

Chapter 1

Event Reconstruction and Source Analysis Methods for Pair-Conversion Telescopes

Eric Charles and James Chiang

*Kavli Institute for Particle Astrophysics and Cosmology, SLAC National
Accelerator Laboratory, Stanford University, Stanford, CA 94305, USA*

The analysis of the pair production data is unique in astronomy and is critically linked with the details of the instrumental hardware. This chapter describes several aspects of event reconstruction and source analysis in pair-conversion telescopes. This includes the methods used to reconstruct the details of individual γ ray interactions; the parametric representations of instrument performance; and the likelihood fitting based analysis techniques used to analyze astrophysical sources from γ -ray event lists. Finally, we describe astrophysical data sets that can be used for calibration and validation of the instrument performance.

1. Introduction

As discussed in Chap. 0, in pair-conversion telescopes individual γ rays convert to e^+e^- pairs, which are recorded by the instrument [1–3]. By reconstructing the e^+e^- pair we can deduce the energy and direction of the incident γ ray. Accordingly, data analysis is event-based: we record and analyze each incident particle separately. Each event consists of a readout of all of the signals deposited in the instrument during a narrow time window, typically $\mathcal{O}(\mu\text{s})$.^a It is worth noting that cosmic-ray rates are high enough that some fraction of events will contain signals from more than one cosmic ray, or from both a γ ray and a cosmic ray.

The field-of-view of pair-conversion telescopes is huge compared to most observatories operating at other wavelengths; typically exceeding a steradian (e.g., the *Fermi*-LAT field-of-view is $\sim 2.4\text{sr}$ at 10 GeV). Pair-conversion telescopes are normally deployed in low-earth orbit with a period of ~ 90 minutes; coupled with the large field-of-view, this makes them excellent sky monitors. In fact, rather than pointing at a series of fixed targets the *Fermi*-LAT is usually operated in a sky-survey mode where it continuously scans across the sky.

In this chapter we will first describe some of the methods used to identify the incoming γ rays and estimate their energy and direction (Sec. 2), discuss the in-

^aDetector technologies are capable of much finer time resolution, down to $\mathcal{O}(\text{ns})$ timescales. However, the power budget available for space-based telescopes significantly limits the achievable time granularity.

strument performance using the *Fermi*-LAT as an example (Sec. 3), describe techniques used to make astronomical measurements from lists of γ ray events (Sec. 4) and finally we will touch briefly on astronomical data samples used to calibrate and validate the instrumental response (Sec. 5).

2. Reconstruction of Pair-Conversion Events

Event reconstruction translates the raw event information from the instrument subsystems into a high-level event description; see Fig. 1 for an illustrative event display of a γ ray in the *Fermi*-LAT [3, 4].

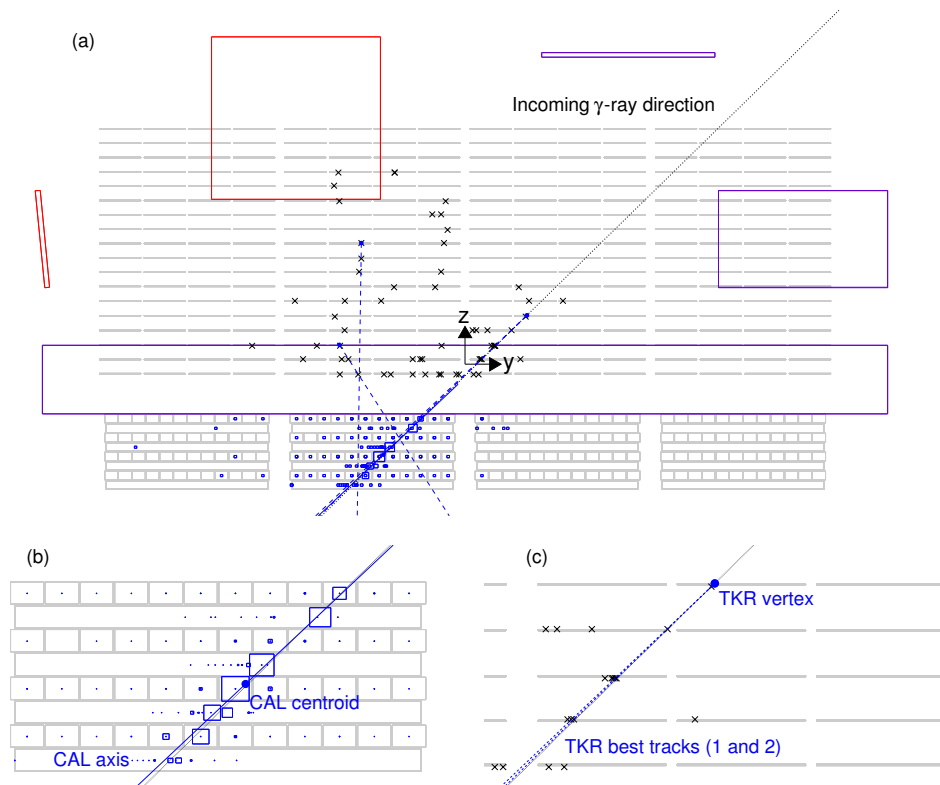


Fig. 1. Event display of a simulated 27 GeV γ ray (a) and zoom over the calorimeter (b) and tracker (c) portions of the event. The small crosses represent the clusters in the tracker (see Sec. 2.1), while the variable-size squares indicate the reconstructed location and magnitude of the energy deposition for every hit crystal in the calorimeter. The dotted line represents the true γ -ray direction, the solid line is the reconstructed calorimeter cluster axis and the dashed lines are the reconstructed tracks. The backscatter from the calorimeter generates tens of hits in the tracker, with two spurious tracks reconstructed in addition to the two associated with the γ ray (note that they extrapolate away from the calorimeter cluster centroid and do not match the calorimeter cluster axis direction). It also generates a few hits in the anti-coincidence detector, which, however, are away from the event direction extrapolation and therefore do not contradict the classification of the event as a γ ray.

Reconstructing the signals in the individual detector channels into a coherent picture of a particle interaction for each of the several hundred events collected every second by a pair conversion telescope is a formidable task. The basic steps of the process are:

- (1) *Digitization*: converting the information about signals in individual channels from the schema used in the electronics readout to more physically motivated schema, including the location of the signal in an instrument-based coordinate system.
- (2) *Event Reconstruction*: applying pattern recognition and fitting algorithms commonly used in high-energy particle physics experiments to reconstruct the event in terms of individual tracks and energy clusters in the detector subsystems and to associate those objects with each other. This step includes separating signals the different particles in multi-particle events.
- (3) *Event analysis*: evaluating quantities that can be used as figures of merit for the event from the collections of tracks, clusters and associated information. Once this information is extracted, multivariate analysis techniques are applied to extract measurements of the energy and direction of the event and to construct estimators that the event is in fact a γ -ray interaction rather than a background cosmic-ray interaction.
- (4) *Event classification*: applying selection criteria to make lists of γ ray events.

While the digitization stage is fairly straightforward, the remaining steps are involved and described further below.

2.1. Event Reconstruction

A wealth of techniques have been developed in high-energy particle physics to reconstruct the momenta and energy of particles observed in a detector (see, e.g., Ref. [5] for an overview of the subject). Furthermore, the specific details of the algorithms depend on the geometry of the detector in question. Therefore, rather than discuss these algorithms in detail we will summarize the overall process.

