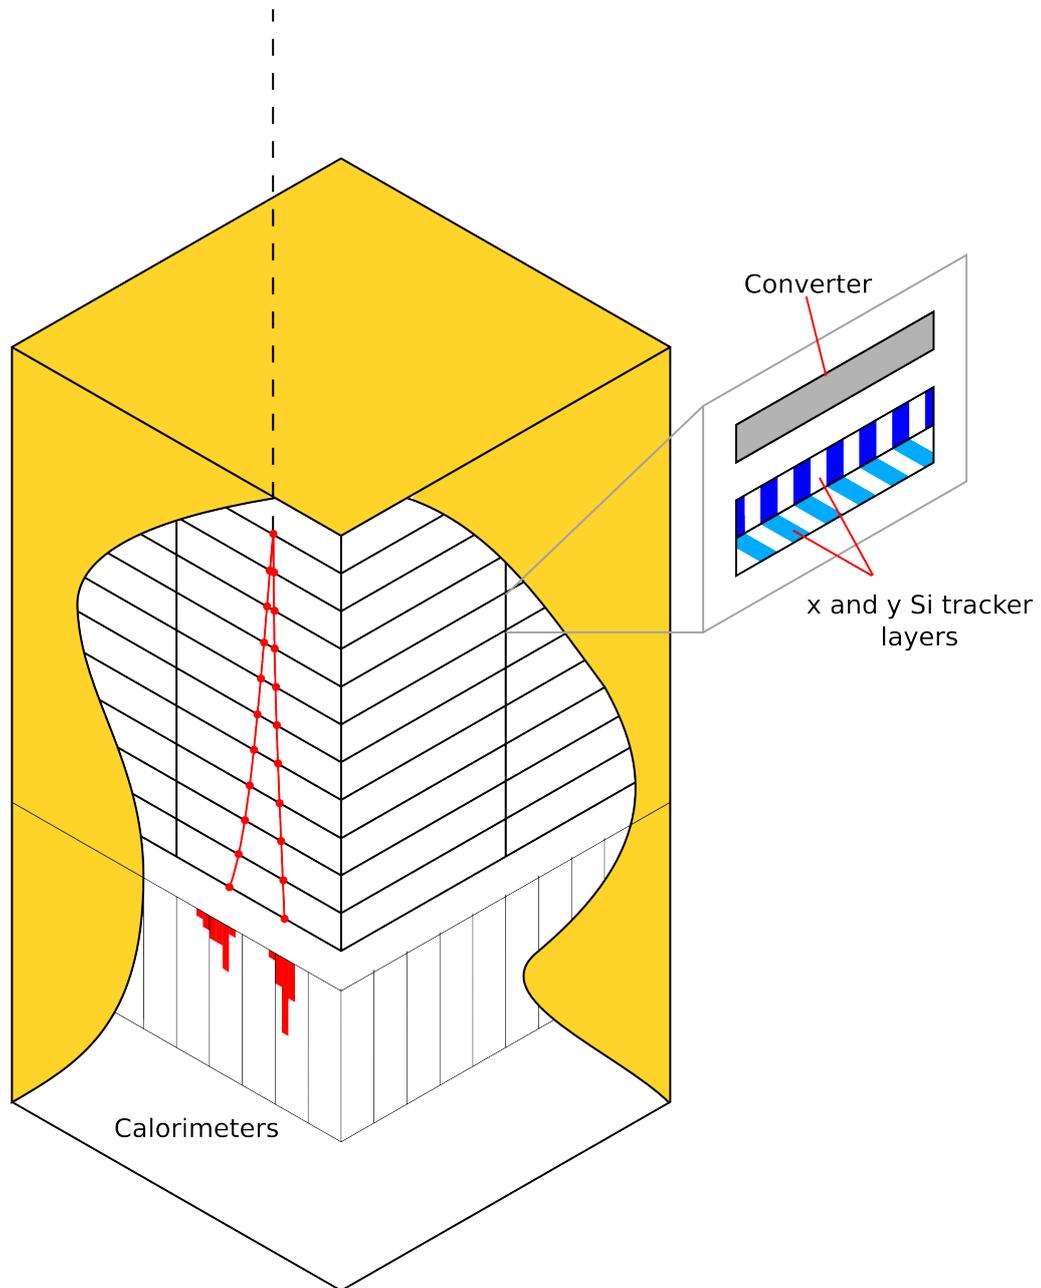


Figure 1.5: A sketch of the *Fermi*-LAT detector. The incoming dotted line represents a gamma-ray photon of energy sufficient to pair-produce an e^+e^- pair whose tracks shown as two red lines.

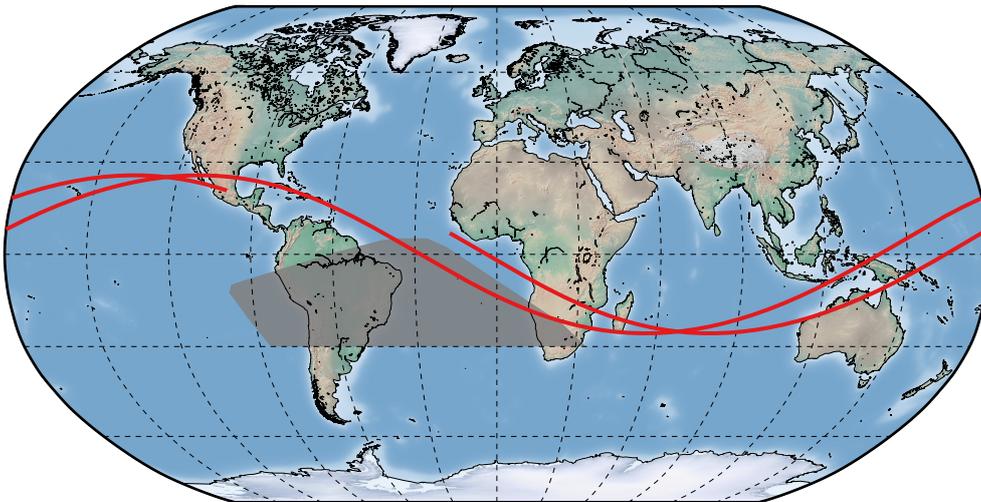


Figure 1.6: The orbit of *Fermi* from 2019-05-01 00:00:00–03:00:00 UTC, providing one exposure of the entire sky. The red line represents the projected position of *Fermi* onto the surface of the Earth in this period and the grey polygon represents the area in which the *Fermi*-LAT does not take data due to the SAA.

AGN.

1.2.4 *Fermi*-LAT Instrument Response

The “response” of the LAT instrument refers to the transformation between the data values (such as incoming direction or energy of a photon) we infer in our measurement and the real physical attributes of the incoming photon. In the case of an ideal detector the response would be simply be a delta function, but no instrument is perfect. There are 3 main sources of uncertainty in the LAT’s response which represent our systematic uncertainty in our results:

- Point Spread Function (PSF) – Given our inability to reflect or refract

gamma-rays, the LAT relies upon a pair conversion and subsequent tracking of electrons to measure the initial direction of an incident gamma-ray. This reconstruction is challenging, and has seen improvements even after the launch of *Fermi* (Atwood et al., 2013b). Each individual photon event has a different effective “PSF” associated with it depending on the scattering it undergoes through the detector, though as a population the “quality” of each reconstruction can assign a confidence region that can be approximated by a King function (also known as the Moffat function) as described by Read et al. (2011); Moffat (1969). The PSF was calibrated both pre-launch and on-orbit through observations of astronomical objects (Ackermann et al., 2013a; Portillo & Finkbeiner, 2014), and a plot of the PSF variation with energy can be found in Figure 1.7.

- Energy Dispersion – The energy reconstruction within the CsI calorimeters is uncertain as inevitably some energy is lost to scattering within the tracker/conversion layers before the electron pair reaches the calorimeters. Despite all corrections that can be made based on our knowledge of the tracks, there is still a $\sim 4\%$ uncertainty in energy even at the 1 GeV ideal energy. Due to the increased Compton scattering below 100 MeV the energy resolution more quickly degrades at lower energies, and difficulties at particularly high energies are related to the “leakage” of energy out of the back of the calorimeters.
- Effective Area – While the size of the detector is an almost trivial measurement at first glance, the effective area of the detector varies strongly with angle of incidence and energy as a result of the spatial extent of the detector and the penetrating ability of the photon. The variation of effective area with these parameters is plotted in

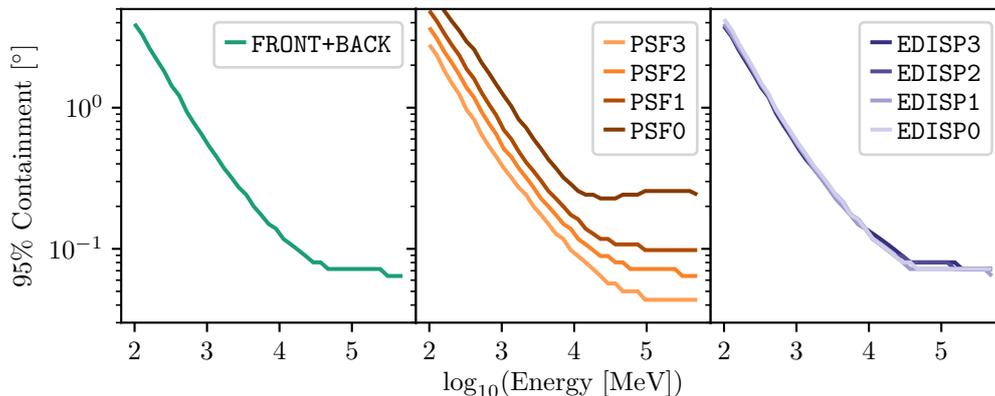


Figure 1.7: The variation of the 95% containment region of the *Fermi*-LAT PSF with energy. Specifically, the PSF-partitioned IRFs show an improvement from ~ 12 arcmin (PSF3) to just below 3 arcmin (PSF0 as a result of the varying reconstruction quality of the photon events). The EDISP classes show little variation in PSF, meaning that energy reconstruction and PSF reconstruction are not particularly interdependent. FRONT+BACK IRF results are shown for comparison. As the PSF varies depending on orbit, these values were calculated for a year-long exposure in the vicinity of NGC 1275 with weighted averaging over angles of incidence.

Figure 1.8.

1.2.5 Types of IRF

As alluded to in the previous Section, there are several different choices of IRF that can be used depending on the needs of a specific analysis. General analyses can use the physically motivated FRONT and BACK IRFs, which define different instrument response functions for photons converting in the front (with less dense converting material) and back (more dense converting material) of the LAT detector. Analyses which require good

FRONT IRF Effective Area @ 1 GeV

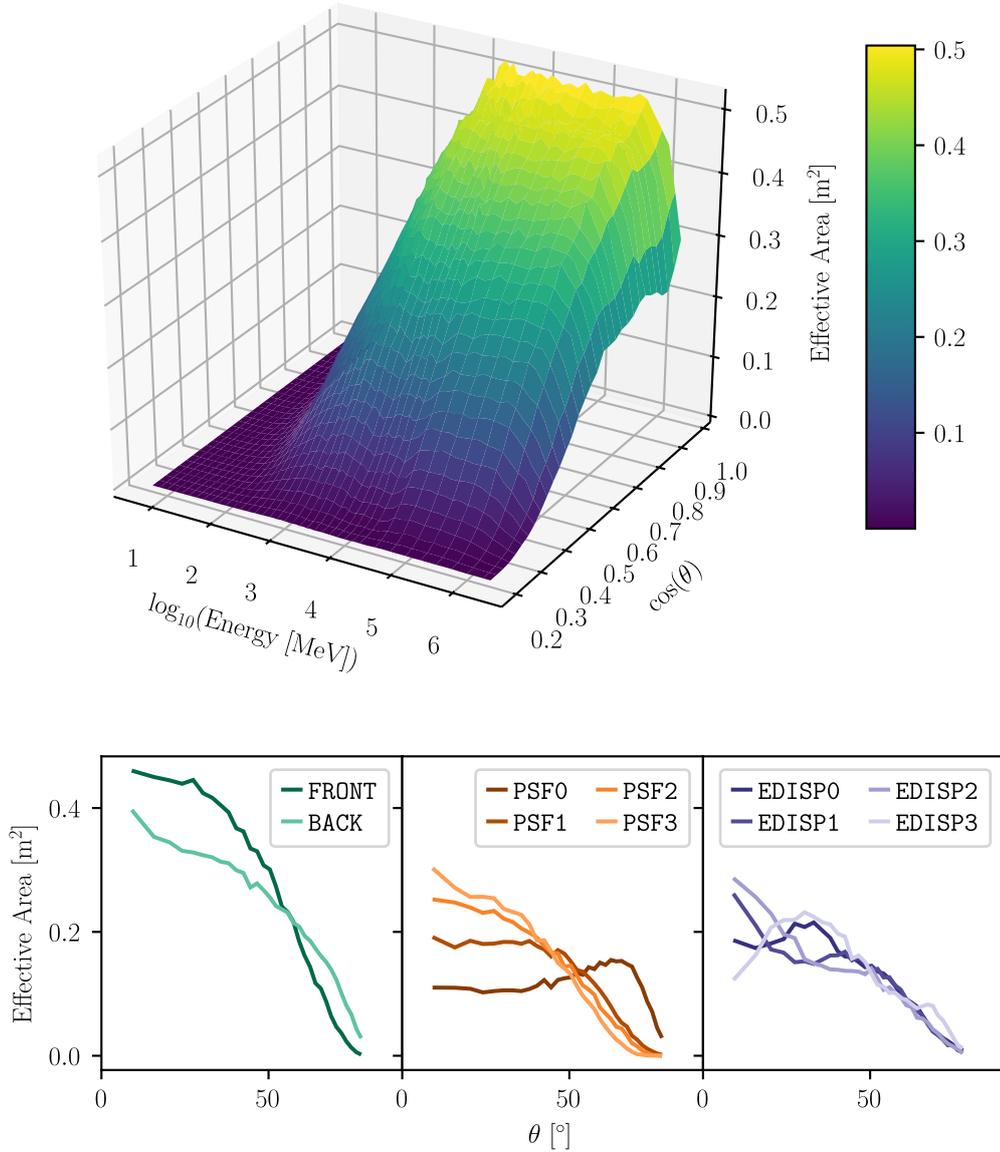


Figure 1.8: Top: FRONT effective area as a function of gamma-ray incidence angle and energy. Bottom: The on-axis effective area with differing type of IRF. The effective area of the PSF and EDISP classes is approximately half that of the FRONT- and BACK-conversion IRFs as they are partitioned into 4 quartiles as opposed to the physically-motivated IRF categories of front- and back-converting photons within the detector.

angular resolution (such as several analyses in this thesis) are more suited towards use of the PSF-partitioned IRFs (namely PSF0, PSF1, PSF2, PSF3 from best to worst angular resolution). Similarly, those analyses requiring good energy dispersion (e.g. Ajello et al., 2016) can use IRFs partitioned by the quality of the energy reconstruction with the EDISP classes.

The PSF- and EDISP-partitioned IRFs are determined through taking quartiles of photons ordered according to several variables resulting from the track reconstruction (in the case of the PSF) or the energy reconstruction (in the case of energy dispersion) which are included in the extended weekly files provided by the *Fermi*-LAT collaboration. An example of this partitioning with events is plotted in Figure 1.9 from one week of *Fermi*-LAT photons as a function of energy. This partitioning is a function of energy and incidence angle as with other IRF dependencies. These terms and IRF types will be referenced frequently when performing analysis in the later chapters of this thesis, and their significant differences are fundamental to the systematic uncertainty budget of the resultant outcomes.

1.2.6 Imaging Atmospheric Cherenkov Telescopes

As discussed in Section 1.2.2, IACTs use a novel approach to pair production from incident gamma-rays. Instead of using high- Z converter material, the atmosphere above a ground-based detector is used instead. Upon an interaction between a VHE gamma-ray and atmospheric particles, an electromagnetic shower results in cascades of relativistic hadrons or leptons. At high enough interaction energies, the daughter particles within the EM shower exceed the phase velocity of the speed of light in air, thus emitting Cherenkov^{***} radiation (Cherenkov, 1934; Frank & Tamm, 1937). The

^{***}As a note, this will be the chosen transliteration of Черенков. The reader may appreciate the variety of transliteration schemes available from Cyrillic to Latin script

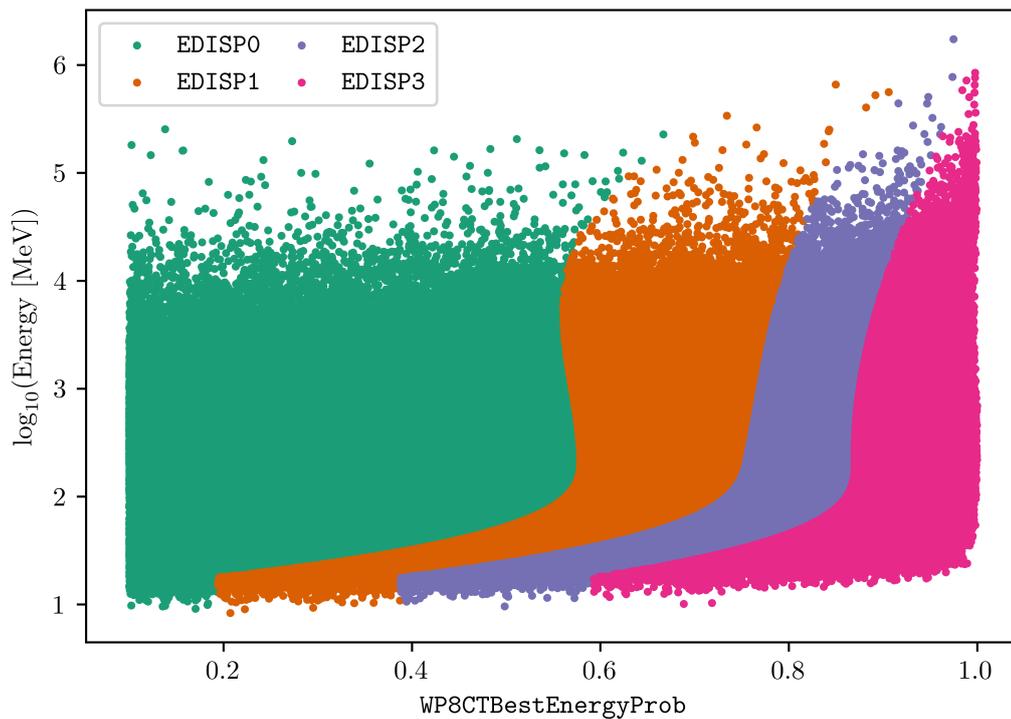


Figure 1.9: One week of photon events observed with the *Fermi*-LAT detector coloured according to the EDISP IRF class associated with it. $\text{WP8CTBestEnergyProb}$ is a parameter resulting from the energy reconstruction algorithms used by *Fermi*-LAT which enumerates the certainty of the reconstruction.

Cherenkov radiation produces a short (~ 10 ns) flash of light in a “pool” approximately 100 m in diameter. While this flash is bright, because of the short duration it is low fluence and requires short-integration cameras using photomultiplier tubes or Silicon photomultipliers as opposed to a traditional CCD solution.

The IACT technique was pioneered at the Whipple Observatory resulting in the first detection of an astronomical source (the Crab nebula) using this technique (Weekes et al., 1989). There are several such instruments currently in operation: H.E.S.S. (Hinton & the HESS Collaboration, 2004), VERITAS (Lessard, 1999), MAGIC (Aleksić et al., 2016b,a), and FACT (Anderhub et al., 2013), with the next-generation Cherenkov Telescope Array (Actis et al., 2011) under construction.

found here. “Cherenkov”, “Tscherenkov”, and “Čerenkóv” (amongst others) are all commonly used, though the prominent “Cerenkov” as an anglicisation of a transliteration is perhaps a stretch.

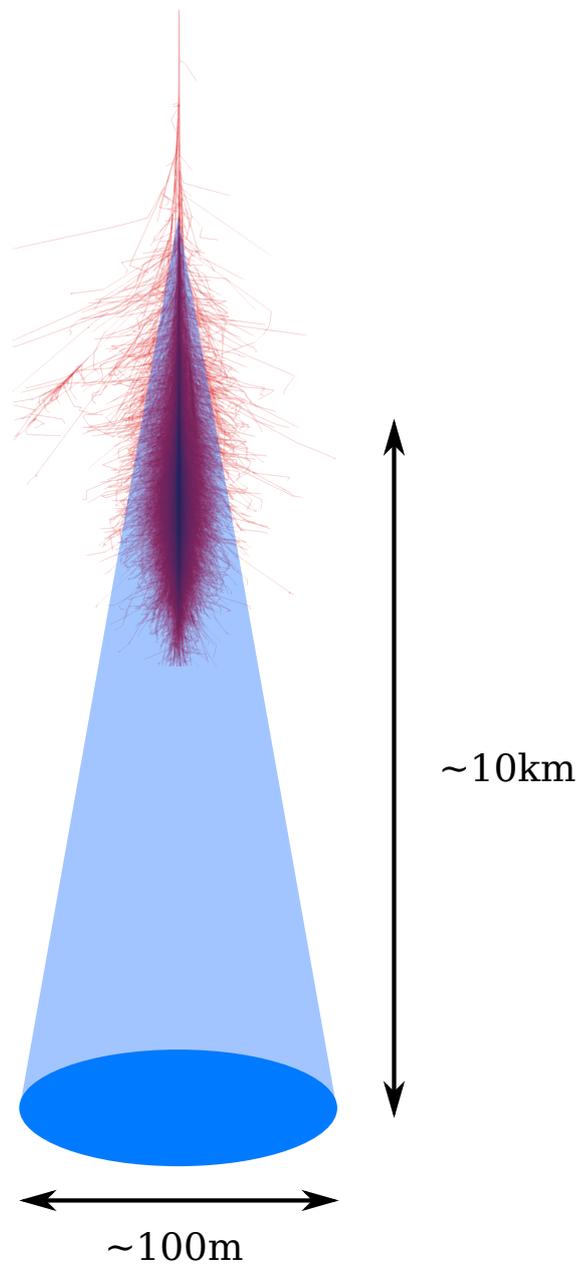


Figure 1.10: A sketch of the electromagnetic cascade resulting from a 1 TeV gamma-ray interacting in the atmosphere, and the resultant Cherenkov radiation as projected onto the ground. The red cascade is taken from simulations using Corsika (Heck et al., 1998) and the image provided by Johannes Knapp available at this URL: <https://www-zeuthen.desy.de/~jknapp/fs/photon-showers.html>.

Chapter 2

Astrophysical Modelling

2.1 Modelling

2.1.1 Overview

As we have explored in Section 1.1, the objects that are bright at gamma-ray energies tend to be bright all the way to the other extreme of the energy spectrum into the radio. This leaves a significant challenge for astronomers attempting to understand the entire picture of AGN emission as there are over 20 orders of magnitude between the lowest and highest energy photons we detect from these objects. Given the huge difference in scale it is perhaps surprising that we have any hope of modelling this emission over such a broad range, but the physical processes involved are thankfully well understood, even if the environments at the centre of the galaxies are far more mysterious. In Sections 2.1.2–2.1.6 we shall describe the relevant physics of AGN emission and how these relate to modelling with the *Fermi*-LAT instrument in Section 2.1.6.

2.1.2 Synchrotron Emission

The preponderance of highly energetic electrons and protons (cosmic rays) in the cosmos was discussed in Section 1.1.2, and their interaction with magnetic fields in the vicinity of AGN is a good starting point for deciphering the broadband emission problem. We discuss all of the physical processes in this section in the context of AGN jets (generally most luminous parts of an AGN at gamma-ray energies, although not always). It should be clarified what we mean when we discuss an AGN’s jet. Far from the “curious straight ray” first discussed by Curtis (1918), we now know that jets are formed close to the central engines of AGN where the magnetic field of the spinning black hole interacts with the magnetic field of the surrounding accretion disc. This results in a beam of collimated plasma which is initially magnetically dominated, which is transported phenomenal distances (kiloparsecs) during which energy is transferred from the electromagnetic field into the kinetic energy of the charged particles via shocks or magnetic reconnection events. The energy in the electrons and protons of the jet plasma is dissipated by several processes which we shall now explore.

2.1.3 Physical Processes

In the case of gamma-ray-detected AGN, the physical processes at work are radiation loss mechanisms from charged particles in magnetic fields. The dominant proportion of radiation at high energies results from electrons, with an uncertain contribution from protons. An AGN jet in its simplest terms is a sparse plasma accelerated within a magnetic field. In the conditions most commonly found in jets, processes such as free-free emission (Bremsstrahlung), ionisation, and adiabatic expansion take a backseat to the more dominant processes of synchrotron and inverse-Compton scatter-

ing when it comes to energy loss from electrons. The energy loss rate for synchrotron emission in the ultra-relativistic limit is:

$$-\left(\frac{dE}{dt}\right) = 2\sigma_T c U_{mag} \gamma^2 \sin^2 \alpha, \quad (2.1)$$

where σ_T is the Thompson cross section, c the speed of light, $U_{mag} = \frac{B^2}{2\mu_0}$ is the magnetic energy density, γ is the Lorentz factor and α is the pitch angle of the electron with respect to the magnetic field. With any irregularities in magnetic field, the pitch angles of electrons travelling in this field are quickly randomised, resulting in an isotropic distribution and the following average energy loss rate:

$$-\left(\frac{dE}{dt}\right) = \frac{4}{3}\sigma_T c U_{mag} \left(\frac{v}{c}\right)^2 \gamma^2. \quad (2.2)$$

The explicit frequency at which the relativistic electrons emit synchrotron radiation requires an involved calculation which is explored by Rybicki & Lightman (Chapter 6, 1986), where we find an electron radiates at its critical frequency $\nu = \gamma^2 \nu_g$, where ν_g is its classical gyrofrequency. When combined with a power-law distribution of electron energies as expected (from e.g. second-order Fermi acceleration) we find the synchrotron emission spectrum in the optically thin regime becomes a power law with index $a = (p - 1)/2$, where p is the power-law index of the electron spectrum.

Synchrotron radiation also results in the significant polarisation that is found in AGN jets, although this is somewhat reduced due to Faraday rotation in the jet plasma itself or along the line of sight to the detector. The jet itself is optically thick below the critical threshold described above, which is often found in the far-IR.

In the presence of low energy ambient photon fields such as the CMB, accretion disc emission or the jet's own synchrotron emission, ultrarelativistic electrons can upscatter photons through inverse Compton scattering (as

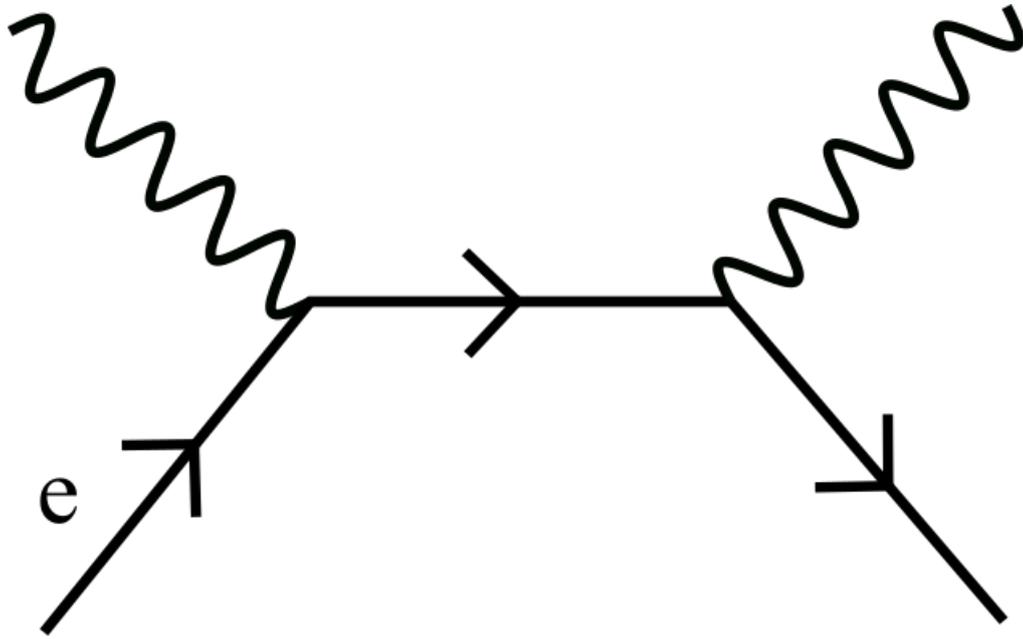


Figure 2.1: A Feynman diagram showing one of two first-order representations of the Compton scattering process.

illustrated in Figure 2.1). The frequency shift is proportional to γ^2 as the photon is Doppler boosted once to the rest frame of the electron and once more back to the original frame. When the energy of the photon and electron are similar and thus the recoil of the electron is significant, dampening the emissivity as the cross section becomes Klein-Nishina and correspondingly sets an upper limit to this process.

Finally, pair-production of electrons as previously discussed in the context of the detection of gamma-rays in Section 1.2.2 is also relevant within jets, especially in the interaction between high energy photons and ambient photon fields which can allow measurements of the cosmological history of the extragalactic background light, which shall be investigated in detail in Chapter 6.

2.1.4 Synchrotron Self-Compton Models

The most commonly applied model of gamma-ray emission in AGN is the single-zone synchrotron self-Compton (henceforth SSC) model. This model explains the double-humped (or “bactrian” as eloquently described by Blandford et al. (2018)) spectrum that is characteristic of AGN. In this formulation, the low frequency hump (peaking around the UV–X-ray) represents the synchrotron emission of high energy electrons in the jet which are relativistically boosted as the jet is collimated towards the observer. The higher energy hump (peaking around the GeV–TeV) then results from jet electrons upscattering the synchrotron photon field by inverse-Compton scattering.

As the inverse-Compton component is directly related to the synchrotron spectrum, some simple but useful relations between the components can be derived. For example, the ratio of the synchrotron and inverse-Compton luminosities can be written:

$$\frac{L_{IC}}{L_{synch}} \approx \frac{U_{synch}}{U_B} = \frac{L_{synch}}{4\pi R^2 c U_B \delta^4}, \quad (2.3)$$

where δ is the Doppler factor, R the size of the emission region, and the other quantities as are defined previously. This relationship allows the calculation of δ and B field from the observable luminosities, peak frequencies, and variability timescale (to calculate R). Full discussion of these quantities and the limits on SSC quantities are derived by Tavecchio et al. (1998). For illustrative purposes, Figure 2.1.4 shows the broadband structure of an SSC model, including an exponential cutoff due to interaction with the extragalactic background light generated using numerical code by Krawczynski et al. (2004). The attenuation due to EBL is calculated assuming a redshift of $z = 0.03$, using the model of Domínguez et al. (2011). This attenuation becomes significant at energies above ~ 100 GeV and imposes a sharp cutoff in the observed spectrum, even when the intrinsic SSC emission continues

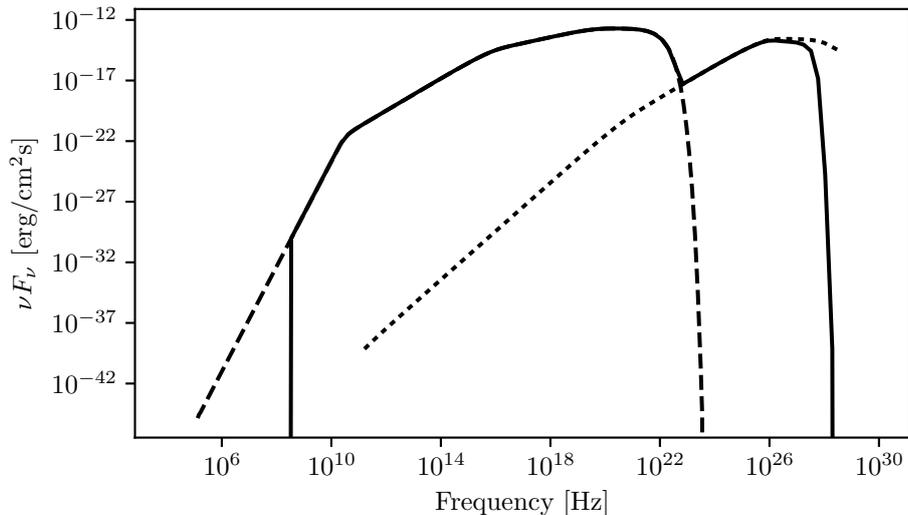


Figure 2.2: An instance of an SSC model with its components. Dashed lines represent the synchrotron flux, dotted lines represent the synchrotron self-Compton flux, and the solid line represents the summed flux after attenuation due to the extragalactic background light (EBL). The attenuation was applied using the EBL model of Domínguez et al. (2011), assuming a source at $z = 0.03$. The model used an electron energy spectrum with a broken power-law, a B field of 0.3G, and a jet oriented at 3° to the line of sight in order to provide a plausible model for a *Fermi*-detected BLL.

to rise. The double-humped structure is evident.

In reality, simple one-zone SSC models often provide accurate descriptions only when applied to the purest BL Lac spectra (e.g. in Mkn 421: Abdo et al., 2011). Other factors such as external Compton emission, where the electron synchrotron emission is not the only relevant ambient photon field, but strong emission from for example the broad line region or accretion disc as in FSRQs. Similarly, these models fail to properly account for VHE emission in objects with misaligned jets such as gamma-ray detected

radio galaxies like Cen A (Petropoulou et al., 2014). Further questions are raised at the highest energies, where flaring events result in significant TeV emission that could possibly require hadronic emission to explain.

2.1.5 Hadronic vs Leptonic Emission

Protons have not yet featured in our discussion which has only concerned the leptonic component of the jet plasma. With regard to the synchrotron emission of jet protons, synchrotron power scales as $1/m^2$ but this lack of synchrotron loss allows hadrons to reach far higher energies than their leptonic counterparts*. As such, proton synchrotron emission results in a peak frequency in the TeV energy range and so hadronic emission can be incredibly important at gamma-ray energies. The Klein-Nishina limit on Compton upscattering places an upper limit around the TeV range on the gamma-ray energies that can be emitted due to the leptonic component, emphasising the importance of hadronic emission. The main issue in differentiating between these two models is that the shape of the proton or muon synchrotron bumps (peaking in the gamma-rays) and the shape of the SSC inverse-Compton bump (roughly the same shape as that of the electron synchrotron bump, but higher in energy) are largely identical within the observational uncertainty in jet parameters (the expected values of emitting region radius and magnetic field). However the existence of photon-proton cascades produces a “bump” in the tail of the high-energy spectral bump, which may be detected with next-generation Cherenkov telescopes (Zech et al., 2017; Actis et al., 2011).

*This mass argument assumes that the dominant hadrons and leptons are protons and electrons respectively.

2.1.6 Phenomenological Models

Gamma-ray spectra of AGN are incredibly data-poor. The average AGN's integrated spectrum over years of observation (at least using space-based instruments with small collection areas) can contain a few hundred to a few thousand detected photon events that are able to be separated from the background. Furthermore, the parameter space for even the simplest physical models such as single-zone SSC emission is huge, given at least 8 physical parameters to fit. With such little information, it is necessary to simplify the physics we have discussed in this section to something more tractable.

One easy simplification with regard to the models describe above uses the fact that below and above the turnover in the synchrotron spectrum (and thus also in the IC spectrum), the emission follows a power law. Assuming that the observation is being made in an energy range that does not include the peak of a hump, a power-law is more than sufficient to model the data. Higher order curvature or features may exist on smaller scales, but this cannot currently be significantly detected in most AGN. In practice, this means that the standard method of measuring gamma-ray spectral energy density is in the form of a phenomenological model in the form:

$$\nu F_\nu = E^2 \frac{dN}{dE} = N_0 \left(\frac{E}{E_0} \right)^{-\gamma}, \quad (2.4)$$

where N_0 represents an overall flux normalisation, E_0 is an energy scaling parameter, and γ is the power-law index.

This simple model may be valid in the majority of cases not just in AGN, but also for many Galactic objects for similar (though not covered in the scope of this thesis) reasons - but as we have already identified a problem exists where the peak of the IC spectrum (in the case of gamma-ray spectra and an SSC model) falls within our measured energy range. Massaro

et al. (2004) showed that a log parabola gives an appropriate description of the resultant energy spectrum around this peak energy and also has direct relevance to the physical parameters of the jet. The model is parameterised as:

$$\frac{dN}{dE} = N_0 \left(\frac{E}{E_0} \right)^{-(\alpha + \beta \log(E/E_0))}, \quad (2.5)$$

where N_0 and E_0 are defined as with the power-law, α represents a power-law index and β the spectral curvature.

Further spectral features, such as interactions with the extragalactic background light which will be explored in Chapter 6, must be approached in a different manner. There are many processes which may result in an exponential attenuation to the two models described above due to effective optical depths between the source and observer. In addition, models of exotic physics such as dark matter predict “bumps” in a spectrum that can be modelled as Gaussian emission lines, or much more complicated cross sections. Finally, broken power-laws which are piecewise defined to allow two separate power-laws below and above a critical energy can be used to model breaks in, for example, the electron energy spectrum that produces the gamma-ray emission. This wide array of alternative phenomenological models will be examined more closely in Section 2.2.5. Most importantly, we shall now describe the methods by which we shall compare and evaluate these phenomenological models.

2.2 Statistics

2.2.1 Overview

Much of this thesis relies upon analysis of data from the *Fermi*-LAT instrument. In order to draw conclusions and determine the statistical certainty

of our hypotheses, we must be able to compare between different theoretical[†] models. Depending on the specific case in question, this can be straight forward or difficult and time consuming. The statistical methods that we use in this thesis will be detailed in this section, and where tractable we shall derive the result. I will also set down the conventions that I will use throughout this work to maximise consistency as far as is possible.

2.2.2 Astronomical Modelling

The reality of astronomy is that the image we receive on our computers (or on our plates, or our chart) is only an abstraction of the real physical object we are studying. In gamma-ray physics this is perhaps more obvious than at other wavelengths due to the poor resolution (by astronomical standards) that is achievable with our instruments. Simply looking at gamma-ray data is usually enough to convince you that the sky probably wouldn't look like that if we could see gamma-rays as we would with an optical telescope[‡]. The intent of modelling is to take the data we see through our telescope and deconvolve the instrument response from the “truth” (the physical object). Inevitably, this can only be an estimation of how the instrument will respond to any given photon, and the systematic uncertainties that result from this must be carefully considered and form a large part of this thesis. Figure 2.3 shows a sketch of this process.

2.2.3 Multivariate Normal and χ^2 Distributions

Starting from a far simpler and more familiar topic, the normal distribution (used interchangeably hereafter with the “Gaussian”), we'll quickly discuss

[†]Though in practice, phenomenological

[‡]e.g. With an instrument with a resolution $\sim \frac{1.22\lambda}{D}$

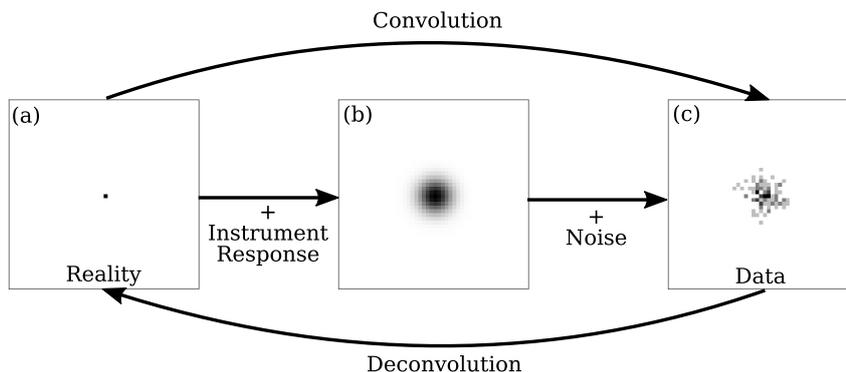


Figure 2.3: A sketch of the modelling process. (a) An idealised point source entirely contained within one pixel. (b) The point source convolved with a Gaussian point spread function (independently normalised) indicative of an instrument response. (c) The instrument’s response when sampled as a Poisson photon-counting process, resulting in the data measured with an instrument.

the χ^2 theorem and how it is relevant for maximum likelihood estimation (a topic we’ll address in Section 2.2.4). The χ^2 distribution is defined as such:

If G_1, \dots, G_k are identical and independently distributed standard normal random variables, the sum of their squares:

$$X = \sum_{i=1}^k G_i^2 \sim \chi_k^2, \quad (2.6)$$

where χ_k^2 [§] is the χ^2 distribution with k “degrees of freedom” (number of independent variables).

Sadly, the data we work with in astronomy are rarely distributed as ideal normal random variables. Firstly, these data are multi-variate. A standard normal distribution G has a probability density following:

[§]As χ_k^2 is actually a function with the parameter k , it would be technically correct to represent this as $\chi^2(k)$, but the dominant notation in the field is the former.

$$\text{p.d.f.}_G(x|\mu = 0, \sigma = 1) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}, \quad (2.7)$$

where x represents the random variable's value, μ the expectation of the distribution, and σ the standard deviation. The equation above in fact represents a general normal distribution, Y , which has been stretched by a factor σ and translated by μ .

In the multivariate case, however a k -dimensional multivariate normal distribution, X , has the probability density:

$$\text{p.d.f.}_X(\mathbf{x}|\mu, \Sigma) = \frac{1}{\sqrt{(2\pi)^k |\Sigma|}} e^{-\frac{1}{2}(\mathbf{x}-\mu)^T \Sigma^{-1}(\mathbf{x}-\mu)}, \quad (2.8)$$

where $\mathbf{x} = (x_1, \dots, x_k)^T$ is a normal random vector, μ is the k -dimensional expectation vector, and Σ is a $k \times k$ covariance matrix which can be represented as $\mathbf{S}\mathbf{S}^T$ [¶], and $|\Sigma|$ is the determinant of Σ . This also allows the representation $\mathbf{X} = \mathbf{S}\mathbf{G} + \mu$.

Given we have a multivariate normal distribution \mathbf{Y} in k dimensions distributed as $\mathbf{Y}_k \sim X(\mathbf{0}, \Sigma)$, then:

$$\mathbf{Y}^T \Sigma^{-1} \mathbf{Y} = \mathbf{G}^T \mathbf{S}^T \Sigma^{-1} \mathbf{S} \mathbf{G} = \mathbf{G}^T \mathbf{G} \sim \chi_k^2, \quad (2.9)$$

as in Equation 2.7. We have now shown that the “stretched” multivariate normal distribution can be “unstretched” to follow a χ_k^2 distribution.

This is where we hit upon our first possible issue with regards to our use of these distributions. Equation 2.8 is defined only when the covariance matrix Σ is non-singular[¶] to maintain unitarity. In cases where the covariance matrix is not positive-definite, the following results in this section will not apply. In Section 2.2.5 we will explore other methods necessary for model comparison in this case.

[¶]Via Cholesky decomposition when any matrix Σ is positive-definite symmetric

^{¶¶} $\text{Det}(\Sigma) \neq 0$

2.2.4 Maximum Likelihood and Wilks's Theorem

Given a model M with parameters Ω and a data set $D = \{d_1, \dots, d_k\}$ of k data points with probability P_i^{**} , the likelihood \mathcal{L} is defined as:

$$\mathcal{L}(D|\Omega) = \prod_{i=0}^k P_i(d_i|\Omega). \quad (2.10)$$

In simple terms, the likelihood is a measure of the probability of achieving this dataset given the model we are using. This is a useful metric when we wish to optimise our model parameters Ω . By finding the maximum likelihood estimate, the point in parameter space at which the likelihood is maximised, we find the most likely parameter values Ω given the data set:

$$\Omega_{\text{MLE}} = \underset{\Omega}{\operatorname{argmax}} \mathcal{L}(D|\Omega), \quad (2.11)$$

or alternately the maximum value of the log likelihood:

$$\Omega_{\text{MLE}} = \underset{\Omega}{\operatorname{argmin}} \log(\mathcal{L}(D|\Omega)). \quad (2.12)$$

But what is the probability distribution of the likelihood? How much does the value of the likelihood have to change before our parameters are “ruled out”?

Initially, we need to decide what models we are comparing. We must also take account of the number of parameters contained in the models being tested, and a suitable penalty for additional parameters must exist to prevent overfitting the data. A sketch of this problem is indicated in Figure 2.4 which highlights the folly of adding degrees of freedom without evidence.

We can define our two models as our *null* hypothesis, H_0 , and *alternative* hypothesis, H_1 , with parameters θ_0 and θ_1 in H_0 and H_1 respectively.

**This probability is a function of the parameters used and the measurement uncertainty of the data, resulting in a different probability distribution for each point.

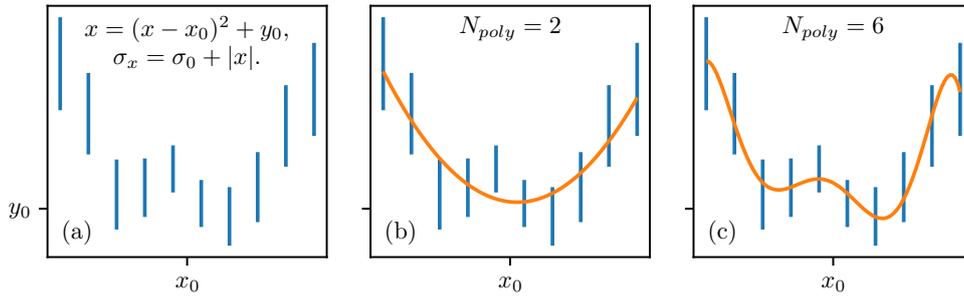


Figure 2.4: An example of overfitting data. (a) A simple quadratic model is used to generate data points with Gaussian errors scaling linearly with $|x|$. (b) The data are fitted with a suitable quadratic model. (c) A sixth-order polynomial is instead used. The data are closer to the model, but the result is clearly absurd.

