


MULTIWAVELENGTH ASTROPHYSICS LABORATORY

Modulo 3: FERMI-LAT data analysis applied to Jetted AGN



Paola Grandi – INAF-OAS, Bologna

paola.grandi@inaf.it , paola.grandi5@unibo.it

A.A. 2025-2026

Table of Contents

I	I: Data Reduction and Analysis Methods	1
1	The Gamma-ray Sky	3
1.1	The <i>Fermi</i> satellite	3
1.2	The Large Area Telescope	4
1.3	FERMI LAT Performance	8
1.3.1	Event Classification	8
1.3.2	Instrument Response Functions	8
2	Introduction to Fermi-LAT data reduction	15
2.1	Similarities and differences between gamma-ray and X-ray analysis workflows	15
2.2	Statistical framework for data reduction and analysis: likelihood overview	16
2.2.1	Likelihood applied to <i>Fermi</i> -LAT data	17
2.3	How to test the statistical significance of a <i>Fermi</i> -LAT source?	18
2.4	LAT background	20
2.4.1	Galactic diffuse emission	20
2.4.2	Isotropic background	22
2.4.3	Background photons produced by Earth limb emission	23
2.5	Spectral models	24
2.5.1	ROI model	24
3	A walkthrough of the data reduction process	26
3.1	Download the data	26
3.2	Data preparation and and Likelihood Analysis	28
3.3	Generation of Scientific Products	29
3.3.1	Spectral Energy Distribution	29
3.3.2	Light Curve	29
3.4	Parameter uncertainty estimation	30
3.4.1	Upper Limits	32
3.4.2	Another use of the TS: Variability Test	32
3.4.3	Variability, Timescales, and Source Size	33

**II II:Basic Astrophysical Concepts on Jetted AGN for the
Lab Session 35**

**4 Supermassive Black Holes and Nuclear Activity: *Accre-
tion Flows and Jets*
36**

4.1	Accretion Processes	36
4.1.1	The Shakura–Sunyaev Disk: A Radiatively Efficient Solution	36
4.1.2	The ADAF Solution: A Radiatively Inefficient Flow	36
4.1.3	How AGN Jets Are Produced: A intuitive explanation	37

5 Relativistic Jets and Blazars 40

5.1	Relativistic Effects in Astrophysical Jets	40
5.1.1	Spectral Energy Distribution of BLAZAR	41
5.1.2	The Two Peaks of the Blazar SED	42
5.2	Hadronic Models and Neutrinos	43

Appendices

A	Even file classification	46
B	Python code for the visualization of a 3D cube file	47
C	Region of Interest Model File	49
D	Data analysis steps	50
E	Connection between angular momentum loss and energy dissipation in a keplerian disk	53

Part I

I: Data Reduction and Analysis Methods

Introduction

In any branch of astrophysics, whether theoretical or observational, it is essential to understand how data are collected, reduced, and analyzed. A solid knowledge of observational procedures, calibration techniques, and data analysis methods is required in order to correctly interpret the results and to place meaningful constraints on theoretical models.

Gamma rays occupy the highest-energy end of the electromagnetic spectrum. They correspond to photons with extremely short wavelengths and very high frequencies, far beyond those of optical or radio photons. The comparison below provides an intuitive sense of scale (the quantities in the table are given as order-of-magnitude estimates rather than exact values):

Band	Wavelength	Frequency (Hz)	Photon Energy (eV)
Radio	30 cm	10^9	10^{-6}
Optical	600 nm	10^{14}	2
X-rays	1 nm	10^{17}	10^3
Gamma rays	10^{-6} nm	10^{23}	10^9

High-energy astrophysics is commonly divided into two main observational regimes:

- ◇ **HE (High Energy)**: from tens of MeV up to ~ 100 GeV, observed with space-based instruments;
- ◇ **VHE (Very High Energy)**: from ~ 100 GeV up to ~ 100 TeV, observed with ground-based facilities, such as Imaging Atmospheric Cherenkov Telescopes (IACTs).

These extreme energies probe non-thermal and relativistic physical processes, offering a unique window into particle acceleration and radiation mechanisms in astrophysical sources.

This lab focuses on the analysis of data acquired by the Large Area Telescope (LAT) onboard the *Fermi* satellite.

The Gamma-ray Sky

The gamma-ray sky, as observed by the *Fermi* satellite, is dominated by two main classes of sources: **blazars** and **pulsars**. However, several other source classes of great astrophysical interest have also been detected, including starburst galaxies, radio galaxies, supernova remnants, binary systems, and unclassified objects (see Fig. 1.1).




Figure 1.1. Sky map, in Galactic coordinates, showing the sources contained in the Fourth Fermi LAT Gamma-ray Source Catalog. For simplicity, all AGN classes are plotted with the same blue symbol. Associations to well-defined classes are plotted in red. Unassociated sources and sources associated with counterparts of unknown nature are plotted in black. The plot is taken from [Abdollahi et al. \(2020\)](#).

1.1 The *Fermi* satellite

Fermi was launched on June 11, 2008, from Cape Canaveral into an initial low-Earth orbit with an altitude of approximately 565 km, an inclination of 25.6° , and an eccentricity < 0.01 . The satellite operates primarily in

survey mode, continuously scanning the entire sky and achieving full sky coverage approximately every three hours.

Fermi carries two instruments: the Large Area Telescope (LAT), the principal instrument, and the Gamma-ray Burst Monitor (GBM). The LAT enables the observation of celestial sources from 20 MeV to above 300 GeV, while the GBM was designed to monitor gamma-ray bursts (see Fig. 1.2).





Figure 1.2. The *Fermi* satellite carries two detectors: the Large Area Telescope (LAT) and the Gamma-ray Burst Monitor (GBM).

The *Fermi* Collaboration produces source catalogs that periodically update the list of gamma-ray sources observed by the LAT. The most recent catalog is the Fourth Fermi LAT Gamma-ray Source (4FGL) Catalog, with subsequent updates (DR1, DR2, DR3, DR4). It can be found [here](#).

1.2 The Large Area Telescope

Unlike optical light and X-rays, gamma rays cannot be focused with conventional mirrors. At these energies, photons typically penetrate matter without interacting via reflection or refraction. Instead, gamma-ray telescopes rely on particle-interaction processes, most notably pair production at high energies. In this process, the incoming photon is converted into an electron–positron pair:

$$\gamma \rightarrow e^+ + e^-.$$

In free space, this process is forbidden because both energy and momentum must be conserved. The conversion must occur in the electric field of a nucleus (N), which absorbs a tiny fraction of the momentum (recoil), allowing the conservation laws to be satisfied (Blumenthal and Gould 1970). The photon must have at least 1.022 MeV (twice the electron rest mass) to produce a pair.

The LAT is composed of **16 identical modules (towers)** arranged in a 4×4 array, each containing two main subsystems: the **Tracker (TKR)** and the **Calorimeter (CAL)**. Surrounding the entire tower array is a single segmented **Anticoincidence Detector (ACD)**, which covers the LAT on the top and sides to reject charged-particle background (see Fig. 1.3).

Readers interested in the hardware specifications are referred to Atwood et al. (2009) and Ackermann et al. (2012).




Figure 1.3. Schematic representation of the Large Area Telescope (LAT). The LAT is composed of a 4×4 array of identical towers, each containing a Tracker (TKR) and an Imaging Calorimeter (CAL) located beneath it. Surrounding the entire array is a segmented Anticoincidence Detector (ACD), used to reject charged-particle background.

TKR — The tracker forms the upper part of each tower. It consists of alternating planes of thin tungsten converter foils and silicon strip detectors. Since the probability (cross section) for pair production (Jackson 1998)

$$\sigma_{\text{pair}}(E_\gamma) \propto Z^2 \ln\left(\frac{E_\gamma}{m_e c^2}\right) \quad [1]$$

increases roughly as Z^2 , tungsten ($Z = 74$), with its high atomic number, provides an efficient material for gamma-ray conversion.

The converter planes alternate with silicon detectors that record the passage of charged particles. These detectors measure the tracks of the electron and positron created when a gamma ray converts into a pair, and this information is used to reconstruct the direction of the incoming photon. A schematic illustration of a photon converting into an electron-positron pair is shown in Figure 1.4.

The tracker is divided into two sections: i) the **front section** and ii) the **back section**.

i) The front section of the tracker (the first 12 x-y planes) contains very thin tungsten sheets. When a photon converts in this region, the resulting electron and positron traverse only a small amount of material and therefore undergo little multiple scattering — that is, small random deflections from atomic nuclei. As a result, their trajectories remain straighter and the reconstructed photon direction has improved angular resolution.

ii) In the much thicker back section (the last four x-y planes), the reconstructed photon direction is less precise due to increased multiple scattering. However, the thicker tungsten layers increase the probability of




Figure 1.4. Gamma rays penetrate into the detector (the tracker) and interact with a high- Z converter material, in this case tungsten, to produce an electron-positron pair. This pair is tracked through the instrument by silicon strip detectors. Since the gamma-ray energy is much larger than the rest mass of the electron and positron, both members of the pair continue predominantly in the direction of the incident gamma ray.

photon conversion, thereby improving the overall detection efficiency of the instrument.

CAL — The Calorimeter is composed of thallium-doped cesium iodide crystals (CsI(Tl)). When the electron-positron pair produced in the tracker enters the calorimeter, it interacts with the crystal material and generates scintillation light in the form of small flashes of visible photons. The intensity of this light is proportional to the energy deposited in each crystal.

Each calorimeter module consists of 96 CsI(Tl) crystals arranged in eight layers, with successive layers rotated by 90° with respect to one another. This X-Y geometry enables the reconstruction of the shower (see Box 1.A) development both longitudinally, along the depth of the calorimeter, and transversely, across its lateral spread.

Each crystal is read out by two photodiodes mounted at the two opposite longitudinal ends of the crystal. The energy deposited in a crystal is proportional to the sum of the signals from the two photodiodes, while the total shower energy is obtained by summing the deposits over all crystals that register scintillation light. By comparing the relative signals collected at the two ends, the position of the energy deposition along the crystal length can be inferred.

Together with the X-Y segmentation, this information allows for a full three-dimensional reconstruction of the electromagnetic shower. This 3D characterization of the shower provides two key benefits:

- ◊ **Energy reconstruction (leakage correction):** High-energy showers may extend beyond the calorimeter, causing part of the energy to escape. By analysing the longitudinal and lateral shower shapes, the LAT software estimates the escaping fraction and corrects the reconstructed photon energy accordingly.
- ◊ **Background rejection:** Electromagnetic showers (from photons or electrons) are compact and well structured, while hadronic showers (from protons or nuclei) are broader and more irregular. Comparing the observed shower profile with the expected electromagnetic one, it is possible to discriminate gamma-ray events from the cosmic-ray background.




Figure 1.5. LAT calorimeter module. The 96 CsI(Tl) scintillator crystal detector elements are arranged in eight layers, with the orientation of the crystals in adjacent layers rotated by 90° . Each crystal is read out by two photodiodes at its ends.

ACD — The anticoincidence detector forms the outer shield covering the top and sides of the instrument. It **rejects charged cosmic-ray particles** (protons, electrons) entering the detector, which are far more numerous than gamma rays. When a charged particle crosses the ACD, it emits light that triggers a veto signal. This signal tells the LAT electronics to reject the event, ensuring that only photons are accepted as valid detections.

Summary

- ◊ A gamma-ray photon enters the LAT without activating the ACD (indicating a neutral particle).
- ◊ It converts into an electron–positron pair in the tungsten foils of the tracker.
- ◊ The silicon planes record the trajectories of these charged particles, determining the direction of the original photon.
- ◊ The particles then enter the calorimeter, where the electromagnetic shower develops and its total energy is measured.
- ◊ Combining the direction (from the tracker) and energy (from the calorimeter) fully characterizes the photon.

1.A — Electromagnetic Shower

An *electromagnetic shower* is a cascade of secondary particles produced when a high-energy photon or a high-energy electron propagates through matter. The shower development is driven by two fundamental electromagnetic processes. High-energy photons interacting with the detector material can undergo *pair production*,

$$\gamma \rightarrow e^+ + e^-,$$

while high-energy electrons and positrons emit photons via *bremsstrahlung*,

$$e^\pm \rightarrow e^\pm + \gamma.$$

The photons produced by bremsstrahlung can in turn convert into new electron–positron pairs, and the newly produced charged particles can emit further photons. The repeated alternation of these two processes leads to a rapid multiplication of particles, forming a cascade.

The longitudinal development of an electromagnetic shower is driven by bremsstrahlung and pair production, which produce secondary particles predominantly in the forward direction. The lateral spread of the

shower, on the other hand, is mainly caused by multiple Coulomb scattering, which progressively deflects charged particles away from the initial direction of propagation. The shower continues to develop until the energies of the particles fall below the threshold for efficient bremsstrahlung and pair production. Below this energy, ionization losses dominate and the cascade gradually dies out.

In the *Fermi*-LAT, electromagnetic showers develop mainly in the calorimeter, which is designed to absorb most of the shower energy. The tracker, on the other hand, is intentionally thin in radiation lengths^a, in order to minimize shower development and preserve precise reconstruction of particle trajectories.

^aThe radiation length X_0 is a characteristic property of a material that governs the development of electromagnetic showers. For electrons and positrons, it corresponds approximately to the distance over which their energy is reduced by a factor $1/e$ due to bremsstrahlung, while for high-energy photons it sets the typical scale for pair production.

1.3 FERMI LAT Performance

In the *Fermi* Large Area Telescope, the interpretation of the detected events relies on extensive Monte Carlo simulations. In these simulations, both gamma-ray photons and background particles are generated and propagated through a detailed model of the instrument, including its geometry, materials, and detector response. By comparing the reconstructed quantities with the known properties of the simulated particles, the performance of the instrument can be characterized and its response to gamma rays can be quantified¹.

By *data processing pass* (PASS), the community refers to a specific version of the complete analysis chain used to transform the raw detector signals into reconstructed physical events. As the understanding of the detector improves and more realistic simulations become available, the data can be reprocessed using updated passes. The most recent version is Pass 8.

1.3.1 Event Classification

Each detected event is classified according to its probability of being a true astrophysical photon and the quality of its reconstruction. A detailed description of the available event classes is provided in the [Fermi-LAT documentation](#). A summary to help navigate the different event classes is shown in Appendix A (Table A.1).

For the study of point-like sources, the Pass 8 **SOURCE** class (`evclass=128`) is commonly adopted, as it provides a good compromise between sensitivity and background rejection.

Events within each class are further subdivided into event types, based on specific event properties, such as whether the photon converted in the Front or Back section of the Tracker (Appendix A, Table A.2).

1.3.2 Instrument Response Functions

Each event class is associated with a specific *Instrument Response Function* (IRF). The IRF characterizes how the LAT converts the incoming photon flux from the sky into detected events. They are assumed to be factorized into three components, with the quantities $\hat{\nu}$, θ , and ϕ defined in Fig. 1.6:

(1) **Effective area**

$$A_{\text{eff}}(E, \hat{\nu}, s),$$

defined as the product of the geometrical collection area, the γ -ray conversion probability, and the efficiency of the event selection s for a photon of true energy E and incident direction $\hat{\nu}$ in the LAT frame.

(2) **Point-spread function (PSF)**

$$P(\hat{\nu}'; E, \hat{\nu}, s),$$

the probability density (see Box 1.B) to reconstruct a direction $\hat{\nu}'$ for a γ ray with true direction $\hat{\nu}$ and true energy E under the event selection s .

(3) **Energy dispersion**

$$D(E'; E, \hat{\nu}, s),$$

the probability density to reconstruct an energy E' for a γ ray with true energy E and incident direction $\hat{\nu}$ for events passing the selection s .

Given a true distribution of γ rays $S(E, \hat{p})$, where \hat{p} refers to the celestial direction of the photons, the expected distribution of observed γ rays for a class-type s can be written as:

¹Details in [The Fermi Large Area Telescope On Orbit: Event Classification, Instrument Response Functions, and Calibration](#).




Figure 1.6. Schematic view of the LAT, showing the layout of the 16 CAL modules and 12 of the 16 TKR modules (the ACD is omitted for clarity). The figure also defines the (θ, ϕ) reference frame. The z axis (LAT boresight) corresponds to the central pointing direction of the instrument. \hat{v} denotes the arrival direction of the photon (credit: Ackermann et al. 2012).

$$M(E', \hat{p}', s) = \iiint S(E, \hat{p}) A_{\text{eff}}(E, \hat{v}(t; \hat{v}), s) P(\hat{v}'(t; \hat{v}'); E, \hat{v}(t; \hat{v}), s) D(E'; E, \hat{v}(t; \hat{v}), s) dE d\Omega dt. \quad [2]$$

The integrals over dE , $d\Omega$, and dt appearing in Eq. (2) are required because the LAT does not measure the *true* photon properties directly. Instead, it reconstructs the photon energy and arrival direction with finite resolution, and its response varies with time as the spacecraft scans the sky.