In principle, near the pair-conversion point the event consists only of the two tracks from electron and positron. (At energies above a few GeV, depending on the granularity of the tracker, the two tracks overlap and become indistinguishable). Within about one radiation length along the direction of travel of the incident γ ray, Bremsstrahlung emission and secondary pair conversions produce the cascade of e^+ , e^- and γ -rays making the EM shower. As discussed in Chap. 0, identifying and correctly reconstructing the e^+ and e^- tracks before the onset of the EM shower is an essential goal of the event reconstruction.

The “clustering” stage of event reconstruction is to combine signals in adjacent detector elements that are likely to have been caused by the same incident particle. In the Silicon strip detectors used in the *Fermi*-LAT and the *AGILE*-GRID, a charged particle often leave signals in two or more adjacent strips in a single layer. The clustering algorithms combines these signals into a set of clusters that will be

used in the later stages of the reconstruction.

Clustering algorithms for signals in calorimeters are different; since the EM showers are fully developed in the calorimeter the goal of the calorimeter clustering algorithm is to gather together all of the signals from a single shower, rather than from the individual tracks in the shower. Calorimeter clustering is important to disentangle the signals caused by incoming γ rays from those caused by cosmic rays that crossed the instrument during the event readout window. The clustering stages often include the application of channel-by-channel calibration algorithms, so as to measure the energy deposition in each cluster.

With highly segmented calorimeters, such as the *Fermi*-LAT calorimeter, and to lesser extent, the *AGILE* Mini-Calorimeter, a dedicated “shower-fitting” algorithm (see, e.g., Refs. [6, 7]) can extract information such as: estimates of the total energy deposited in the cluster as well as the energy that leaked out of the calorimeter, topology of the clusters, e.g., the axes of the calorimeter clusters and the transverse and longitudinal extent of the showers with respect to those axes. Topological information is useful for discriminating between γ rays and cosmic rays, in particular baryons such as protons and heavy nuclei. Baryons do not generate electromagnetic (EM) showers, but rather deposit energy by a combination of ionization and nuclear interactions, resulting in significantly different shower topologies than γ rays, electrons or positrons.

In the “track-finding” stage of the event analysis, clusters in the tracker are linked together into tracks representing the path of individual particles. Far too many track-finding algorithms exist to be discussed here; four different algorithms are used in the reconstruction of *Fermi*-LAT data alone. Probably the most important points to make are that no single track-finding algorithm is a panacea, and different algorithms have often have complementary strengths. E.g., one algorithm may be very fast, and excellent for picking out the simple and straight tracks often left by high-energy cosmic rays, while another may be significantly slower, but better at disentangling the more complex events caused by γ rays with their EM showers.

Once individual tracks have been found, Kalman filter-based “track-fitting” algorithms are used to extract estimates of the direction of the individual particles [8–10]. The Kalman filter technique successively updates the best-estimate track parameters and uncertainty estimates to account for the information gained by each hit along the track trajectory and the information lost as the particles experiences multiple scattering in the detector material. The output of the Kalman-filter track fit is a series of best-fit track parameters and associated covariance matrices, one set for each signal observed or volume of material traversed by the tracked particle. Each set of parameters gives the best-fit estimate of the track trajectory, using information up to that point in the track. In practical terms, however, the track parameters at the beginning of the track are the most important, as they can be used both to determine the direction of the incoming particle and to propagate

the track to test for the existence of associated signals in other detector sub-systems (e.g., an associated hit in an anti-coincidence veto system that might indicate the track was most likely caused by a charged particle starting outside the tracking volume).

The final stage of the event reconstruction is to associate information from the various instrument sub-systems. For example, tracks can be extrapolated to the calorimeter and matched up with calorimeter clusters, or extrapolated to a surrounding anti-coincidence veto detector and associated with signals there.

In fact, the description given here is something of an oversimplification. Information from various stages of event reconstruction can be useful in other stages. E.g., for very complicated events, typically with thousands of individual signals in the tracker, the axis of the EM shower in the calorimeter can be used to seed the track-finding algorithms, to reduce the combinatorial challenge to a manageable level. Typically the later stages of the event reconstruction have some amount of overlap, and part or all of the process might be iterated more than once, often using information from previous iterations. Also, the tracks and calorimeter clusters from the different particles in multi-particle events must be separated during the event reconstruction.

2.2. Event Analysis

Once each event has been fully reconstructed, the reconstructed event must be analyzed to extract information that will be used to determine the energy, direction and species of the incoming particle. It is also useful to estimate how well the energy and direction were measured, e.g., to select sub-samples of events that are particularly well reconstructed, or to reject events that are poorly reconstructed.

It is possible to design literally hundreds of different quantities to characterize each event. Here we will list the most important questions these quantities are designed to answer.

- (1) *Which track or tracks came from the original e^+e^- pair?* In practice this is a matter of finding the *longest* and *straightest* tracks, as those are the ones from the highest energy particles. These are most likely to be from the original pair, as opposed to back-scattered particles from the EM shower. Furthermore, if the two longest and straightest tracks from a vertex, i.e., if they come from a single point, that is a strong indicator that the event is in fact from a γ -ray pair conversion.
- (2) *How accurate are the estimates of the energy and direction?* Events for which a large fraction of the event energy leaked out of the calorimeter naturally have poorer energy resolution than events for which the energy was well-contained. Likewise, high-energy events for which the pair-conversion occurred near the top of the tracker and signals were observed in every tracker plane naturally provide a more accurate direction estimate than events for which the pair conversion happened at the bottom of the track and only three or four track planes have

signals.

- (3) *Does the event start well inside the instrument?* Incoming cosmic rays will leave signals from the point at which they enter the instrument; for converted e^+e^- pairs the signals will only start after the conversion point. If a likely conversion point can be identified, quantifying the number of detector signals “missing” before that point along the track extrapolation is a powerful discriminator between γ rays and cosmic rays.
- (4) *Is the topology of the event consistent with an EM shower?* “Shower-shape variables” distinguish between EM showers caused by γ rays and the hadronic showers caused by cosmic rays proton and heavy ions. However these do not discriminate against cosmic-ray electrons and positrons. Similarly, the specific ionization, i.e., the energy deposition per unit length in the instrument, scales as the square of the particle charge, making it a powerful discriminator against heavy ions. On the other hand, because of Bremsstrahlung processes, EM shower events tend to have a larger number of extra signals in the tracker near the e^+e^- pair.
- (5) *Do the instrument sub-systems provide a consistent picture of the event?* Because of the huge level of background rejection needed, low-probability events can pose a significant analysis challenge. This is especially true as these low-probability events might partially mimic a γ ray event. A specific example would be a cosmic-ray electron that enters the side of the calorimeter, creating an EM shower, from which a single particle escapes into the tracker and leaves a short track before running out of energy. The tracker reconstruction might mis-identify the event as a track that started in the middle of the tracker. The calorimeter reconstruction might well identify the event as an EM shower. However, the axis of the calorimeter shower would probably not match the track direction particularly well. Only by putting together the entire event would it be identified as a likely cosmic ray interaction.

2.3. Event Classification

Since pair-conversion telescopes are subjected to a flux of cosmic rays orders of magnitude larger than the γ -ray flux, it is not feasible to design a single set of event selection criteria that remove all the cosmic-ray backgrounds while preserving efficiency for selecting the γ rays. In fact, optimization with sophisticated machine learning algorithms is required to reach the signal to noise levels needed for astronomy.