Our test is required to distinguish the case where the parameters θ_0 is a subset of parameter space $\Theta^{\dagger\dagger}$, and the alternative hypothesis has θ_1 in the complement of θ_0 .

To find the behaviour of this function at its minimum, a second order Taylor approximation to the log likelihood function l around the maximum likelihood estimate in 1D:

$$l(\theta) \approx l(\theta_{\text{MLE}}) + \frac{dl(\theta_{\text{MLE}})}{d\theta}(\theta - \theta_{\text{MLE}}) + \frac{1}{2} \frac{d^2l(\theta_{\text{MLE}})}{d\theta^2}(\theta - \theta_{\text{MLE}})^2, \quad (2.13)$$

noting that $\frac{dl(\theta_{\text{MLE}})}{d\theta} = 0$ at the point of maximum likelihood. Thus:

$$l(\theta) - l(\theta_{\text{MLE}}) = \frac{1}{2} \frac{d^2l(\theta_{\text{MLE}})}{d\theta^2}(\theta - \theta_{\text{MLE}})^2. \quad (2.14)$$

Most importantly this shows that $2(l(\theta) - l(\theta_{\text{MLE}}))$ is distributed as the square of a normally distributed random variable $(\theta - \theta_{\text{MLE}})$. By definition (Equation 2.7), this product of two normal random variables is distributed as χ^2 and this correspondence between composite models is known

^{††}A confidence region in parameter space

as Wilks’s theorem. The proof that the quantity $(\theta - \theta_{\text{MLE}})$ is asymptotically normally distributed relies on the central limit theorem, and is succinctly explained by e.g. Casella & Berger (2001). The fact that these models are “nested” is fundamental, which means that the alternative model can be reduced to the null hypothesis with certain parameter values. The original treatment by Wilks (1938) is of course the origin of this theorem’s name, but the mathematics is far less friendly.

To state this theorem explicitly: *as the sample size $n \rightarrow \infty$, the test statistic $-2\log(\Lambda) = -2\log\left(\frac{\mathcal{L}(\boldsymbol{\theta}_0|x)}{\mathcal{L}(\boldsymbol{\theta}_1|x)}\right)$ for a nested model will be asymptotically χ^2 -distributed with degrees of freedom equal to $\dim(\boldsymbol{\theta}_1) - \dim(\boldsymbol{\theta}_0)$, when H_0 holds true.*

This distribution now allows the calculation of the significance in composite model comparisons as shown in Figure 2.5, which neatly shows the penalty paid for additional degrees of freedom when the above calculation is expanded to multiple parameters^{‡‡}.

2.2.5 Wilks’s Theorem in Practice

There are many phenomenological models used for *Fermi*-LAT analysis, which I have summarised and defined in Appendix C. The most relevant for use in this thesis are constant valued functions (CV), power-laws (PL), log parabolae (LP), smoothly-broken power-laws (S-BPL) and power-laws with exponential cutoffs (PLEC). Now that we have the ability to compare between certain (nested) models as a result of Section 2.2.5, how does this help in our specific case? Using *Fermi*-LAT data uses a maximum likelihood method^{§§} through which a model of the sky is created. The log likelihood is

^{‡‡}which also holds once nuisance parameters are taken account of, something not discussed in our treatment

^{§§}With the choice of a binned or unbinned likelihood

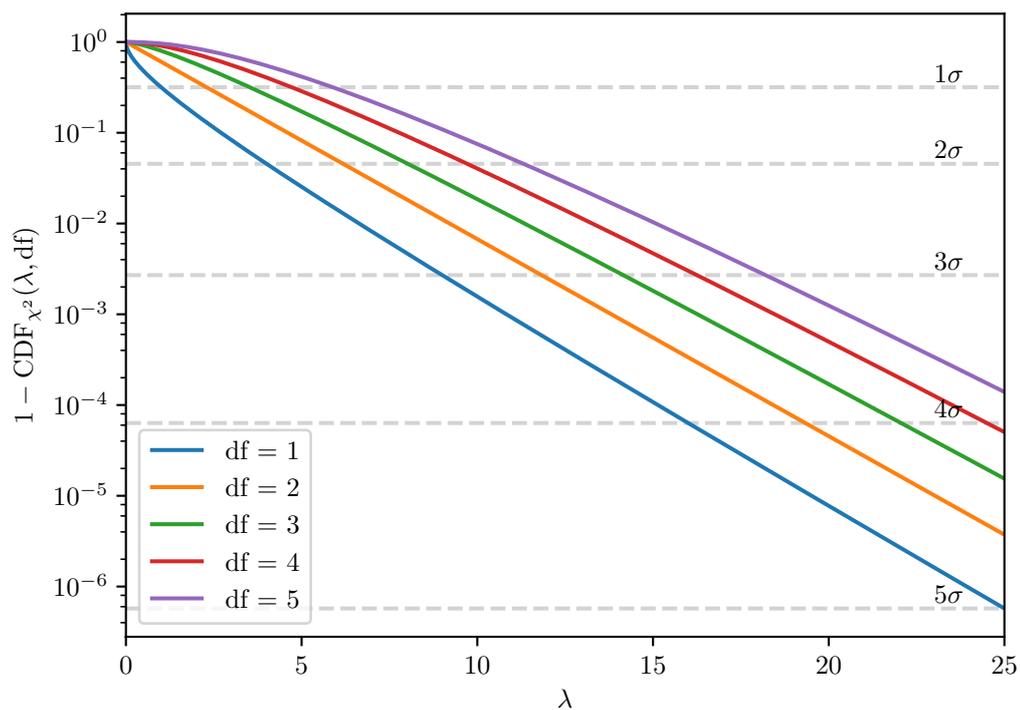


Figure 2.5: The reverse cumulative density function for a χ^2 -distributed variable, λ , with degrees of freedom, df, between 1 and 5. Dashed horizontal lines show the probabilities relevant for the sigma values often used as a reference point in this thesis.

minimised within this software using standard minimisation algorithms (e.g. MINUIT, James & Roos (1975)) after models are convolved with measured instrument response functions.

There are several inherent issues with this choice of analysis as the model parameters do not quite meet the specifications for the application of Wilks’s theorem. Both spatial (right ascension, declination, extension) and spectral (power law index, curvature indices, cutoff values) parameters are optimised using the *Fermi* Science Tools^{¶¶}. However there are several physical limitations that restrict our parameter space in an *a priori* fashion, for example the fact that the parameter space for spectral normalisations must be positive (there is no such thing as a gamma-ray sink) and that the two position parameters are degenerate in the null hypothesis^{***}.

While these issues persist, Mattox et al. (1996) instead performed Monte Carlo simulations of *EGRET*^{†††} data showing that despite these awkward boundary conditions, the test statistic, $TS = -2(\log(\mathcal{L}_0) - \log(\mathcal{L}_1))$ is distributed as a χ_n^2 with n degrees of freedom as is appropriate for simple power law models. This statistical basis is now used for *Fermi*-LAT analysis, though the assumption that the TS is distributed as a χ^2 is often made for models that are not nearly so well-studied. Taking the models discussed earlier in Section 2.1.6, Figure 2.6 shows a breakdown of which commonly used models are nested within each other and can be compared using Wilks’s theorem and the likelihood ratio test.

While we have noted some of the shortcomings of this method of analysis, Wilks’s theorem and maximum likelihood estimation has proven robust for point-source analysis using the *Fermi*-LAT instrument. Those shortcomings

^{¶¶}<https://fermi.gsfc.nasa.gov/ssc/data/analysis/software/>

^{***} H_0 : The source does not exist

^{†††}The “Energetic Gamma Ray Experiment Telescope”

	CV	PL	LP	PL- SEC	S- BPL
CV		1	2	3	4*
PL			1	2	3*
LP				X	X
PL- SEC					X
S- BPL					

Figure 2.6: Diagram of the number of parameters needed to compare between different models. The diagonal is greyed out due to the self-referential nature of these comparisons and due to symmetry only the top independent values have been listed. An X corresponds to those models which are not nested whatsoever. A number in a cell represents the number of degrees of freedom are required to compare the model in the top row to the model in the first column. An asterisk highlights the issue of parameter degeneracies within the SBPL model where if two power law indices are identical, the break energy becomes entirely degenerate with infinite covariance, breaking the positive-definiteness required for Wilks's theorem to hold.

can as well be solved using other statistical methods (e.g. the Akaike information criterion, Akaike (1974)) or the brute-force Monte Carlo method as used by Mattox et al. (1996). In fact we will make use of both of these extensions in this thesis.

Spectral Breaks in Centaurus A

3.1 Introduction

Centaurus A (Cen A) is a giant radio galaxy at a distance of approximately 3.7 Mpc (Saviane & Jerjen, 2007) with a central black hole of mass $\sim 5 \times 10^7 M_{\odot}$ (Cappellari et al., 2009; Neumayer, 2010) and is the second brightest misaligned AGN (MAGN) within the 3FGL (Acero et al., 2015) after NGC 1275. In the optical waveband, Cen A is the fifth-brightest galaxy in the sky, noted for its prominent obscured dust lane and evidence of a prior galaxy merger (Israel, 1998).

Data from other wavelengths is vital in untangling the complicated structure of Cen A, which can be seen from Figure 1.2. The radio-emitting core of Cen A has been constrained below a size of 10^{-2} pc (Horiuchi et al., 2006), and the subparsec jet-counter-jet system whose motion has been mapped to high precision over long periods of time, whose radio knots and X-ray knots have both been investigated by Hardcastle et al. (2003) and Müller et al. (2014). This allowed measurements of jet components with apparent speeds of 0.1–0.3 c and a constraint of the jet angle to the line of sight to $12^{\circ} < \theta < 45^{\circ}$ through the jet to counter-jet ratio. Infrared

parsec-scale measurements of Cen A's core have found evidence for a central synchrotron source separate from the thermal dust emission present in the disc (Brookes et al., 2006; Meisenheimer et al., 2007). On yet larger scales, the giant radio lobes of Cen A stretch out up to 500 kpc in diameter with an age of 10–80 Myr resulting from a population of relativistic particles emitting synchrotron radiation (and inverse-Compton scattering at high energies) (McKinley et al., 2013). Modelling of these giant lobes includes a precessing central core having undergone several separate outbursts (Morganti et al., 1999) which interact with the surrounding medium, although alternative models do exist (Gopal-Krishna & Wiita, 2010).

Cen A is classified as a Fanaroff & Riley type I radio source, corresponding to a BL Lac object observed at large angles to the jet (Urry & Padovani, 1995) as we described in Section 1.1.3, which is supported by the viewing angles measured by Müller et al. (2014). As discussed in Section 2.1, BL Lac objects can be modelled using SSC models and despite the large viewing angles, Cen A is no different. Broadband spectral modelling based on simple SSC, however, fails to reproduce the observed flux above a few GeV in *Fermi*-LAT data (Chiaberge et al., 2001; Petropoulou et al., 2014; Sahakyan et al., 2013), and is further disfavoured when taking into account more recent H.E.S.S. observations (Aharonian et al., 2009; H. E. S. S. Collaboration et al., 2018). Lacroix et al. (2015) suggested that the hearts of MAGN such as Cen A could be caused by annihilation from relic WIMP dark matter, though even with a spiky dark matter profile (Gondolo & Silk, 1999) a velocity-dependent annihilation cross section would be required. While M87 is an object with extraordinary variability in the gamma-rays (Aharonian et al., 2006), Cen A shows no significant variability whatsoever in the 3FGL (Acero et al., 2015), although this misses several years of data before the release of 4FGL light curves.

With the advent of the Pass 8 data release by the *Fermi*-LAT collaboration Atwood et al. (2013a) in mid-2015, we looked to investigate Cen A to much greater statistical certainties and higher energies than could be achieved previously by Sahakyan et al. (2013) and *Fermi* Collaboration et al. (2010). The results achieved during this study were published in Brown et al. (2017), though this thesis contains up-to-date information and far greater detail into the specifics of the data analysis with regard to variability and quantification of systematic uncertainty.

3.2 *Fermi*-LAT Data Selection, Reduction & Initial Analysis

We use 8 years of *Fermi*-LAT Pass 8 data accrued since August 2008, extracting all photon events within a 15° radius around Cen A Core's catalog position (RA, Dec: 201.37° , -43.03° ; Acero et al., 2015) between 100 MeV and 300 GeV. This region safely encompasses Cen A's core emission, the extended radio lobes, and all nearby bright point sources, although Cen A is still by far the brightest source in the field of view. We used a summed likelihood analysis to take advantage of the PSF-partitioned instrument response functions within the *Fermi* Science Tools in which PSF quartiles were reduced separately but parameter values maintained parity. The event class chosen for this analysis is the **SOURCE** class defined by the *Fermi* collaboration for studies on point sources, which avoids the effects of misclassified cosmic ray events present in the **TRANSIENT** classes. Further cuts are made to the zenith angle of the photons used in this analysis to $< 90^\circ$ such that gamma-rays resulting from cosmic ray interactions with the Earth limb are excluded from our dataset, and our criterion

“DATA_QUAL>0&&LAT_CONFIG==1” ensures that the LAT is in data-taking mode outside of the SAA at the times of our analysis.

The initial likelihood analysis used the 3FGL point source catalogue as a seed point source model, importing the spectral models and associated parameters for the sources within 25° of Cen A, which extends 10° beyond our data in order to take into account bright sources outside of our RoI but which still possibly contribute photon flux. We use a likelihood analysis which is log binned in energy such that we have 10 bins per decade in energy. The output of the “sky model” is found by calculating the “livetime” (pointing time) and exposure (livetime convolved with the effective area as a function of time) of the LAT across our RoI, and thus producing a mapping from our sources’ spatial extent when convolved with the instrument response functions and finding the expected counts map for our model parameters. This “model map” formed by adding the contributions of each source in our model in the relevant energy band is then directly compared to the binned data to calculate the likelihood assuming Poisson statistics. The likelihood minimization over our sources’ spectral parameters itself is performed using the MINUIT minimizer (James & Roos, 1975). No spatial parameters (e.g. the RA and Dec of sources) are freed in this process, primarily because each source has a $>4\sigma$ detection from the 3FGL usually with a counterpart at other wavelengths, but also because the *Fermi*-LAT tools do not support this (for very good reason).

The spectral parameters of the sources in our model are iteratively freed in order of significance (Cen A being the brightest point source in the RoI, before moving on to dimmer objects) to find an approximate minimum before a simultaneous fit is performed with all parameters free. The parameters of objects outside the RoI are however never freed in this process and instead fixed to the catalogue values. A model of Galactic interstel-

lar emission, `gll_iem_v06.fit` as provided by Acero et al. (2016) is used to compensate for Galactic flux, for which only the normalisation is freed. Similarly, models of isotropic emission, `iso_P8R2_SOURCE_V6_PSF0-3.txt`, are used to correspond with the expected isotropic background emission from unresolved populations of galaxies and misclassified cosmic ray events specified for our combination of chosen IRFs again whose spectral shape is not allowed deviate. Furthermore, the giant radio lobes of Cen A are a significant diffuse source of gamma-rays, and the spatial model created by Abdo et al. (2010a) using the template of Ka-band WMAP observations of Cen A's lobes with a 1° cut-out to avoid contamination by Cen A's core emission is used for our analysis. The spectral model used for the lobes is that of a simple power-law as found to be sufficient by Abdo et al. (2010a). At this stage, we remove point sources for which the likelihood ratio test statistic (henceforth TS) is < 1 , as these models did not contribute whatsoever to the overall photon count, but did provide extra degrees of freedom which were essentially degenerate.

Once this initial minimisation has occurred, we create a TS map whereby a grid of positions 0.1° apart across the RoI is created and a test source is placed at each node of the grid. The significance of a point source with an assumed power-law index of 2.0 is calculated using Wilks's theorem by comparing the optimised log likelihood values before and after removing the point source, creating an image whose intensity is proportional to the significance of a point source not included in our sky model. A new point source is introduced in our sky model if the peak of our TS map is $> 5\sigma$, and the optimisation described in the previous paragraph is re-run. This process of optimising model parameters, creating TS maps, and adding new point sources is repeated until there are no peaks in our TS map representing a point source with a significance greater than 5σ , giving us some sense of

certainty that we have accounted for all point sources in our RoI that could misattribute flux to another object in our model.

In the immediate surrounding of Cen A we find 3 new point sources, 2 of which are plotted in the sky region in Figure 3.1. The new sources result from the greater integration time of our analysis compared to the analysis of the 3FGL catalogue. Intriguingly, 2 of these point sources are coincident or bordering the radio lobe model, whose extent is shown, suggesting the spatial distribution of flux in the WMAP template is insufficient to map the electron population at gamma-ray energies, but given the radically different electron cooling timescales at microwave and gamma-ray energies this is perhaps unsurprising. Thankfully, these point sources are several degrees away from our source of interest (Cen A's core) and the covariance between the parameters of these new point sources and the core emission is effectively zero and comparable to that of the two most equidistant sources within our RoI.

Due to the proximity (or indeed overlap) of the core and lobe models, we plot the likelihood profile as a function of the core and lobe normalisations (with the spectral shape of each optimised in each bin). The angle of the major axis of this error ellipse in Figure 3.2 is indicative of some covariance between the two models, which will be taken into consideration when detailed spectral modelling is performed.

3.3 Centaurus A's Core Spectrum

Our initial analysis presented in the previous section is satisfactory to begin detailed study of Cen A's core itself, making note of the complicated situation due to the interplay of the core and lobes also faced by other analyses of this region (*Fermi* Collaboration et al., 2010; Yang et al., 2012;

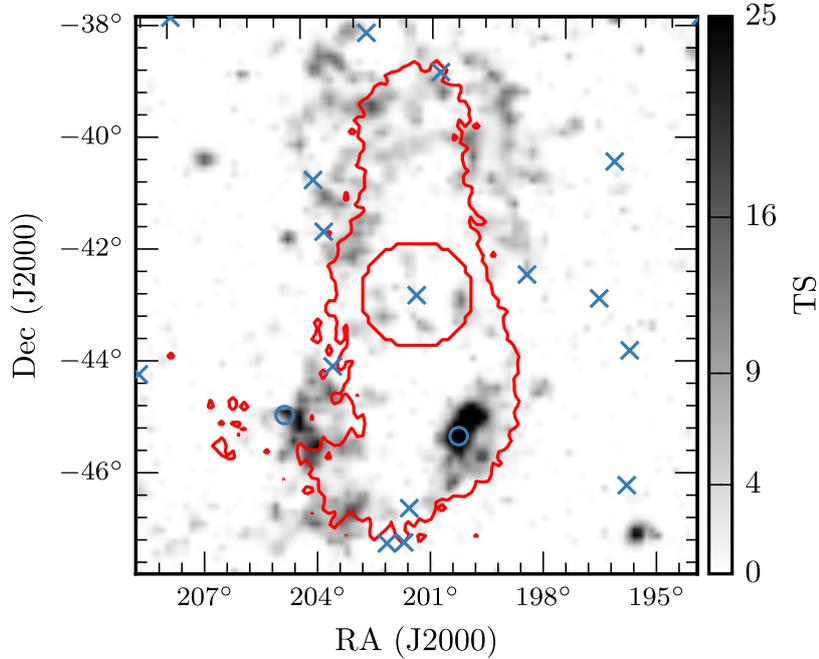


Figure 3.1: The TS map of our RoI and the sources contained within. The red contours represent the boundary of the extended template used for the spatial map of the giant radio lobes, and blue crosses represent point sources contained within the RoI, including Cen A in the very centre. Blue circles represent two of the three point sources discovered within this analysis, with the third being outside of the bounds of this image. As noted previously, excess significance (photon counts that are unaccounted for within our model) is present on the edges and within the large-scale structure of the lobes.

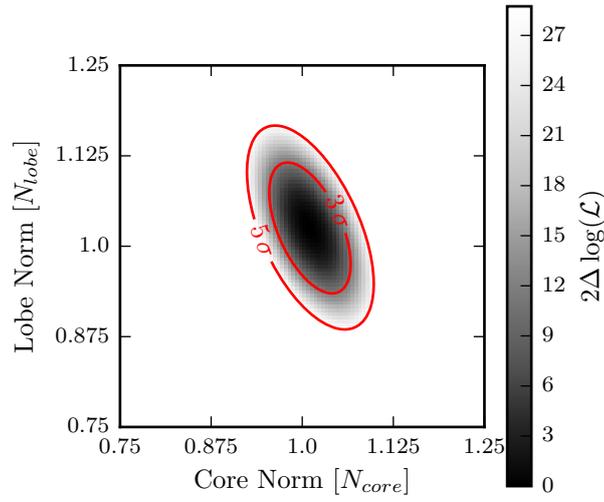


Figure 3.2: The 2D likelihood profile as a function of the core and lobe model normalisations in Cen A. The contours represent $\Delta \log$ likelihood values of 11.829 and 28.744 corresponding to the 3σ and 5σ uncertainty regions respectively. In each bin all other parameters in the model are optimised.

Sahakyan et al., 2013, 2018; H.E.S.S. Collaboration et al., 2018). The analysis described in Section 3.2 is binned across the whole energy range, and so the resultant model parameters are optimised across this entire range. The 3FGL assumes a simple power-law spectrum of the form:

$$\frac{dN}{dE} = N_0 \times \left(\frac{E}{E_0} \right)^{-\gamma} \quad (3.1)$$

where N_0 is a normalization factor, E_0 an arbitrary scale energy fixed at 1 GeV such that parameters are of order ~ 1 to aid the minimization routine, and γ is the power-law index. In order to assess the quality of the fit of a simple power-law, we re-analyse the data in 9 discrete energy bins from 300 MeV to 500 GeV (but still maintaining the 10 bins per decade in energy in their individual binned likelihood analyses) and allow the normalisation of Cen A's core to vary, keeping background sources constant. This results in the data shown in Figure 3.3. Evidently, the power-law model used does not seem to fit the data satisfactorily, with an upturn present at around the GeV scale.

Due to the obvious shortcomings of the power-law model in this case, we look to fit this data with a piecewise broken power-law model defined as:

$$\frac{dN}{dE} = N_0 \times \begin{cases} \left(\frac{E}{E_b} \right)^{\gamma_1} & \text{if } E < E_b, \\ \left(\frac{E}{E_b} \right)^{\gamma_2} & \text{otherwise.} \end{cases} \quad (3.2)$$

Where N_0 is defined as above, E_b is the break energy (now a parameter and not fixed as with E_0), γ_1 and γ_2 are spectral indices above and below the break energy.

Fitting this new model, we find a $E_b = 2.6 \pm 0.3$ GeV, with $\gamma_1 = 2.73 \pm 0.02$ and $\gamma_2 = 2.29 \pm 0.07$, and is plotted in Figure 3.4. A log likelihood ratio test comparing the power-law and broken power-law models produces a TS of 28.6, although as noted in Section 2.2.5, these models cannot be compared

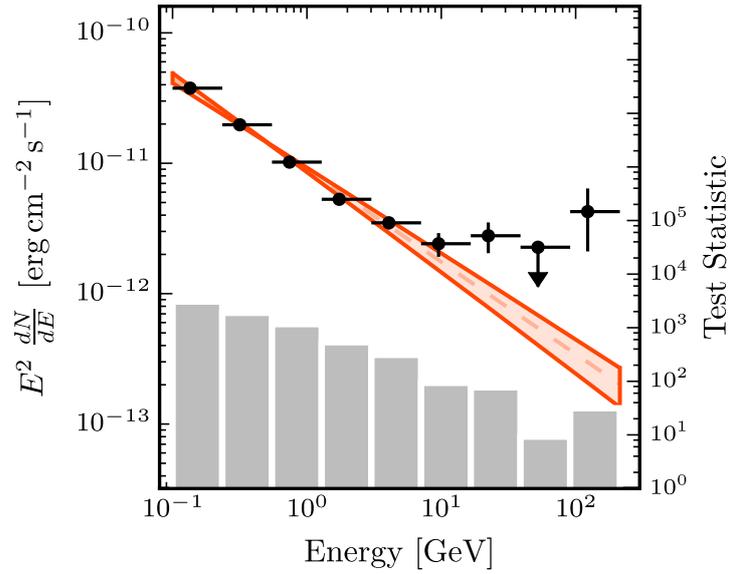


Figure 3.3: The spectral energy distribution (SED) of Centaurus A’s core fitted with a power-law spectrum. The red shaded area represents the 1σ confidence region resulting from the uncertainties on a power-law model fit to this data, with the dashed red line representing the best-fit parameters. The black data points represent the flux predicted in each bin when allowing Cen A’s normalization to vary independently in each energy bin independently. The grey histogram represents the TS of Cen A in each energy bin analysed.

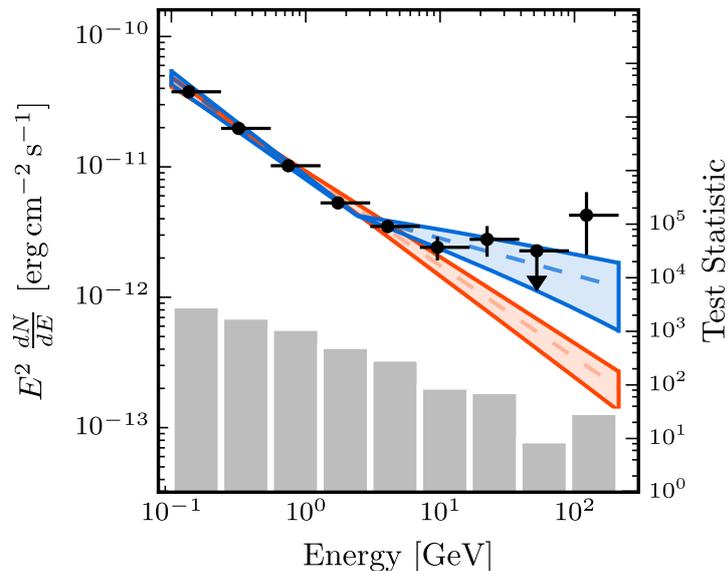


Figure 3.4: The SED of Centaurus A from Figure 3.3 with the additional broken power-law model shown in blue.

in a simple manner using Wilks’s theorem. A simple understanding of this can be gained when noticing that were a model indeed a perfect power-law, $\gamma_1 = \gamma_2$, the parameters E_b and N_0 become entirely degenerate and thus breaching the assumptions of Wilks’s theorem. We shall explore the possible significance of this value in Section 3.4.

3.4 Model Comparison

As Wilks’s theorem is inapplicable in this case, we must instead find another way to understand the distribution of the likelihood ratio test statistic. There are several common methods of comparing models in such cases, and we will begin with the Akaike Information Criterion (AIC; Akaike, 1974).

The AIC, defined as:

$$\text{AIC} = 2k - 2 \ln(\mathcal{L}), \quad (3.3)$$

where k is the number of parameters optimised in the model, and \mathcal{L} is the maximum likelihood value of this model fit. The relevant quantity for the comparison of two models is the difference in AIC (ΔAIC), or more usefully $\exp(\Delta\text{AIC}/2)$ is proportional to the probability that the alternative model minimizes the “information loss” compared to the “true” model, and is sometimes referred to as the relative likelihood. This criterion can be seen as penalising the addition of degrees of freedom without gaining a substantially improved fit. A similarly defined criterion, the Bayes Information Criterion (BIC; Schwarz, 1978):

$$\text{BIC} = \ln(n)k - 2 \ln(\mathcal{L}), \quad (3.4)$$

where n is the number of data points or observations is included in addition to k and \mathcal{L} . There is substantial discussion as to when each of these criteria is the most suitable for use (Vrieze, 2012; Aho et al., 2014), but in cases where the “true” model is not amongst the options (such as our phenomenological models), the AIC is preferable.

Furthermore, the AIC has been shown to be inaccurate when the sample size is small, and incorrectly penalising extra degrees of freedom. The AIC_c (the subscript c represents a “correction”) can be used in this case, though its penalty term differs depending on the distribution of residuals and linearity of parameters. Thankfully, Cen A is an particularly bright gamma-ray object and we are very much in the asymptotic regime of the AIC.

Applying the AIC to the model comparison discussed in the previous section, we find a $\Delta\text{AIC} = 26.6$ in preference of a broken power-law, which is interpreted that a power-law is 1.67×10^{-6} times as probable as a broken power-law to minimize the information loss, i.e. a $\sim 5.5\sigma$ preference.

We can also attempt to establish the significance of this preference by Monte Carlo simulation of our RoI to calibrate the TS numerically. This has been approached for model comparisons for *Fermi*-LAT in catalogs of gamma-ray bursts (Ackermann et al., 2013c) which found deviations from simple χ^2 assumptions, and was also initially calculated for likelihood analysis of EGRET data by Mattox et al. (1996) which found that in the case of simple power law sources, χ^2 assumptions generally hold. In this case we simulate a power-law model approximating that of Cen A (a power-law index of 2.6 and comparable normalisation) with the Cen A lobe model and Galactic interstellar emission model as backgrounds. Other sources in the immediate ROI were not simulated, as Cen A is the brightest source in the region (along with the lobes) by some margin. The distribution of the TS was thus calibrated and in Figure 3.5 we show that this method of fitting at least approximately follows the expected relationships. Leaving all parameters free (including the break energy) leads more often to complete failure of the fit due to the degeneracy, preventing this equivalent simulation from being of use.

Finally, we can also choose an alternative phenomenological model which is nested and contains an upturn. The obvious method to fit in this case is the log parabola:

$$\frac{dN}{dE} = N_0 \left(\frac{E}{E_0} \right)^{-(\alpha + \beta \log(E/E_0))}, \quad (3.5)$$

where N_0 and E_0 are defined as in Section 3.2, α is effectively a power law index, and β represents the curvature in the spectrum. A negative value of β would result in an upwards curve to the fitted model, which we would expect given the data. Indeed, fitting with this model provides a value of $\beta = -0.04 \pm 0.01$, at $\text{TS} = 16.64$ ($\sim 4.1\sigma$). The downside of fitting with this model is that extrapolating above our 500 GeV energy limit provides

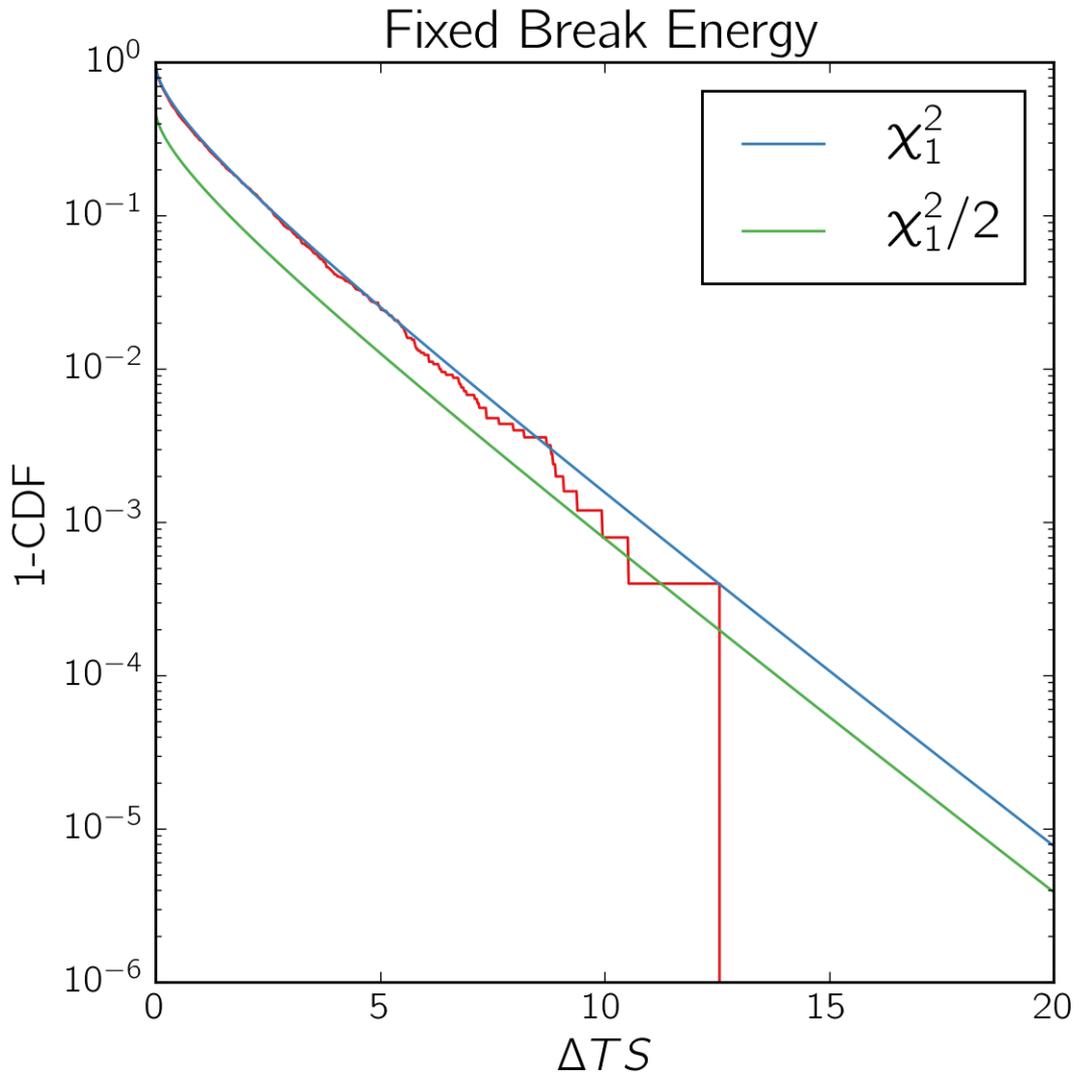


Figure 3.5: Distribution of the source TS when fitting a broken power-law to a simulated (unbroken) power-law signal with a fixed break energy parameter. As only one degree of freedom has been added, this should follow a χ_1^2 distribution.

an unphysical prediction that is clearly incorrect when comparing with the H.E.S.S. results published by (Aharonian et al., 2009), and only holds towards the top of our energy range due to our lack of constraining data (and thus large uncertainty). Comparing with the 3FGL catalogue, 395 sources are modelled using log parabolas in preference to a power-law (determined at the 4σ level), of which none have a negative β parameter. Figure 3.6 shows this catalogue distribution of curvature. In general, negative curvature using this model should really be discarded *a priori*, or in practice set a suitable prior to avoid these parameter values due to their unphysical nature, despite the value it can provide in this case to set a lower bound on the significance of the upturn. The aphorism “all models are wrong, but some are useful” certainly applies in the case of the log parabola.

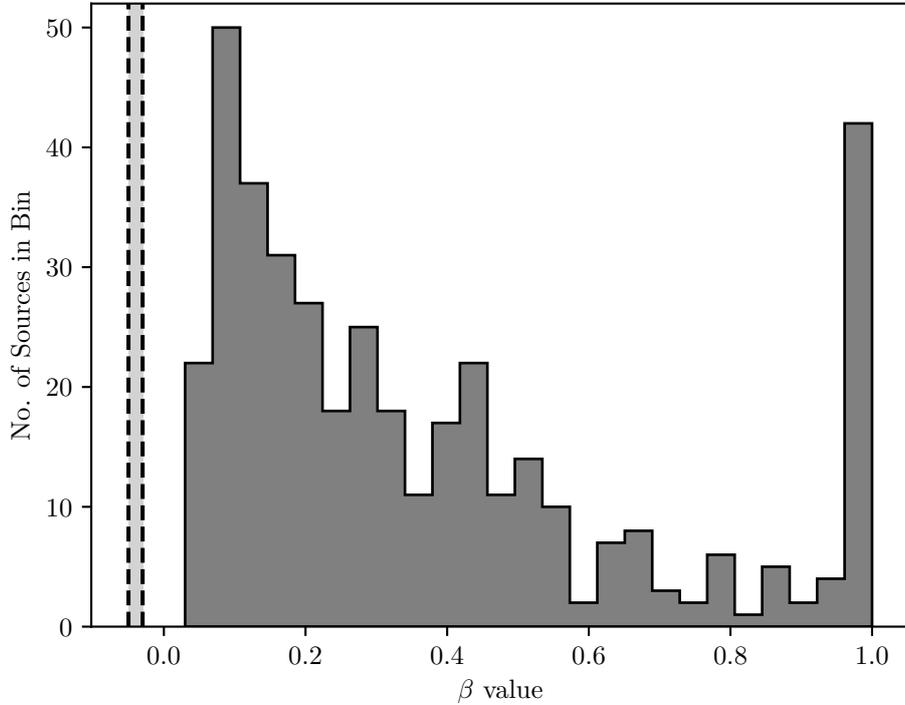


Figure 3.6: A histogram of the β values of 3FGL sources with log parabola spectral models, compared with the value calculated for Cen A whose uncertainty region is shown by the dashed lines and light grey region. The pile-up of sources at $\beta = 1$ is due to the parameter constraints introduced by Acero et al. (2015) to avoid excessive curvature. It should also be noted that values of $\beta=0$ suffer a selection bias as the log parabola reduces to a power-law at $\beta = 0$, meaning these sources do not meet the requirement for a 4σ preference over a power-law. What is most relevant is the lack of any sources with negative curvature, as is found in Cen A.

Model	Norm [MeV cm ⁻² s ⁻¹]	Shape Parameters	Scale E [MeV]	DoF	$2\Delta \log(\mathcal{L})^a$	ΔAIC
PL	$(4.36 \pm 0.06) \times 10^{-9}$	$\gamma = 2.690 \pm 0.014$	$E_0 = 481.17$	2	–	–
BPL	$(4.21 \pm 1.38) \times 10^{-11}$	$\gamma_1 = 2.73 \pm 0.02$ $\gamma_2 = 2.29 \pm 0.07$	$E_b = 2593 \pm 314$	4 ^b	28.6	24.6
LP	$(4.15 \pm 0.08) \times 10^{-11}$	$\alpha = 2.70 \pm 0.01$ $\beta = -0.04 \pm 0.01$	$E_0 = 481.17$	3	16.637	14.637

Table 3.1: A summary of the spectral fits using several models. PL: Power-Law, BPL: Broken Power-Law, LP: Log Parabola. The definitions of each model can be found in Equations 3.1, 3.2 and 3.5 respectively.

^a This is compared to that of the best-fit power-law model.

^b The broken power-law model does not satisfy Wilks’s theorem as it is not nested with a power-law model. The large discrepancy between normalisations is due to the scale energy varying between models, whose fractional difference is then raised to a power of ~ 2.5 .

3.5 Variability

Central to the interpretation of the deviation from a power-law model is how variable the source is, and especially if this variability is consistent between the two components above and below the break energy at ~ 2.6 GeV. For instance, if the emission is astrophysical in origin, variability might be detected in both components (independently or simultaneously, depending on the specific model), but for more exotic interpretations such as annihilating dark matter or populations of millisecond pulsars, emission would be entirely constant.

Variability is assessed by applying the analysis steps of Section 3.2 to 16 separate time bins (each 6 months in duration) as shown in Table 3.2, and is done in two separate energy ranges from 100 MeV–2.6 GeV and 2.6 GeV–300 GeV. 16 bins were chosen as a compromise between maintaining photon statistics and temporal resolution for a somewhat faint source such as Cen A (at least compared to the many bright blazars detected with *Fermi*). The seed model is instead replaced by our optimised model from the 8 yr dataset, which is yet again optimised as in the previous analysis and a scan is performed by varying the normalisation of Cen A to provide a likelihood profile corresponding to the flux of Cen A in each bin. The resulting data are plotted in Figure 3.7.

Start Time (UTC)	End Time (UTC)
2008-08-04 15:43:36	2009-02-03 12:07:11
2009-02-03 12:07:11	2009-08-05 08:30:47
2009-08-05 08:30:47	2010-02-04 04:54:22
2010-02-04 04:54:22	2010-08-06 01:17:58
2010-08-06 01:17:58	2011-02-04 21:41:34
2011-02-04 21:41:34	2011-08-06 18:05:10
2011-08-06 18:05:10	2012-02-05 14:28:46
2012-02-05 14:28:46	2012-08-06 10:52:21
2012-08-06 10:52:21	2013-02-05 07:15:56
2013-02-05 07:15:56	2013-08-07 03:39:32
2013-08-07 03:39:32	2014-02-06 00:03:08
2014-02-06 00:03:08	2014-08-07 20:26:44
2014-08-07 20:26:44	2015-02-06 16:50:20
2015-02-06 16:50:20	2015-08-08 13:13:54
2015-08-08 13:13:54	2016-02-07 09:37:30
2016-02-07 09:37:30	2016-08-08 06:01:06

Table 3.2: Table containing the start and end times of the temporal binning used for the lightcurve in Section 3.5

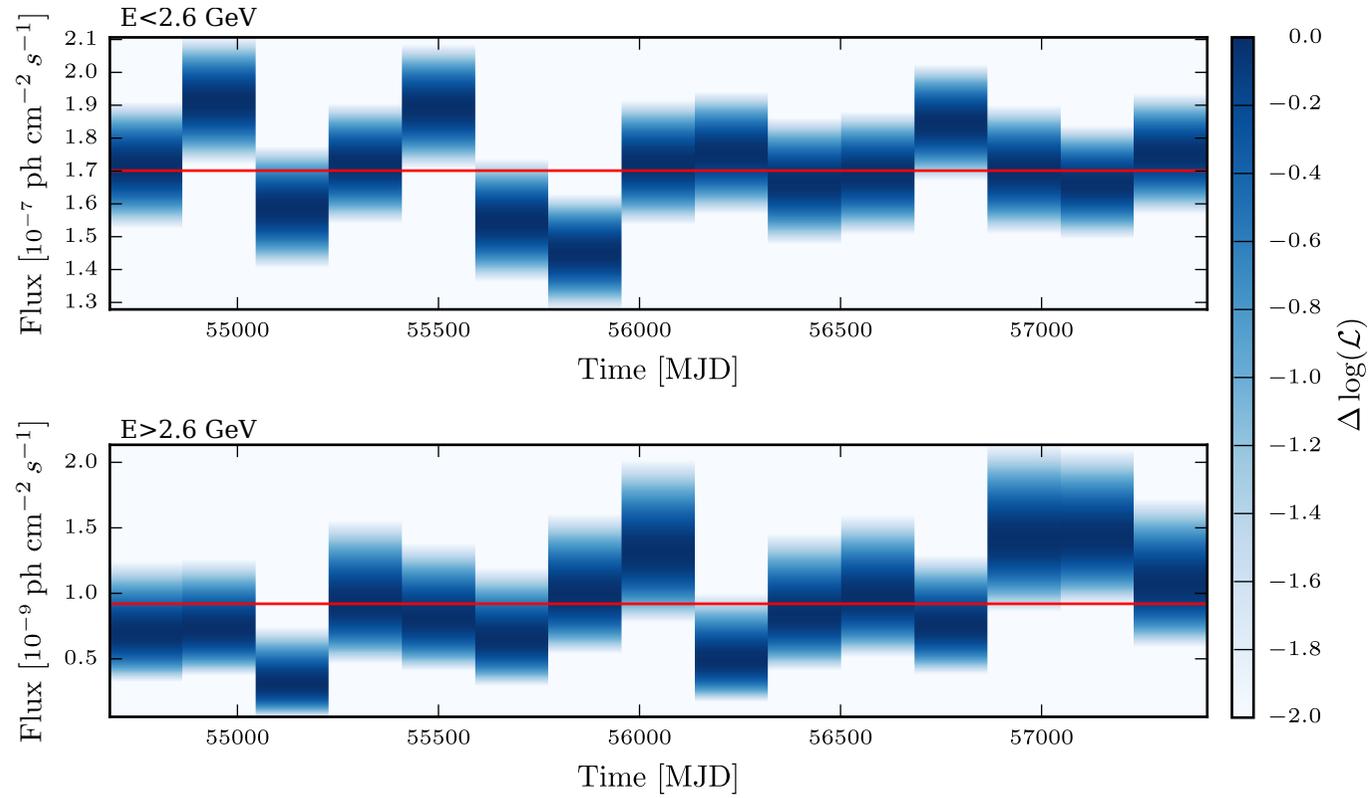


Figure 3.7: Lightcurves for energies $E < 2.6$ GeV (top) and $E > 2.6$ GeV (bottom). The blue colourmap represents the change in likelihood from the optimised position in each bin, and the red line represents the overall average flux.

Significance of the variability in each energy range is assessed using a χ^2 test with symmetrised uncertainties and alternately the TS_{var} as defined by Nolan et al. (2012):

$$\text{TS}_{\text{var}} = 2 \sum_i^n [\log(\mathcal{L}_i(F_i)) - \log(\mathcal{L}_i(F_{\text{const}}))], \quad (3.6)$$

where the index i represents the n time bins, $\mathcal{L}_i(F_i)$ is the likelihood at the optimised flux for bin i , and $\mathcal{L}_i(F_{\text{const}})$ is the likelihood value of the averaged flux value in bin i . The TS_{var} is distributed as χ_{n-1}^2 and can therefore be easily assessed for significance. It should be noted that in the high-statistics regime such that the fluxes are distributed normally, the TS_{var} reduces to the χ^2 , and in this case both methods find almost identical results consistent with no variability either above or below the break energy at the 95% level.

3.6 Systematic Uncertainty in IRFs

Our discussion of the IRFs in Section 1.2.5 did not extend to the evaluation of systematic uncertainties on modelling from uncertainty in instrument response. In this section we will mainly consider the uncertainty of the broken power-law of Cen A due to uncertainty in the LAT instrument's effective area A_{eff} . While the LAT was calibrated in the lab using test beams and with Monte Carlo simulations before launch (Rando & the *Fermi*-LAT Collaboration, 2009), further calibration post-launch has been assessed with on-sky data by Ackermann et al. (2012) which also aims to identify detector degradation over time.