- ◊ **Integration over dE (true energy).** A photon reconstructed with energy E' may have arrived with a range of true energies E , due to the finite energy resolution of the detector and possible energy losses (e.g. shower leakage). The integral over dE accounts for all true photon energies that can contribute to the measured energy E' , weighted by the source spectrum and the energy dispersion function.
- ◊ **Integration over $d\Omega$ (true direction).** Because of the finite angular resolution of the LAT, a photon reconstructed from direction \hat{p}' may have originated from nearby true sky directions \hat{p} . The integral over $d\Omega$ sums the contributions from all true directions that, through the point-spread function (PSF), can be reconstructed as \hat{p}' .
- ◊ **Integration over dt (time).** The LAT operates on a satellite that continuously changes its orientation while surveying the sky. As a result, the same source is observed under different incidence angles and detector configurations, and the instrument response varies with time. The integral over dt accounts for all observing conditions during the selected time interval, effectively incorporating the accumulated exposure.

In summary, the expected number of reconstructed events is obtained by *forward folding* the true sky model through the instrument response and the observation history:

$$\text{true sky model} \xrightarrow{\text{IRFs} + \text{scanning law}} \text{predicted reconstructed counts.}$$

1.B — Probability density functions

A *probability density function* (PDF) describes how a continuous variable x (such as an angle or an energy) is distributed. The value of the PDF at a given point, $p(x)$, is *not* a probability by itself. The probability that the variable lies within a finite interval $[x_1, x_2]$ is obtained by integrating the PDF over

that interval,

$$P(x_1 \leq x \leq x_2) = \int_{x_1}^{x_2} p(x) dx.$$

A probability is therefore a dimensionless quantity between 0 and 1.

Example: Gaussian distribution. A common example of a probability density function is the Gaussian (normal) distribution,

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}},$$

where μ is the mean value and σ is the standard deviation. The probability that the measured value lies within one standard deviation from the mean is given by

$$P(|x - \mu| \leq \sigma) = \int_{\mu-\sigma}^{\mu+\sigma} p(x) dx \simeq 0.68.$$

This means that approximately 68% of the measurements are expected to fall within $\pm 1\sigma$ of the mean value.

Effective Area

The **effective area**, $A_{\text{eff}}(E, \hat{v}, s)$, quantifies the probability that a gamma-ray photon with energy E and direction \hat{v} in the LAT reference frame is detected and selected as a valid event. It can be expressed conceptually as the product of three factors: the geometrical collection area of the instrument (A_{geom}), the probability that the photon converts within the detector (P_{conv}), and the efficiency (ε_{sel}) of a given event selection (s):

$$A_{\text{eff}}(E, \hat{v}, s) = A_{\text{geom}} P_{\text{conv}}(E, \hat{v}) \varepsilon_{\text{sel}}(E, \hat{v}, s).$$

The effective area depends primarily on the photon energy E and on the direction of incidence, usually described by the polar angle θ and the azimuthal angle ϕ (see Fig. 1.6). Since the effective area is routinely averaged over the azimuthal angle in scientific analyses, plots are generally shown as a function of photon energy and incidence angle.

In **Fig. 1.7—left panel**, the effective area is shown as a function of photon energy, assuming $\theta = 0^\circ$, i.e. for photons arriving *on-axis*. The maximum efficiency is reached between approximately 1 GeV and 100 GeV; below and above this range, the effective area gradually decreases.

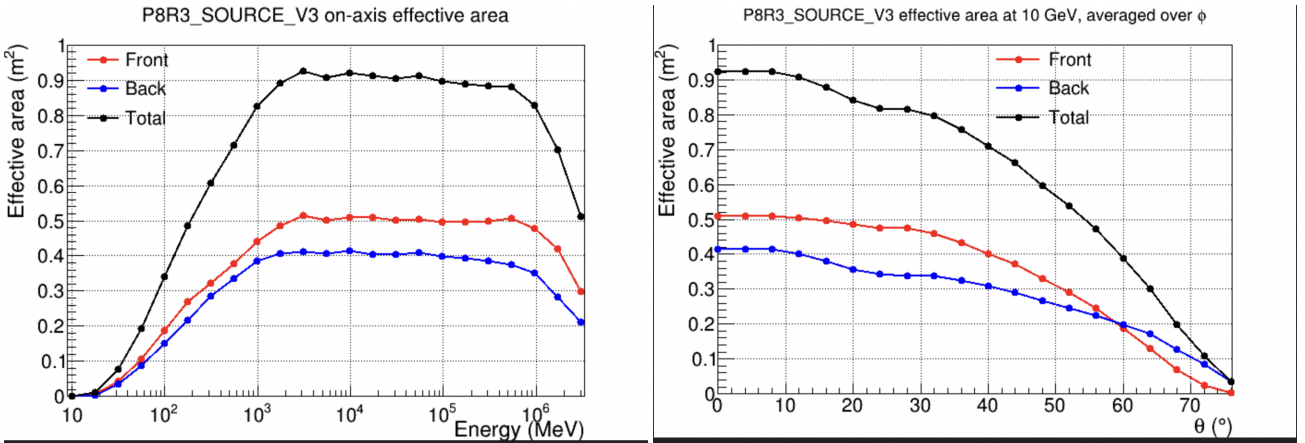


Figure 1.7. On-axis effective area as a function of energy (*left panel*) and angular dependence (*right panel*) of the effective area at 10 GeV (credit: **FERMI LAT-SLAC**).

As the pair-production cross section increases approximately with the logarithm of the photon energy (eq. 1), the LAT effective area rises rapidly at low energies and then becomes nearly constant around 1 GeV, when almost all photons interact and convert efficiently in the detector.

At high photon energies (above about 100 GeV), the situation is different. Here, the probability of pair production is very high, and most gamma rays convert into an electron–positron pair. The resulting relativistic electrons and positrons emit high-energy photons through bremsstrahlung. These photons can subsequently undergo pair production, generating new electron–positron pairs and initiating an extended electromagnetic shower, which at high energies may exceed the physical size of the LAT calorimeter. To maintain data quality and reject background from charged particles, the LAT team applies stricter selection cuts at these energies. Consequently, many high-energy γ -ray events are discarded, and the effective area drops again.

In **Fig. 1.7—right panel**, the effective area is plotted for photons of 10 GeV observed at different incidence angles θ . As θ increases, the effective area decreases because the projected area shrinks, the photon is less likely to convert, and reconstruction efficiency decreases.

The integral of the effective area over the solid angle, called the *acceptance*, is defined as

$$A(E) = \int A_{\text{eff}}(E, \theta, \phi) d\Omega = \int_0^{2\pi} \int_0^{\pi/2} A_{\text{eff}}(E, \theta, \phi) \sin \theta d\theta d\phi, \quad [12]$$

and represents another widely used performance measure. The acceptance as a function of energy is shown in **Fig. 1.8**.




Figure 1.8. Acceptance, i.e. the integral of the effective area over solid angle, as a function of energy.

Point-spread function

The **point-spread function (PSF)**, defined as $P(\hat{\mathbf{v}}'; E, \hat{\mathbf{v}}, s)$, is the probability density to reconstruct an incident direction. Because the LAT cannot reconstruct the incoming direction perfectly, the reconstructed vector $\hat{\mathbf{v}}'$ does not coincide exactly with the true direction $\hat{\mathbf{v}}$. To quantify this uncertainty, detailed Monte Carlo simulations are used: for each simulated photon the true direction is known, and by comparing it with the reconstructed one, the LAT team obtains a distribution of angular errors for each energy and incidence angle.

The PSF therefore gives the probability density for the LAT to reconstruct a photon at an angular separation $\Delta\theta$ from its true direction, i.e.

$$\text{PSF}(\Delta\theta | E, \hat{\mathbf{v}}, s) \equiv P(\hat{\mathbf{v}}'; E, \hat{\mathbf{v}}, s).$$

This equality assumes an approximate axial symmetry of the instrument response around the true photon direction. Under this assumption, the probability density depends only on the angular separation $\Delta\theta$ between the true and reconstructed directions and not on the azimuthal angle. As a result, all points located at the same angular distance from the true direction have the same probability density.

Since the LAT continuously scans the sky, photons from a given source are detected with many different incidence angles θ . To describe the overall, time-averaged angular resolution of the LAT during its all-sky survey, the PSF must therefore be averaged over all possible incidence angles. Not all directions contribute equally: the LAT is more sensitive at some angles (larger effective area $A_{\text{eff}}(E, \theta)$) and it spends more time observing certain directions than others. The resulting quantity, known as the *acceptance-weighted PSF*, represents the effective, observation-averaged angular resolution of the LAT in survey mode.

A convenient single-number summary of the LAT angular performance is the (68% or 95%) containment radius, i.e. the angular separation that contains 68% or 95% of the reconstructed directions for a given true photon. Plotting this containment radius as a function of energy provides a simple and intuitive summary of the LAT angular resolution.

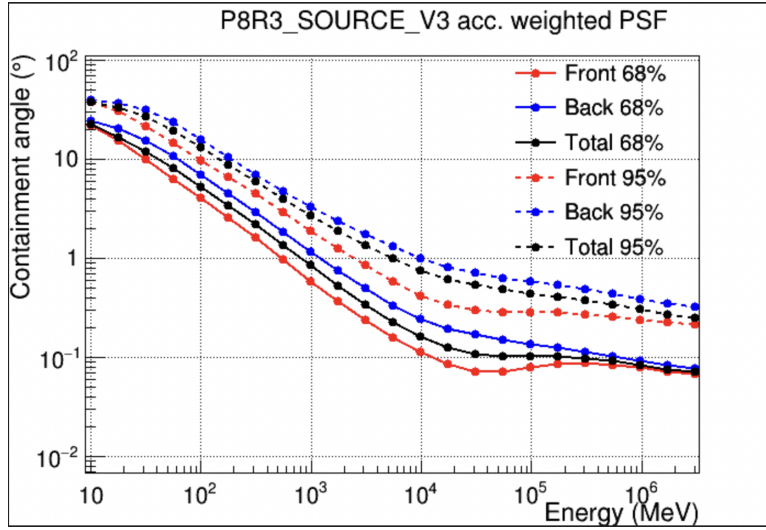


Figure 1.9. 68% and 95% containment radii of the acceptance-weighted PSF from the P8R3_V3 instrument response functions, as a function of photon energy. Results are shown for both FRONT and BACK conversion events, as well as for the total (combined) PSF. The containment radii represent the angular distance within which 68% or 95% of the reconstructed photons from a point source are included, providing a direct measure of the LAT angular resolution.

Figure 1.9 illustrates the energy dependence of the LAT point-spread function. At an energy of 100 MeV, the PSF is relatively broad, indicating a large angular uncertainty in the reconstruction of the photon direction. The 68% containment angle is on the order of $\sim 3^\circ$. At low energies (< 1 GeV), the PSF is dominated by *multiple Coulomb scattering* of the e^+e^- pair in the tungsten converter foils. In this regime, the electrons and positrons have relatively low momentum (p), so each scattering from a tungsten nucleus produces a sizeable deflection. The root-mean-square scattering angle follows approximately $\theta_0 \propto \frac{1}{p}$, so particles with low momentum are deflected more easily. As the pair traverses several radiation lengths in the tracker, the accumulation of many such small-angle deflections produces a characteristic “zig-zag” trajectory and thus a broadened reconstructed photon direction.

As the energy increases, the situation changes. The electrons and positrons produced in the pair conversion carry higher momenta, so multiple Coulomb scattering becomes much less effective. At these higher energies, however, the charged particles can lose a significant fraction of their energy through *bremsstrahlung* as they pass through the tungsten foils. Bremsstrahlung predominantly affects the particle’s energy rather than its direction: the emission of a high-energy photon typically occurs at very small angles with respect to the particle’s trajectory and therefore does not noticeably change the track direction measured by the tracker.

Finally, it should be noted that at large off-axis angles (θ), the PSF broadens because the e^+e^- pair traverses more material in the tracker, increasing the contribution of multiple scattering and making track reconstruction more difficult.

1.C Coulomb Scattering vs. Bremsstrahlung

When a charged particle such as an electron or positron produced in a pair conversion travels through matter, it interacts with the electric field of atomic nuclei. Two distinct physical processes arise from this interaction: *multiple Coulomb scattering* and *bremsstrahlung*. Although both are caused by the nuclear field, they affect the particle’s trajectory and energy loss in very different ways.

- ◊ At **low energies (below a few hundred MeV)**: Coulomb scattering dominates. The e^+e^- pair follows a zig-zag path that broadens the point-spread function (PSF).
- ◊ At **intermediate energies**: the effect of multiple scattering decreases as the particle momentum increases ($\theta_0 \propto 1/p$, where p is the particle momentum), while bremsstrahlung starts to become non-negligible.
- ◊ At **high energies (GeV and above)**: bremsstrahlung dominates the interaction probability, but it affects mainly the particle's energy rather than its direction.

In summary, Coulomb scattering controls the LAT angular resolution at low energies, while bremsstrahlung becomes increasingly important at high energies, governing energy losses but not significantly worsening the PSF.

Energy Dispersion Function

The **energy dispersion function (EDF)**, defined as $D(E'; E, \hat{\mathbf{v}}, s)$, gives the probability density to reconstruct an energy E' for a γ ray with true energy E , incident direction $\hat{\mathbf{v}}$, and event selection s .

When a γ ray enters the LAT, its energy cannot be measured perfectly. Instead, the detector reconstructs a value E' that fluctuates around the true photon energy E . To characterize these fluctuations, the LAT team relies on detailed Monte Carlo simulations in which photons of known energy and direction are generated and processed through the same reconstruction algorithms used for real LAT data. For each true energy E , these simulations produce a distribution of reconstructed energies E' . This distribution, i.e. the *energy dispersion function* (EDF), describes the probability of measuring a given E' when the true photon energy is E .

Since the LAT observes the sky while continuously scanning, photons do not always arrive from the same direction. The shape of the EDF changes with the incidence angle, because the containment of the electromagnetic shower is better for some angles than for others. To represent the average performance of the instrument in survey mode, the LAT uses an *acceptance-weighted* EDF, obtained by averaging the EDFs from all angles while giving more weight to those directions where the instrument is more sensitive and spends more observing time.

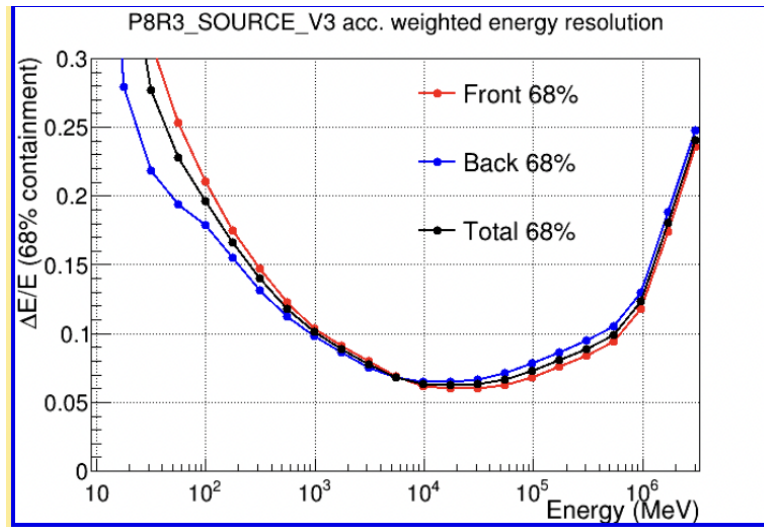


Figure 1.10. Acceptance-weighted energy resolution (i.e. 68% containment half-width of the reconstructed incoming photon energy) as a function of energy for FRONT, BACK, and FRONT+BACK events.

The EDF energy performance is often summarized using the *energy resolution*, usually expressed as $\Delta E/E$, where ΔE is an interval that contains a fixed fraction (e.g. 68%) of the reconstructed energies. Plotting $\Delta E/E$ as a function of energy provides a simple and intuitive way to evaluate how well the LAT measures the photon energy across its full energy range. The resulting curve has a characteristic “U-shape”, as shown in Fig. 1.10, reflecting the different physical effects that dominate at different energies.