It is also important to note that the optimal point of the background rejection versus signal efficiency curve (also known as the receiver operation characteristic or ROC curve) depends on the scientific measurement being made. Accordingly, instrument teams often provide a small number of different sets of γ -ray event lists, made with more or less stringent background-rejection criteria. E.g., in each of their data releases the *Fermi*-LAT collaboration has provided separate selections optimized for the study of: (i) short time-scale transients such as gamma-ray bursts,

(ii) persistent point sources, (iii) large scale diffuse emission and (iv) all-sky studies of the isotropic gamma-ray background.^b Similarly, the *AGILE* collaboration provides one data selection for most analyses and a looser selection for the analysis of gamma-ray bursts and pulsar timing [2].

3. Instrument Response Functions

The standard for analysis of pair-conversion telescope data is the maximum likelihood analysis formalism as described in Ref. [11]. A critical component of this formalism is the parametric representation of instrument performance: the instrument response functions (IRFs). In practice, analyses assume that the IRFs can be factorized into three parts:

- (1) *Effective Area*, $A_{\text{eff}}(E, \hat{v}, s)$, the product of the cross-sectional geometrical collection area, γ ray conversion probability, and the efficiency of a given event selection (denoted by s) for a γ ray with energy E and direction \hat{v} , in the instrument frame;
- (2) *Point-spread Function* (PSF), $P(\hat{v}'; E, \hat{v}, s)$, the probability density to reconstruct an incident direction \hat{v}' for a γ ray with (E, \hat{v}) in the event selection s ;
- (3) *Energy Dispersion*, $D(E'; E, \hat{v}, s)$, the probability density to measure an event energy E' for a γ ray with (E, \hat{v}) in the event selection s .

Note that the IRFs can change markedly across the instrument field-of-view. In practice, because of their high level of azimuthal symmetry, existing pair-conversion telescopes have parametrized the IRFs in terms of the incidence angle with respect to the bore-sight, θ , and either ignored the azimuthal dependence of the IRFs or treated it as a small correction factor.

Generally, IRFs are created by making Monte Carlo simulations of millions of γ rays incident from a variety of directions, locations and energies, fitting the resulting distributions to parametrized functions at each energy and incidence angle, and storing tables of the fit parameters.

It is useful to define the IRFs in the instrument frame when discussing instrument performance; but when analyzing celestial γ -ray data it is more practical to work in sky-coordinates, and we must therefore account for the pointing history of the instrument. So we denote the IRFs in celestial coordinates, \hat{p} and in terms of a generalized vector $\vec{L}(t)$ describing, as a function of time, the instrument attitude and other relevant degrees of freedom, such as the instrument mode, or if the instrument is even taking data (“live”). The IRFs are then:

$$\begin{aligned} A_{\text{eff}} &= A_{\text{eff}}(E, \hat{p}, \vec{L}(t), s); \\ P &= P(\hat{p}'; E, \hat{p}, \vec{L}(t), s); \\ D &= D(E'; E, \hat{p}, \vec{L}(t), s). \end{aligned} \tag{1}$$

^bSee, e.g., http://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation/Cicerone/Cicerone_Data_Exploration/Data_preparation.html.

Clearly, this implies that the “live”-time pointing history of the instrument must be tracked and saved.

Because the instrument response has to be evaluated over the constantly changing orbit and attitude of the telescope, experience has shown that it is useful to pre-compute and store the distribution of observing time in the instrument reference frame of any given direction in the sky. For a single direction this is referred to as the “observing profile” or “off-axis histograms”, and written $t_{\text{obs}}(\hat{v}; \hat{p})$; and we refer to a collection of observing profiles for a pixelization of the entire sky as a “livetime cube”.

The exposure in a given direction can be calculated by integrating the product of the observing profile and the effective area over the instrument field-of-view:

$$\mathcal{E}(E, \hat{p}, s) = \int d\hat{v} A_{\text{eff}}(E, \hat{v}, s) t_{\text{obs}}(\hat{v}; \hat{p}). \quad (2)$$

Figure 2 shows examples of the observing profile for two different sky locations and a map of the exposure over the entire sky at 1 GeV.

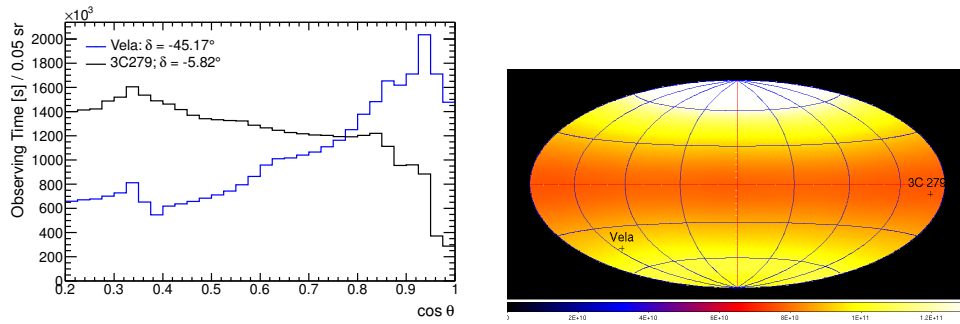


Fig. 2. Observing profiles and exposure at 1 GeV for 5 years (2008 Aug. 03 to 2013 Aug. 03) of *Fermi*-LAT observations. Left: observing profiles towards the Vela pulsar and the blazar 3C279. Right: exposure at 1 GeV in cm^2s for the P7REP_SOURCE_V15 IRFs; shown in a Hammer-Aitoff projection in equatorial coordinates along with the locations of Vela and 3C 279. The difference in observing profile between the two sources is a consequence of the different declinations (δ) of those sources.

3.1. Effective Area

When describing the performance of a pair-conversion telescope, we commonly show the effective area at normal incidence as a function of the energy and the angular dependence of the effective area for a given energy. Figure 3 shows these performance curves for the *Fermi*-LAT.

In the middle of the energy band the effective area is primarily determined by the geometrical cross section of the instrument and the pair-conversion efficiency, and is fairly uniform with energy. At low energies the effective area falls off because the converted electron and positron may not have enough energy to leave a long