Following the recommendations of Ackermann et al. (2012), we discuss methods of quantifying systematic errors in analysis of *Fermi*-LAT data. The main contribution to systematic error is uncertainty in the effective area A_{eff} of the detector, which matches Monte Carlo simulations to around

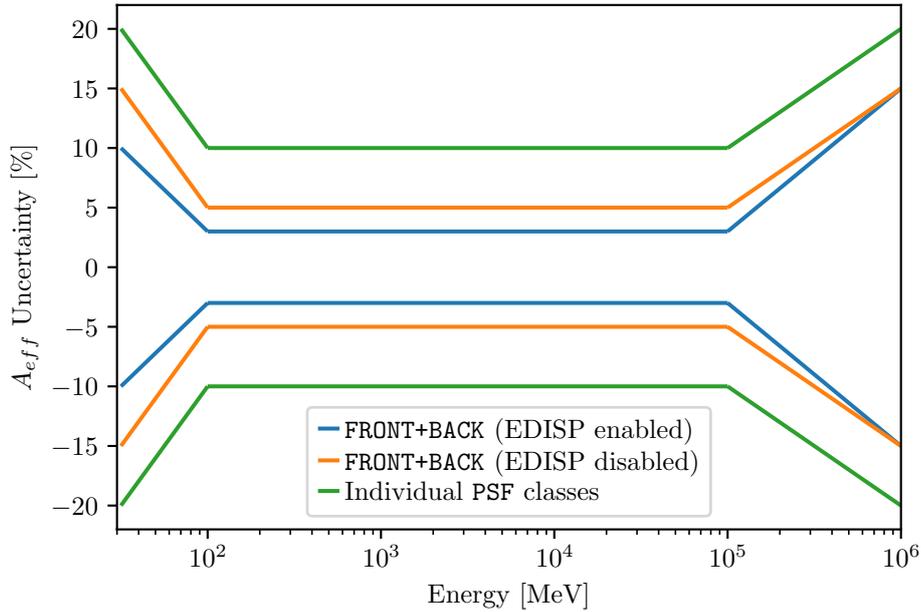


Figure 3.8: The $\pm 1\sigma$ uncertainty regions as determined by Monte Carlo simulations and on-orbit calibration for 3 different IRF combinations. The point-to-point correlation within this range cannot be shown in this plot, but the correlation is strong in neighbouring energy bins, becoming uncorrelated after half a decade in energy (as this calculation is performed logarithmically).

10% (and in most cases far better) as plotted in Figure 3.8. Taking this 10% value, and applying it to the IRFs, we have several options to customize the energy dependence of A_{eff} .

If we have a fractional systematic error:

$$\epsilon(E) = \frac{\delta A_{eff}}{A_{eff}}, \quad (3.7)$$

where δA_{eff} is the relative uncertainty on A_{eff} measured by Ackermann et al. (2012) and plotted in Figure 3.8. By defining some bracketing functions $B(E)$ which exist in the range $[-1, 1]$, we can define a modified A'_{eff}

as such:

$$A'_{eff}(E, \theta) = A_{eff}(E, \theta) \cdot (1 + \epsilon(E)B(E)). \quad (3.8)$$

There are two limiting cases to consider here:

- Limiting Fluxes – If we change A_{eff} by a flat multiplicative factor independent of energy, we can find the maximised integrated flux error of any given source. This corresponds to a bracketing function $B(E) = \pm 1$
- Limiting Index – We can also alter A_{eff} maximally on either side of the decorrelation energy E_0 , to maximise the error on the observed index. The bracketing function recommended by Ackermann et al. (2012) is a tanh function as such:

$$B(E) = \pm \tanh\left(\frac{1}{k} \log\left(\frac{E}{E_0}\right)\right). \quad (3.9)$$

The constant k determines the severity of the transition around E_0 . A value of $k = 0.13$ corresponds to smoothing over twice the LAT energy resolution.

These two limiting cases are plotted in Figure 3.9, along with the *Fermi* IRF for an angle of $\cos(\phi) = 1$.

The decorrelation energy as used by Ackermann et al. (2012) is calculated using:

$$E_0 = E \exp\left(\frac{\sigma_{N\gamma}}{N\sigma_\gamma^2}\right), \quad (3.10)$$

where E_0, N, γ have their usual definitions, σ_γ^2 is the variance in γ and $\sigma_{N\gamma}$ is the covariance between N and γ . These two limiting cases are plotted in figure 3.9, and their application to the effective area IRFs for an angle of $\cos(\phi) = 1$ is shown in Figure 3.11.

We now apply these altered IRFs to our analysis of the spectrum of Centaurus A (Cen A). Applying the same transformations of the IRFs in

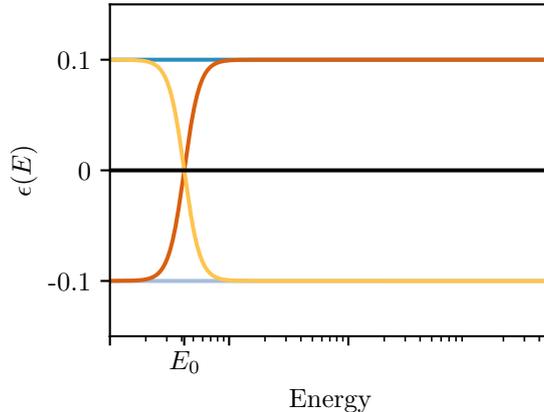


Figure 3.9: A sketch of the limiting bracketing IRF functions recommended by Ackermann et al. (2012). Red lines represent tanh alterations and blue lines represent the linear offsets. The black line represents the case assumed by the analysis tools by default.

Figure 3.9 and re-analysing the central point source of Cen A, we can see how the best fit power law transforms by finding the limiting fits for flux and index of a power law fit. These are plotted in Figures 3.12 and 3.13 for $\epsilon(E) = 0.1$. A summary of the fitted results from Cen A with various IRF alterations can be found in Tables 3.6 and 3.4.

While this provides an upper limit (the *worst case scenario*, in reality) to our systematic error, we may also want to be slightly more rigorous. In the majority of cases, the LAT systematics are more likely to be on the order of 2–3% rather than the worst case 10% (when compared to Monte Carlo simulations, Ackermann et al., 2012). To acknowledge this, we can instead create a distribution of $\epsilon(E)$, reanalysing the data for each of these IRFs. From these analyses, we can apply a weighted bootstrap to find our best fitting flux, index and their associated errors. The linear alterations to the IRFs had no effect on the model comparison between power-law and broken

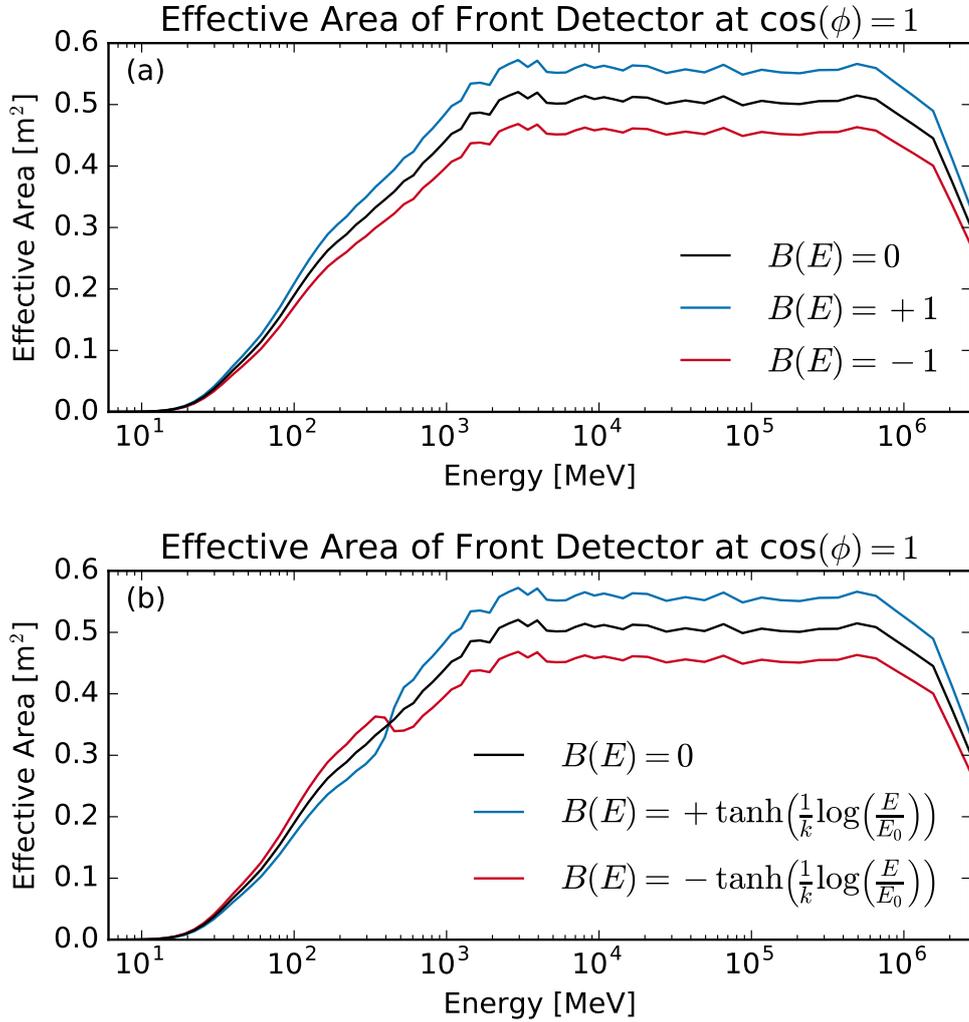


Figure 3.10: A comparison between the linear and tanh bracketing IRFs with baseline A_{eff} data provided with the *Fermi* Science Tools. The effective area for the front detector at $\cos(\phi) = 1$ is plotted with various bracketing functions, $B(E)$ for relative uncertainty $\epsilon(E) = 0.1$. **(a)** Linear $B(E)$ for relative uncertainty $\epsilon(E) = 0.1$. **(b)** Tanh bracketing functions as shown in the legend, with a decorrelation energy of 420 MeV.

power-law, with the likelihood ratio test resulting in $TS = 26.6 \pm 0.2$. By comparison, the tanh bracketing IRFs had a far larger effect on the model comparison, resulting in $TS = 26.2 \pm 3.1$. The uncertainties in TS values is calculated by the 68% uncertainty region following bootstrap Monte Carlo, a selection of which are plotted in Figures 3.14 & 3.15.

We now apply these altered IRFs to our analysis of the spectrum of Centaurus A (Cen A). Applying the same transformations of the IRFs in figure 3.11 and re-analysing the central point source of Cen A, we can see how the best fit power law transforms by finding the limiting fits for flux of a power law fit. This is plotted in Figure 3.13 for $\epsilon(E) = 0.1$. A summary of the fitted results from Cen A with various IRF alterations can be found in figure 3.4.

While this provides an upper limit (the *worst case scenario*, in reality) to our systematic error, we may also want to be slightly more rigorous. In the majority of cases, the LAT systematics are more likely to be on the order of 2–3% rather than the worst case 10% (when compared to Monte Carlo simulations, Ackermann et al., 2012). To acknowledge this, we can instead create a distribution of $\epsilon(E)$, reanalysing the data for each of these IRFs. From these analyses, we can apply a bootstrap to find our best fitting flux, index and their associated errors using the $\pm 34.1\%$ uncertainty region obtained from our sample. Most relevantly, the linear alterations to the IRFs had no effect on the model comparison between power-law and broken power-law, with the likelihood ratio test resulting in $TS = 26.6 \pm 0.2$ (Figure 3.14). By comparison, the tanh bracketing IRFs (Figure 3.15) had a larger effect on the model comparison, resulting in $TS = 29.2 \pm 4.1$. Understandably, a broken power-law fit becomes more significant with the tanh function as the decorrelation energy at which the tanh fit is of the order of the break energy we find. This creates an artefact that either “flattens

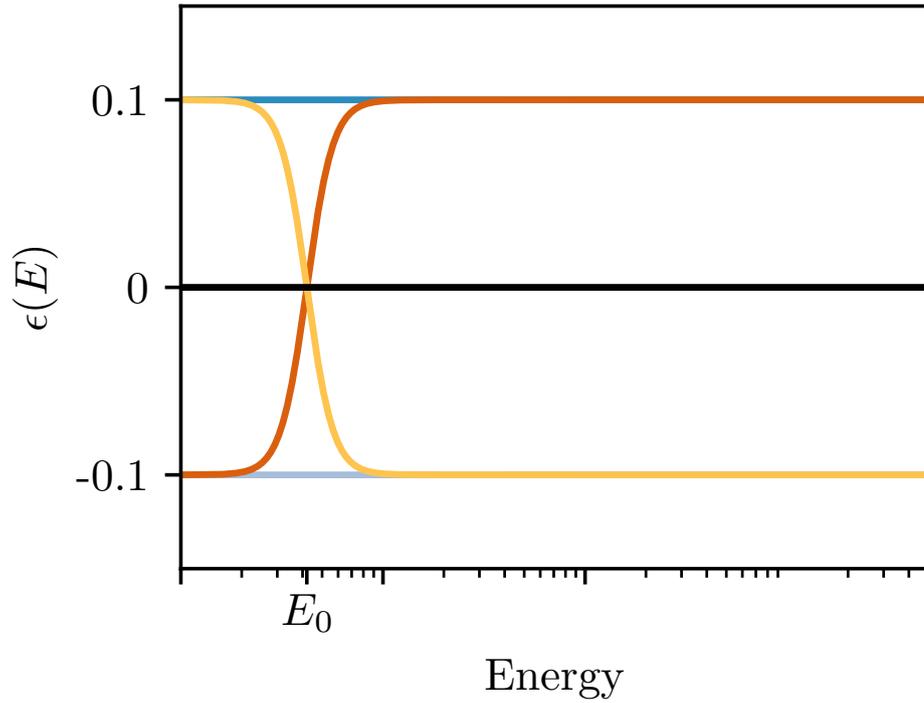


Figure 3.11: A comparison between the linear and tanh bracketing IRFs. The effective area for the front detector at $\cos(\phi) = 1$ is plotted with various bracketing functions, $B(E)$ for relative uncertainty $\epsilon(E) = 0.1$. Blue: Linear $B(E)$ for relative uncertainty $\epsilon(E) = 0.1$. Red: Tanh bracketing functions as shown in the legend, with a decorrelation energy of 420 MeV.

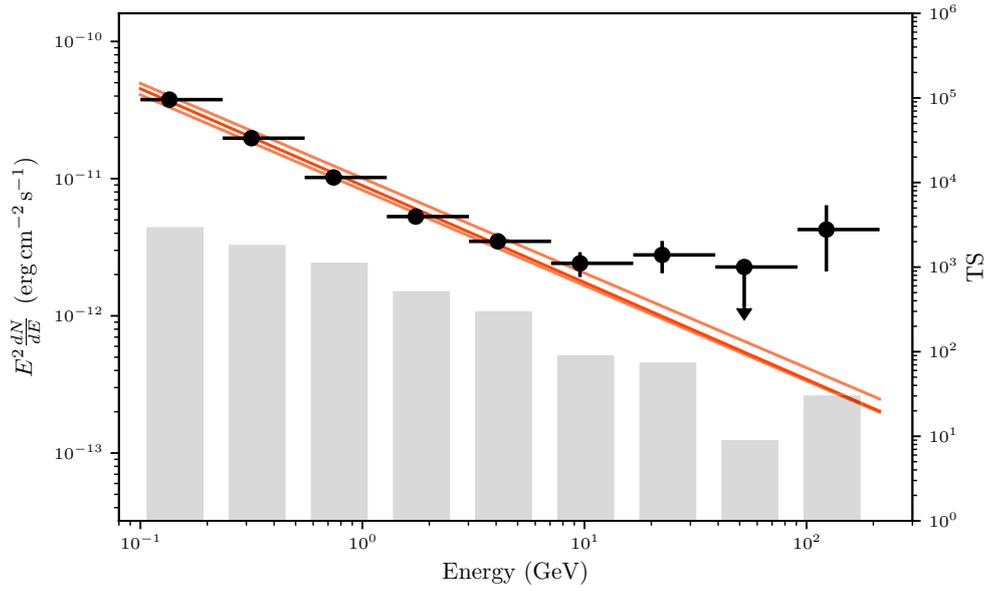


Figure 3.12: The $B(E) = \pm 1$ IRFs plotted either side of the naive $B(E) = 0$ fit.

out” the break, or significantly emphasises it. Thankfully, we can convince ourselves that this is likely not the case in reality, as such a break is not seen in the population of AGN with the LAT.

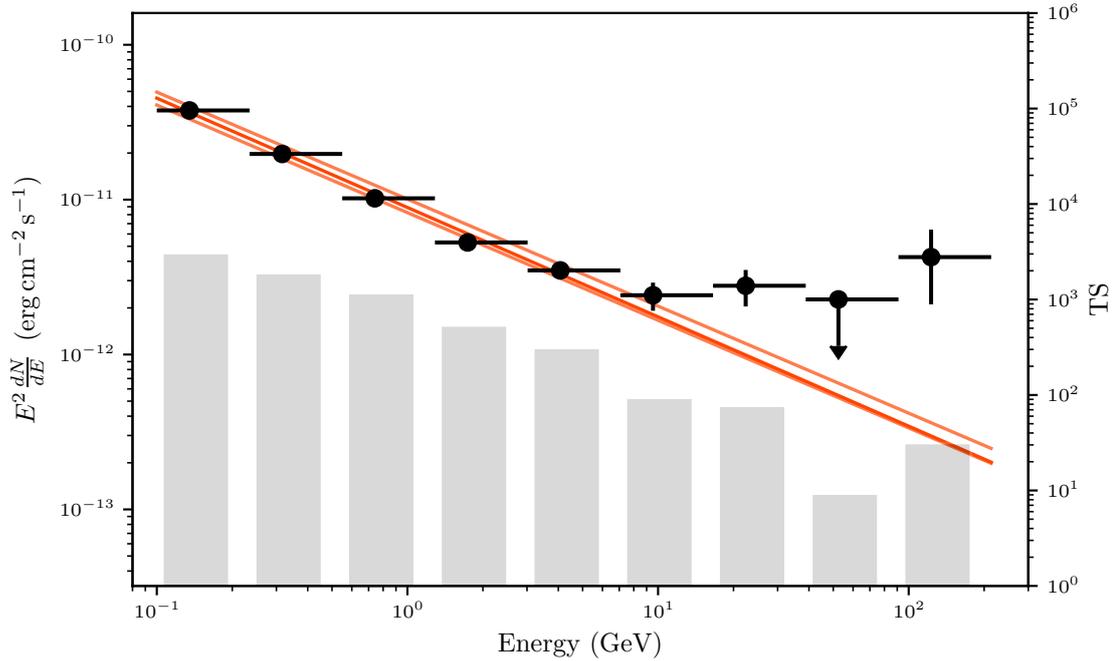


Figure 3.13: The $B(E) = \pm 1$ IRFs plotted either side of the naive $B(E) = 0$ fit (darker red). The bar chart represents the TS of each binned data point.

Power Law			
$B(E)$	Prefactor [$\text{MeV}^{-1} \text{cm}^{-2} \text{s}^{-1}$]	Index	Notes
0	4.020×10^{-11}	2.706	Unaltered IRFs
+1	3.660×10^{-11}	2.708	Lower Flux
-1	4.460×10^{-11}	2.704	Higher Flux
$+\tanh\left(\frac{1}{k} \log\left(\frac{E}{E_0}\right)\right)$	4.044×10^{-11}	2.770	Softer Index
$-\tanh\left(\frac{1}{k} \log\left(\frac{E}{E_0}\right)\right)$	4.058×10^{-11}	2.629	Harder Index

Table 3.3: Summary of the broken power law fits with various bracketing IRFs for $\epsilon(E) = 0.1$, $k = 0.13$ and $E_0 = 420 \text{ MeV}$ in the case of the power-law.

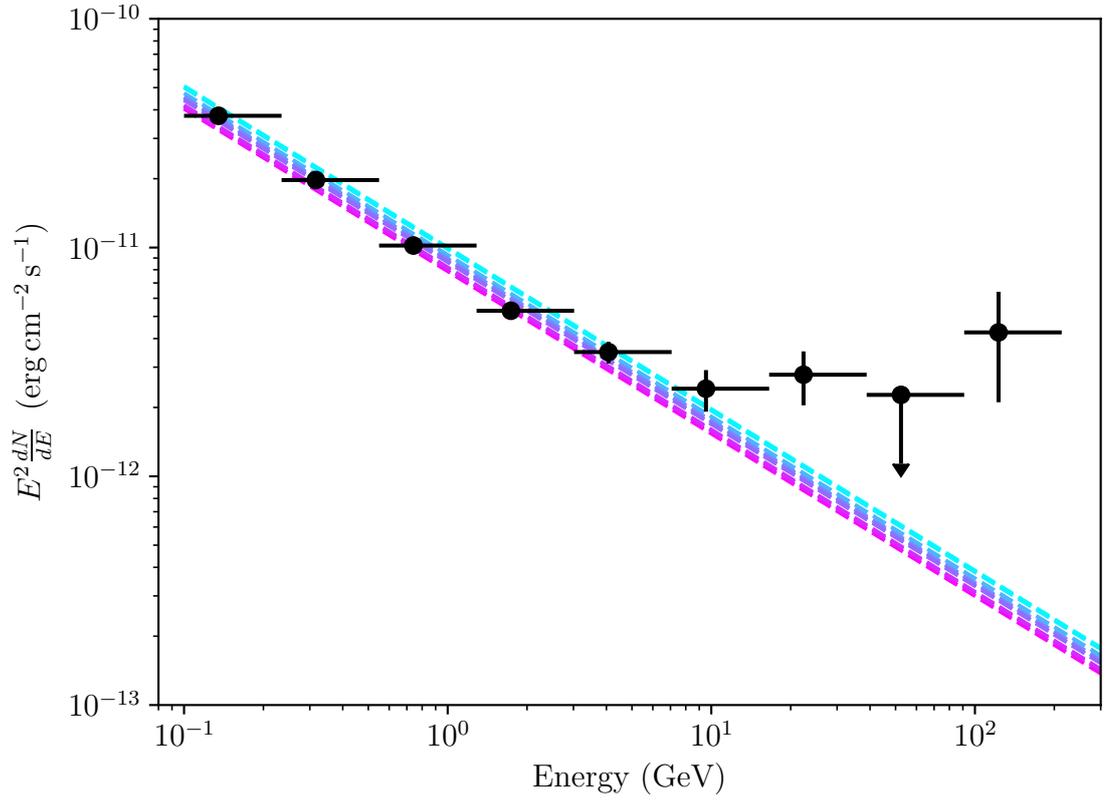


Figure 3.14: A subset of the power-law fits to the bootstrapped linear IRF alterations. The dashed lines represent individual best-fit power laws and are coloured according to the magnitude of $\epsilon(E)$ (more purple corresponds to more negative).

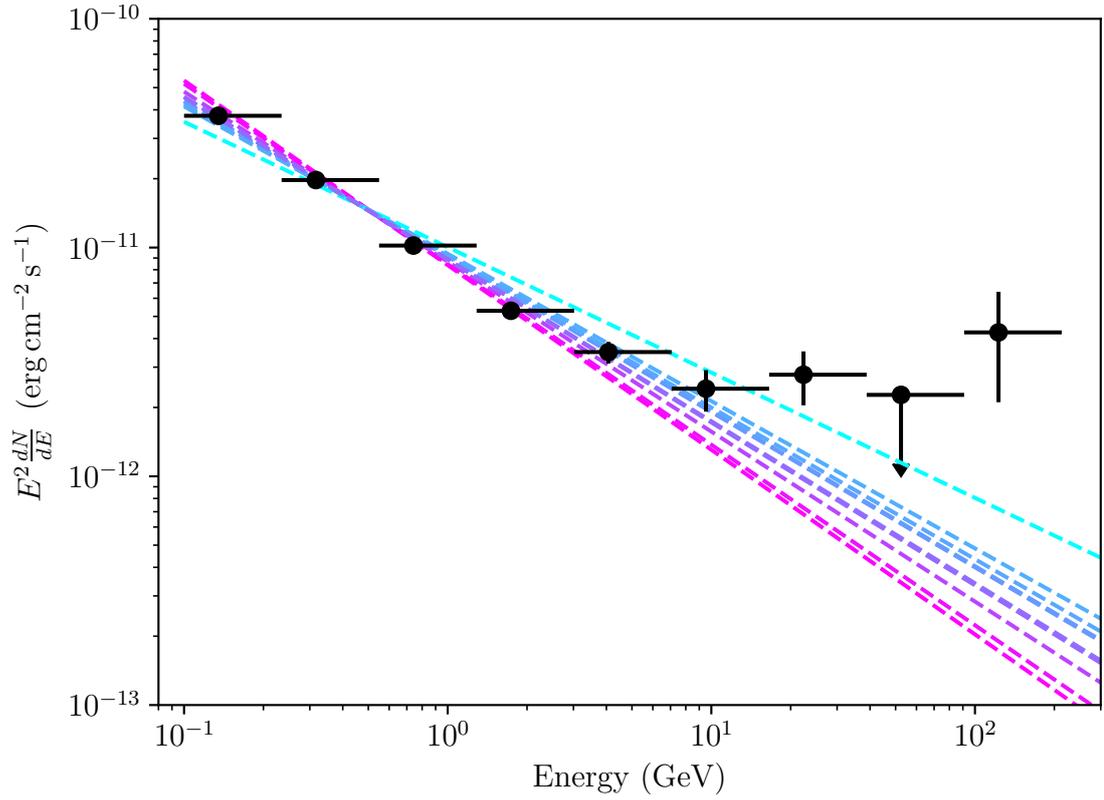


Figure 3.15: A subset of the power-law fits to the bootstrapped tanh IRF alterations. The dashed lines represent individual best-fit power laws and are coloured according to the magnitude of $\epsilon(E)$ (more purple corresponds to more negative).

Broken Power Law						
$B(E)$	Prefactor [$\text{MeV}^{-1} \text{cm}^{-2} \text{s}^{-1}$]	Index 1	Index 2	Break Energy [GeV]	Notes	
0	4.503×10^{-13}	2.766	2.282	2.420	Unaltered IRFs	
+1	4.459×10^{-13}	2.769	2.289	2.342	Lower Flux	
-1	5.362×10^{-13}	2.766	2.289	2.358	Higher Flux	
$+\tanh\left(\frac{1}{k}\log\left(\frac{E}{E_0}\right)\right)$	4.335×10^{-13}	2.830	2.207	2.370	Softer Index	
$-\tanh\left(\frac{1}{k}\log\left(\frac{E}{E_0}\right)\right)$	3.134×10^{-13}	2.683	2.238	2.928	Harder Index	

Table 3.4: Summary of the power law fits with various bracketing IRFs for $\epsilon(E) = 0.1$, $k = 0.13$ and $E_0 = 420 \text{ MeV}$ in the case of the power-law.

3.7 Physical Interpretations

As discussed in Section 2.1, a simplistic model for gamma-ray emission from AGN is the single-zone synchrotron self-Compton model in which a single (Why just one?) sphere (Why a sphere? Of what size?) of high-energy electrons (Just electrons? With what energy distribution? In the steady state?) travelling (At what angle to the line of sight?) relativistically (At what speed?) in a magnetic field (Of what strength?). This model is certainly not the last word in astrophysical modelling but as our (incredibly glib) questions show, even this simple model has a large number of parameters (~ 10) to be fit to the data. While information (read: statistics) at radio, optical, and X-ray is plentiful, this is not so in the gamma-ray regime. the *Fermi*-LAT is a survey detector that must integrate years of data for statistically relevant information. Our 8-year analysis of Cen A in Section 3.3 contains only 9 bins in 3 decades of energy, of which one is an upper limit. Frustratingly, many of these parameters depend heavily on the “second peak” of the spectrum caused by the inverse-Compton emission which is mostly constrained by gamma-ray data. Perhaps VHE data, more readily available using IACTs could be useful, but integrations over years such as with the LAT are effectively impossible, resulting in uncorrelated data. Nonetheless, the SSC model is at least an approachable first estimate of the astrophysical emission of AGN for comparison. Indeed, Petropoulou et al. (2014) found that a single-zone SSC model of Cen A was insufficient to model the high-energy tail and instead propose a multi-zone SSC model with 23 separate parameters (with several constrained with other observations). These models also make predictions of the size of the emitting region of $R \sim 10^{15}$ cm, and a resultant timescale of variability t_{var}

via the light-crossing time:

$$t_{var} \geq \frac{R(1+z)\Gamma(1-\beta\cos\theta)}{c} \quad (3.11)$$

where the redshift $z = 0.0018$, $\Gamma(1-\beta\cos\theta) = 1$, and c is the speed of light, giving $t_{var} \geq 37$ hrs. It is safe to say that the largest weakness with this interpretation is the huge number of parameters needed to fit models given the limited data available.

As foreshadowed in Section 3.1, there are alternative explanations that could explain the emission from a nearby galaxy such as Cen A. Due to this proximity, MAGN such as NGC 1275 can be visible despite jet angles that are incongruent with our line of sight. This same fact means that other non-beamed sources of emission should be considered. For example, in our own Galaxy we do not detect particularly luminous jet emission from Sgr A*, although it is a point source at gamma-ray energies. Nonetheless, once this is modelled out there is a significant residual emission originating from the centre of the Galaxy. This emission was originally suggested as evidence of self-annihilating WIMP dark matter within the dense inner region of the Galaxy (Goodenough & Hooper, 2009; Gordon & Macías, 2013). This result has seen scepticism and great debate over what proportion of the signal can be explained through astrophysical emission processes and what can be ascribed to beyond the standard model (BSM) physics (Abazajian, 2011; Bartels et al., 2018). Two analyses in particular used wavelet-based analyses (Bartels et al., 2016) and non-Poissonian statistics (Lee et al., 2016) to provide evidence against the predicted “smoothness” of the dark matter distribution, pointing instead to a population of unresolved point sources present in the inner Galaxy, thought to be millisecond pulsars (MSPs). Further questions have arisen as to whether this population of MSPs exists, as it has not been observed at other wavelengths and there are some doubts

as to the origin of this population (Hooper & Mohlabeng, 2016).

It should be noted before we explore DM and MSP scenarios that in both cases some jet emission is required to model the broadband spectrum of Cen A and is used to explain the gamma-ray spectrum up the break energy found in Section 3.3. This jet emission is comparable to that of single-zone SSC models as published by Petropoulou et al. (2014) and *Fermi* Collaboration et al. (2010). Physical modelling of the DM and MSP interpretations that are discussed in the remainder of this chapter was performed by collaborators for our joint publication Brown et al. (2017), and so will be given as context, but are not entirely the author’s original work.

To evaluate the plausibility of a millisecond pulsar (MSP) origin for the observed spectral bump in the gamma-ray spectrum of Cen A, a simple order-of-magnitude sanity check was performed. Integrating the fitted broken power-law model over the bump energy range (1–10 GeV), we estimate the gamma-ray luminosity associated with this feature to be $L_{\text{bump}} \approx 1.9 \times 10^{42} \text{ erg s}^{-1}$, assuming a distance of 3.7 Mpc.

By comparison, typical MSPs observed in the Milky Way with *Fermi*-LAT have gamma-ray luminosities of $L_{\text{MSP}} \sim 10^{33}\text{--}10^{34} \text{ erg s}^{-1}$ (e.g. Abdo et al., 2013; Winter et al., 2016). Using a conservative average of $3 \times 10^{33} \text{ erg s}^{-1}$, this implies that more than 6×10^8 MSPs would be required within the core of Cen A to account for the bump’s luminosity. This is many orders of magnitude greater than even the total MSP population estimated for the Milky Way, rendering this explanation implausible.

Another possible explanation for the observed high-energy emission is inverse-Compton (IC) scattering from high-energy electrons within the kiloparsec-scale jet of Cen A. Such electrons are known to exist based on radio and X-ray observations, and could plausibly upscatter ambient photons (e.g. CMB, starlight) into the GeV–TeV regime (e.g. Abdo et al., 2010b; Hard-

castle et al., 2009). A detailed treatment of this scenario is beyond the scope of this chapter, but it represents a physically motivated alternative to MSP or dark matter interpretations.

First-order estimates for DM annihilation can be found assuming a standard NFW (Navarro et al., 1996) DM density profile and annihilation into $b\bar{b}$ with a (thermally averaged) annihilation cross section :

$$\langle\sigma v\rangle = 3 \times 10^{-26} \text{ cm}^2 \text{ s}^{-1},$$

the relevant quantity to reproduce relic abundance and the expected prompt emission (ignoring acceleration processes and secondary emission). This under-predicts the data by several orders of magnitude and is shown in Figure 3.16.

Spiky dark matter distributions with inner slopes much sharper than that of the NFW profile as discussed in Section 3.1 can reproduce the expected flux easily as shown in Figure 3.17, although the rather low :

$$\langle\sigma v\rangle \approx 1.6 \times 10^{-32} \text{ cm}^2 \text{ s}^{-1},$$

cannot explain the observed fraction of DM in the universe. This would require other interpretations such as an expanded set of DM particles (Boehm et al., 2004; Zurek, 2009) or a velocity-dependent cross section. Furthermore, only single-channel annihilations are considered and DM particles may annihilate to different states. In these models, both the DM particle mass and cross section are fitted parameters, and the $t\bar{t}$ model obtains a $\chi^2/\text{d.o.f.} = 0.62$, or alternately with a decay to $b\bar{b}$ we find $\chi^2/\text{d.o.f.} = 0.87$.

The MSP model used relies upon the standard pulsar emission mechanisms (Harding et al., 2007) through e^+e^- production and inverse-Compton processes to generate gamma-ray emission up to TeV energies which has previously been applied to the galactic centre (Bednarek & Sobczak, 2013) and

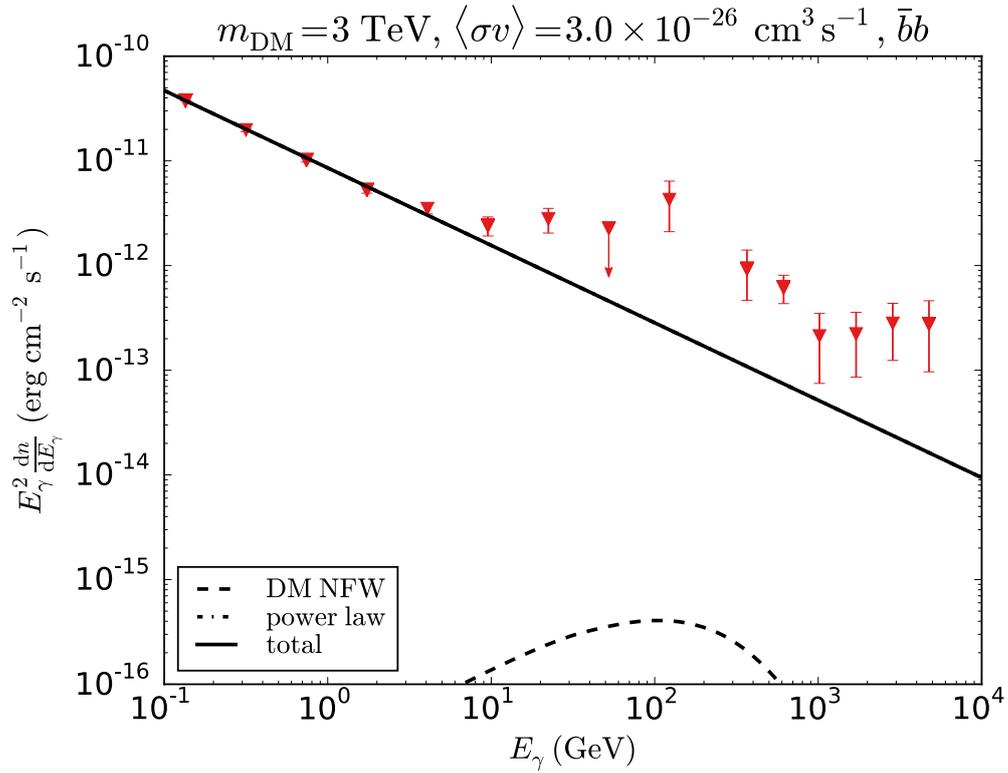


Figure 3.16: Prompt emission from an NFW-distributed WIMP dark matter annihilation to $b\bar{b}$ with a relic cross section. The behaviour before the break is modelled with a simple power-law extrapolated from the broken power-law in Section 3.3. The data above 300 GeV are the results published by Aharonian et al. (2009).

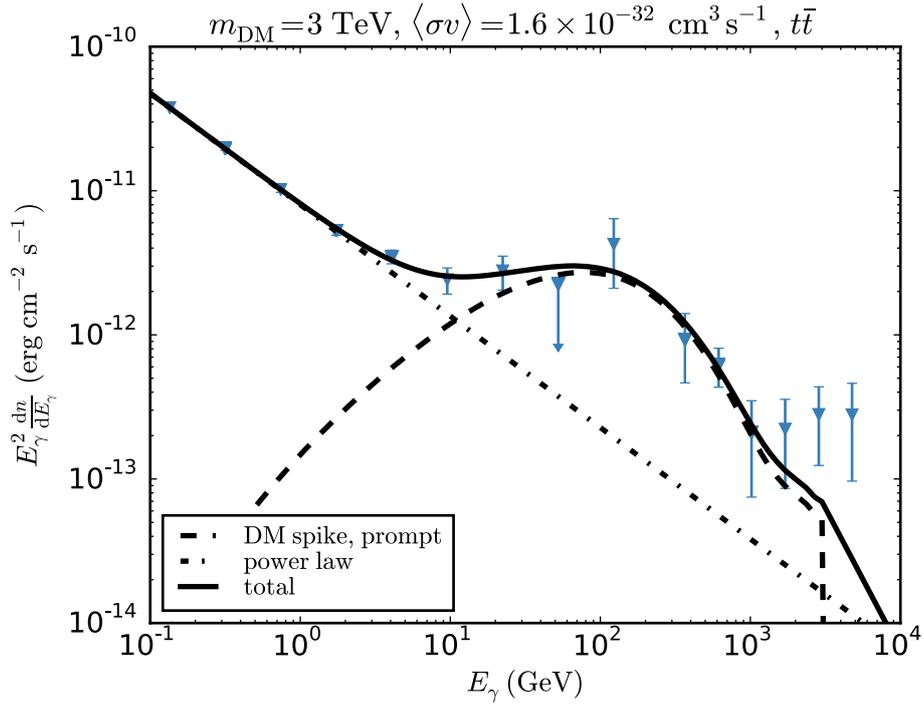


Figure 3.17: DM fit with a spiky dark matter profile, with a DM mass of 3 TeV annihilating to $t\bar{t}$ results in a best-fit cross section of $1.6 \times 10^{-32} \text{ cm}^2 \text{ s}^{-1}$. The extrapolated power-law jet from the best-fit spectrum below the break energy is shown with the dot-dashed line. The data above 300 GeV are the results published by Aharonian et al. (2009).

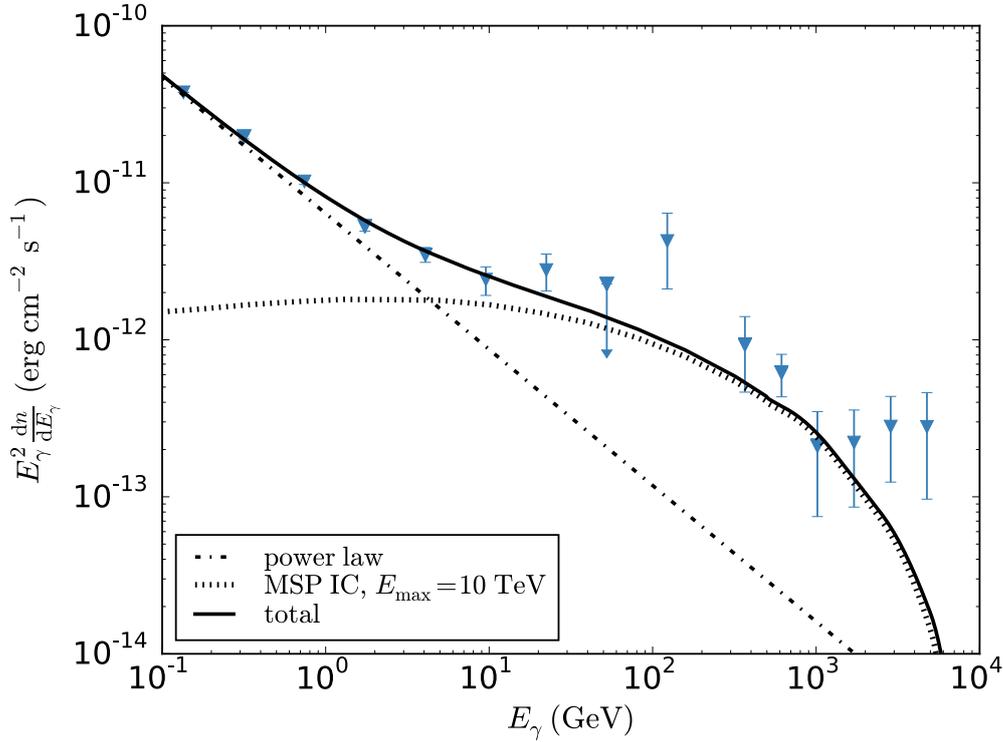


Figure 3.18: MSP population fit with $\rho_{\text{MSP}} \propto r^{-2.4}$ and an extrapolated power-law jet from the best-fit spectrum below the break energy. The data above 300 GeV are the results published by Aharonian et al. (2009).

M31 (Ackermann et al., 2017a). The main assumptions of the model by Bednarek & Sobczak (2013) are that the electrons can be accelerated to TeV energies by the MSPs, and that the local radiation fields allow inverse-Compton losses to dominate synchrotron losses. Using a constant $10 \mu\text{G}$ magnetic field and a density power-law of $\rho_{\text{MSP}} \propto r^{-2.4}$, we find a fit to the data with $\chi^2/\text{d.o.f.} = 1.4$, which is also consistent with the data as plotted in Figure 3.18. It should be noted in this case that our poor knowledge of the local magnetic field and photon density is a far larger source of uncertainty than the constraints of the gamma-ray data.

A combined model with both DM and MSP contributions was also con-

sidered with a much heavier DM particle to explain the VHE tail and a good χ^2 fit, but given the uncertainties surrounding the inputs to the models under consideration (photon field density, magnetic field strength, multiple DM branching ratios, DM mass uncertainty, etc.) this is undoubtedly overfitting the limited data.

3.8 Conclusions

I have presented an in-depth analysis of Centaurus A using the *Fermi* Large Area Telescope, finding the following:

1. There is a significant ($\sim 5\sigma$) hardening in the spectrum away from a simple power-law, with a break energy of 2.6 GeV, and this agrees with VHE gamma-ray measurements with H.E.S.S..
2. This result is robust to systematic effects due to the IRFs used in the analysis, and is significant in all but the most extreme assumptions for the LAT response. Similarly, the upturn is present and statistically significant despite the difficulties arising from the phenomenological models used.
3. No significant variability is detected over the 8 yr dataset either above or below the break energy, whereas in other similar galaxies such as M87, this variability is strong (Aharonian et al., 2006).
4. Non-jet origins of this emission are explored. WIMP dark matter models cannot explain this upturn with standard DM distributions, nor can this occur with spiky profiles and relic cross sections. MSP populations in Cen A's bulge could also explain this emission, but there is little constraint on photon fields or magnetic fields to make accurate predictions of this flux.

5. Combinations of DM and MSP models can be considered but in general will overfit the data.

These results put the previous hints of an upturn (Sahakyan et al., 2013; Aharonian et al., 2009) onto a solid statistical footing, but only skim the surface of the physical interpretations possible for Cen A. While we have highlighted the shortcoming of multi-component SSC models due to the vast quantity of parameters needed with such sparse data, DM interpretations can seem attractive due to their low dimensionality. What has not been considered is the incredible Bayes factor associated with the DM interpretation without independent corroboration of these results. If similar DM properties could be measured through other methods (colliders or direct detection), or even other analyses with astronomical signals (the galactic centre excess or isotropic line emission) then this Bayes factor could be greatly reduced, but certainly multiple corroborating experiments would be required to close the book on this matter. MSP interpretations sit in the middle, with both significant parameter uncertainties and source confusion in the inner galaxy contributing to the Bayes factor. This work does however show that all three physical interpretations are valid within the constraints of the data under plausible assumptions. Further study at VHE such as with H.E.S.S. (see H. E. S. S. Collaboration et al., 2018, for one example), or in future with CTA will undoubtedly help constrain these models further. One of the greatest difficulties with this analysis is the scale of the systematics at VHE.

Spectral Variability in IC 310

4.1 Introduction

IC 310 is an enigmatic AGN. It has been variously classified as a head-tail radio galaxy (Acerro et al., 2015; Sijbring & de Bruyn, 1998), a high-frequency peaked BL-Lac object (Rector et al., 1999; Kadler et al., 2012), and is currently recognised as an AGN of unknown type (Rieger & Levinson, 2018). A part of the Perseus galaxy cluster, it is located at $(\text{RA}, \text{Dec}) = (49.18^\circ, 41.32^\circ)$, a short angular distance away from NGC 1275, itself a particularly bright radio galaxy known for its variability in the gamma-rays (Brown & Adams, 2011). IC 310 was discovered at VHE by MAGIC during an extended exposure on the Perseus cluster during a particularly bright flaring period, in which variability on the scale of a day was found (Neronov et al., 2010). Subsequent observations found far shorter variability timescales of < 5 min, which correspond to light crossing times far smaller than the size of the black hole, and hence speculation about a possible magnetospheric origin of the emission (Aleksić et al., 2014).