At **low energies** ($< 200\text{--}300$ MeV), the electromagnetic shower is only partially developed when it enters the calorimeter. As a result, the measured energy fluctuates strongly from event to event, producing a broad

EDF and a poor energy resolution.

At **intermediate energies** (hundreds of MeV to a few GeV), the calorimeter contains most of the shower, and the intrinsic segmentation of the detector allows a more precise reconstruction of the total deposited energy. This is the regime where the LAT achieves its best energy resolution.

At **high energies** ($\gtrsim 10\text{--}20$ GeV), the shower becomes too extended to be fully contained within the calorimeter volume. A growing fraction of the energy leaks out of the back or sides of the detector, and the correction algorithms cannot recover this missing energy perfectly.

Introduction to Fermi-LAT data reduction

This chapter introduces basic concepts and key information required for *Fermi*-LAT data analysis.

2.1 Similarities and differences between gamma-ray and X-ray analysis workflows

X-ray and gamma-ray data reduction share several conceptual steps.

In both energy bands, the analysis requires:

- ◊ an **event file**, containing the arrival time, reconstructed energy, and sky position of each photon;
- ◊ a **spacecraft/housekeeping file**, describing the configuration and status of the satellite;
- ◊ **background models/files**;
- ◊ **instrument response functions** describing the detector performance.

Dedicated processing software cleans the raw data by filtering out spurious or non-physical events and prepares the resulting dataset for the creation of scientific data products:

- ◊ **images**;
- ◊ **light curves**;
- ◊ **spectra**.

Despite these similarities, the two analyses differ in fundamental ways because of the different detector technologies, the number of detected photons, and the associated statistical treatment.

- ◊ X-ray photons are **focused by grazing-incidence mirrors** and concentrated onto **CCD detectors**, which directly measure their arrival position and energy with good precision. This focusing capability leads to a narrow and well-defined point-spread function (PSF) and allows source photons to be confined within a small extraction region. The analysis is typically performed on a small region around the source, and a nearby, source-free region is selected to estimate the background. The data products are relatively straightforward to handle, since they depend mainly on the well-characterized properties of the focusing optics and CCD detectors, which are generally stable and well calibrated.
- ◊ Gamma-ray photons cannot be focused by an optical system. When they enter the detector, they interact with the material and convert into an **electron–positron pair**. The direction of the incoming gamma ray is then **reconstructed** from the tracks of these secondary particles in the silicon strip tracker, while the energy is estimated from the development of the electromagnetic shower in the calorimeter. Because the measurement is indirect and relies on particle tracking rather than optical focusing, the angular uncertainty is much larger than in X-ray instruments, and the response is intrinsically probabilistic. Consequently, gamma-ray analyses cover a much larger **Region of Interest (ROI)**, often several degrees across. Because of the broad PSF and the presence of a strong and structured diffuse background, all sources inside the ROI must be modeled simultaneously. The analysis therefore requires a global approach, where the entire region contributes to the construction and optimization of the model, rather than focusing solely on the target source.

The statistical methods are necessarily different. The number of detected γ -ray photons is often low, making the Gaussian approximation invalid. The χ^2 statistic relies on the assumption that each bin contains a sufficiently large number of events (typically at least 15–20) so that the Poisson distribution approaches a Gaussian. Since this condition is often not satisfied in gamma-ray analysis (see Box B.5), the χ^2 test would provide biased parameter estimates and unreliable significances.

The Cash statistic (C-stat), although suitable for low-count Poisson data, also cannot be applied in the same way as in X-ray astronomy. Unlike X-ray observations, the gamma-ray background cannot be measured from a local, source-free region; instead, it must be modeled globally and includes several components such as Galactic diffuse emission, isotropic emission, and all point sources within the ROI. Furthermore, the instrumental response of gamma-ray telescopes is substantially more complex: the effective area, the broad and energy-dependent PSF, and the energy dispersion must all be folded into the model in a non-linear way. Gamma-ray data are also inherently multi-dimensional, involving spatial and spectral dimensions simultaneously. For these reasons, the statistically correct framework for standard LAT source analyses is the **Poisson likelihood**.

2.A Example: Expected LAT counts from an integrated flux

Consider the gamma-ray source MKN 421 with an integrated photon flux between 1 and 100 GeV of

$$F = 3.3 \times 10^{-8} \text{ ph cm}^{-2} \text{ s}^{-1}.$$

Assume an average LAT effective area (front + back) of about

$$A_{\text{eff}} \simeq 1.0 \text{ m}^2 = 10^4 \text{ cm}^2,$$

obtained from the quoted value $A_{\text{eff}} \simeq 2.5 \text{ m}^2 \text{ sr}$ divided by the field of view ($\Delta\Omega_{\text{LAT}} \simeq 2.4 \text{ sr}$).

The expected count rate over the whole energy band is

$$R_{\text{tot}} = F A_{\text{eff}} = 3.3 \times 10^{-8} \times 10^4 = 3.3 \times 10^{-4} \text{ ph s}^{-1}.$$

This corresponds to $\sim 1 \text{ ph hour}^{-1}$ and $\sim 29 \text{ ph day}^{-1}$.

Thus, for an integrated flux of $10^{-8} \text{ ph cm}^{-2} \text{ s}^{-1}$ in the 1–100 GeV band, the LAT detects only a few tens of photons in one day of observations. This clearly shows that gamma-ray astronomy typically operates in the **Poisson regime**, not in the Gaussian regime.

2.2 Statistical framework for data reduction and analysis: likelihood overview

The likelihood L is defined as the probability of obtaining the observed data given an assumed model. In the low-count regime typical of gamma-ray sources, the number of detected photons per bin follows a Poisson distribution. If the model predicts an average number of counts m_i in the i -th bin, and the data contain n_i observed counts, the probability of observing n_i is

$$p(n_i | m_i) = \frac{m_i^{n_i}}{n_i!} e^{-m_i}.$$

The likelihood L is the product of these probabilities over all bins:

$$L = \prod_i \frac{m_i^{n_i}}{n_i!} e^{-m_i}.$$

The exponential factors can be combined as

$$\prod_i e^{-m_i} = e^{-\sum_i m_i} = e^{-N_{\text{exp}}},$$

where N_{exp} is the total number of events predicted by the model. Thus the likelihood becomes

$$L = e^{-N_{\text{exp}}} \prod_i \frac{m_i^{n_i}}{n_i!}.$$

Taking the logarithm yields

$$\log L = -N_{\text{exp}} + \sum_i n_i \log(m_i) - \sum_i \log(n_i!).$$

If the bin sizes are made extremely small, each bin contains at most one event. In this limit the likelihood becomes

$$L = e^{-N_{\text{exp}}} \prod_{j=1}^{N_{\text{obs}}} m_j,$$

where the product runs over the detected photons only¹. This is the **unbinned likelihood**, which uses the full information (energy, direction, time) of each event. It is the most accurate form of the likelihood, although it can become computationally expensive when the number of photons is large.

Further information can be found in the *Fermi* [Cicerone](#) documentation.

2.2.1 Likelihood applied to *Fermi*-LAT data

The *Fermi*-LAT data analysis is based on a likelihood approach, in which the observed gamma-ray events are statistically compared with the predictions of a model describing the gamma-ray sky. The sky model contains all gamma-ray sources contributing to the Region of Interest (ROI), together with diffuse background components. These components are folded with the Instrument Response Functions (IRFs) to predict the number of detected events. The likelihood is then maximized by varying model parameters until the best agreement with the data is reached.

The total model in the ROI is the sum of all point-like or extended sources plus diffuse background components:

$$S_{\text{tot}}(E, \hat{v}) = \sum_{i=1}^N S_i(E, \hat{v}; \eta_i) + \beta_{\text{gal}} B_{\text{gal}}(E, \hat{v}) + \beta_{\text{iso}} B_{\text{iso}}(E).$$

Components:

- ◊ $S_i(E, \hat{v}; \eta_i)$: spectral (and spatial) model of source i ;
- ◊ $B_{\text{gal}}(E, \hat{v})$: Galactic diffuse emission (normalization β_{gal});
- ◊ $B_{\text{iso}}(E)$: isotropic diffuse emission (normalization β_{iso}).

The parameter vector for each source can be written schematically as

$$\eta_i = \{\text{normalization}, \alpha, \beta, \dots\}.$$

As described in Section 1.3.2, the expected distribution of reconstructed events is $M(E', \hat{p}', s)$, obtained by folding the sky model with the IRFs (see Eq. 2 in Chapter 1). For a binned analysis, the expected number of counts in bin k is then

$$m_k = \int_{\text{bin } k} M(E', \hat{p}', s) dE' d\Omega'.$$

$d\Omega' \neq d\Omega$

The solid angle $d\Omega$ appearing in the definition of $M(E', \hat{p}', s)$ (Eq. 2 in Chapter 1) refers to the *true* photon directions on the sky and is integrated over when folding the sky model with the instrument response. The solid angle $d\Omega'$ instead refers to the *reconstructed* photon directions and is the variable over which the expected counts are integrated when computing the number of events in a given analysis bin.

Assuming Poisson statistics for the number of detected events in each bin, the likelihood comparing predicted counts m_k with observed counts n_k is

$$\mathcal{L}(\{\eta_i\}, \beta_{\text{gal}}, \beta_{\text{iso}}) = \prod_k \frac{m_k^{n_k} e^{-m_k}}{n_k!}.$$

The best-fit parameters maximize the likelihood:

$$\{\eta_i^*, \beta_{\text{gal}}^*, \beta_{\text{iso}}^*\} = \arg \max_{\eta_i, \beta_{\text{gal}}, \beta_{\text{iso}}} \mathcal{L}.$$

¹The factorial function takes the value $n! = 1$ for $n = 0$ and $n = 1$.

Summary

The likelihood method applied to *Fermi*-LAT data involves:

- (1) building a sky model including all sources and diffuse components;
- (2) folding this sky model with the IRFs to predict detected counts;
- (3) comparing predicted and observed counts with a Poisson likelihood;
- (4) varying source and background parameters until the likelihood is maximized.

2.3 How to test the statistical significance of a *Fermi*-LAT source?

In *Fermi*-LAT analyses, the detection of a γ -ray source is commonly based on the likelihood ratio test. The null hypothesis H_0 corresponds to a sky model in which the source is not included, while the alternative hypothesis H_1 includes the source with one or more free parameters (e.g. flux normalization and possibly the spectral index). For each hypothesis, we compute the maximum likelihood of the observed photon data assuming Poisson statistics. The Test Statistic (TS) used by the LAT is defined as

$$TS = -2 \ln \left(\frac{L_0}{L_1} \right) = 2 (\ln L_1 - \ln L_0), \quad [2.1]$$

where L_0 and L_1 are the maximized likelihoods under H_0 and H_1 .

The interpretation of TS relies on *Wilks' theorem*. When two statistical models are *nested*—that is, when the more complex model H_1 reduces to the simpler one H_0 by imposing q constraints on its parameters—the likelihood-ratio statistic in Eq. [2.1] converges in distribution to a chi-square distribution with q degrees of freedom. Therefore, if only the normalization k of a power-law model,

$$F(E) = k E^{-\Gamma},$$

is left free to vary (with fixed Γ), the detection statistic is asymptotically distributed as

$$TS \sim \chi_1^2.$$

In this case, the detection significance can be approximated by

$$\text{Significance} \approx \sqrt{TS},$$

an approximation introduced by [Mattox et al. \(1996\)](#). A source is typically considered detected when $TS \geq 25$, corresponding to roughly a 5σ significance (see Table 2.1).

When two independent parameters differ between H_0 and H_1 , the TS follows approximately a chi-square distribution with two degrees of freedom (χ_2^2). In this case, the significance must be computed from the χ_2^2 distribution rather than from the \sqrt{TS} rule.

Confidence level	p-value	χ_1^2	χ_2^2
68% ($\sim 1\sigma$)	0.32	1.00	2.30
95% ($\sim 2\sigma$)	0.05	3.84	5.99
99.7% ($\sim 3\sigma$)	0.003	9.00	11.83
99.9937% ($\sim 4\sigma$)	6.3×10^{-5}	16.00	18.42
99.99994% ($\sim 5\sigma$)	5.7×10^{-7}	25.00	28.74

Table 2.1. Critical values of the χ^2 distribution for selected confidence levels, with 1 and 2 degrees of freedom.

2.B How to read Table 2.1

The chi-square table provides the *critical values* that a test statistic (TS or χ^2) must exceed in order to reach a given significance level.

What does the p-value represent?

In a chi-square test, the p-value is defined as the *integral of the tail of the χ^2 distribution* to the right of the observed test statistic (see Fig. 2.1), i.e.

$$p = \int_{\chi_{\text{observed}}^2}^{\infty} f_{\chi^2}(x) dx.$$

This is the probability of obtaining a value of the test statistic at least as large as the observed one *if the null hypothesis were true*.

1. Identify the degrees of freedom.

- ◇ 1 d.o.f. → testing one parameter (e.g. source normalization);
- ◇ 2 d.o.f. → testing two parameters (e.g. normalization and spectral index).

2. Locate your test statistic in the appropriate column.

For example, when testing one parameter, if $\chi^2 \simeq 25$, then the tail probability is approximately

$$p \approx 5.7 \times 10^{-7},$$

corresponding to a 5σ Gaussian-equivalent significance.

Figure 2.1 compares the χ_1^2 and χ_2^2 probability density functions and illustrates the “tail probability” for a given TS value (here $TS = 5$).




Figure 2.1. Tail probabilities for $TS = 5$ under χ_1^2 (orange) and χ_2^2 (blue). The shaded regions represent the probability that a noise fluctuation produces $TS \geq 5$ when the null hypothesis is true.

- (1) One free parameter (χ_1^2). A $TS = 5$ corresponds to a tail probability of about 2.3%, i.e. roughly a $\sim 2\sigma$ detection.
- (2) Two free parameters (χ_2^2). A $TS = 5$ corresponds to a tail probability of about 8.2%, i.e. roughly a $\sim 1.4\sigma$ detection.

Likelihood fit and TS evaluation in *Fermi*-LAT analysis

During a standard *Fermi*-LAT analysis, the likelihood optimization and the computation of the Test Statistic (TS) are performed in two steps.

1. Global likelihood fit

All sources within the ROI, together with the diffuse components, are included in the model. The likelihood is maximized by fitting all free parameters simultaneously, producing a global best-fit model.

2. Test Statistic (TS) computation

For each source i , its significance is evaluated by comparing two models:

- ◊ H_1 : full model including source i ,
- ◊ H_0 : model with source i removed.

All other sources remain in the model with parameters fixed at their best-fit values. The TS is

$$TS_i = 2 [\ln L(H_1) - \ln L(H_0)].$$

Summary

- ◊ Global fit: simultaneous optimization of all parameters.
- ◊ TS evaluation: only the source under test is removed; all others are kept fixed.

2.C Example of nested models

In likelihood analysis, two models are said to be *nested* when the simpler model can be obtained from the more complex one by fixing (or constraining) some of its parameters.

A typical example in *Fermi*-LAT analyses is the search for a new gamma-ray source in a given region of the sky.

- ◊ **Simple model** H_0 : no source is present; the model includes only the diffuse backgrounds.
- ◊ **Complex model** H_1 : a point source is added, described by a power-law spectrum with parameters N_0 (normalization) and fixed Γ (spectral index).

The simple model H_0 is recovered from H_1 by imposing the constraint

$$N_0 = 0,$$

i.e. by “switching off” the source. In this case the two models are nested, and the Test Statistic follows approximately a chi-square distribution with one degree of freedom (Wilks’ theorem).

2.4 LAT background

The analysis of gamma-ray data requires a proper description of the background. The *Fermi*-LAT team provides two files for modeling it: one for the Galactic diffuse emission (`gll_iem_v07.fits`), and one for the isotropic/extragalactic component (`iso_P8R3_SOURCE_V3_v1.txt`).

2.4.1 Galactic diffuse emission

At *Fermi*-LAT energies, diffuse γ -ray emission from the Milky Way dominates the sky. The Galactic diffuse emission observed by the LAT is produced when cosmic rays interact with the interstellar medium. High-energy cosmic-ray protons produce γ rays primarily through inelastic collisions with interstellar gas. In these interactions, the proton-proton collision leads to the production of neutral pions (π^0), which subsequently decay into pairs of gamma-ray photons:

$$p + p \rightarrow p + p + \pi^0 + X \quad \pi^0 \rightarrow \gamma + \gamma.$$

Here, X denotes all additional secondary particles produced in the inelastic proton–proton interaction (e.g. nucleons and charged pions), which are not explicitly written. Electrons contribute through bremsstrahlung and inverse Compton (IC) scattering on the interstellar radiation field ([Galactic Interstellar Emission Model for the 4FGL Catalog Analysis](#)).

