FACULTÉ DES SCIENCES

UNIVERSITÉ LIBRE DE BRUXELLES

Searches for neutrino emission from blazar flares with IceCube

Thesis presented by Christoph RAAB

in fulfilment of the requirements of the PhD Degree in Sciences
("Docteur en Sciences")
Année académique 2020-2021

Supervisor : Professor Juan Antonio AGUILAR SÁNCHEZ
Inter-University Institute For High Energies

Thesis jury :

Laurent FAVART (Université libre de Bruxelles, Chair)
Simona TOSCANO (Université libre de Bruxelles, Secretary)
Juan Antonio AGUILAR SÁNCHEZ (Université libre de Bruxelles)
Markus AHLERS (Københavns Universitet)
Krijn DE VRIES (Vrije Universiteit Brussel)

LA LIBERTÉ DE CHERCHER
Fonds David et Alice Van Buuren
Fondation Jaumotte – Demoulin

SEARCHES FOR NEUTRINO EMISSION FROM BLAZAR FLARES WITH ICECUBE

PHD THESIS

CHRISTOPH RAAB

IIHE BRUSSELS 2021

Abstract

Cosmic rays reach Earth from beyond the Milky Way and with energies up to 10^{20} eV. The responsible accelerators have to date not been discovered. However, multi-messenger astronomy can shed light on the question, based on the principle that protons and nuclei accelerated in dense and energetic environments would also produce gamma rays and neutrinos. Such environments may be found in **blazars**, which are therefore cosmic ray accelerator candidates. Their gamma-ray emission has been observed to increase, sometimes by orders of magnitude, during **flares** as observed in light curves taken by the Large Angle Telescope on the Fermi satellite. When the latter was launched in 2008, the IceCube Neutrino Observatory had also started taking data, detecting the Cherenkov light from high-energy neutrino interactions in the glacier ice under the geographic South Pole. These two experiments have enabled multi-messenger searches for neutrinos in time correlation with the gamma-ray emission from blazars. This work builds on this principle and extends it by **stacking** the signal from multiple blazar flares. Thus, their individually undetectable neutrino emission could still be discoverable. One first analysis focused on the blazar TXS 0506+056, whose flare in 2017 coincided with arrival of the neutrino IceCube 170922-A. Extending into a lower energy range than the alert, the search found no additional excess neutrinos associated with the flare. A second analysis used 179 bright and variable blazars. They were divided in two specific blazar classes and weighted relatively to each other, with two weighting schemes motivated physically using the observed gamma-ray luminosity and a third, generic weighting to cover unconsidered scenarios. No significant neutrino excess was found in the unblinded likelihood fits for any of the source catalogues and weighting schemes. Their combined trial-corrected p-value was $p = (79.1 \pm 0.3)\%$. The limits derived from this analysis are also discussed and its relation with other searches considered. Since that was the first **blazar flare stacking**, this work also proposes further improvements to the analysis which will help advance the search for cosmic ray accelerators.