Analysis of Pair-Conversion Telescope Data

9

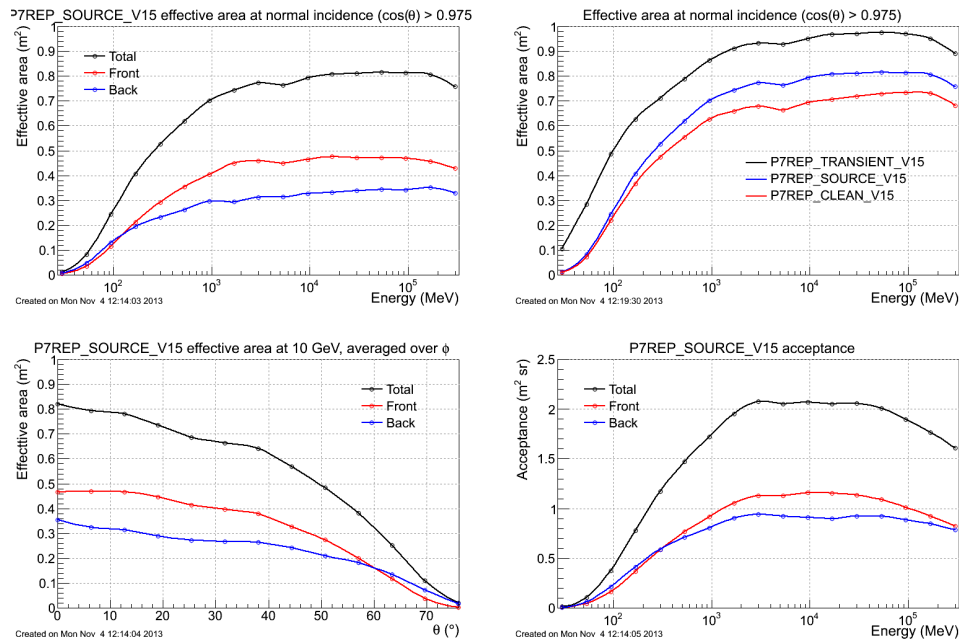


Fig. 3. Effective area and acceptance of the *Fermi*-LAT, for the P7REP_SOURCE version of the event selections and the associated P7REP_SOURCE_V15 IRFs. Top left: A_{eff} for γ -rays entering normal to the orientation of the tracker planes. Top right: comparison of the A_{eff} of the P7REP_SOURCE event selection and two other selections. Bottom left: variation of the A_{eff} with incidence angle (θ) for 10 GeV γ rays. Bottom right: acceptance, Eq. (3), as a function of energy. “Front” and “back” refer to events that convert in the front (first 12 converter/sensor layers) and back (last 6 layers) sections of the LAT tracker. The Tungsten converters in the back layers are thicker and induce more γ -ray conversions, resulting in roughly the same total acceptance as the front section with fewer converter layers.

enough track to be properly reconstructed. At high energies the effective area falls off because some fraction of the huge number of particles created in the EM shower in the calorimeter scatter back into the tracking volume and complicate the event to the point that it becomes exceedingly difficult to properly reconstruct.

The incidence angle dependence of the effective area is primarily a geometric effect; as the incidence angle increases, the chance of a particle leaving enough signals in both the tracker and calorimeter such that both the direction and energy can be reconstructed decreases. The variation of the effective area with incidence angle would be quite different for instruments that did not have the stacked planar geometry of existing pair-conversion telescopes.

When discussing instrument performance, it is often useful to integrate over the field-of-view to define the “acceptance” of the instrument:

$$A(E, s) = \int d\hat{v} A_{\text{eff}}(E, \hat{v}, s). \quad (3)$$

The acceptance is plotted in the lower-right panel of Figure 3.

3.2. Point-Spread Function

As discussed in Chap. 0, at low energies the PSF of a pair-conversion telescope is determined by multiple scattering, while at high energies it is determined by the spatial resolution of the individual signals in the tracker and the distance between successive signals. Furthermore, the increasing complexity of higher energy events makes it more likely that the track-finding algorithms will associate one or more signals with the wrong track. In general this results in widening the tails of the PSF without much change to the core of the distribution. Figure 4 shows two measures of the PSF as a function of energy for the *Fermi*-LAT.

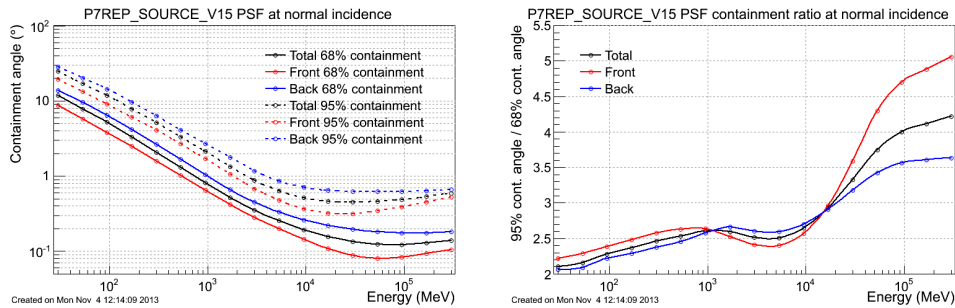


Fig. 4. Containment radius of the point-spread function of the *Fermi*-LAT for the P7REP_SOURCE_V15 IRFs as a function of energy for γ -rays entering at normal incidence. Left: 68% and 95% containment radii. Right: ratio of the 95% to 68% containment radii. The expected ratio for a two-dimensional Gaussian is 1.61; i.e., at high energies the tails of the PSF are very non Gaussian. The thicker Tungsten converters in the back section result in more multiple scattering of the tracks and worse direction resolution.

For instruments with planar geometries such as the *Fermi*-LAT or the *AGILE*-GRID the PSF varies with the incidence angle of the incoming γ ray; at larger incidence angles the particles will cross more material (and experience more multiple scattering) per tracking plane. For a particular direction in the sky, the time- and observing profile-averaged effective PSF is:

$$\bar{P}(\hat{p}'; E, \hat{p}, s) = \frac{\int d\hat{v} P(\hat{p}'; E, \hat{p}, \hat{v}, s) A_{\text{eff}}(E, \hat{v}, s) t_{\text{obs}}(\hat{v}; \hat{p})}{\int d\hat{v} A_{\text{eff}}(E, \hat{v}, s) t_{\text{obs}}(\hat{v}; \hat{p})}. \quad (4)$$

3.3. Energy Dispersion

As discussed in Chap. 0, the energy resolution of existing pair-conversion telescopes is determined by how much of the energy of the incident γ ray is deposited in the calorimeter. At low energies the energy resolution degrades because a large fraction of the energy is typically lost in the tracker, before the particles reach the calorimeter; at very high energies the energy resolution degrades because of energy leakage out the back and sides of the calorimeter. Figure 5 shows the energy

resolution^c of the *Fermi*-LAT as function of energy and incidence angle.




Fig. 5. Energy resolution of *Fermi*-LAT as a function of Energy and incidence angle. Left: energy resolution as of function of energy for γ rays arriving at normal incidence. Right: energy resolution for 10 GeV γ rays as a function incidence angle.

For specific analyses where the energy resolution is important, it can be useful to define time- and observing profile-averaged effective energy dispersion for a source of interest, in analogy to Eq. (4):

$$\bar{D}(E'; E, \hat{p}, s) = \frac{\int d\hat{v} D(E'; E, \hat{v}, s) A_{\text{eff}}(E, \hat{v}, s) t_{\text{obs}}(\hat{v}; \hat{p})}{\int d\hat{v} A_{\text{eff}}(E, \hat{v}, s) t_{\text{obs}}(\hat{v}; \hat{p})}. \quad (5)$$

4. Likelihood Formalism

Our knowledge of the GeV sky is quite limited. As of 2015 only ~ 3000 γ -ray sources have been detected [12]. In general, the limited statistics of the available data only allow us to answer simple questions about any particular source. Often we can measure only the total flux and overall spectral index. For brighter sources we can measure spectral curvature, temporal variability or spatial extension if they are pronounced enough.

Furthermore, the single event-based data acquisition, broad PSF and the typical sky-survey operating mode make the analysis of pair-conversion telescope data quite different from the analysis of data from many other types of instruments.

The maximum likelihood technique has been seen as the way to get the most precise estimates of the relatively small number of source parameters we can measure with the available data. Maximum likelihood fitting was first used to analyze pair-conversion telescope data in the COS-B era [13]. The methodology has been refined and adapted first by the EGRET [11], and then the *Fermi*-LAT and *AGILE* collaborations.