The Galactic Interstellar Emission Model (IEM) describes this emission using spatial templates that trace the distribution of interstellar gas. Neutral atomic hydrogen (HI) is mapped using 21-cm radio surveys, molecular hydrogen (H_2) is inferred from CO emission, and the dark neutral medium (DNM) is traced through infrared dust emission. Together, these components trace the gas column density that determines where π^0 decay and bremsstrahlung occur.

The spectral shapes of each emission component are computed with GALPROP, a cosmic-ray propagation code that models the spectra of protons and electrons and their interactions with gas and radiation fields throughout the Galaxy. In the final IEM, the spatial templates (gas and IC) are combined with GALPROP spectral predictions, and their normalizations are fitted to LAT data in different Galactic regions.

Note: The spectral shapes of the diffuse components are provided by GALPROP, but their normalizations are fitted to LAT data. This is necessary because the cosmic-ray density, gas distribution, and interstellar radiation field vary across the Galaxy and are not known with perfect accuracy. Moreover, additional large-scale residual structures that cannot be directly modeled from external tracers may be included as smoothed residual maps or “patches”.




Figure 2.2. Example of a large-scale diffuse structure (often referred to informally as a “patch”): the Fermi bubbles. A giant gamma-ray structure emerges when processing *Fermi* all-sky data at GeV energies. The dumbbell-shaped feature extends tens of degrees north and south from the Galactic plane. Credit: NASA/DOE/*Fermi* LAT/D. Finkbeiner et al.

The Galactic background file used in our analysis is `gll_iem_v07.fits`. It contains a 3D data cube representing diffuse gamma-ray emission as a function of:

- ◊ Energy (28 layers/slices);
- ◊ Galactic latitude (b) (1441 pixels);
- ◊ Galactic longitude (l) (2880 pixels).

A Python example that plots a slice of the Galactic background is provided in Appendix B. By changing the variable `index`, you can display different slices. An example output is shown in Fig. 2.3.

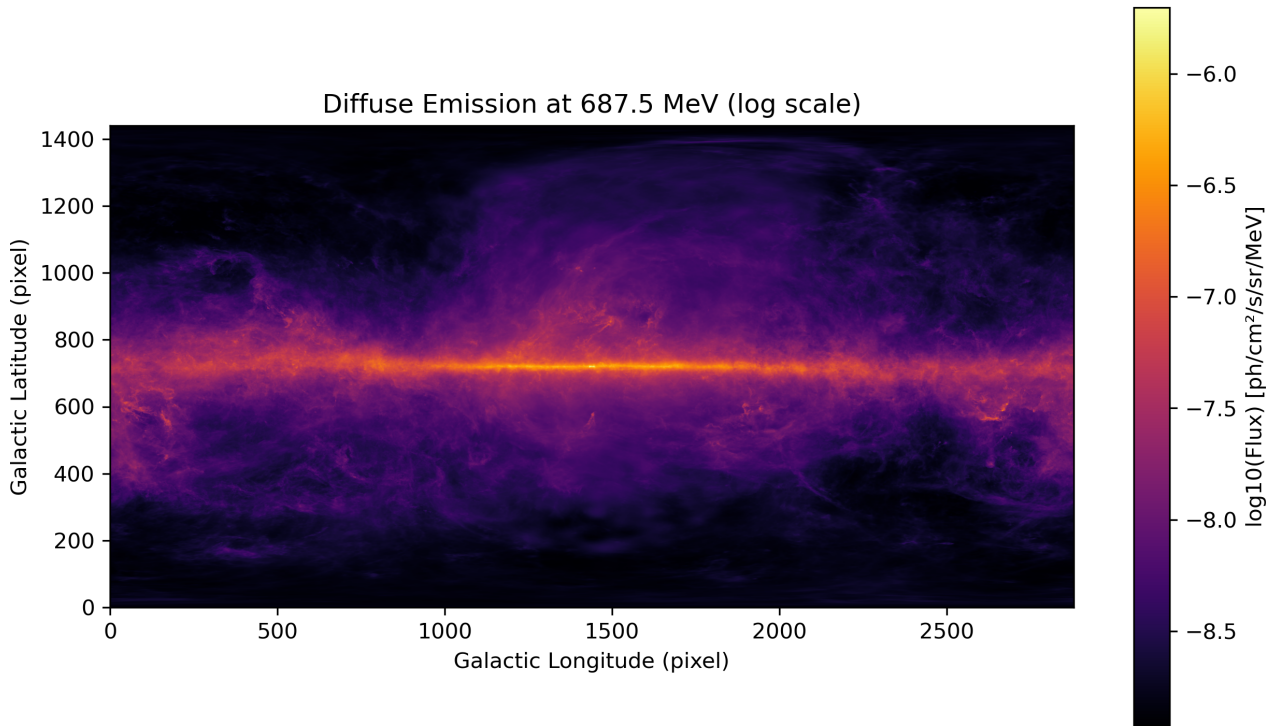


Figure 2.3. All-sky map of the gamma-ray emission observed by the *Fermi* Large Area Telescope (LAT) at ~ 7000 MeV, in Galactic coordinates (cylindrical projection). The bright horizontal band corresponds to the diffuse Galactic emission, mainly produced by cosmic rays interacting with interstellar gas and radiation fields.

A direct way to inspect the file is to use `ds9`.

2.4.2 Isotropic background

This background includes diffuse gamma rays produced outside our Galaxy, many faint extragalactic sources that cannot be resolved individually, and a small fraction of cosmic-ray particles misclassified as gamma rays. Because these components are approximately uniform across the sky, they are referred to as *isotropic emission* and must be included in almost all likelihood analyses.

To construct isotropic spectral templates, the *Fermi*-LAT team fits the all-sky gamma-ray emission while excluding the brightest regions, such as the Galactic plane ($|b| < 10^\circ$) and areas near the celestial poles ($|\text{Dec}| > 60^\circ$). These cuts reduce contamination from Galactic diffuse emission and from gamma rays produced at the Earth limb (see Section 2.4.3).

The isotropic templates are created using data in the energy range 30 MeV to 1 TeV and are extrapolated up to 3 TeV using a power-law model. Each template must be used together with the *same* Galactic diffuse model, event class, and event type with which it was generated, to ensure consistency.

Each isotropic template file contains three columns: (1) the central energy of the bin (MeV), (2) the differential flux ($\text{ph cm}^{-2} \text{s}^{-1} \text{MeV}^{-1} \text{sr}^{-1}$), (3) the $1\text{-}\sigma$ uncertainty on the differential flux. The isotropic file used in our analysis is `iso_P8R3_SOURCE_V3_v1.txt` (Fig. 2.4).

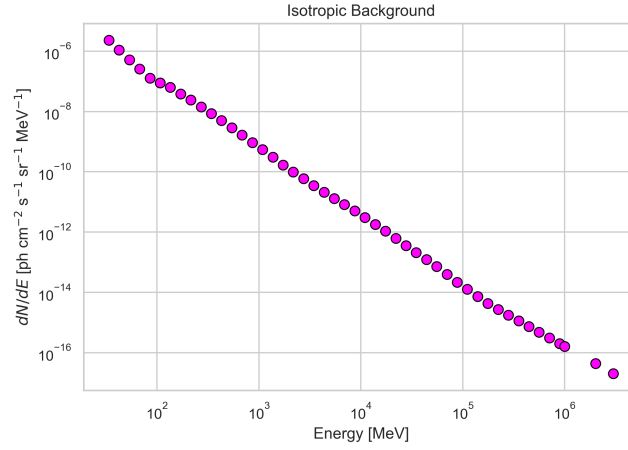


Figure 2.4. Differential flux of the isotropic background as a function of energy. The strong decline with energy indicates that the isotropic component contributes mainly at low energies and becomes much less important at higher energies.

2.4.3 Background photons produced by Earth limb emission

The *Earth albedo* gamma-ray emission is predominantly produced by inelastic interactions of cosmic-ray protons with atmospheric nuclei, leading to neutral pion production and subsequent decay into gamma-ray pairs ($\pi^0 \rightarrow \gamma\gamma$). To suppress this background contribution, events from the Earth limb are excluded by applying a maximum zenith-angle cut. Figure 2.5 illustrates the spacecraft geometry and the definition of the zenith angle used in data selection.




Figure 2.5. Geometry of the *Fermi*-LAT with respect to the Earth and to incoming photon directions. The **zenith** is the radial direction pointing away from the Earth's center through the spacecraft; it serves as the reference direction for defining the **zenith angle** θ_z , i.e. the angle between the photon direction and the zenith. Photons with large θ_z originate near the bright atmospheric limb and are removed through a zenith-angle cut in standard LAT analyses. The LAT z -axis (**boresight**) is the central pointing direction of the instrument, and the *incidence angle* θ is the angle between the photon direction and this axis. The incidence angle affects the instrument response, including the effective area and the point-spread function. The **rocking angle** θ_r is the angle between the boresight and the zenith and indicates how much the spacecraft is tilted during survey-mode operations. The **Earth limb** corresponds to lines of sight tangential to the atmosphere, where cosmic-ray interactions generate intense gamma-ray emission.

2.5 Spectral models

Gamma-ray sources in *Fermi*-LAT analyses are often described by a power-law spectral model:

$$\frac{dN}{dE} = N_0 \left(\frac{E}{E_0} \right)^{-\Gamma},$$

where dN/dE is the differential photon flux in units of $\text{ph cm}^{-2} \text{s}^{-1} \text{MeV}^{-1}$, N_0 is the normalization at the pivot energy E_0 , and Γ is the photon index.

The parameter E_0 , referred to as the *pivot energy* (or **Scale** in many LAT model files), is not a physical characteristic of the source. It is a reference energy at which the normalization N_0 is defined. The choice of E_0 is made to minimize the statistical correlation between N_0 and Γ . Selecting an appropriate pivot energy improves the numerical stability of the likelihood fit and leads to more reliable parameter uncertainties. Importantly, the value of E_0 does *not* change the physical shape of the spectrum; it only sets the energy at which the normalization is evaluated.

For bright sources with high statistics, a simple power law may be insufficient and a more complex spectral model may be required. Available models are listed [here](#). In the case of TXS 0506+056, the LAT spectrum is often modeled with a log-parabola:

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

where N_0 is the normalization at E_0 , α is the photon index at E_0 , and β quantifies the curvature (larger β means stronger curvature). Here, \log denotes the natural logarithm.

Figure 2.6 shows two sets of log-parabola parameters with different curvature.




Figure 2.6. Log-parabola spectra illustrating the effect of curvature in LAT spectral models. Both curves are computed for a pivot energy $E_0 = 1 \text{ GeV}$, normalization $N_0 = 2 \times 10^{-11} \text{ ph cm}^{-2} \text{s}^{-1} \text{MeV}^{-1}$, and $\alpha = 2.0$. The two values of the curvature parameter, $\beta = 0.1$ (mild curvature) and $\beta = 0.3$ (strong curvature), show how the spectrum steepens more rapidly with increasing energy as β increases.

2.5.1 ROI model

The initial model for a Region of Interest (ROI) of radius R , centered on the target source, can be generated using the script `make4FGLxml.py` included in the *Fermi* Science Tools. Once the target coordinates are provided (thus defining the ROI center), the script queries the 4FGL catalog and extracts all gamma-ray sources within the radius R (see Fig. 2.7). For each selected source, the script imports the catalog best-fit spectral and spatial model and then adds the diffuse background components. The resulting model is written to an XML file. An example of a LAT XML model file is shown in Appendix C.

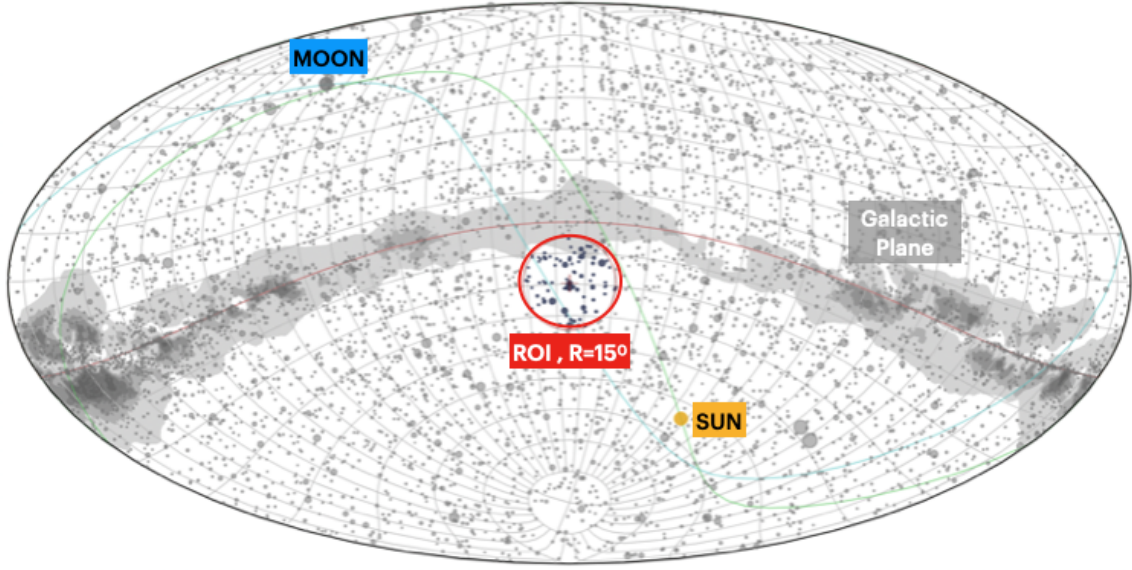


Figure 2.7. All-sky map in equatorial coordinates showing the region of interest (ROI; red circle) with a radius of 15° centered on TXS 0506+056. The grey structure traces the Galactic plane. The apparent trajectories of the Sun (green) and the Moon (cyan) are also shown (from the [Fermi LAT Light Curve Repository](#)).

2.D From the differential spectrum dN/dE to the spectral energy distribution (SED)

- (1) $dN/dE \rightarrow$ photon spectrum $[\text{ph cm}^{-2} \text{s}^{-1} \text{MeV}^{-1}]$
- (2) $E dN/dE \rightarrow$ photon SED $[\text{ph cm}^{-2} \text{s}^{-1}]$
- (3) $E^2 dN/dE \rightarrow$ energy SED in MeV $[\text{MeV cm}^{-2} \text{s}^{-1}]$
- (4) $E^2 dN/dE \times 1.602 \times 10^{-6} \rightarrow$ energy SED in erg $[\text{erg cm}^{-2} \text{s}^{-1}]$ ($1 \text{ MeV} = 1.602 \times 10^{-6} \text{ erg}$)

Figure 2.8 shows an example photon spectrum and the corresponding energy SED. A broken power law is used to highlight the presence of a spectral break (a change in photon index), which becomes more evident in the SED representation.

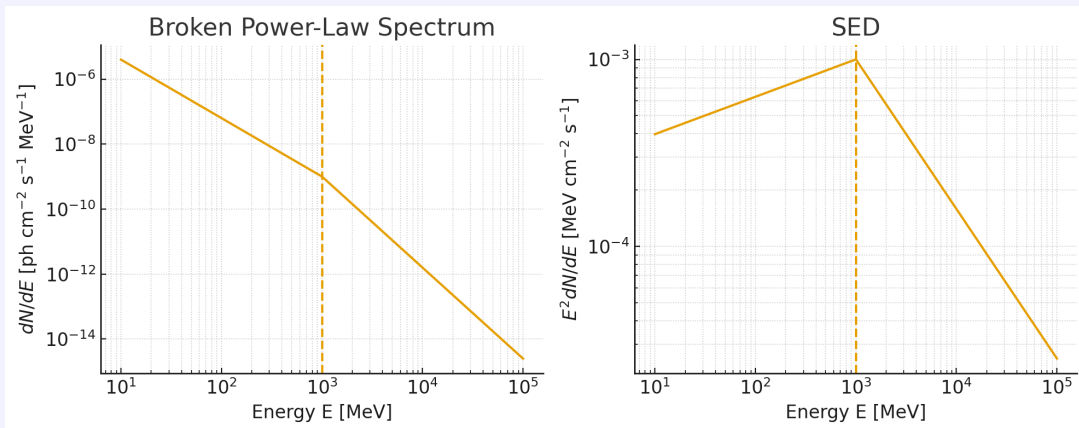


Figure 2.8. Example of a photon spectrum described by a broken power law (left panel) and the corresponding energy SED (right panel). The spectrum follows $dN/dE = N_0 (E/E_{\text{break}})^{-\Gamma_1}$ for $E < E_{\text{break}}$ and $dN/dE = N_0 (E/E_{\text{break}})^{-\Gamma_2}$ for $E \geq E_{\text{break}}$, with $E_{\text{break}} = 1 \text{ GeV}$, $\Gamma_1 = 1.8$, and $\Gamma_2 = 2.8$. The right panel shows the corresponding spectral energy distribution, computed as $E^2 dN/dE$.

