Variability of Point Sources of Gamma Rays Detected by the *Fermi* Large-Area Telescope

Eric Wallace

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2016

Reading Committee:

Thompson Burnett, Chair

Miguel Morales

R. Jeffrey Wilkes

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Variability of Point Sources of Gamma Rays
Detected by the *Fermi* Large-Area Telescope

Eric Wallace

Chair of the Supervisory Committee:
Professor Thompson Burnett
Physics

The *Large Area Telescope* (LAT) aboard the *Fermi Gamma-Ray Space Telescope* has revolutionized gamma-ray astronomy, allowing the detection of thousands of point sources of gamma rays. Variability studies are of significant interest as a potential source of information about the emission mechanisms, and as a means to identify gamma-ray sources with known sources in other wavelengths and to improve detection sensitivity in searches for new sources. The inclusion of temporal resolution, however, adds to the already considerable complexity of the required analysis, and as a result, variability studies have generally been limited either in scope or in detail, or both, compared to time-integrated spectral analyses. *pointlike* is a software package designed for fast maximum likelihood analysis of LAT data, allowing for interactive and large-scale analyses. Here, we present an application of *pointlike* to the characterization of the variability of the full sample of known gamma-ray point sources.

We describe the construction of light curves in one-month time bins, spanning the first 42 months of the *Fermi* mission, for a sample of 2652 sources. We discuss the use of the detection significance in individual months to improve the significance of detection of marginal sources, and show that including that measure of significance increases the set of significantly detected sources by nearly 20% compared to using only the average significance. We describe a statistical measure of the significance of variability in a light curve, and exam-

ine the variability of the sample as whole, and of subsets associated with particular source types, especially pulsars. We discuss the use of pulsars, which are generally non-variable on long timescales, to calibrate variability statistics, and to assess the importance of systematic errors in estimates of variability. Finally, we discuss the potential to extend the method to produce light curves of longer duration and finer time binning, and to search for transient sources.

TABLE OF CONTENTS

Page

List of Figures iv

List of Tables vi

Chapter 1: Introduction 1

Chapter 2: Gamma-ray Astronomy and *Fermi* 4
2.1 Principles of γ-ray astronomy 4
2.1.1 Astrophysical production of γ rays 4
2.1.2 γ-ray interactions in matter 6
2.2 Astrophysical γ-ray Sources 8
2.2.1 Active Galactic Nuclei 8
Blazars 11
2.2.2 Pulsars 12
2.3 The *Fermi* Large-Area Telescope 12
2.3.1 Converter-tracker (TKR) 13
2.3.2 Calorimeter (CAL) 15
2.3.3 Anticoincidence Detector (ACD) 16
2.4 Event reconstruction and Instrument Response Functions 17
2.4.1 Reconstruction and classification 17
2.4.2 Instrument Response Functions 17

Chapter 3: LAT data analysis 24
3.1 Source Models 25
3.1.1 Point Sources 25
3.1.2 Extended Sources 29
3.1.3 Diffuse Emission 29

3.2 Instrument Response . . . 30
3.3 Construction of the likelihood function . . . 32
3.3.1 Likelihood Ratio Tests and Detection Significance . . . 34
3.4 *pointlike* . . . 36
3.4.1 Event binning . . . 36
3.4.2 Point Source Localization and Detection . . . 39

Chapter 4: Light Curves . . . 42
4.1 3-year All-sky Analysis . . . 42
4.1.1 Data Selection and Instrument Response Functions . . . 42
4.1.2 Source Detection and localization . . . 44
4.1.3 Regions of Interest . . . 44
4.1.4 Source Models . . . 44
Point and Extended Sources . . . 44
Diffuse Sources . . . 45
4.1.5 Summary of All-sky Fitting Procedure . . . 46
4.1.6 Population Statistics . . . 46
4.2 Light Curve Analysis . . . 46
4.2.1 Data Selection and Instrument Response Functions . . . 47
4.2.2 Time Binning . . . 47
4.2.3 Source Models . . . 50
Diffuse Models . . . 50
Point and Extended Sources . . . 52
4.2.4 Single-source Profile Likelihood . . . 52
4.2.5 Light Curve Results . . . 54
Example light curves . . . 55
Accumulated significance . . . 58
4.2.6 Light Curve Correlations . . . 61
4.2.7 Summary . . . 61

Chapter 5: Variability Analysis . . . 63
5.1 Quantifying Variability . . . 63
5.1.1 Flux-dependence . . . 64

5.1.2 Pulsars 66

Chapter 6: Conclusions 71

Bibliography 75

LIST OF FIGURES

Figure Number — Page

2.1 Scattering cross section for light in carbon and lead . . . 9
2.2 LAT effective area as a function of energy . . . 20
2.3 LAT effective area as a function of incidence angle . . . 20
2.4 LAT PSF as a function of energy . . . 21
2.5 LAT energy dispersion as a function of energy . . . 22
2.6 LAT energy dispersion as a function of incidence angle . . . 23

3.1 Example power-law spectrum . . . 26
3.2 Example exponential cutoff spectrum . . . 27
3.3 Example log-parabolic spectrum . . . 28
3.4 Example TS maps for a well-localized source (left) and a poorly-localized source (right). The black solid contours are at TS=1.51,2.45, and 3.01, corresponding to 68%, 95%, and 99% containment for a gaussian likelihood surface. The green dashed contours show the corresponding contours from the elliptical fit. . . . 40

4.1 TS distribution for 3-year source list . . . 47
4.2 Point source flux distribution for 3-year source list . . . 48
4.3 3C 279 light curve . . . 55
4.4 Markarian 501 light curve . . . 56
4.5 PSR J0030+4510 light curve . . . 57
4.6 2FGL J0858.2-3129 light curve . . . 59
4.7 Point source flux distribution for 3-year source list . . . 60
4.8 Cross-correlation of flux estimates as a function of source separation. . . . 62

5.1 Variability index distribution . . . 64
5.2 Variability index vs. mean flux . . . 65
5.3 Crab (nebula+pulsar) light curve . . . 67
5.4 PSR J2021+4026 light curve . . . 68

5.5 Variability index distribution 70
5.6 Light curves for the Vela and Geminga pulsars 70

LIST OF TABLES

Table Number Page

3.1 Event binning for 3-year all-sky analysis 38

4.1 Excluded time intervals 43

4.2 Monthly light curve time bins 51

ACKNOWLEDGMENTS

The list of people to whom I owe some gratitude is, as always, too long to enumerate, but particular thanks to:

My advisor, Toby Burnett, from whom I've learned a good deal about experimental physics, and who has offered years' worth of advice, encouragement, and patience.

The rest of my thesis committee, Jeff Wilkes, Miguel Morales, Ann Nelson, and Scott Anderson, who have been as helpful and encouraging a committee as I could have asked for.

The Fermi-LAT Collaboration, which is impressively full of excellent scientists who are also great fun to work with.

All of the great friends I've made over the course of graduate school. In particular, my officemates and fellow *Fermi* students, Matthew Kerr and Marshall Roth; and Laura Bodine and Will Terrano, who provided a great deal of feedback and discussion on my dissertation, along with countless interesting discussions about all things physics-related and most things not.

My wife, Zoë, for her constant love and support through all the ups and downs of gruaduate school and life in general.

And last, but certainly not least, my parents, Vernon and Meredith, who inspired my curiosity about the world, and provided the support, of all forms, that allowed me to pursue it.

DEDICATION

For my parents.

Chapter 1

INTRODUCTION

Astrophysical sources of γ rays include the most extreme environments and phenomena in the known universe and provide cosmic laboratories for the study of fundamental physics at high energies and densities. Many of these sources are highly variable, exhibiting large changes in flux on timescales ranging from seconds to years. Measuring and understanding this variability is essential to understanding the physical processes at work in these environments, and provides important tools for the detection and characterization of new sources. Variability analyses also present significant challenges, beyond those common to other analyses of γ-ray data. This work presents an analysis of the variability of a large sample of γ-ray sources detected by the *Fermi* Large Area Telescope (LAT) in the first three years of its mission.

Since its launch in 2008, the Large Area Telescope aboard *Fermi* Gamma-ray Space Telescope has revolutionized the field of γ-ray astronomy, detecting a wealth of new sources of γ rays, and allowing the study of previously-known sources in unprecedented detail. In its primary survey mode, the LAT observes the full sky approximately once every three hours. This exceptional coverage has allowed the construction of the deepest catalogs to date of γ-ray sources in the range from 100 MeV to greater than 1 TeV in photon energy. Furthermore, it allows near-continuous monitoring of the γ-ray flux across the full sky. The second *Fermi* LAT source catalog (2FGL), constructed based on data from the first two years of the mission, contains 1873 sources. More than half are associated with active galactic nuclei (AGN), mostly blazars.

In Chapter 2, we review the principles of γ-ray astronomy. We describe the primary mechanisms by which γ rays are produced in astrophysical environments, and briefly describe

a few of the important classes of astrophysical γ-ray sources. We discuss the nature of γ-ray interactions with matter, and explore how those interactions inform and constrain the design of γ-ray detectors.

Finally, we discuss the design of the *Fermi* Large Area Telescope in particular, with some comparison to earlier γ-ray telescopes, particularly the *EGRET* instrument on the Compton Gamma-Ray Observatory. We describe the main instrument subsystems, and the processing of detector-level data to extract photon data. We then examine the calibration and performance of the instrument, and present the standard parameterizations of the instrument response.

Chapter 3 looks in some detail at the maximum likelihood analysis used for most analyses of LAT data. We discuss the construction of the likelihood function and statistics derived from it, and their use in extracting observables of interest from the data. We then present the *pointlike* package and its particular implementation of the likelihood analysis. In particular, we emphasize the particular aspects of the implementation designed to simplify the likelihood calculation, improving the tractability of maximum-likelihood analyses. We compare the performance of *pointlike* to the more rigorous implementation provided by the *Fermi* collaboration, and argue that for a variety of applications, it provides significantly improved computational performance without significant impact on analysis results.

Chapter 4 presents the analysis that forms the heart of this work, a study of the temporal variability of all sources detectable in three years of Large-Area Telescope data. In contrast to most previous studies of variability of LAT-detected γ-ray sources, this analysis aims to produce variability statistics for the full set of known sources, rather than a subset of relatively bright objects. This breadth comes at the cost of depth, and the analyses of the behaviour of individual sources are not as detailed as those produced in (2LAC, etc); rather they pursue a complementary approach of attempting to produce a statistical characterization of the variability properties of the entire population.

The starting point for the light curve analysis is the construction of a source list based on three years of data. We describe the detection procedure, and the analysis of the average

properties of the source sample. We present a summary of the resulting source list, and compare the results of the analysis to those of the 2FGL catalog. This analysis results in a list of sources roughly 20% larger than the 2FGL, when selected by similar criteria. The properties of the sample are generally in good agreement with expectations based on the 2FGL.

The light curve analysis consists of the a simplfied version of the 3-year analysis, applied to 42 one-month intervals, producing estimates of the flux as a function of time for each of the sources in the 3-year list. We describe the simplifying assumptions made to render the computation tractable, and discuss their potential impact on the results of the analysis. We present a method to improve the significance of marginal detections in the time-averaged analysis using the light curve information, and argue that incorporating this additional significance estimate in our selection criteria expands our sample of confidently-detected sources by approximately 10%. Finally, we present some example light curves, and summarize the properties of the set.

In Chapter 5, we discuss the assessment of the significance of observed variability in a light curve. We present an analysis of the variability of the set of light curves constructed in Chapter 4. We describe the use of pulsars as calibration sources to assess the reliability of our measure of variability and the impact of systematic errors. We explore the dependence of estimated variability on source flux, and discuss the implications for the study of variability in faint sources.

Chapter 2

GAMMA-RAY ASTRONOMY AND *Fermi*

Gamma-ray astronomy is the study of extraterrestrial light in the highest-energy band of the electromagnetic spectrum. Since its inception in the early 1960s, the field has seen successive advances as improvements in detector technology have allowed the design of more sensitive and accurate γ-ray telescopes. In this chapter we discuss the physical considerations that drive the design of γ-ray telescopes and provide a brief history of past γ-ray observatories. We then describe the design and operation of the *Fermi* Large-Area Telescope, with some discussion of particular advances with respect to previous γ-ray observatories.

2.1 Principles of γ-ray astronomy

The goals and tools of astronomy in any waveband are driven by the interactions of light in that band with matter. Those interactions determine what sorts of astrophysical bodies might produce said light and what sorts of telescopes and detectors might be used to observe them. Accordingly, we shall give a brief overview of the interactions by which γ rays are produced and detected. We shall then consider the implications of their interaction properties for the design of γ-ray telescopes.

2.1.1 Astrophysical production of γ rays

Photon energies in the γ-ray band are too high to be produced in any significant quantity by persistent sources. Non-thermal γ rays can be produced by electromagnetic interactions of charged particles or by certain particle decays. The only decay process that contributes significantly to the flux of astrophysical γ rays is the decay of the π^0 to two photons, which is the dominant decay for neutral pions produced in interactions of cosmic-ray protons with the

gas and dust of the Milky Way. The spectrum for this decay peaks near half the pion mass, about 67.5 MeV, and falls off with energy approximately as a power-law at higher energies. There is also the interesting possibility of particle dark matter annihilating to photons, but all that can be said with certainty about that scenario is that any resulting γ-ray flux is much too small to be worth considering outside the context of dedicated searches.

Charged particles produce γ rays primarily by three processes: in gas and dust, primarily by bremsstrahlung; in regions with magnetic fields, by synchrotron radiation; and in radiation fields, by inverse Compton scattering of soft photons. In most regions of interest, some combination of these will be present. The spectra of bremsstrahlung and synchrotron radiation are directly related to the spectra of the population of charged particles producing them. The inverse Compton spectrum is primarily determined by the spectrum of the target radiation field. In many cases of interest, the primary target radiation field is synchrotron radiation produced by the same population of particles, a situation referred to as "synchrotron self-Compton"; in this case, the spectrum is again tied to that of the charged particle population. In particular, in all three cases, if the charged-particle energies follow a power-law distribution, the resulting radiation spectrum is another power law, generally with a somewhat steeper slope.

The most familiar non-thermal population of charged particles are cosmic rays, which are, in fact, well-described by a power-law distribution, albeit with at least two well-defined regions of different power-law indices. The low-energy portion of the population (below the "knee" of the spectrum) is thought to originate primarily from supernova remnants within our galaxy. When a supernova explodes into a surrounding medium, the ambient matter is swept up in the blast wave, relaxing on a characteristic time scale into an adiabatically expanding shell of plasma. Charged particles near the shock front can be reflected back and forth across it many times, gaining a little energy in each crossing, in the process known as Fermi acceleration. The resulting energy spectrum is a power law, as seen in the cosmic-ray spectrum (Blandford & Ostriker, 1978). Similar shock acceleration processes likely also account for a significant fraction of particle acceleration in extragalactic objects such as

blazars. Because of the prevalance of shock acceleration as a source of charged particles and the nature of the interactions by which they produce γ rays, power-law spectra are ubiquitous in the γ-ray sky. They are by no means the only type of spectra observed, and they are often modified by absorption, cooling, and other processes. But they are a well-motivated default assumption, and are quite often an accurate fit to within the experimental uncertainty.

2.1.2 γ-ray interactions in matter

The detection of γ rays, and thereby the designs of γ-ray telescopes, are governed by their interactions in matter. The most basic consideration is the opacity of the atmosphere. Very fortunately from a health standpoint, but quite unfortunately from a telescope design standpoint, the atmosphere is opaque to light of frequencies above the high end of the ultra-violet, including the entire γ-ray band. Direct observations of γ rays therefore necessitate balloon-borne or space-based observations. In the TeV γ-ray range, γ rays interacting in the upper atmosphere can produce electromagnetic showers that can be observed by ground-based optical telescope arrays. Here, we are concerned mainly with the GeV range, and with sattelite-based telescopes.

Figure 2.1 shows the cross-sections for interactions of light in carbon and in lead. In both cases, the dominant interaction at low energies is the photo-electric effect, giving way to Compton scattering in the low-energy γ-ray régime. Above about 1 MeV (twice the electron rest-mass), it is possible for for a photon to interact with the Coulomb field of an atomic nucleus and produce an electron-positron pair. The cross-section for pair production increases quickly as a function of energy, and for energies $E \gtrsim 10\,\mathrm{MeV}$ this is the dominant process by which γ rays interact with matter. Accordingly, GeV γ-ray telescopes are, as a rule, pair-production telescopes[1] . The advantages of high-Z elements as target material are clear from the the cross sections. The total cross section in the high-energy γ-ray is at least an order of magnitude greater for lead than for carbon. Furthermore, the pair-production

[1]Here and elsewhere, I use "GeV γ ray" to indicate photon energies in the range from a few tens of MeV to a few hundred GeV.

cross section rises much more quickly with energy in the heavier elements than in the lighter ones, strongly dominating the total cross section above ∼10 MeV. By 100 MeV, no more than a few percent of photons interact by Compton scattering in lead. In carbon, by contrast, Compton scattering remains important at least up to 1 GeV.

As noted in the previous section, the dominant interaction of high-energy electrons in matter is bremsstrahlung on the nuclear Coulomb fields. The rate of energy loss for photons due to pair-production and for electrons due to bremsstrahlung are both approximately linear in the distance traversed in the material. The characteristic length scales for the two interactions are closely related. For electrons, the e-folding distance for energy loss to bremsstrahlung in a given material is the *radiation length*, X_0. That is,

$$|\frac{dE}{dx}|_{brems} \approx \frac{E}{X_0}. \tag{2.1}$$

For γ rays, the mean free path for pair-production is $\frac{7}{9}X_0$.

The combination of these two processes leads to electromagnetic showers. A shower begins when an incident high-energy photon produces an electron-positron pair. The electron and positron each radiate high-energy bremsstrahlung photons, which can in turn produce additional pairs, and so on. The ensuing cascade develops exponentially, continuing as long as the product particles are sufficiently energetic. For photons, the limiting energy is the pair-production threshold. For electrons, it is the critical energy, below which ionization losses begin to dominate over bremsstrahlung; this critical energy, where energy loss from ionization is equal to that from bremmstrahlung, is again energy dependent, and occurs around 10 MeV for heavy elements. The shower profile is peaked around the incident trajectory of the initial photon, so that reconstruction of the profile allows for estimation of the incidence direction.

Though the primary source of energy loss for the electrons is bremsstrahlung, produced in close approaches to the nuclear Coulomb fields, the electrons are also affected by the fields of more distant nuclei, which perturb their paths in a process called multiple Coulomb scattering. The small deviations produced by these interactions are well-described by a

random walk, and for sufficient thickness of traversed material, the angle between the initial and perturbed trajectories is distributed approximately as a Gaussian, with a width

$$\theta_0 \approx 8^\circ \sqrt{\frac{x}{X0}} \frac{100MeV}{E} \tag{2.2}$$

in the relativistic limit. This multiple scattering serves to broaden the distribution of reconstructed photon directions relative to the true incident direction (the point-spread function, or PSF).

2.2 Astrophysical γ-ray Sources

2.2.1 Active Galactic Nuclei

By far the most common persistent point sources of γ rays, comprising at least half the sample of point sources seen by the LAT, are active galactic nuclei (AGN). These are accreting supermassive black holes at the centers of galaxies, which efficiently convert gravitational potential energy into radiation over the entire electromagnetic spectrum through the process of accretion. They display enormous luminosities, often outshining the integrated emission from their host galaxies by orders of magnitude. Since the discovery of Seyfert galaxies in the early 20th century, and of quasars in the 1960s, several types of AGN were discovered, displaying a broad range of luminosities and spectral types, leading eventually to unification schemes explaining the diversity of observed properties in terms of a now generally accepted model for the structure of such nuclei, with various observational differences arising from a combination of intrinsic differences and orientation effects.