This discovery was followed up at other wavelengths, including with *Fermi*-LAT (Neronov et al., 2010) where it was found toward the upper

energy range of the instrument where the improving PSF allowed disentanglement from the nearby bright radio galaxy NGC 1275. Notably the spectrum found with this analysis denoted a particularly hard source with a power law index $\gamma = 1.3 \pm 0.5$ (Aleksić et al., 2014a). This is understandable given MAGIC spectra predicted an inverse-Compton peak within the VHE regime, consistent with the HBL behaviour that had been hinted at by Kadler et al. (2012). It should be noted that even as an HBL, IC 310 would be an outlier. As noted by Aleksić et al. (2014a), it is significantly less luminous than other HBLs, and the kiloparsec-scale radio structure is unusual for these objects.

At a redshift of $z = 0.0189$, attenuation by interaction with the extragalactic background light is minimal until TeV energies, but can be significant. Previous work by Aleksić et al. (2014a) provided both raw and absorption-corrected spectra, the latter of which shall be assumed when referred to later in this Chapter.

X-ray studies of IC 310's spectra have found a similarly mysterious object. Aleksić et al. (2014b) found strong variability in X-ray spectral power law indices, finding both a soft ($\gamma = 2.55 \pm 0.07$) state and a hard ($\gamma = 1.76 \pm 0.07$) state. This was also accompanied with an increase in N_H explained by the presence of material close to the black hole. Such an increase was also reported in Cen A by Benlloch et al. (2001).

In this chapter, we'll attempt to re-analyse IC 310 using an expanded dataset compared to that used by Neronov et al. (2010) and contextualise this with the X-ray and VHE observations discussed so far. The analysis and results presented in this Chapter were published in the article Graham et al. (2019).

Science Tools Version	v10r0p5
Pass	P8R2
ROI Width	20°
Event Class	SOURCE (128)
Event Type	PSF (4,8,16,32)
Zenith Cut	< 90°
Time Selection Filter	(DATA QUAL>0)&&(LAT CONFIG==1)
Energy	100 MeV–500 GeV
Time	(2008-08-04 15:43:38)–(2016-08-08 06:01:06) UTC
Galactic Model	gll_iem_v06.fit
Isotropic Model	iso_P8R2_SOURCE_V6_PSF{0..3}_v06.txt

Table 4.1: A summary of the data selection performed for the preliminary analysis on IC 310.

4.2 Data Reduction & Preliminary Analysis

We use many of the same choices in our data selection here as we did in our analysis of Cen A, as described in Section 3.2. We shall discuss the differences in this Section, but an exhaustive table can be found in Table 4.1. The largest change from our last analysis is that the energy range has been slightly extended from 100 MeV–500 GeV as IC 310 is particularly significant at higher energies, and is more easily distinguished from NGC 1275 with the improved PSF and relative brightness. A slightly longer time period was considered as more data were available when the analysis began.

Our preliminary analysis was more involved than in the case of Cen A due to the proximity of NGC 1275 (0.6° away). In the 3FGL, IC 310 is detected with a significance just under 8σ and a hard power-law index ($\gamma = 1.902 \pm 0.143$) with no significant variability ($TS_{var} = 38.742$, corresponding

4.2. DATA REDUCTION & PRELIMINARY ANALYSIS

Source	RA (J2000) (degrees)	Dec. (J2000) (degrees)	Offset (degrees)	TS
PS J0312.8+4121	48.21	41.35	0.716	42.07
PS J0355.3+3910	58.82	39.16	7.680	39.79 ^a
PS J0302.5+3353	45.63	33.88	7.958	51.36
PS J0257.1+3360	44.28	33.99	8.289	30.12
PS J0344.6+3433	56.15	34.55	8.725	36.32
PS J0233.0+3740	38.25	37.66	9.181	40.86
PS J0410.7+4218	62.68	42.30	10.107	74.66
PS J0344.2+3202	56.04	32.04	10.795	34.55
PS J0359.4+5052	59.86	50.87	12.055	359.99 ^b

Table 4.2: Summary of the new sources with $TS > 25$ found in our ROI. We have followed the naming convention of the *Fermi* catalogues for the source column with “PS” identifying the source as newly discovered.^a PS J0355.3+3910 was no longer a significant source when an analysis centred directly on this source was performed.^b While not in the 3FGL, this source was identified as 4C+50.11 in Carpenter et al. (2014) and is present in the 4FGL.

to 1.3σ with 47 degrees of freedom). The 3FGL was used as a seed catalogue in the model of the ROI, before the parameters of all sources were fitted in an iterative process from most significant to least significant before a final fit with all parameters free. As our data set contains several more years of data than were used for the 3FGL survey, we then applied a source finding algorithm using *fermipy* (Wood et al., 2017). In total, 9 new sources with significance $> 5\sigma$ ($TS > 25$) were discovered in this 20° ROI as summarised in Table 4.2.

One source, designated as PS J0312.8+4121, was detected at a signif-

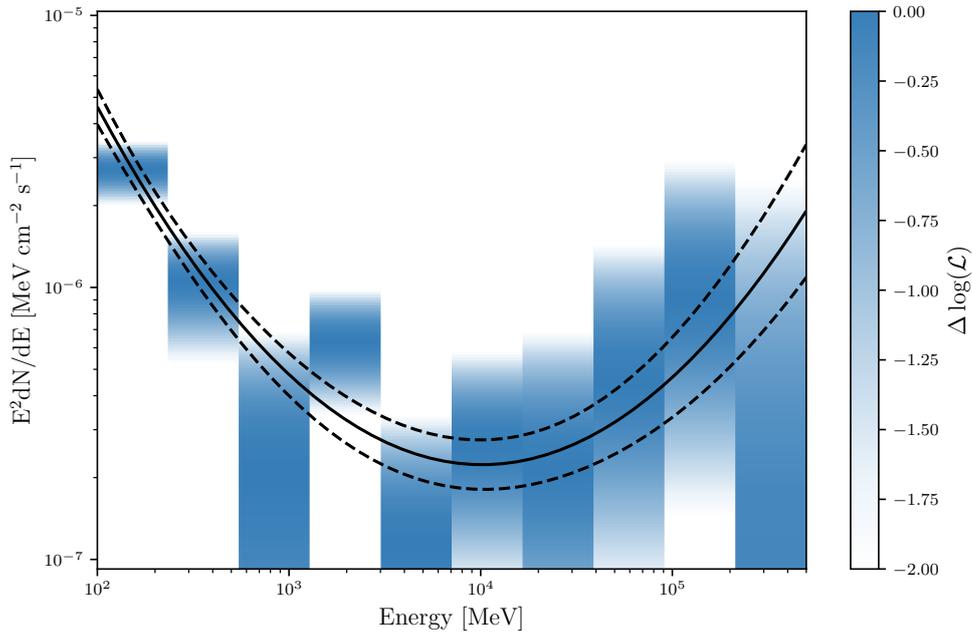


Figure 4.1: The spectral energy distribution of IC 310 when fitted with a log parabola model. This fit is performed over the entire 8-year time period and is jointly optimised with the nearby point sources and diffuse background emission templates.

ificance $>5\sigma$ at a distance $<1^\circ$ away from IC 310 and NGC 1275. This complicates the analysis of the region with regard to the number of possible models to be investigated, which we shall explore in Section 4.3. The initial SED of IC 310 as found from this preliminary analysis is plotted in Figure 4.1, and appears once again to have an upturn similar to Cen A in the previous Chapter. The log parabola fit in this initial model was preferred by $\text{TS}_{\text{curve}} > 16$, which was considered to be significant curvature by Acero et al. (2015). This concludes the initial modelling steps taken to obtain a rough fit to the region which we can now explore more fully.

4.3 Spectral Modelling

As we have found in even just our basic modelling in Section 4.2, this region of the gamma-ray sky is particularly complex. This initial analysis found a possible spectral upturn in IC 310, but unlike our previous analysis of Cen A, we have the confounding factors of two nearby sources (one strong, one relatively weak). Two factors will aid us in this analysis - firstly NGC 1275, while strongly variable in pure flux, is considered to vary only in normalisation and not spectrally (Ajello et al., 2016), although we do allow for this possibility. Secondly, the degree of freedom due to the LAT's energy range allows us to overcome the limitations of its PSF, as was found by Neronov et al. (2010).

The models considered are as follows:

- Model A: The model for IC 310 is a power law with the index fixed to its 3FGL value. This serves as the reference model for our work compared to the null hypothesis, which in this case is the 3FGL.
- Model B: The spectral index parameter for IC 310 is varied in addition to its normalisation.
- Model C: Spectral curvature is introduced by substituting a log parabola spectral model for the power law description of IC 310. These models are nested and so can provide a direct comparison. The $\Delta \log(\mathcal{L})$ between models C and B is the TS_{curve} as defined in the 2FGL (Nolan et al., 2012), though distinct from `Signif_Curve` as we do not take into account systematic uncertainties.
- Model D: We now add the point source PS J0312.8+4121 to model B. The spectral model is assumed to be a power law with spectral index 2.0. In this model, we have also allowed the spectral index of

IC 310 to vary. This results in 4 additional degrees of freedom from model A, namely the spectral index of IC 310, and the normalisation and position (RA and Dec.) of PS J0312.8+4121. The 3 parameters associated with PS J0312.8+4121 are not nested (there is a degeneracy when the normalisation of the source is zero), however, assuming the calibration performed by Mattox et al. (1996) persists, all significance test values remain accurate.

- Model E: The spectral index of PS J0312.8+4121 is allowed to vary from model D. The additional free parameters are now the two spectral indices, and the normalisation and position of PS J0312.8+4121.
- Model F: IC 310 is now replaced with the log parabola spectral model from model C, along with a fixed power law spectral model for PS J0312.8+4121 as in model D.
- Model G: The spectral index of PS J0312.8+4121 in model E is now allowed to vary. This was the first additional parameter that did not result in an overall increase in significance.

It should be noted that only models (A, B, C) and (D, E, F, G) are truly comparable in terms of TS. This is because the position was not re-optimised in each case for models D, E, F, and G. As the parameter was not left free to vary in the fit, the different TS values are only approximations (local minima). As the models may not be fully optimised, the TS may be lower than the fully optimised values. We can at least be assured that the optimal value would result in larger TS values, not smaller ones (when directly compared to model D). Comparisons between models E, F, and G assume the position of PS J0312.8+4121, and may be larger or smaller when the position is minimised (unlike comparisons with model D), as the

relative changes in $\log(\mathcal{L})$ are not calculated.

The summary in Table 4.3 shows model F to be the most likely model given the observed data. In each case the sky model was optimised and the log likelihood extracted. Each additional degree of freedom must decrease the $\log(\mathcal{L})$ by 1.0 to improve the significance, the relationship defining the Akaike Information Criterion (AIC, Akaike, 1974). We also include the AIC as an independent test of significance assuming that approximate nesting is insufficient as with our analysis in Section 3.4. The significance (in units of σ) is calculated as:

$$\text{Significance} = \sqrt{2} \operatorname{erf}^{-1}(F_{\chi^2}(TS, k)), \quad (4.1)$$

where erf^{-1} is the inverse error function, $F_{\chi^2}(x, k)$ is the cumulative distribution function of a χ^2 -distributed statistic at a value of x with k degrees of freedom, and σ refers to one standard deviation of a normal distribution.

The AIC is defined as in Section 3.4:

$$AIC \triangleq 2k - 2\log(\mathcal{L}), \quad (4.2)$$

where k is the number of degrees of freedom in the model, and \mathcal{L} is defined as above. The model with the minimum AIC value is the preferred model. Table 4.4 contains the optimised spectral parameter values for each model tested in this Section.

Name	Model	$-\log(\mathcal{L})$	Δk ^a	TS ^b	Signif. ^b (σ)	AIC ^c
A	IC 310 Catalog Index	1886808.45	0	-	-	-
B	IC 310 Power Law (PL)	1886805.72	1	5.47	2.34	-3.46
C	IC 310 Log Parabola (LP)	1886798.19	2	20.52	4.14	-16.52
D	IC 310 (PL) + PS (PL, Index=2.0)	1886791.06	4	34.79	5.02	-26.78
E	IC 310 (PL) + PS (PL, Index Free)	1886789.63	5	37.66	5.05	-27.64
F	IC 310 (LP) + PS (PL, Index=2.0)	1886784.43	5	48.04	5.91	-38.04
G	IC 310 (LP) + PS (PL, Index Free)	1886784.20	6	48.50	5.74	-36.5

Table 4.3: A summary of the sky models tested in Section 4.3. ^a The relative number of degrees of freedom (k) in each model as compared to model A. ^b These values are approximate, and are discussed further in Section 4.3. ^c As the $\log(\mathcal{L})$ values returned from the *Fermi* tools is only relatively correct due to the suppression of model-independent terms, we show only the change in AIC relative to model A.

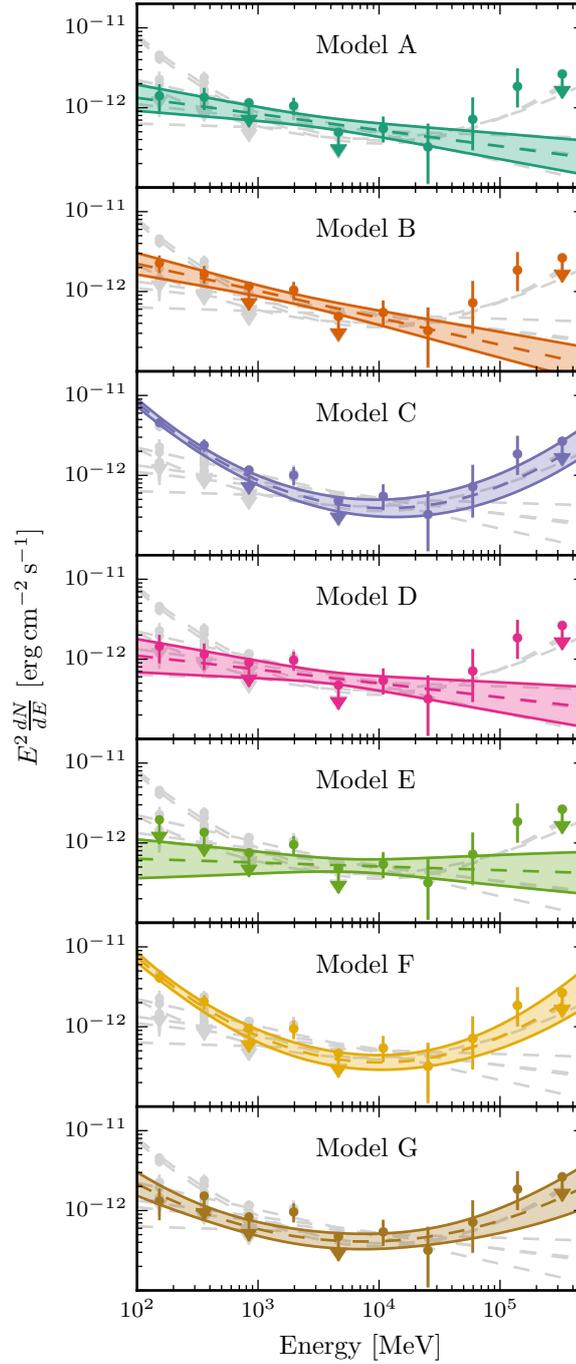


Figure 4.2: The variation of the SED of IC 310 for each model considered in Table 4.3. Each SED was calculated by profiling the normalisation of IC 310 in 10 logarithmically-spaced bins between 100 MeV and 500 GeV. It is readily seen that only the 3 low energy bins are significantly affected by this modelling, the source of our scepticism in Sects. 4.4.1, 4.4.2 and 4.3. Upper limits are shown if the flux density in an energy bin is consistent with 0 within 2σ .

Power Law Models				
Model	Prefactor [MeV cm ⁻² s ⁻¹]	Index	Scale [MeV]	
A	$(2.96 \pm 0.60) \times 10^{-13}$	1.9	6.42×10^3	
B	$(6.38 \pm 1.47) \times 10^{-13}$	2.34 ± 0.163	6.42×10^3	
D	$(8.13 \pm 1.57) \times 10^{-15}$	2.17 ± 0.19	6.42×10^3	
E	$(7.87 \pm 1.59) \times 10^{-15}$	2.05 ± 0.193	6.42×10^3	

Log Parabola Models				
Model	Prefactor [MeV cm ⁻² s ⁻¹]	Alpha	Beta	Eb [MeV]
C	$(5.4 \pm 1.01) \times 10^{-13}$	2.66 ± 0.08	$-(0.136 \pm 0.027)$	1×10^3
F	$(4.78 \pm 1.41) \times 10^{-13}$	2.66 ± 0.11	$-(0.142 \pm 0.033)$	1×10^3
G	$(3.66 \pm 0.81) \times 10^{-14}$	2.36 ± 0.09	$-(0.089 \pm 0.028)$	1×10^3

Table 4.4: The parameter values obtained for the initial log parabola and power-law fits to IC 310 within the integrated dataset.

The best-fit model (both when assuming nesting and the AIC) is shown to be that with IC 310 with a log parabola fit and a negative β (upward curvature), preferred by $\sim 3\sigma$ over a power law with index 2.01. In all cases, a negatively-curved spectrum is preferred for IC 310 over any models without such curvature. Curvature in IC 310 is preferred to a softer spectrum of PS J0312.8+4121 by calculating the difference between models D and E, and models F and E. It should be noted that with its index left to vary, PS J0312.8+4121 does not achieve 5σ significance (a TS of 28.744 with 2 degrees of freedom) between models G and C. Due to these facts, and the overall significance of each model, we find that model F allows the great-

est overall significance accounting for the number of parameters introduced. Freeing the index of PS J0312.8+4121 between models F and G is shown to be overfitting, reducing the overall significance of the model in both statistical metrics. The SED of IC 310 given each model is shown in Figure 4.2. Similarly, we plot the TS map of the ROI for each of the models discussed in this Section in Figure 4.3, which shows a convincing case for localised excess emission in the region of PS J0312.8+4121, and thus the necessity of including this source in our models.

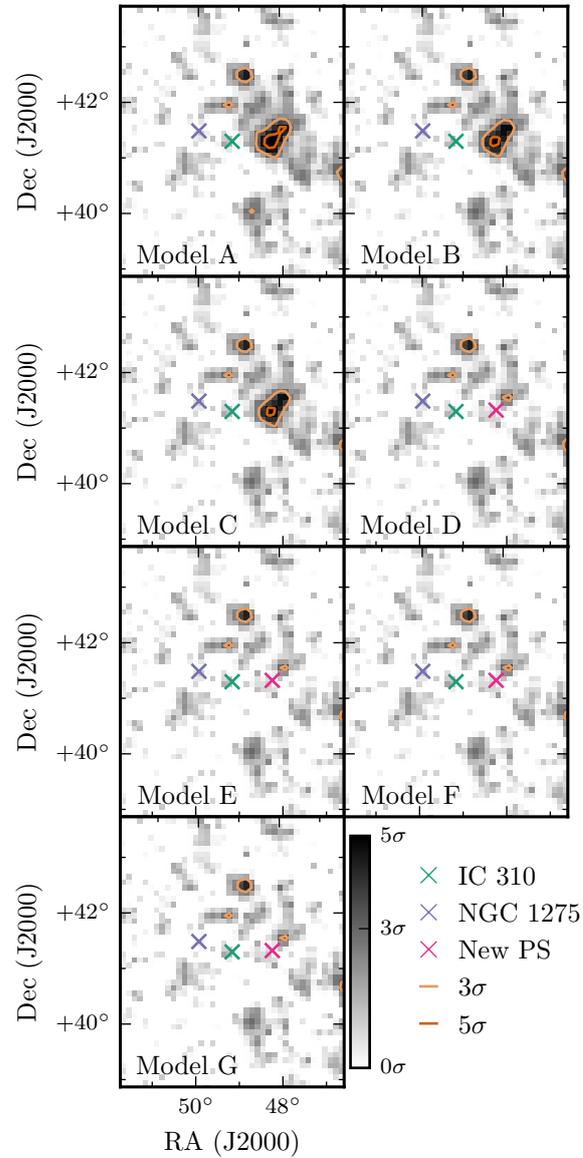


Figure 4.3: A TS map of the ROI in a 3° radius around IC 310, showing the relevant point sources for each model assessed in Section 4.2. The statistical significance of PS J0312.8+4121 can be seen as a peak before a point source model is included from model D onwards. Other potential point sources manifesting as smaller peaks can be seen, but we limit ourselves to a 5σ threshold for new point sources.

4.4 Constraining the Influence of NGC 1275

Having noted the difficulty of separating the emission of NGC 1275 and IC 310, we can first look to establish whether our detection of seemingly low-energy emission at the position of IC 310 is an artefact of the fringes of NGC 1275’s emission. The PSF of the LAT is highly energy-dependent, and the low energy photons could be “leaking” into our model for IC 310. If this were to be the case, we would expect similarly unmodelled photons unidirectionally from NGC 1275 (after accounting for the fish-eye effect this might appear distorted slightly, but not over an 8-year integration). To see if this is the case, we look to measure the TS of a test source with a soft spectrum corresponding to the spectral index measured for our model of IC 310 below 3 GeV ($\gamma = 2.4$).

In Figure 4.4 we test the TS at 120 positionings at various bearings from NGC 1275 on a circle with radius 0.6° around the 3FGL position of NGC 1275 (corresponding to a change of 3° in bearing between each test point). We constrain the energy range for this analysis to ensure all signal is associated with the low-energy behaviour and assume a soft power law index of 2.4 for IC 310 as our initial analysis indicated. We choose a lower energy bound of 578.5 MeV by calculating the 95% containment region on PSF3 class photons calculated for our ROI using the `gtpsfc` tool, and choose 3 GeV as the upper energy bound as the approximate upturn of the SED found in Section 4.2. We find a very strong peak (TS=21.73) in the direction of IC 310, along with two smaller peaks (TS = 13.2 and 3.3) at approximately (RA, Dec): $(49.403^\circ, 41.948^\circ)$ and $(50.279^\circ, 40.964^\circ)$ respectively. While this strong peak can be interpreted as the low-energy emission of a quiescent IC 310, the provenance of the two weak peaks is unclear. Given their angular offset (slightly less than 180°) and the fact

that extended emission has been observed in the radio galaxies Centaurus A and Fornax A (Abdo et al., 2010a; The Fermi-LAT Collaboration, 2016), we investigate the prospect of extended emission in Section 4.4.1. Dark matter present in the Perseus cluster could cause some extended emission, but the peaked nature of this emission contrasts with the large-scale smooth emission expected from DM annihilation in galaxy clusters (see e.g. Han et al., 2012, Fig. 1).

The lack of significance found in this plot other than coincident with IC 310 provides strong evidence that the low-energy emission previously discovered is almost certainly not an artefact of the vicinity of NGC 1275, and instead originates solely from IC 310’s immediate locale.

4.4.1 Extended emission from NGC 1275

Source confusion between IC 310 and NGC 1275 presents a problem in our analysis given their 0.6° separation. Especially at low energies where the LAT’s PSF is larger, this may present problems if NGC 1275 is spatially resolvable. Other radio galaxies have shown significant extension such as Fornax A (The Fermi-LAT Collaboration, 2016) and Cen A, as discussed in the previous Chapter. As discussed in Section 4.2, extension (if elliptical or extended along one axis) could explain the smaller peaks seen in Figure 4.4.

To investigate the systematic uncertainty posed if this were the case, we can search for extension in the point source of NGC 1275 by replacing it with a flat disc of various radii in the range $0.1^\circ - 2.0^\circ$. The disc model is defined as:

$$Disc(x, y) = \begin{cases} \frac{1}{\pi\sigma^2} & x^2 + y^2 \leq \sigma^2 \\ 0 & x^2 + y^2 > \sigma^2. \end{cases} \quad (4.3)$$

As shown by Lande et al. (2012), both disc and Gaussian extension tem-

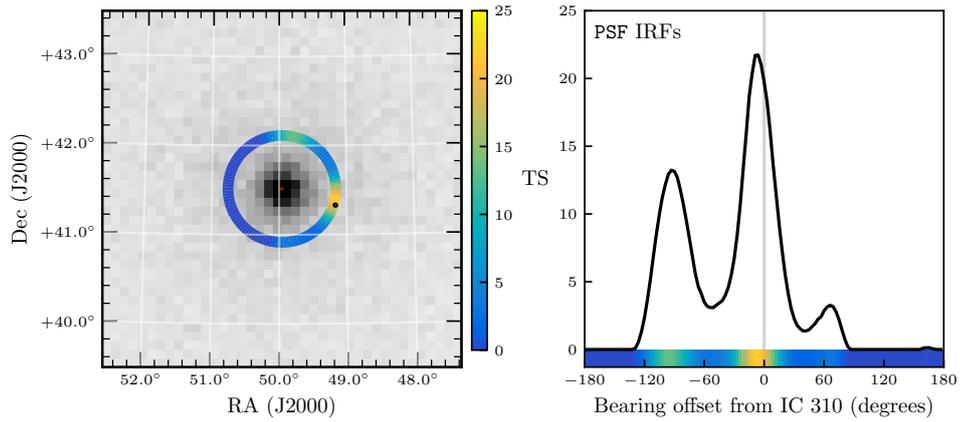


Figure 4.4: Left: The *Fermi*-LAT counts map centred upon NGC 1275 with a 1° radius using 0.1° bins. The annulus is coloured according to the TS of a point source with index $\Gamma = 2.45$ at this position in an analysis between 578.5 MeV and 3 GeV using a composite (PSF0 + PSF1 + PSF2 + PSF3) likelihood. The red ellipse represents the localised position of NGC 1275 from Section 4.4.2 and the black ellipse represents the localised position of IC 310 obtained from the 3FGL. Right: The annulus TS values unwrapped and plotted as a function of bearing relative to IC 310. The corresponding colourmap value is plotted underneath.

plates provide minimal bias when the size of the model is comparable to the size of the *Fermi*-LAT point spread function. Again, following Lande et al. (2012), we assess the significance of extension using the TS_{ext} statistic:

$$\text{TS}_{\text{ext}} = 2 \log(\mathcal{L}_{\text{ext}}/\mathcal{L}_{\text{ps}}), \quad (4.4)$$

which was found to follow a $\chi_1^2/2$ distribution. It should be noted that this distribution was only validated up to a TS_{ext} of 16 (a limitation due to statistics), thus we extrapolate no significance beyond 4σ for greater TS_{ext} values. Radial extension from 0.01° up to 2.0° were fitted, with a TS_{ext} of 0.006 and an upper limit of 0.03° extension. This is well within the *Fermi*-LAT point spread function, even at the highest energies, and rules out any significant extension in NGC 1275. There are no significant point sources (at the level of 5 standard deviations) immediately surrounding NGC 1275 that could bias this result.

At other wavelengths, far higher resolution studies of NGC 1275 could reveal the directionality of any existing extension. Previous radio observations (e.g. Lister et al. (2013)) find that the inner jet axis is aligned in the North-South plane, incompatible with asymmetric extension in a westerly direction that could explain residual emission towards IC 310. This independent evidence suggests that extension is unlikely to be the cause of our signal or even the two smaller peaks found in Section 4.4.

4.4.2 Localisation of NGC 1275

Alternatively, asymmetry as shown in Figure 4.4 could be symptomatic of poor localisation of NGC 1275. We perform a maximum likelihood position estimation on NGC 1275 after removing IC 310 from the model. While this may bias us towards IC 310's position, the resulting error is far smaller than assuming IC 310's low-energy presence given the dominance of NGC 1275's

photon statistics. Such a procedure finds a shift towards IC 310 of only 0.004° which is well within the 0.1° pixel size used in this analysis.

Using this updated position, the preference for an integrated log parabola model over a power-law remains unchanged, though the localisation will not be repeated in each time bin, in Section 4.5 as *a priori* we assume this will remain unaltered. We keep the 3FGL catalogue position of IC 310 for each time bin when exploring the variability.

4.5 Temporal Evolution

IC 310 is highly variable, not just in terms of flux, but also spectrally. This was noted by Ahnen et al. (2017) at VHE, but this should also be the case within the *Fermi*-LAT energy range. To establish the temporal effects of integrating over quiescent and flaring periods observed by MAGIC, we can break the SED from Section 4.3 down into smaller time periods. Using 360-day bins, we re-fit Model E from Section 4.3 to the data. This model was used in preference to Model F, as in only two of the eight time bins was Model F preferred to Model E at the 3σ level. We plot these SEDs in Figure 4.5, which clearly shows the variation in power-law index with time. Once again, close attention was paid to ensure that NGC 1275 and PS J0312.8+4121 are jointly optimised with IC 310. The parameter space is quite broad and some bins show only upper limits in flux. While the parameter space investigated is not large enough to fully encompass some of the time bins, it does show that in several periods IC 310 is incompatible with the harder emission previously detected. Indeed, this apparent change from soft to hard spectrum provides an explanation for the upturn in the SED found in Figure 4.1.

If we want to compare the spectra in each bin, it is more useful to look

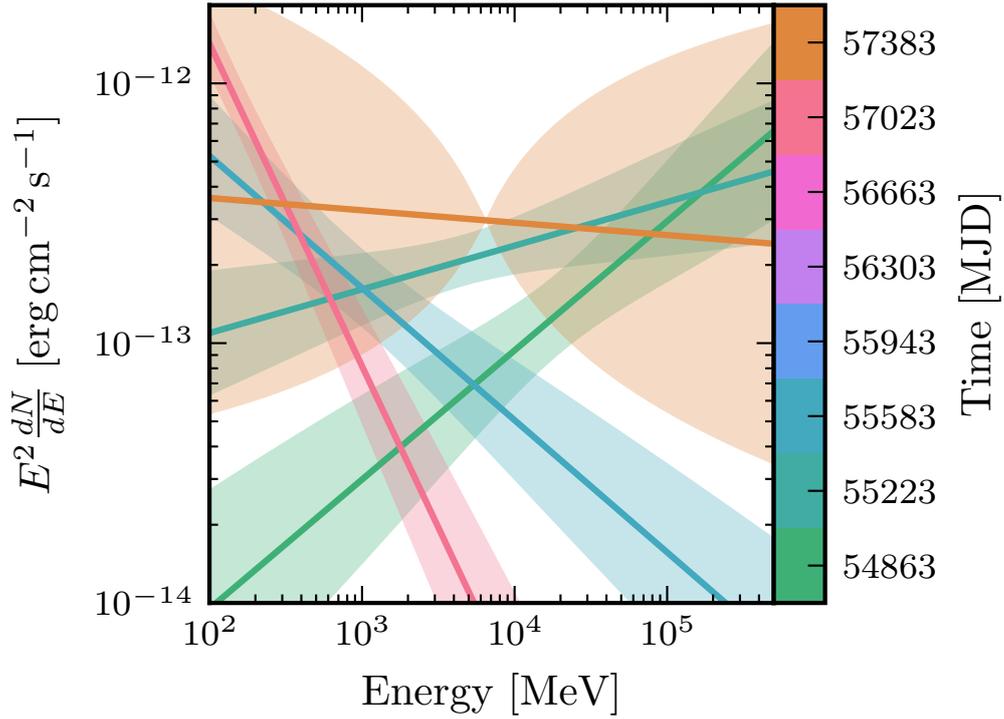


Figure 4.5: The SED for each annual bin, where each spectrum shown has a TS of at least 19.33 representing an approximate 4σ cut defined with two degrees of freedom: normalization and spectral index. The shaded areas around each optimised model represent the 1σ flux uncertainty propagated from the uncertainties in spectral model parameters after the likelihood fit.

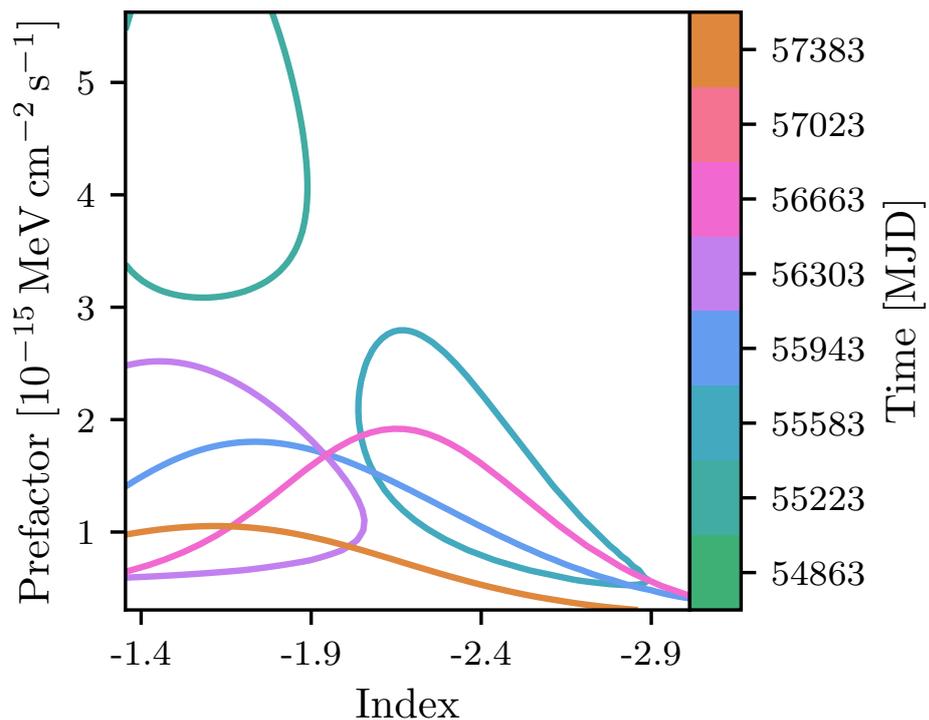


Figure 4.6: The 1σ uncertainty contour in the parameter space of a power law for each time bin assessed in Section 4.5. The normalisation (but not spectral shape) of NGC 1275 has been left free to vary for each parameter combination. Bins with optimised parameters outside of this parameter space are excluded.

at a grid scan of the two spectral parameters in each case. This is performed by exhaustively calculating the likelihood given all combinations of parameter values. The parameter space is plotted in Figure 4.6, which shows a strong separation between hard and soft states. The limits on the parameter space grid scan were prescribed as half the minimum value of the parameter in the integrated fit and double the maximum value of the parameter in the integrated fit. In each case, the parameter value is set within the `pyLikelihood` instance of our total model fit and the ROI accordingly optimised with background sources free, including a completely free spectral model of NGC 1275. This variation in spectral index is a convincing origin of the curved shape of the overall spectrum found in Sect. 4.2.

Finding the maximum likelihood in each normalisation slice of this parameter space leads to a profile of the index parameter for IC 310. These profiles are shown in Figure 4.7, and by summing the profiles with and without VHE flare detection from MAGIC (as shown in black), we find a distinct bimodality in index.

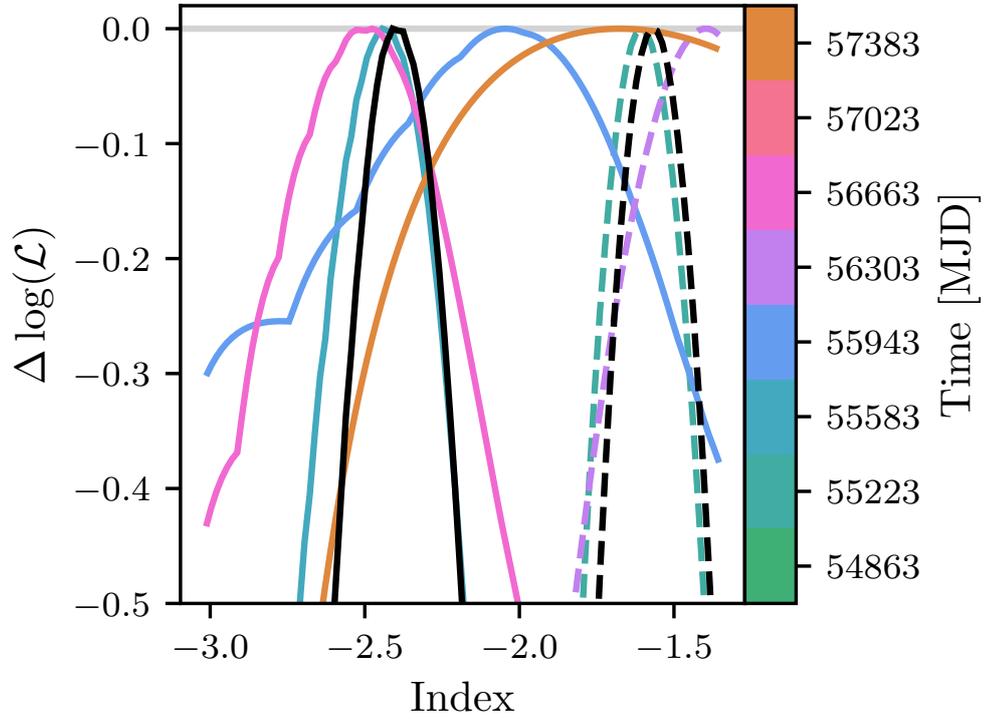


Figure 4.7: The likelihood space from Figure 4.6 profiling over normalisation. Solid profiles represent bins in which IC 310 was not detected by VHE telescopes (as determined from Table 4.5 and dashed lines represent VHE detected bins. The black lines represent the sum of the VHE detected and non-VHE detected profiles with the minimum subtracted, and thus the combined maximum. Artefacts are introduced due to the limited region over which the grid scan was performed and also due to the resolution chosen for the grid scan. Bins are chosen for comparison with Figure 4.6.

Table 4.5: Detailed information on the binnings assessed in Sect. 4.5.

Start [ISO (MJD)]	End [ISO (MJD)]	Index	TS _{IC 310}	VHE Detection
2001-08-04 (54682.66)	2009-07-30 (55042.66)	-1.50	27.45	–
2009-07-30 (55042.66)	2010-07-25 (55402.66)	-1.84	57.78	Aleksić et al. (2010)
2010-07-25 (55402.66)	2011-07-20 (55762.66)	-2.71	25.82	–
2011-07-20 (55762.66)	2012-07-14 (56122.66)	-2.53	25.03	–
2012-07-14 (56122.66)	2013-07-09 (56482.66)	-1.65	9.68	Cortina (2012) Ahnen et al. (2017)
2013-07-09 (56482.66)	2014-07-04 (56842.66)	-2.55	19.37	–
2014-07-04 (56842.66)	2015-06-29 (57202.66)	-2.05	20.37	–
2015-06-29 (57202.66)	2016-06-23 (57562.66)	-1.57	26.84	–

IRF	Event Types	TS _{Soft 1}	TS _{Soft 2}
PSF	(4 + 8 + 16 + 32)	24.22	21.68
EDISP	(64 + 128 + 256 + 512)	2.15	~ 0
FB	3	2.58	~ 0

Table 4.6: The significance of this result when varying the IRFs used in the analysis. The event types as defined by the *Fermi*-LAT collaboration are used in a composite likelihood analysis for PSF and EDISP classes, whereas the combined event class is used for the FB analysis. TS_{Soft 1} refers to the third time bin (MJD: (55402.66, 55762.66)) and TS_{Soft 2} refers to the seventh time bin (MJD: (56842.66, 57202.66)).

4.5.1 Choice of IRFs

As noted in Section 4.2, in order to decouple IC 310’s emission from that of NGC 1275, we chose to use the PSF IRFs. This choice increased our systematic error on the effective area, especially at the extreme ends of the *Fermi*-LAT energy range, by around a factor of two. Our analysis is inherently spectral, so there is also an argument that the EDISP IRFs (in which the energy dispersion is minimised) would be most appropriate. Alternatively, the combined FRONT and BACK IRFs could have provided a suitable middle-ground between energy dispersion and point spread function. In each case we will reanalyse this data while choosing a different set of IRFs and using the TS of IC 310’s soft spectral state as our figure of merit (as the hard spectral state has been independently studied by Neronov et al. (2010)). Table 4.6 shows the effect on the significance of the soft state with each set of IRFs.

The significance of the result drops significantly as the PSF of the set of IRFs increases. This is due to the inability to differentiate between IC 310

and NGC 1275 with a degraded PSF at low energies. While this vindicates our decision to use the PSF-partitioned IRFs, it does show that this analysis is sensitive to our choice of IRFs. This is not a cause for concern as the PSF-partitioned IRFs have been well-validated (to the point of being used in the recent 4FGL), but should be noted in any case. As a visual representation of the PSF of the LAT instrument at various energies, Figure 4.8 shows the variation of PSF (and data) with IRF choice. Figure 4.9 further shows that the energy at which the 95% containment radius exceeds 0.6° (the separation between NGC 1275 and IC 310) is severely reduced using the PSF3 IRFs. With the FB IRFs this energy is 928.6 MeV, but with PSF3 IRFs, this is reduced to only 578.5 MeV, almost a factor of two improvement.

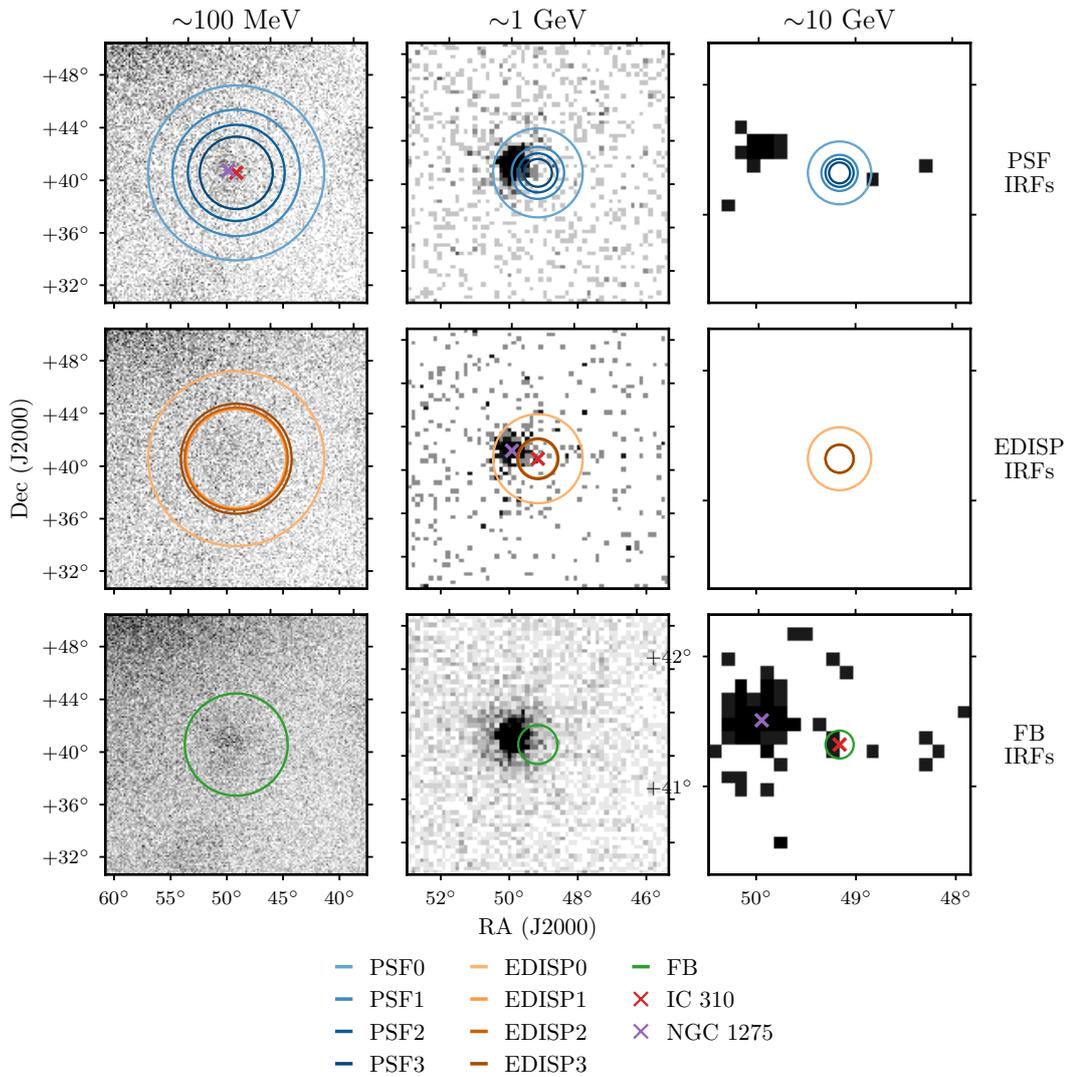


Figure 4.8: The variation in PSF with IRF choice and energy. The data in the top row correspond to data partitioned into the PSF3 IRF class, the data in the middle row into EDISP3, and the data in the bottom row in the summed FB class (all photons would fall into this category). Data in the first column are binned in energy from 100–126 MeV, data in the second column from 1 GeV to 1.26 GeV, and data in the final column from 10–12.6 GeV. The image is zoomed between energy ranges for easier viewing, decreasing from 10°, to 3°, to 1° with increasing energy. These data represent the first annual bin of our lightcurve. The circles plotted show the 95% containment angles of the PSF of the *Fermi*-LAT instrument using the IRFs as listed in the legend. Containment angles were calculated by numerically integrating the PSFs returned by the `gtpsf` tool within the *Fermi* science tools. Asymmetry due to the fisheye effect and inclination angles to the LAT boresight is not considered. The positions of NGC 1275 and IC 310 are marked along the diagonal to show the scales relevant to our analysis.

4.6 Discussion

We have reported the detection of a soft emission state in addition to the very energetic emission commonly associated with IC 310. While noted for these bright flaring periods, by inspecting the SED of IC 310 over the lifetime of *Fermi* to date we find this soft emission is concentrated between periods in which strong flares have been observed.