A walkthrough of the data reduction process

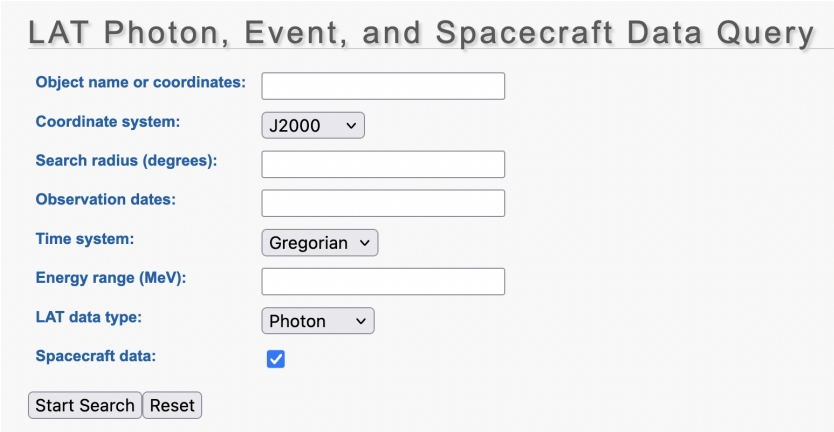
In this chapter, the main steps of gamma-ray data reduction and analysis (spectrum and light curve) are presented. The analysis threads are documented on the corresponding [website](#). Since this exercise uses the `fermipy` software within a Jupyter notebook, a useful starting point for learning how to perform a Fermi-LAT analysis with `fermipy` is the official documentation, available at [Fermipy](#). Below, only the essential steps of the analysis workflow are summarized.

Source Analysis with Fermi LAT Data

- (1) Download Event data file and Spacecraft data file
- (2) Prepare the data
- (3) Run the likelihood analysis
- (4) Produce the SED
- (5) Produce a light curve

3.1 Download the data

The event and spacecraft data are retrieved from the LAT data [server](#). Figure 3.2 illustrates the layout of the data access portal.



The screenshot shows a web form titled "LAT Photon, Event, and Spacecraft Data Query". It contains several input fields and dropdown menus for specifying search parameters. The fields are: "Object name or coordinates:" (text input), "Coordinate system:" (dropdown menu with "J2000" selected), "Search radius (degrees):" (text input), "Observation dates:" (text input), "Time system:" (dropdown menu with "Gregorian" selected), "Energy range (MeV):" (text input), "LAT data type:" (dropdown menu with "Photon" selected), and "Spacecraft data:" (checkbox, which is checked). At the bottom of the form are two buttons: "Start Search" and "Reset".

Figure 3.1. Layouy of the LAT data access

- ◇ The “search radius” field in the portal specifies the circular sky region in which the analysis will be performed. A value of 15–20 degrees is typically recommended for the region of interest.
- ◇ Observation dates may be specified in the Gregorian calendar system, as a number of seconds in the MET (Mission Elapsed Time) system, or as Modified Julian Dates (MJD).

The Fermitools interface uses **Mission Elapsed Time (MET)**, defined as the number of seconds since 00:00:00 UTC on January 1, 2001. During analysis it is often necessary to convert MET to other time

3.1. Download the data

expressions such as calendar date, day of year, or Modified Julian Date (MJD). A convenient online conversion tool is provided by HEASARC:

<https://heasarc.gsfc.nasa.gov/cgi-bin/Tools/xTime/xTime.pl>

- ◇ Energy range (MeV): an interval of 100 MeV to 1000000 MeV (1 TeV) is typically recommended for this type of analysis.
- ◇ Photon data include the “Source” event class, which is the one we are interested in.

The output of the download procedure generally consists of files similar to those listed below.

Filename	Number of Entries	Size (MB)	Status
L150622221409E3651F7F79_SC00.fits	445088	64.54	Available
L150622221409E3651F7F79_PH00.fits	94727	8.90	Available
L150622221409E3651F7F79_PH01.fits	167258	15.68	Available

If you would like to download the files via wget, simply copy the following commands and paste them into a terminal window. The files will be downloaded to the current directory in the terminal window.

```
wget http://fermi.gsfc.nasa.gov/FTP/fermi/data/lat/test/queries/L150622221409E3651F7F79_SC00.fits
wget http://fermi.gsfc.nasa.gov/FTP/fermi/data/lat/test/queries/L150622221409E3651F7F79_PH00.fits
wget http://fermi.gsfc.nasa.gov/FTP/fermi/data/lat/test/queries/L150622221409E3651F7F79_PH01.fits
```

Figure 3.2. Downloaded files

Several *PH0*.fits files may be generated, each containing part of the event data, together with a housekeeping file *SC*.fits. All *PH*.fits files must be listed in a single *.txt file, which will then be used as input for the data reduction (see the Jupyter notebook).

A quick look at the data can be obtained using TOPCAT. The *PH* FITS files obtained from the Fermi portal contain an EVENTS table in which each row corresponds to a detected photon. The most relevant columns for the analysis are listed below.

- ◇ **ENERGY (MeV)**: reconstructed photon energy.
- ◇ **RA, DEC (deg)**: equatorial coordinates of the event.
- ◇ **THETA (deg)**: angle between the reconstructed photon direction and the LAT boresight.
- ◇ **PHI (deg)**: azimuthal angle in instrument coordinates.
- ◇ **ZENITH_ANGLE (deg)**: angle between the photon direction and the Earth zenith (important for Earth–limb filtering).
- ◇ ...
- ◇ **TIME (s)**: event arrival time in MET (Mission Elapsed Time).
- ◇
- ◇ **EVENT_CLASS**:
- ◇ **EVENT_TYPE**
- ◇ **CONVERSION_TYPE**: 0 for “Front” conversion, 1 for “Back”.
- ◇ ...

The *SC*.fits file stores all quantities describing the spacecraft attitude, position, magnetic environment, data quality, and live time. The spacecraft file collects the information for 30 second intervals, but some intervals may be shorter. The column live time gives the effective observing time (in seconds) for each interval. These parameters are needed to compute exposure, apply the IRFs correctly, exclude bad intervals (e.g., SAA passages), and track how the LAT was oriented while collecting data.

3.A Van Allen radiation belts and South Atlantic Anomaly (SAA)

The Earth is surrounded by regions of charged particles trapped by the geomagnetic field, known as the **radiation belts**. They consist of two main regions, commonly referred to as the *inner* and *outer* Van Allen belts. The inner belt is dominated primarily by energetic protons, whereas the outer belt contains mainly high-energy electrons (*upper panel*).

The South Atlantic Anomaly (SAA) is a region over the South Atlantic where the inner Van Allen radiation belt approaches unusually close to Earth's surface (*lower panel*). As a result, the particle flux in this area is extremely high. When the spacecraft passes through the SAA, the LAT is exposed to an intense background of charged particles and must temporarily suspend scientific data taking.

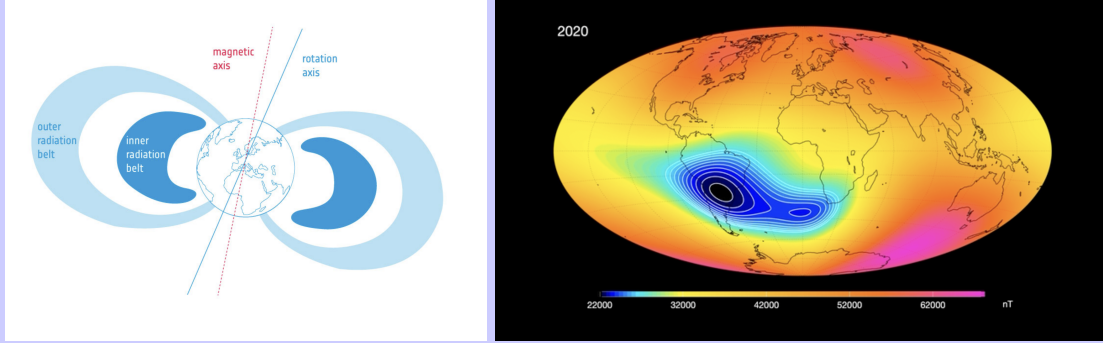


Figure 3.3. Van allen belts (*left panel*) and South Atlantic Anomaly (*right panel*)

3.2 Data preparation and and Likelihood Analysis

After applying the event selection cuts to the event file according to the configuration specified in the YAML input file, the subsequent steps of the analysis (details in Appendix D) aim to the construction of the (`ccube.fits`) count cube and the (`srcmap_00.fits`).

The count cube is a three-dimensional histogram of the LAT photon data. The two spatial axes define the Region of Interest (ROI), while the third axis consists of the energy bins specified by the user. Each element of the cube, Counts (i, j, k), represents the number of photons detected in sky pixel (i, j) and energy bin k . The ccube contains only observed counts. It is possible to generate a sequential visualization of the individual slices at different energies by means of DS9.

The srcmap map is defined on the same spatial and energy grid as the counts cube (`ccube.fits`). It represents the model prediction of the expected counts. For every source listed in the XML file, the software calculates how many photons should be detected in each sky pixel and in each energy bin. This calculation uses several ingredients: the source position and spectrum (from the XML model), the LAT exposure, the instrument response functions (IRFs), the point-spread function (PSF), and the selected event type.

During the likelihood analysis, the predicted counts are compared with the observed ones in order to determine the best-fitting parameters of the sources.

The likelihood analysis allows us to derive the best fit parameters of the sources in the ROI (global fit) as well as the TS for the source of interest. The model of the best fit is generally written in a new XML. Below the example of very faint source is shown. Each row of the output for the analyzed source is listed in Table 3.1

A detailed tutorial on the data analysis procedure can be found [here](#).

```
print(gta.roi['4FGL J0730.5+6720'])
Name          : 4FGL J0730.5+6720
Associations  : ['4FGL J0730.5+6720', 'WISE J073125.40+671847.0']
RA/DEC       : 112.642/ 67.348
GLON/GLAT    : 148.550/ 28.623
TS           : 4.29
Npred        : 55.99
Flux         : 1.254e-08 +/- 6.75e-09
EnergyFlux   : 2.934e-06 +/- 1.58e-06
SpatialModel : PointSource
SpectrumType : PowerLaw
Spectral Parameters
b'Prefactor' : 1.949e-12 +/- 1.05e-12
b'Index'     : -2.741 +/- nan
b'Scale'     : 558.9 +/- nan
```

Parameter	Meaning
Name	Identifier of the source in the model.
Associations	Catalog objects associated with the γ -ray source.
RA/DEC	Right Ascension and Declination (J2000) of the source position.
GLON/GLAT	Galactic longitude and latitude of the source position.
TS	Test Statistic, measuring the detection significance of the source (see section 2.3)
Npred	Number of photons predicted by the model for this source over the full energy range (see section 2.2).
Flux	Integrated photon flux ($\text{ph cm}^{-2} \text{s}^{-1}$).
EnergyFlux	Integrated energy flux ($\text{MeV cm}^{-2} \text{s}^{-1}$) (see section 2.4.4).
SpatialModel	Spatial model used to describe the source (e.g. <i>PointSource</i>).
SpectrumType	Spectral model adopted for the source (e.g. <i>PowerLaw</i> , <i>Log-Parabola</i>).
Prefactor	Normalization of the spectral model at the pivot energy.
Index	Spectral index of the model (slope of the power law).
Scale	Pivot energy at which the normalization is defined.
Note: A value of nan indicates that the parameter is frozen (not varied during the fit).	

Table 3.1. Description of the parameters appearing in the likelihood fit output.

3.3 Generation of Scientific Products

There are two main ways to study a gamma-ray source: the Spectral Energy Distribution (SED) and the light curve. Both rely on maximumlikelihood fitting, but they divide the data in different dimensions.

3.3.1 Spectral Energy Distribution

When computing an SED, the entire observation time is used. The analysis is divided into energy bins, and in each energy bin the fit estimates only the normalization of the source. This allows us to reconstruct how the flux depends on energy. In principle, one could also allow the spectral index to vary in each bin. However, this requires very high statistics: in most realistic cases there are not enough photons per energy bin to constrain both the normalization and the index. For this reason, the index is usually kept fixed when constructing an SED. An example of a spectral energy distribution derived with is shown in Figure 3.3.2-(*Upper Panel*).

SED: Basic Concep

In simple terms:

- ◊ the time interval is fully integrated;
- ◊ the data are divided into energy bins;
- ◊ a separate likelihood fit is performed for each energy bin.

3.3.2 Light Curve

For light curves, the logic is similar to that used for building the SED: the data are divided into time bins, and in each interval the full energy range is considered. Only the normalization of the model is left free to vary, so that the flux can change with time. Allowing the spectral index to vary independently in each time bin would again require a large number of photons per bin and is therefore usually not feasible.

Light Curve: Basic Concept

In simple terms:

- ◊ the **energy range is fully integrated**;

- ◇ the data are divided into **time intervals**;
- ◇ a separate **likelihood fit** is performed for each time bin.

Final note to keep in mind – During the computation of both the SED and the light curve, only the parameters of the target source are left free in each energy or time bin (in most cases, only the normalization is allowed to vary). In general, all the other sources in the Region of Interest keep the parameters obtained from the global likelihood fit and remain fixed.

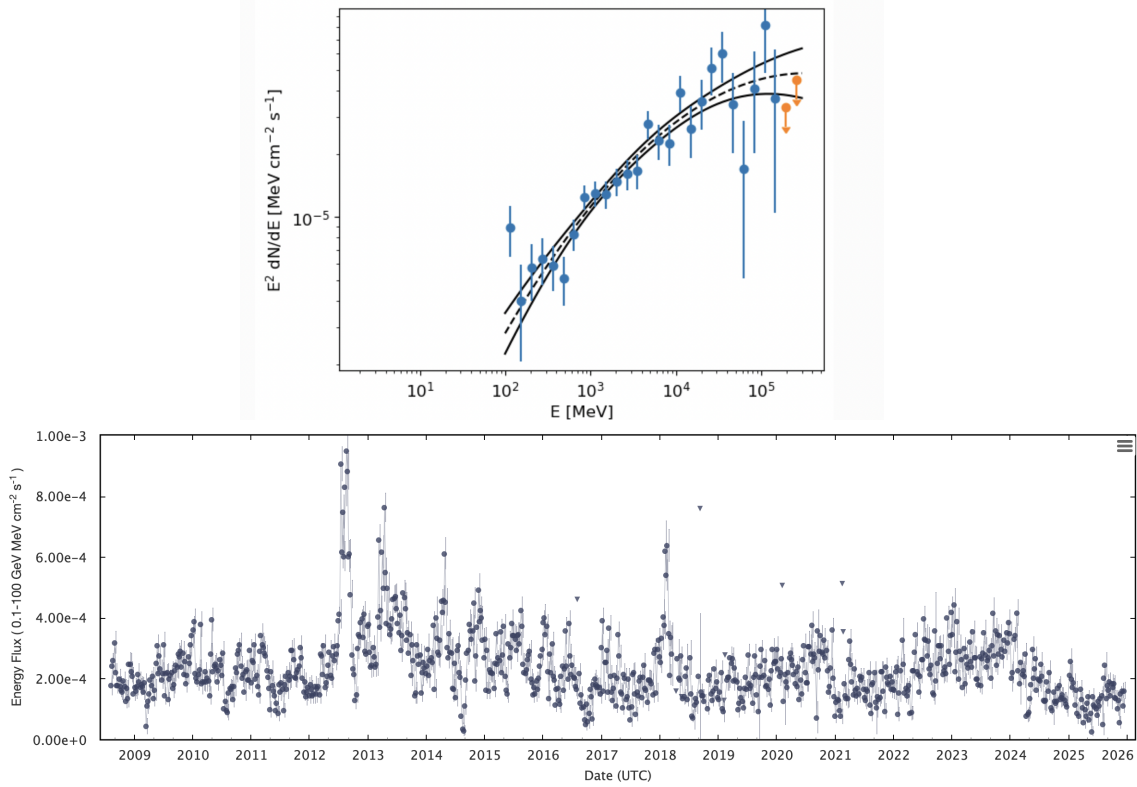


Figure 3.4. *UpperPanel* – Spectral Energy Distribution.