The central engine consists of a supermassive black hole surrounded by an accretion disk. Thermal emission from the disk ionizes matter in the clouds of matter in the surrounding broad-line and narrow-line regions (BLR and NLR), producing optical spectra with strong emission lines. In the BLR, the clouds are virialized in the gravitational potential of the black hole and move at high velocities, producing emission lines that are strongly doppler broadened. The slower moving NLR clouds produce narrower lines. A toroidal region of warm

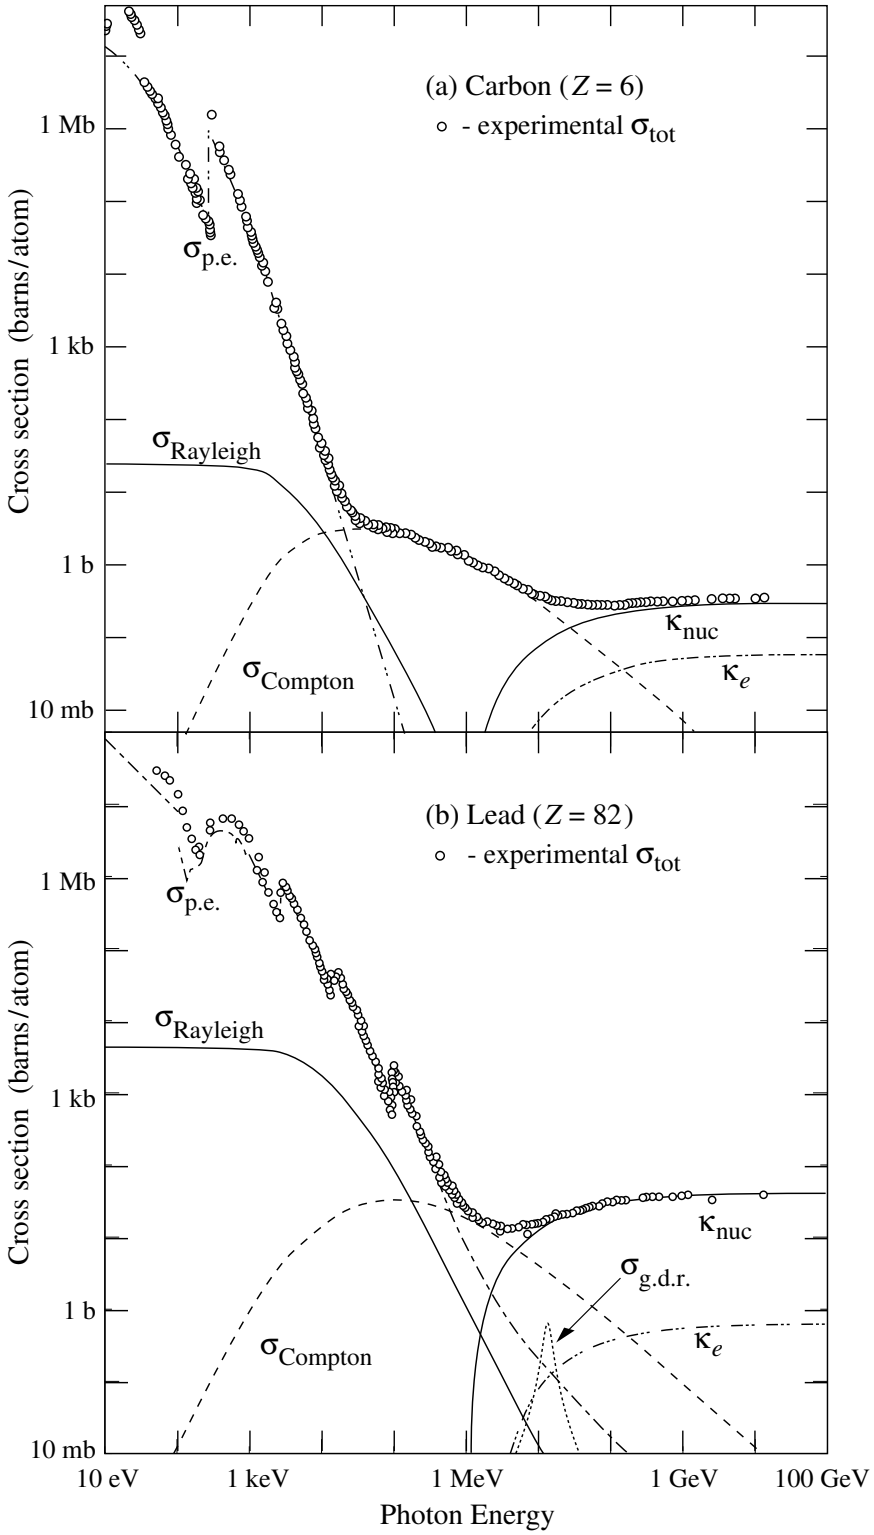

Figure 2.1 Total cross section for scattering of a photon in carbon and in lead. Compton scattering remains important to higher energies in carbon, but it is clear that pair production dominates in both cases at energies much above 100 MeV. The pair-production cross section remains approximately flat up to at least $\sim$ 10 TeV. Reproduced from Olive (2014)

dust outside the BLR obscures radiation from BLR and central engine in obects viewed edge-on, and reprocesses disk and BLR radiation into infrared-peaked thermal emission. Finally, outflow of charged particles forms collimated jets of relativistically moving plasma roughly along the axis of the accretion disk.

The axisymmetric structure of this model leads to anisotropic emission, with the observed properties of the source varying depending on the orientation of a particular AGN with respect to the observer's line of sight, in addition to differences in the intrinsic properties of the source. This leads to various taxonomies of AGN based on the observed emission. Most AGN can be grouped into two broad groups depending on the width of observed emission lines in their optical spectra, with type 1 AGN showing both broad and narrow emission lines and bright continuum emission in their optical spectra and type 2 AGN showing weak continuum emission and only narrow emission lines; explained as an orientation effect, type 1 AGN are taken to be those in which the line of sight is relatively close to the axis, so that the emission from the disk and the BLR is visible, while type 2 AGN are those viewed close to edge-on, such that the emission from the BLR and the central engine is obscured by the dusty torus. Further distinctions can be made within these groups based on bolometric luminosity and radio loudness, i.e. the ratio of radio to optical flux. Some small fraction ($\sim 1\%$) of AGN are viewed at very nearly on-axis, so that the relativistic jet points very close to the line of sight. In this case, the jet emission is relativistically beamed toward the observer, leading to large enhancements in the luminosity of the jet emission proportional to δ^4 relative to the rest frame luminosity, where δ is the doppler beaming factor $\delta = [\Gamma(1 - \beta cos(\theta))]^{-1}$, β the bulk velocity of the emitting plasma, and $\Gamma = (1 - \beta^2)^{-\frac{1}{2}}$. This class of AGN viewed nearly on-axis are collectively known as blazars; their emission is dominated by the beamed jet emission, leading to their defining observational characteristics of high variability in all wavebands, broadband spectra dominated by the non-thermal continuum jet emission, and high polarization. Blazars constitute the majority of the AGN observed in γ rays.

Blazars

As with the other classes of AGN, blazars can be further classified according to their specific observational properties. Perhaps the most common such division is based on optical spectra, by which they are divided into two sub-types, called BL Lacertae objects (BL Lacs) and flat spectrum radio quasars (FSRQs). BL Lacs are characterized by nearly featureless optical spectra, with only weakly detectable emission lines, or none at all; the spectra are strongly dominated by the non-thermal continuum emission from the jet. A typical definition of "weak" in this context, and the one we will adopt for this work where relevant is a line equivalent width EW < 5 Å.

The spectral energy distributions of blazars are generally dominated by two broad peaks (in the usual νF_ν representation), one in the IR-UV range and one in the γ ray band. The low-energy peak is thought to be produced by synchrotron emission of charged particles accelerating in the magnetic field of the jet, while the γ ray peak is produced by the inverse Compton scattering of the same population of charged particles on ambient photon fields. The seed photons for this IC scattering can be provided by the synchrotron radiation responsible for the low-frequency peak (the synchrotron self-Compton or SSC model), or by external photon fields (the external Compton or EC model) including the thermal emission from the accretion disk and the BLR, infrared emission from the dusty torus, and microwave photons from the CMB. In FSRQs, an optical-UV component from the accretion disk is sometimes visible over the falling edge of the synchrotron peak.

The two-peaked structure of blazar SEDs leads to a second classification, based on the location of the synchrotron peak. Blazars are classed as low synchrotron peaked (LSP) if their synchrotron luminosity peaks at $\nu_{peak} \lesssim 10^{14} Hz$, intermediate synchrotron peaked (ISP) for $10^{14} \lesssim \nu_{peak} \lesssim 10^{15} Hz$, or high synchrotron peaked (HSP) for $\nu \gtrsim 10^{15}$. In this scheme, nearly all FSRQs are LSPs, while the BL Lacs run the full range from LSP to HSP. The locations of the two peaks are generally correlated, with the IC peaks typically falling in the low end of the GeV band in LSPs and reaching the TeV range in some HSPs.

2.2.2 Pulsars

The second largest class of point sources of γ rays, and the largest class of galactic sources, is the pulsars. These are rapidly rotating neutron stars (rotation periods ranging from tens of milliseconds to a few seconds) with extremely powerful magnetic fields (ranging from about $10^{10} - 10^{14}$ Gauss). The electric fields induced by the rotation of these magnetic fields accelerate charged particles to very high energies. That population of charged particles interact with the magnetic field, producing pair cascades and curvature radiation, which are responsible for the bulk of the emission from pulsars.

The radiation so produced tends to be concentrated at relatively small angles to the axis of the magnetic fields. The misalignment of the magnetic dipole axis with respect to the rotation axis produces a lighthouse effect, with the beam of radiation passing our line of sight once per rotation period, and producing the characteristic pulsations.

The spectrum of emission from pulsars is cut off at energies above a few GeV, and their γ-ray emission is limited to lower energies. They can, however, be extremely luminous at those low energies, and the brightest of them are the brightest persistent sources of γ rays in the sky. Averaged over many rotation periods, the observed flux from pulsars tends to be very steady. Their distinctive spectral features, brightness, and steadiness, along with the ability to firmly identify them by their pulsations, makes pulsars quite useful as calibration sources in many contexts, in addition to their value as objects of study in their own right.

2.3 The Fermi Large-Area Telescope

The *Fermi* Gamma-ray Space Telescope was launched 11 June 2008, on a mission to provide the best characterization yet of the GeV γ-ray sky. It orbits in a nearly circular orbit at an altitude of $\approx$565 km and an orbital inclination of $\approx 26°$. The spacecraft typically operates in a sky survey mode, in which the boresight rocks north and south of the zenith direction by an angle $\pm$D on alternate orbits. For the first year of the mission, the nominal rocking angle was set at $\mathrm{D} = 35°$; it was increased to $\mathrm{D} = 39°$ for much of July and August 2009,

and then to D = 50° on 2 September 2009, where it has remained since. The spacecraft can also be inertially pointed to increase exposure toward particularly interesting targets, and for periods of pointed observation the viewing angle with respect to the zenith can exceed the nominal rocking angle; in practice, such periods are rare.

The primary instrument onboard the *Fermi* observatory is the Large Area Telescope (LAT). It is a large field of view ($> 2\,\mathrm{sr}$), imaging γ-ray telescope providing unprecedented sensitivity to and resolution of γ rays in the 20 MeV to 300 GeV energy range. The secondary instrument, the Gamma-ray Burst Monitor (GBM) is a scintillator-based detector sensitive to X rays and γ rays in the 8 keV to 40 MeV energy range, designed to provide sensitivity to transient sources, especially γ-ray bursts. In this work, we will be concerned solely with the LAT, which is described in more detail in the next section.

The design of the LAT is modular, consisting of three primary subsystems: a converter-tracker (TKR), a calorimeter (CAL), and an anticoincidence detector (ACD). These basic components are common to essentially all pair-production telescopes. In the following sections, we outline the physics considerations driving the design of each subsystem, and describe the implementation of each in the LAT, with some comparison to earlier telescopes, particularly EGRET.

2.3.1 Converter-tracker (TKR)

The TKR serves a dual purpose. First, it provides a target for the initial conversion of an energetic photon to an electron-positron pair. Second, it provides the position-sensitive detectors which track the trajectories of the resulting charged particles. The characteristics of the interactions discussed in § 2.1.2 place fairly strict constraints on the design of the detector. It must contain sufficient radiation lengths of converter material to ensure conversion of nearly all incident γ rays. The converter material should be a high-Z element, to ensure that the electron critical energy is well below the energy range of interest, that the photon cross section is relatively flat and dominated by pair production, and that incident photons are much more likely to interact in the converter material than in other parts of the detector. In

order to mitigate the effects of multiple scattering, detection of the initial electron-positron pair should take place as close as possible to the conversion point, minimizing the amount of material the charged particles must traverse before they are detected. This leads naturally to a design in which layers of thin, high-Z converter foils are interleaved with active detector layers. Very thin converter layers, while reducing the effects of multiple scattering and the chances of the initial electrons losing much of their energy to bremsstrahlung before detection, naturally reduce the efficiency of conversion; this can be balanced by using a large number of layers to provide sufficient radiation lengths.

Most previous γ-ray telescopes have used spark chambers for the active detector component. The most advanced of the previous spark-chamber detectors was *EGRET*, in which the tracker-converter used 28 layers of spark chambers interleaved with tantalum converter foils, providing a high-energy angular resolution of .4°. In the LAT, spark chambers are replaced with solid-state detectors, which carry a few significant advantages over spark chambers. Whereas the gas in a spark chamber degrades with use and needs to be flushed and replenished occasionally, solid-state detectors require no consumables for their operation. In addition, the finer pitch achievable by solid-state detectors allows for improved angular resolution at high energies.

The LAT TKR is composed of a 16 identical modules, in a 4x4 array, each containing 18 tracker planes, interleaved with tungsten converter foils. Each tracker plane consists of two orthogonally-oriented layers of silicon-strip detectors (SSDs), with 228 μm strip spacing, allowing for a two-dimensional (x-y) measurement of the position of a charged particle. Directly above each tracker plane is a layer of tungsten foil, which provides the target for pair-conversion of incident γ rays. In order to balance the requirements of fine angular resolution and large effective area for pair production, the TKR is divided into two sections, called the *front* and *back*, with different thicknesses of converter foil. The *front* section consists of the 12 layers at the top of the detector, and has foils of thickness 0.028 X_0; the *back* section consists of the next four layers, with foils of thickness 0.18 X_0. The final two layers consist of tracker layers with no converter foil, allowing the products of conversion in

the last foil to satisfy the three-in-a-row trigger criterion required by the track reconstruction algorithm.

The thin foils in the *front* section provide better low-energy angular resolution than the thick foils of the *back*, due to reduced effects of multiple scattering, at the expense of comparatively lower probability for conversion in each layer. The thick foils of the *back* section offer improved effective area, especially important at high energies, with a concomitant loss of angular resolution (a factor of ≈ 2 at 1 GeV). The total depth of each section in radiation lengths is chosen to give roughly equal total numbers of converted photons in the two sections.

2.3.2 Calorimeter (CAL)

The primary function of the calorimeter (CAL) is to provide a measurement of the total energy of the reconstructed γ rays. It must therefore provide a sufficient depth of target material to completely stop the shower particles across the energy range of the instrument, and means of measuring the energy loss in the stopping process. Generally, this is acheived using some variety of scintillator, in which the amount of light produced by the passage of a charged particle is proportional to the energy loss; the light yield in the calorimeter can then be used to reconstruct the total energy of the shower.

The LAT calorimeter consists of a hodoscopic array of 2.7 cm x 2.0 cm x 32.6 cm CsI(Tl) scintillator crystals, with an overall depth of 8.6 radiation lengths. Like the TKR, the CAL is modular, with 16 modules corresponding to the 16 modules of the TKR. Each module contains 96 crystals, arranged in 8 layers of 12 crystals each; the orientation of the long dimension of the crystals alternates from layer to layer. The crystals are optically isolated, with PIN photodiode readouts at each end. The segmentation provides a physical location for the energy deposited in each crystal, with two coordinates provided by the location in the array, and the third determined by the asymmetry in the light yield at the ends of the crystal along its long dimension.

The segmentation of the CAL allows a reconstruction of the spatial profile of a shower,

and carries several advantages over a monolithic calorimeter such as that used in EGRET. Measurement of the longitudinal shower profile allows for effective estimates of the initial electron energy even at energies up to $1TeV$ where energy leakage is becomes significant. Furthermore, measurement of the shower profile provides a powerful tool for rejecting hadronic backgrounds, since the mean profiles of hadronic showers differ significantly from leptonic ones. Finally, the spatial distribution of energy deposition in the CAL provides an independent estimate of the incident particle direction, which, combined with the estimate from the TKR, provides improved reconstruction efficiency and quality.

2.3.3 Anticoincidence Detector (ACD)

The primary function of the anticoincidence detector (ACD) is to discriminate charged-particle backgrounds from the desired γ ray signal. Since the flux of charged cosmic rays can exceed that of celestial γ rays by as much as a factor of 10^5, very efficient rejection of these backgrounds is necessary to extract useful γ ray data. As noted in § 2.3.2, the profiles of hadronic showers differ significantly from those initiated by γ rays, allowing hadronic backgrounds to be rejected by the TKR and CAL. Cosmic ray electrons, in contrast, produce showers very similar in shape to those initiated by γ rays and cannot be rejected with the required efficiency using only the TKR and CAL. The ACD solves this problem by detecting charged particles incident on the LAT, thereby providing a signal that can be used to veto events produced by charged cosmic rays.

The LAT ACD consists of 89 tiles of plastic scintillator, each read out by two photomultiplier tubes for redundancy. The tiles are overlapped in one dimension, and gaps in the other dimension are covered with high-efficiency scintillator ribbons. The overall detection efficiency for singly-charged particles in the ACD is $> 99.97\%$.

Earlier γ ray telescopes generally used monolithic ACDs, which were prone to significant self-veto effects from the so-called backsplash effect. This occurs when interactions in the calorimeter produce large quantities of isotropically distributed secondary particles, mostly soft X rays. Some fraction of these X rays can Compton scatter in the ACD, producing a

spurious veto and causing an otherwise acceptable γ ray event to be rejected. In EGRET, this effect reduced the detection efficiency by a factor of two at 10 GeV compared to 1 GeV, and to nearly zero above 50 GeV. The segmentation of the LAT ACD mitigates this effect by allowing potential ACD vetoes to be restricted to triggers in the ACD tile nearest to the path of the candidate photon. Furthermore, onboard rejection of events via the ACD veto signal is turned off above 20 GeV, allowing for ground-based analysis by more complex software; this has the beneficial side-effect of preserving the high-energy cosmic-ray electron data for science analysis.

2.4 Event reconstruction and Instrument Response Functions

2.4.1 Reconstruction and classification

Signals in the detector subsystems are subjected to a series of onboard trigger criteria which, when met, trigger a readout of the detector state that is sent to ground as an event for more detailed processing. The ground data processing pipeline fits the low-level detector data, including the "hits" in the tracker, the calorimeter energy-deposition profile, and ACD signals to produce estimates of the incident photon direction and energy. The triggers and data processing are described in detail in Atwood et al. (2009). Fiducial cuts on the various low-level fit quantities define several hierarchical classes of events, varying in purity and efficiency of reconstruction, suitable for different types of analysis.