This likelihood analysis formalism compares the observed distributions of γ rays, $N(\hat{p}', E', t, s)$, to an expected distribution, $M(\hat{p}', E', t, s)$, that is created by con-

^cDefined as the half width of the energy window containing 34% + 34% (i.e., 68%) of the energy dispersion on both sides of its most probable value, which typically gives slightly larger values of energy resolution than using the *smallest* 68% containment window.

volving source models, $S(\hat{p}, E, t; \vec{\alpha})$, where $\vec{\alpha}$ denotes the free parameters of the model, with the IRFs. Specifically, the likelihood model is

$$\begin{aligned} M(E', \hat{p}', t, s) &= \int dE d\hat{p} A_{\text{eff}}(E, \hat{p}, \vec{L}(t), s) P(\hat{p}'; E, \hat{p}, \vec{L}(t), s) \times \\ &\quad D(E'; E, \hat{p}, \vec{L}(t), s) S(\hat{p}, E, t; \vec{\alpha}) \\ &= \int dE d\hat{p} R(E', \hat{p}'; E, \hat{p}, t, s) S(\hat{p}, E, t; \vec{\alpha}), \end{aligned} \quad (6)$$

where the latter relation defines R , the ‘‘total response’’ of the instrument. The fitting procedure then finds the values, $\hat{\alpha}$, of the fit parameters that maximize the likelihood, $\mathcal{L}(M(\vec{\alpha}); N)$.

When testing the significance of additional model components or parameters, a test-statistic (TS) can be defined in terms of the likelihood ratio of the best-fit model with the additional component included with respect to the null hypothesis (i.e., the best fit model with those components omitted):

$$TS = 2 \log \frac{\mathcal{L}(M(\hat{\alpha}); N)}{\mathcal{L}_{\text{null}}(M(\hat{\alpha}); N)}. \quad (7)$$

In most analyses of γ -ray data the statistics are sufficient that the TS is well-described by χ^2 distribution with degrees of freedom equal to the number of additional parameters in the test-hypothesis model with respect to the null-hypothesis model (i.e., Wilks' [14] or Chernoff's [15] theorem applies).

The likelihood ratio test can be used in a number of ways, e.g., to estimate the significance of a source candidate at a particular location, to establish the presence of spectral curvature or additional spectral components, or to test for source variability.

We will first discuss the formalism used to construct the likelihood for binned data (Sec. 4.1), then by considering the limit of small bins, we can move to the formalism for likelihood analysis of unbinned data (Sec. 4.2). In both cases we will discuss approximations that can be used to simplify and vastly speed up the analysis.

4.1. Binned Likelihood Analysis

For binned data the Poisson likelihood is:

$$\mathcal{L} = \prod_j \frac{m_j^{n_j} e^{-m_j}}{n_j!}, \quad (8)$$

where n_j is observed number of events in bin j and m_j is the number of events predicted to lie within the bin given the model. In practice, the negative log-likelihood is minimized with respect to the model parameters, $\vec{\alpha}$:

$$-\log \mathcal{L} = - \sum_j (n_j \log m_j - m_j - \log n_j!) \quad (9)$$

$$= - \sum_j n_j \log m_j - N_{\text{pred}}. \quad (10)$$

Here the $\log n_j!$ term is neglected since it does not depend on the model parameters, and we have defined

$$N_{\text{pred}} \equiv \sum_j m_j \quad (11)$$

which is the total number of counts predicted by the model.

Typically, data are binned both in the apparent direction, \hat{p}' , and apparent energy, E' ; in some analyses they are also binned in time. Most analyses only use γ rays for a relatively small part of the sky. This region of interest is dictated by the size of the source of interest and the PSF and is typically of $\mathcal{O}(10^\circ)$. To avoid excessive loss of information from binning, the bin size must be smaller than the PSF, typically $\mathcal{O}(0.1^\circ)$, resulting in binned data images of roughly $\mathcal{O}(100 \times 100)$ pixels for each energy bin.

The general expression for the expected counts in bin j is:

$$\begin{aligned} m_j &= \int_j M(E', \hat{p}', t, s) \\ &= \int_j dE' d\hat{p}' \int dt \int_{\text{SR}} dEd\hat{p}R(E', \hat{p}'; E, \hat{p}, t, s)S(\hat{p}, E, t; \vec{\alpha}), \end{aligned} \quad (12)$$

where SR denotes the “source region”, i.e., the region of the sky for which sources contribute to the region of interest. Because of the finite PSF, the source region must be larger than the region of interest to account for contributions from sources near the edge of the region of interest. Figure 6 shows example counts and model maps for two different energy ranges for a region of interest centered on the blazar 3C 279. For analyses of all but the smallest regions of the sky for very short time scales it is impractical to actually perform all these integrals. Fortunately, in most cases a number of simplifying approximations can be found.

Constant sources and time-integrated IRFs. So long as both the source flux and the IRFs are constant, the precomputed livetime cube can be used to calculate contributions to the total response for any direction, e.g., the exposure, Eq. (2), or the average PSF, Eq. (4). Because of the paucity of photons, non-periodic time-dependence in form of flares or longer time scale secular changes in flux or spectral shape will be better characterized by performing constant-source likelihood analyses using sub-intervals of the data rather than trying to parametrize and fit a time-dependent model. Similarly, discrete, time-dependent, changes in the IRFs are best handled by treating the sub-intervals as separate observations.

Neglecting energy dispersion. In general, the physical processes responsible for astrophysical γ -ray emission do not produce narrow spectral features. Furthermore, the IRFs for pair-conversion telescopes vary slowly as compared to the typical energy resolution of 10%. Therefore, in many analyses neglecting the energy dispersion results in negligible biases. This vastly speeds up the analysis by avoiding the




Fig. 6. Counts (left) and model (right) maps for a $14^\circ \times 14^\circ$ region centered on the blazar 3C 279. The maps were made using 5 years of *Fermi*-LAT data from the P7REP_SOURCE event selection and the associated P7REP_SOURCE_V15 IRFs. All of the maps use 0.1° pixels and are smoothed with a 0.3° Gaussian kernel. The model maps were built using sources from the 2FGL catalog [16], including some sources located just outside the region of interest. Top: maps for the 1 GeV to 1.78 GeV energy range. Bottom: maps for the 10 GeV to 17.8 GeV energy range.

convolution integrals over the true energy. Technically, this is equivalent to using a delta-function to represent the energy dispersion $D(E'; E, \hat{p}, \vec{L}(t), s) = \delta(E' - E)$. We will refer only to true energies, E , henceforth.