Data points represent the flux measured in each energy bin and the vertical bars show the statistical uncertainties (1σ). Upper limits (95%) are plotted for bins in which the detection significance is low. The solid curve corresponds to the best fit spectral model obtained from the likelihood analysis (global fit), and the shaded band illustrates the 1σ confidence region of the model prediction.

Lower Panel – Light curve, computed with 1-week temporal bins in the 100 MeV~100 GeV band taken from the [Fermi-LAT Light Curve Repository](#)

3.4 Parameter uncertainty estimation

In a given region of interest (ROI), the likelihood function depends on many parameters describing the target source, the surrounding sources, and the diffuse background components. As discussed in Sect. 2.2.1, the likelihood \mathcal{L} is maximized by varying the free parameters within the ROI. In practice, most analysis software works with the transformed quantity

$$-2\log \mathcal{L},$$

which is numerically more stable and has a direct connection with the χ^2 distribution (see Sect. 2.3).

When the goal is to estimate the uncertainty on a *single* parameter of the target source—such as its normalization—the problem can be reduced to the study of that parameter alone. Let us denote this parameter by θ . The uncertainty on θ is determined by examining how the likelihood changes as θ is displaced from its best-fit value $\hat{\theta}$. Near the maximum, the function $-2\log \mathcal{L}$ is well approximated by a parabola, which allows confidence intervals to be defined in terms of changes in its value. For each trial value of θ , all other free parameters in the model are re-optimized. This procedure defines the *profile likelihood*.

According to Wilks' theorem, the difference

$$\Delta(-2\log \mathcal{L})$$

follows a χ^2 distribution under fairly general conditions. This result provides simple and widely used criteria for the construction of confidence intervals. For a single parameter of interest,

$$\Delta(-2\log \mathcal{L}) = 1 \quad \Rightarrow \quad 68\% \text{ confidence interval,}$$

while for two parameters varied simultaneously the corresponding threshold is

$$\Delta(-2\log \mathcal{L}) = 2.30 \quad \Rightarrow \quad 68\% \text{ confidence region.}$$

These values are standard and follow directly from the properties of the χ^2 distribution. An example of a likelihood profile is shown in Fig. 3.5, which illustrates the behaviour of $\Delta(-2\log \mathcal{L})$ as a function of the parameter θ (e.g. the normalization of the target source spectrum).

Profile likelihood and symmetric 68% C.L. errors (good fit case)

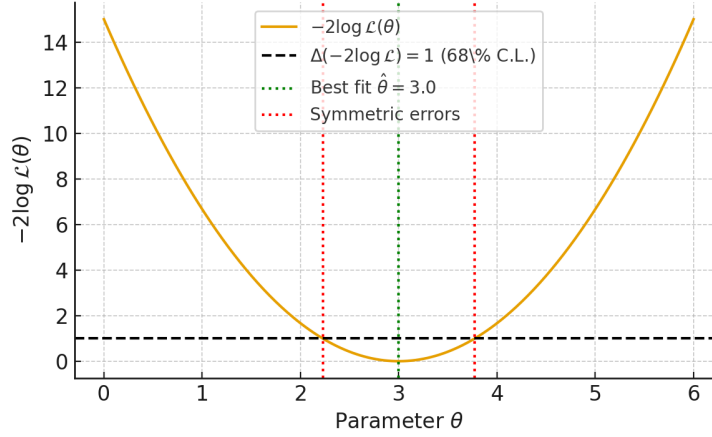


Figure 3.5. Example of a likelihood profile expressed as $\Delta(-2\log L)$ as a function of the parameter θ . The curve reaches its minimum at the best-fit value θ_0 . The horizontal dashed line marks the level $\Delta(-2\log L) = 1$, which corresponds to the 68% confidence interval for one fitted parameter according to Wilks' theorem. The intersection of this line with the likelihood curve defines the lower and upper bounds of θ .

3.B Different Uses of the Test Statistic (TS)

The quantity $\Delta(-2\log \mathcal{L})$ used to define confidence intervals is formally the same structure that appears in the definition of the Test Statistic (TS) widely used in Fermi-LAT analyses. The TS is defined as

$$\text{TS} = -2\log \left(\frac{\mathcal{L}_0}{\mathcal{L}_1} \right) = 2 [\log \mathcal{L}_1 - \log \mathcal{L}_0],$$

where \mathcal{L}_1 is the maximum likelihood of the model including the source and \mathcal{L}_0 is the likelihood under the null hypothesis (e.g. without the source). In this case the TS measures how much the likelihood improves when a parameter associated with the source is allowed to vary from its null value.

The quantity, $\Delta(-2\log \mathcal{L})$ used to define confidence intervals, measures how much the likelihood worsens when the parameter is displaced from its best-fit value. Both quantities follow a χ^2 distribution according

to Wilks' theorem, and both quantify the curvature of the likelihood surface, providing complementary tools for source detection and parameter uncertainty estimation.

3.4.1 Upper Limits

In some bins of the SED (or of the light curve), the source is not significantly detected if

- ◇ the Test Statistic (TS) is low (typically $TS \leq 4$),
- ◇ the best-fit flux is compatible with zero,
- ◇ the likelihood profile does not show a clear minimum.

In these cases the likelihood curve does not define a reliable symmetric error estimate. The data point is then substituted by an upper limit. The upper limit estimation uses exactly the same likelihood-based logic as two-sided error estimates, but with a *one-sided* confidence threshold:

$$\Delta(-2\log \mathcal{L}) = 2.71,$$

which corresponds to a 95% confidence upper limit for one free parameter.

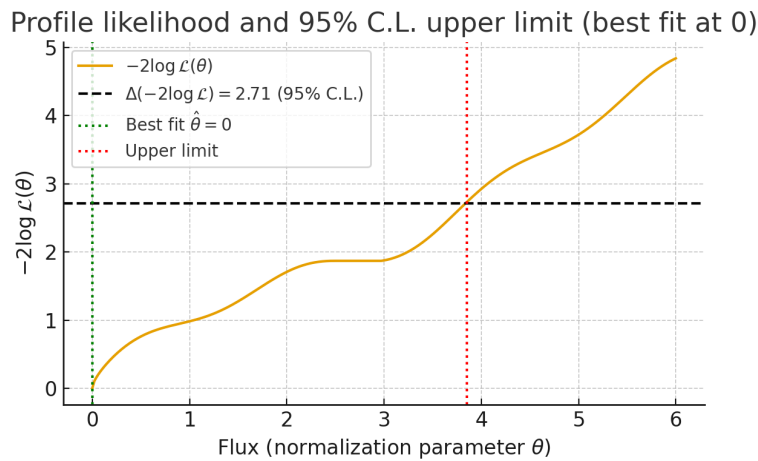


Figure 3.6. Example of an upper-limit determination using the profile likelihood. When the best-fit flux is $\hat{\theta} = 0$, the likelihood has no symmetric minimum and only an upper limit can be derived. The 95% confidence limit corresponds to the point where the curve rises by $\Delta(-2\log \mathcal{L}) = 2.71$ from its minimum value.

3.4.2 Another use of the TS: Variability Test

In the analysis of gamma-ray light curves from the Fermi-LAT, it is often necessary to determine whether a source is consistent with a constant flux or shows significant variability. To address this question, the LAT team uses the variability test statistic TS_{var} . The idea is to compare two models: a constant-flux model, in which the source has a single flux value F_{const} for all time bins, and a variable model, in which each time bin has its own flux F_i . If the source is intrinsically constant, the constant-flux model should fit the data almost as well as the variable one; if the source is variable, the variable model will provide a substantially higher likelihood. The variability statistic is defined as

$$TS_{\text{var}} = 2 \sum_{i=1}^N [\ln \mathcal{L}_i(F_i) - \ln \mathcal{L}_i(F_{\text{const}})] + C$$

where the first term quantifies the improvement in likelihood obtained when each time bin is allowed to vary independently, while C is a corrective term introduced ad hoc to stabilise the statistic (see [4FGL catalog paper](#)). Under the assumption of a constant source flux, the variability index TS_{var} is expected to follow a chi-square distribution with $N - 1$ degrees of freedom. The Fermi community has chosen to adopt a p -value of

$p = 0.01$ as the threshold to distinguish between sources that are consistent with being constant and sources that show significant variability. For example, if $N = 50$ (50 time bins), this p -value threshold corresponds to $TS_{\text{var}} \geq 76.78$. Therefore, if the variability test statistic TS_{var} of a light curve exceeds this value, the source is considered to be inconsistent with a constant flux. Other examples are in Table 3.2.

Table 3.2. Example of the variability-index thresholds corresponding to $p = 0.01$ for different time bins.

N_{bin}	dof	TS_{var} (for $p = 0.01$)
10	9	21.67
50	49	76.78
100	99	135.81

3.4.3 Variability, Timescales, and Source Size

Variability is a powerful diagnostic tool in high-energy astrophysics, as changes in the observed flux provide direct constraints on the size and physical nature of the emitting region. These constraints follow from causality: photons emitted from different parts of an extended region reach the observer at different times, and rapid variability cannot be produced by regions larger than the distance light can travel over the observed variability timescale. This idea can be illustrated by considering a flare that rises or decays significantly within a short time interval. If the emitting region were very large, differences in light-travel time across the region would smear out the variability, making the flare appear broader and less sharply defined. The observation of a rapid and well-defined change in flux therefore implies that the emitting region must be sufficiently compact for light-travel-time effects to be negligible on the observed timescale.

This is the essence of the light-crossing argument, which yields the fundamental constraint

$$R \lesssim c \Delta t'.$$

The timescale we measure, Δt_{obs} , is not equal to the intrinsic one $\Delta t'$. The observed variability timescale is affected both by relativistic Doppler effects (see section 5.1) and by cosmological time dilation. The relativistic motion of the emitting region contracts intrinsic timescales by a factor δ^{-1} , while the expansion of the Universe dilates any timescale by a factor $(1 + z)$.

As a result, the observed and intrinsic timescales are related by

$$\Delta t_{\text{obs}} = \frac{\Delta t'}{\delta} (1 + z).$$

Solving for the intrinsic timescale yields

$$\Delta t' = \frac{\Delta t_{\text{obs}} \delta}{1 + z}.$$

Substituting this into the light-crossing argument provides the frequently used source-size constraint:

$$R \lesssim \frac{c \Delta t_{\text{obs}} \delta}{1 + z}.$$

Estimating the variability timescale from the light curve

The variability timescale Δt_{obs} is estimated from the observed light curve. A commonly used definition is the doubling time, based on an exponential model of the flux:

$$F(t) = F_0 e^{t/t_{\text{var}}}.$$

For two flux measurements F_1 and F_2 separated by an interval Δt ,

$$\frac{F_2}{F_1} = e^{\Delta t/t_{\text{var}}},$$

which yields

$$t_{\text{var}} = \frac{\Delta t}{\ln(F_2/F_1)}.$$

For example, if the flux doubles in one hour,

$$\ln(F_2/F_1) = \ln(2) \approx 0.693$$

the characteristic exponential timescale is

$$t_{\text{var}} \approx \frac{1}{0.693} \approx 1.44 \text{ hours.}$$

NOTE that the estimate of the variability timescale t_{var} is reliable only when the difference between the two flux measurements is statistically significant:

$$|F_2 - F_1| \geq 3\sqrt{\sigma_1^2 + \sigma_2^2},$$

otherwise, the inferred timescale may simply reflect statistical fluctuations rather than intrinsic source variability.

In the optional laboratory exercise, we assume that, due to the intrinsic uncertainties in our estimates, the flux-doubling timescale can be approximated as $t_{\text{var}} \sim \Delta T$.

Part II

II: Basic Astrophysical Concepts on Jetted AGN for the Lab Session

Supermassive Black Holes and Nuclear Activity: *Accretion Flows and Jets*

Although most galaxies host a supermassive black hole, the vast majority of them are presently inactive. Only about **1% of galaxies** show a phase of nuclear activity strong enough to be classified as an AGN. Among these active systems, only a small fraction is capable of producing relativistic plasma jets. In the local Universe, roughly **10% of AGN** launch powerful relativistic outflows extending from parsec to megaparsec scales. These jet-producing systems are found primarily in massive elliptical galaxies and are referred to as **jetted AGN**.

4.1 Accretion Processes

An AGN is powered by matter accreting onto a supermassive black hole. The structure of an accretion flow around a black hole is governed by four fundamental conservation equations: conservation of mass, radial and vertical momentum, angular momentum, and energy. These equations admit different solutions, the two most relevant for AGN being the **Shakura–Sunyaev disk** and the **ADAF** (Advection-Dominated Accretion Flow).

4.1.1 The Shakura–Sunyaev Disk: A Radiatively Efficient Solution

In the **Shakura–Sunyaev model** regime, the disk is:

- ◇ **geometrically thin**, with vertical height H much smaller than the radius R ($H/R \ll 1$);
- ◇ **optically thick**, so that matter and radiation remain in local thermal equilibrium;
- ◇ **radiatively efficient**, meaning that most of the gravitational energy released by accretion is immediately emitted.

Accretion can only proceed if angular momentum is transported outward. In the thin disk, this effect is described by the viscosity (encoded in the α parameter) that causes neighbouring annuli of the disk to exchange angular momentum. Inner annuli lose angular momentum and drift inward, while outer annuli gain it and move outward. As gas spirals inward, it loses gravitational potential energy. This energy is efficiently *dissipated as heat* by viscous stresses and then *radiated locally*. These concepts are discussed intuitively in Appendix E.

Because the disk is optically thick, each radius emits roughly as a black body, with the temperature increasing towards smaller radii. The full disk therefore produces a *multi-temperature black-body spectrum*. This radiatively efficient thin disk is characterized by accretion rates $\dot{m} \gtrsim 10^{-2}$. A schematic view of an AGN with a Shakura–Sunyaev disk and its nuclear environment is shown in Figure 4.1

4.1.2 The ADAF Solution: A Radiatively Inefficient Flow

In an **Advection-Dominated Accretion Flow**, the particle density is very low. Coulomb interactions between ions and electrons are extremely inefficient, so the plasma cannot thermalise. As a consequence:

$$T_{\text{ions}} \sim 10^{12} \text{ K}, \quad T_e \sim 10^9 - 10^{10} \text{ K},$$

and the flow develops a *two-temperature structure* (see Figure 4.2). The electrons can still cool via synchrotron radiation, bremsstrahlung, and inverse Compton scattering, whereas the ions retain most of the heat. Most of the energy is thus advected into the black hole rather than emitted. Due to the high internal pressure of the hot gas, the flow inflates vertically, so that the scale height becomes comparable to the radial extent,

$$H/R \sim 1.$$




Figure 4.1. A schematic diagram of the current paradigm for radio-loud AGN with a Shakura–Sunyaev disk (not to scale). A luminous accretion disk surrounds the central black hole. Broad emission lines are produced in clouds orbiting above the disk and possibly by the disk itself. A thick dusty torus (or warped disk) obscures the broad-line region for transverse lines of sight; some continuum and broad-line emission can be scattered into those directions by hot electrons permeating the region. A hot corona above the accretion disk may also contribute to the hard X-ray continuum. Narrow emission lines are produced in clouds located much farther from the central source. Radio jets emanate from the region near the black hole, travelling at relativistic speeds. For a $10^8 M_\odot$ black hole, the event-horizon radius is $\sim 3 \times 10^{13}$ cm; the accretion disk emits predominantly from $\sim (1 - 30) \times 10^{14}$ cm; the broad-line clouds lie within $\sim (2 - 20) \times 10^{16}$ cm; and the inner radius of the dusty torus is around 10^{17} cm. The narrow-line region extends from $\sim 10^{18}$ to 10^{20} cm, while radio jets have been detected on scales from 10^{17} cm to more than several $\times 10^{24}$ cm—about an order of magnitude larger than the biggest galaxies (Urry and Padovani 1995).

ADAFs are relevant for accretion rates $\dot{m} \lesssim 10^{-3}$.

Unlike AGN powered by standard, radiatively efficient accretion (i.e. Shakura–Sunyaev disk), AGN hosting an ADAF show no observational evidence for the presence of either a broad-line region (BLR) or a dusty torus. However, both accretion regimes are capable of launching jets, although the fraction of gravitational energy partitioned between jet power and radiative emission differs substantially between these systems.