2.4.2 Instrument Response Functions

The efficiency and quality of the reconstruction is described by a set of instrument response functions (IRFs), designed for use in maximum likelihood analyses of the sort that will be described in the next chapter. The instrument response is assumed to be factorable into three parts: the effective area,$A_{\rm eff}(E,\hat{v})$, describing the probability to detect a photon with energy E and incident direction $\hat{v}$; the point-spread function, $P(\hat{v}';E,\hat{v})$, describing the distribution of reconstructed photon incidence angle, $\hat{v}'$, for photons of energy E and true

incident direction $\hat{v}$; and the energy dispersion, $D(E';E,\hat{v})$, describing the distribution of reconstructed energies, E', for photons with true energy and incidence angle E, $\hat{v}$. In all the preceding definitions, $\hat{v}$ and $\hat{v}'$ are unit vectors describing photon directions in the reference frame of the LAT.

A set of IRFs describe the efficiency and accuracy of a particular reconstruction algorithm and event class. Major versions of the reconstruction software are labeled as "Pass N" (e.g., Pass 6, Pass 7), each of which have several minor versions. An event class refers to a particular set of selection cuts (SOURCE, CLEAN, etc) and the conversion region in the TKR (*front* or *back*). All analyses in this work use the Pass 7, version 6 SOURCE class: P7SOURCE_V6.

The specific forms of the IRFs are developed from extensive Monte Carlo simulations. In addition, the PSF and effective area have been validated and modified based on on-orbit data; there are no known astrophysical sources of monoenergetic γ rays to use as calibration sources for the energy dispersion. The factorability of the instrument response rests on the assumption that any correlations between the reconstructed photon energy and direction are negligible. Such correlations have also been evaluated with Monte Carlo simulations, and found to be neglibigle for source fluxes averaged over several orbital periods. These validation studies for the Pass 7 event analysis are described in detail in Ackermann et al. (2012).

In the above description, we have presented the various IRFs as a function of a general incident direction, $\hat{v}$. While they are all, in general, functions of both the angle with respect to the boresight, θ, and the azimuthal angle, ϕ, the azimuthal dependence ϕ is generally much weaker, and tends to average out over longer exposures. The errors introduced by ignoring the azimuthal dependence and using IRFs averaged over ϕ are typically small compared with other sources of uncertainty, and that is the approach we take in the analyses described here. We express the resulting IRFs as functions only of E and cos(θ).

The effective area, $A_{\rm eff}(E, cos(\theta))$, is the product of the geometric cross section of the LAT and the probabilities for a photon to pair produce in the TKR, for the resulting shower

to trigger a readout, and for the reconstructed event to pass the relevant selection cuts. The probability to trigger an event readout is dependent on the instrumental deadtime resulting from the instrument readout time of $\sim 25\mu s$. The deadtime depends on the trigger rate, which varies with the orbital environment of the spacecraft due to changes in the charged particle background rate. The livetime fraction can vary from $\sim 80\%$ to $\sim 92\%$, but typically averages $\sim 90\%$. The reconstruction efficiency is dependent on both incidence angle and energy, with the quality of reconstruction generally being better for higher energies and closer to normal incidence. It is also affected by the presence of residual signals from charged particles passing through the detector a short time before a readout ("ghost tracks"), which result in an energy-dependent reduction in effective area ranging from $\sim 10\%$ above 1 GeV to $\sim 20\%$ near 100 MeV. The rate of ghost events is also proportional to the charged-particle background rate, and the resulting loss of effective area can be estimated from the trigger rate. Figures 2.2 and 2.3 show the dependence of A_{eff} on energy (at normal incidence) and on incidence angle (at 10 GeV), respectively, for the P7SOURCE_V6 IRF version.

The point-spread function (PSF), $P(|\hat{v}'-\hat{v}|; E, cos(\theta))$ is a normalized probability density function angular deviation $|\hat{v}' - \hat{v}|$ photons incident on the detector with energy E and incidence angle $\hat{v}$ in the instrument frame. It is strongly energy dependent, being dominated at low energies by the $1/E$ dependence of multiple Coulomb scattering, and limited at high energies by the ratio of the pitch (center-to-center separation) of the TKR strips to the length of the particle track. The PSF for a given energy and incidence angle is parameterized as a sum of two analytic functions,

$$P(x) = f_{\rm core}K(x, \sigma_{\rm core}, \gamma_{\rm core} + (1 - f_{\rm tail})K(x, \sigma_{\rm tail}, \gamma_{\rm tail}), \tag{2.3}$$

where $K(x, \sigma, \gamma)$ is the so-called King function,

$$K(x, \sigma, \gamma) = \frac{1}{2\pi\sigma^2}\left(1 - \frac{1}{\gamma}\right)\left(1 + \frac{1}{2\gamma}\frac{x^2}{\sigma^2}\right)^{-\gamma}. \tag{2.4}$$

The parameters σ and γ represent, respectively, the width of the function and the relative

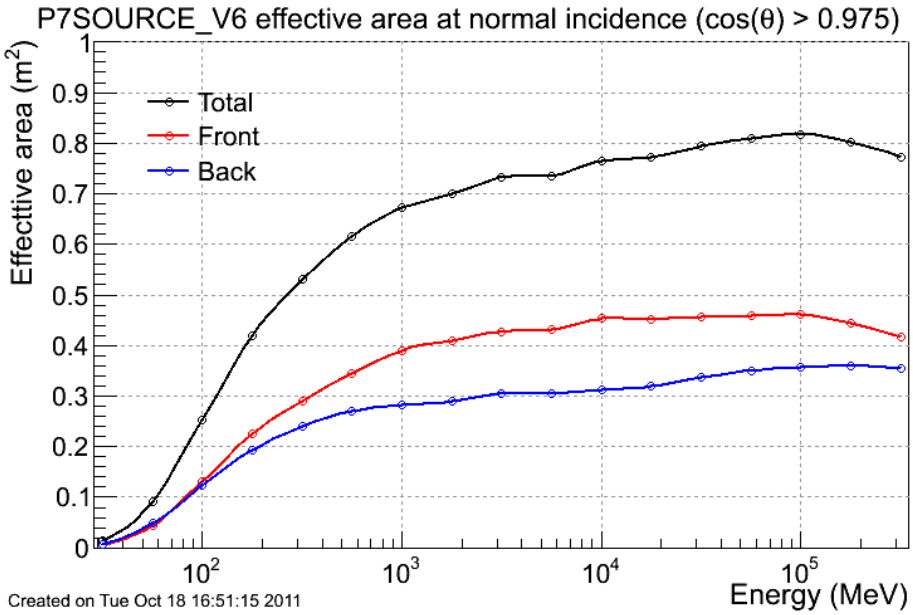

Figure 2.2 LAT effective area at normal incidence ($cos(theta) > 0.975$) as a function of energy. Reproduced from http://www.slac.stanford.edu/exp/glast/groups/canda/archive/pass7v6/lat_Performance.htm

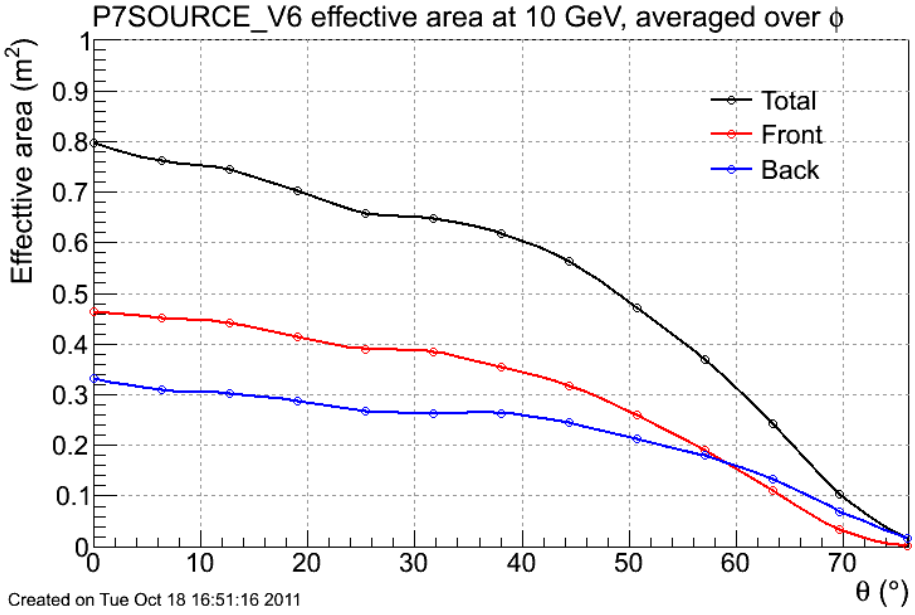

Figure 2.3 LAT effective area as a function of incidence angle θ for 10 GeV photons, averaged over azimuth. Reproduced from http://www.slac.stanford.edu/exp/glast/groups/canda/archive/pass7v6/lat_Performance.htm

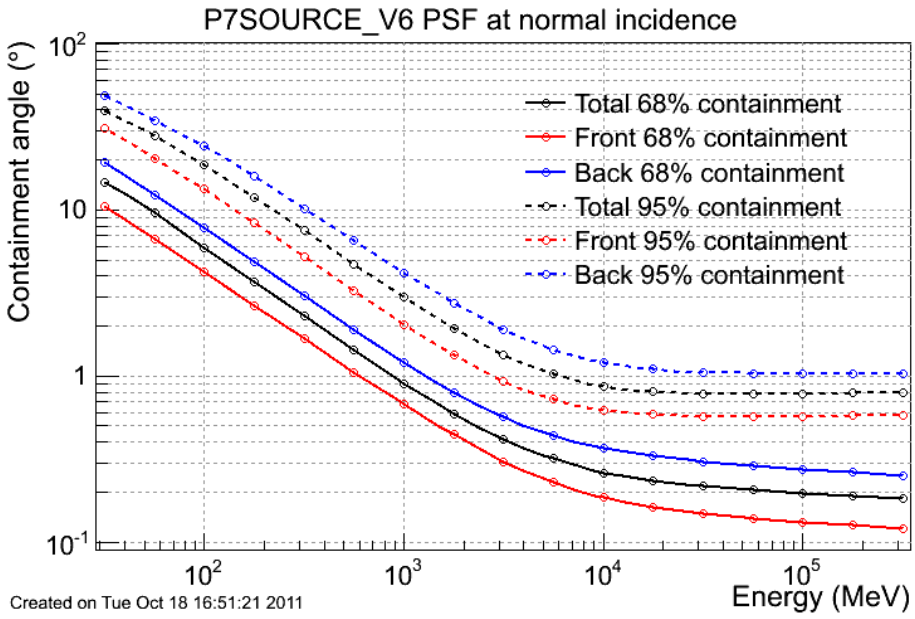

Figure 2.4 Containment radii (68% and 95%) of reconstructed incidence angles for normally incident photons ($cos(\theta) > 0.975$). Reproduced from http://www.slac.stanford.edu/exp/glast/groups/canda/archive/pass7v6/lat_Performance.htm

contribution of the tails, with the function approaching a Gaussian of width σ as $\gamma \to$ inf. The variable x is the scaled angular deviation

$$x = \frac{|\hat{v}' - \hat{v}|}{S(E)} S(E) = \sqrt{\left(c_0 * \left(\frac{E}{100MeV}\right)^{-\beta}\right)^2 + c_1^2}. \quad (2.5)$$

The parameters of the scaling function $S(E)$ are chosen so to keep the peak of the distribution of scaled deviations near 1 for all energies, so that $S(E)$ contains most of the energy-dependence of the PSF. Figure 2.4 shows the 68% and 95% containment angles, r_{68} and r_{95} as a function of energy for the P7SOURCE_V6 PSF.

The energy dispersion, $D(E'; E, cos(\theta))$ is the likelihood to measure an energy E' for a photon with energy E incident at angle θ. Although the energy dispersion is itself quite significant, its effect on spectral analyses is typically minor, due to its relatively weak energy dependence combined with the near constancy of the effective area over most of the LAT

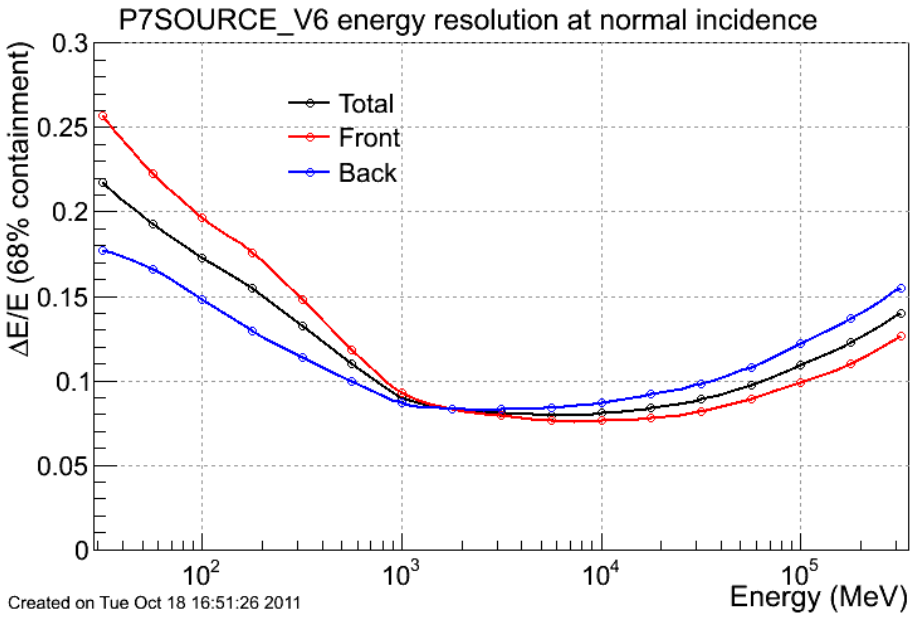

Figure 2.5 Energy resolution as a function of energy at normal incidence ($cos(\theta) > 0.975$). Reproduced from http://www.slac.stanford.edu/exp/glast/groups/canda/archive/pass7v6/lat_Performance.htm

energy range. Even at low energies, where the effective area does change rapidly as a function of the energy, the resulting biases on estimated spectral parameters, namely a slight softening of spectra and a sight increase in flux, are no larger than a few percent for the great majority of cases (Ackermann et al., 2012). It is therefore common to neglect the energy dispersion, taking $D(E'; E, cos(\theta)) = \delta(E - E')$, except in particular situations, such as searches for spectral lines.

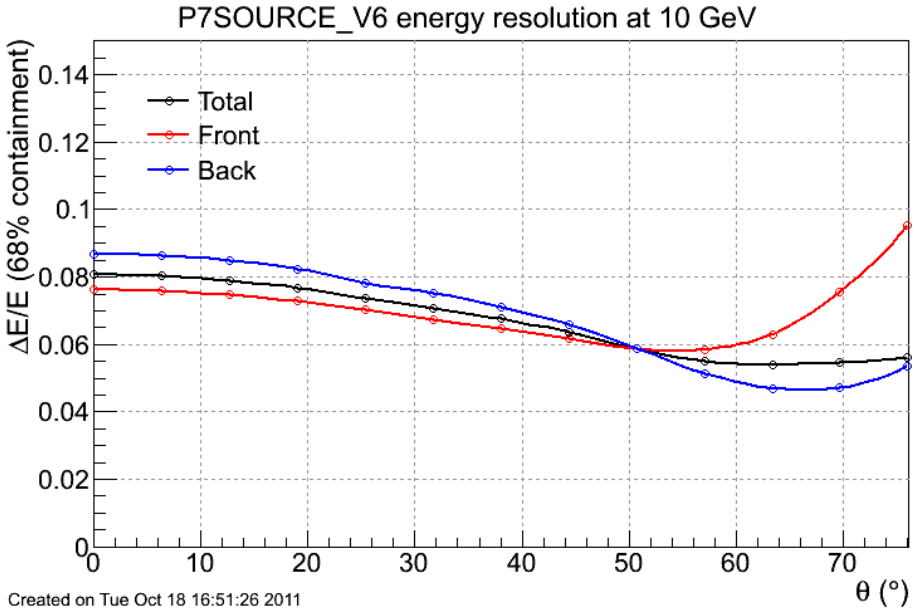

Figure 2.6 Energy resolution as a function of incidence angle at 10 GeV. Reproduced from http://www.slac.stanford.edu/exp/glast/groups/canda/archive/pass7v6/lat_Performance.htm

Chapter 3

LAT DATA ANALYSIS

The nature of γ-ray data makes the applications of many of the techniques of conventional astronomy, particularly the extraction of physical source properties directly from images, impractical. The broad low-energy point-spread function leads in many cases to significant spatial overlap of the observed emission from neighboring sources. This source confusion, combined with the strong energy-dependence of the PSF, precludes any straightforward selection of regions of the sky in which the γ-ray intensity can be attributed to a single source over a background. This problem is especially pronounced in the Galactic plane, where both the density of point sources and the intensity of the diffuse emission are far greater than at higher galactic latitudes. Furthermore, the typical photon fluxes of γ-ray sources are sufficiently low that significant detection may require very long exposures, during which the orientation of the spacecraft is constantly changing. Since the instrument response varies significantly with incidence angle, it is therefore not a simple matter to convert a count of photons attributed to a particular source to a flux measurement.

In light of these difficulties, analysis of LAT data generally proceeds through maximum likelihood (ML) analysis, in which the photon data (i.e., incident directions and energies of individual reconstructed photons) are combined with models of the instrument response (the IRFs described in § 2.4.2) and of the physical properties of the sources (position and energy spectrum of point sources, and spatial intensity profiles and energy spectra of diffuse sources) to compute a likelihood function describing the probability to observe the data given a particular model. The ML estimators for the model parameters are the values that maximize the likelihood function (or, equivalently, its logarithm).

The likelihood describes the probability for a given set of reconstructed events (the data)

to be observed given a model of the spatial distribution and energy spectrum of a set of γ-ray sources (the model). Analysis proceeds in a "forward-folded" space; that is, the model of the physical source properties is convolved with the IRFs describing the efficiency and accuracy of event reconstruction to produce a model for the emission in a coordinate system fixed to the instrument geometry, and it is this folded source model that is used to construct the likelihood. In the following sections, we first describe the spatial and spectral source models employed in typical analyses. We then describe the computation of the likelihood itself, as well as some common simplifying assumptions used to make that computation more tractable. Finally, we describe the *pointlike* package, which is designed to provide a more efficient framework for analyses equivalent to those provided by the standard *Fermi Science Tools*.

3.1 Source Models

The first input to the maximum likelihood procedure outlined above is a parameterized model of the γ-ray intensity from all sources, $S(E, \hat{p})$, as a function of energy E and celestial position $\hat{p}$. We describe here our typical models for different classes of sources, partitioned by angular size.

3.1.1 Point Sources

The vast majority of astrophysical sources of γ rays have angular sizes much smaller than the angular resolution of the LAT and are therefore point sources for our analyses. Their spatial distributions are simply Dirac delta functions in two celestial coordinates, $\delta\left(\hat{p} - \hat{p}_0\right)$, with $\hat{p}_0$ the direction to the source. The emission of a point source is characterized completely by its position along with a spectral flux density, $f(E, t; \vec{\lambda})$, where $\vec{\lambda}$ is a set of parameters specific to the chosen spectral model. While we have included the time-dependence explicitly here, in practice, for a given likelihood calculation, we will use a time-independent spectrum, performing separate calculations for each subinterval in a piecewise constant model in cases where we want to include time dependence; we therefore suppress any time dependence in

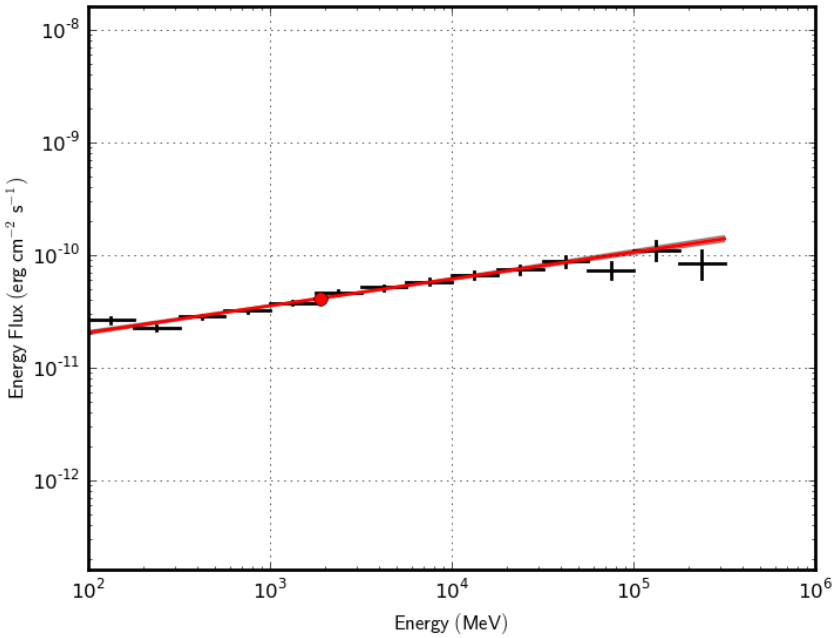

Figure 3.1 Spectral energy density of Markarian 421 (2FGL J1104.4+3812). The red line is the best-fit power-law, with a photon spectral index $\gamma = 1.75$.

the spectral models presented here.