Applying these two approximations makes it possible to express the expected counts in bin j as a convolution of the source model with the exposure and average PSF:

$$m_j = \int_j dE d\hat{p}' \int_{\text{SR}} d\hat{p} \mathcal{E}(E, \hat{p}, s) \bar{P}(\hat{p}'; \hat{p}, E, s) S(\hat{p}, E, t; \vec{\alpha}). \quad (13)$$

Factorization of spatial and spectral dependence in source models. Up to this point we have treated the source model as a single entity. In fact, the source model is typically a combination of many point sources (primarily pulsars and AGN), contributions from Galactic diffuse γ -ray emission, isotropic extragalactic emission, and a few spatially extended sources, such as supernova remnants or nearby galaxies. It is often very useful to compute the distribution of expected counts for each source component (indexed by i) separately and combine the expected counts distributions:

$$\begin{aligned} m_j &= \sum_i m_{ij} \\ &= \sum_i \int_j dE d\hat{p}' \int_{\text{SR}} d\hat{p} \mathcal{E}(E, \hat{p}, s) \bar{P}(\hat{p}'; \hat{p}, E, s) S_i(\hat{p}, E, t; \vec{\alpha}_i). \end{aligned} \quad (14)$$

In many cases it is possible to factor the source model for a given source into a spatial part, $\tilde{S}_i(\hat{p})$, and a spectral part, $s_i(E; \vec{\alpha}_i)$. This is explicitly true for point sources, where $\tilde{S}_i(\hat{p}) = \delta(\hat{p})$, but is also the case when fitting to a spatial template obtained from the analysis of data at other wavelengths. In fact, as the notation here suggests, in many cases the spatial description of the source is assumed to be known, and only the spectral parameters are being fit.

For point sources the expected counts distribution is simply the spectral term convolved with the average PSF at the source location, \tilde{p} . For sources with extensions of $\mathcal{O}(1^\circ)$ the average PSF changes little across the source, so using any location, \tilde{p}_{nom} within the source will result in negligible bias and speed up the procedure dramatically by making it possible to invoke the convolution theorem to perform the spatial integrals:

$$m_{ij,\text{pt}} = \int_j dE d\hat{p}' \mathcal{E}(E, \tilde{p}) \bar{P}(\hat{p}'; \tilde{p}, E) s_i(E; \vec{\alpha}_i); \quad (15)$$

$$m_{ij,\text{ex}} = \int_j dE d\hat{p}' \int_{\text{SR}} d\hat{p} \mathcal{E}(E, \tilde{p}) \bar{P}(\hat{p}'; \tilde{p}_{\text{nom}}, E) s_i(E; \vec{\alpha}_i) \tilde{S}_i(\hat{p}). \quad (16)$$

If only the spectral parameters are being fit it is worthwhile to group every other term in Eq. (15) into an energy dependent spectral prefactor, $d_{ij}(E)$, and precompute those values:

$$d_{ij,\text{pt}}(E) = \int_{\Delta\hat{p}'_j} d\hat{p}' \mathcal{E}(E, \tilde{p}) \bar{P}(\hat{p}'; \tilde{p}, E); \quad (17)$$

$$d_{ij,\text{ex}}(E) = \int_{\Delta\hat{p}'_j} d\hat{p}' \int_{\text{SR}} d\hat{p} \mathcal{E}(E, \tilde{p}) \bar{P}(\hat{p}'; \tilde{p}_{\text{nom}}, E) \tilde{S}_i(\hat{p}). \quad (18)$$

In both cases the expected counts are simply

$$m_{ij} = \int_{\Delta E_j} dE d_{ij}(E) s_i(E; \vec{\alpha}_i). \quad (19)$$

Maps consisting of the precomputed $d_{ij}(E)$ values for each source in the model for a particular set of IRFs are often referred to as “source maps”. In fact, they are

a combination of the spatial component of the source models and the instrumental response.

4.2. Unbinned Likelihood Analysis

As stated earlier, the formalism for unbinned likelihood analysis can be obtained by considering the limit of very small bins, such that any bin either has zero or one event. In that case, the log of the likelihood is

$$\log \mathcal{L} = \sum_j \log M(E', \hat{p}', t, s) - N_{\text{pred}}, \quad (20)$$

where the index j now runs over all the bins with one count (cf. eq. 10); or, equivalently, over all of the counts. The total number of predicted γ ray counts is

$$N_{\text{pred}} = \int dE' d\hat{p}' M(E', \hat{p}', t, s). \quad (21)$$

N_{pred} does not depend on the data, and must be calculated by numerical integration. Effectively this is equivalent to summing the binned model counts over the region and energy bands of interest, and the approximations discussed in Sec. 4.1 may be used (eqs. 10 and 11). Additionally, if a source is well contained in the region of interest, or if the entire sky is being analyzed, the spatial convolution may be dispensed with entirely, since only the total number of expected counts is needed, rather than the distribution of those counts.

When evaluating the sum, $\sum_j \log M(E', \hat{p}', t)$, it is impractical to convolve the complete instrument response with every source in the model for every observed γ ray, so simplifying approximations are required. These approximations are similar to, but slightly different than the those used for the analysis of binned data and are discussed below.

Constant sources and time-integrated IRFs. As was the case for binned data, is generally more practical to perform constant-source likelihood analyses using sub-intervals of the data rather than trying to parametrize and fit a time-dependent model. Powerful techniques, such as the Bayesian Blocks method [17, 18], can be used with unbinned time data to identify likely sub-intervals of constant flux.

Neglecting energy dispersion. Again, neglecting the energy dispersion results in negligible biases for many measurements.

Applying these two approximations makes it possible to express the model likelihood to observe each γ ray, indexed by j , as

$$M_j(E, \hat{p}', t, s) = \int d\hat{p} A_{\text{eff}}(E, \hat{p}, \vec{L}(t), s) P(\hat{p}'; E, \hat{p}, \vec{L}(t), s) S(\hat{p}, E; \vec{\alpha}). \quad (22)$$

Note that the effective area and the PSF are to be evaluated given the state and orientation of the instrument *at the time the j^{th} γ ray was recorded.*

Factorization of spatial and spectral dependence in source models. For the analysis of unbinned data, we obtain these expressions for the model likelihood to observe a given γ ray from point source and extended sources, respectively:

$$M_{ij,\text{pt}} = A_{\text{eff}}(E, \tilde{p}, \vec{L}(t), s)P(\hat{p}'; E, \tilde{p}, \vec{L}(t), s)s_i(E; \vec{\alpha}_i); \quad (23)$$

$$M_{ij,\text{ex}} = \int_{\text{SR}} d\hat{p}A_{\text{eff}}(E, \tilde{p}, \vec{L}(t), s)P(\hat{p}'; E, \tilde{p}, \vec{L}(t), s)s_i(E; \vec{\alpha}_i)\tilde{S}_i(\hat{p}). \quad (24)$$

When only the spectral parameters of a given source are free in a fit, it is useful to precompute the spectral prefactors, d_{ij} for each γ ray in the sample:

$$d_{ij,\text{pt}} = A_{\text{eff}}(E, \tilde{p}, \vec{L}(t), s)P(\hat{p}'; E, \tilde{p}, \vec{L}(t), s); \quad (25)$$

$$d_{ij,\text{ex}} = \int_{\text{SR}} d\hat{p}A_{\text{eff}}(E, \tilde{p}, \vec{L}(t), s)P(\hat{p}'; E, \tilde{p}, \vec{L}(t), s)\tilde{S}_i(\hat{p}). \quad (26)$$

Precomputing the “diffuse response”. Most studies focus on individual γ -ray sources. However, the d_{ij} terms will be the same for every analysis that uses the same data selection, IRFs and source models. Clearly, it would be impractical to compute and save the d_{ij} for every combination of point source and γ ray. However, a model of large-scale diffuse emission is needed for every analysis. More specifically, analyses of individual point sources typically use pre-existing models of the Galactic diffuse [19] and isotropic [20] γ -ray emission (note that the isotropic spectrum also includes a residual cosmic-ray component). Accordingly, it can be useful to compute the $d_{\text{gal},j}$ and $d_{\text{iso},j}$ for the most commonly used IRFs for each γ ray in the data sample, and to include those values as part of the data release.