Figure 4.2. ADAF schematic view as proposed by Rees et al 1982.

4.1.3 How AGN Jets Are Produced: A intuitive explanation

Active galactic nuclei (AGN) can produce powerful, relativistic jets that extend for thousands or even millions of light-years. The key ingredients are: a spinning black hole, a magnetized accretion flow, and the interaction between the magnetic field and the curved spacetime around the black hole.

The gas spiraling toward the black hole carries magnetic field lines with it because the magnetic field is frozen into the plasma. As the accretion continues, the poloidal (vertical) magnetic field becomes concentrated

near the black hole. This vertical field is essential: it provides the “backbone” along which the eventual jet will be guided and collimated. If the black hole is rotating, general relativity predicts that the spacetime around it is dragged into rotation as well (frame dragging). This means that anything located in the inner region—gas, electrons, or magnetic field lines—is forced to rotate. As a consequence, the initially vertical magnetic field lines become twisted by this differential rotation.




Figure 4.3. Schematic representation of the poloidal and toroidal of the magnetic field.

When a poloidal field rotated, it naturally develops a toroidal (azimuthal) component, $B_{tor} = B_\phi$ (from now on, we use B_ϕ instead of B_{tor} to align with the notation commonly adopted in the literature). Near spinning black hole, the twisting is extremely strong, so the toroidal field can grow to very large values. In magnetohydrodynamics (MHD), the toroidal field represents stored magnetic energy (see Figure 4.3). A strong toroidal field produces magnetic pressure proportional to

$$P_{\text{mag}} \propto \frac{B_\phi^2}{8\pi}.$$

Plasma can move freely along magnetic field lines but not across them. Therefore, the poloidal field lines act like rails that guide the outflow. As the magnetic pressure pushes outward, plasma is accelerated along these lines, forming a narrow, collimated jet (Figure 4.1).

4.A Energy Extraction from a Rotating Black Hole: The Blandford–Znajek Mechanism

When a rotating black hole is embedded in a magnetized environment, the interaction between spacetime rotation and magnetic fields can give rise to a powerful electromagnetic outflow. A key role in this process is played by the black hole *event horizon*. The event horizon is the boundary surrounding a black hole beyond which no information or radiation can escape to a distant observer.

For a rotating (Kerr) black hole, the event horizon has a finite radius, r_H , and is characterized by a well-defined angular velocity, Ω_H . Due to frame dragging, spacetime in the vicinity of the horizon is forced to rotate, and any matter or magnetic field threading the horizon must co-rotate with it. This property allows the horizon to interact electromagnetically with the surrounding plasma and makes it possible to extract rotational energy from the black hole.

If a large-scale poloidal magnetic field threads the event horizon, the magnetic flux associated with this field is defined as

$$\Phi_{\text{BH}} \equiv \int_{\text{horizon}} \mathbf{B}_{\text{pol}} \cdot d\mathbf{A},$$

where \mathbf{B}_{pol} is the poloidal component of the magnetic field and $d\mathbf{A}$ is the oriented surface element of the horizon. The integral is taken over the entire horizon surface. [The surface element $d\mathbf{A}$ represents an infinitesimal patch of the event horizon surface (which can be approximated, for simplicity, as a sphere), with magnitude equal to the area of the patch and direction normal to the horizon. The magnetic flux Φ_{BH} is obtained by summing the contributions of the magnetic field crossing all such surface elements over the entire horizon]. The magnetic flux Φ_{BH} therefore measures the total amount of poloidal magnetic field anchored to the black hole and quantifies how strongly the black hole is magnetically coupled to its environment.

The rotation of spacetime twists the poloidal magnetic field lines threading the horizon, generating a toroidal magnetic field component, B_ϕ . In the Blandford–Znajek framework, the strength of this toroidal component scales as

$$B_\phi \sim \Omega_H r_H B_{\text{pol}},$$

where Ω_{H} is the angular velocity of the horizon. For a Kerr black hole,

$$\Omega_{\text{H}} = \frac{a c}{2r_{\text{H}}},$$

where a is the dimensionless black hole spin parameter, defined as $a \equiv cJ/(GM^2)$, with J the black hole angular momentum and M its mass. This relation shows explicitly that the generation of a toroidal magnetic field is directly linked to the black hole spin: in the absence of rotation ($a = 0$), $\Omega_{\text{H}} = 0$ and no toroidal field is produced.

The presence of a strong toroidal magnetic field implies a magnetic pressure

$$P_{\text{mag}} \propto \frac{B_{\phi}^2}{8\pi},$$

which accelerates plasma away from the black hole. While plasma motion across magnetic field lines is inhibited, motion along them is allowed; as a result, the poloidal magnetic field lines act as guiding channels that collimate the outflow into a narrow jet. The emerging flow is therefore initially magnetically dominated.

In this picture, the energy powering the jet is extracted directly from the rotational energy of the black hole. The Blandford–Znajek mechanism provides a quantitative description of this process, showing that the jet power depends on both the magnetic flux threading the horizon and the black hole spin. A commonly used scaling relation for the Blandford–Znajek power is

$$P_{\text{BZ}} \propto \Phi_{\text{BH}}^2 a^2.$$

This dependence reflects the fact that the magnetic flux determines how much magnetic field is available to couple to the black hole, while the spin controls how efficiently this field is twisted and converted into outgoing electromagnetic power.

The extracted energy is transported away from the black hole predominantly in the form of Poynting flux, that is, electromagnetic energy carried by large-scale electric and magnetic fields. As the jet propagates outward, this electromagnetic energy can be gradually converted into the kinetic energy of particles and, eventually, into the observed radiation.

Relativistic Jets and Blazars

When a relativistic jet is closely aligned with our line of sight, its radiation is strongly amplified by relativistic Doppler boosting. In this configuration, the jet emission dominates the entire observed spectrum, essentially hiding any thermal contribution from the accretion disk or the host galaxy. AGN observed under these conditions are known as **blazars**.

Blazars are divided into two main subclasses:

- (1) **BL Lacertae objects (BL Lacs)**: These systems are powered by *radiatively inefficient accretion flows*. Their emission is dominated by the jet, and their spectrum is **smooth and featureless at all frequencies**. Significant optical emission lines are absent, and the observed radiation is almost entirely non-thermal.
- (2) **Flat Spectrum Radio Quasars (FSRQs)**: These systems host *radiatively efficient accretion disks*. Their spectra show strong optical emission lines produced by gas near the central black hole: the **Broad Line Region (BLR)** and the more extended **Narrow Line Region (NLR)** (see Figure 4.1) The prominent emission lines observed in FSRQs are broad and originate in the broad-line region (BLR).

5.1 Relativistic Effects in Astrophysical Jets

Relativistic jets in Active Galactic Nuclei (AGN) consist of plasma moving at velocities very close to the speed of light. Because the emitting material travels relativistically and is often oriented near our line of sight, the observed radiation is amplified and blue-shifted. This is why blazars — jets observed nearly face-on — appear much brighter than misaligned radio galaxies, even when their intrinsic power is similar.

The strength of these relativistic effects is quantified by the Doppler factor,

$$\delta = [\gamma (1 - \beta \cos \theta)]^{-1},$$

where $\beta = v/c$ is the jet speed in units of the speed of light, $\gamma = (1 - \beta^2)^{-1/2}$ is the Lorentz factor, and θ is the angle between the jet direction and our line of sight (Figure 5.1).

The dependence of δ on the viewing angle is extremely steep. For jets seen at small angles, the Doppler factor can reach values of 10–30 or more. As θ increases, the boosting rapidly diminishes. Radio galaxies, observed at large angles ($\theta \gtrsim 30^\circ$), typically have $\delta \approx 1$ and therefore show little or no amplification. Blazars, viewed at $\theta < 10^\circ$, experience strong Doppler boosting.

The Doppler factor modifies both the observed flux and the observed time-scales. For an intrinsic power-law spectrum,

$$F'_{\nu'} = K \nu'^{-\alpha},$$

the observed flux density becomes

$$F_{\nu} = \delta^{3+\alpha} F'_{\nu'},$$

showing that even a moderately high δ can make the jet appear far more luminous than it intrinsically is. Relativistic boosting also affects the observed time-scales. Time intervals are shortened according to

$$\Delta t_{\text{obs}} = \frac{\Delta t_{\text{int}}}{\delta},$$

so that variability occurring over days in the jet rest frame may appear compressed into a few hours in the observer's frame. Moreover, the observed frequency is blue-shifted by relativistic Doppler boosting and is related to the intrinsic frequency by $\nu_{\text{obs}} = \delta \nu_{\text{int}}$,




Figure 5.1. Dependence of the Doppler factor on the angle to the line of sight. Different curves correspond to different Lorentz factors: from top to bottom, $\gamma = 15$, 10, 5, and 2. The expanded scale in the inset shows the angles for which $\delta = 1$ (Urry and Padovani 1995)

Summary

AGN jets move at relativistic speeds ($\beta \sim 0.99$), and their observed emission is strongly boosted when the jet is oriented toward us. The amplification depends critically on γ and θ : blazars, with jets nearly aligned to the observer, show very strong Doppler boosting, while radio galaxies, viewed at larger angles, show little or none. The Doppler factor enhances the flux, compresses the observed time-scales, and modifies the observed spectrum, making orientation a fundamental factor in determining how AGN jets appear.

5.1.1 Spectral Energy Distribution of BLAZAR

The spectral energy distribution (SED) of blazars shows **two broad peaks** (Figure 5.2): a low-energy peak and a high-energy peak.




Figure 5.2. Blazar sequence (Fossati et al 1998)

For BL Lac objects, the position of the first peak (i.e. of the **synchrotron peak**) is used to define three subclasses:

- ◊ **Low-synchrotron-peaked BL Lacs (LBL)**: peak in the infrared/optical;
- ◊ **Intermediate-synchrotron-peaked BL Lacs (IBL)**: peak in the optical/UV;
- ◊ **High-synchrotron-peaked BL Lacs (HBL)**: peak in the UV/X-ray.

In this broader context, **FSRQs** represent the most luminous end of the blazar family, with synchrotron peaks typically in the infrared. This so-called **blazar sequence**, which links luminosity to the peak frequency of the SED, has been **strongly debated**, mainly because it may be affected by substantial **observational biases**, especially those related to the selection of the samples used to construct it. Despite these limitations, it remains a *useful phenomenological framework* for organising and classifying blazars.

5.1.2 The Two Peaks of the Blazar SED

The characteristic double-peaked shape of the blazar spectral energy distribution (SED) is produced by two distinct physical mechanisms operating in the relativistic jet. A standard reference for radiative processes in astrophysics is *Radiative Processes in Astrophysics* by Rybicki and Lightman.

(1) **FIRST PEAK: Synchrotron Radiation**

The first peak of the blazar SED is produced by **synchrotron radiation**, emitted by relativistic electrons spiralling around magnetic field lines. A commonly used description of the electron population is a power-law distribution in Lorentz factor (γ_e):

$$N(\gamma_e) = K \gamma_e^{-p}, \quad \gamma_{\min} < \gamma_e < \gamma_{\max},$$

where $N(\gamma_e)$ is the number of particle with γ_e for unit of volume, K is a normalization constant and p is the electron energy index. The index p is related to the observed synchrotron spectral index α by:

$$p = 1 + 2\alpha.$$

For such a distribution, the synchrotron emissivity at frequency ν is:

$$\epsilon_{\nu}^{\text{syn}} \propto K m_e^{-(\alpha+1)} B^{\alpha+1} \nu^{-\alpha}$$

expressed in units of $\text{erg cm}^{-3} \text{ s}^{-1} \text{ sr}^{-1}$.

(2) **SECOND PEAK: Inverse Compton Radiation**

When an electron is not at rest, but instead has an energy much larger than that of an incoming photon, their interaction can lead to a transfer of energy **from the electron to the photon**. This process is known as inverse Compton scattering. For a seed photon of frequency ν_{seed} scattered by an electron with Lorentz factor γ_e , the characteristic frequency of the up-scattered photon in the Thomson regime is

$$\nu_{\text{IC}} \approx \gamma_e^2 \nu_{\text{seed}}.$$

Consider a power-law distribution of relativistic electrons:

$$N(\gamma_e) = K \gamma_e^{-p}, \quad \gamma_{\min} < \gamma_e < \gamma_{\max}.$$

The emissivity at the characteristic inverse-Compton frequency ν_c can then be written as

$$\epsilon_c(\nu_c) \propto K \nu_c^{-\alpha} m_e^{-2} \int \frac{U_r(\nu) \nu^{\alpha}}{\nu} d\nu,$$

where $U_r(\nu)$ is the radiation energy density of the seed photon field and $p = 1 + 2\alpha$, as in the case of synchrotron emission.

The total radiation energy density is given by

$$U_r = \int n(\epsilon) \epsilon d\epsilon,$$

where $n(\epsilon)$ is the photon number density. Relevant sources of seed photons include:

- ◇ synchrotron photons produced within the jet (SSC: synchrotron self-Compton);
- ◇ external photon fields from the accretion flow, disk, BLR, NLR, or dusty torus (EC: external Compton);
- ◇ cosmic microwave background (CMB) photons, which are important for large-scale jets.

Because BL Lacs and FSRQs have fundamentally different circumnuclear environments, the dominant inverse Compton process also differs between the two subclasses. **FSRQs** possess a rich external photon field supplied by the accretion disk, the BLR, the dusty torus, and the NLR. In these sources, *external Compton (EC)* scattering is typically the favoured mechanism and provides the best fit to the observed γ -ray spectra (Figure 5.3 - *Left Panel*). In **BL Lac objects**, where the environment is photon-poor and lacks strong external radiation fields, the observed high-energy emission is generally well reproduced by *synchrotron self-Compton (SSC)* models (Figure 5.3 - *Right Panel*).




Figure 5.3. Examples of spectral energy distributions (SEDs) of a FSRQ (right, 3C 454.3) and a BL Lac object (left, Mkn 501) from multi-epoch observations. In the SED of 3C 454.3, in addition to the jet emission, we can identify an accretion disc component, the X-ray corona contribution and the IR emission from the torus (dashed black lines). In Mkn 501, instead, only the contributions from the jet and host galaxy are evident (Prandini and Ghisellini 2022.)

5.2 Hadronic Models and Neutrinos

In hadronic models of AGN jets, relativistic protons (and possibly higher- Z nuclei) are present together with electrons. As in the leptonic case, the protons are often assumed to follow a non-thermal energy distribution described by a power law between a minimum and a maximum Lorentz factor:

$$N_p(\gamma_p) = N_{0,p} \gamma_p^{-p}, \quad \gamma_{p,\min} \leq \gamma_p \leq \gamma_{p,\max}.$$

Relativistic protons follow the same synchrotron and Inverse Compton radiation laws as electrons. The key difference is that the particle mass entering the synchrotron formulas is the proton mass rather than the electron mass. Since the proton mass is approximately 10^3 times larger than the electron mass, proton cooling is much less efficient.

- ◇ *Proton synchrotron emission versus electron synchrotron emission*

For the same magnetic field strength and particle energy to make proton synchrotron emission relevant in AGN jets, one must require either very strong magnetic fields, very high proton densities, or both.

- ◇ *Proton Inverse Compton emission versus electron Inverse Compton emission*

Even if a proton reaches very large Lorentz factors, it remains an extremely inefficient IC radiator. A proton must satisfy

$$\gamma_p \sim (m_p/m_e) \gamma_e \approx 1000 \gamma_e$$

to produce photons of the same energy as an electron, but its IC power is still suppressed by the factor $(m_e/m_p)^2$.

Despite their inefficiency, relativistic protons can interact very effectively with low-energy photons through photo-hadronic processes. The following reactions occur:

$$p + \gamma \rightarrow p + \pi^0, \quad p + \gamma \rightarrow n + \pi^+, \quad p + \gamma \rightarrow p + \pi^+ + \pi^-.$$

These pions then decay into secondary particles. Neutral pions produce gamma rays,

$$\pi^0 \rightarrow 2\gamma,$$

while charged pions produce muons and neutrinos,

$$\pi^+ \rightarrow \mu^+ + \nu_\mu, \quad \pi^- \rightarrow \mu^- + \bar{\nu}_\mu.$$

The muons subsequently decay,

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu, \quad \mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu,$$

injecting electrons and positrons that radiate through synchrotron or inverse-Compton emission and may trigger electromagnetic cascades within the source.

A key consequence is that photo-hadronic interactions naturally produce **high-energy neutrinos**. These neutrinos escape the source without further interactions, providing a clear and unique signature of the presence of relativistic protons in AGN jets. Figure 5.4 displays the SED of the BL Lac Mkn 421, with the high-energy peak reproduced by a hadronic model. Remarkably, the model also allows us to estimate the luminosity carried by the associated neutrinos.