The processes by which γ rays are emitted in astrophysical sources are generally non-thermal and produce broadband energy spectra which extend over several orders of magnitude in energy, with few or no spectral features (e.g., lines, sharp breaks, or cutoffs). In many cases, such spectra are adequately described by a simple power-law (PL) model:

$$f_{PL}(E; N_0, \gamma) = N_0 \left(\frac{E}{E_0}\right)^{-\gamma}, \tag{3.1}$$

where E_0 is an arbitrary reference energy, and N_0 and γ are the two parameters of the model, representing the differential flux at E_0 and the photon spectral index respectively. Figure 3.1 shows an example of a spectral energy distribution with the best-fit power-law model.

Pulsars tend to exhibit spectra in which an intrinsic power-law spectrum is exponentially

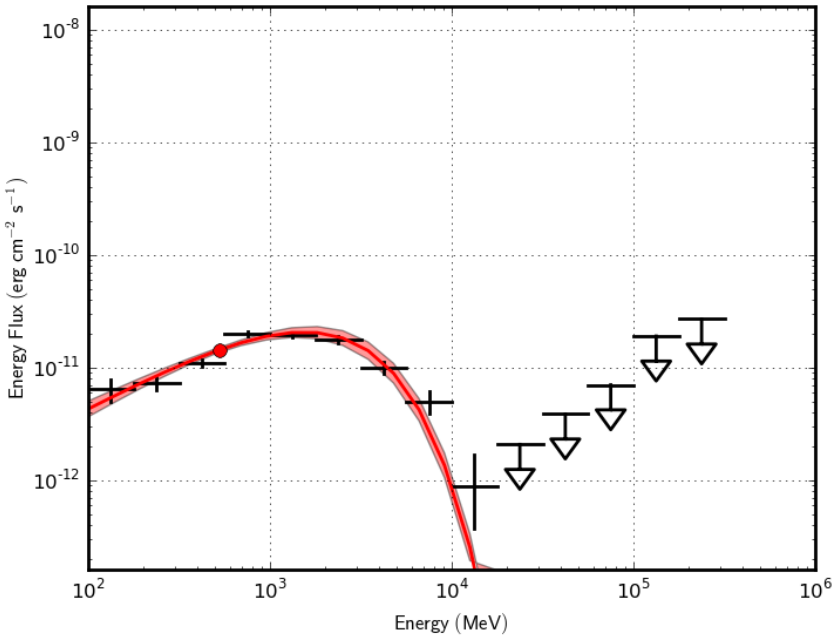

Figure 3.2 Spectral energy density of PSR J0030+0451 (2FGL J0030.4+0450). The red line is an exponentially attenuated power-law with $\gamma = 1.14$ and $E_c = 1799$ MeV.

attenuated above some characteristic energy, due to special relativistic effects near the light cylinder. The resulting exponential-cutoff (EC) model has the form:

$$f_{EC}(E; N_0, \gamma, E_c) = N_0 \left(\frac{E}{E_0}\right)^{-\gamma} e^{\left(-\frac{E}{E_c}\right)}. \tag{3.2}$$

The free parameters are the same N_0 and γ of the power-law model, along with the cutoff energy E_c. Figure 3.2 shows the spectrum of the millisecond pulsar J0030+0451, with an exponential cutoff model. The hard ($\gamma = 1.14$) low-energy spectrum is typical of pulsars. The upper limits on the flux above the cutoff energy are clearly inconsistent with an extension of the low-energy power-law component.

Finally, many sources exhibit some softening of the spectral index γ (i.e. γ transitions from a smaller value at low energies to a larger value at high energies), but without a sharp cutoff. Such spectral breaks are generally observed to be relatively smooth, rather

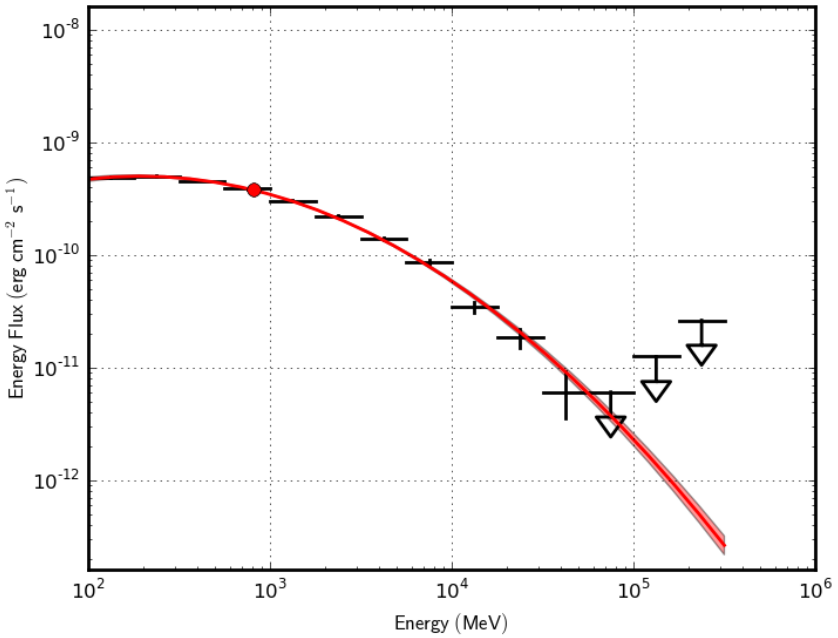

Figure 3.3 Spectral energy density of 3C 454.3 (2FGL J2253.9+1609). The red line is a log-parabola with $\gamma = 2.32$, $\beta = 0.138$ and $E_b = 620$ MeV.

than discontinuous. These cases are generally adequately described by a log-parabolic (or log-normal) model, with the spectrum given by a Guassian in $log(E)$:

$$f_{LP}(E; N_0, \gamma, \beta, E_b) = N_0 \left(\frac{E}{E_b}\right)^{-(\gamma+\beta log(\frac{E}{E_b}))} . \tag{3.3}$$

The extra parameters here, compared to the simple power law, are the break energy E_b, and the curvature parameter β. Unlike the EC model for pulsar spectra, the particular form of the log-parabolic model is not physically motivated; it is simply a convenient parameterization of spectral curvature. Such spectral softening at high energies is seen relatively commonly in bright quasars, where it may indicate the transition to a different scattering regime for the charged-particle populations responsible for the high-energy emission (i.e., the increased importance of higher-order terms in the Klein-Nishina cross section). Figure 3.3 shows the energy spectrum of the quasar 3C 454.3, fit with a log-parabola spectrum.

3.1.2 Extended Sources

A small subset of astrophysical sources, mostly supernova remnants, while identifiable as discrete sources, have angular extents large enough to be resolved by the LAT. In this case, we need to model the specific intensity of the emission as a function of position as well as energy. In practice, we typically assume that the spectrum is independent of the position, so that the intensity is separable into a spatial profile and a spectral flux density, $S(E, \hat{p}) = s(\hat{p})f(E)$. We model the spectra of the extended sources with the same set of models described above for the point sources.

Since there are so few of these sources, it is feasible to choose spatial models for their emission on a case-by-case basis, but they generally fall into two classes. For sources with angular extent resolvable by the LAT, but still small enough that the angular distribution of the emission appears smooth, it is sufficient to use a simple analytical model for the angular distribution, such as a uniform disk or a two-dimensional Gaussian. For sources that are sufficiently extended that substructure is resolvable in the spatial distribution (for example, the nearby radio galaxy Cen A), the γ-ray emission may be assumed to trace the emission in other wavebands, such as radio, and templates of the emission from observations in those wavebands may be used as to model the spatial distribution of the γ-ray intensity. In this work, extended sources are modeled using the spatial models given in (Nolan et al., 2012).

3.1.3 Diffuse Emission

The diffuse γ-ray background is separated into a Galactic component (originating within the Milky Way) and an extragalactic component. The extragalactic component is isotropic, while the Galactic component is, of course, highly structured. Although the Galactic emission is strongly concentrated along the plane of the Milky Way, it is not negligible even near the Galactic poles, where its intensity is comparable to that of the extragalactic background.

The diffuse Galactic emission arises from the interaction of cosmic-ray electrons and protons with the interstellar gas and radiation in the galaxy. The primary modes of γ ray

production in such interactions are the decay of neutral pions produced in hadronic interactions, inverse Compton (IC) scattering of cosmic-ray electrons with interstellar radiation, and Bremsstrahlung radiation of cosmic-ray electrons in the interstellar gas. The standard template provided by the LAT collaboration for modeling the diffuse Galactic emission is derived from gas column densities in galactocentric annuli, extracted from spectral line surveys of H_I and CO, and a template for the spatial distribution of IC γ rays modeled with GALPROP (Strong et al., 2007). The gas emissivities, inverse Compton normalization, and isotropic intensity (accounting for local inverse Compton), are fit to the observed γ ray signal in each of several energy bands to produce a template for the diffuse Galactic emission. Details on the procedure for and inputs to this fit are available through the *Fermi* Science Support Center website.

The remaining diffuse emission at high galactic latitudes ($b \gtrsim 30^\circ$) not modeled by the galactic template is taken to represent the combination of diffuse extragalactic emission and residual cosmic-ray backgrounds (i.e., cosmic-rays misclassified as photons). This component is modeled as an isotropic flux of γ rays with a spectrum fit to the residuals between the observed high-latitude counts and the Galactic diffuse template. Since the degree of cosmic-ray contamination varies depending on the event selection criteria, separate spectral templates are produced for each event class, and for front- and back-converting photons.

3.2 Instrument Response

The models of the previous section provide a prediction of the flux of γ rays incident on the LAT as a function of direction, energy, and time. This predicted flux then needs to be translated into a predicted rate of observed counts in the detector, which we do by convolving it with a response matrix describing the efficiency and accuracy of event reconstruction. In all that follows, we will use primed variables to denote the reconstructed properties of a detected photon, and unprimed variables to denote the true properties of the incident photon. The LAT measures the photon arrival times with accuracy better than $1\mu s$, which is negligible for our purposes, so we do not distinguish between the true and measured photon arrival

times.

Given a model $S(E, \hat{p}; \vec{\lambda})$ of the intensity of γ rays on sky, as a function of photon energy E and incident direction $\hat{p}$, the predicted distribution of reconstructed events in a given time range can be expressed as a convolution of the model with a response function describing the efficiency and accuracy of photon detection and reconstruction:

$$M(E', \hat{p}'; \vec{\lambda}) = \iiint dE d\hat{p} dt S(E, \hat{p}; \vec{\lambda}) R(E', \hat{p}'; E, \hat{p}, t). \tag{3.4}$$

$M(E', \hat{p}'; \vec{\lambda})$ is the expected number of events in an infinitessimal bin centered at energy E' and celestial position $\hat{p}'$. The integrals in Eqn. (3.4) are over the energy range of the LAT, the time range for the observation (excluding any detector deadtime), and a fiducial region of the sky; the integral over solid angle in principle covers the full sky, but in practice it only need be large enough to include any sources that may contribute an appreciable photon flux to the region of interest for the analysis, as determined by the point-spread function. The response function $R(E', \hat{p}'; E, \hat{p}, t)$ is the probability for a photon with energy E incident from the celestial position $\hat{p}$ to be detected with apparent energy and incidence angle $E', \hat{p}'$. Under the assumptions discussed in § 2.4.2 regarding the separability of the instrument response functions (IRFs), the response function is the product of the IRFs presented there:

$$R(E', \hat{p}'; E, \hat{p}, t) = A_{\rm eff}(E, \hat{p}, t) P(\hat{p}'; E, \hat{p}, t) D(E'; E, \hat{p}, t). \tag{3.5}$$

Though we have expressed the response in terms of the celestial direction $\hat{p}$, the IRFs are more naturally formulated as functions of the incident photon direction in the instrument reference frame, $\hat{v}(\hat{p}, t)$. The time dependence of the PSF and energy dispersion arises solely from the time-dependence of the transformation between the celestial and instrument-centered reference frames, due to the changing position and orientation of the spacecraft. The effective area term includes the efficiency of event detection, which depends in part on the charged-particle background rate and thereby the geomagnetic environment of the spacecraft; it therefore carries an explicit time dependence in addition to that associated with the

coordinate transformation. The calculation of Eqn. (3.4) could in principle be carried out using a time-dependent model $S(E, \hat{p}, t; \vec{\lambda})$; in practice, however, we generally treat any time dependence using a piecewise-constant model, evaluating the likelihood separately in each of a series of time bins.

As noted in § 2.4.2, it is common to neglect the effects of energy dispersion, eliminating the integral over energy in Eqn. (3.4). With that simplification, Eqn. (3.4) becomes

$$M(E', \hat{p}'; \vec{\lambda}) = \iint d\hat{p}dtS(E, \hat{p}; \vec{\lambda})A_{\text{eff}}(E', \hat{p}, t)P(\hat{p}'; \hat{p}, E, t). \tag{3.6}$$

For a model containing multiple sources, indexed by j, the total modeled photon distribution is

$$M(E', \hat{p}'; \vec{\lambda}) = \iint d\hat{p}dt \sum_{j=1}^{N_{sources}} S_j(E, \hat{p}; \vec{\lambda_j})A_{\text{eff}}(E', \hat{p}, t)P(\hat{p}'; \hat{p}, E, t), \tag{3.7}$$

where $\vec{\lambda_j}$ is the vector of parameters specifying the model for the jth source.

3.3 Construction of the likelihood function

To compare the predicted photon distributions of the previous section to the observed counts, we bin the observed counts in (reconstructed) energy and position. The probability of the observed count distribution under a given model, with parameters $\vec{\lambda}$, is the product of Poisson probabilities for each bin,

$$\mathcal{L}(\vec{\lambda}) = \prod_{i=1}^{N_{bins}} \frac{\nu_i(\vec{\lambda})^{n_i} e^{-\nu_i(\vec{\lambda})}}{n_i!}, \tag{3.8}$$

where n_i and $\nu_i(\vec{\lambda})$ are the observed and predicted numbers of counts, respectively, in the ith bin. The predicted number of counts in bin i is

$$\nu_i(\vec{\lambda}) = \iint dE'd\hat{p}'M(E', \hat{p}'; \vec{\lambda}). \tag{3.9}$$

It is common to use, instead of the likelihood itself, its logarithm

$$log(\mathcal{L})(\vec{\lambda}) = \sum_{i=1}^{N_{bins}} \left[n_i \, log(\nu_i(\vec{\lambda})) \right] - N_{pred}, \tag{3.10}$$

where N_{pred} is the total number of photons predicted by the model, and additive terms with no dependence on the parameters $\vec{\lambda}$ have been dropped.

In the limit of small bins, such that the photon count for each bin is either 0 or 1, the likelihood in equation Eqn. (3.8) reduces to the so-called unbinned likelihood. In this approach, each photon is accounted for individually, and no information is lost to binning. The unbinned likelihood is optimal in the sense of making the best use of the data to estimate the model parameters, but is more computationally intensive than the binned version(Nolan et al., 2012).

The maximum likelihood estimators $\hat{\lambda}$ for the model parameters are the values that maximize the likelihood (or, equivalently, the log likelihood):

$$\left. \frac{\partial \mathcal{L}}{\partial \lambda_i} \right|_{\lambda_i = \hat{\lambda}_i} = 0. \tag{3.11}$$

Statistical uncertainties on the model parameters can be estimated from the second derivatives of the likelihood,

$$V_{ij}^{-1} = - \left(\frac{\partial^2 \mathcal{L}}{\partial \lambda_i \partial \lambda_j} \right) \Bigg|_{\lambda_{i,j} = \hat{\lambda}_{i,j}}, \tag{3.12}$$

where V_{ij} denotes the estimated covariance of the estimators $\hat{\lambda}_i$ and $\hat{\lambda}_j$. Equivalently, a 68.3% confidence interval can be defined by the contour

$$log\mathcal{L}(\lambda) = log(\mathcal{L}(\hat{\lambda})) - \frac{1}{2}. \tag{3.13}$$

In the large sample limit, these estimates correspond exactly to the classical confidence intervals.

3.3.1 Likelihood Ratio Tests and Detection Significance

To assess the significance of a detection, we use a likelihood ratio test. The likelihood ratio test statistic comparing two models with parameters $\vec{\lambda_0}$ and $\vec{\lambda_1}$ is

$$\mathrm{TS} = -2log\left(\frac{\mathcal{L}_{max}(\hat{\lambda_0})}{\mathcal{L}_{max}(\hat{\lambda_1})}\right). \tag{3.14}$$

We consider the particular case where the models are nested, meaning that the parameters $\vec{\lambda_0}$ are a proper subset of the parameters $\vec{\lambda_1}$, so that model 0 (the null hypothesis) can be considered as a special case of model 1 (the alternate hypothesis) with one or more of the $\vec{\lambda_1}$ fixed to particular values. Provided that none of the extra parameters take on values on the boundary of their allowed range under the null hypothesis (e.g., a strictly non-negative parameter that is zero under the null hypothesis), and that the likelihood satisfies certain regularity conditions, Wilks' theorem (Wilks, 1938) states that under the null hypothesis, the test statistic TS is asymptotically distributed as

$$\mathrm{TS} \sim \chi^2_n, \tag{3.15}$$

where n is the difference in the number of free parameters between the two models. In the case where the alternate hypothesis contains one more parameter than the null hypothesis, and its value under the null hypothesis lies on the boundary of its allowed range, the related Chernoff's theorem (Chernoff, 1954) states that the resulting asymptotic null distribution is

$$\mathrm{TS} \sim \frac{1}{2}(\chi^2_0 + \chi^2_1), \tag{3.16}$$

where χ^2_0 denotes the degenerate distribution localized at zero[1]; that is, under the null hypothesis, the test statistic is expected to be zero half the time and distributed as chi-

[1]The χ^2_0 notation is chosen in anticipation of the following discussion of sums of random variables with this distribution. An alternative statement, perhaps more familiar to physicists, is that the PDF of χ^2_0 is a delta function at zero: $P_{\chi^2_0}(X = x) = \delta(x)$.

squared with one degree of freedom the other half.

In this form, the likelihood ratio test provides our most commmon estimate for the significance of detection of a particular source. Considering the comparison between a model containing parameters for some set of sources to a null model in which the normalization parameter in the spectral model of a source of interest is fixed to zero, the likelihood ratio test statistic comparing these models is distributed according to Eqn. (3.16). Quantiles of this distribution can be found in terms of the cumulative distribution function of the normal distribution by

$$P(\mathrm{TS} \geq n^2) = 1 - CDF_{normal}(n|0,1), \tag{3.17}$$

so that TS ≥ 9 corresponds to a detection at a 3-σ confidence, TS ≥ 25 to 5-σ, and so on. TS ≥ 25 is the standard employed by the LAT team for the published point-source catalogs.