5. Calibration and Validation Control Data

Because of the complexity of pair-conversion telescopes and of the physics simulations of particle interactions in them we cannot expect Monte Carlo simulations to perfectly reproduce the flight data. For this reason it is important to calibrate and validate the IRFs using flight data. Detailed descriptions of calibration and validation processes have been published by several instrument teams [4, 21–24]. Here we will briefly describe the most useful data sets for on-orbit calibration and validation studies.

Although no astrophysical source has perfectly known properties, in practice there are several sources for which accurate background subtraction allows extracting a clean γ -ray sample that can be used to validate the Monte Carlo predictions. Figure 7 shows examples of the definitions of the signal and background regions used for background subtraction from four such samples. Taken together, the four samples span the energy band accessible to pair-conversion telescopes. The samples are described in more detail in the rest of this section.

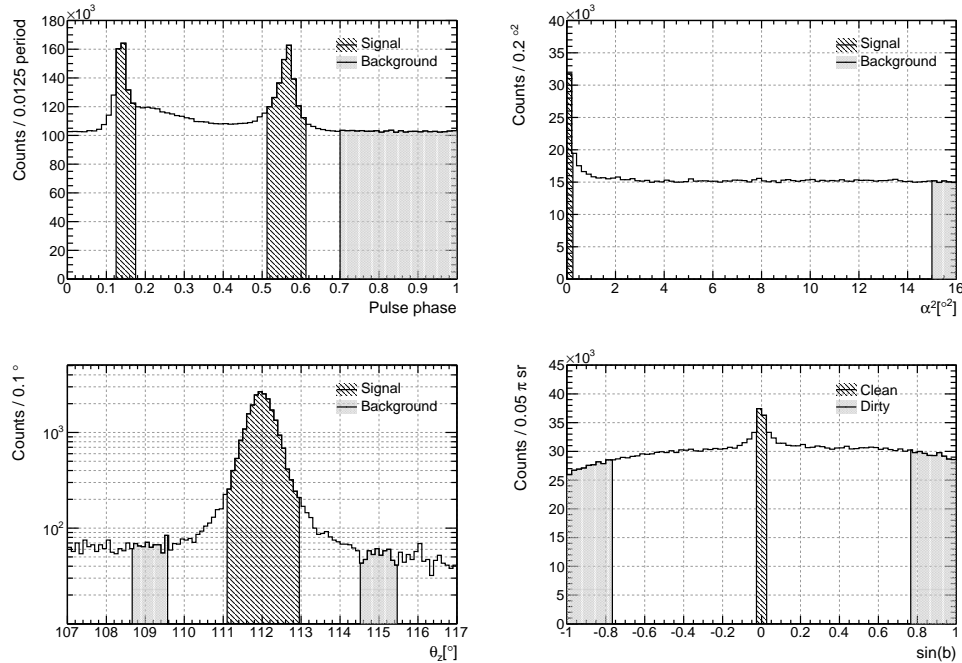


Fig. 7. Examples of background subtraction methods used to define pure γ -ray calibration samples. In each case signal and background regions in the discriminating variable are shown. Top left, pulse phase-selection for the Vela pulsar; top right, square of the angular separation, α , from the nearest AGN; bottom left, zenith-angle, θ_z , selection to isolate the Earth's Limb; bottom right, distribution of Galactic latitudes for γ rays above 17.8 GeV. The figures are illustrative only and use looser event selections than typical analyses, details of the selection criteria, the data sets used and the definitions of the signal and background regions are given in Ref. [4].

5.1. Bright Pulsars

The pulsed high-energy γ -ray emission from pulsars allows for almost perfect control of the background subtraction. The very short ($< \text{ms}$) timescales of pulse substructures imply emission regions of $< \mathcal{O}(300 \text{ km})$, making them essentially perfect point sources. Furthermore, the Vela pulsar (PSR J0835–4510) has the largest integral flux $> 100 \text{ MeV}$ of any γ -ray source, and the Crab (PSR J0534+2200) and Geminga (PSR J0633+1746) pulsars are also extremely bright γ -ray sources [16]. These factors make pulsars very good calibration sources. Unfortunately, the pulsed emission of pulsars typically cuts off above several GeV, and pulsations are nearly undetectable above 30 GeV by pair-conversion telescopes such as the *Fermi*-LAT [25–27].

5.2. Bright Active Galactic Nuclei

Given the density of bright γ -ray sources in the sky, at energies where the 95% containment radius of the PSF is less than $\sim 1^\circ$ we can use the angular distance

between a γ ray and the nearest source as a good discriminator for background subtraction, particularly at high Galactic latitudes where there are fewer sources and the interstellar diffuse emission is less pronounced. Unfortunately, no single source is bright enough to provide adequate statistics to serve as a good calibrator. However, by considering γ rays from a sample of bright and/or hard spectrum Active Galactic Nuclei (AGN) that are isolated from other hard sources it is possible to make good calibration sample for energies up to ~ 30 GeV. It is worth noting that some AGN might exhibit spatial extension because of the “pair-halo” effect where TeV γ rays interact the extra-galactic background light and are reprocessed into a “halo” of GeV γ rays around the original source direction. [28] Therefore, caution is required when using AGN to calibrate the PSF with flight data; see e.g., Ref. [29].

5.3. The Earth's Limb

The Earth's atmosphere is a very bright γ -ray source. Furthermore, at energies above a few GeV the γ -ray flux seen from space is dominated by γ rays from the interactions of primary cosmic-ray protons with the upper atmosphere. This consideration, together with the typical narrowness of the PSF at energies > 10 GeV, causes the Earth limb to appear as a very bright and sharp feature, with a very smooth spectrum, which provides an excellent calibration source. When using the Earth limb as a calibration source we generally limit the energy range to energies > 10 GeV, primarily because at lower energies orbital variations in the geomagnetic field significantly affect the γ -ray fluxes [30].

5.4. The Galactic Ridge

At energies above ~ 30 GeV no single source provides enough γ rays for a good comparison between flight data and MC simulations. However, the combination of bright Galactic sources and Galactic diffuse backgrounds means that there is a very large excess of γ rays coming from the Galactic plane relative to high Galactic latitudes.

The intensity of the γ -ray emission at low latitudes in the inner Galaxy is more than an order of magnitude greater than at high latitudes in the outer Galaxy. This can be used to validate event selections by comparing their effect on γ -ray rich samples at low Galactic latitudes with γ -ray poor samples high Galactic latitudes.

6. Looking Forward

In this chapter we have given an introduction to event reconstruction and source analysis in pair-conversion telescopes, and given several specific examples involving the *Fermi*-LAT.

Prior to the advent of the *Fermi*-LAT, astrophysical observations in the 100 MeV to 100 GeV range were largely limited by statistics. However, with the data ob-

tained from the *Fermi*-LAT after more than six years of continuous operation, we now have sufficient data that systematic effects, in terms of modeling of the astrophysical sources, understanding the residual charged-particle backgrounds, and characterizing the instrument response, are starting to dominate our measurements for the brighter sources at the 1% to 10% level.

Therefore, we expect that next generation pair conversion telescopes will require important refinements to the data analysis methodology described here, but the payoff from those refinements will be the ability to provide much more detailed answers about the γ -ray sky.