Résumé

Les rayons cosmiques proviennent d'au-delà de la Voie lactée et atteignent la Terre avec des énergies pouvant aller jusqu'à 10^{20} eV. Les objets qui accélèrent ces rayons cosmiques n'ont toujours pas été découverts. Toutefois, l'astronomie multimessager peut apporter un élément de réponse à cette question, en supposant que les protons et les noyaux accélérés dans des environnements denses et énergétiques pourraient également produire des rayons gamma et des neutrinos. Les « blazars » sont de possibles candidats pour les accélérateurs de rayons cosmiques. Une augmentation de leurs émissions de rayons gamma, parfois de plusieurs ordres de grandeur, a été observée lors de phénomènes qu'on appelle « éruption », comme le montrent les courbes de lumière prises par le télescope spatial Fermi-LAT. Lorsque ce dernier a été lancé en 2008, l'observatoire de neutrinos IceCube avait également commencé à prendre des données, détectant la lumière Tcherenkov provenant d'interactions de neutrinos à haute énergie dans la glace qui se trouve sous le Pôle Sud géographique. Ces deux expériences ont permis de mener à bien des recherches multi-messagers de neutrinos en corrélation temporelle avec l'émission de rayons gamma des blazars. Ce principe est le point de départ de cette thèse, qui va plus loin en employant la méthode du « stacking », qui consiste à combiner les signaux provenant de plusieurs éruptions de blazars. Ainsi, leurs émissions individuelles de neutrinos, habituellement indétectables, pourraient être découvertes après combinaison. Une première analyse s'est concentrée sur le blazar TXS 0506+056, dont l'éruption en 2017 a coïncidée avec l'arrivée de l'évènement IceCube 170922-A. En considérant une gamme d'énergie inférieure à celle de l'alerte 170922-A, pas d'autres neutrino excédentaire n'a été associé à l'éruption. Une deuxième analyse est basée sur 179 blazars lumineux et variables. Ces blazars ont été répartis en deux classes spécifiques, et chacun d'entre eux a reçu un poids relatif. Trois schémas de pondération ont été considérés : les deux premiers étant motivés par des observations, le troisième étant plus générique. Aucun excès significatif de neutrinos n'a été observé après avoir effectué des ajustements par maximum de vraisemblance sur les données non masquées, pour les différents catalogues de sources et schémas de pondération. Leur valeur-p combinée est de $p = (79.1 \pm 0.3)\%$. Les limites dérivées de cette analyse sont discutées ainsi que leur rapport avec les résultats d'autres recherches. Puisqu'il s'agit du premier stacking d'éruptions de blazars, nous suggérons également des améliorations à apporter à l'analyse afin de permettre la poursuivre de la recherche d'accélérateurs de rayons cosmiques.

Samenvatting

Kosmische straling afkomstig van buiten de Melkweg bereikt de Aarde met energieën tot wel 10^{20} eV. De astrofysische bronnen waarin deze deeltjes worden versneld zijn tot op heden nog niet ontdekt. De multi-boodschapperastronomie kan een nieuw licht werpen op de oorsprong van kosmische straling, aangezien protonen en atoomkernen die worden versneld in een dichte en energetische omgeving ook gammastralen en neutrino's produceren. „Blazars" zijn mogelijke kandidaat-versnellers. Observaties van blazars, gemaakt met de ruimtetelescoop Fermi-LAT , tonen aan dat hun gammastraling tijdens zogenaamde „flakkers" toeneemt. Rond de tijd dat deze werd gelanceerd, begon het IceCube Neutrino Observatorium ook gegevens te verzamelen. Deze laatste detecteert hoog-energetische neutrino's aan de hand van het Cherenkovlicht dat geproduceerd wordt tijdens hun interacties met de ijskap bij de geografische zuidpool. Deze twee experimenten hebben het mogelijk gemaakt om een multi-booschapperzoektocht te verrichten naar neutrino's van blazars die een tijdscorrelatie hebben met diens flakkers van gammastraling. Dit is het uitgangspunt van dit proefschrift, waarbij er ook een zogenaamde „stapelmethode" wordt toegepast. Op deze manier kan de neutrino-emissie van indivuele blazarflakkers, die afzonderlijk te zwak is om te detecteren, gecombineerd worden en mogelijks toch worden ontdekt. Een eerste analyse legt de focus op de blazar TXS 0506+056, waarvan een flakker in 2017 samenviel met de aankomst van het neutrino IceCube 170922-A. In een relatief lager energiebereik wordt er geen surplus aan neutrino's gevonden gecorreleerd met de flakker. In een tweede analyse maken we gebruik van de stapelmethode om neutrino's te zoeken afkomstig van 179 heldere en variabale blazars. Deze worden onderverdeeld in twee specifieke klassen en krijgen elks een zeker gewicht in de stapelanalyse. Hiervoor worden twee wegingsschema's gebruikt die gemotiveerd zijn door de geobserveerde gammastraling, alsook een derde generieke weging. Ook hier wordt er geen significant neutrinosignaal geobserveerd. De gecombineerde p-waarde is $p = (79.1 \pm 0.3)\%$. Hieruit worden limieten afgeleid, en worden de verbanden met andere zoekacties besproken. Aangezien dit werk de eerste analyse omvat naar neutrino's afkomstig van blazarflakkers gebruik makende van een stapelmethode, worden er in dit werk ook verdere verbeteringen van de analyse voorgesteld. Deze zullen als een startpunt dienen voor toekomstige zoektochten naar de nog onbekende bronnen van kosmische straling.