Taking the source flux to be constant across each individual band, the test statistic defined above provides a measure of the detection significance for each band. The sum of the band TS values, TS_{band}, is a measure of the overall detection significance, independent of the spectral shape. In the null case, the TS value for each band is distributed according to Eqn. (3.16). To find the distribution of the sum of N such random variables, we can consider the moment-generating function for the distribution of Eqn. (3.16),

$$M(t) = \frac{1}{2}(1 + (1-2t)^{-\frac{1}{2}}). \tag{3.18}$$

If $\vec{X} = X_0, X_1, ..., X_N$ are independently distributed random variables with moment-generating functions $M_0(t), M_1(t), ..., M_N(t)$, then the moment-generating function for the sum $\sum X_i$ is $M_{sum}(t) = M_0(t)M_1(t)...M_N(t)$. For the particular case where the X_i are identically distributed according to Eqn. (3.16), the moment-generatiing function for their

sum is

$$\begin{aligned} M_{sum}(t) &= \frac{1}{2^N}(1+(1-2t)^{-\frac{1}{2}})^N \\ &= \frac{1}{2^N}\sum_{k=0}^{N}\binom{N}{k}(1-2t)^{-\frac{k}{2}}, \end{aligned} \tag{3.19}$$

where the second line follows from application of the binomial theorem to the first. Recognizing $(1-2t)^{k/2}$ as the moment-generating function for the χ^2 distribution with k degrees of freedom, we conclude that the asymptotic distribution of a sum of N random variables TS_i distributed according to Eqn. (3.16), $TS_{sum} = \sum TS_i$, is a weighted mixture of χ^2 distributions with different numbers of degrees of freedom:

$$\mathrm{TS}_{sum} \sim \frac{1}{2^N}\sum_{k=0}^{N}\left(\binom{N}{k}\chi^2_i\right). \tag{3.20}$$

In particular, we expect that the asymptotic null distribution for the statistic TS_{band}, calculated from a spectrum in N_{band} energy bands, is that of Eqn. (3.20) with $N = N_{band}$.

3.4 pointlike

The *pointlike* package was originally developed as a fast tool for identifying and localizing point sources, and eventually extended to provide a full range of tools for likelihood analysis of LAT data. It takes advantage of several approximations in order to optimize for speed without sacrificing significant accuracy compared to the officially distributed *Science Tools*.

3.4.1 Event binning

Photons are first binned in energy, typically into bands of equal logarithmic width. Energy band sizes are chosen for a given analysis to balance spectral resolution and photon counts per bin. A typical value is four logarithmic energy bands per decade, which for most point source analyses provides sufficient resolution to capture most interesting spectral features of

γ-ray point sources, and sufficient statistics to characterize the source flux in most energy bins with reasonable accuracy. Bands of width $.25dex$ are broader than the energy dispersion at all energies, further minimizing the effect of neglecting it in Eqn. (3.6). *Front* and *back* photons are binned separately, resulting in two bins per energy band.

The angular binning is done using a HEALPix tesselation (Górski et al., 2005). The HEALPix resolution is chosen based on the width of the PSF for each combination of energy band and conversion type. Provided that the pixel size is small compared to the PSF, any effects on the likelihood due to the angular binning are swamped by the limited angular resolution of the detector, and the analysis is essentially identical to the unbinned version. Within that constraint, the pixel size can be increased freely without significant negative effects on parameter estimation.

The effects of this energy-dependent binning scheme can be seen in Table 3.1, which shows a typical binning of events from 30 MeV to 1 TeV in bands of width $\Delta(log(E)) = .25$. For each combination of energy band and conversion type, the 68% containment radius of the PSF (R_{68}), pixel size, number of occupied pixels, and number of photons are shown. The effect of the energy-dependent pixelisation is greatest at low energies where the PSF is broadest, as shown by the small numbers of pixels and large numbers of photons per pixel. Conveniently, due to the typical energy spectra of astrophysical γ-ray sources combined with the energy-dependence of the instrument effective area, this is also the portion of the LAT energy range where the most photons are detected, further enhancing the data compression. At very high energies (above about 10 GeV), where photons are sparse and the energy-dependence of the PSF flattens out, a minimum pixel size is chosen which is small enough that multiple counts per pixel are very unlikely (one photon per pixel for each of those bands), and the likelihood calculation for these bands approaches the optimal unbinned limit; the unbinned calculation is quite efficient at these energies due to the sparseness of occupied pixels.

Table 3.1 Event binning for 3-year all-sky analysis
Events are separated into energy bands from E_{min} to E_{max}, with separate bands for *front*- or *back*-converting photons (indicated by the "type" column). N_{side} is the HEALPix N_{side} parameter for the band. "Pixels" and "photons" indicate the total number of pixels and photons in each band.

conversion type	E_{min}	E_{max}	N_{side}	pixels	photons
front	100	178	86	88752	4684790
back	100	178	52	32448	5898981
front	178	316	135	218684	5657728
back	178	316	82	80688	5874750
front	316	562	216	541221	4812286
back	316	562	132	209008	4417866
front	562	1000	355	966602	3133136
back	562	1000	218	495639	2645093
front	1000	1778	644	1080515	1786068
back	1000	1778	401	701341	1421306
front	1778	3162	1543	802119	910804
back	1778	3162	1007	597559	704476
front	3162	5623	5490	383445	391871
back	3162	5623	4514	294785	298931
front	5623	10000	8181	165703	166605
back	5623	10000	8173	121626	121840
front	10000	17783	8192	68736	68816
back	10000	17783	8192	48666	48681
front	17783	31623	8192	28787	28811
back	17783	31623	8192	20236	20237
front	31623	56234	8192	12202	12208
back	31623	56234	8192	9441	9441
front	56234	100000	8192	4985	4987
back	56234	100000	8192	5127	5127
front	100000	177828	8192	2060	2061
back	100000	177828	8192	2923	2925
front	177828	316228	8192	706	706
back	177828	316228	8192	1381	1381

3.4.2 Point Source Localization and Detection

In the models of § 3.1 and § 3.3, the positions of the point sources are treated on the same footing as the spectral parameters. In a straightforward implementation of the likelihood described in § 3.3, one would maximize the likelihood simultaneously over the joint set of spectral and positional parameters. In the *pointlike* implementation, localization (that is, the optimization with respect to source position) is treated separately from the spectral fitting, with the underlying assumption that the spatial and spectral components of the model are independent. Spectral fits are performed with the source position fixed, and vice versa.

To localize a point source, we begin with a model of the sources in the region of interest, based on an existing list of known sources[2], and perform an initial spectral fit. We then scan the over a grid of nearby positions, evaluating the likelihood-ratio test statistic (TS) for detection of the source in question at each point. This results in a map of the likelihood surface for the source location, which we refer to as a "TS map". Contours of constant TS near a well-defined maximum of such a map are elliptical in the Gaussian limit. The parameters (center, semi-major and semi-minor axes, and position angle of one axis) of the ellipse corresponding to the TS = 1 ($\Delta log(\mathcal{L}) = .5$) contour provide an estimate of the position parameters and their covariance matrix. Localizing point sources in this way is therefore equivalent to simply optimizing the position parameters jointly with the spectral parameters, up to the neglect of any correlations between the position and spectrum.

A simple measure of the quality of the elliptical fit is the mean-squared deviation of the fit function from the calculated TS along a contour of the ellipse. Large deviations indicate significant departure from ellipticity. Typically, we evaluate this quality measure along the 2σ contour, though the choice is somewhat arbitrary. Such a departure from ellipticity can arise for faint sources in regions with relatively strong, structured backgrounds, such as near the galactic plane. It can also be indication of source confusion, such as from the presence of

[2]If for some reason we wanted to create a catalog *de novo*, without reference to prior analyses, a list of seed sources could be produced from the data by image processing methods, based on clustering algorithms, wavelets, etc.

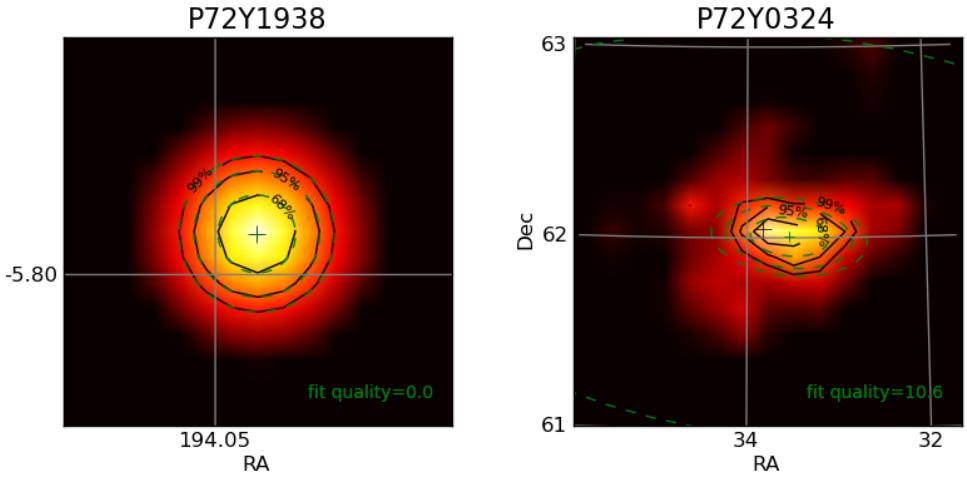

Figure 3.4 Example TS maps for a well-localized source (left) and a poorly-localized source (right). The black solid contours are at TS=1.51,2.45, and 3.01, corresponding to 68%, 95%, and 99% containment for a gaussian likelihood surface. The green dashed contours show the corresponding contours from the elliptical fit.

multiple unresolved point sources, or simply that the detection is spurious and the putative point source is just a fluctuation in the background. Visual inspection of the TS map can help to distinguish between these cases. Figure 3.4 shows example TS maps for a well-localized source and a poorly-localized source.

Treating the localization separately in this way has two main advantages, however. One is that visual inspection of the TS maps for poorly-fit sources can provide a useful troubleshooting tool. More important is their usefulness in detection of new sources. In a common scenario, we have an existing catalog of known sources that we want to update in light of new data, improvements in event reconstruction, revised models of the instrument response or the diffuse backgrounds, or some combination of such. To search a region for new potential sources, we begin by optimizing the existing model, including both positions and spectra of the point sources. We then add a test point source to the model, typically with a simple power-law spectrum with photon spectral index Γ=2. Evaluating the TS for detection of this source at each position on a grid, just as in the localization procedure above,

produces a map of the "residual significance"; that is, the model residual expressed in units of point source significance, rather than counts or flux. Candidate sources can be identified from this map by a simple clustering algorithm or similar, and their quality assessed from an elliptical fit to the TS contours. A putative source meeting appropriate thresholds for peak significance and eccentricity of the error ellipse is added to a list of candidate sources, and this procedure is repeated for each region on the sky. The list of candidates is added to the existing list, and the procedure iterates to equilibrium.

Chapter 4

LIGHT CURVES

In this chapter, we describe the construction of monthly light curves for all point sources detected in the first three years of *Fermi* data collection, using the *pointlike* package. The optimizations described in chapter 3, making reasonable approximations to acheive major improvements in performance with little if any effect on accuracy, make *pointlike* particularly well-suited to this type of large-scale analysis. For this reason, as of the second *Fermi* point-source catalog (2FGL; Nolan et al. (2012)), an all-sky *pointlike* analysis was used for the detection and localization steps, and the initial spectral analysis step in the construction of the *Fermi* catalogs. In this work, we extend the *pointlike* catalog analysis to the time domain, constructing light curves for the full set of sources. We shall begin, then, by describing the fundamental full-sky analysis, including the source models and fitting procedure. We then describe the details of the construction of the source list and the main light curve analysis. Finally, we present some sample light curves, and a summary of the properties of the full set.

4.1 3-year All-sky Analysis

4.1.1 Data Selection and Instrument Response Functions

For the 3-year analysis, we analyzed data spanning the first three years of the mission, in the range from August 4, 2008 through July 31, 2011 (Mission Elapsed Time (MET) 239557417-333849602[1]. The data are selected according to the standard Pass 7 Source class selections,

[1]Mission Elapsed Time (MET) is defined as seconds since 00:00:00 UTC on 2001 January 1.

Table 4.1 Excluded time intervals
Beginning and end (t_{start} and t_{stop}) of excluded time intervals, in MET. The reason for excluding each interval (all particular γ-ray) bursts) are given in the "reason" column.

t_{start}	t_{stop}	$\Delta t(s)$	Reason
243216749	243217979	1230	GRB 080916C
263607771	263625987	18216	GRB 090510
273582299	273586600	4301	GRB 090902B
275631598	275632048	450	GRB 090926A
333803358	333804842	1484	GRB 110731A

as provided by the FSSC Data Catalog [2]. Further cuts are applied on the photon incidence angle ($\theta < 66.4°$) and zenith angle ($z < 100°$). Events are binned spectrally in energy bands of equal width in log(E), with 4 bands per decade of energy, from 100 MeV to 316 GeV. Spatially, events are binned in a HEALPix pixelization with energy-dependent pixel size, using the same binning described in Table 3.1.

We mask out a few short time ranges corresponding to particularly bright GRBs, to avoid skewing results for nearby sources; the time ranges involved are short enough that these omissions will not affect the results for steady sources. The masked intervals are listed in Table 4.1.

We use the P7SOURCE_V6 IRF version, which includes the on-orbit calibration of the PSF described in Roth (2012). Due to an observed discrepancy between the *front* and *back* effective area at low energies with the standard IRF, we adjust the low-energy *back* effective area slightly, scaling it down by a factor increasing linearly from .82 at 100 MeV to 1 at 300 MeV. The *front* effective area is unchanged. We neglect the effects of energy dispersion.

[2] http://www.slac.stanford.edu/exp/glast/groups/canda/archive/pass7v6/lat_Performance.htm

4.1.2 Source Detection and localization

The initial set of point sources for this analysis is the source list from the 2FGL catalog, comprising 1873 sources. To that list, we added seeds derived by the residual TS method described in § 3.4.2. We discard those seeds for which the localization procedure did not converge or resulted in an elliptical fit quality less than 4. The resulting list contains 3517 candidate point sources. The model also includes 22 sources previously found to show significant spatial extension(Lande et al., 2012; Ackermann et al., 2011a).

4.1.3 Regions of Interest

In order to make joint optimization of the model over the full sky computationally tractable, we divide the sky into 1728 regions of interest (ROIs), each containing between 0 and 25 sources, which are optimized independently. The ROIs are concentric with the pixels of a HEALPix pixelization with $N_{\rm side} = 12$, corresponding to a pixel size of $\sim 5^\circ$. In each ROI fit, sources within the corresponding HEALPix boundary are fit to data extracted in a circle of radius 5° concentric with the HEALPixel[3]. Since the models of sources in one the outer regions of one ROI comprise part of the background for neighboring ROIs, several iterations of the ROI fits are performed to allow convergence between ROIs.

4.1.4 Source Models

Point and Extended Sources

We model the spectra of the point and extended sources using the set of basic models described in chapter 3. Sources from the 2FGL catalog are assigned the model provided in the catalog. For the 119 known γ-ray pulsars (i.e., those from which pulsations have been detected in γ rays), we use the exponentially-attenuated power-law model; all of the pulsars have been subject to more detailed dedicated studies, and we initialize them with the best spectral parameters from previous analyses, generally those from the second LAT

[3]Extended sources are assigned to the ROI containing the centroid of the spatial model.

pulsar catalog(Abdo et al., 2013). The remaining sources are initially assigned power-law spectra. After an initial fit, sources modeled with power-law spectra are tested for spectral curvature, and the power-law model is replaced by a log-parabolic one in cases where the change significantly improves the likelihood.

Diffuse Sources

The galactic diffuse emission is modeled using the standard template distributed by the *Fermi* Science Support Center. The isotropic component of the diffuse emission is modeled with independent spectral models for *front* and *back* events. Additionally, the *front* spectrum has a correction applied at low energies to account for the deviation of the true effective area from the standard IRF described above.

The diffuse models are optimized with respect to the full-sky γ-ray distribution, and the observed diffuse flux in a particular region will naturally show some deviation from the model, due to expected statistical fluctuations and possibly to unmodeled astrophysical sources (e.g., molecular clouds). Because the diffuse emission comprises a large fraction of the total γ-ray flux, especially in regions near the galactic plane, local residuals in the diffuse emission can disproportionately affect fits to the spectra of point sources in the region. To minimize such effects, we therefore scale the spectra of the galactic and isotropic diffuse components in each ROI by an overall normalization factor, and the galactic component by a small corrective slope. The normalization for the isotropic emission is common to the *front* and *back* components.

Finally, the emission from the Earth limb is modeled using a template derived from observed residuals in two years' worth of data compared to the fluxes predicted from the 2FGL model. Since the limb emission is most prominent near the celestial poles, where the off-axis exposure is more dominant, we only include this component in the models for ROIs whose centers lie at declinations $\delta > 45°$. The degree of contamination by limb emission is greater in the *back* events than in the *front*, owing to the broader PSF. We therefore scale the limb component by separate normalizations for the *front* and *back*.

4.1.5 Summary of All-sky Fitting Procedure

Starting with the model constructed as described above, we independently optimize the spectral models for each of the ROIs, recording the value of the log-likelihood before and after the fit. On each iteration, any ROI with a log-likelihood difference $\Delta\log\mathcal{L} \geq 10$ in the previous iteration is refit, along with its neighbors. When the spectral fits have converged, we repeat the localization step and search the residual TS map for additional seed sources, and repeat the test for spectral curvature, changing model types as needed. The iteration of localization and detection, model updating, and spectral fitting continues until the model converges.

4.1.6 Population Statistics

Of the 3539 candidate sources in the model, 2247 are detected with TS $\geq$ 25, and 3179 with TS $\geq$ 10. Figure 4.1 shows the distribution of TS values. Figure 4.2 shows the distribution of photon flux from 100 MeV to 100 GeV for sources with TS $\geq$ 25 in the 3-year analysis. For comparison, the same distribution is shown for the reported 2-year average fluxes for the point sources in 2FGL. The distributions are qualitatively similar, with the main difference being the increased detection sensitivity at moderate to low source fluxes. Of the 1873 point sources in the 2FGL catalog, 1641 are detected in our 3-year analysis with TS $\geq$ 25, and 1731 with TS $\geq$ 10.

4.2 Light Curve Analysis

The basic procedure for the light curve analysis is the same as for the all-sky analysis described above, performed independently for each month. In light of the limited statistics available for each individual fit and the need to keep the overall computation time reasonable, the monthly fits are simplified by restricting the set of parameters left free in the monthly fits, as described in the § 4.2.3. In order to more accurately describe the uncertainty in the flux of individual sources, we introduce a profile likelihood, described in § 4.2.4.