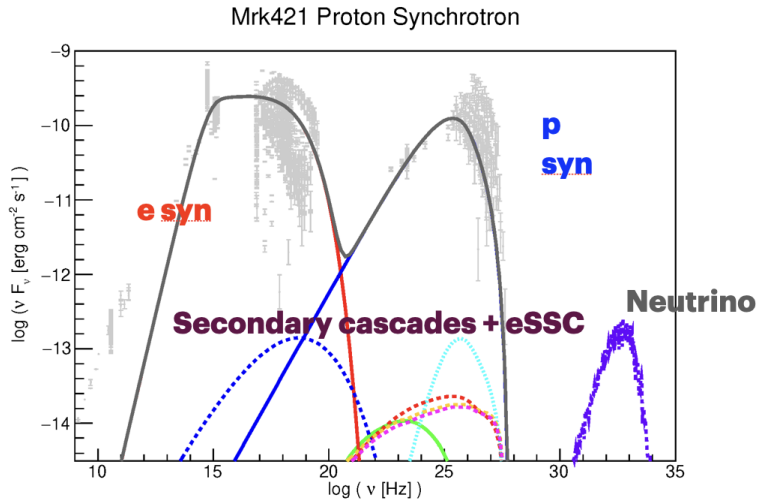


Figure 5.4. SED model of MKN421. The high peak is produced by proton synchrotron. The contribution of the cascaded are also shown as well as the neutrinos emission

5.B Pros and cons of hadronic models

Cons: in hadronic scenarios, the low radiative efficiency of relativistic protons requires a large proton energy density and/or very strong magnetic fields to produce high-energy radiation observed in jetted AGN. As a result, the jet power can be extremely (too much) large.

Pros: hadronic models provide a natural explanation for neutrino production.

Appendices

APPENDIX A

Even file classification

A.1 - Fermi-LAT *Pass 8 (P8R3)* event classes with their corresponding `evclass` identifiers. Each class represents a different balance between signal purity, photon statistics, and background contamination.

Event Class	Hierarchy	evclass	Description / Main Use
P8R3_TRANSIENT020	Standard	16	Transient event class with background rate $\sim 2 \times A10$ IGRB, being A10 IGRB the isotropic gamma-ray background (IGRB) intensity defined in the Ackermann et al. (2010) model. Used for short-duration phenomena (GRBs, flares) where maximizing photon counts is essential.
P8R3_TRANSIENT010	Standard	64	Transient event class with background rate comparable to A10 IGRB. Suitable for bright, short events.
P8R3_SOURCE	Standard	128	Recommended for most analyses. Balanced between purity and photon statistics; ideal for point-like or moderately extended sources.
P8R3_CLEAN	Standard	256	Identical to SOURCE below 3 GeV, but with $1.3\text{--}2 \times$ lower background above 3 GeV. Suitable for high-latitude and hard-spectrum sources.
P8R3_ULTRACLEAN	Standard	512	Very low background and excellent PSF. Similar background rate to ULTRACLEANVETO; suited for diffuse studies.
P8R3_ULTRACLEANVETO	Standard	1024	The cleanest Pass 8 event class. Background $\sim 15\text{--}20\%$ lower than SOURCE below 10 GeV and $\sim 50\%$ lower at 200 GeV. Best for diffuse or low-background studies.
P8R3_SOURCEVETO	Standard	2048	Similar to SOURCE below 10 GeV; at higher energies the background matches ULTRACLEANVETO but with $\sim 15\%$ more acceptance.
P8R3_TRANSIENT020E	Extended	8	Extended version of TRANSIENT020, with less restrictive fiducial cuts on track length through the calorimeter.
P8R3_TRANSIENT010E	Extended	32	Extended version of TRANSIENT010, optimized for low-energy photons and large incidence angles.
P8R3_TRANSIENT015S	No-ACD	65536	Excludes Anti-Coincidence Detector (ACD) variables. Used when ACD pile-up occurs, e.g. during solar flares.

A.2 - Fermi-LAT event types used in data analysis.

Event Type	evtype	Description
FRONT	1	Photons converting in the front section of the tracker, where the tungsten layers are very thin. These events have better angular resolution (narrower PSF) because the electron–positron pair undergoes little multiple scattering, but the probability of conversion is lower.
BACK	2	Photons converting in the back section of the tracker, which contains thicker tungsten layers. These events have poorer angular resolution due to stronger multiple scattering and bremsstrahlung, but a higher probability of conversion .
FRONT + BACK	3	Combination of both FRONT and BACK events. This is the default choice in most LAT analyses, since it maximizes photon statistics while accounting for both angular resolution and detection efficiency.

APPENDIX B

Python code for the visualization of a 3D cube file

Python Code: Log-scale Visualization of a FITS Cube Slice

```
import numpy as np

import matplotlib.pyplot as plt

from astropy.io import fits

# Carica il file

with fits.open('/Users/paola/Desktop/LABx/1H0717+714/models/gll_iem_v07.fits') as
    hdul:

    cube = hdul[0].data # shape (28, 1441, 2880)

    energies = hdul['ENERGIES'].data['energy']

# Seleziona uno slice (ad esempio il bin 10)

index = 10

image = cube[index]

energy = energies[index]

# Applica log10 in sicurezza (aggiungendo un offset minimo)

log_image = np.log10(image + 1e-10) # evitare log(0)

# Plot in scala log

plt.figure(figsize=(10, 6))

plt.imshow(log_image, origin='lower', cmap='inferno')

plt.title(f"Diffuse Emission at {energy:.1f} MeV (log scale)")

plt.xlabel("Galactic Longitude (pixel)")

plt.ylabel("Galactic Latitude (pixel)")

plt.colorbar(label='log10(Flux) [ph/cm²/s/sr/MeV]')

plt.savefig("diffuse_emission_logscale.png", dpi=300, bbox_inches='tight') #
    formato PNG ad alta risoluzione

plt.show()
```

Region of Interest Model File

Example of a source model definition file used for likelihood analysis. Each parameter in the file includes a flag (free) that determines whether it is allowed to vary during the likelihood fit: **free** = 1 indicates that the parameter is free and will be optimized, whereas **free** = 0 indicates that the parameter remains fixed throughout the analysis.

```
<?xml version="1.0" standalone="no"?>
<source_library title="source library">
  <source name="4FGL J0557.2+7705" type="PointSource">
    <spectrum type="PowerLaw">
      <parameter free="0" max="1000" min="1e-05" name="Prefactor" scale="1e-14" value="0.1143700987" />
      <parameter free="0" max="5" min="0" name="Index" scale="-1" value="2.154637575" />
      <parameter free="0" max="2717.082275" min="2717.082275" name="Scale" scale="1" value="2717.082275" />
    </spectrum>
    <spatialModel type="SkyDirFunction">
      <parameter free="0" max="360" min="-360" name="RA" scale="1" value="89.3129" />
      <parameter free="0" max="90" min="-90" name="DEC" scale="1" value="77.0854" />
    </spatialModel>
  </source>
  <source name="4FGL J0559.5+7026" type="PointSource">
    <spectrum type="LogParabola">
      <parameter free="0" max="1000" min="1e-05" name="norm" scale="1e-12" value="0.003304105393" />
      <parameter free="0" max="5" min="-5" name="alpha" scale="1" value="2.531776905" />
      <parameter free="0" max="2" min="-2" name="beta" scale="1" value="0.2594358921" />
      <parameter free="0" max="666.7203369" min="666.7203369" name="Eb" scale="1" value="666.7203369" />
    </spectrum>
    <spatialModel type="SkyDirFunction">
      <parameter free="0" max="360" min="-360" name="RA" scale="1" value="89.8827" />
      <parameter free="0" max="90" min="-90" name="DEC" scale="1" value="70.4434" />
    </spatialModel>
  </source>
  <source name="galdiff" type="DiffuseSource">
    <spectrum apply_edisp="false" type="PowerLaw">
      <parameter free="0" max="10" min="0.1" name="Prefactor" scale="1" value="0.8482565568" />
      <parameter free="0" max="1" min="-1" name="Index" scale="-1" value="0.07383095621" />
      <parameter free="0" max="1000" min="1000" name="Scale" scale="1" value="1000" />
    </spectrum>
    <spatialModel file="/Users/paola/Desktop/LABx/background/gll_iem_v07.fits" type="MapCubeFunction">
      <parameter free="0" max="10" min="0" name="Normalization" scale="1" value="1" />
    </spatialModel>
  </source>
  <source name="isodiff" type="DiffuseSource">
    <spectrum apply_edisp="false" file="/Users/paola/Desktop/LABx/background/iso_P8R3_SOURCE_V3_v1.txt" type="FileFunction">
      <parameter free="0" max="1000" min="0.001" name="Normalization" scale="1" value="1.208439851" />
    </spectrum>
    <spatialModel type="ConstantValue">
      <parameter free="0" max="10" min="0" name="Value" scale="1" value="1" />
    </spatialModel>
  </source>
</source_library>
```


APPENDIX D

Data analysis steps

- (1) **Event selection.** The event files listed in the `*.txt` input list are combined to produce a first-level filtered event file, `ft1_00.fits`. This file contains only the photons that fall within the selected Region of Interest (ROI), time range, energy range, zenith-angle cut, and event class/type.
- (2) **Application of mission-quality cuts.** Using the information stored in the spacecraft file, the software removes all events recorded during intervals of poor data quality, invalid instrument configuration, or other non-scientific conditions. The resulting file, `ft1_filtered_00.fits`, contains only events acquired during valid Good Time Intervals (GTIs).
- (3) **Construction of the livetime cube (`ltcube_00.fits`).** The livetime cube summarizes how much observing time the LAT has accumulated in each direction of the sky. The sky is divided into small pixels in right ascension (RA) and declination (DEC), with each pixel corresponding to one row of the cube. Because the LAT continuously rocks and reorients, each sky pixel is observed under many different incidence angles. For this reason, the livetime for each pixel is stored in 40 separate $\cos \theta$ bins, where θ is the incidence angle between the incoming photon direction and the LAT boresight. Summing these 40 bins gives the total livetime for a pixel, which is used to compute the exposure in the subsequent likelihood analysis.
- (4) **Construction of the Count Cube (`ccube.fits`)** The count cube is a three-dimensional histogram of the LAT photon data. The two spatial axes define the Region of Interest (ROI), while the third axis consists of the energy bins specified by the user.

Each element of the cube,

$$\text{Counts}(i, j, k),$$

represents the number of photons detected in sky pixel (i, j) and energy bin k . The `ccube` contains only observed counts. It is possible to generate a sequential visualization of the individual slices (RA, DEC) at different energies by means of DS9.

- (5) **Construction of the bexmap (`bexmap_00.fits`)** The effective area of the LAT depends on the photon energy and on the incidence angle with respect to the instrument boresight, $A_{\text{eff}}(E, \theta)$. It is a property of the detector and therefore **does not depend on time**. What *does* change with time is the orientation of the spacecraft. As the LAT scans the sky, each sky direction \hat{v} is observed for different amounts of time and at different incidence angles. This information is stored in the **livetime cube**, which contains the observing time accumulated in each sky pixel as a function of incidence angle, $T(\hat{v}, \theta)$. To compute the exposure, these two quantities are combined:

$$\text{Exposure}(E, \hat{v}) = \sum_i A_{\text{eff}}(E, \theta_i) T(\hat{v}, \theta_i) \Delta(\cos \theta).$$

The exposure therefore represents how effectively the LAT has observed each point in the sky at each energy. It is measured in $\text{cm}^2 \text{s}$.

Note: The accumulated observing time $T(\hat{v}, \theta)$ are evaluated at representative incidence angles θ_i , while the angular integration is performed using bins of equal width in $\cos \theta$ (see point 3). A simple way to remember the meaning of these quantities is:

1. **Effective area:** the instrument's "ability to see".
2. **Livetime:** the "time spent looking".
3. **Exposure:** "how much was effectively seen" (ability \times time).

It is possible to generate a sequential visualization of the individual slices (RA, DEC) at different energies by means of DS9

(6) **Construction of the bexpmap (bexpmap_roi.fits)**

The file `bexpmap_roi.fits` contains the binned exposure map restricted to the region of interest. It is derived by cropping the full binned exposure map and does not involve any recomputation.

Square Exposure Map in gtexpcube2

`gtexpcube2` does not generate a circular ROI. Instead, it produces a **square map** centered on the target position, defined by:

- ◊ **number of pixels:** $\text{nxpix} \times \text{nypix}$
- ◊ **pixel size:** `binsz` (degrees per pixel)

The total angular size of the map is:

$$\text{Size} = \text{nxpix} \times \text{binsz}.$$

To ensure that a circular ROI of radius R is fully contained within the square map, the following condition must be satisfied:

$$\text{nxpix}, \text{nypix} \geq \frac{2R}{\text{binsz}}.$$

- (7) **Construction of the *source maps* (src_00.fits)** In this step, the XML model of the Region of Interest (ROI) is used to compute the expected number of photons from each source in the sky. For every source listed in the XML file, the software calculates how many photons should be detected in each sky pixel and in each energy bin. This calculation uses several ingredients: the source position and spectrum (from the XML model), the LAT exposure, the instrument response functions (IRFs), the point-spread function (PSF), and the selected event type. The resulting file, `srcmap_00.fits`, is defined on the same spatial and energy grid as the counts cube (`ccube.fits`). It represents the model prediction of the expected counts.

In the table above, the different analysis steps are summarized together with the corresponding input and output files for each step, along with a brief description of the task performed.

Summary of a FERMI-LAT analysis step by step				
Step	Input	Process	Output	Short Description
1	PH files (list)	Event selection	ph_00.fits	Select events within ROI, time range, energy, zenith cut, event type.
2	ph_00.fits + SC	Mission-quality cuts	ph_filtered_00.fits	Remove data taken during bad-quality intervals; keep only valid GTIs.
3	ph_filtered_00.fits + SC	Livetime cube	ltcube_00.fits	Compute observing time for each sky pixel and incidence-angle bin.
4	ph_filtered_00.fits	Count cube	ccube.fits	Bin photons in RA, Dec, and energy to form a 3D histogram.
5	ltcube_00.fits + IRFs	Exposure map	bexpmap_roi.fits	Compute sky exposure combining effective area and livetime.
6	bexpmap_roi.fits + ccube.fits + XML	Source map	srcmap_00.fits	Predict expected counts per pixel/energy bin for each model source.
7	srcmap_00.fits + ccube.fits + XML	Likelihood analysis	Best-fit parameters	Fit model to data and compute statistical significance (TS).

Object	Depends on events?	Uses IRFs?	What it represents
Counts cube (ccube)	yes	no	reconstructed events
Livetime cube	no	no	observing geometry and time
Exposure map	no	yes (A_{eff})	instrument sensitivity
Source map (srcmap)	no	yes ($A_{\text{eff}} + \text{PSF}$)	model-predicted counts

Connection between angular momentum loss and energy dissipation in a keplerian disk

The specific angular momentum in a Keplerian disk is

$$\ell = \sqrt{GM r}.$$

If a ring of gas loses an amount $\Delta\ell < 0$, it cannot remain at the same orbital radius. The new radius must satisfy

$$r_{\text{new}} < r_{\text{old}}.$$

The orbital velocity in a Keplerian potential is

$$v_{\phi} = \sqrt{\frac{GM}{r}}.$$

Thus, when r decreases, v_{ϕ} increases:

$$r \downarrow \Rightarrow v_{\phi} \uparrow.$$

Therefore, losing angular momentum makes the inner gas rotate faster. However, this increase in velocity is not sufficient to balance the change in orbital energy.

Angular momentum and orbital energy scale differently with radius:

$$\ell \propto r^{1/2}, \quad E = -\frac{GM}{2r} \propto -r^{-1}.$$

Because of these scaling relations:

- ◇ angular momentum decreases slowly as r decreases,
- ◇ orbital energy becomes much more negative very rapidly.

When a fluid element moves inward, its orbital speed increases modestly, but its binding energy becomes significantly larger in magnitude. The net effect is a substantial loss of orbital energy.

The outer ring does receive the angular momentum, but only a very small amount of energy, because it resides at a large radius. Thus, most of the orbital energy lost by the inner gas cannot be transferred outward.

This “excess” energy must be dissipated locally by turbulent and magnetic stresses, which convert it into heat. The heated gas then cools by emitting photons, producing the characteristic radiation of accretion disks.