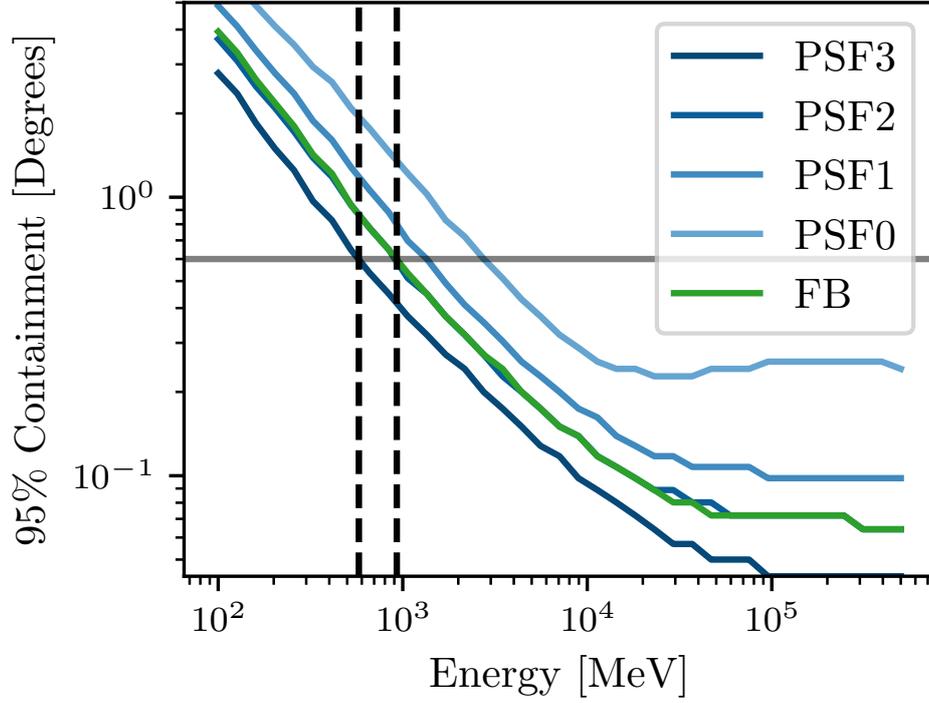


Figure 4.9: 95% containment radii for our region of interest as a function of energy and instrument response function partition. Each line represents an individual IRF class, coloured as described in the legend. Vertical lines represent the 95% containment regions of the PSF3 and FB classes, and the horizontal line represents 0.6° , the separation of IC 310 and NGC 1275. Artefacts are due to the finite resolution used in the `gtpsfr` tool used to calculate these values.

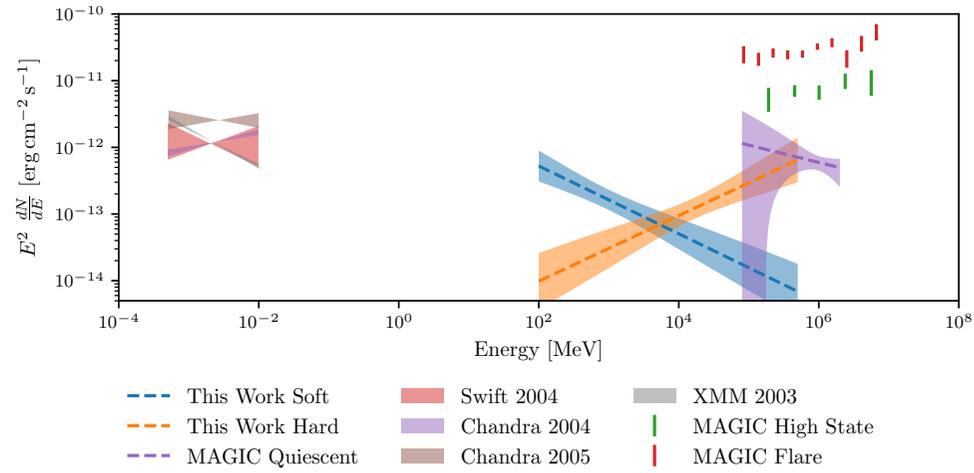


Figure 4.10: The analysis from this work plotted with the data presented in Aleksić et al. (2014a). The ‘hard’ spectrum plotted from this work represents the first bin of our analysis (MJD: 54682.66, 55042.66), and the ‘soft’ spectrum represents the third bin (MJD: 55402.66, 55762.66). These data are not simultaneous.

The two states, flaring and quiescent, are most recognisably distinct in their spectra. The peak of the inverse-Compton emission during flaring periods was undetected even above 10 TeV, while our analysis has found the inverse-Compton peak to be lower than 100 MeV during quiescent flux states under the assumption of a synchrotron self-Compton model. This represents an increase of more than 5 orders of magnitude in inverse-Compton peak energy on time-scales of a year. Similar spectral variance has been seen over multiple X-ray observations as shown in Figure 4.10, which could represent the same behaviour for the synchrotron peak.

Without simultaneous multiwavelength observations to constrain the broadband SED it is almost impossible to make conclusions as to how this bimodal behaviour can be modelled. Recent simultaneous observations by Ahnen et al. (2017) find that a simple SSC model can sufficiently explain the broadband behaviour of IC 310 using different viewing angle estimates. Interestingly, the spectral indices of the analysis presented in this paper and the index of the VHE quiescent emission detected by Ahnen et al. (2017) are consistent. However they predict a peak IC energy around 80 GeV, which would disagree with the downward-going spectrum detected here with *Fermi*-LAT. It should be noted that the measurements presented in this work and in Ahnen et al. (2017) are not contemporaneous, so cooling of the electron population could explain this change in peak energy.

To gain simultaneous observations of the Compton and inverse-Compton peaks, X-ray instruments must integrate exposure on IC 310 over the year-long time periods necessary for significant detection using *Fermi*. Assuming the SSC paradigm, we can make some physical inferences. The cooling time-scales of the electrons emitting gamma-rays at GeV energies are far shorter than the time periods (year-long bins) that we consider, and the flares could be explained (at least in the *Fermi* energy range) by a Doppler

factor increase due to changes in the jet. Without simultaneous broadband observations, however, these inferences remain speculation but can be explained by a plausible mechanism.

Whilst we have shown that the integrated spectrum of IC 310 does indeed show a negative β parameter, this should not be interpreted as an intrinsic property of the electron energy distribution of IC 310. This would have to result from an upturn in energy spectrum of accelerated particles moving down the jet, an unlikely prospect at the level required for these results. This is unambiguously an artefact of time integration.

4.7 Conclusions

The main conclusions of this analysis are as follows:

- IC 310's integrated spectral energy distribution as measured by the *Fermi*-LAT instrument shows an upwardly curving spectrum as in the case of Cen A in Chapter 3.
- Unlike Cen A, IC 310 shows spectral variability in the 365-day bins used in this analysis. Two states of soft ($\gamma \approx 2.5$) and hard ($\gamma < 2$) emission are evident in this analysis, and are supported by non-contemporary measurements in the X-ray and VHE gamma rays.
- Taking the contemporaneous data into account, this bimodality represents a change in the second peak of the broadband SED from <100 MeV to >10 TeV, over 5 orders of magnitude.

While we have discovered extremely interesting behaviour in the spectrum of IC 310 which is corroborated at other wavelengths, the lack of long-term multiwavelength monitoring does not allow meaningful physical

modelling of this source. Interestingly, our finding that multiple episodes of flaring activity superimposed with the quiescent flux of IC 310 has resulted in a significant spectral artefact (an upwards curve to the SED), poses further questions about the validity of long-term integrations of LAT-detected sources that are used for studies of cosmological quantities and tests of beyond the standard model physics using LAT data. This will be explored further in Chapter 6.

Spectral Variability in Blazars

5.1 Introduction

So far in Chapters 3 and 4, evidence has been presented of odd spectral characteristics in two different AGN. While in the first case, there was a complete lack of variability in the source, in the second case this upturn in the spectrum was an artefact of a bimodality in the spectrum. In this chapter, we will look to measure the spectral variability in bright LAT-detected AGN and attempt to characterise the variability in both flux and SED seen amongst this population.

Many attempts to catalogue the gamma-ray sky have been made in both high and very high energy regimes. Those made using the *Fermi*-LAT instrument are the previously mentioned point source catalogues, namely the 0FGL–4FGL, the high-energy catalogues 1FHL–3FHL, the inner galaxy catalogue FGES, and the extended source catalogue FHES, which are summarised in Table 5.1 along with their respective references. In addition to these source-finding catalogues, additional, more focused, catalogues such as the LAT-detected AGN catalogues (the LAC catalogues: Abdo et al. (2010a), Ackermann et al. (2011), and Ackermann et al. (2015a)) also exist.

These contain no new analysis, but incorporate other useful information such as multiwavelength measurements and redshifts.

Name	Energy [GeV]	Data [Month]	# Sources	Event Selection	Lightcurve	Reference
0FGL	0.2–100	3	205 (Point)	P6V1 DIFFUSE, FB	–	<i>Fermi</i> -LAT Collaboration (2008)
1FGL	0.1–100	11	1451 (Point)	P6V3 DIFFUSE, FB	Monthly	Abdo et al. (2010b)
2FGL	0.1–100	24	1873 (Point)	P7V6 SOURCE, FB	Monthly	Nolan et al. (2012)
1FHL	10–500	36	511 (Point)	P7V6 CLEAN, FB	–	Ackermann et al. (2013b)
3FGL	0.1-300	48	3033 (Point)	P7V15 SOURCE, FB	Monthly	Acero et al. (2015)
FGES	10–2000	72	46 (Extended)	P8 SOURCE, PSF	–	Ackermann et al. (2017b)
2FHL	50–2000	80	360 (Point)	P8 SOURCE, FB	–	Ackermann et al. (2016)
3FHL	10–2000	84	1556 (Point)	P8 SOURCE, PSF	–	Ajello et al. (2017)
FHES	1–1000	78	24 (Extended)	P8 SOURCE, PSF	–	Ackermann et al. (2018)
4FGL	0.05–1000	96	5098 (Point)	P8 SOURCE, PSF	–	The Fermi-LAT collaboration (2019)

Table 5.1: A catalogue of catalogues. Relevant metadata of the released LAT collaboration catalogues.

The original catalogue of *Fermi* AGN to investigate variability (Abdo et al., 2010c) looked in great depth at the power spectral density (PSD), structure function, and individually flaring periods in the first two years of *Fermi*-LAT exposure. A comparison between those AGN classified as FSRQs and BLLs within initial week-long lightcurves was made, and while the two populations had slightly different power-law indices, they could not be significantly separated. Ackermann et al. (2011) looked again at variability using the 2FGL catalogue (Nolan et al., 2012) and the monthly lightcurves that were published alongside. In this case, “little evidence” could be found to separate these populations. Additionally, the statistics to define variability, in this case the TS_{var} as used in Section 3.4, was developed and used for the first time. As we consider further extensions of the concept of variability in this Chapter, we shall look to extend these definitions now that a fourfold increase in data has allowed deeper analysis.

5.2 Data Analysis

The analysis procedure uses many of the same choices described in Chapters 3 and 4 when it comes to the instrument response functions applied and data reduction parameters. Table 5.2 is provided for a full summary of the choices made for all data analysis in this Chapter. The sole difference at this initial data selection stage regards the smaller size of ROI in each case (10° as opposed to 20°). Large ROI (15° and above) are often used to avoid the emission from sources outside of the ROI “leaking” emission (usually when flaring and not accounted for by the catalogue model) into the ROI and causing spurious signals when searching for weak sources. In this case, our objects are extraordinarily bright and almost always the brightest object in the region. In this case, such effects will be small unless the flaring event in

Science Tools Version	v10r0p5
Pass	P8R2
ROI Width	10°
Event Class	SOURCE (128)
Event Type	PSF (4,8,16,32)
Zenith Cut	< 90°
Time Selection Filter	(DATA QUAL>0)&&(LAT CONFIG==1)
Energy	100 MeV–500 GeV
Time	(2008-08-04 15:43:38)–(2016-08-08 06:01:06) UTC
Galactic Model	gll_iem_v06.fit
Isotropic Model	iso_P8R2_SOURCE_V6_PSF{0..3}_v06.txt

Table 5.2: A summary of the data selection parameters used within the *Fermi* science tools in Chapter 5.

an object $> 10^\circ$ away is of truly colossal magnitude. For comparison, the 99% containment region is approximately 1° at 1 GeV, the energy at which these measurements are most sensitive.

While in the previous two chapters we considered only one AGN at a time, we now look to a population. The aim is to study the variability of extragalactic objects. Taking the 3FGL catalogue as a base sample, we then select point sources associated with extragalactic object. This is performed by selecting the following source categories from the 3FGL catalogue: *agn*, *bcu*, *bll*, *fsrq*, *gal*, *nlsy1*, *rdg*, *sbg*, *sey*, *ssrq* and their capitalised counterparts (lower case designations represent weak associations, while upper case represents corroborated firm identifications). The definitions of these categories are summarised in Table 5.3. This represents a complete selection of extragalactic objects from the catalogue.

The first criterion ensures that contamination from the Galactic inter-

Designator	Description
AGN	Non-blazar active galaxy
BCU	Blazar candidate of uncertain type
BLL	BL Lac type of blazar
FSRQ	FSRQ type of blazar
GAL	Normal galaxy (or part)
NLSY1	Narrow line Seyfert 1
RDG	Radio galaxy
SBG	Starburst galaxy
SEY	Seyfert galaxy
SSRQ	Soft spectrum radio quasar

Table 5.3: Galaxy classification designations used in the 3FGL catalogue as defined by Acero et al. (2015) (Section 4.1, Table 6). Only those designations used in this analysis have been extracted to present here.

stellar emission is only a very minor concern, especially given the prominence of our source selection. The vast majority of sources at $|b| < 10^\circ$ are densely packed Galactic sources, complicating analysis.

Our second criterion on detection significance in the 3FGL is used to ensure that when a monthly lightcurve with a free spectral model is produced, the source is still likely to be detectable on the timescale of a month. Even so, this may result in some bins in which the source is undetectable despite this strong limit on significance.

The final criterion ensures that correlations between lightcurves do not exist in the final dataset due to two sources in close proximity. While our method can distinguish between two nearby sources, the cross correlations between lightcurves can be significant unless the systematic uncertainty is cared for properly. For the sake of a clean sample, we exclude the 8 sources listed in Table 5.4 which were found to be within 1° of each other (in this situation, *both* sources are omitted), leaving 294 other sources.

The analysis procedure for each source begins with a simple analysis using individual ROI selections as described in Table 5.2. As in Chapters 3 and 4, we run an initial optimisation, followed by a source finding step, followed by a further optimisation and a removal of weak point sources with $TS < 2$.

Monthly lightcurves for each source are calculated for each source in the sample by scanning the log likelihood of the model when profiling out the spectral parameters of all nearby sources (within 5°). Sources further than 5° away from the centre of the ROI are fixed to their values from the 8-year fit and not optimised in each bin individually. We also calculate separate lightcurves, leaving the spectral parameters of the source free to vary in each bin. This final step allows a full investigation of the spectral variability of the AGN in the sample.

Sep. ($^{\circ}$)	Source 1	Source 2
0.462	NRAO 512	3C 345
0.712	1ES 2322-409	PKS 2325-408
0.760	1ES 1215+303	PG 1218+304
0.895	PMN J2345-1555	PKS 2345-16

Table 5.4: A list of the sources removed from the catalogue selection due to their proximity.

5.3 Integrated Spectra

When this analysis was first completed, the 8-year *Fermi*-LAT catalogue had not been published. As this has since been released (The Fermi-LAT collaboration, 2019), we are provided with an opportunity to compare our integrated 8-year results to an independent analysis before we move on to the results of the variability study. The 4FGL uses largely the same analysis method as in this work, but includes a completely redeveloped Galactic interstellar emission model. Through this comparison, we can estimate the remaining systematic uncertainty introduced to our results because of the use of an older Galactic map despite our $b > 10^{\circ}$ requirement.

The spectral indices of the 174 sources modelled as power laws in the 3FGL (about 60% of our sample) are plotted in Figure 5.1. The mean fractional deviation $(\gamma_{4\text{FGL}} - \gamma_{\text{This Work}})/\gamma_{4\text{FGL}} = 0.0027$. Sub-% disagreement is as expected from small changes in background model and slightly different data ranges and selections (which are on the percent level). We looked for any variation with Galactic latitude, but no such proportionality could be found. This result gives us confidence going forward that we can reproduce the quality of *Fermi* catalogue results through our detailed spectral analysis.

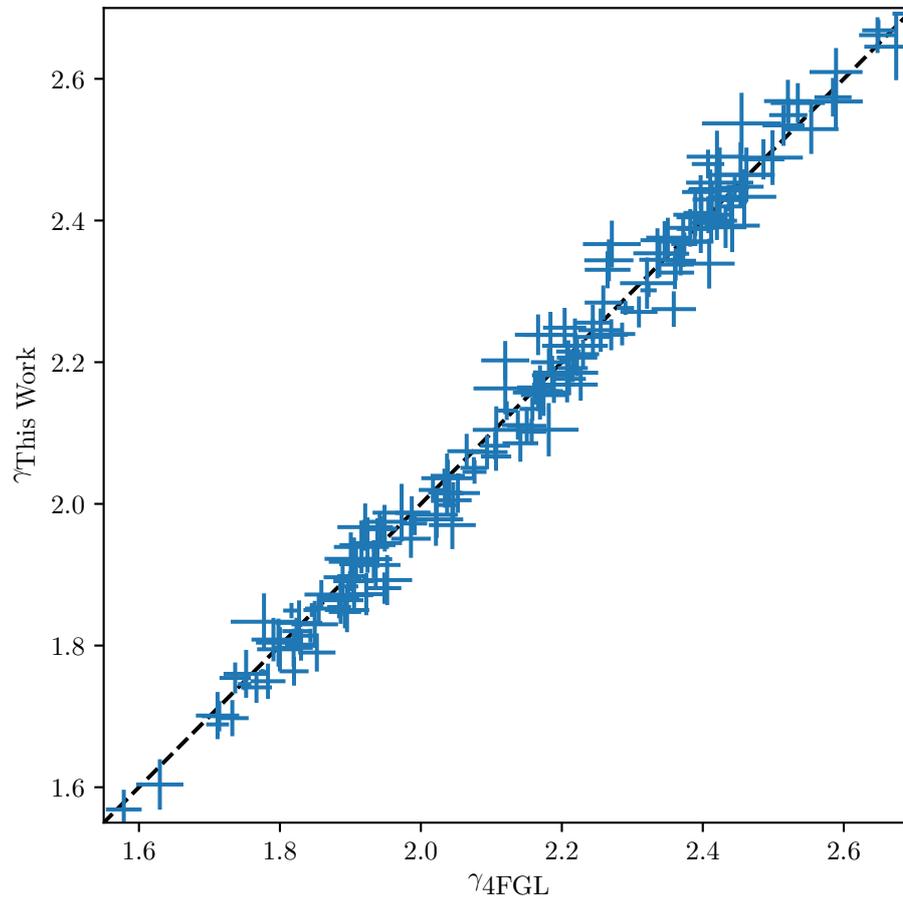


Figure 5.1: The measurement of photon index as measured in this sample compared to the same sample from the more recent 4FGL catalogue. The dashed black line represents a one-to-one correspondence between the 4FGL and this work.

More in depth spectral analysis can also provide further points of comparison. Taking the sources in the 3FGL that were assigned power-law models, we applied a model comparison between a power-law and log parabola with the expanded data set. Using the $TS_{\text{curve}} = 9$ limit now applied in the 4FGL (relaxed from $TS_{\text{curve}} = 16$ in the 3FGL), we find 108 of the 176 sources originally modelled as power laws in the 3FGL are better suited by a log parabola. The 4FGL itself finds that 122 of these sources meet its $TS_{\text{curve}} = 9$ criterion. In order to reproduce 122 sources with our comparison, this criterion needs to be relaxed only to $TS_{\text{curve}} = 8$. While there is a slight difference in significance of curvature in this case, it can again easily be explained by the small offset in data periods. Most importantly, this shows our improved ability to observe and resolve spectral features with the LAT using Pass 8 data and improved instrument response functions (and of course more integration time).

5.4 AGN Variability

Blazar variability is typically studied in a series of sub-populations (e.g. by Ackermann et al., 2015a), with the most common dividing line being that between flat spectrum radio quasars (FSRQs) and BL Lac objects (BLLs). The distinction of a BLL from an FSRQ is generally defined by the following characteristics:

- The rest frame equivalent width of the strongest optical emission line is $< 5 \text{ \AA}$,
- The optical spectrum shows a Ca II H/K break ratio $C < 0.4$ (Landt et al., 2002),
- Wavelength coverage of the optical spectrum satisfies $(\lambda_{\text{max}} - \lambda_{\text{min}}) / \lambda_{\text{max}} >$

1.7.

These are the conventions of Stocke et al. (1991) and Urry & Padovani (1995) which have been applied to the data of the 3FGL in Ackermann et al. (2015a). This provides optical classifications which will be used for our survey. The characteristics defined above generally ensure that those objects we describe as BLLs are more completely jet-dominated than FSRQs, which show stronger line emission. The final criterion simply ensures that the optical spectrum used for the determination is of sufficient quality that no strong lines have been missed.

While the power spectral density (PSD) of variability in AGN at gamma-ray wavelengths is that of a red (or pink) power spectrum regardless of the optical characterisation, the exact timescale of variability (and hence gradient of the PSD power law) depends entirely on the injection of new particles into the jet or internal shocks. The increased activity in the accretion discs of FSRQs could provide a measurably different distribution in flaring timescales compared to that of BLLs. Whether the previously described distinction in equivalent width and C is truly sufficient to distinguish between objects with intrinsically weak-excitation and high-excitation lines simply outshone by particularly bright non-thermal emission is a bone of contention (Ghisellini et al., 2011). Hence it is possible that our classifications in this case are inaccurate in the case of very strongly jetted objects, and it should be noted that this shortcoming cannot be overcome through our gamma-ray analysis alone.

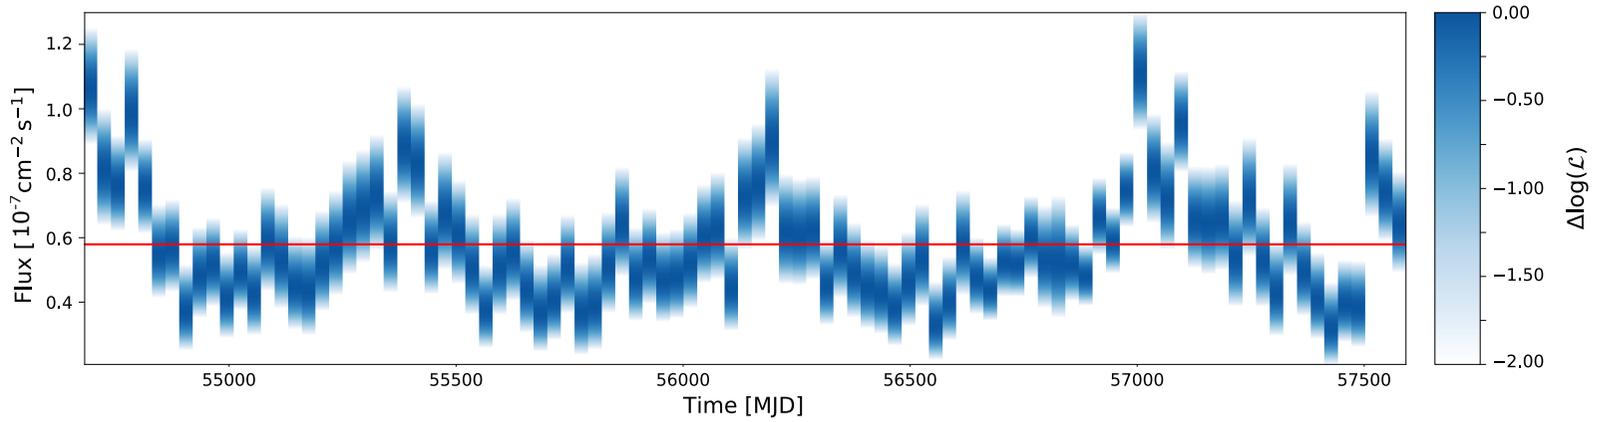


Figure 5.2: The lightcurve of PG 1555+113 from our catalogue analysis as an example. The colour gradients show the $\Delta \log(\mathcal{L})$ from the maximum likelihood value. A change in $\log(\mathcal{L})$ of 0.5 represents the 1σ uncertainty interval, and a change of 2.0 represents the 2σ uncertainty interval. This is true as our scan of the normalisation (a direct mapping to flux) is distributed as $\chi^2_{n=1}$. The red line represents the measured (average) flux value over the 8-year integrated exposure.

We study the timescale of variability of our selection first by taking the lightcurves calculated in Section 5.2, and calculating the TS_{var} as defined by Nolan et al. (2012) and used previously in Section 3.5. Over the course of the 8-year dataset, we find that 268 of 294 sources are variable at the 5σ level of certainty over the 96 month period.

The most direct calculation of variability given that our dataset is binned evenly is that of the power density spectrum (P), calculated as:

$$P(f_k) = \frac{1}{N^2} [|C_k|^2 + |C_{N-k}|^2] \quad k = 1, 2, \dots, \left(\frac{N}{2} - 1 \right), \quad (5.1)$$

where N is the number of bins, f_k is the frequency in the k th bin, and C_k is the discrete Fourier transform:

$$C_k = \sum_{j=0}^{N-1} c_j e^{2\pi i j k / N} \quad k = 0, \dots, N - 1, \quad (5.2)$$

where in our case c_j are our flux values. This PSD is normalised to variance divided by the squared mean flux per unit frequency to allow comparisons with previous *Fermi* collaboration studies. These values were summed to account for the populations of FSRQs and BLL objects as a whole in Figure 5.3, and then subsequently to compare between LSP, ISP, and HSP BLL objects in Figure 5.4. This PSD analysis follows the procedure followed by Abdo et al. (2010d), except for the longer time periods considered in this work. Uncertainty regions shown result purely from the statistical uncertainty on flux in each bin for each object, and do not account for the systematic uncertainties inherent due to the limited sample size and bias due to bright AGN, assuming that as a population any individual periodic behaviour is averaged out. Similarly, the stochastic nature of the measurement due to the finite time range is unaccounted for by the presented uncertainty.

AGN power density spectra result from stochastic processes related to accretion disc activity and interactions within the AGN jet, which usually

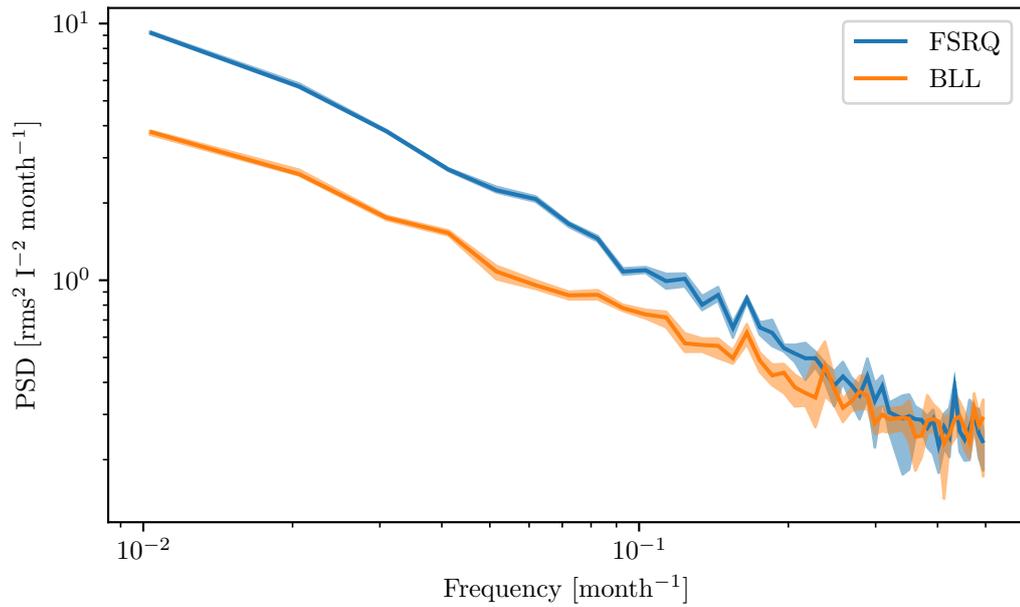


Figure 5.3: Power spectral density of FSRQ and BLL objects with their associated uncertainty. Uncertainties were calculated using Monte Carlo simulations using the uncertainty in flux measurements, and the 32nd and 68th percentiles of the resultant PSD distributions were taken as the $\pm 1\sigma$ uncertainty region.

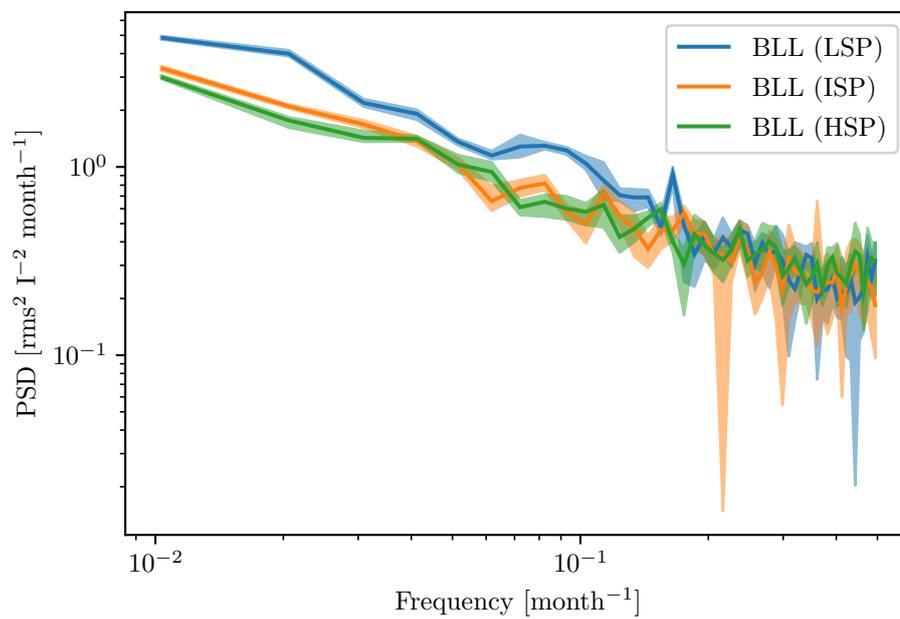


Figure 5.4: As Figure 5.3, but showing BLL objects classified by their synchrotron peak frequency.

Class	PL Index
FSRQ	-0.948 ± 0.017
BLL	-0.867 ± 0.017
BLL (LSP)	-0.904 ± 0.035
BLL (ISP)	-0.886 ± 0.028
BLL (HSP)	-0.814 ± 0.027

Table 5.5: Summary of linear fits to PSD slopes in Figures 5.3 and 5.4.

show $P(f) \propto 1/f^\alpha$ dependencies representing a “pink” power spectrum. Fitting a log-linear model to each classification as above results in a significant difference in PSD slope for every case. These slope values are summarised in Table 5.5. Not only are FSRQs and BLLs significantly different in slope, where FSRQs have far more power over long timescales (but similar short-term flaring behaviour). We also find significant differences between subclasses of BLL objects sorted according to their synchrotron peak frequency. Note that in Table 5.5 and in Figures 5.3 and 5.4 the 31st bin is excluded as it contains the 53.4 day precession period of the *Fermi* spacecraft.

These power-law values show that FSRQ objects show stronger variability on long timescales than BLL objects, corresponding to a more variable baseline level of activity, but the relative frequency of short flare-like events is similar between the two classes. When approaching the BLL class by peak synchrotron frequency, we find a significant progression from LSP to HSP objects, whereby LSPs similarly show stronger long-term variability compared to HSPs. The ISP spectrum however can only show some evidence ($\sim 2\sigma$) of separation from HSPs.

5.5 A Statistical Basis for Spectral Variability with *Fermi*-LAT

As we have previously discussed, the techniques for identifying variability with the LAT are well-established. Most commonly this takes the form of the TS_{var} although prior to this, F_{var} (fractional rms variability amplitude; Vaughan et al., 2003) and $\sigma_{\chi^2_{\text{S}}}$ (excess variance; Vaughan et al., 2003) were both commonly used. As the TS_{var} is particularly simple to apply given the analysis procedure used for LAT data (a likelihood approach), we shall construct a similar statistic which can be applied to assess spectral variability.

The TS_{var} is defined as the hypothesis test between the following hypotheses:

- Null hypothesis: The normalisation of the spectral model is constant in each bin of the lightcurve.
- Alternative hypothesis: The normalisation of the spectral model varies with time.

This hypothesis test corresponds to the following test statistic:

$$\text{TS}_{\text{var}} = 2 \sum_i^n [\log(\mathcal{L}_i(F_i)) - \log(\mathcal{L}_i(F_{\text{const}}))], \quad (5.3)$$

where the subscript i represents each individual bin, and F_i and F_{const} are the models with best fitting flux in the bin i and the “constant” value respectively. This statistic is distributed as χ^2_{n-1} as we are testing a hypothesis with $n - 1$ parameters more than a constant-valued lightcurve.

We can define two further statistics, which tell us two different properties of our sources. Firstly, we have the hypothesis test:

- Null hypothesis: The normalisation and spectral parameters of the source are constant in each bin of the lightcurve.

5.5. A STATISTICAL BASIS FOR SPECTRAL VARIABILITY WITH *FERMI-LAT*

- Alternative hypothesis: The parameters of the spectral model vary significantly with time.

Translating this hypothesis test into the form of Eqn 5.3, we obtain:

$$\text{TS}_{\text{full var}} = 2 \sum_i^n [\log(\mathcal{L}_i(N_i, \mathbf{\Gamma}_i)) - \log(\mathcal{L}_i(N_{\text{const}}, \mathbf{\Gamma}_{\text{const}}))], \quad (5.4)$$

where N_i, N_{const} represent the normalisation of the model when optimised in each bin and the averaged value respectively. The vectors $\mathbf{\Gamma}_i, \mathbf{\Gamma}_{\text{const}}$ represent the spectral parameters optimised in each bin and the averaged spectral model respectively. I have named this $\text{TS}_{\text{full var}}$ as it represents a statistic that takes into account variability not just in flux, but also spectrally.

In this case, $\text{TS}_{\text{full var}}$ is distributed as $\chi_{(y+1)(n-1)}^2$ where y is the number of spectral parameters in the model. In the case of a power-law this would be $\chi_{2(n-1)}^2$ and for a log parabola would be $\chi_{3(n-1)}^2$.

While this case is certainly interesting, it is not exactly what I initially stated to be the aim - the detection specifically of spectral variability. This is in fact specified by the following statistic:

- Null hypothesis: The spectral parameters specifying spectral shape are constant in each bin of the lightcurve.
- Alternative hypothesis: The spectral parameters specifying shape vary with time.

Again, translating this into the form of a test statistic, we find:

$$\text{TS}_{\text{spec var}} = 2 \sum_i^n [\log(\mathcal{L}_i(N_i, \mathbf{\Gamma}_i)) - \log(\mathcal{L}_i(N_i, \mathbf{\Gamma}_{\text{const}}))], \quad (5.5)$$

whereby only variation in spectral parameters that specify the shape of the SED are taken into account. Inspecting the distribution of the statistic once

again, we find this is simply $\chi_{y(n-1)}^2$ where y is defined as above. The useful identity:

$$\text{TS}_{\text{full var}} - \text{TS}_{\text{var}} = \text{TS}_{\text{spec var}}, \quad (5.6)$$

also becomes apparent from Equations 5.3, 5.4, and 5.5. It should be noted that our notation has slightly changed with the introduction of spectral parameters. For simplicity, stating the equality $\mathcal{L}_i(F_i) = \mathcal{L}_i(N_i, \mathbf{\Gamma}_{\text{const}})$ should clarify the validity of Eqn 5.6.

The distribution of each of the statistics described here assumes that when a source is detected, the spectral parameters of nested models can be treated with the standard χ^2 apparatus. This distribution was confirmed explicitly for the index of power-law models by Monte Carlo simulation by Mattox et al. (1996), and for the log-parabola follows directly from Wilks’s theorem (Wilks, 1938).

5.6 Spectral Variability of AGN

Applying these statistics to the lightcurves we created in Section 5.2 and plotting the results in Figure 5.5, we find that the vast majority of sources is variable in flux at greater than 5σ significance (irrespective of their optical class or synchrotron peak). Similarly this holds for the “full variability” we defined in the previous section with $\text{TS}_{\text{full var}}$. A divergence is apparent, however with spectral variability ($\text{TS}_{\text{spec var}}$); only approximately half of these sources show significant spectral variability 5σ level.

The likely reason that spectral variability is seen far more rarely than flux variability in our sources is statistics. The LAT is far more sensitive to changes in flux than it is to precise measurements of spectral index, and so much more integration time (or far brighter sources) is necessary for a robust detection of spectral variation. Indeed, when we investigate the

5.6. SPECTRAL VARIABILITY OF AGN

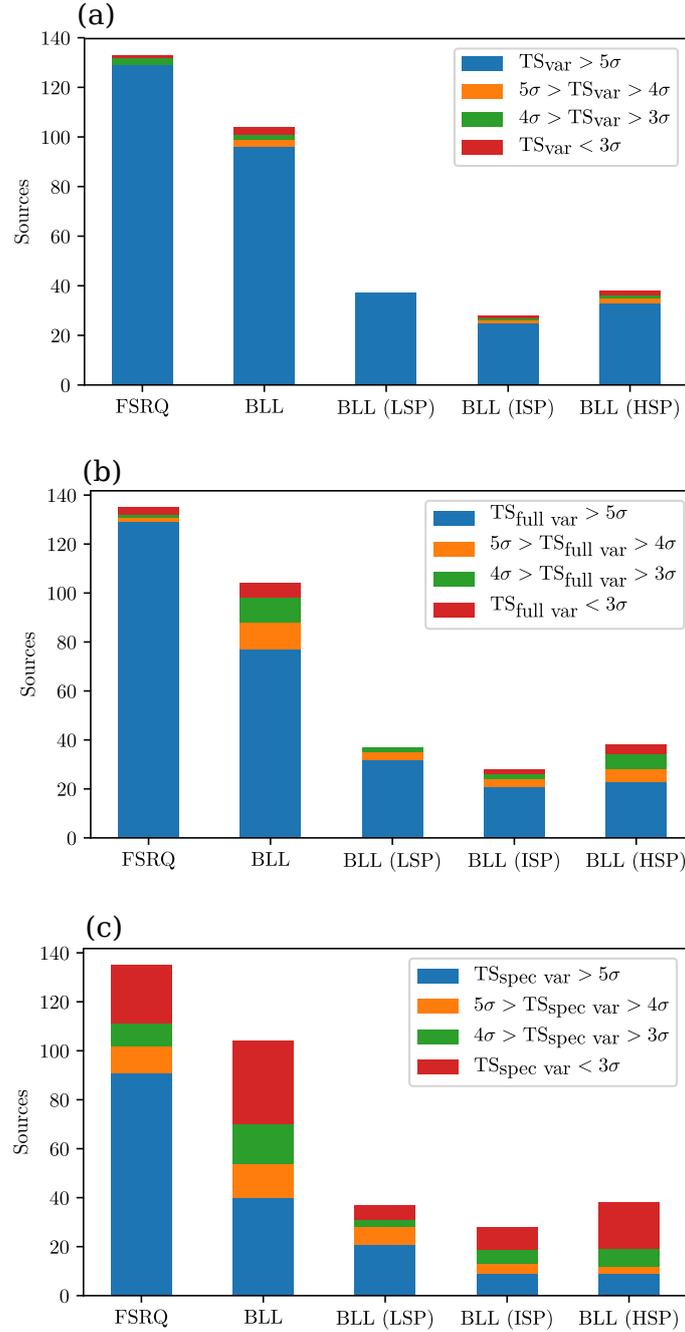


Figure 5.5: Significance of (a) TS_{var} (b) $TS_{\text{full var}}$ and (c) $TS_{\text{spec var}}$ by optical class and in the case of BLL, the peak synchrotron frequency. The corresponding values are not plotted for FSRQs as they are dominantly LSPs.

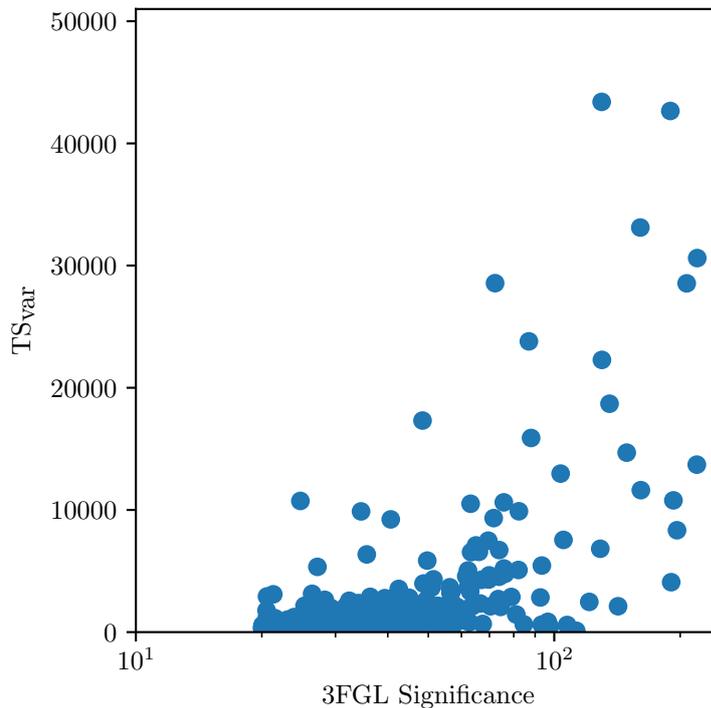


Figure 5.6: Flux variability in the 8-year lightcurves compared to their significance in the seed 3FGL catalogue.

relationship between significance of detection of our source (as determined from the 3FGL) and the significance of variability in the 8-year monthly lightcurves in Figures 5.6 and 5.7, there is some correlation, but not as strong as we might suspect. Being a bright AGN is not necessarily the most important factor in the detection of variability.

Instead, we might look towards how the $TS_{\text{spec var}}$ varies as a function of integration time in each AGN. Figure 5.8 shows the growth of $TS_{\text{spec var}}$ with increasing time (shown in units of σ given the appropriate degrees of freedom by bin). Clearly while each AGN may be in a different state, the population unambiguously approaches significant spectral variability with further observation. We must conclude from this that spectral variability

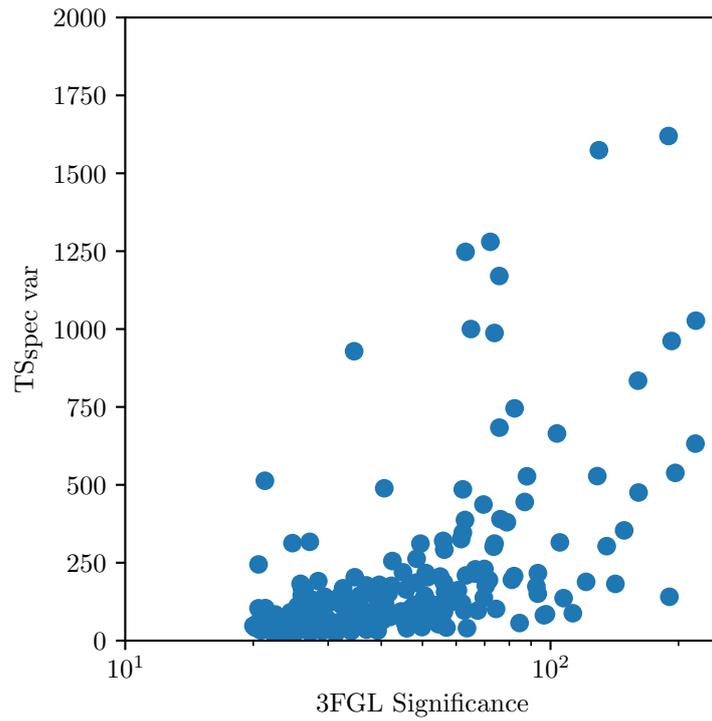


Figure 5.7: As in Figure 5.6 but showing spectral variability. This plot is however only illustrative as each point will have differing degrees of freedom due to their varied spectral models.

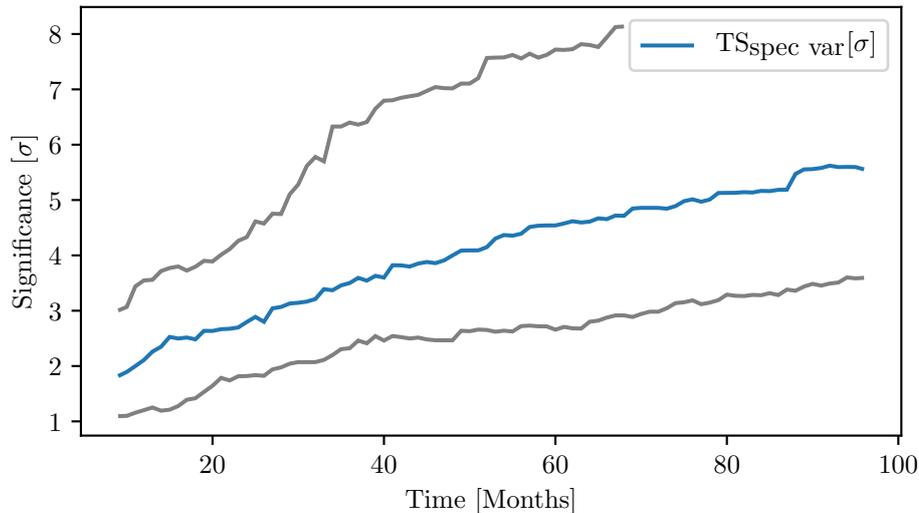


Figure 5.8: Significance of $TS_{\text{spec var}}$ with time. Blue: The median $TS_{\text{spec var}}$ value of the AGN sample. Grey: The 32nd and 68th percentile ($\approx 1\sigma$) $TS_{\text{spec var}}$ values amongst the AGN sample. Values greater than 8σ are excluded due to the limitations of the `scipy erf-1` routine.

is most definitely present in gamma-ray AGN, but can only be robustly detected with the long-term measurements that are now possible given 8+ years of data. With further integration time meaning not only greater statistics, but also a longer period in which to probe the AGN duty cycle, it seems inevitable that spectral variability is as common as flux variability (effectively ubiquitous) in gamma-ray AGN.