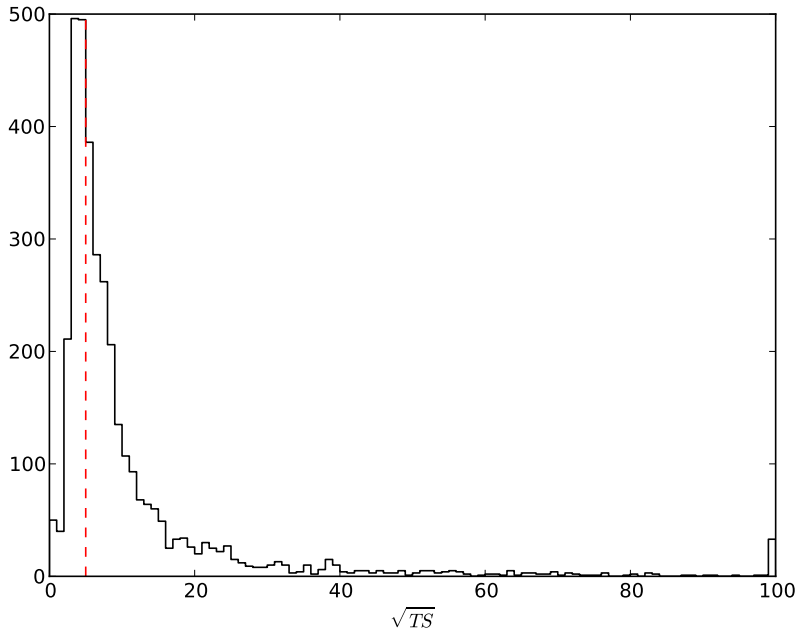

Figure 4.1 Distribution of $\sqrt{\mathrm{TS}}$ in the 3-year analysis.

The flux estimates from the profile likelihood are the basis of the light curves used in the remainder of the chapter.

4.2.1 Data Selection and Instrument Response Functions

The light curves presented here span 42 months, from August 4, 2008 through Jan 31, 2012 (MET 239241601-349747202). The data selections and instrument response functions are the same as those used for the 3-year analysis described in § 4.1. The same intervals listed in Table 4.1; no additional intervals are added in the final six months.

4.2.2 Time Binning

The optimal time binning for a light curve is, in general, dependent on the properties of the source being studied. Fainter sources, and those in regions of higher background, require longer integration times for significant detection, and therefore call for coarser binning to

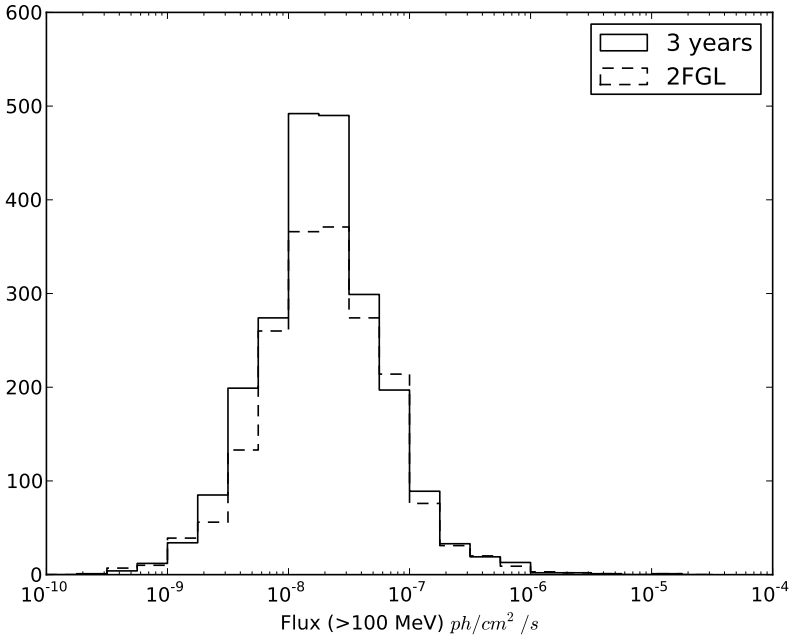

Figure 4.2 Distribution of point source flux in the range 100 MeV to 100 GeV, for sources with TS > 25 in the 3-year analysis (solid line) and the 2FGL(dashed line) sample.

ensure that the source flux can be adequately measured in most time bins, while brighter sources allow for finer binning and improved resolution for short-term variability. If a source is highly variable, different time binnings may be appropriate for different time periods. If it is primarily the episodes of stronger emission that are of interest, one may simply restrict the analysis to those intervals, choosing an appropriate fixed time binning. An alternative approach, used for example by Lott et al. (2012), is to employ a variable binning, choosing the length of each time bin individually, such that the detection significance or fluence in each bin is approximately constant.

In the analysis at hand, we are simultaneously constructing light curves for a large, heterogenous set of sources, which introduces further complications. Since we are jointly fitting all of our sources in each bin, we require commensurate binning for all light curves; we can neither select a variable binning based on the behavior of a particular source, which would be inappropriate for all the rest, nor select different binnings for each source, which would preclude the joint fitting. We must then balance the conflicting desires for sufficient resolution to distinguish interesting features in the light curves of the brighter sources, and sufficient integration time that the fainter sources are detectable in enough time bins to be useful. Furthermore, since each fit of the full-sky model is considerably more complex than that for a single-source analysis, we are also constrained by computation time in how fine a binning we can practically use.

For this analysis, we use time bins of one month. This is too coarse to resolve the particularly interesting flares of bright blazars, but provides sufficient exposure to significantly detect a large sample of sources in enough bins for a broad characterization of their variability. Since blazars (by far the largest population of variable gamma-ray sources) are variable on all observed time scales, monthly binning should be sufficient for discriminating variable sources in general.

There is no need for the time bins to be exactly the same length. The strong spatial and time dependence of the exposure requires that it be computed for each point on the sky for each time interval, and guarantees that no uniform time binning could produce uniform

exposure to different regions of the sky. We therefore chose to use calendar months for our time binning, rather than dividing the full period of observation into equal-duration subintervals. This is purely for convenience of interpretation and data-handling. It is a different convention than that used in the 2FGL, which has minor implications for comparisons with the results of that analysis. The exact intervals used are given in Table 4.2.

Within each interval, data are selected and binned according to the same criteria as for the 3-year analysis, as described in § 4.1.1.

4.2.3 Source Models

The source models used in the one-month fits are the same as those in the 3-year fit. This includes the diffuse emission models and normalizations, the limb emission template, locations and spectral models of point sources and extended sources, and spatial models of extended sources. The essential difference is that the monthly fits use a restricted set of free parameters, in order to reduce the complexity of the fits and keep the processing time down.

Diffuse Models

Since we expect the intensity of the diffuse emission to remain essentially constant on the timescales we are concerned with, we leave the per-ROI normalization factors for the Galactic and isotropic emission fixed in the light curve analysis. As noted in chapter 3, the model of the isotropic emission includes a component due to irreducible backgrounds, which are expected to vary somewhat in time due to their dependence on the geomagnetic environment of the spacecraft; however, since these backgrounds represent a small fraction of the true isotropic γ-ray emission any such variation is expected to have a negligible effect on the fits to values of interest, and the reduced statistics imposed by the shorter time intervals used in the light-curve analysis are insufficient to constrain such variations.

The emission from the Earth limb is concentrated at large zenith angles. This leads to increased limb contamination in intervals with greater exposure toward the celestial poles. We therefore expect variation in the observed flux from the limb, and we leave the limb

Table 4.2 Time intervals used in the light curve construction, as calendar months and corresponding start and end times in Mission Elapsed Time

Interval Number	Month	$t_{start}(MET)$	t_{end} (MET)
0	Aug 2008	239241601	241920001
1	Sep 2008	241920001	244512001
2	Oct 2008	244512001	247190401
3	Nov 2008	247190401	249782401
4	Dec 2008	249782401	252460802
5	Jan 2009	252460802	255139202
6	Feb 2009	255139202	257558402
7	Mar 2009	257558402	260236802
8	Apr 2009	260236802	262828802
9	May 2009	262828802	265507202
10	Jun 2009	265507202	268099202
11	Jul 2009	268099202	270777602
12	Aug 2009	270777602	273456002
13	Sep 2009	273456002	276048002
14	Oct 2009	276048002	278726402
15	Nov 2009	278726402	281318402
16	Dec 2009	281318402	283996802
17	Jan 2010	283996802	286675202
18	Feb 2010	286675202	289094402
19	Mar 2010	289094402	291772802
20	Apr 2010	291772802	294364802
21	May 2010	294364802	297043202
22	Jun 2010	297043202	299635202
25	Sep 2010	304992002	307584002
26	Oct 2010	307584002	310262402
27	Nov 2010	310262402	312854402
28	Dec 2010	312854402	315532802
29	Jan 2011	315532802	318211202
30	Feb 2011	318211202	320630402
31	Mar 2011	320630402	323308802
32	Apr 2011	323308802	325900802
33	May 2011	325900802	328579202
34	Jun 2011	328579202	331171202
35	Jul 2011	331171202	333849602
36	Aug 2011	333849602	336528002
37	Sep 2011	336528002	339120002
38	Oct 2011	339120002	341798402
39	Nov 2011	341798402	344390402
40	Dec 2011	344390402	347068802
41	Jan 2012	347068802	349747202

normalizations free in the monthly fits. These are the only free parameters related to diffuse emission in the monthly fits.

Point and Extended Sources

We fix the positions of the point and extended sources to those found in the 3-year fit. We do not expect any of our sources to exhibit detectable proper motion, and the position estimates based on the full data set should, in general, be the most accurate. It is possible that some sources could be more accurately localized based on a particular subset of the data; for example, a faint source with a bright, variable neighbor may be subject to considerable confusion in the full fit, which could be ameliorated by considering only intervals where the bright source was relatively quiescent. However, such cases are unlikely to have any major effect on the conclusions of this study, and would be better left to dedicated studies of the relevant regions.

We use the same spectral models for the point and extended sources as in the 3-year fit, but fix all parameters except the normalization. This results in an estimate of the total γ-ray flux as a function of time under the assumption that the spectral shape remains constant. Previous studies have, in fact, revealed spectral variability in blazars. For example, (Abdo et al., 2010e) reported a distinct "harder-when-brighter" trend in the flat-spectrum quasar PKS 1502+106 during a flaring state, and (Abdo et al., 2010f) reported a weak correlation in the same direction on one-week time scales for the brightest FSRQs and LSP- and ISP-BLLacs. The same results show, however, that such effects are quite small, especially relative to the overall flux variability, and that it is less apparent on longer time scales. We therefore expect that neglecting spectral variability will not significantly affect our results.

4.2.4 Single-source Profile Likelihood

The joint likelihood described in § 4.1 provides our best estimate of the flux of each of the sources in our model. In order to characterize the behavior of a single source, however, we would like to reduce the joint likelihood to a function in a single parameter describing the

flux of that source. The simplest single-parameter likelihood for the parameter λ_i would be the conditional likelihood,

$$\mathcal{L}_{cond}(\lambda_i) = L_{joint}(\lambda_i; (\lambda_j = \hat{\lambda_j} | j \neq i)), \tag{4.1}$$

that is, the joint likelihood evaluated as a function of λ_i with the rest of the parameters fixed to the values that maximize the joint likelihood. This is similar in spirit to using the diagonal elements of the covariance matrix as estimates of the parameter uncertainties, ignoring any correlations between the parameters, which will tend to underestimate the statistical variance. While this is perfectly acceptable in some contexts, in studies of variability, we are primarily interested in whether the variance in the measured source fluxes is larger than we would expect from statistical fluctuations. Since source confusion is significant in γ-ray data, especially in shorter exposures, it is important to account for correlations between sources.

To this end, we define the profile likelihood,

$$\mathcal{L}_{prof}(\lambda_i; (\lambda_j = \hat{\lambda_j}) = \sup_{j \neq i} \left[L_{joint}(\lambda_i; (\lambda_j = \lambda_j | j \neq i)) \right], \tag{4.2}$$

which is again the full joint likelihood treated as a function of a single parameter λ_i, but now maximized over the rest of the $\vec{\lambda}$ at each value of λ_i. This procedure naturally accounts accounts for the correlations between the parameters. Unfortunately, it is much more difficult to compute; for an N-parameter each evaluation of $L_{prof}(\lambda_i)$ involves an $(N-1)$-dimensional optimization. To mitigate this problem, inspired by (Tompkins, 1999), we fit the profile likelihood to a simple analytic function,

$$P(s; s_{peak}, b, E) = \left(E(s+b)^{E(s_{peak}+b)} e^{-E(s+b)} \right), \tag{4.3}$$

where s is now the single flux parameter of interest. The functional form is that of a Poisson distribution for counts observed from the sum of a signal and background, with s and b corresponding to the signal and background fluxes, E playing the role of an exposure,

and $N_{eff} = E(s_{peak} + b)$ corresponding to an effective number of observed counts. It should be noted, however, that the parameters E and b do not correspond directly to the true exposure or the fluxes of other sources in the ROI. Finding an appropriate fit function generally requires relatively few evaluations of the likelihood, after which the fit function can be used for subsequent likelihood-based calculations, making the use of the profile likelihood much less computationally intensive.

By construction, the single-source profile likelihood and the full joint likelihood for the ROI are identical at their shared maximum. The estimates of source fluxes derived from the two should therefore be identical, in principle. In practice, since the ROI models are only optimized to within some tolerance, there may be some room for minor improvements in the profile likelihood fits, and the flux estimates may show small, but not significant, differences. We estimate uncertainties in the source fluxes from the profile likelihood in the usual way, by finding the points where $\log(\mathcal{L}_{prof}(F_i))$ decreases by a .5 from the maximum. The error bars so produced will be at least as large as those estimated from the diagonal elements of the covariance matrix.

In cases where a source is not detected significantly in a particular time bin, the above uncertainty estimate provides a 1σ (i.e., 68% confidence) upper limit on the source flux. If a different confidence level for upper limits is desired, the appropriate decrease in the log likelihood can be found from the assumed $\chi^2/2$ distribution of the log likelihood. For example, a 2σ upper limit corresponds to the point where $\log(\mathcal{L}_{prof}(F_i))$ decreases by 2 from its maximum.

4.2.5 Light Curve Results

Of the 3179 sources with good localizations (elliptical quality ≤ 4), and TS ≥ 10 in the 3-year fit, 1004 (1683) are not detected with TS $\geq 9(25)$ in any of the monthly fits. Conversely, 40 (8) sources are detected with TS $\geq 9\,(25)$ in all 42 months. On average, 437 sources (14%) are detected with TS ≥ 9 in each month, and 292 (9%) with TS ≥ 25.

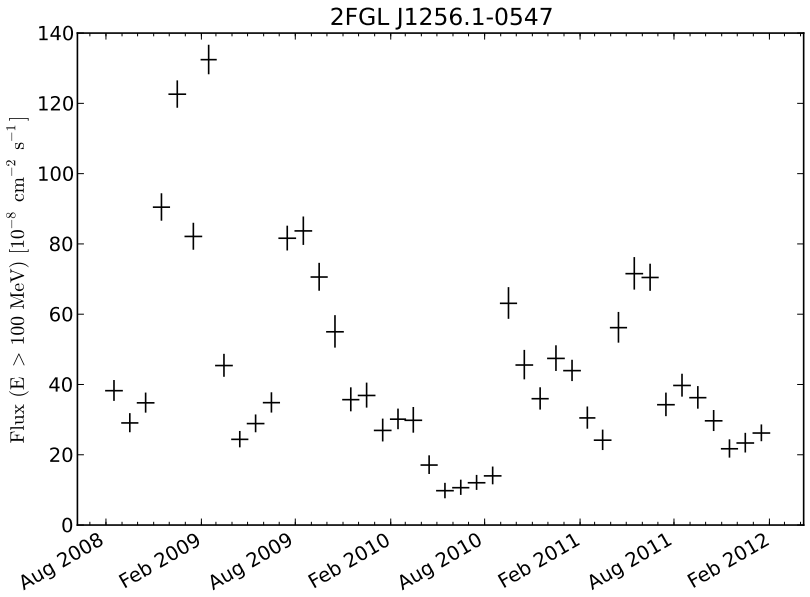

Figure 4.3 Flux as a function of time for the quasar 3C 279. The red one-sided error bars represent 1σ upper limits.

Example light curves

Figures 4.2.5, 4.2.5, and 4.2.5 show light curves for the quasar 3C 279 (2FGL J1256.1-0547), the BL Lac object Markarian 501 (2FGL J1653.9+3945), and the pulsar PSR J0030+0451 (2FGL J0030.0+0450), respectively. 3C 279 shows the dramatic flaring behavior typical of flat-spectrum radio quasars. Markarian 501 is also clearly variable, with some significant flares, but with a steadier quiescent flux, which is relatively common behavior in BL Lacs. PSR J0030+0451 displays the very steady flux characteristic of pulsar light curves.

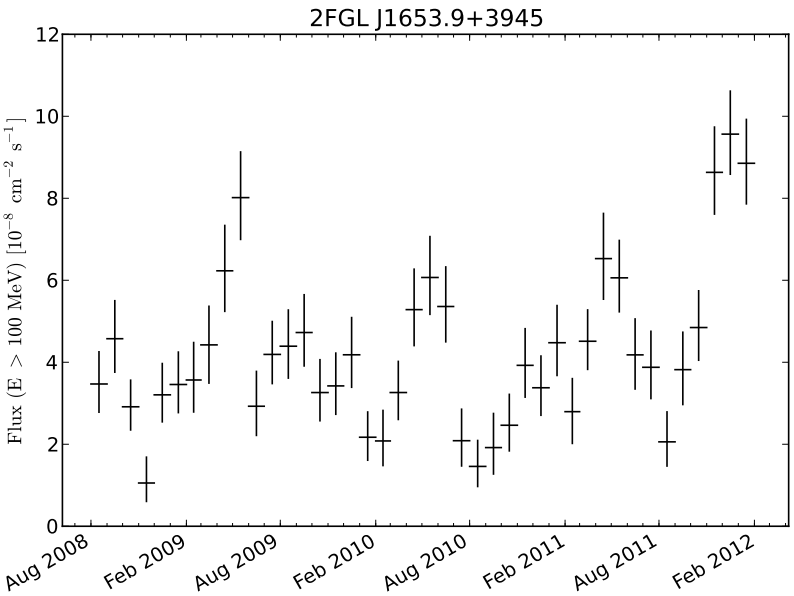

Figure 4.4 Flux as a function of time for the BL Lac object Markarian 501. The red one-sided error bars represent 1σ upper limits.

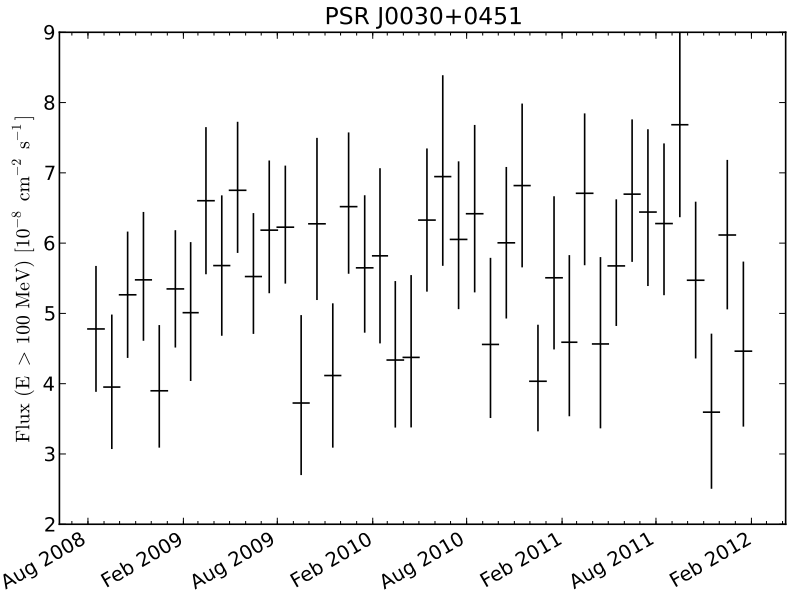

Figure 4.5 Flux as a function of time for the pulsar PSR J0030+4510. The red one-sided error bars represent 1σ upper limits.

Accumulated significance

The monthly likelihoods offer a measure of detection significance accounting for variability of the source on monthly scales. Consider the sum of TS values for each time bin,

$$\mathrm{TS}_{sum} = \sum_{i=0}^{N} 2 * \left(log(\mathcal{L}_i(F_i)) - log(\mathcal{L}_i(0)) \right), \tag{4.4}$$

where $\mathcal{L}_i$ is the likelihood function for the source flux in the ith time bin, calculated from the data in that time range, and F_i is the estimated source flux for the ith bin (i.e., the flux value that maximizes $\mathcal{L}_i$).