Acknowledgments

Much of our current understanding of how best to analyze data from pair-conversion telescopes was the result of the work of many people within the *Fermi*-LAT collaboration. This work included many refinements to the IRFs, the implementation of the likelihood formalism and the development of calibrations samples. In particular, we would like to acknowledge the contributions of Patrick Nolan to the formulation and implementation of the maximum likelihood analysis used for *Fermi*-LAT data.

References

- [1] G. Kanbach et al., The project EGRET (Energetic Gamma-Ray Experiment Telescope) on NASA's Gamma-Ray Observatory (GRO), *Space Sci. Rev.* **49**, 69–84, (1988).
- [2] M. Tavani et al., The AGILE Mission, *A&A*. **502**, 995–1013 (Aug., 2009). doi: 10.1051/0004-6361/200810527.
- [3] W. B. Atwood et al., The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission, *ApJ*. **697**, 1071–1102 (June, 2009). doi: 10.1088/0004-637X/697/2/1071.
- [4] M. Ackermann et al., The Fermi Large Area Telescope on Orbit: Event Classification, Instrument Response Functions, and Calibration, *ApJS*. **203**, 4 (Nov., 2012). doi: 10.1088/0067-0049/203/1/4.
- [5] R. Mankel, Pattern recognition and event reconstruction in particle physics experiments, *Reports on Progress in Physics*. **67**, 553–622 (Apr., 2004). doi: 10.1088/0034-4885/67/4/R03.
- [6] C. Labanti et al., Design and construction of the Mini-Calorimeter of the AGILE satellite, *Nucl. Inst. Meth.* pp. 470–479 (Jan., 2009). doi: 10.1016/j.nima.2008.09.021.
- [7] P. Bruel, Gamma rays, electrons and positrons up to 3 TeV with the Fermi Gamma-ray Space Telescope, *Journal of Physics Conference Series*. 404(1):012033 (Dec., 2012). doi: 10.1088/1742-6596/404/1/012033.
- [8] R. E. Kalman, A new approach to linear filtering and prediction problems, *Transactions of the ASME—Journal of Basic Engineering*. **82**(Series D), 35–45, (1960).
- [9] P. Billoir, R. Fruhwirth, and M. Regler, Track Element Merging Strategy and Vertex Fitting in Complex Modular Detectors, *Nucl. Inst. Meth.* **A241**, 115–131, (1985). doi: 10.1016/0168-9002(85)90523-6.
- [10] R. Fruhwirth, Application of Kalman filtering to track and vertex fitting, *Nucl. Inst. Meth.* **A262**, 444–450, (1987). doi: 10.1016/0168-9002(87)90887-4.

- [11] J. R. Mattox et al., The Likelihood Analysis of EGRET Data, *ApJ*. **461**, 396 (Apr., 1996). doi: 10.1086/177068.
- [12] F. Acero et al., Fermi Large Area Telescope Third Source Catalog, *ApJS*. 218:23 (June, 2015). doi: 10.1088/0067-0049/218/2/23.
- [13] A. M. T. Pollock et al., Search for gamma-radiation from extragalactic objects using a likelihood method, *A&A*. **94**, 116–120 (Jan., 1981).
- [14] S. S. Wilks, The large-sample distribution of the likelihood ratio for testing composite hypotheses, *Ann. Math. Statist.* **9**, 60–62, (1938). doi: 10.1214/aoms/1177732360.
- [15] H. Chernoff, On the distribution of the likelihood ratio, *Ann. Math. Statist.* **25**, 573–578, (1954). doi: 10.1214/aoms/1177728725.
- [16] P. L. Nolan et al., Fermi Large Area Telescope Second Source Catalog, *ApJS*. 199:31 (Apr., 2012). doi: 10.1088/0067-0049/199/2/31.
- [17] J. D. Scargle, Studies in Astronomical Time Series Analysis. V. Bayesian Blocks, a New Method to Analyze Structure in Photon Counting Data, *ApJ*. **504**, 405–418 (Sept., 1998). doi: 10.1086/306064.
- [18] J. D. Scargle, J. P. Norris, B. Jackson, and J. Chiang, Studies in Astronomical Time Series Analysis. VI. Bayesian Block Representations, *ApJ*. 764:167 (Feb., 2013). doi: 10.1088/0004-637X/764/2/167.
- [19] J.-M. Casandjian. Diffuse galactic radiation. In eds. F. A. Aharonian, W. Hofmann, and F. M. Rieger, *American Institute of Physics Conference Series*, vol. 1505, *American Institute of Physics Conference Series*, pp. 37–45 (Dec., 2012). doi: 10.1063/1.4772218.
- [20] A. A. Abdo et al., Spectrum of the Isotropic Diffuse Gamma-Ray Emission Derived from First-Year Fermi Large Area Telescope Data, *Phys. Rev. Lett.*. 104(10):101101 (Mar., 2010). doi: 10.1103/PhysRevLett.104.101101.
- [21] D. J. Thompson et al., Calibration of the Energetic Gamma-Ray Experiment Telescope (EGRET) for the Compton Gamma-Ray Observatory, *ApJS*. **86**, 629–656 (June, 1993). doi: 10.1086/191793.
- [22] J. A. Esposito et al., In-Flight Calibration of EGRET on the Compton Gamma-Ray Observatory, *ApJS*. **123**, 203–217 (July, 1999). doi: 10.1086/313227.
- [23] A. A. Abdo et al., The on-orbit calibration of the Fermi Large Area Telescope, *Astroparticle Physics*. **32**, 193–219 (Oct., 2009). doi: 10.1016/j.astropartphys.2009.08.002.
- [24] A. W. Chen et al., Calibration of AGILE-GRID with in-flight data and Monte Carlo simulations, *A&A*. 558:A37 (Oct., 2013). doi: 10.1051/0004-6361/201321767.
- [25] A. A. Abdo et al., The Vela Pulsar: Results from the First Year of Fermi LAT Observations, *ApJ*. **713**, 154–165 (Apr., 2010). doi: 10.1088/0004-637X/713/1/154.
- [26] A. A. Abdo et al., Fermi-LAT Observations of the Geminga Pulsar, *ApJ*. **720**, 272–283 (Sept., 2010). doi: 10.1088/0004-637X/720/1/272.
- [27] A. A. Abdo et al., Fermi Large Area Telescope Observations of the Crab Pulsar And Nebula, *ApJ*. **708**, 1254–1267 (Jan., 2010). doi: 10.1088/0004-637X/708/2/1254.
- [28] F. A. Aharonian, P. S. Coppi, and H. J. Voelk, Very high energy gamma rays from active galactic nuclei: Cascading on the cosmic background radiation fields and the formation of pair halos, *ApJL*. **423**, L5–L8 (Mar., 1994). doi: 10.1086/187222.
- [29] M. Ackermann et al., Determination of the Point-spread Function for the Fermi Large Area Telescope from On-orbit Data and Limits on Pair Halos of Active Galactic Nuclei, *ApJ*. 765:54 (Mar., 2013). doi: 10.1088/0004-637X/765/1/54.
- [30] A. A. Abdo et al., Fermi large area telescope observations of the cosmic-ray induced γ -ray emission of the Earth's atmosphere, *Phys. Rev. D*. 80(12):122004 (Dec., 2009). doi: 10.1103/PhysRevD.80.122004.