5.7 Characteristics of Spectral Variability

Spectral variability in AGN is usually found to follow the “harder when brighter” paradigm, and this behaviour has indeed been seen in some individual LAT-detected AGN (Brown & Adams, 2011; Angioni, 2018). Inspect-

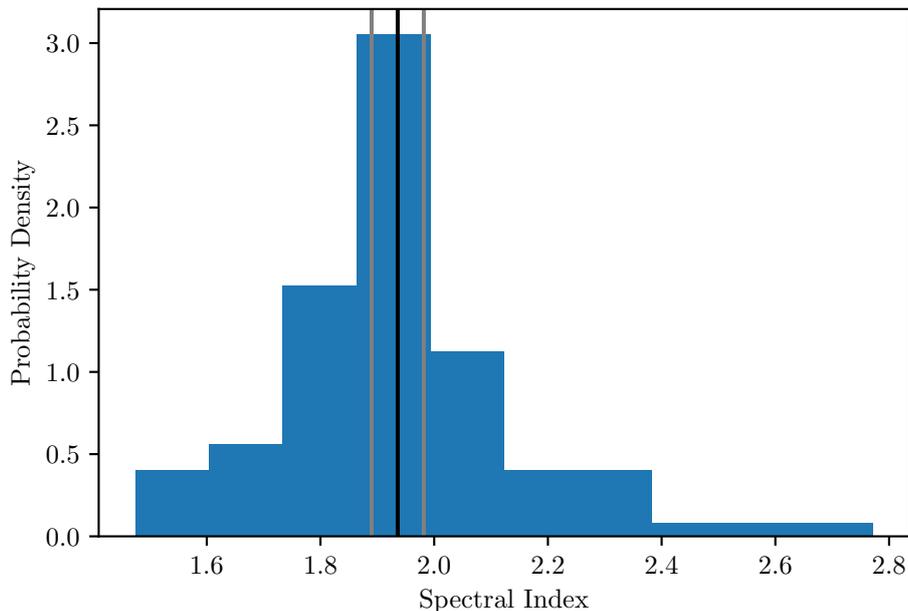


Figure 5.9: Histogram of spectral index for the monthly lightcurve of 3FGL J1309.3+4304 modelled by a power-law. The black line represents the optimised index in the integrated 8-year dataset, and the grey is the associated 1σ uncertainty.

ing the variation in spectral index of a single AGN (3FGL J1309.3+4304) in Figure 5.9, we find the expected wide distribution. Most notably, the distribution of power-law index values is far wider than would be expected from the uncertainty in the integrated 8-year spectrum. This illustrates the fundamental issue with using integrated spectra for investigation of spectral features in these objects. Integrated spectra can easily contain features from these brief periods of particularly hard or soft emission, possibly leading to mischaracterisation. The analysis of IC 310 in Chapter 4 is one such example.

Having shown the existence of spectral variability in the AGN population, and seen how wide the distribution in the case of power-law spectra can

be for an individual AGN, we now wish to characterise this behaviour for a population. In Figure 5.10, we plot a histogram of the normalised (relative to integrated values) variation in spectral parameters to form a picture of the duty cycle in spectral variability. For this simpler case, we look to characterise only those objects fitted with power-law spectra, and separate this into the two cases of FSRQs and BLL classified objects. The duty cycles shown are remarkably similar for both BLL and FSRQs objects, although FSRQs have a much more significant tail towards softer normalised indices. Both sets of objects also show very wide distributions, with BLL and FSRQs having a skewness of 0.993 and 1.048 (both positively skewed towards soft indices) and excess kurtosis of 3.995 and 4.166 respectively (both are heavy-tailed distributions compared to a Gaussian). While these values are similar for both classes, FSRQs show a slightly longer, heavier tail towards softer spectra.

5.7.1 Harder When Brighter?

As we have discussed previously, the “harder when brighter” behaviour of AGN is seen in several wavelength ranges (Emmanoulopoulos et al., 2012), and so this spectral variability in gamma-rays might be expected. For the avoidance of doubt, “harder when brighter” refers here to the tendency for AGN spectra to harden (decreased spectral index) when the integrated flux increases. Our dataset can provide an expanded view of this phenomenon for the objects in question, although this is limited to those objects modelled with a power-law to make a direct flux-index comparison, rather than a far more complex view that must be taken when considering curved spectra which would require much more data to have similar statistical power. This does sadly reduce the number of AGN available for this analysis to 41 BLL

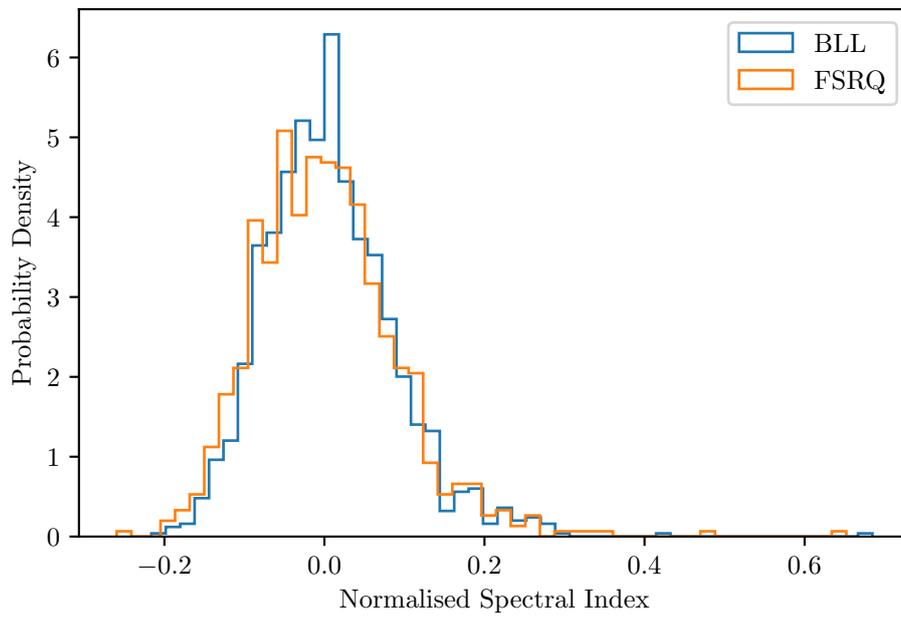


Figure 5.10: As with Figure 5.9 except spectral indices are normalised to the 8-year integrated value, and summed for all BLL and FSRQ objects.

objects and 52 FSRQ objects.

We approach this by applying a linear fit to the spectral index and the log of the integrated flux. As our data have uncertainties in both flux and index and substantial covariance, we use the `LinMix` algorithm (Kelly, 2007). This algorithm uses a hierarchical Bayesian model in which the intrinsic distribution of the parameters is that of a mixture of Gaussians. Applied in our case, we plot in Figure 5.11 the distribution of slope parameters (negative slopes represent harder-when-brighter behaviour). The BLL objects have a wider distribution than that of FSRQs, but both populations generally follow the harder-when-brighter paradigm. The uncertainty in the slopes of those BLL objects showing softer-when-brighter behaviour is such that harder-when-brighter behaviour is not excluded at the 1σ level. A representative example of the uncertainties on the data in each lightcurve and the fit parameters obtained is presented in Figure 5.12.

5.7.2 Discussion

In this Chapter I performed a wide-ranging analysis of the several hundred brightest gamma-ray AGN based upon seeds derived from the *Fermi*-LAT point source catalogues. I compared my 8-year integrated analysis with that performed by the *Fermi*-LAT collaboration (using slightly different background models and data) and found strong agreement, confirming that our spectral analysis produces reliable and reproducible results. This analysis was then extended beyond that performed by the *Fermi* collaboration to calculate monthly lightcurves, and within each lightcurve bin allowed the spectral parameters to vary. I have defined new statistics that allow for measurements of not just variability in flux, but also spectral variability and separately for combinations of spectral and flux variability.

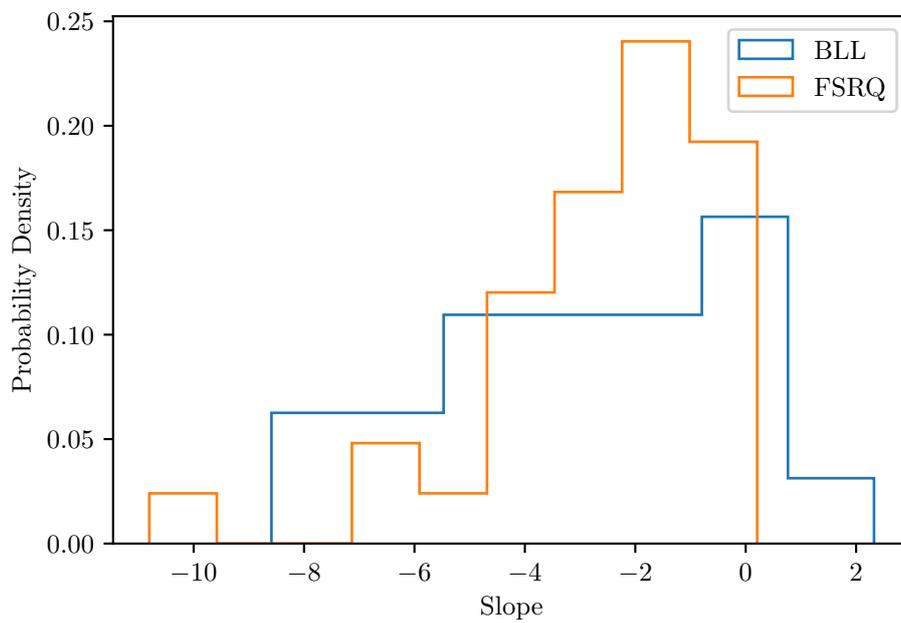


Figure 5.11: Histogram of linear fit slopes to BLL flux-index relationships. A negative slope represents a hardening with increased flux.

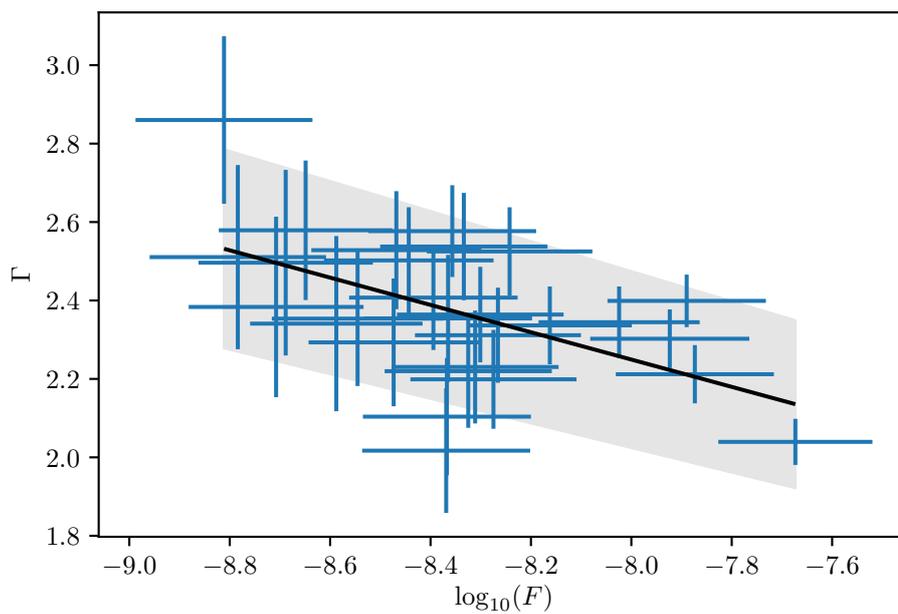


Figure 5.12: An example of one BLL object in the $\log_{10}(F)$ – Γ plane. The uncertainty region on the fit is presented as the 32nd and 68th percentiles of the posterior distribution of the parameters obtained using `LinMix`.

Applying this to our dataset, spectral variability is shown to be far rarer than flux variability. Roughly half of the AGN in our sample are spectrally variable, but over 90% are flux variable at the 5σ level. With this dataset, we also look to characterise the variability of gamma-ray AGN, finding significant differences in the PSD of flux variability between FSRQ and BLL objects, and also between low, intermediate, and high synchrotron-peaked BLLs. This transition represents increased long-term (baseline) variability in FSRQs than BLLs, which can plausibly result from the more active accretion behaviour of FSRQ objects leading to a much greater occurrence of low frequency variability from accretion compared to the gas-starved nuclei of a typical BLL. As the synchrotron peak frequency increased, the same transition from high baseline power to more flare-dominated behaviour was found. This progression similarly follows the transition from high-luminosity to low-luminosity galaxies and the transition between external photon field Compton scattering to pure synchrotron self-Compton emission from LBL to HBL objects (Fossati et al., 1998). None of the classes studied showed evidence of a break in the power density spectrum on the frequency range probed in this case (a Nyquist frequency of 0.5 month^{-1}).

A critical caveat to the interpretation of the average PSDs presented here is that the process involves averaging temporal properties across multiple sources. This approach carries a similar conceptual risk to that highlighted earlier in the thesis — that averaging can obscure meaningful, time-dependent physical processes. In this case, averaging PSDs over a population may suppress the variability signatures of individual sources, especially if a small number of bright AGN dominate the signal. As such, these averaged PSDs should be interpreted cautiously, and not as representative of any single object’s temporal behaviour. A full treatment of this issue would require per-source PSD modelling and weighted statistical combina-

tion, which lies beyond the scope of this chapter.

As expected, it was also found also find that spectral variability grows as more data is available. This is not directly related to the simple scaling of measurement uncertainty with statistics, but also an effect of the duty cycle of spectral variability. The growth in significance of $TS_{\text{spec var}}$ over time in AGN as a population hints that spectral variability, as with flux variability, is inherent to the physical processes at play in accelerating and cooling the charged particle populations in these jets.

I also looked to characterise the spectral variability of these objects, finding that FSRQs as a population tend to have a wider variation in spectral parameters than that of BLL objects and are slightly more heavy-tailed, meaning that FSRQs are slightly more likely to be found in softer states (relative to their integrated state) more often than BLL objects within our sample. As the spectral index in our energy range is strongly related to the peak frequency of the SED's second hump, we might expect BLLs to have a wider distribution as the peak frequency is doubly-dependent on the energy of the electron population in the jet, as opposed to the case of external Compton emission which comprises a relatively larger fraction of the FSRQ emission. However, this could also relate to a selection effect as only power-law modelled AGN were considered, and brighter AGN will generally be resolved into curved spectra.

Gamma-ray FSRQs have previously been found to exhibit strong harder-when-brighter characteristics, whereas BLL objects have traditionally had a far wider spread that has been both harder-when-brighter and softer-when-brighter in roughly equal measure (Lott & Larsson, 2013). While we do find that the flux-index relationship of BLLs is more widely spread than in FSRQs, a strong bias towards harder spectral index with greater flux is seen, unlike this previous work. It should be noted that Lott &

Larsson (2013) analysed a factor of 4 fewer BLL objects than this work, and that greater statistics and exposure are the likely cause of this discrepancy. Nevertheless, this work does indeed find several BLL objects with softer-when-brighter behaviour. The reason for this discrepancy between BLL and FSRQ objects can be explained in terms of the processes responsible for their gamma-ray emission. The environment of an FSRQ jet results in a far greater component of external Compton scattering from the BLR, dusty torus, accretion disc and X-ray corona, which results in a population that is far more homogeneous than that of BLL objects.

5.8 Conclusions

An analysis of the brightest 291 AGN above $|10^\circ|$ in galactic latitude, calculating SEDs and monthly binned lightcurves has been performed. A statistical method of measuring spectral variability using Wilks's theorem has been developed and applied to these data. The analysis has shown:

- Power density spectra of FSRQs and BLLs have significantly different power-law spectra, with FSRQs showing significantly more long-term variability than BLL objects.
- Similarly, a significant difference between LSP, ISP, and HSP BLLs is found, with LSPs showing more power in long-term variability.
- Spectral variability is not significantly found (at $> 5\sigma$) in as large a fraction as flux variability, but this fraction is growing with time. Over 50% of objects are spectrally variable at the 5σ level.
- FSRQs and BLL objects have generally similar characteristics in terms of spectral variability. Although we find FSRQs tend to inhabit softer

spectral states relative to their integrated spectra more often than BLLs, this difference is marginal.

- Harder-when-brighter behaviour is found in both FSRQs and BLLs when flux is taken into account as well as spectral index, but BLLs have far more scatter than FSRQs, including some BLLs showing (non-significant) softer-when-brighter behaviour. No FSRQs were found to exhibit softer-when-brighter behaviour.

With further integration time, many of the shortcomings of the current analysis can be overcome. For example, greater resultant flux sensitivity will provide a larger selection of AGN to sample, but also remove uncertainty due to the inherently stochastic nature of variability studies. Given that such a high fraction of the AGN show strong spectral variability, in the next Chapter we look to see how this can affect astrophysical measurements.

Chapter 6

Spectral Variability and Measurement of the Extragalactic Background Light

6.1 Introduction

The extragalactic background light (EBL) is the cosmic diffuse radiation field emitted by stars (Inoue et al., 2014) and active galaxies (Shankar et al., 2016) throughout the history of the Universe. The EBL thus varies strongly as a function of cosmic time (or more conveniently, z) and wavelength. This radiation field represents the second most energetic diffuse photon field after the cosmic microwave background, and so measurements of this field are important contributions to the energy balance of the Universe. The EBL's SED shows a characteristic double peak, with one in the optical to near-UV as a result of summed stellar emission, and a second lower-frequency peak due to absorption and re-emission by interstellar dust.

Direct measurements of the EBL are very difficult and plagued by many irreducible foreground components (e.g. zodiacal light and the Milky Way's own emission), which lead to large biases in the results (Mattila et al., 2017;

Bernstein, 2007). Furthermore, this direct measurement also precludes the possibility of measuring the *evolution* of the EBL intensity as a function of cosmic history as it can only provide today’s snapshot. Far easier to calculate, however, is a lower limit on the EBL SED from photon counting in deep galaxy surveys (Madau & Pozzetti, 2000; Fazio et al., 2004; Keenan et al., 2010; Tsumura et al., 2013; Driver et al., 2016), although this isn’t without its own shortcomings (Bernstein et al., 2002; Somerville et al., 2004).

The EBL is particularly important to the study of gamma-ray astronomy due to the energy dependence of the $\gamma\gamma$ cross section above ~ 100 GeV. This interaction, shown in Figure 6.1, generally results in the creation of an electron-positron pair, although other products are possible at higher interaction energies and smaller cross sections. The existence of this interaction manifests as an effective “optical depth” that varies with energy. The threshold for this interaction exists at $E_\gamma \times E_{EBL} = 2(m_e c^2)^2$ and provides a critical energy for gamma-ray interaction of $E_{crit}(z) \approx 170(1+z)^{-2.38}$ GeV (Franceschini et al., 2008). This means that below 100 GeV, even the most distant discovered gamma-ray AGN appear almost entirely unattenuated, but in the TeV range the exponential attenuation results in precipitous cut-offs in an AGN’s SED.

By measuring the degree of attenuation as a function of energy, and assuming we know the redshift of a sufficiently bright AGN (not always a given in the case of BLLs) then the intensity spectrum of the EBL can be calculated (Biteau & Williams, 2015). More often however, one of many models of the EBL intensity is scaled by a normalisation parameter and optimised for a given dataset (Fermi-LAT Collaboration et al., 2018b; Acciari et al., 2019). While these analyses are often considered to be measurements in their own right, it is more accurate to say that they can place upper limits on the EBL intensity, as again the provenance of each source of radiation

is uncertain. In this sense such analyses can be seen as the counterpart of the photon counting analyses to exclude upper and lower intensity ranges, if not an exact measurement.

These gamma-ray analyses of course come with their own set of biases and systematic uncertainties, and in particular we aim to investigate the systematics introduced due to spectral variability in this Chapter. Previous analyses by e.g. Domínguez & Ajello (2015) and the Fermi-LAT Collaboration et al. (2018b) simply use the integrated SEDs calculated over the entire *Fermi*-LAT dataset. As we have discovered throughout this thesis, most relevantly in Chapter 4, spectral variability can easily distort an integrated spectrum into a shape that could possibly hide (or create) an EBL-like signal.

We explore the EBL models which shall be used in the analysis in Section 6.2, before using the dataset produced in Chapter 5 to analyse the chosen models, and investigate the effect of accounting for spectral variability upon a measurement of the EBL intensity from LAT data.

6.2 EBL Models

In general, there are four main types of EBL models, each of which have their own assumptions and predictions:

- Forward evolution models: These models begin with a Λ CDM hierarchical structure formation simulation to form the initial cosmological conditions, which is then evolved using a semi-analytical model (SAM). We use the model created by Gilmore et al. (2012) for our analysis, which took WMAP5 cosmological conditions seeded with a Chabrier initial mass function (Chabrier, 2003) before applying their SAM. The model is too complex to summarise briefly here, but

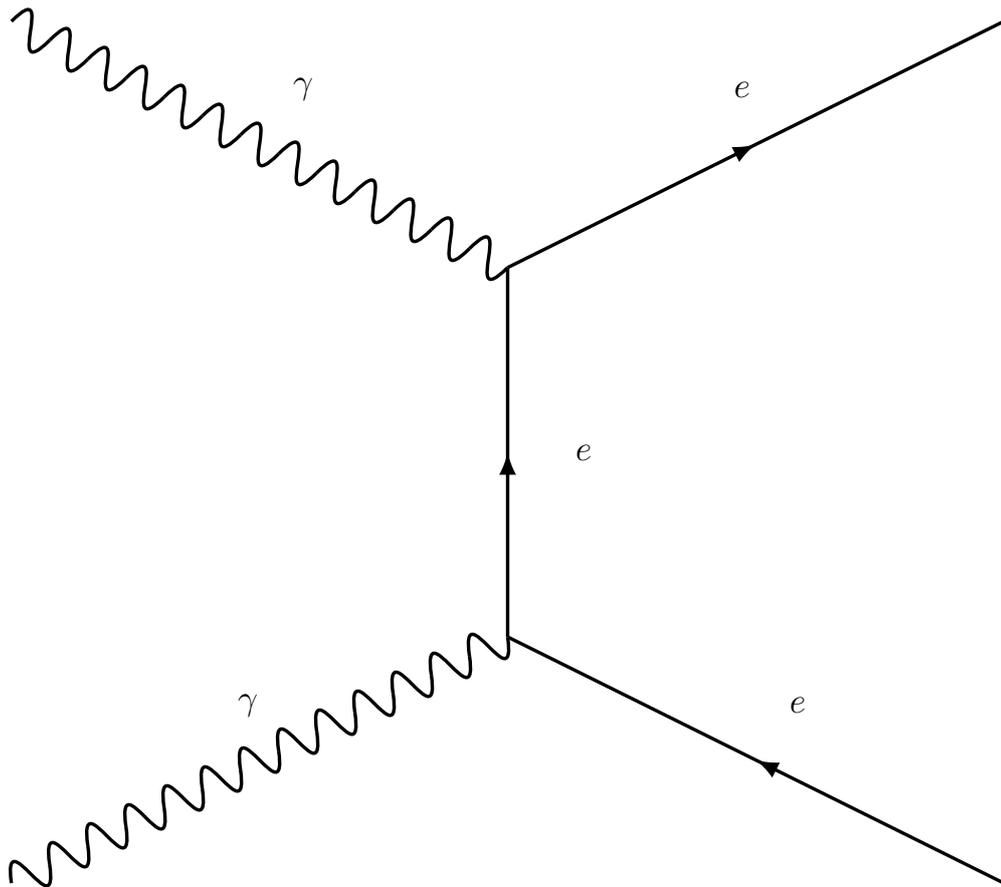


Figure 6.1: Feynman diagram for pair production to electrons.

CHAPTER 6. SPECTRAL VARIABILITY AND MEASUREMENT OF THE EXTRAGALACTIC BACKGROUND LIGHT

Somerville et al. (2012) provides a deep discussion of the specific elements of the SAM used in this case. Inoue et al. (2014) uses an alternative SAM method specifically taking into account Pop-III stars that are particularly relevant to EBL measurements at $z > 6$.

- Backward evolution models: Models such as Franceschini et al. (2008) take empirical galaxy data and evolve these measurements back in time with estimated redshift scaling depending on wavelength (according to the different associated processes) to create a prediction of the EBL at increasing redshift.
- Evolution of galaxy populations inferred from various star formation rate and initial mass functions: Finke et al. (2010) uses this method with 5 different SFR and IMF combinations, predicting the associated starlight emission, including absorption and reradiation by dust to create a combined luminosity density resultant from these contributions and then extrapolate this to calculate an EBL intensity as a function of redshift. Kneiske & Dole (2010) uses this same method but applies a simple stellar population (SSP) model more directly (Finke et al. (2010) does reproduce these SSPs despite using a different method).
- Evolution of galaxy populations from measured galaxy SEDs: Domínguez & Ajello (2015) takes much the same approach as Finke et al. (2010) and Kneiske & Dole (2010), except instead of approximating galaxy evolution through SFR density, instead uses galaxy SEDs up to redshift 4 to provide an empirical measurement of galaxy evolution, before calculating the commensurate EBL contribution.

Of course, none of these models is perfect, but in general they are in good agreement. The EBL models predicted by the work cited above are plotted

in Figure 6.2; these models (namely Gilmore et al., 2012; Inoue et al., 2014; Franceschini et al., 2008; Finke et al., 2010; Kneiske & Dole, 2010; Domínguez & Ajello, 2015) will be used in the analysis.

6.3 Data Selection

6.3.1 *Fermi*-LAT Analysis

The analysis in the section is largely identical to that in Chapter 5 but for a revised method specifically relating to the likelihood analysis once an SED has been fit to each AGN. Some changes in source selection are however different due to the availability (or lack thereof) of redshifts for BLL objects. For the sake of clarity an exhaustive description of the data analysis pipeline for this Chapter is given below

We use 8 years of Fermi Large Area Telescope (*Fermi*-LAT) data taken between 239557417 and 492328870 MET. For the majority of this observation time, the Fermi observatory was in all-sky survey mode. In all-sky survey mode, 30 minutes of livetime are obtained for the whole sky over a 3-hour period. All base analysis is performed using the *Fermi*-LAT science tools v9r33p0.

As in Chapter 5, we consider only data in the energy range 100 MeV–500 GeV in order to make use of the Fermi collaboration-provided Galactic emission model `gll_iem_v06.fit`, and to avoid issues with likelihood fitting and large systematic error at low energies. Similarly, we applied the filter `DATA_QUAL>0 && LAT_CONFIG==1` in order to ensure the telescope was in data-taking mode. We have chosen to use `SOURCE`-class photons and the `FRONT+BACK` event type and accordingly applied the `P8R2_SOURCE_V6` instrument response functions and used the isotropic emission template

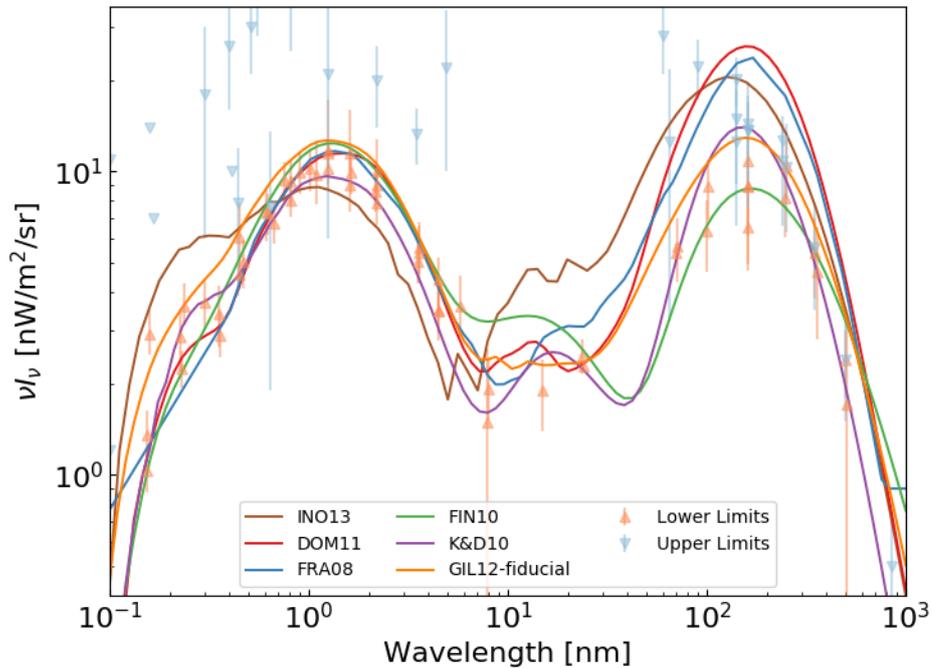


Figure 6.2: The EBL intensity predicted by various models (named in legend) at $z = 0$. We use the the following abbreviations for the published predictions for comparison: INO13 (Inoue et al., 2014), FIN10 (Finke et al., 2010), DOM11 (Domínguez & Ajello, 2015), K&D10(Kneiske & Dole, 2010), FRA08 (Franceschini et al., 2008), GIL12-fiducial (Gilmore et al., 2012). Upper limits from gamma-ray observations and lower-limits from galaxy count predictions are included.

`iso_P8R2_SOURCE_V6_v06.txt`. To avoid contamination from cosmic ray-induced gamma-ray emission from the Earth limb, we also apply a 90° zenith cut.

The *Fermi*-LAT third source catalogue (3FGL; Acero et al., 2015) is used to choose the sources used in the analysis. To obtain the largest possible selection of AGN at a variety of distances, all AGN in the 3FGL are selected and both those at low Galactic latitudes $|b| < 10^\circ$ and those with an average significance in the 3FGL of less than 20σ (derived from the test statistic) are discarded. At this point, no AGN are excluded based on their variability. Of the 301 sources selected, 259 of these had associated redshifts in the 3LAC (Ackermann et al., 2015a) and 8 objects showed lightcurve covariance due to their proximity, so were also removed from this sample (see Table 5.4 in the previous chapter for a summary of removed sources). Objects without associated redshifts were similarly discarded. The redshift distribution by source classification is shown in Figure 6.3.1. Due to the inherent difficulty of measuring distances to BLLs (their very definition as a class prohibits the high quality spectra required for measuring redshift), the fraction of FSRQ objects increases at high redshift. Furthermore, as the calculation for spectral variability is only defined according to Mattox et al. (1996) when the lightcurve contains no upper limits (individual bins with $< 1\sigma$ uncertainty), this sample is reduced greatly to only 100 sources.

For each source, we choose a region of interest (ROI) with a 7.1° radius centred on the source is chosen. This is then cut down to a $10^\circ \times 10^\circ$ square for a binned likelihood analysis using 10 energy bins per decade. The standard *Fermi*-LAT binned analysis chain follows, calculating livetime, exposure maps and source maps. The initial source map consists of the diffuse models discussed earlier and is otherwise seeded with 3FGL sources within a 20° radius as sources outside of the ROI can contribute significant

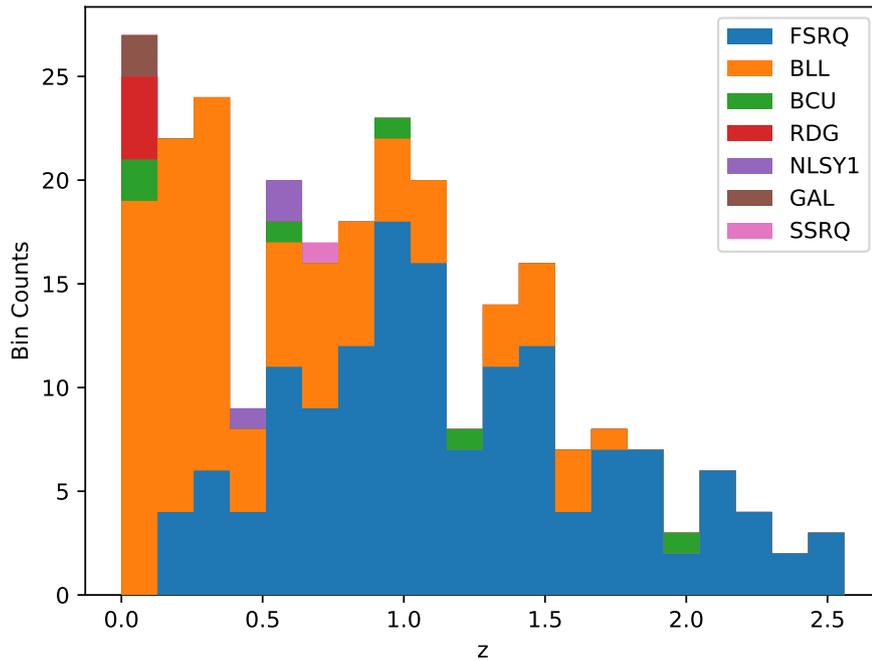


Figure 6.3: A stacked histogram showing the distribution of redshifts by 3FGL source classification in our sample. The overall height of a bar represents the number of sources in a redshift bin, and the colours are stacked to represent the breakdown within a bin by source type. In general, BL Lac objects dominate at lower redshifts, followed by FSRQs at higher redshifts.

flux.

An iterative fitting process using the tool *fermipy* (Wood et al., 2017) optimised the model by fitting the sources individually in order of significance to get close to the global minimum before freeing all sources for the final fit. The spectral shape of a source is only fitted if the test statistic (TS), defined as:

$$\text{TS} = 2(\log(\mathcal{L}_1) - \log(\mathcal{L}_0)), \quad (6.1)$$

is greater than 25. \mathcal{L}_0 and \mathcal{L}_1 refer to the likelihood of the null and alternative models, which in this case refer to the fit of the model without the source, and with the source included.

As all sources are expected to have significances well in excess of our imposed $\text{TS} > 25$ limit, their spectral energy distributions (SEDs) were optimised using the spectral model assigned to them in the 3FGL (depending on the degree of spectral curvature, either a power law or log parabola). This spectral model is specifically used for the global fit.

6.3.2 Redshift Data

To obtain redshift values, we initially take those obtained by the *Fermi* collaboration in the 3LAC (Ackermann et al., 2015a). As these values do not have associated uncertainties (at least in the catalogue itself), it was necessary to look to the NASA/IPAC Extragalactic Database and SIMBAD using `astroquery` for automation. A prior for the redshift is constructed before profiling out this parameter using the previous measurements and uncertainties assuming that the uncertainty is distributed normally. If no uncertainty is available, the uncertainty is assumed to be zero.

6.4 Analysis Methods

6.4.1 EBL Fitting

The EBL model is optimized using likelihood profiles extracted from the *Fermi* tools. This method means that the spectral fit does not feed back into the global fit of the ROI, which leads to a small systematic error. We can rest assured that this error is small, as the EBL attenuation only significantly affects a source's emission at energies greater than 10 GeV, even for the most distant gamma-ray sources (around $z=2$). The point spread function of the LAT detector is substantially smaller at energies greater than 1 GeV than below 1 GeV due to the absence of multiple scattering. Consequently, such photons are very unlikely to be associated with other point sources such that the global fit is substantially altered.

To extract likelihood profiles in each of the energy bins defined in Section 6.3, we assume that spectral curvature within a bin is minimal, and use a power law model with spectral index 2.0 to approximate the flux in each bin. We then scan the normalisation parameter of the power law solely in this bin and calculate the likelihood of the ROI model for each normalisation value. This relation between flux and likelihood is calculated until a $\Delta \log(\mathcal{L})$ of 2 (corresponding to 2σ) is found in each direction, or until 0 flux is reached, representing a bin in which $TS < 4$.

These likelihood profiles are interpolated using spline functions of order 2, which are extrapolated beyond the 2σ limits. This extrapolation has been tested by extending the profile to 10σ for the brightest sources and calculating the mean squared error at each point. This was found to introduce a $< 0.1\%$ systematic error, which is far smaller than that of the detector's effective area (around 10% or more throughout the energy range). According to Wilks's theorem (Wilks, 1938) this parameter must be free of constraints

to ensure the statistical accuracy of the likelihood ratio test (the asymptotic nature of the approximation relies upon finding the global minimum). Our extrapolation of likelihood values allows us to compare log likelihoods for even the most extreme cases of spectral parameter values. These profiles do not extend to negative flux values, as extreme values of spectral indices, spectral curvature parameters or EBL attenuation normalisations can never result in negative fluxes. Only negative overall flux normalisations can cause this effect, and given the statistical significance of these sources (all $> 20\sigma$), we can safely assume that they are not gamma-ray sinks.

Once interpolating functions for each source are calculated, a log parabola spectrum with EBL attenuation is then fitted,

$$\frac{dN}{dE} = N \left(\frac{E}{E_0} \right)^{-(\alpha + \log(\beta))} \times e^{-a\tau_{\gamma\gamma}(E,z)}, \quad (6.2)$$

where N represents the overall normalisation of the model, α and β represent the spectral index and curvature, E_0 the scale energy, a the EBL scaling factor, and $\tau_{\gamma\gamma}(E, z)$ the theoretically predicted opacity at a given z . The uncertainty on the flux in this model (assuming Gaussian uncertainties) is a complex function, and so the calculation is derived in Appendix B for the interested reader.

This model is a representation of the intrinsic spectrum and its curvature as parameterised by the log parabola, with an exponential decay term due to the effective gamma-ray opacity of the universe at a given energy and distance. The addition of the exponential decay and our scaling factor a represents a model nested with that of a log parabola, allowing us to represent the statistical significance of the addition of a by

$$\text{TS} = 2(\log(\mathcal{L}(a)) - \log(\mathcal{L}(a = 0))), \quad (6.3)$$

whose value is then described by a χ^2 distribution with one degree of freedom.

The fit is performed using the fitting package MINUIT via the python `iminuit` interface. Firstly we fit a log parabola (with parameters N , α , and β) to approximate the intrinsic spectrum, then introduce the attenuation and fit all four parameters simultaneously.

6.4.2 Systematic Fitting Uncertainty

As a slightly unusual approach of using an external fitting package to the *Fermi* science tools has been taken - although it should be noted that the *Fermi* tools themselves use MINUIT - we may introduce a systematic bias by doing so. This is unavoidable as the *Fermi* tools do not have the functionality to apply EBL attenuation to spectral models. The data that we are fitting is an SED constructed from likelihood profiles in flux as a function of energy that is extracted using the *Fermi* tools. This means that while we have profiled out the effect of nuisance parameters (e.g. the Galactic diffuse emission and the parameters of nearby 3FGL point sources), we effectively remove this portion of the “feedback loop” inherent to the fitting process. We are confident to do this as the changes to the AGN spectrum will be small (usually less than a 10% effect at energies > 10 GeV, and smaller than 1% below this energy), and dominantly at high energies where the PSF is small, meaning the covariance with other models is restricted to smaller spatial areas and those models in the immediate vicinity. In practice, no background sources will fall within even the 99% PSF containment region of any of our sample AGN at this energy.

To perform a rigorous test of this fitting routine, we simulated 12 AGN detectable by *Fermi*-LAT, of which 6 had power-law spectra attenuated by the EBL (using the Domínguez & Ajello (2015) model), and 6 had log-parabola spectra similarly attenuated by the EBL to form a “data challenge”

for our method. Applying the MINUIT fitting routine once energy-binned likelihood profiles in flux were extracted, good agreement was found with the simulated dataset. The simulated parameters and reconstructed values with uncertainties are plotted in Figures 6.4 and 6.5. In all cases, the optimised EBL normalisation parameter is within the 2σ uncertainty region of the correct value, and all other parameters show good correlation to the “true” simulated signal. Larger uncertainties in the case of the log parabola result from the unavoidable covariance between the β curvature parameter in a log parabola and the exponential cutoff due to the EBL and is entirely expected in our application to the data.

As a further check to the methodology, I also re-fit the entire sample of log parabola-modelled AGN in both the *Fermi* tools and my own fitting method (with the EBL normalisation fixed to zero) and compare the resultant parameter values as plotted in Figure 6.6. The spectral index parameter (α) and normalisation (N) both have distributions that fall within 1% of the value obtained using the *Fermi* tools, whereas β has a small bias ($\sim 2.5\%$) towards smaller curvature. As these uncertainties are small compared to the magnitude of the effect we are measuring (the EBL has been detected far beyond 5σ in previous studies using the LAT), this is evidence that my fitting routines are suitable for measuring the EBL.

6.4.3 Joint Likelihood Fit

A binned *Fermi*-LAT likelihood analysis approaches source fitting by applying the maximum likelihood method by binning the data in both space and energy, where the summed counts in each bin follow a Poisson distribution around the predicted count rate according to the theoretical model (spectrally parameterised by $\boldsymbol{\mu}$) convolved with the IRFs. The likelihood

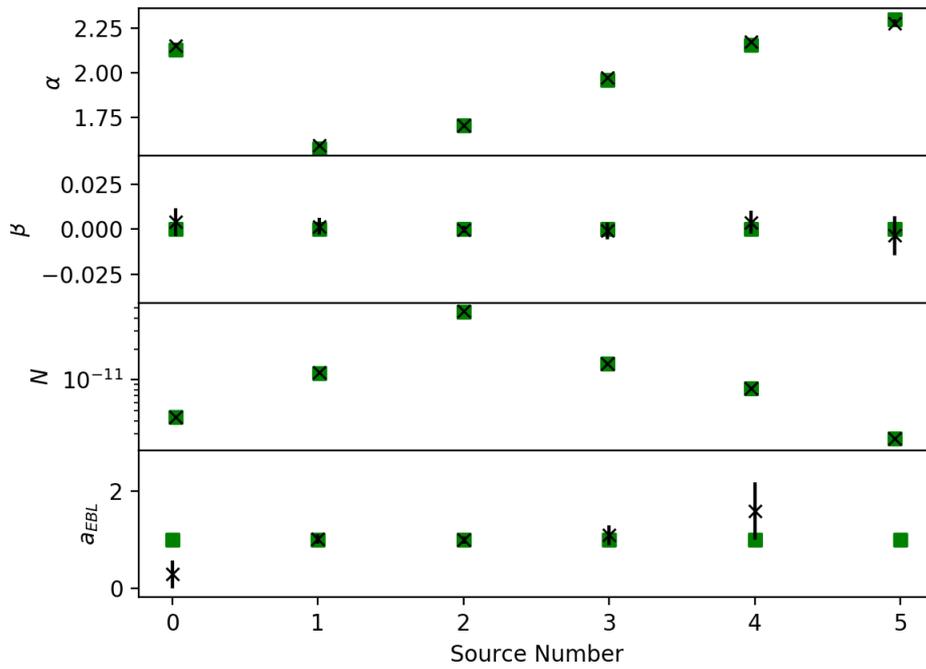


Figure 6.4: Results of the data challenge set in Section 6.4.2. Simulated AGN power-law spectra parameters are denoted by green squares which effectively represents the “truth” to which Gaussian noise was added. Fitted parameters using the likelihood fitting routine are denoted by black crosses with associated uncertainties.

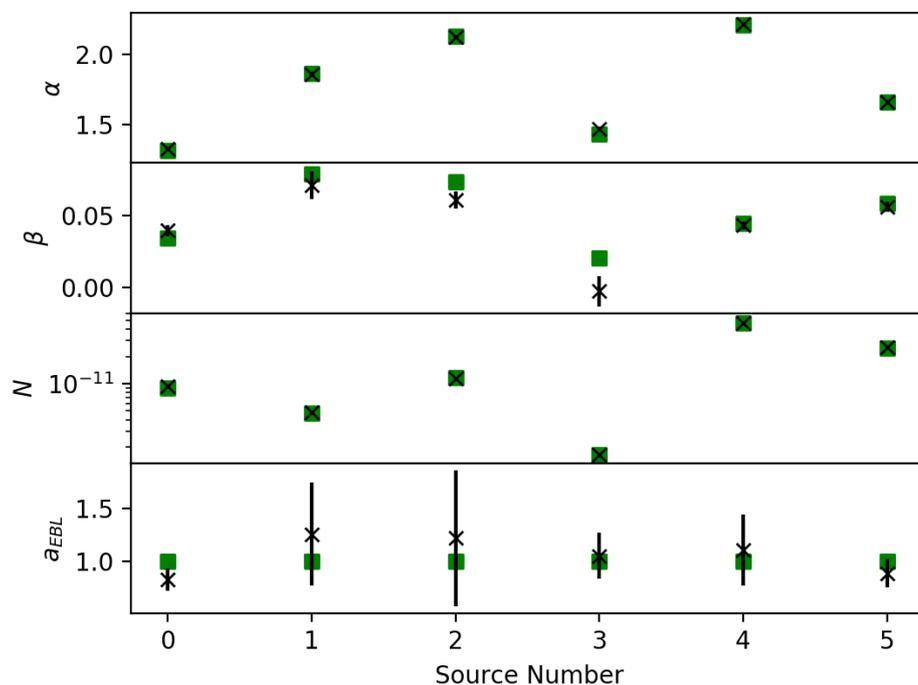


Figure 6.5: Parameters measured as in Figure 6.4 except for sources simulated as log parabolae.