For a steady source against a steady background, this statistic should be approximately equal to the TS for the full time period, differing only by statistical fluctuations. On the other hand, for a variable source or a steady source against a variable background, the significance may be much higher in some time periods, producing a larger TS_{sum}. This can produce a more significant detection of sources detected with only marginal average significance over the full time range of the analysis.

In the null case (i.e., no detectable source flux in any of the monthly bins), TS_{sum} is expected to be asymptotically distributed according to Eqn. (3.20). The corresponding CDF is, similarly, a weighted sum of CDFs of χ^2 distributions, and the quantiles of the distribution are easily found from this. Formal 1σ, 3σ, and 5σ detection significances correspond to TS_{sum} ≈28.1, 48.6, and 75.5, respectively. Of the 1292 well-localized sources with TS < 25, 405 have $\mathrm{TS}_{sum} \geq 75.5$.

Figure 4.2.5 shows an example light curve of a source that does not reach the $TS > 25$ threshold in the three-year analysis, but is included based on TS_{sum}. In this case, the source is significantly detected in several months within the first two years of the light curve, and was included in the 2FGL catalog. However, it is only significantly detected in 3 of the last 20 months, resulting in a reduced average significance.

Figure 4.7 shows the distribution of the three-year mean photon flux from 100 MeV to 100 GeV for the 2652 sources with either TS ≥ 25 in the three-year analysis or $\mathrm{TS}_{sum} \geq 75.5$. The

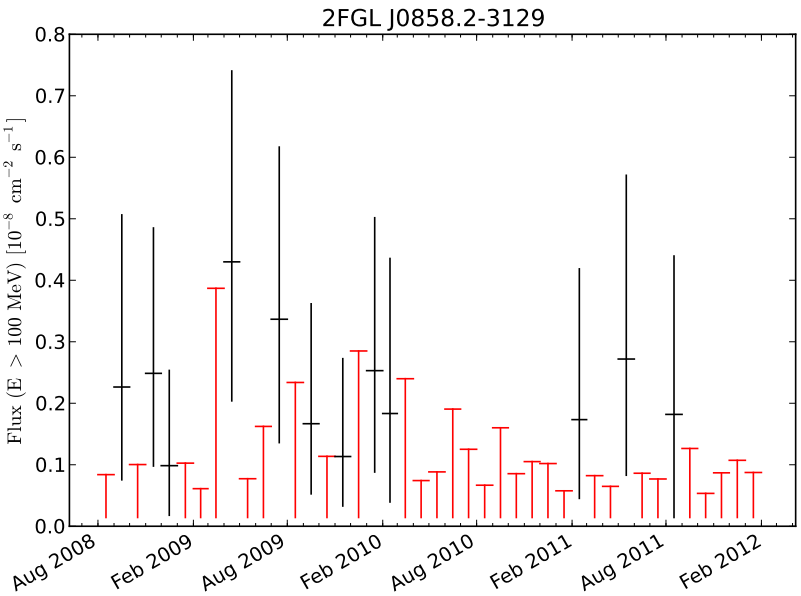

Figure 4.6 Flux as a function of time for the source 2FGL J0858.2-3129. The red one-sided error bars represent 1σ upper limits.

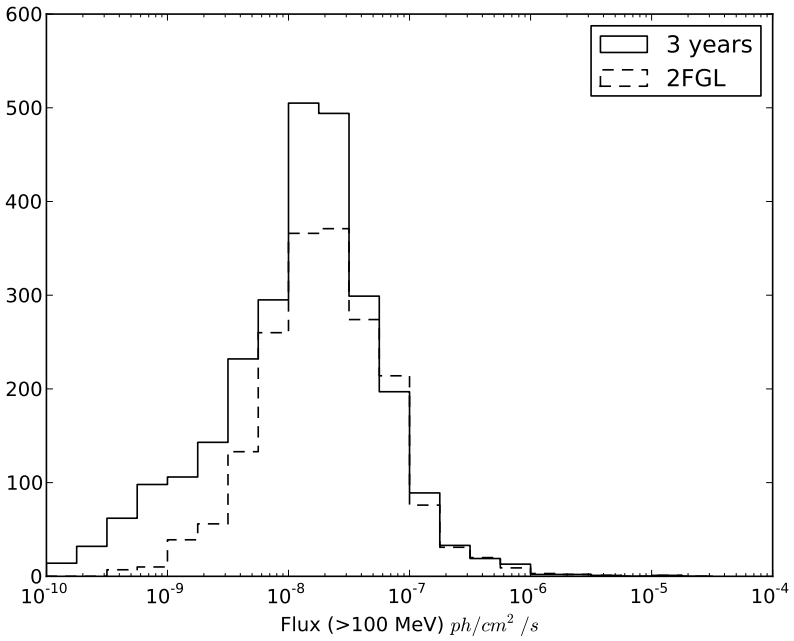

Figure 4.7 Distribution of point source flux in the range 100 MeV to 100 GeV, for sources with TS $\geq$ 25 or TS$_{sum} \geq$ 75.5 in the 3-year analysis and the 2FGL sample.

distribution of 2FGL fluxes is again shown for comparison. The most prominent difference in the histograms is the significant tail of sources at low flux in the present sample, in contrast to the relatively sharp low-flux cutoff in the 2FGL population around $10^{-8}\, photons/cm^2/s$. The sharper cutoff in the 2FGL population corresponds to the cut on detection significance, suggesting that 10^{-8} photons/cm^2/s marks a reasonably robust minimum integrated flux for significant detection of a point source. In the present sample, there is a similarly sharp drop at slightly lower flux, reflecting the improved detection sensitivity with a longer observation. The longer tail extending from approximately $(30 - .1)$x10^{-10} photons/cm^2/s consists of sources with TS $<$ 25 in the 3-year analysis, but are included based on either reaching TS $>$ 25 in at least one month, or a sum of monthly TS values TS$_{sum} \geq$ 75.

4.2.6 Light Curve Correlations

Figure 4.8 shows the cross-correlation between light curves for pairs of sources, computed as the sample covariance of the two light curves divided by the product of their standard deviations, as a function of their separation. It is clear that the flux estimates for sources within about a degree of each other are anticorrelated. This is an example of source confusion due to the large PSF, and the effect is much smaller for harder sources. At somewhat larger separations, there is a slight bias toward positive correlation, indicating that the fits to sources within several degrees are not completely independent. These correlations are likely largely due to residuals in the diffuse emission being attributed to the point sources. Since the profile likelihoods used to derive the fluxes are only optimized over the set of sources within the ROI containing the source of interest, there can also be correlations between pairs of sources straddling pixel boundaries. This is primarily a problem when one of the sources in such a pair is much brighter than the other.

4.2.7 Summary

We described the construction of 42-month light curves for a sample of 2839 sources detected in three years of data. Source detection, localization, and initial spectral fitting are performed using the full three years of data, resulting in a full-sky model of the three-year average γ-ray flux. Optimization over the full-sky is acheived by iterative fits to 1728 regions of interest. Each light curve comprises a fit to a piecewise-constant model with one-month time bins. Fits for each month are produced by the same full-sky optimization as the 3-year average fit, with the diffuse models and the spectral shape parameters of the source models fixed to their 3-year average values.

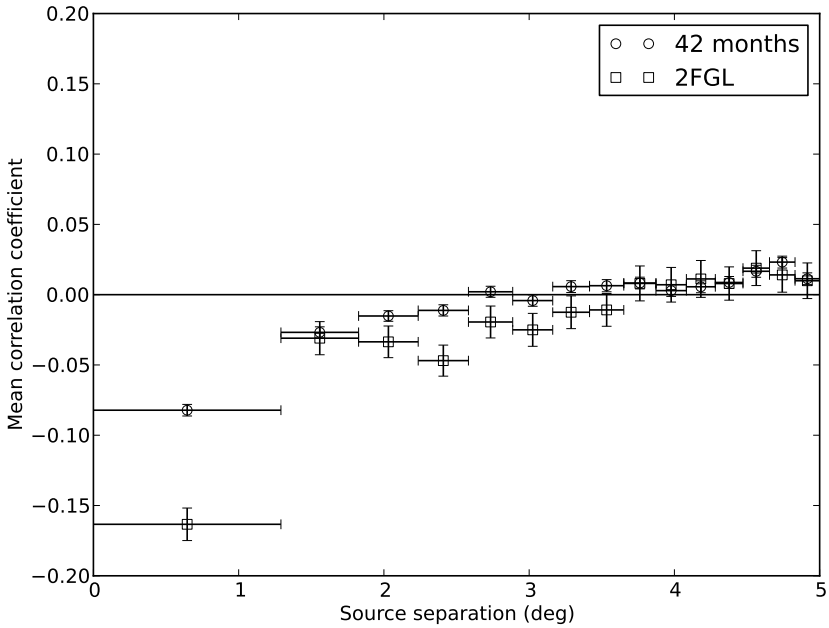

Figure 4.8 Cross-correlation of flux estimates as a function of source separation.

Chapter 5

VARIABILITY ANALYSIS

In the previous chapter, we described the construction of 42-month light curves for the the set of sources detected in 3 years of data. We now turn to a quantitative discussion of the variability of the sources in the sample. We describe a likelihood-based statistic describing the significance of observed variability in a light curve. We examine the pulsars, which are expected to be non-variable, and as calibration sources to assess the impact of unmodeled systematic error on the apparent variability of sources. We examine the observed variability as a function of mean source flux, and discuss the feasibility of detecting significant variability in faint sources.

5.1 Quantifying Variability

To construct a measure of the variability of a source, we use a likelihood ratio test of the form described in § 3.3.1. The null model is one in which the source of interest has a constant flux across the entire light curve, and the alternate model is one in which this the flux is piecewise-constant, with a different value in each time bin. The test statistic is

$$\mathrm{TS}_{var} = \sum_{i=0}^{N} 2 * \left(log(\mathcal{L}_i(F_i)) - log(\mathcal{L}_i(F_{mean})) \right), \tag{5.1}$$

with $\mathcal{L}_i$ the likelihood corresponding to bin i, F_i the best fit flux for bin i, and F_{mean} the best fit flux for the full light curve, assuming a constant flux. F_{mean} is estimated from the monthly likelihoods; while it should be comparable to the estimated mean flux from an analysis of the full time period, it is not, in general, exactly the same, due to the monthly variation of nearby sources. We expect this statistic to be distributed, in the null case, as $\chi^2_{N-1}(\mathrm{TS}_{var})$. A formal 3σ significance for variability then corresponds to $\mathrm{TS}_{var} \gtrsim 73.5$, and

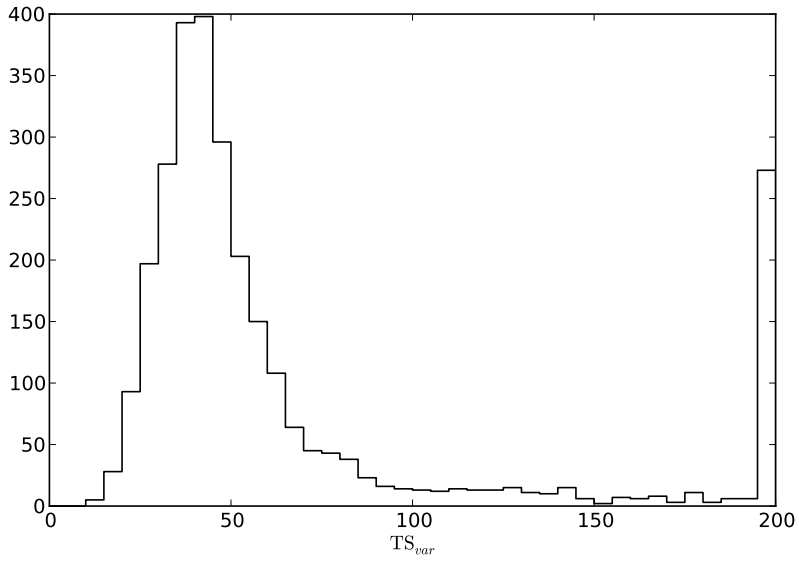

Figure 5.1 Distribution of TS_{var} for all sources. Sources with $TS_{var} > 200$ are included in the last bin.

5σ to $TS_{var} \gtrsim 103$. Figure 5.1 shows the distribution of TS_{var} for all sources with either $TS \geq 25$ in the 3-year analysis or $TS_{sum} \geq 75.5$ (2652 sources in all). 585 sources have $TS_{var} \geq 73.5$, and 437 have $TS_{var} \geq 103$.

5.1.1 Flux-dependence

Figure 5.2 shows the variability statistic TS_{var} as a function of average source flux for the 2475 sources in our sample. At high flux, the sources clearly divide into distinct high- and low-variability populations. The bright, low-variability sources are mostly pulsars, which will be discussed in more detail below. Among the significantly variable sources, TS_{var} tends to increase with increasing flux. This is primarily a reflection of the Poisson statistics governing the emission. The fractional uncertainty in the flux decreases approximately as the square root of the flux, so that a doubling of the flux of a brighter source contributes more to the

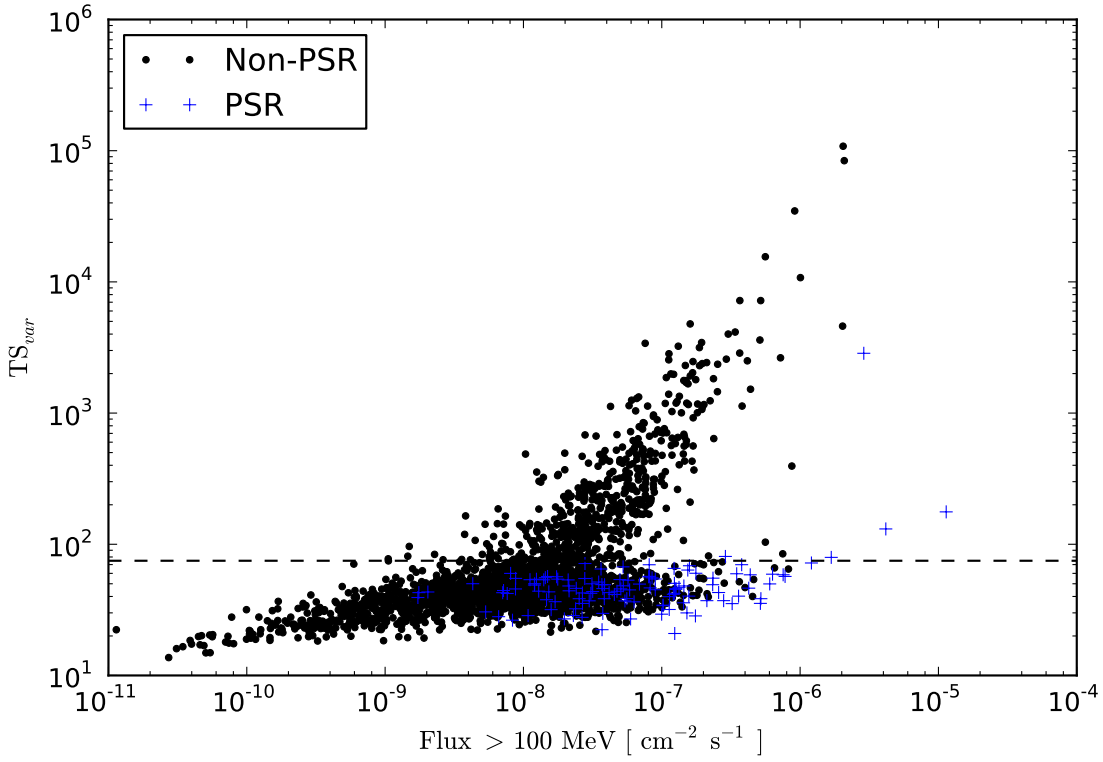

Figure 5.2 Variability index as a function of mean integrated flux.

statistical significance of the variability than would the same doubling in a fainter source.

The fraction of sources identified as significantly variable decreases sharply at lower mean flux. As the source flux decreases, fewer and fewer of the monthly fluxes surpass the threshold for significant detection, with the light curves containing only upper limits on the fluxes in those bins. In those cases, the likelihood differences between the mean flux and the estimated monthly fluxes are quite small; that is, downward variations in the flux tend to be less significant than upward variations, and the overall significance of the variability of the source is reduced. Very few sources with mean flux below about $10^{-8} photons/cm^2/s$ show detectable variability.

5.1.2 Pulsars

In general, we expect the γ-ray flux from pulsars, averaged over many pulsation periods, to be steady. This expectation is borne out by past observations, with a couple of notable exceptions. Several large flares have been observed in the flux from the vicinity of the Crab pulsar, J0534+2200 (Abdo et al., 2011b). The largest of them is dramatically visible in the monthly light curve shown in Figure 5.3. The γ-ray flares from the Crab are believed to originate from the pulsar wind nebula, rather than the pulsar itself; the Crab is unique among LAT-detected pulsars in that the γ-ray flux from the pulsar and its associated pulsar wind nebula are of comparable magnitude, and the two cannot be spatially resolved by the LAT. On the other hand, a sharp drop in the observed flux from the radio-quiet pulsar J2021+4026 in October, 2011 is robustly associated with the pulsar itself; the flux change is seen in the pulsed emission, and is coincident with a timing glitch (i.e., a sudden change in the pulsation period)(Allafort et al., 2013). While such events are a possible source of true variability in pulsars, they seem to be quite rare, and there is no evidence for any similar flux changes associated with other observed timing glitches in LAT-detected pulsars (Dormody & Fermi Collaboration, 2011). The flux change in PSR J2021+4026, illustrated in Figure 5.4, is only detected at about 3σ in this analysis, and the source is not flagged as significantly variable ($\mathrm{TS}_{var} = 72$).

Figure 5.5 shows the distribution of TS_{var} for only the known pulsars in the sample (those for which pulsations have been observed in γ rays). Though the distribution is skewed somewhat to higher values than the predicted χ^2 distribution with 41 degrees of freedom, the relatively narrow distribution at low TS_{var} generally supports the picture of pulsars as non-variable sources. There are three outliers with $\mathrm{TS}_{var} > 103$. One is the Crab, which is in fact variable, as explained above. The other two are the Vela and Geminga pulsars, which are the brightest pulsars in the sample other than the Crab.

A trend of increasing formal variability with increasing mean flux is apparent for the pulsars in Figure 5.2, and Vela and Geminga represent the extreme of that trend. The

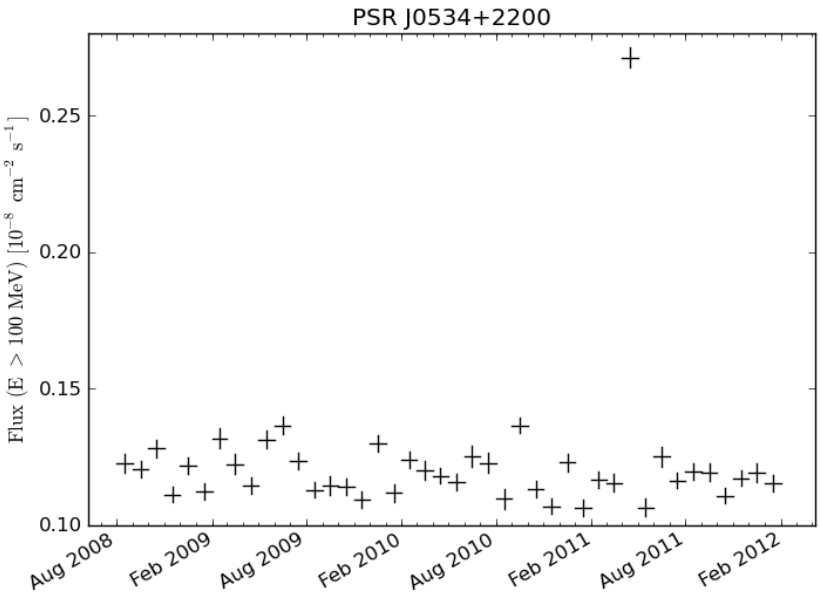

Figure 5.3 Integral flux above 100 MeV from the Crab as a function of time, in 1-month time bins.