Contents

"The Universe is very big [. . .]"
Jocelyn Bell Burnell [1]

1
Cosmic Rays & Multimessenger Astronomy

This chapter describes the acceleration, propagation and observation of cosmic rays. General conditions for cosmic accelerators are also discussed, while the description of the specific sources studied in this thesis is deferred to chapter 5. This chapter will also explain the interactions of cosmic rays leading to neutral messenger particles, whose detection principle is described in chapter 2 and chapter 4. The chapter concludes with a description of multimessenger astronomy, one of the currently leading paradigms of cosmic ray studies.

1.1. Cosmic ray observation

1.1.1. Spectrum and composition

Cosmic rays are charged particles that reach Earth from space, a definition which excludes secondary particles produced in Earth's atmosphere, as well as all neutrons, neutrinos, and photons. Cosmic ray particles of the lowest detected energies have solar origin and consist of the thermal solar wind made up of electrons, protons, alpha particles and heavier ions up to E/nucleon $=10^4$ eV, as well as a non-thermal power-law tail mostly due to rarer acceleration events up to E/nucleon $=10^8$ eV [2, p. 15]. Beyond this energy, nuclei from outside the

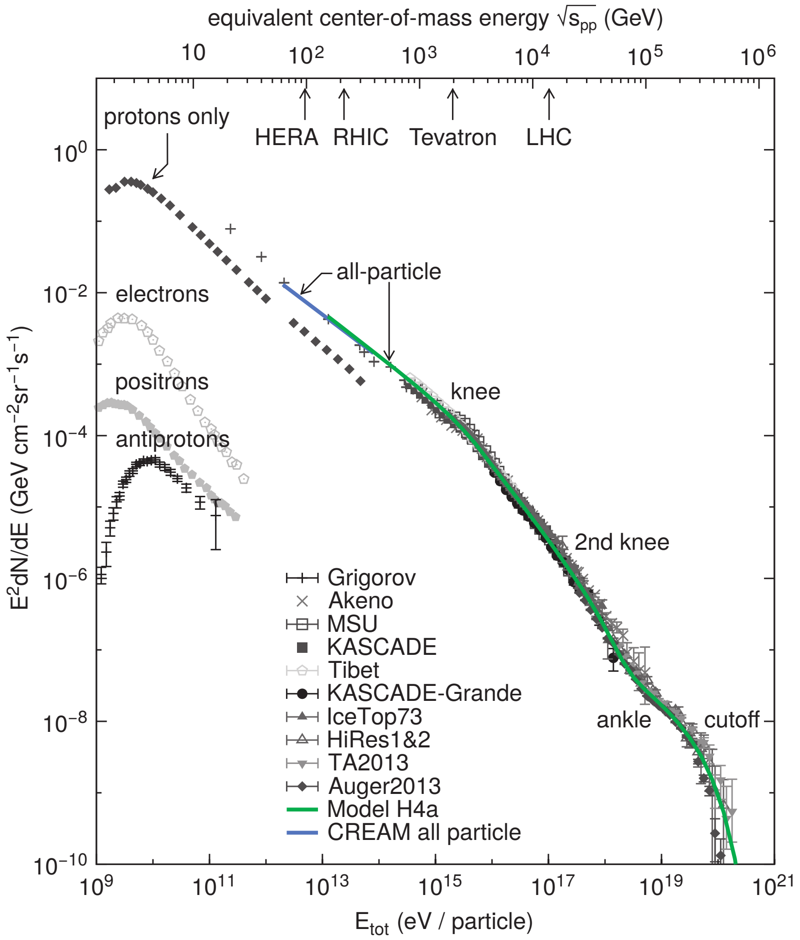

Figure 1.1.: The cosmic ray spectrum versus per-particle energy spanning 11 orders of magnitude, adapted from [2] with [3] and [4]. The flux is scaled by E^2 in order to highlight the anatomical features of the spectrum. Lines show models, and points individual measurements. Towards higher energies, the apparent discrepancies in normalization result from the difficulties calibrating the absolute energy scale of an air shower experiment. These can be understood by considering that no laboratory or accelerator experiment probes the hadronic interactions at such high centre-of-mass energies as are the origin of the observed air showers.

solar system dominate the cosmic-ray spectrum. Ignoring the 1% of electrons, whose fraction diminishes with energy due to their steeper spectrum until they almost disappear above $\approx 10^{12}$ eV [5], future discussions of cosmic rays focus on protons and heavier nuclei.