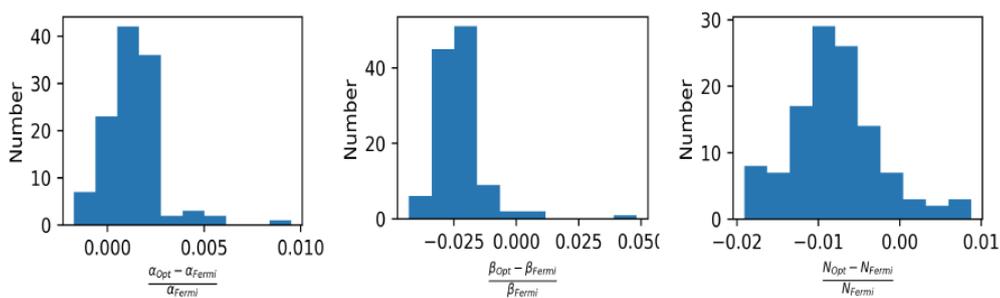


Figure 6.6: Relative uncertainty between fitting SEDs with my MINUIT-based routine and that returned by the *Fermi* science tools.

function is described by

$$\mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\theta}|D) = \prod_{jk} \frac{\lambda_{jk}^{n_{jk}} e^{-\lambda_{jk}}}{n_{jk}!}, \quad (6.4)$$

where λ is the expected number of counts for particular values of $\boldsymbol{\mu}$ and the set of nuisance parameters $\boldsymbol{\theta}$. In this case $\boldsymbol{\theta}$ are the modelled parameters for all other point sources and diffuse sources in the ROI. The index k represents the product over spatial bins, and the index j represents the product over energy bins.

In Section 6.4.1, I described how the normalisation of a source in each energy bin is profiled, resulting in a set of $\{\boldsymbol{\mu}_j\}$ for each spatial bin k . A global energy fit as is usually assumed in analyses would assume that $\{\boldsymbol{\mu}_j\} \approx \boldsymbol{\mu}$, however we only apply this approximation after extracting profiled likelihood values for a range of $\boldsymbol{\mu}_j$.

As we will be using a combined analysis across all of our sample, we must also take into account statistical uncertainties on the redshifts used in Equation 6.2. As with the analyses of dark matter signals in dwarf spheroidal galaxies by Ackermann et al. (2015b), we add the redshift, z , of each source as a nuisance parameter. The form of the likelihood function for the combined fit is,

$$\mathcal{L}(\boldsymbol{\mu}, \{\boldsymbol{\alpha}_i\}|D) = \prod_i \mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\theta}_i|D_i) \times \frac{1}{\sigma_i \sqrt{2\pi}} e^{-(z_i - \bar{z}_i)^2 / (2\sigma_i^2)}, \quad (6.5)$$

where the index i represents each source considered in our sample, \bar{z}_i and σ_i are the redshift and its associated statistical error, $\{\boldsymbol{\alpha}_i\}$ is the set of nuisance parameters for each ROI including redshift uncertainty, $\boldsymbol{\theta}$ are the background source parameters, and $\boldsymbol{\mu}$ are the ROI-independent quantities whose value we wish to calculate. In this case, $\boldsymbol{\mu}$ is simply the normalisation a of the EBL in Equation 6.2.

6.5 Method Validation

6.5.1 Sources of Systematic Uncertainty

The main systematic errors in a spectral analysis of this complexity originate with the applicability of a simple log parabola as a description of the intrinsic spectrum of an AGN. There are two main ways in which this assumption can break down. Firstly, the intrinsic spectrum may simply not be shaped as a log parabola. If the intrinsic spectrum is instead a power law, thankfully the log parabola model is nested with a power law, and so would simply result in zero curvature. Unfortunately, if the intrinsic spectra exhibited an exponential cutoff, this would appear in our data as a false EBL detection. All but one AGN in the 3FGL catalogue are modelled by power laws and log parabolae, the exception being 3C 454.3, which is modelled using a power-law with a super exponential cutoff. As this contributes only one data point amongst 250 other fits, we assume that this will have a minimal effect on the final result.

Secondly, the dataset that we use is integrated over 8 years; if any spectral variability were present in a source, integrating over the variability would distort the shape of its spectrum. While it is possible in the case of particularly bright AGN (generally the 8-year $TS > 10,000$) to fit the spectrum in each monthly bin of the lightcurve, for the vast majority of sources the spectral parameter fits provide insufficient data for a meaningful answer once systematic uncertainties are taken into account. Instead, I take a different approach as described in Section 6.5.2.

6.5.2 Variability Analysis

To solve the second issue outlined in Section 6.5.1, we use the test statistic $TS_{\text{spec var}}$ as defined in Section 5.5 to make a cut on sources whose spectral index is varying in time at a level greater than 5σ . While the most recent studies by the *Fermi*-LAT collaboration (Fermi-LAT Collaboration et al., 2018a) instead made this cut based on TS_{var} , this method explicitly removes sources with spectral variation yet allows for spectrally stable flux variability. This requires a cut on $TS_{\text{spec var}}$ at 180.43 (5σ based on a χ^2 with 97 degrees of freedom) for a power law model, or 305.35 (5σ based on a χ^2 with 194 degrees of freedom) for those sources modelled with a log parabola, or in the case of 3C 454.3, a power law with super exponential cutoff. While this 5σ cut may seem excessive, we are in fact considering 301 sources. Assuming my variability tests are independent hypothesis tests on the existence of variability, we in fact find that the 5σ level is reduced to a p-value of 0.017% given the trials factor.

To confirm the accuracy of my method in calculating TS_{var} , the first 48 months of the monthly lightcurves are compared with the values calculated in the 3FGL in Figure 6.7. These values are impossible to directly compare as the 3FGL used `Pass 7 Reprocessed` data, a smaller energy range (100 MeV to 300 GeV), and an unbinned likelihood analysis method (as opposed to a binned analysis in this work). Despite these differences in analysis, the relation between the TS_{var} calculated in this work and the 3FGL have a strong linear relationship in log space.

These values could be biased by brighter sources scaling similarly by the metric calculated in this work's TS_{var} as in the 3FGL's TS_{var} . In order to check that we are matching our dataset correctly, we can compare the results of the linear regression as a function of the number of binned months

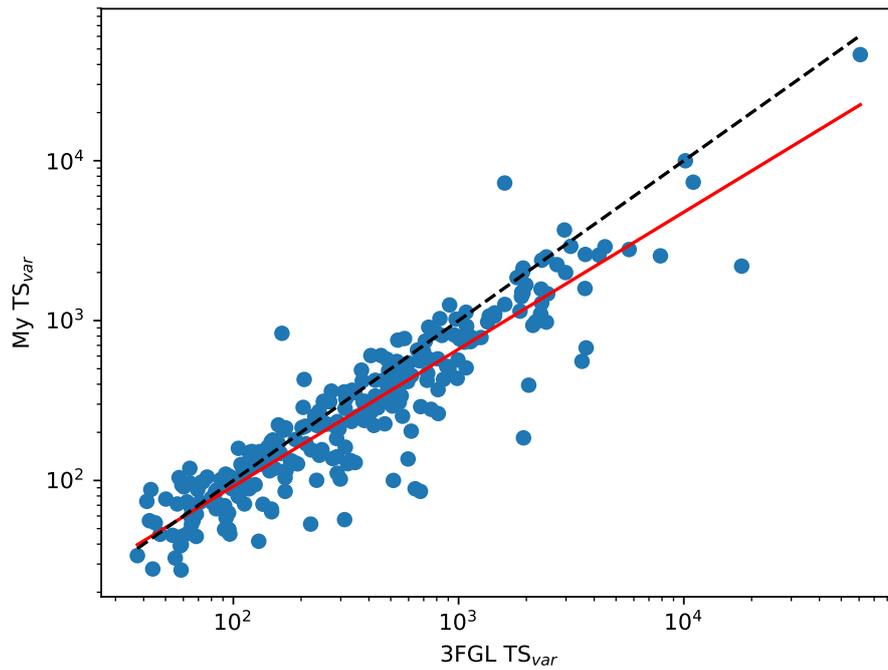


Figure 6.7: The TS_{var} calculated in this work plotted against the TS_{var} calculated for a source in the 3FGL. The black dotted line represents a one-to-one relationship between the two values, and the red line represents a linear regression fit ($r^2=0.842$) to the log of the values. The spread is largely due to different data selection processes.

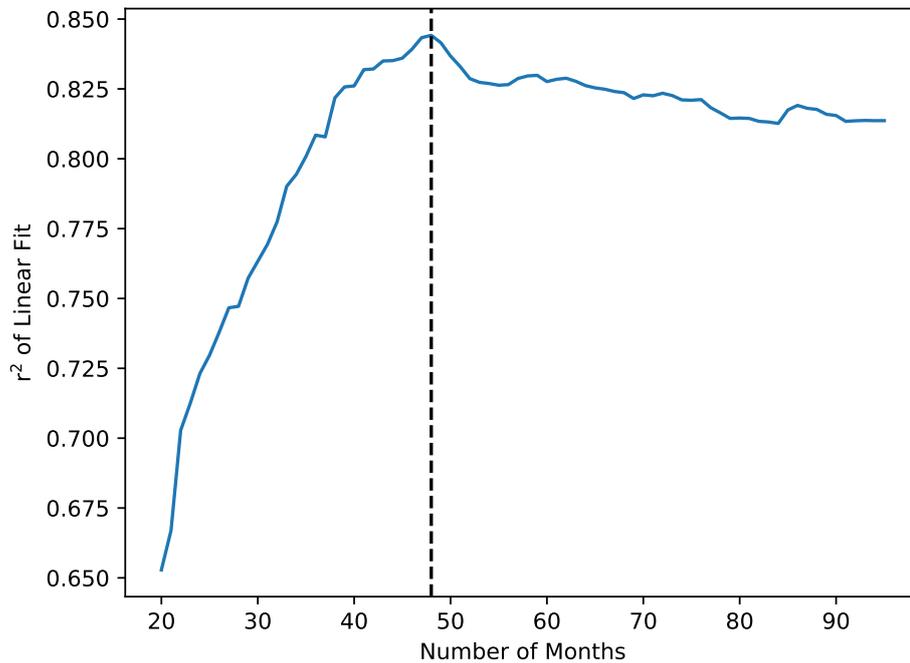


Figure 6.8: The r^2 value of the linear fit in log space as a function of the number of months of data used in the calculation of the TS_{var} . This is convincing evidence that the correlations are not simply due to similar scaling, but inherent correlation between the sources in the dataset.

of data I used to calculate the TS_{var} , as in Figure 6.8. This shows that the smallest spread (measured by the r^2 value) between the two datasets occurs at 48 months of integrated data, the time corresponding to the 3FGL's data.

6.6 Results

Applying the methods described in Sections 6.4.1 and 6.4.3 to the data to obtain best-fit normalisations for each of the six EBL models, we obtain the results presented in Figures 6.9 and 6.10. Figure 6.9 shows the full dataset,

irrespective of the detected spectral variability of the sources. Figure 6.10 shows this measurement when only non-variable sources contribute.

Investigating the differences between previous studies and this work, we must compare the relevant samples. Ackermann et al. (2012) presents the most direct comparison, as while Biteau & Williams (2015) did use *Fermi*-LAT data, this was used when combined with IACT data. As we are interested in the performance of our measurement with regards to variability, in this case we shall compare solely the results reported by Ackermann et al. (2012).

As is apparent from Figure 6.9, that when the EBL normalisation is calculated in individual redshift bins (minimum of 10 sources per bin), there is a huge amount of scatter compared to Figure 6.10. Of course, deviation from the expected normalisation of 1.0 is not in itself indicative of an incorrect result. The models themselves are purely theoretical, and deviation as a function of redshift may simply imply that the EBL intensity does not evolve as predicted with redshift. Similarly, while we can compare with Ackermann et al. (2012)’s values and that of the IACT measurements, judging our results in this context simply causes a bias towards “herding” of the results. Given this, the most reliable indicator of our method’s sensitivity compared to previous methods is to examine the relative size of the uncertainty regions when all analysis is held constant.

With the complete unfiltered sample of 251 AGN that had valid $TS_{\text{spec var}}$ values using the Domínguez & Ajello (2015) model as a baseline for comparison, we find a value of $a = 1.33 \pm 0.08$. Taking the sample of 160 AGN without 5σ -detected spectral variability, we instead find $a = 1.32 \pm 0.09$. Further reducing this sample to those 61 sources without any spectral variability *or* flux variability, we find $a = 0.98 \pm 0.11$. This is the value represented in the grey bar in Figures 6.9 and 6.10. All uncertainties are calculated using

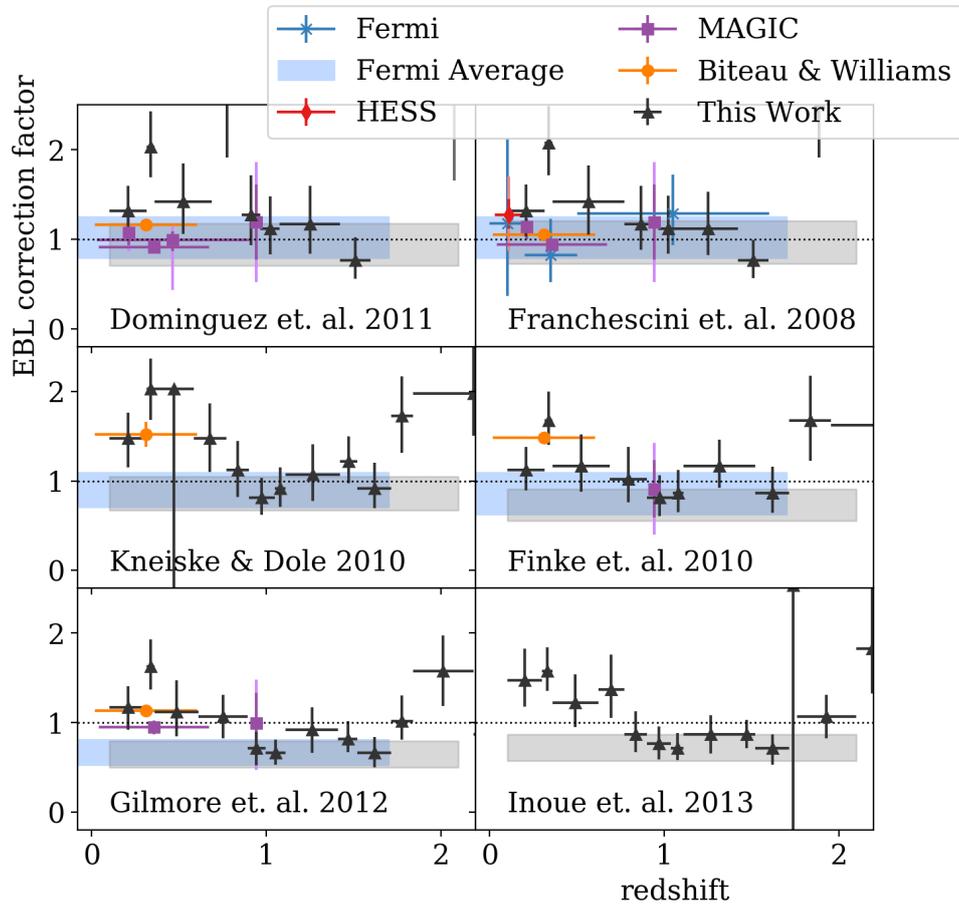


Figure 6.9: EBL normalisation as a function of redshift (where available) for the 6 models considered in this work. As no other work considered the Inoue et al. (2014) model, only our results are presented. Independent measurements are sourced from HESS (H.E.S.S. Collaboration et al., 2013), MAGIC (Ahnen et al., 2016b,a; Mazin et al., 2017), *Fermi*-LAT (Ackermann et al., 2012) and combined (Biteau & Williams, 2015). The grey bar represents the overall measurement excluding spectrally variable sources.

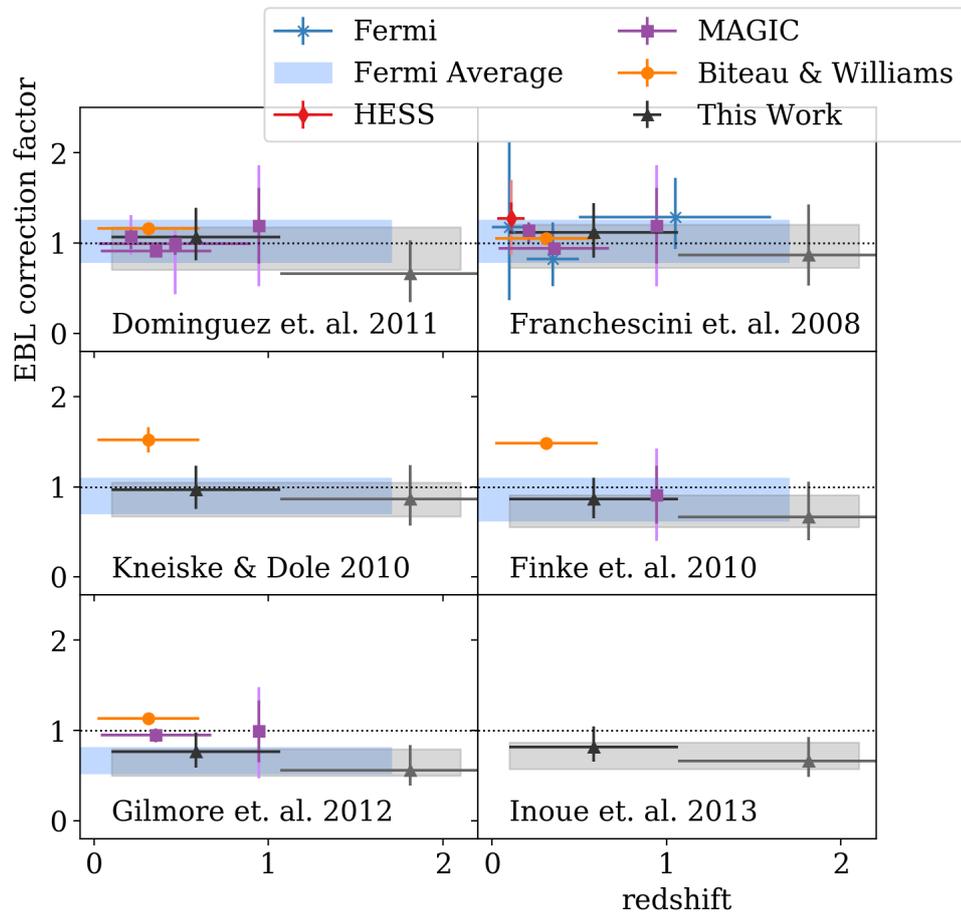


Figure 6.10: EBL normalisation for our selection of models as described in Figure 6.9, except with sources restricted to only those sources without significant variability, resulting in fewer available redshift bins.

the method suggested by Barlow (2003) (Eqn. 21) for symmetrising uncertainties with likelihood profiles, which was shown to perform appropriately with Poissonian likelihood profiles (which is approximately the case here). Impressively, despite reducing our sample by a factor of four, the statistical uncertainty has only marginally increased by 30%. Of course, this scaling would not occur if the effect was purely statistical, meaning that the intrinsic scatter due to variability causes a significant systematic uncertainty.

Finally, when repeating the calculation without the redshift uncertainties included, negligible change is found in the uncertainty or optimised value. Given fewer than half of the AGN considered had associated uncertainties from our simple search of NED and SIMBAD, this is perhaps unsurprising. Indeed, it seems that even in the case where the redshift uncertainty is significant, this uncertainty is far smaller than that inherent to the *Fermi* statistical uncertainty when fitting the SEDs, as the relevant changes in critical energy are almost always within a single energy bin. We must conclude that this is certainly not a limiting systematic effect.

6.7 Discussion

These results provide evidence that despite reducing the dataset by more than 70%, the statistical uncertainty in the measurement only increased by $\sim 30\%$. Obviously there is intrinsic scatter caused by variability that is not being taken into account in a purely statistical estimation of the uncertainty. One main outstanding question is whether the scatter caused by AGN variability in this case causes simply a larger than stated uncertainty in EBL intensity, or in fact a bias towards higher or lower intensities. In our simple estimation, when removing variability completely (both spectral and flux variability), we found a value that is inconsistent with the complete

dataset at the 95% level. This represents a possible bias towards higher EBL intensities when variability is not accounted for.

The bias inherent to this situation stems directly from the distribution of spectral variability in AGN, as it has been previously assumed (e.g. by Ackermann et al. (2012)) that AGN are likely to be found equally in harder or softer states than their integrated value, leading to an “averaging out”. Our results would suggest that this is not the case, in fact that allowing variable sources to contribute to this measurement causes an overestimate of EBL attenuation in such studies.

Comparing this to our measurement in Chapter 5, where we found that AGN (both BLL and FSRQ) have a far longer tail towards softer spectra when their deviation from the integrated value is considered, we can see that a propensity towards softer spectra would cause a dearth of flux at high energies when averaged over a long time period. This would result in a “softening” at higher energies, or higher spectral curvature at high energies when modelled using a log parabola. Comparing this to the modelling of the EBL, we can see that introducing an exponential cutoff term at these high energies would result in a spurious increase in the cutoff parameter, leading to a larger EBL signal than is present intrinsically. This effect could entirely explain the bias reported in Section 6.6.

The exclusion of a large fraction of the dataset to perform measurements of the EBL is of course not ideal. There is a substantial amount of information that is lost in this exercise, but this is largely recoverable given sufficient effort. Attempting to fit the EBL intensity within the monthly lightcurve bins results poses a difficult challenge due to the sparsity of the data in these small time periods. However, if instead an adaptive binning algorithm is applied, taking into account both spectral and flux variability, such as that developed by Scargle et al. (2013) and the similar method ap-

plied to *Fermi*-LAT data by Lott et al. (2012), then good quality SEDs can be generated for relatively quiescent time periods, allowing data from these variable sources to be re-included in the EBL analysis, effectively removing the systematic effects of spectral variability but maintaining statistical power.

Indeed, the recent work by Fermi-LAT Collaboration et al. (2018a), published since the completion of this work, separates out flaring periods of LAT AGN and considers the EBL attenuation separately with a greatly increased sample size, reducing the errant contribution of spectrally varying sources to the EBL. Thankfully, the restrictions placed upon IACTs (short observation times compared to the LAT) mean that these effects are already catered for as short time periods are considered separately (e.g. in Biteau & Williams (2015)). Depending on the particular tools chosen to analyse the EBL attenuation with the advent of CTA, care may have to be taken to avoid including spectral variability scatter if a *Fermi*-like analysis package is to be used. Indeed, IACT measurements are necessary to probe the lower wavelength ranges of the EBL intensity that cannot be reached with the LAT above \sim TeV energies.

6.8 Conclusions

In this Chapter I have studied the effect of spectral variability on measurements of the extragalactic background light using the *Fermi*-LAT telescope. The main conclusions are:

- I provide new and updated measurements of the EBL intensity in the *Fermi* energy range with smaller uncertainties than this effect's discovery by Ackermann et al. (2012), covering an extended redshift range up to $z = 2.1$.

- Applying these measurements to the more recent theoretical modelling of Inoue et al. (2014), I find good agreement ($< 2\sigma$ distance) with all current theoretical models.
- Systematic uncertainties resulting from spectral variability of AGN are substantial, and there is evidence at the 95% level that a bias is introduced by including spectrally variable sources in this type of analysis.

Further improvements to this measurement are immediately possible given a larger sample size more consistent with a complete AGN catalog (e.g. Fermi-LAT Collaboration et al., 2018a), but care must be taken to avoid periods of significant variability. These findings in general corroborate our conclusion that spectral variability causes a distortion (or other artefacting) of gamma-ray SEDs that precludes simple physical measurements from integrated SEDs. Other areas where such variability will produce systematic bias are in integrated measurements of dwarf spheroidals in DM searches (Ackermann et al., 2015b) or measurements of axion-like particle fluctuations in integrated spectra (Ajello et al., 2016). Great care must be taken to produce meaningful physical results in these areas.

Conclusions

This thesis has explored the question of the existence of spectral variability in gamma-ray detected active galactic nuclei. Using the *Fermi*-LAT telescope, we have performed deep studies of individual AGN in the case of Centaurus A and IC 310 which both showed unusual spectral energy distributions that exhibited an upturn in their SEDs with increasing energy. In the case of IC 310 (Chapter 4), a particularly active high-frequency peaked BL Lac object, I have shown that this upturn was in fact an artefact of bimodal (or possibly multi-modal) emission with a particularly hard flaring state that can be detected on short timescales and a soft steady state that only becomes apparent over long baseline measurements. As with many gamma-ray analyses, there were many possible confounding systematic uncertainties due to the localisation of nearby objects and the type of instrument response functions used in this specific case. After we dealt with these factors and ruled out their influence to cause our signal, our temporal analysis showed beyond doubt that this spectral behaviour could be entirely explained by spectral variability. The exact physical processes responsible for such huge changes in peak emission are unknown, but have been suggested to result from magnetospheric effects from near the central

black hole’s event horizon, allowing such short-timescale flaring.

Yet in the case of Centaurus A (Chapter 3), the case proved more complicated. The SED of Cen A showed a particularly sharp broken power law (as opposed to the gradually curved spectrum of IC 310). In this case however, the low-energy spectrum showed slight variability (long-term variations of a baseline, not flaring behaviour as shown by BLL objects), whereas the high-energy spectrum showed no significant variability. This is expected to some extent since Cen A is a different type of object than the vast majority of AGN detected at gamma-ray energies. Its jet is not pointed directly towards the Earth, and so the observation angle of the jet is large and thus not conducive to fast variability that is often emphasised by relativistic effects. What this does mean, however, is that the break in the spectrum is not the result of a temporal artefact because of the long integration periods used in *Fermi*-LAT analysis. Instead this means the break is an intrinsic feature of Cen A’s emission. Several plausible scenarios exist, whether this is due to multi-component SSC models with at least two different populations of emitting particles in Cen A’s jet, or as a result of more exotic interpretations such as dark matter self-annihilation or populations of millisecond pulsars is yet unknown.

In Chapter 5, I developed the statistical methodology to measure spectral variability in the likelihood-based analysis used for *Fermi*-LAT data. I analysed the brightest 291 *Fermi*-LAT AGN to calculate spectral energy distributions and monthly lightcurves throughout 8 years of data. I found significant ($> 5\sigma$) spectral variability in almost half of the AGN studied at a 5σ level. As a population, both BL Lac objects and flat spectrum radio quasars both showed similar characteristics in their distribution of spectral variability as both had a preference for softer emission more often than harder emission. Looking further into the characteristics of variabil-

ity, I calculated power spectral densities for the populations of BL Lacs at different peak synchrotron emission (LSP, ISP, HSP), and also for BL Lacs as a whole and FSRQs. I found significant evidence (at the 2σ level) of a difference in PSD power-laws between all of the classes listed. This showed that BL Lacs have far less variable activity at lower frequencies (timescales of years) compared to FSRQs, and so the archetypical BL Lac is more dominated by short, bright flaring events than FSRQs. Similarly this trend is reproduced within the BL Lac population as the synchrotron peak energy increased, the dominance of short flaring events increased (lower power on long timescales). This is largely to be expected as those objects that are more FSRQ-like have far more active accretion discs to fuel long-term changes of FSRQs. It was also established that the harder-when-brighter paradigm holds for the bright AGN that were selected, though the BL Lac objects showed a far greater variance in their behaviour, with the FSRQ population being far more clustered.

Applying these results to a practical calculation of the attenuation to gamma-ray SEDs due to interactions from the extragalactic background light in Chapter 6, it is clear that that sources with significant spectral variability contribute to a bias in this measurement if not excluded or accounted for. Despite reducing the dataset by as much as 70% in order to exclude variability, the statistical uncertainty in the result increased by only 30%. This is due to the intrinsic scatter introduced by integrated spectra whose EBL attenuation is overwhelmed by unaccounted-for systematic uncertainty due to variability. We found that not only was spectral variability important to exclude this effect, but also flux variability. As we discovered in Chapter 5, spectral variability can only be detected in particularly bright sources, but the harder-when-brighter relation shows that spectral variability does correlate with flux variability, and so this is an entirely expected

consequence. Furthermore, the results of Chapter 5 showed that spectral variability is not normally distributed. AGN are more likely to be found in relatively softer states than in relatively harder states compared to their integrated average. This behaviour in itself explains the bias that was discovered in Chapter 6 whereby once the variable sources were removed from the sample, the amount of EBL attenuation is reduced. The measurement of the EBL intensity is notoriously complex at all wavelengths due to the irreducible foregrounds and confounding factors, and while measurements at GeV–TeV have a more approachable background, this is still not a simple calculation.

The range of phenomena that can be measured using gamma-ray data is huge. Many fundamental problems in particle physics can be approached due to the extreme energies generated in the jets of AGN, pulsars, and even in diffuse interactions within nearby galaxies. As I have shown, spectral variability must be carefully taken into account in such measurements, especially when AGN are the vehicle by which the measurements are made. One particular example is that searches for photon-axion-like particle interactions within gamma-ray spectra, whereby such beyond the standard model interactions are expected to cause oscillations within a particularly bright AGN’s SED. As shown in Section 4, variability can easily cause hugely “significant” spectral artefacts when only considered over long-term integrations. Ajello et al. (2016) uses a 6-year integration of *Fermi*-LAT data to perform this calculation on NGC 1275 to set limits on photon-ALP couplings. NGC 1275 has well-documented evidence of spectral variability during its flaring periods which were not accounted for in this work. No doubt the systematic bias we found in Chapter 6 applies moreso to measurements as sensitive as photon-ALP couplings.

Certainly, future measurements must take care when dealing with AGN

in measurements of fundamental physics. Thankfully, the tool set for this analysis is ever-expanding with the advent of `fermipy` (Wood et al., 2017) amongst others allowing complex analysis with less code required from the user. Once complicated and fallible tasks such as scripting lightcurve analysis and SEDs have been made simpler during the time the research presented here was undertaken. As similar tools are developed for use with future instruments such as the Cherenkov Telescope Array’s software `ctools` and `gammapy` (Knödlseder et al., 2016), it is necessary to ensure effects such as spectral variability are being included in the survey simulations to provide accurate estimations of sensitivity to such fundamental physical measurements. Importantly, the greatly reduced integration times inherent with IACT measurements does hugely reduce the systematic uncertainty that can be introduced by spectral variability on the monthly timescale investigated in this work.

At the timescales we have probed in this work, in only the brightest objects were curved spectra significant enough to be modelled accurately. This extra degree of spectral freedom adds complexity, but contains far more context about the physical behaviour of the objects in question. Future characterisation of the variability of source SEDs will open a new chapter on modelling AGN flaring behaviour, specifically on the exact acceleration and cooling scenarios happening within the jet. These measurements will also be key to studying the mixture of hadronic vs leptonic emission that is still unclear with current measurements. The hadronic emission expected to be present in astrophysical jets and thought to produce the majority of ultra-high energy cosmic rays detected over the past several decades of research result in very subtle (and fast) alterations of the SED at high energies (Böttcher et al., 2013).

Inevitably the solution to all of the outstanding problems discussed here

will require observations not just over the gamma-ray energy range but across the entire electromagnetic spectrum, if not from multi-messenger astronomy. The emission of AGN expands across such a huge energy spectrum that multiwavelength campaigns are required for truly unassailable conclusions to these questions. Similarly, the problems established in particle physics will require subsequent investigation by other means, whether by direct detection or in collider studies for prospective WIMP dark matter detections or by microwave cavity in the search for axion-like particles. The combined weight of observations from broadband EM observations, with the new addition of gravitational wave astronomy and cosmic ray detections shall provide a new window to the Universe.

The analysis presented in Chapter 5 showed the huge improvements that have been made with eight years of data compared to the two years in the previous studies by the *Fermi*-LAT collaboration. As the LAT has now been on orbit for over 11 years, the growing dataset (and developing analysis tools) has been a crucial window to the MeV–TeV wavelength range that is unobserved by all other major instruments. The mission has been extended long past its original 5-year lifespan and while the science that is still ongoing is monumental, only a much longer exposure would have capability to solve the problems posed in this thesis once and for all, and this will require a new GeV–TeV instrument.

The construction of CTA is currently ongoing, with work now having started on the Northern array and expected to start soon on the Southern array by 2021. Active Galactic Nuclei are the subject of one of the key science projects, with both long-term monitoring providing lightcurve information and high-quality spectra of 40 blazars and radio galaxies to be published. These data will be instrumental for further analysis of spectral variability and its physical origin. Similarly, the planned deep exposure of

CHAPTER 7. CONCLUSIONS

Cen A (along with M 87) will be vital for deconvolving the physics of the spectrum explored in Chapter 3.

7.1 Acknowledgements

The work presented herein has relied upon the dedicated work of generations of astronomers, collaborations, and funding bodies.

I have been supported through my studies by an STFC studentship, grant reference ST/N50404X/1.

The Digitized Sky Surveys were produced at the Space Telescope Science Institute under U.S. Government grant NAG W-2166. The images of these surveys are based on photographic data obtained using the Oschin Schmidt Telescope on Palomar Mountain and the UK Schmidt Telescope. The plates were processed into the present compressed digital form with the permission of these institutions.

The UK Schmidt Telescope was operated by the Royal Observatory Edinburgh, with funding from the UK Science and Engineering Research Council (later the UK Particle Physics and Astronomy Research Council), until 1988 June, and thereafter by the Anglo-Australian Observatory. The blue plates of the southern Sky Atlas and its Equatorial Extension (together known as the SERC-J), as well as the Equatorial Red (ER), and the Second Epoch [red] Survey (SES) were all taken with the UK Schmidt.

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.

Part of this work is based on archival data, software or online services provided by the Space Science Data Center - ASI.

I would like to thank my fantastic supervisory team, Paula Chadwick and Anthony Brown, without whom I would have been flying blind. Their support has been unwavering, and their guidance valuable.

Appendix A

Catalogue of Spectral Variability

A.1 Catalogue Data Selection

Table A.1: The AGN selected for use in Section 5.2. Definitions as described in text.

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J0008.0+4713	2.0166	47.2281	24.7448	bll	MG4 J000800+4712
J0009.3+5030	2.3337	50.5104	36.4933	bll	NVSS J000922+503028
J0022.5+0608	5.643	6.1378	28.1753	bll	PKS 0019+058
J0030.3-4223	7.5895	-42.3954	28.7538	fsrq	PKS 0027-426
J0033.6-1921	8.4022	-19.3638	39.216	bll	KUV 00311-1938
J0043.8+3425	10.9692	34.4294	22.6455	fsrq	GB6 J0043+3426
J0045.3+2126	11.3494	21.4459	22.3391	bll	GB6 J0045+2127
J0050.6-0929	12.6744	-9.4849	41.9779	BLL	PKS 0048-09
J0051.2-6241	12.8196	-62.6881	23.4147	bcu	1RXS J005117.7-624154
J0058.0-3233	14.5128	-32.5666	20.5244	bll	PKS 0055-328
J0059.0-7242e	14.75	-72.7	27.9457	GAL	SMC
J0100.2+0745	15.0678	7.7662	23.0848	bll	GB6 J0100+0745
J0108.7+0134	17.1804	1.579	71.654	fsrq	4C +01.02
J0112.1+2245	18.0266	22.759	81.1326	BLL	S2 0109+22

Continued on next page

A.1. CATALOGUE DATA SELECTION

Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J0112.8+3207	18.21	32.127	44.5418	fsrq	4C +31.03
J0118.8-2142	19.707	-21.7079	28.4291	fsrq	PKS 0116-219
J0120.4-2700	20.1239	-27.015	45.6367	bll	PKS 0118-272
J0132.6-1655	23.1707	-16.9179	34.5813	fsrq	PKS 0130-17
J0136.5+3905	24.1305	39.0963	56.9168	bll	B3 0133+388
J0137.0+4752	24.2569	47.8764	55.9446	fsrq	OC 457
J0137.6-2430	24.4147	-24.5104	22.526	fsrq	PKS 0135-247
J0141.4-0929	25.367	-9.4851	29.4619	bll	PKS 0139-09
J0144.6+2705	26.1566	27.0888	50.6147	bll	TXS 0141+268
J0145.1-2732	26.2921	-27.5339	29.4705	fsrq	PKS 0142-278
J0154.0+0824	28.5104	8.4096	24.812	bll	GB6 J0154+0823
J0203.6+3043	30.9227	30.7172	36.8994	bll	NVSS J020344+304238
J0205.2-1700	31.3233	-17.0054	24.0475	fsrq	PKS 0202-17
J0209.4-5229	32.3707	-52.4979	28.0731	bll	1RXS J020922.2-522920
J0210.7-5101	32.6972	-51.0282	62.9476	BCU	PKS 0208-512
J0211.2+1051	32.817	10.8578	50.7517	BLL	MG1 J021114+1051
J0217.8+0143	34.4681	1.7328	46.2922	fsrq	PKS 0215+015
J0221.1+3556	35.2801	35.9407	48.4395	FSRQ	B0218+357
J0222.6+4301	35.6735	43.0329	121.259	BLL	3C 66A
J0236.7-6136	39.1955	-61.6004	31.7846	fsrq	PKS 0235-618
J0237.9+2848	39.4801	28.8047	105.192	FSRQ	4C +28.07
J0238.6+1636	39.6672	16.6152	88.0341	BLL	AO 0235+164
J0245.9-4651	41.4838	-46.8527	73.4918	fsrq	PKS 0244-470
J0252.8-2218	43.2007	-22.3105	69.6904	fsrq	PKS 0250-225
J0303.4-2407	45.871	-24.1226	63.625	BLL	PKS 0301-243
J0309.0+1029	47.2586	10.498	25.8408	fsrq	PKS 0306+102
J0309.9-6057	47.478	-60.9628	26.0419	fsrq	PKS 0308-611
J0316.1+0904	49.0419	9.0773	23.5387	bll	GB6 J0316+0904
J0319.8+4130	49.9586	41.5137	108.538	RDG	NGC 1275
J0325.5+2223	51.3995	22.3987	22.1763	fsrq	TXS 0322+222
J0334.3-4008	53.577	-40.1478	58.3261	bll	PKS 0332-403

Continued on next page

APPENDIX A. CATALOGUE OF SPECTRAL VARIABILITY

Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J0334.3-3726	53.582	-37.4441	36.6941	bll	PMN J0334-3725
J0339.5-0146	54.8826	-1.7731	45.0043	fsrq	PKS 0336-01
J0349.9-2102	57.4778	-21.0486	25.373	fsrq	PKS 0347-211
J0403.9-3604	60.9938	-36.0772	103.562	FSRQ	PKS 0402-362
J0407.1-3825	61.7958	-38.42	29.077	fsrq	PKS 0405-385
J0423.2-0119	65.8247	-1.3321	56.1391	FSRQ	PKS 0420-01
J0424.7+0035	66.1781	0.5951	22.0226	bll	PKS 0422+00
J0428.6-3756	67.1684	-37.9357	135.569	bll	PKS 0426-380
J0433.6+2905	68.4183	29.0905	29.1413	bll	MG2 J043337+2905
J0442.6-0017	70.6617	-0.2989	50.8425	fsrq	PKS 0440-00
J0449.0+1121	72.2514	11.3589	35.6363	fsrq	PKS 0446+11
J0449.4-4350	72.3618	-43.8364	96.5072	bll	PKS 0447-439
J0453.2-2808	73.3142	-28.147	32.0458	fsrq	PKS 0451-28
J0455.7-4617	73.9447	-46.2845	30.0136	fsrq	PKS 0454-46
J0457.0-2324	74.2639	-23.412	161.081	FSRQ	PKS 0454-234
J0501.2-0157	75.3182	-1.9659	32.4879	fsrq	S3 0458-02
J0505.3+0459	76.3432	4.9979	24.7344	fsrq	PKS 0502+049
J0507.1-6102	76.7803	-61.05	20.0192	fsrq	PKS 0506-61
J0508.0+6736	77.0035	67.6143	39.2811	bll	1ES 0502+675
J0509.4+0541	77.3639	5.699	54.5026	bll	TXS 0506+056
J0516.7-6207	79.1889	-62.1211	38.5485	bll	PKS 0516-621
J0522.9-3628	80.7416	-36.4715	70.4297	BCU	PKS 0521-36
J0526.6-6825e	81.65	-68.42	83.8399	GAL	LMC
J0530.8+1330	82.7187	13.5148	35.9048	FSRQ	PKS 0528+134
J0532.0-4827	83.0004	-48.4585	73.8326	BCU	PMN J0531-4827
J0532.7+0732	83.1769	7.5446	50.9016	FSRQ	OG 050
J0538.8-4405	84.7091	-44.0859	219.161	BLL	PKS 0537-441
J0543.9-5531	85.9775	-55.5318	25.3013	bll	1RXS J054357.3-553206
J0601.2-7036	90.3132	-70.6086	39.5656	fsrq	PKS 0601-70
J0607.4+4739	91.8598	47.6538	30.6481	bll	TXS 0603+476
J0608.1-1522	92.031	-15.3795	22.3166	fsrq	PMN J0608-1520

Continued on next page

A.1. CATALOGUE DATA SELECTION

Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J0612.8+4122	93.2079	41.3712	45.8131	bll	B3 0609+413
J0627.0-3529	96.7727	-35.4893	20.1217	rdg	PKS 0625-35
J0629.4-1959	97.3581	-19.9857	33.6274	bll	PKS 0627-199
J0630.9-2406	97.7469	-24.1101	38.3969	bll	TXS 0628-240
J0635.7-7517	98.9417	-75.2926	22.7864	fsrq	PKS 0637-75
J0644.3-6713	101.088	-67.2227	47.8154	bcb	PKS 0644-671
J0650.7+2503	102.697	25.0632	29.4866	bll	1ES 0647+250
J0654.4+4514	103.604	45.2469	39.3749	FSRQ	B3 0650+453
J0654.4+5042	103.609	50.702	26.3108	fsrq	GB6 J0654+5042
J0700.6-6610	105.158	-66.173	36.9975	bll	PKS 0700-661
J0701.4-4634	105.35	-46.5735	21.1885	fsrq	PKS 0700-465
J0707.0+7741	106.755	77.6982	25.3069	bll	NVSS J070651+774137
J0712.6+5033	108.172	50.5563	21.2015	bll	GB6 J0712+5033
J0713.9+1933	108.496	19.5591	33.6123	fsrq	MG2 J071354+1934
J0718.7-4319	109.69	-43.33	28.7444	bll	PMN J0718-4319
J0719.3+3307	109.829	33.1253	66.5521	fsrq	B2 0716+33
J0721.9+7120	110.49	71.3452	196.552	BLL	S5 0716+71
J0725.2+1425	111.323	14.4312	69.8997	FSRQ	4C +14.23
J0738.1+1741	114.536	17.698	55.3463	bll	PKS 0735+17
J0739.4+0137	114.851	1.6274	38.9864	fsrq	PKS 0736+01
J0742.6+5444	115.66	54.7337	61.6222	fsrq	GB6 J0742+5444
J0757.0+0956	119.274	9.9475	24.4734	bll	PKS 0754+100
J0805.4+7534	121.352	75.5815	26.6485	bll	RX J0805.4+7534
J0808.2-0751	122.072	-7.8519	73.8415	fsrq	PKS 0805-07
J0809.8+5218	122.454	52.3077	56.1573	bll	1ES 0806+524
J0811.2-7529	122.811	-75.4915	30.7794	bll	PMN J0810-7530
J0811.3+0146	122.848	1.7812	28.747	bll	OJ 014
J0814.7+6428	123.677	64.4799	26.6705	bll	GB6 J0814+6431
J0816.4-1311	124.116	-13.1908	26.4811	bll	PMN J0816-1311
J0818.2+4223	124.562	42.399	67.438	bll	S4 0814+42
J0830.7+2408	127.681	24.1378	21.3038	FSRQ	OJ 248

Continued on next page

APPENDIX A. CATALOGUE OF SPECTRAL VARIABILITY

 Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J0831.9+0430	127.991	4.5032	42.2809	bll	PKS 0829+046
J0841.4+7053	130.373	70.8972	34.5315	FSRQ	S5 0836+71
J0849.9+5108	132.482	51.1445	32.3593	NLSY1	SBS 0846+513
J0850.0+4855	132.52	48.9197	34.5004	bll	GB6 J0850+4855
J0850.2-1214	132.557	-12.234	40.3552	fsrq	PMN J0850-1213
J0854.8+2006	133.7	20.1121	56.513	BLL	OJ 287
J0856.7-1105	134.18	-11.0862	30.7993	bll	PMN J0856-1105
J0909.1+0121	137.298	1.358	34.0632	fsrq	PKS 0906+01
J0909.8-0229	137.458	-2.4997	26.3509	fsrq	PKS 0907-023
J0915.8+2933	138.965	29.5604	30.0152	bll	Ton 0396
J0920.9+4442	140.23	44.7048	82.1091	fsrq	S4 0917+44
J0921.8+6215	140.453	62.2665	25.8382	fsrq	OK 630
J0930.2+8612	142.562	86.202	21.066	bll	S5 0916+864
J0946.5+1017	146.636	10.2881	23.3297	fsrq	TXS 0943+105
J0948.8+0021	147.223	0.3629	39.3908	NLSY1	PMN J0948+0022
J0953.0-0839	148.255	-8.6635	29.8033	bll	PMN J0953-0840
J0957.6+5523	149.413	55.3866	112.857	fsrq	4C +55.17
J0958.6+6534	149.66	65.5667	27.1564	bll	S4 0954+65
J1006.7-2159	151.7	-21.9854	21.6524	fsrq	PKS 1004-217
J1012.6+2439	153.152	24.6595	40.842	fsrq	MG2 J101241+2439
J1015.0+4925	153.771	49.4326	84.5748	bll	1H 1013+498
J1016.0+0513	154.003	5.2167	27.1731	fsrq	TXS 1013+054
J1024.1-3232	156.041	-32.539	21.0536	fsrq	PKS 1021-323
J1026.9-1750	156.729	-17.8399	23.0653	bll	1RXS J102658.5-174905
J1033.2+4116	158.303	41.2763	24.6545	fsrq	S4 1030+41
J1033.8+6051	158.461	60.8593	51.4352	FSRQ	S4 1030+61
J1037.5+5711	159.389	57.1949	51.2286	bll	GB6 J1037+5711
J1040.4+0615	160.116	6.2604	20.5135	bcu	GB6 J1040+0617
J1048.4+7144	162.122	71.7422	40.6629	FSRQ	S5 1044+71
J1054.5+2210	163.646	22.18	22.6867	bll	87GB 105148.6+222705
J1058.5+0133	164.626	1.5643	55.0161	BLL	4C +01.28