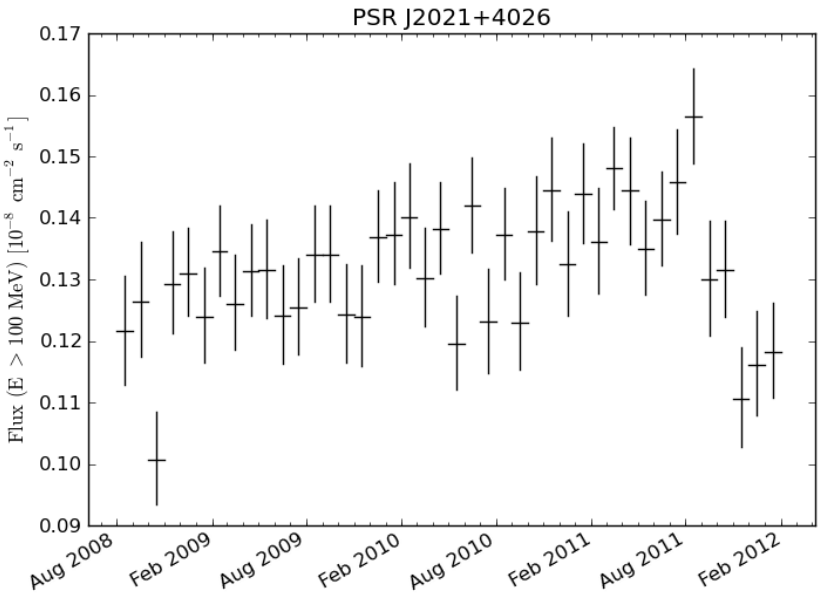

Figure 5.4 Integral flux above 100 MeV from PSR J2021+4026 as a function of time, in 1-month time bins.

light curves for the Vela and Geminga pulsars are shown in Figure 5.6. Both show greater scatter than would be expected based on the uncertainties in their monthly fluxes, but no prominent flares or shifts in flux. This suggests that systematic errors not accounted for in the statistical uncertainties derived from the likelihood make a significant contribution to the relative errors in the flux measurements of these sources, whereas for fainter sources, the statistical errors dominate. Under the assumption that the pulsars are indeed steady, the formal variability of the brightest of them offers a rough measure of the importance of such systematics. Based on visual inspection of Figure 5.2, we estimate that the observed variation is dominated by statistics for source fluxes below $\sim 10^{-6} photons/cm^2/s$.

One way of accounting for the contribution of systematics to TS_{var} for brighter sources is to scale the monthly contributions to TS_{var} by a factor depending on the mean source flux,

$$\mathrm{TS}_{var} = \sum_{i=0}^{N} 2 * (log(\mathcal{L}_i(F_i)) - log(\mathcal{L}_i(F_{mean}))) \left(\frac{\Delta F_i^2}{\Delta F_i^2 + f^2 F_{mean}^2} \right). \tag{5.2}$$

Eqn. (5.2) with $f = .02$, defines the variability index used in the 2FGL. This is approximately equivalent to adding a systematic uncertainty component of 2% of the mean source flux in quadrature with the statistical uncertainty estimated from the likelihood. Applying the same weighting to our sample reduces TS_{var} to below the nominal 3σ threshold, as well, suggesting that the systematic error contribution in this analysis is comparable to that in the 2FGL analysis. The weighting does not significantly affect the variability estimates for fainter or significantly variable sources; only one non-pulsar source drops below the 3σ threshold (from $\mathrm{TS}_{var} = 73.54$ to $\mathrm{TS}_{var} = 73.40$) when applying the weighting.

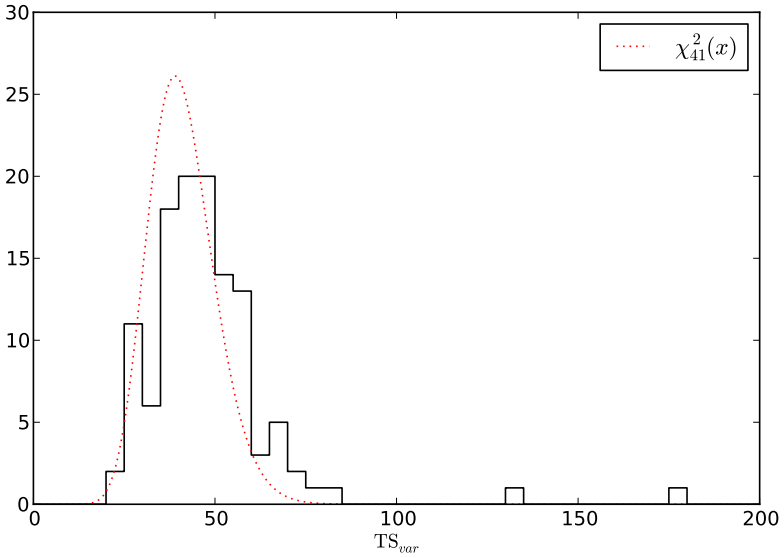

Figure 5.5 Distribution of TS_{var} for pulsars only. Sources with $TS_{var} > 200$ are included in the last bin. The dashed curve is a χ^2 distribution with 41 degrees of freedom.

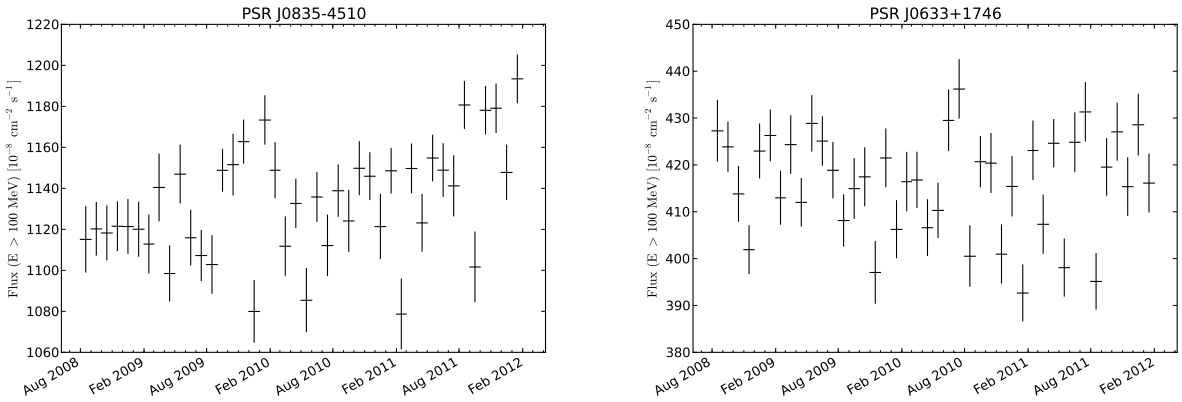

Figure 5.6 Flux as a function of time for the pulsars J0835-4510 (Vela, left) and J0633+1746 (Geminga, right).

Chapter 6

CONCLUSIONS

Temporal variability is very common in γ-ray sources, and understanding and characterizing their behavior in the time-domain is an important aspect of the effort to understand them. Analysis of the variability of the γ-ray flux from individual sources provides insight into the physical processes responsible for their emission. Correlations of γ-ray variability with observations at other wavelengths allows the identification of γ-ray sources with multiwavelength counterparts, and augments our picture of the broadband behavior of such objects. Time-resolved analysis also allows the detection of variable and transient sources that would not be significantly detected in the integrated data set. Furthermore, a time-dependent of the γ-ray sky potentially allows the discrimination of variable sources from backgrounds and nearby point sources, allowing for more accurate characteristization of all sources, variable or not.

The amount of data produced by the *Fermi* Large Area Telescope, and the large number of sources it has allowed us to detect, present considerable challenges for population-level analyses compared to earlier instruments. Extending such analyses to the time domain compounds those challenges. Large-scale analyses of the time-domain behavior of the γ-ray sky are necessarily lacking in detail compared to time-averaged analyses, due to both the additional computation required for time-resolved analysis and the reduced statistics available in the shorter time ranges involved. Such studies are therefore generally restricted in complexity, providing coarse detail for a large set of sources, such as those presented in the *Fermi*-LAT source catalogs; in scope, examining a smaller set of the brightest sources in more detail, such as those presented in the LAT AGN catalogs; or some combination.

The analysis presented in this work represents an effort to mitigate those restrictions,

producing an analysis broad in scope, including the full set of known γ-ray sources, while improving on the level of detail in typical analyses of similar scope. The limitation of low statistics is outside of our control, but the computational challenges are more surmountable, and that is our focus. We demonstrated an efficient framework for all-sky variability analysis, producing an analysis of comparable detail and accuracy to the equivalent analysis from the 2FGL catalog, extended to an expanded source list derived from 3 years of data.

In chapter 3, we discussed the analysis of LAT data. We motivated the use of likelihood methods for LAT analysis and described the essential likelihood computations. We then presented the *pointlike* package. *pointlike* provides a framework for efficient likelihood analysis of LAT data, and provides the basis for the variability analysis presented here. The improvement in efficiency relative to the standard LAT collaboration analysis tools is achieved by taking advantage of our understanding of the instrument response to simplify the likelihood calculations by discarding information that only provides computational precision below the resolution of the detector (e.g., binning data at lower resolution at energies where the point-spread function is large). We described the approximations employed by *pointlike* to achieve this improved efficiency, and discussed the accuracy of the resulting calculations relative to the more formally rigorous implementation of the standard tools.

In chapter 4, we described the application of pointlike to the study of variability. The source list for the analysis was constructed using the same basic procedure for source detection and localization used in the construction of the 2FGL catalog, applied to three years of data. The resulting list comprised 3539 sources, with 2247 meeting the significance threshold used for inclusion in the 2FGL analysis. The properties of this sample are generally comparable to those of the 2FGL, with a lower flux threshold for detection as expected based on the longer observation time.

We then constructed light curves for this source list by applying the same fundamental all-sky fit to each time bin. The models for these fits were simplified by leaving the positions and spectral shapes of the sources fixed. We produced light curves in 42 one-month bins, extending the full time range somewhat beyond that used in the construction of the source

list. The monthly fits were performed for the full list of candidate sources, with no cut on the detection significance in the 3-year fit. This allowed us to construct a separate statistic for detection significance based on the accumulated significance in the monthly fits, which gives a measure of detection significance accounting for flux variations in both the source of interest and any nearby sources. Accepting sources that fail the usual significance cut in the 3-year analysis, but pass a comparable cut on accumulated significance from the light curves expands the final source sample by 405 sources.

In chapter 5, we discussed the quantification of variability. We presented a likelihood-based statistic to assess the significance of variability in a light curve, and applied it to the light curves described in the previous chapter. Approximately 20% of the sources in our sample showed significant variability. The significance of observed variability depends strongly on the mean flux of the source. For our analysis, we found that the fraction of sources flagged as significantly variable dropped off sharply for mean fluxes below about $10^{-8} photons/cm^2/s$. This limit likey corresponds to the threshold for significant detection in an individual time bin.

We discussed the use of pulsars, which are expected to have steady fluxes on the timescales we are considering, as calibration sources to assess the reliability of our variability statistic and the contribution of systematic error to the observed significance of variability. We concluded that the formal significance of variability is insensitive to systematic error except for the brightest pulsars.

The analysis presented here is simple to extend to longer time ranges, and should be practical to update to include as much data as the telescope can provide. Beyond simply extending the time range of the analysis, improving the time resolution is also straightforward; halving the size of the time bins will result in the same increase in computation time as doubling the length of the light curves at the same binning, *ceteris paribus*. However, there is a limit to the time resolution at which this analysis will remain useful. As the size of he time bins decreases, the threshold for significant detection on those time scales increases. Already at monthly time scales we saw that the availibility of useful variability information

was quite limited for fluxes much below $10^{-8} photons/cm^2/s$. At significantly shorter time scales, it may be preferable to restrict the analysis to the brighter sources.

Another interesting extension of the analysis would be to search for transient sources. We saw that accounting for variability allowed for significant detections of what would otherwise have been marginal sources. However, those marginal detections were still limited to signals detectable, at least faintly, when averaged over the full data set. With all-sky models for each time bin in hand, repeating the source detection procedure for each would likely turn up yet more sources either too faint, or active for too short a period, to be detectable in the time-averaged procedure. Such a search would potentially extend the already impressive discovery potential of the LAT even further.

BIBLIOGRAPHY

Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009a, ApJ, 700, 597

—. 2009b, Phys. Rev. D, 80, 122004

Abdo, A. A., Ackermann, M., Atwood, W. B., et al. 2009c, ApJ, 696, 1084

Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009d, ApJS, 183, 46

—. 2009e, Astroparticle Physics, 32, 193

—. 2010a, ApJS, 188, 405

—. 2010b, ApJ, 713, 146

—. 2010c, ApJ, 720, 272

—. 2010d, ApJ, 722, 520

—. 2010e, ApJ, 710, 810

—. 2010f, ApJ, 710, 1271

—. 2010g, ApJ, 715, 429

Abdo, A. A., Ackermann, M., Agudo, I., et al. 2010h, ApJ, 716, 30

Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010i, ApJ, 713, 154

—. 2011a, ApJ, 736, 131

—. 2011b, Science, 331, 739

—. 2012, ApJ, 758, 140

Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, The Astrophysical Journal Supplement Series, 208, 17

Ackermann, M., Ajello, M., Allafort, A., et al. 2011a, Science, 334, 1103

Ackermann, M., Ajello, M., Baldini, L., et al. 2011b, ApJ, 726, 35

Ackermann, M., Ajello, M., Allafort, A., et al. 2011c, ApJ, 743, 171

Ackermann, M., Ajello, M., Albert, A., et al. 2012, ApJS, 203, 4

Allafort, A., Baldini, L., Ballet, J., et al. 2013, ApJ, 777, L2

Allison, J., Amako, K., Apostolakis, J., et al. 2006, IEEE Transactions on Nuclear Science, 53, 270

Atwood, W. B., Bagagli, R., Baldini, L., et al. 2007, Astroparticle Physics, 28, 422

Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, ApJ, 697, 1071

Baldini, L., Barbiellini, G., Bellazzini, R., et al. 2007, AIP Conference Proceedings, 921, 190

Bignami, G. F., Boella, G., Burger, J. J., et al. 1975, Space Science Instrumentation, 1, 245

Blandford, R. D., & Ostriker, J. P. 1978, ApJ, 221, L29

Boinee, P., Cabras, G., de Angelis, A., et al. 2003, in Science with the New Generation of High Energy Gamma-Ray Experiments : Between Astrophysics and Astroparticle Physics, ed. S. Ciprini, A. de Angelis, P. Lubrano, & O. Mansutti, 141

Byrd, R., Lu, P., Nocedal, J., & Zhu, C. 1995, SIAM Journal on Scientific Computing, 16, 1190

Chernoff, H. 1954, Annals of Mathematical Statistics, 25, 573

Dermer, C. D., Cavadini, M., Razzaque, S., et al. 2011, ApJ, 733, L21+

Dormody, M., & Fermi Collaboration. 2011, in Bulletin of the American Astronomical Society, Vol. 43, American Astronomical Society Meeting Abstracts #217, 434.05

Fichtel, C. E., Hartman, R. C., Kniffen, D. A., et al. 1975a, ApJ, 198, 163

Fichtel, C. E., Kniffen, D. A., Thompson, D. J., & Ogelman, H. B. 1975b, in International Cosmic Ray Conference, Vol. 1, International Cosmic Ray Conference, 106–111

Fossati, G., Maraschi, L., Celotti, A., Comastri, A., & Ghisellini, G. 1998, MNRAS, 299, 433

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759

Grove, J. E., & Johnson, W. N. 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7732, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series

Healey, S. E., Romani, R. W., Cotter, G., et al. 2008, ApJS, 175, 97

Hermsen, W., Bloemen, J. B. G. M., Bennett, K., et al. 1981, in International Cosmic Ray Conference, Vol. 1, International Cosmic Ray Conference, 230–233

Jones, E., Oliphant, T., Peterson, P., et al. 2001–, SciPy: Open source scientific tools for Python, [Online; accessed 2015-04-13]

Kanbach, G., Bertsch, D. L., Fichtel, C. E., et al. 1988, Space Sci. Rev., 49, 69

Kerr, M. 2010, Likelihood methods for the detection and characterization of gamma-ray pulsars with the Fermi large area telescope, copyright - Copyright ProQuest, UMI Dissertations Publishing 2010; Last updated - 2014-01-08; First page - n/a

Kirsch, M. G. F., Altieri, B., Chen, B., et al. 2004, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 5488, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. G. Hasinger & M. J. L. Turner, 103–114

Klebesadel, R. W., Strong, I. B., & Olson, R. A. 1973, ApJ, 182, L85

Lande, J., Ackermann, M., Allafort, A., et al. 2012, ApJ, 756, 5

Lott, B., Escande, L., Larsson, S., & Ballet, J. 2012, A&A, 544, A6

Mattox, J. R., Bertsch, D. L., Chiang, J., et al. 1996, ApJ, 461, 396

Moiseev, A. A., Hartman, R. C., Ormes, J. F., et al. 2007, Astroparticle Physics, 27, 339

Neyman, J., & Pearson, E. S. 1928, Biometrika, 20A, 175

Nolan, P. L., Tompkins, W. F., Grenier, I. A., & Michelson, P. F. 2003, ApJ, 597, 615

Nolan, P. L., Abdo, A. A., Ackermann, M., et al. 2012, ApJS, 199, 31

Oegelman, H., Fichtel, C. E., & Kniffen, D. A. 1975, Nature, 255, 208

Olive, K.A., e. a. 2014, Chin. Phys., C38, 090001

Pollock, A. M. T., Masnou, J. L., Bignami, G. F., et al. 1981, A&A, 94, 116

Pollock, A. M. T., Bennett, K., Bignami, G. F., et al. 1985, A&A, 146, 352

Protassov, R., van Dyk, D. A., Connors, A., Kashyap, V. L., & Siemiginowska, A. 2002, The Astrophysical Journal, 571, 545

Rando, R. 2009, arXiv:0907.0626, arXiv:0907.0626

Read, A. M., Rosen, S. R., Saxton, R. D., & Ramirez, J. 2011, A&A, 534, A34

Roth, M. 2012, A Likelihood Method for Determining the On-orbit Point-Spread Function of the Fermi Large-Area Telescope, copyright - Copyright ProQuest, UMI Dissertations Publishing 2012; Last updated - 2014-01-19; First page - n/a

Strong, A. W., Moskalenko, I. V., & Ptuskin, V. S. 2007, Annual Review of Nuclear and Particle Science, 57, 285

The Fermi-LAT collaboration. 2010, ArXiv e-prints, arXiv:1007.0483

Tibaldo, L. 2010, in American Institute of Physics Conference Series, Vol. 1223, American Institute of Physics Conference Series, ed. C. Cecchi, S. Ciprini, P. Lubrano, & G. Tosti, 89–98

Tompkins, W. F. 1999, Applications of likelihood analysis in gamma-ray astrophysics, copyright - Copyright UMI - Dissertations Publishing 1999; Last updated - 2015-05-06; First page - n/a

Urry, C. M., & Padovani, P. 1995, Publications of the Astronomical Society of the Pacific, 107, 803

Wilks, S. S. 1938, Annals of Mathematical Statistics, 9, 60