Being charged particles, not all cosmic rays approaching the solar system manage to diffuse upstream against the solar wind in an effect called **solar modulation**, which is relevant up to $\approx 10^{10}$ eV [2], as can be seen in fig. 1.1. At this point, the flux of cosmic rays arriving at Earth is roughly 1000 particles per square metre per second*. However, cosmic rays have been detected with energies up to $\approx 10^{20}$ eV [6, 7], i.e. 10 orders of magnitude higher. Over the same range, the flux decreases to 1 particle per square kilometre per century [8, sec. 10.4.1.1], i.e. a reduction of roughly 18 orders of magnitude. This large dynamic range, in both energy and flux, indicates that the sources of cosmic rays and physics governing their propagation are very diverse. The all-particle spectrum, shown in fig. 1.1, can be approximated by a repeatedly broken power law, with an anatomy as follows [9, 8, 2]:

- At 2×10^{10} eV, solar modulation is overcome and the spectrum can be approximated to [9]:

$$\frac{\mathrm{d}N}{\mathrm{d}E\,\mathrm{d}A\,\mathrm{d}t\,\mathrm{d}\Omega} \approx 1.8 \times 10^4 \left(\frac{E}{\mathrm{GeV}}\right)^{-2.7} \frac{\text{nucleons}}{\mathrm{m}^2\ \mathrm{s}\ \mathrm{sr}\ \mathrm{GeV}} \tag{1.1}$$

  where E is the energy per nucleon and $\mathrm{d}N$ the number of nucleons with energies $[E, E + \mathrm{d}E]$ passing through an area $\mathrm{d}A$ into a solid angle $\mathrm{d}\Omega$. For this part of the spectrum, direct measurements of cosmic rays estimate that cosmic rays are mostly made up of protons at 74%, another 18% nucleons are bound in Helium nuclei, and the rest are heavier nuclei [9]. These fractions remain more or less constant until energies of 100 TeV [10, sec. 30.1], with observations up to 6×10^{12} eV showing a steadily increasing Helium abundance [11]. The fact that the cosmic-ray spectrum is featureless over this energy range suggests that a single class of sources is responsible for the acceleration of the bulk of cosmic rays.

- Around 5×10^{15} eV, the spectrum begins to steepen (to $E^{-3.1}$) in a first break called the "knee" [8]. At these energies chemical composition studies

*numerical integral of fig. 1.1

become difficult as direct detection of cosmic-rays is not possible and measurements rely on the physics of air showers (see section 1.1.2). There are, however, indications that above the "knee" heavier elements become more relevant [12].

- After 5×10^{17} eV, a "second knee" brings it gradually to $\propto E^{-3.3}$ [8]. If the knee and second knee can be explained as a series of cutoff energies at $Z \cdot E_{\max}$ (the so-called Peters cycle) [2, p. 24], then the second knee would mark the end of a population of galactic cosmic rays.
- Around 3×10^{18} eV the spectrum flattens again, back to $\propto E^{-2.6}$, resembling an "ankle" in its anatomy [2]. Above this energy, the flux is ≈1 particle/km^2/year [13]. At these energies cosmic rays are believed to have an extragalactic origin since their gyroradius is too large for easy containment in the Galaxy (see section 1.1.3). The "ankle" itself is sometimes believed to be the transition towards extra-galactic cosmic rays [14].
- With 10^{20} eV the spectrum begins to cut out. This feature could be due to protons losing energy in interactions with the cosmic microwave background (CMB) (see section 1.3.2) or it can simply reflect the exhaustion of the sources of ultra-high-energy cosmic rays, depending on their composition. The highest energy cosmic rays observed lie within a factor 3 of here [6, 7], although more precise numbers can not be given due to the difficulty in calibrating the energy scale of air shower detection methods (see section 1.1.2).