Continued on next page

A.1. CATALOGUE DATA SELECTION

Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J1058.5-8003	164.631	-80.0512	38.9524	bll	PKS 1057-79
J1058.6+5627	164.665	56.4549	54.6339	BLL	TXS 1055+567
J1059.2-1133	164.803	-11.5514	32.9315	bll	PKS B1056-113
J1104.4+3812	166.118	38.2099	190.35	BLL	Mkn 421
J1112.4+3449	168.116	34.8233	24.4789	fsrq	TXS 1109+350
J1117.0+2014	169.274	20.2334	26.4147	bll	RBS 0958
J1120.8+4212	170.21	42.2052	23.5287	bll	RBS 0970
J1121.4-0554	170.367	-5.9016	29.3266	fsrq	PKS 1118-05
J1127.0-1857	171.763	-18.9558	76.2633	fsrq	PKS 1124-186
J1129.9-1446	172.499	-14.7737	22.0746	fsrq	PKS 1127-14
J1132.7+0034	173.186	0.5749	27.4725	bll	PKS B1130+008
J1136.6+7009	174.153	70.1517	27.2228	bll	Mkn 180
J1146.8+3958	176.724	39.9695	48.7004	fsrq	S4 1144+40
J1147.0-3811	176.764	-38.1846	22.4033	fsrq	PKS 1144-379
J1150.5+4155	177.635	41.9196	33.9734	bll	RBS 1040
J1153.4+4932	178.35	49.5358	42.4854	FSRQ	OM 484
J1159.5+2914	179.882	29.2377	93.3978	fsrq	Ton 599
J1208.7+5442	182.188	54.7067	37.2308	fsrq	TXS 1206+549
J1217.8+3007	184.465	30.1178	54.7458	bll	1ES 1215+303
J1221.3+3010	185.341	30.1766	38.9234	bll	PG 1218+304
J1221.4+2814	185.374	28.2401	49.8504	bll	W Comae
J1222.4+0414	185.611	4.234	30.0163	fsrq	4C +04.42
J1224.9+2122	186.226	21.3823	219.674	FSRQ	4C +21.35
J1229.1+0202	187.287	2.041	149.012	FSRQ	3C 273
J1231.7+2847	187.946	28.7923	36.957	bll	B2 1229+29
J1239.5+0443	189.895	4.7192	64.9973	fsrq	MG1 J123931+0443
J1243.1+3627	190.797	36.4537	31.2353	bll	Ton 116
J1246.7-2547	191.689	-25.7912	92.644	fsrq	PKS 1244-255
J1248.2+5820	192.07	58.3479	55.5287	bll	PG 1246+586
J1253.2+5300	193.301	53.0133	47.5189	bll	S4 1250+53
J1256.1-0547	194.041	-5.7912	160.596	FSRQ	3C 279

Continued on next page

APPENDIX A. CATALOGUE OF SPECTRAL VARIABILITY

Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J1258.6-1800	194.65	-18.0039	22.4215	fsrq	PKS B1256-177
J1303.0+2435	195.752	24.5948	28.5482	bll	MG2 J130304+2434
J1304.3-4353	196.096	-43.8995	33.5281	bll	1RXS J130421.2-435308
J1307.6-4300	196.918	-43.0007	23.7	bcu	1RXS J130737.8-425940
J1309.3+4304	197.345	43.0749	29.0478	bll	B3 1307+433
J1310.6+3222	197.657	32.3697	51.9702	fsrq	OP 313
J1312.5-2155	198.126	-21.9322	20.8226	bll	PKS 1309-216
J1312.7+4828	198.188	48.474	66.0435	agn	GB 1310+487
J1314.8+2349	198.714	23.8256	21.8897	bll	TXS 1312+240
J1316.0-3338	199.0	-33.6486	25.3009	fsrq	PKS 1313-333
J1321.0+2215	200.269	22.2628	20.5839	fsrq	TXS 1318+225
J1325.4-4301	201.367	-43.0305	43.8557	RDG	Cen A Core
J1326.8+2211	201.722	22.1844	26.2161	fsrq	B2 1324+22
J1332.0-0508	203.015	-5.1476	49.7059	fsrq	PKS 1329-049
J1344.2-1724	206.051	-17.401	33.1461	fsrq	PMN J1344-1723
J1345.6+4453	206.401	44.8942	75.773	fsrq	B3 1343+451
J1350.8+3035	207.704	30.5985	20.2155	fsrq	B2 1348+30B
J1408.8-0751	212.204	-7.8565	20.841	fsrq	PKS B1406-076
J1418.4-0233	214.609	-2.5638	33.0228	bll	NVSS J141826-023336
J1418.5+3543	214.65	35.7186	36.3181	BCU	NVSS J141828+354250
J1419.9+5425	214.984	54.4195	20.08	bll	OQ 530
J1427.0+2347	216.757	23.7998	107.355	BLL	PKS 1424+240
J1427.9-4206	216.991	-42.1106	129.82	FSRQ	PKS B1424-418
J1438.7+3710	219.688	37.1679	28.0608	fsrq	B2 1436+37B
J1444.0-3907	221.009	-39.1299	33.9898	bll	PKS 1440-389
J1448.0+3608	222.01	36.1445	23.0564	bll	RBS 1432
J1454.5+5124	223.63	51.4084	33.6076	bll	TXS 1452+516
J1457.4-3539	224.355	-35.6567	62.1986	FSRQ	PKS 1454-354
J1500.9+2238	225.241	22.6357	20.4039	bll	MS 1458.8+2249
J1504.4+1029	226.105	10.4936	189.471	FSRQ	PKS 1502+106
J1505.1+0326	226.278	3.4488	21.9727	nlsy1	PKS 1502+036

Continued on next page

A.1. CATALOGUE DATA SELECTION

Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J1512.2+0202	228.065	2.0334	23.6835	fsrq	PKS 1509+022
J1512.8-0906	228.213	-9.1032	207.272	FSRQ	PKS 1510-08
J1517.6-2422	229.421	-24.3756	45.3534	bll	AP Librae
J1522.1+3144	230.543	31.7365	192.683	fsrq	B2 1520+31
J1522.6-2730	230.662	-27.5082	25.9911	bll	PKS 1519-273
J1532.7-1319	233.188	-13.3249	35.6348	bcu	TXS 1530-131
J1542.9+6129	235.738	61.4953	62.6687	bll	GB6 J1542+6129
J1549.4+0237	237.372	2.6309	23.3843	fsrq	PKS 1546+027
J1553.5+1256	238.386	12.9457	32.5297	fsrq	PKS 1551+130
J1555.7+1111	238.936	11.1939	97.5447	BLL	PG 1553+113
J1604.6+5714	241.17	57.2413	26.1786	fsrq	GB6 J1604+5714
J1607.0+1551	241.764	15.8624	27.6048	fsrq	4C +15.54
J1610.8-6649	242.706	-66.8222	26.2702	bll	PMN J1610-6649
J1625.7-2527	246.445	-25.4556	56.2572	fsrq	PKS 1622-253
J1626.0-2951	246.518	-29.8624	20.9228	FSRQ	PKS 1622-29
J1630.6+8232	247.652	82.535	20.9523	rdg	NGC 6251
J1635.2+3809	248.812	38.1525	86.9665	FSRQ	4C +38.41
J1637.7+4715	249.431	47.264	31.9848	fsrq	4C +47.44
J1640.6+3945	250.159	39.7553	28.2796	FSRQ	NRAO 512
J1642.9+3950	250.748	39.8487	20.5269	FSRQ	3C 345
J1647.4+4950	251.873	49.8372	29.3034	bcu	SBS 1646+499
J1653.9+3945	253.478	39.754	93.3205	BLL	Mkn 501
J1657.7+4807	254.428	48.1316	28.0647	fsrq	4C +48.41
J1700.1+6829	255.05	68.4849	49.4491	fsrq	TXS 1700+685
J1703.6-6211	255.917	-62.1956	34.5591	fsrq	MRC 1659-621
J1709.6+4318	257.424	43.3125	62.2035	fsrq	B3 1708+433
J1719.2+1744	259.811	17.7421	20.5789	bll	PKS 1717+177
J1723.9+4004	260.986	40.0684	21.7139	fsrq	S4 1722+40
J1725.0+1152	261.262	11.8833	32.7156	bll	1H 1720+117
J1727.1+4531	261.8	45.5272	29.9227	fsrq	S4 1726+45
J1730.5+0023	262.647	0.3971	22.7763	fsrq	PKS 1728+004

Continued on next page

APPENDIX A. CATALOGUE OF SPECTRAL VARIABILITY

 Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J1733.0-1305	263.256	-13.0858	36.8071	FSRQ	PKS 1730-13
J1734.3+3858	263.586	38.9775	34.6523	fsrq	B2 1732+38A
J1739.4+4955	264.873	49.9269	22.5296	fsrq	S4 1738+49
J1740.3+5211	265.085	52.1909	29.2728	fsrq	4C +51.37
J1748.6+7005	267.169	70.0968	52.6903	bll	S4 1749+70
J1751.5+0939	267.882	9.6578	35.3532	bll	OT 081
J1754.1+3212	268.546	32.2082	41.6282	bll	RX J1754.1+3212
J1800.5+7827	270.131	78.4618	74.4577	bll	S5 1803+784
J1806.7+6949	271.681	69.8179	46.1173	bll	3C 371
J1824.2+5649	276.058	56.8329	34.5206	bll	4C +56.27
J1825.2-5230	276.321	-52.5093	21.5724	bcu	PKS 1821-525
J1829.6+4844	277.414	48.7488	24.1047	ssrq	3C 380
J1838.8+4802	279.702	48.0371	20.1531	bll	GB6 J1838+4802
J1849.2+6705	282.324	67.0928	63.2795	FSRQ	S4 1849+67
J1852.4+4856	283.108	48.9469	29.3405	fsrq	S4 1851+48
J1903.2+5541	285.821	55.6852	46.9488	bll	TXS 1902+556
J1911.2-2006	287.802	-20.1117	32.1197	fsrq	PKS B1908-201
J1912.9-8008	288.234	-80.1495	21.3465	fsrq	PKS 1903-80
J1917.7-1921	289.443	-19.3658	28.7994	bll	1H 1914-194
J1918.2-4110	289.569	-41.176	27.7793	bll	PMN J1918-4111
J1923.5-2104	290.885	-21.0726	41.0319	fsrq	TXS 1920-211
J1924.8-2914	291.219	-29.2348	27.4199	fsrq	PKS B1921-293
J1926.8+6154	291.707	61.9017	26.0934	bll	1RXS J192649.5+615445
J1958.0-3847	299.518	-38.7843	37.6616	fsrq	PKS 1954-388
J1959.1-4245	299.781	-42.7604	22.5223	fsrq	PMN J1959-4246
J2000.0+6509	300.016	65.1544	61.9246	bll	1ES 1959+650
J2001.0-1750	300.259	-17.8451	21.0597	fsrq	PKS 1958-179
J2005.2+7752	301.31	77.8811	25.4181	bll	S5 2007+77
J2009.3-4849	302.349	-48.8277	40.2979	BLL	PKS 2005-489
J2025.6-0736	306.422	-7.6013	63.1062	fsrq	PKS 2023-07
J2039.0-1047	309.752	-10.7984	25.8845	bll	TXS 2036-109

Continued on next page

A.1. CATALOGUE DATA SELECTION

Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J2056.2-4714	314.067	-47.2482	78.948	fsrq	PKS 2052-47
J2103.9-6233	315.986	-62.561	21.148	bcu	PMN J2103-6232
J2116.1+3339	319.047	33.6504	36.0159	bll	B2 2114+33
J2121.0+1901	320.269	19.0235	27.4131	fsrq	OX 131
J2126.5-4605	321.628	-46.087	34.0277	FSRQ	PKS 2123-463
J2139.4-4235	324.852	-42.5985	45.7504	bll	MH 2136-428
J2143.5+1744	325.896	17.7337	60.5729	fsrq	OX 169
J2147.2+0929	326.81	9.498	29.4183	FSRQ	PKS 2144+092
J2147.3-7536	326.84	-75.6088	75.8167	FSRQ	PKS 2142-75
J2151.8-3025	327.953	-30.4286	27.6853	fsrq	PKS 2149-306
J2157.5+3126	329.38	31.4497	56.6653	fsrq	B2 2155+31
J2158.8-3013	329.72	-30.2268	142.102	bll	PKS 2155-304
J2202.4-8339	330.61	-83.656	23.0741	fsrq	PKS 2155-83
J2202.7+4217	330.687	42.2835	128.827	BLL	BL Lacertae
J2203.4+1725	330.864	17.4274	54.1404	fsrq	PKS 2201+171
J2225.8-0454	336.468	-4.9007	26.9563	fsrq	3C 446
J2229.7-0833	337.434	-8.5511	58.5735	FSRQ	PKS 2227-08
J2232.5+1143	338.139	11.7203	72.1985	FSRQ	CTA 102
J2236.3+2829	339.077	28.4889	52.1409	bll	B2 2234+28A
J2236.5-1432	339.135	-14.5455	69.547	BLL	PKS 2233-148
J2243.4-2541	340.864	-25.6945	25.8581	bll	PKS 2240-260
J2243.9+2021	340.991	20.3654	45.9231	bll	RGB J2243+203
J2250.7-2806	342.68	-28.1161	48.3886	bcu	PMN J2250-2806
J2254.0+1608	343.502	16.1459	480.741	FSRQ	3C 454.3
J2258.0-2759	344.512	-27.9853	56.7738	fsrq	PKS 2255-282
J2311.0+3425	347.764	34.4231	66.8727	FSRQ	B2 2308+34
J2321.9+3204	350.479	32.0704	26.3757	fsrq	B2 2319+31
J2323.5-0315	350.898	-3.2647	39.3481	fsrq	PKS 2320-035
J2323.9+4211	350.981	42.1839	25.1158	bll	1ES 2321+419
J2324.7-4040	351.184	-40.6754	22.2179	bll	1ES 2322-409
J2325.2+3957	351.32	39.9543	23.3974	bll	B3 2322+396

Continued on next page

APPENDIX A. CATALOGUE OF SPECTRAL VARIABILITY

Table A.1 – *Continued from previous page*

Name	RA	Dec.	Sig. (σ)	Class	Alt. Name
J2325.3-3557	351.345	-35.9546	32.0572	fsrq	CTS 0490
J2327.7+0941	351.944	9.6863	21.5131	fsrq	PKS 2325+093
J2328.4-4034	352.113	-40.5797	23.8975	bcu	PKS 2325-408
J2329.3-4955	352.331	-49.9233	130.01	fsrq	PKS 2326-502
J2330.8-2144	352.709	-21.746	26.3731	fsrq	PMN J2331-2148
J2340.7+8016	355.199	80.2707	32.1353	bll	1RXS J234051.4+801513
J2345.2-1554	356.307	-15.9131	82.3238	FSRQ	PMN J2345-1555
J2348.0-1630	357.009	-16.5019	26.0016	fsrq	PKS 2345-16

Appendix B

Error on EBL Attenuated Flux

The relative error on the flux $\frac{dN}{dE} = \Phi$ presented in Equation 6.2 and shown as shaded areas when plotting SEDs after fitting are calculated as follows,

$$\begin{aligned}
\sigma_{\Phi}^2 = & \left(\frac{\partial \Phi}{\partial N_0} \right)^2 \sigma_{N_0}^2 + \left(\frac{\partial \Phi}{\partial \alpha} \right)^2 \sigma_{\alpha}^2 + \left(\frac{\partial \Phi}{\partial \beta} \right)^2 \sigma_{\beta}^2 + \left(\frac{\partial \Phi}{\partial a} \right)^2 \sigma_a^2 \\
& + 2 \left(\frac{\partial \Phi}{\partial N_0} \right) \left(\frac{\partial \Phi}{\partial \alpha} \right) \sigma_{N_0 \alpha} + 2 \left(\frac{\partial \Phi}{\partial N_0} \right) \left(\frac{\partial \Phi}{\partial \beta} \right) \sigma_{N_0 \beta} \\
& + 2 \left(\frac{\partial \Phi}{\partial N_0} \right) \left(\frac{\partial \Phi}{\partial a} \right) \sigma_{N_0 a} + 2 \left(\frac{\partial \Phi}{\partial \alpha} \right) \left(\frac{\partial \Phi}{\partial \beta} \right) \sigma_{\alpha \beta} \\
& + 2 \left(\frac{\partial \Phi}{\partial \alpha} \right) \left(\frac{\partial \Phi}{\partial a} \right) \sigma_{\alpha a} + 2 \left(\frac{\partial \Phi}{\partial \beta} \right) \left(\frac{\partial \Phi}{\partial a} \right) \sigma_{\beta a}.
\end{aligned} \tag{B.1}$$

The values σ_{xy} represent the covariance between parameters x and y , as calculated in the error matrix. Calculating the partial derivatives of Φ with respect to the parameters, we find:

$$\frac{\partial \Phi}{\partial N_0} = \frac{\Phi}{N_0} \tag{B.2}$$

$$\frac{\partial \Phi}{\partial \alpha} = -\Phi \log\left(\frac{E}{E_0}\right) \tag{B.3}$$

$$\frac{\partial \Phi}{\partial \beta} = -\Phi \log^2\left(\frac{E}{E_0}\right) \tag{B.4}$$

$$\frac{\partial \Phi}{\partial a} = -\Phi \tau(E, z), \tag{B.5}$$

in which for plotting purposes we have ignored the error on z as this error is external to our fitting. As such, the fractional variance is calculated as:

$$\begin{aligned}
\left(\frac{\sigma_\Phi}{\Phi}\right)^2 &= \left(\frac{\sigma_{N_0}}{N_0}\right)^2 + \log^2\left(\frac{E}{E_0}\right)\sigma_\alpha^2 + \\
&+ \log^4\left(\frac{E}{E_0}\right)\sigma_\beta^2 + \tau(E, z)^2\sigma_a^2 + \\
&- 2\log\left(\frac{E}{E_0}\right)\frac{1}{N_0}\sigma_{N_0\alpha} - 2\frac{1}{N_0}\log^2\left(\frac{E}{E_0}\right)\sigma_{N_0\beta} + \\
&- 2\frac{\tau(E, z)}{N_0}\sigma_{N_0a} + 2\log^3\left(\frac{E}{E_0}\right)\sigma_{\alpha\beta} + \\
&+ 2\log\left(\frac{E}{E_0}\right)\tau(E, z)\sigma_{\alpha a} + \log^2\left(\frac{E}{E_0}\right)\tau(E, z)\sigma_{\beta a}.
\end{aligned} \tag{B.6}$$

Phenomenological Spectral Models

In this Appendix, we will summarise the most relevant spectral models that are used in *Fermi*-LAT analysis and built into the *Fermi* science tools.

Constant value functions of the form:

$$\frac{dN}{dE} = C, \quad (\text{C.1})$$

where C is a constant flux value independent of energy. Power-laws of the form:

$$\frac{dN}{dE} = N \left(\frac{E}{E_0} \right)^{-\gamma} \quad (\text{C.2})$$

where N is a normalisation in units equivalent to $\text{MeV}^{-1} \text{cm}^{-2} \text{s}^{-1}$, E is energy, E_0 is the scale energy (an arbitrary value in the units of E in this case, but set close to the decorrelation energy for fitting purposes), and γ is the dimensionless power-law index. Log parabolae of the form:

$$\frac{dN}{dE} = N \left(\frac{E}{E_0} \right)^{-\alpha + \beta \log(\frac{E}{E_0})} \quad (\text{C.3})$$

where N , E , and E_0 are defined as before, α is a linear spectral index, and β is a curvature parameter. Broken power-laws of the form:

$$\frac{dN}{dE} = N_0 \times \begin{cases} \left(\frac{E}{E_b} \right)^{\gamma_1} & \text{if } E < E_b, \\ \left(\frac{E}{E_b} \right)^{\gamma_2} & \text{otherwise.} \end{cases} \quad (\text{C.4})$$

where γ_1 and γ_2 are the power-law indices above and below the break, and E_b is the break energy. The broken power-law above can also be represented as a “smoothly”

broken power-law:

$$\frac{dN}{dE} = N \left(\frac{E}{E_0} \right)^{-\gamma_1} \left(1 + \left(\frac{E}{E_b} \right)^{\frac{\gamma_1 - \gamma_2}{\beta}} \right)^{-\beta}, \quad (\text{C.5})$$

with the introduction of a curvature parameter β . Finally, a power-law with a super-exponential cutoff:

$$\frac{dN}{dE} = N \left(\frac{E}{E_0} \right)^{-\gamma} \exp \left(- \left(\frac{E}{E_c} \right)^{\gamma_2} \right), \quad (\text{C.6})$$

where E_c is the cutoff energy.

Bibliography

- Abazajian, K. N. 2011, , 3, 010
- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010a, ApJ, 715, 429
- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010b, ApJ, 715, 429
- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010c, ApJ, 722, 520
- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010d, ApJ, 710, 1271
- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010a, The Astrophysical Journal, 719, 1433
- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2011, ApJ, 736, 131
- Abdo, A. A. et al. 2010b, The Astrophysical Journal, 719, 1433
- Abdo, A. A. et al. 2013, The Astrophysical Journal Supplement Series, 208, 17
- Acciari, V. A., Ansoldi, S., Antonelli, L. A., et al. 2019, MNRAS, 486, 4233
- Acero, F., Ackermann, M., Ajello, M., et al. 2015, ApJS, 218, 23
- Acero, F., Ackermann, M., Ajello, M., et al. 2016, The Astrophysical Journal Supplement Series, 223, 26
- Ackermann, M., Ajello, M., Albert, A., et al. 2012, ApJS, 203, 4
- Ackermann, M., Ajello, M., Albert, A., et al. 2017a, ApJ, 836, 208

- Ackermann, M., Ajello, M., Allafort, A., et al. 2011, *ApJ*, 743, 171
- Ackermann, M., Ajello, M., Allafort, A., et al. 2013a, *ApJ*, 765, 54
- Ackermann, M., Ajello, M., Allafort, A., et al. 2013b, *ApJS*, 209, 34
- Ackermann, M., Ajello, M., Allafort, A., et al. 2012, *Science*, 338, 1190
- Ackermann, M., Ajello, M., Asano, K., et al. 2013c, *ApJS*, 209, 11
- Ackermann, M., Ajello, M., Atwood, W. B., et al. 2016, *ApJS*, 222, 5
- Ackermann, M., Ajello, M., Atwood, W. B., et al. 2015a, *ApJ*, 810, 14
- Ackermann, M., Ajello, M., Baldini, L., et al. 2018, *ApJS*, 237, 32
- Ackermann, M., Ajello, M., Baldini, L., et al. 2017b, *ApJ*, 843, 139
- Ackermann, M., Albert, A., Anderson, B., et al. 2015b, *Phys. Rev. Lett.*, 115, 231301
- Actis, M., Agnetta, G., Aharonian, F., et al. 2011, *Experimental Astronomy*, 32, 193
- Aharonian, F., Akhperjanian, A. G., Anton, G., et al. 2009, *ApJ*, 695, L40
- Aharonian, F., Akhperjanian, A. G., Bazer-Bachi, A. R., et al. 2006, *Science*, 314, 1424
- Ahnen, M. L., Ansoldi, S., Antonelli, L. A., et al. 2016a, *A&A*, 595, A98
- Ahnen, M. L., Ansoldi, S., Antonelli, L. A., et al. 2016b, *A&A*, 590, A24
- Ahnen, M. L., Ansoldi, S., Antonelli, L. A., et al. 2017, *A&A*, 603, A25
- Aho, K., Derryberry, D., & Peterson, T. 2014, *Ecology*, 95, 631
- Ajello, M., Albert, A., Anderson, B., et al. 2016, *Phys. Rev. Lett.*, 116, 161101
- Ajello, M., Atwood, W. B., Baldini, L., et al. 2017, *ApJS*, 232, 18
- Akaike, H. 1974, *IEEE Transactions on Automatic Control*, 19, 716
- Akaike, H. 1974, *IEEE Transactions on Automatic Control*, 19, 716
- Aleksić, J., Ansoldi, S., Antonelli, L. A., et al. 2016a, *Astroparticle Physics*, 72, 76

BIBLIOGRAPHY

- Aleksić, J., Ansoldi, S., Antonelli, L. A., et al. 2016b, *Astroparticle Physics*, 72, 61
- Aleksić, J., Ansoldi, S., Antonelli, L. A., et al. 2014, *Science*, 346, 1080
- Aleksić, J., Antonelli, L. A., Antoranz, P., et al. 2014a, *A&A*, 563, A91
- Aleksić, J., Antonelli, L. A., Antoranz, P., et al. 2014b, *A&A*, 563, A91
- Aleksić, J., Antonelli, L. A., Antoranz, P., et al. 2010, *ApJ*, 723, L207
- Alfvén, H. & Herlofson, N. 1950, *Physical Review*, 78, 616
- Anderhub, H., Backes, M., Biland, A., et al. 2013, *Journal of Instrumentation*, 8, P06008
- Angioni, R. 2018, *The Astronomer's Telegram*, 11251
- Antonucci, R. 1993, *Annual Review of Astronomy and Astrophysics*, 31, 473
- Arp, H. C., Bolton, J. G., & Kinman, T. D. 1967, *ApJ*, 147, 840
- Atwood, W., Albert, A., Baldini, L., et al. 2013a, *ArXiv e-prints* [[arXiv]1303.3514]
- Atwood, W. B., Baldini, L., Bregeon, J., et al. 2013b, *ApJ*, 774, 76
- Baade, W. & Minkowski, R. 1954, *ApJ*, 119, 206
- Balázs, C., Conrad, J., Farmer, B., et al. 2017, *Phys. Rev. D*, 96, 083002
- Barlow, R. 2003, in *Statistical Problems in Particle Physics, Astrophysics, and Cosmology*, ed. L. Lyons, R. Mount, & R. Reitmeyer, 250
- Bartels, R., Krishnamurthy, S., & Weniger, C. 2016, *Phys. Rev. Lett.*, 116, 051102
- Bartels, R., Storm, E., Weniger, C., & Calore, F. 2018, *Nature Astronomy*, 2, 819
- Bednarek, W. & Sobczak, T. 2013, *Monthly Notices of the Royal Astronomical Society: Letters*, 435, L14
- Benlloch, S., Rothschild, R. E., Wilms, J., et al. 2001, *A&A*, 371, 858
- Bennett, A. S. 1962, *MmRAS*, 68, 163
- Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, *ApJS*, 208, 20

- Bernstein, R. A. 2007, *ApJ*, 666, 663
- Bernstein, R. A., Freedman, W. L., & Madore, B. F. 2002, *ApJ*, 571, 107
- Biteau, J. & Williams, D. A. 2015, *ApJ*, 812, 60
- Blandford, R., Meier, D., & Readhead, A. 2018, arXiv e-prints, arXiv:1812.06025
- Boehm, C., Fayet, P., & Silk, J. 2004, *Phys. Rev. D*, 69, 101302
- Bolton, J. G. & Stanley, G. J. 1948, *Nature*, 161, 312
- Böttcher, M., Reimer, A., Sweeney, K., & Prakash, A. 2013, *ApJ*, 768, 54
- Brookes, M. H., Lawrence, C. R., Keene, J., et al. 2006, *ApJ*, 646, L41
- Brown, A. M. & Adams, J. 2011, *MNRAS*, 413, 2785
- Brown, A. M., Böhm, C., Graham, J., et al. 2017, *Phys. Rev. D*, 95, 063018
- Burbidge, E. M. 1967, *ApJ*, 147, 845
- Burbidge, G. R. 1959, in *IAU Symposium, Vol. 9, URSI Symp. 1: Paris Symposium on Radio Astronomy*, ed. R. N. Bracewell, 541
- Burbidge, G. R. 1961, *Nature*, 190, 1053
- Burbidge, G. R. & Hoyle, F. 1957, *AJ*, 62, 9
- Cappellari, M., Neumayer, N., Reunanen, J., et al. 2009, *MNRAS*, 394, 660
- Carpenter, B., Ojha, R., D'Ammando, F., Orienti, M., & Cheung, C. C. 2014, *The Astronomer's Telegram*, 5838
- Casella, G. & Berger, R. 2001, *Statistical Inference* (Duxbury Resource Center)
- Chabrier, G. 2003, *PASP*, 115, 763
- Cherenkov, P. A. 1934, *Doklady Akademii Nauk SSSR*, 2, 451+
- Chiaberge, M., Capetti, A., & Celotti, A. 2001, *MNRAS*, 324, L33
- Cortina, J. 2012, *The Astronomer's Telegram*, 4583

BIBLIOGRAPHY

- Crnojević, D., Sand, D. J., Spekkens, K., et al. 2016, *ApJ*, 823, 19
- Curtis, H. D. 1918, *Publications of Lick Observatory*, 13, 9
- Domínguez, A. & Ajello, M. 2015, *The Astrophysical Journal*, 813, L34
- Domínguez, A. et al. 2011, *Monthly Notices of the Royal Astronomical Society*, 410, 2556
- Driver, S. P., Andrews, S. K., Davies, L. J., et al. 2016, *ApJ*, 827, 108
- Edge, D. O., Shakeshaft, J. R., McAdam, W. B., Baldwin, J. E., & Archer, S. 1959, *Memoirs of the Royal Astronomical Society*, 68, 37
- Emmanoulopoulos, D., Papadakis, I. E., McHardy, I. M., et al. 2012, *MNRAS*, 424, 1327
- Fanaroff, B. L. & Riley, J. M. 1974, *Monthly Notices of the Royal Astronomical Society*, 167, 31P
- Fath, E. A. 1909, *Lick Observatory Bulletin*, 5, 71
- Fazio, G. G., Ashby, M. L. N., Barmby, P., et al. 2004, *ApJS*, 154, 39
- Feain, I. J., Cornwell, T. J., Ekers, R. D., et al. 2011, *ApJ*, 740, 17
- Fender, R. P., Belloni, T. M., & Gallo, E. 2004, *Monthly Notices of the Royal Astronomical Society*, 355, 1105
- Fermi-LAT Collaboration, Abdollahi, S., Ackermann, M., et al. 2018a, *Science*, 362, 1031
- Fermi-LAT Collaboration, Abdollahi, S., Ackermann, M., et al. 2018b, *Science*, 362, 1031
- Fermi* Collaboration, Falcone, A., Hase, H., Pagoni, C., & Ploetz, C. 2010, *The Astrophysical Journal*, 719, 1433, arXiv: 1006.5463
- Fermi*-LAT Collaboration. 2008, *The Fermi-LAT Bright Sources List (0FGL)*
- Ferrarese, L. & Merritt, D. 2000, *The Astrophysical Journal Letters*, 539, L9
- Finke, J. D., Razzaque, S., & Dermer, C. D. 2010, *ApJ*, 712, 238
- Fossati, G., Maraschi, L., Celotti, A., Comastri, A., & Ghisellini, G. 1998, *Monthly Notices of the Royal Astronomical Society*, 299, 433

- Franceschini, A., Rodighiero, G., & Vaccari, M. 2008, *A&A*, 487, 837
- Frank, I. M. & Tamm, I. E. 1937, *Compt. Rend. Acad. Sci. URSS*, 14, 109, [*Usp. Fiz. Nauk*93, no.2, 388 (1967)]
- Ghisellini, G., Tavecchio, F., Foschini, L., & Ghirland a, G. 2011, *MNRAS*, 414, 2674
- Gil de Paz, A., Boissier, S., Madore, B. F., et al. 2007, *The Astrophysical Journal Supplement Series*, 173, 185
- Gilmore, R. C., Somerville, R. S., Primack, J. R., & Domínguez, A. 2012, *MNRAS*, 422, 3189
- Gondolo, P. & Silk, J. 1999, *Phys. Rev. Lett.*, 83, 1719
- Goodenough, L. & Hooper, D. 2009, arXiv e-prints, arXiv:0910.2998
- Gopal-Krishna & Wiita, P. J. 2010, *New Astronomy*, 15, 96
- Gordon, C. & Macías, O. 2013, *Phys. Rev. D*, 88, 083521
- Graham, J. A., Brown, A. M., & Chadwick, P. M. 2019, *MNRAS*, 485, 3277
- Greenstein, J. L. 1963, *Nature*, 197, 1041
- Greenstein, J. L. & Schmidt, M. 1967, *ApJ*, 148, L13
- H. E. S. S. Collaboration, Abdalla, H., Abramowski, A., et al. 2018, *A&A*, 619, A71
- Han, J., Frenk, C. S., Eke, V. R., et al. 2012, *MNRAS*, 427, 1651
- Hanbury Brown, R., Jennison, R. C., & Gupta, M. K. D. 1952, *Nature*, 170, 1061
- Hansen, R. C. 1981, *Proceedings of the IEEE*, 69, 170
- Hard Townes, C. 1947, *ApJ*, 105, 235
- Hardcastle, M. J., Kraft, R. P., Sivakoff, G. R., et al. 2007, *The Astrophysical Journal Letters*, 670, L81
- Hardcastle, M. J., Worrall, D. M., Kraft, R. P., et al. 2003, *The Astrophysical Journal*, 593, 169

BIBLIOGRAPHY

- Hardcastle, M. J. et al. 2009, *Monthly Notices of the Royal Astronomical Society*, 393, 1041
- Harding, A. K., Grenier, I. A., & Gonthier, P. L. 2007, *Astrophysics and Space Science*, 309, 221
- Hazard, C., Mackey, M. B., & Shimmins, A. J. 1963, *Nature*, 197, 1037
- Heck, D., Knapp, J., Capdevielle, J. N., Schatz, G., & Thouw, T. 1998, CORSIKA: a Monte Carlo code to simulate extensive air showers.
- Heney, L. G. & Keenan, P. C. 1940, *ApJ*, 91, 625
- H.E.S.S. Collaboration, Abdalla, H., Abramowski, A., et al. 2018, *A&A*, 619, A71
- H.E.S.S. Collaboration, Abramowski, A., Acero, F., et al. 2013, *A&A*, 550, A4
- H.E.S.S. Collaboration, Abramowski, A., Aharonian, F., et al. 2016, *Nature*, 531, 476
- Hinton, J. A. & the HESS Collaboration. 2004, *New Astronomy Reviews*, 48, 331
- Hooper, D. & Mohlabeng, G. 2016, *Journal of Cosmology and Astro-Particle Physics*, 2016, 049
- Horiuchi, S., Meier, D. L., Preston, R. A., & Tingay, S. J. 2006, *Publications of the Astronomical Society of Japan*, 58, 211
- Hoyle, F. & Fowler, W. A. 1963a, *Nature*, 197, 533
- Hoyle, F. & Fowler, W. A. 1963b, *MNRAS*, 125, 169
- Hubble, E. P. 1926, *ApJ*, 64
- Inoue, Y., Tanaka, Y. T., Madejski, G. M., & Domínguez, A. 2014, *ApJ*, 781, L35
- Israel, F. 1998, *The Astronomy and Astrophysics Review*, 8, 237
- James, F. & Roos, M. 1975, *Comput. Phys. Commun.*, 10, 343
- Jansky, K. G. 1932, *Proc. IRE*, 20, 1920
- Jansky, K. G. 1933, *Proc. IRE*, 21, 1387

- Jansky, K. G. 1935, Proc. IRE, 23, 1158
- Jennison, R. C. & Das Gupta, M. K. 1953, Nature, 172, 996
- Kadler, M., Eisenacher, D., Ros, E., et al. 2012, A&A, 538, L1
- Keenan, R. C., Barger, A. J., Cowie, L. L., & Wang, W. H. 2010, ApJ, 723, 40
- Kelly, B. C. 2007, ApJ, 665, 1489
- Kneiske, T. M. & Dole, H. 2010, A&A, 515, A19
- Knödlseder, J., Mayer, M., Deil, C., et al. 2016, A&A, 593, A1
- Kormendy, J. & Richstone, D. 1995, ARA&A, 33, 581
- Kormendy, J. & Richstone, D. 1995, Annual Review of Astronomy and Astrophysics, 33, 581
- Krawczynski, H., Hughes, S. B., Horan, D., et al. 2004, ApJ, 601, 151
- Krolik, J. H. & Begelman, M. C. 1988, ApJ, 329, 702
- Lacroix, T., Boehm, C., & Silk, J. 2015, Phys. Rev. D, 92, 043510
- Lande, J., Ackermann, M., Allafort, A., et al. 2012, ApJ, 756, 5
- Landt, H., Padovani, P., & Giommi, P. 2002, MNRAS, 336, 945
- Lasker, B. M., Sturch, C. R., McLean, B. J., et al. 1990, AJ, 99, 2019
- Lee, S. K., Lisanti, M., Safdi, B. R., Slatyer, T. R., & Xue, W. 2016, Phys. Rev. Lett., 116, 051103
- Lessard, R. 1999, Astroparticle Physics, 11, 243 , teV Astrophysics of Extragalactic Sources
- Lister, M. L., Aller, M. F., Aller, H. D., et al. 2013, AJ, 146, 120
- Lott, B., Escande, L., Larsson, S., & Ballet, J. 2012, A&A, 544, A6
- Lott, B. & Larsson, S. 2013

BIBLIOGRAPHY

- Lynden-Bell, D. 1969, *Nature*, 223, 690
- Madau, P. & Pozzetti, L. 2000, *MNRAS*, 312, L9
- Majaess, D. 2010, *Acta Astronomica*, 60, 121
- Massaro, E., Perri, M., Giommi, P., & Nesci, R. 2004, *A&A*, 413, 489
- Matthews, T. A. & Sandage, A. R. 1963, *ApJ*, 138, 30
- Mattila, K., Väisänen, P., Lehtinen, K., von Appen-Schnur, G., & Leinert, C. 2017, *MNRAS*, 470, 2152
- Mattox, J. R., Bertsch, D. L., Chiang, J., et al. 1996, *ApJ*, 461, 396
- Mazin, D., Domínguez, A., Fallah Ramazani, V., et al. 2017, in *American Institute of Physics Conference Series*, Vol. 1792, 6th International Symposium on High Energy Gamma-Ray Astronomy, 050037
- McKinley, B., Briggs, F., Gaensler, B. M., et al. 2013, *MNRAS*, 436, 1286
- Meisenheimer, K., Tristram, K. R. W., Jaffe, W., et al. 2007, *A&A*, 471, 453
- Michelson, P. F. 1996, in *Proc. SPIE*, Vol. 2806, *Gamma-Ray and Cosmic-Ray Detectors, Techniques, and Missions*, ed. B. D. Ramsey & T. A. Parnell, 31–40
- Moffat, A. F. J. 1969, *A&A*, 3, 455
- Moiseev, A. A., Ajello, M., Buckley, J. H., et al. 2015, *arXiv e-prints*, arXiv:1508.07349
- Morganti, R., Killeen, N. E. B., Ekers, R. D., & Oosterloo, T. A. 1999, *MNRAS*, 307, 750
- Müller, C., Kadler, M., Ojha, R., et al. 2014, *A&A*, 569, A115
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, *ApJ*, 462, 563
- Neronov, A., Semikoz, D., & Vovk, I. 2010, *A&A*, 519, L6
- Neumayer, N. 2010, *Publications of the Astronomical Society of Australia*, 27, 449
- Nolan, P. L., Abdo, A. A., Ackermann, M., et al. 2012, *ApJS*, 199, 31

- Oke, J. B. 1963, *Nature*, 197, 1040
- Petropoulou, M., Lefa, E., Dimitrakoudis, S., & Mastichiadis, A. 2014, *Astronomy & Astrophysics*, 562, A12
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, *A&A*, 571, A1
- Portillo, S. K. N. & Finkbeiner, D. P. 2014, *ApJ*, 796, 54
- Rando, R. & the *Fermi*-LAT Collaboration. 2009, arXiv e-prints, arXiv:0907.0626
- Read, A. M., Rosen, S. R., Saxton, R. D., & Ramirez, J. 2011, *A&A*, 534, A34
- Reber, G. 1944, *ApJ*, 100, 279
- Rector, T. A., Stocke, J. T., & Perlman, E. S. 1999, *ApJ*, 516, 145
- Rieger, F. & Levinson, A. 2018, *Galaxies*, 6, 116
- Robinson, I., Schild, A., & Schucking, E. L., eds. 1965, *Quasi-stellar sources and gravitational collapse*
- Rybicki, G. B. & Lightman, A. P. 1986, *Radiative Processes in Astrophysics*, 400
- Ryle, M. & Smith, F. G. 1948, *Nature*, 162, 462
- Ryle, M., Smith, F. G., & Elsmore, B. 1950, *MNRAS*, 110, 508
- Sahakyan, N., Baghmanyan, V., & Zargaryan, D. 2018, *A&A*, 614, A6
- Sahakyan, N., Yang, R., Aharonian, F. A., & Rieger, F. M. 2013, *ApJ*, 770, L6
- Sahakyan, N., Yang, R., Aharonian, F. A., & Rieger, F. M. 2013, *ApJ*, 770, L6
- Salpeter, E. E. 1964, *ApJ*, 140, 796
- Saviane, I. & Jerjen, H. 2007, *AJ*, 133, 1756
- Scargle, J. D., Norris, J. P., Jackson, B., & Chiang, J. 2013, *The Astrophysical Journal*, 764, 167
- Schmidt, M. 1963, *Nature*, 197, 1040

BIBLIOGRAPHY

- Schwarz, G. 1978, *Ann. Statist.*, 6, 461
- Seyfert, C. K. 1943, *ApJ*, 97, 28
- Shankar, F., Bernardi, M., Sheth, R. K., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 460, 3119
- Sijbring, D. & de Bruyn, A. G. 1998, *A&A*, 331, 901
- Slipher, V. M. 1917, *Lowell Observatory Bulletin*, 3, 59
- Smith, F. G. 1951, *Nature*, 168, 555
- Soltan, A. 1982, *Monthly Notices of the Royal Astronomical Society*, 200, 115
- Somerville, R. S., Gilmore, R. C., Primack, J. R., & Domínguez, A. 2012, *Monthly Notices of the Royal Astronomical Society*, 423, 1992
- Somerville, R. S., Lee, K., Ferguson, H. C., et al. 2004, *ApJ*, 600, L171
- Spitzer, Lyman, J. & Baade, W. 1951, *ApJ*, 113, 413
- Spitzer, Jr., L. & Saslaw, W. C. 1966, *ApJ*, 143, 400
- Stoeckl, J. T., Morris, S. L., Gioia, I. M., et al. 1991, *ApJS*, 76, 813
- Tavecchio, F., Maraschi, L., & Ghisellini, G. 1998, *ApJ*, 509, 608
- The Fermi-LAT Collaboration. 2016, *ArXiv e-prints* [[arXiv]1606.04905]
- The Fermi-LAT collaboration. 2019, *arXiv e-prints*, arXiv:1902.10045
- Tran, H. D. 2001, *The Astrophysical Journal*, 554, L19
- Tsumura, K., Matsumoto, T., Matsuura, S., Sakon, I., & Wada, T. 2013, *PASJ*, 65, 121
- Ulrich, M.-H., Maraschi, L., & Urry, C. M. 1997, *Annual Review of Astronomy and Astrophysics*, 35, 445
- Urry, C. M. & Padovani, P. 1995, *Publications of the Astronomical Society of the Pacific*, 107, 803

- Vaughan, S., Edelson, R., Warwick, R. S., & Uttley, P. 2003, MNRAS, 345, 1271
- Vrieze, S. 2012, Psychological methods, 17 2, 228
- Weekes, T. C., Cawley, M. F., Fegan, D. J., et al. 1989, ApJ, 342, 379
- Wilks, S. S. 1938, Ann. Math. Statist., 9, 60
- Winter, M. et al. 2016, arXiv e-prints [1602.03332]
- Wood, M., Caputo, R., Charles, E., et al. 2017, International Cosmic Ray Conference, 35, 824
- Wu, Y.-Z., Zhang, E.-P., Liang, Y.-C., Zhang, C.-M., & Zhao, Y.-H. 2011, The Astrophysical Journal, 730, 121
- Yang, R.-z., Sahakyan, N., Wilhelmi, E. D. O., Aharonian, F., & Rieger, F. 2012, in American Institute of Physics Conference Series, Vol. 1505, American Institute of Physics Conference Series, ed. F. A. Aharonian, W. Hofmann, & F. M. Rieger, 590–593
- Zech, A., Cerruti, M., & Mazin, D. 2017, A&A, 602, A25
- Zel'dovich, Y. B. & Novikov, I. D. 1964, Soviet Physics Doklady, 9, 246
- Zensus, J. A. 1997, Annual Review of Astronomy and Astrophysics, 35, 607
- Zurek, K. M. 2009, Phys. Rev. D, 79, 115002