1.1.2. Overview of current detection methods

Since cosmic rays at the lowest energies are absorbed by Earth's atmosphere or deflected by its magnetic field, they need to be detected with **satellite or balloon experiments**. Satellite-borne detectors have typically high duty cycles near 100%, but are limited in weight by the price of rocket fuel and feasible size of a spacecraft. To make the best of the resulting small acceptance (e.g. m^2 sr), they are compact and complex composites of highly sensitive detectors from the world of accelerator-based experiments, such as transition radiation detectors, time-of-flight detectors, calorimeters, silicon trackers and Ring-imaging Cherenkov (Rich) detectors. Further augmentation can be anti-coincidence units

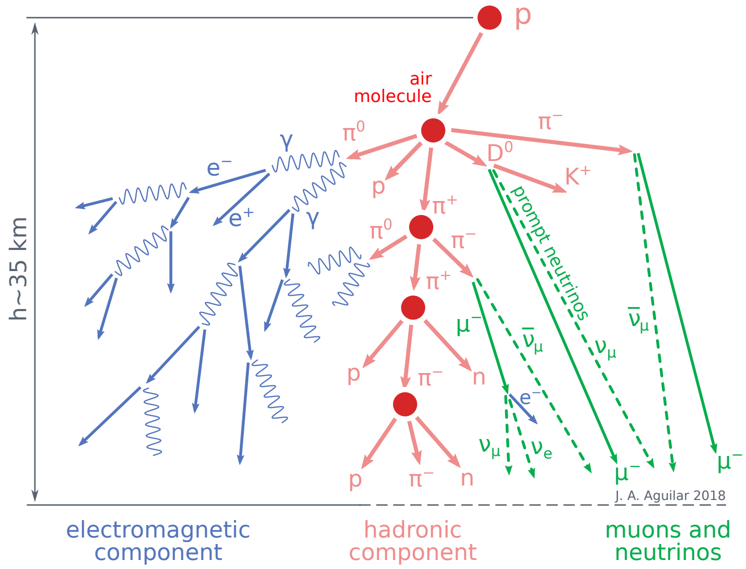

Figure 1.2.: Sketch of an air shower initiated by a proton, the subsequent interactions and decays, and the resulting particle populations at ground level.

for background suppression, or magnets e.g. in AMS-02 [15, 16] which turn it into a spectrometer which measures particle momentum, allowing to identify it. Among those focusing on nuclei, AMS and Pamela [17] are spectrometers, while ISS-Cream [18] is not. All together, satellite experiments cover an energy range of 10^8 eV to 10^{12} eV. Balloon experiments often use similar detector technology as satellite experiments. They are cheaper to launch with potentially higher payloads and therefore acceptances; on the other hand, they only reach as high as the stratosphere, and take data only for the duration of their flight which typically lasts from 10 up to 50 days. Examples of balloon experiments are Cream [18] and Atic-1/2 [19].

Due to the steeply falling cosmic ray spectrum, observation of cosmic rays at energies above the knee requires large detection areas, and therefore cosmic ray detectors on Earth's surface. While the atmosphere remains opaque to cosmic rays, at these energies it can itself act as a detector medium for ground-based observatories. When a cosmic ray nucleus interacts with an air molecule, a burst of new hadronic particles such as pions carries away the original momentum and energy. These may interact again, particularly if their energy is high, which creates an avalanche, called a **shower**, of hadronic particles. Those $\pi^{\pm}$, along with other mesons, which do not get to interact eventually decay into muons and neutrinos, whose long interaction lengths means they are lost from the shower. However, π^0 practically immediately decay into a pair of photons, each of which can undergo pair production $\gamma \xrightarrow{\text{nucleus}} e^+e^-$ and Compton scattering $\gamma e^-_{\text{at rest}} \rightarrow \gamma e^-_{\text{accelerated}}$ to transfer their energy onto electrons and/or positrons. These in turn lose energy in the atmosphere via Bremsstrahlung, in which a photon is radiated which can then repeat the process.

This electromagnetic shower is continually fed by its hadronic parent, which it quickly eclipses in terms of the total number of particles. As these numbers at first grow exponentially, the per-particle share of energy diminishes accordingly [4, sec. 2.3]. As a function of this, the cross section for Bremsstrahlung decreases [10, sec. 34.4]. Eventually, there are shower particles which instead lose most of their energy by ionization, which does not radiate further photons, and so halts the multiplication of the shower. Along with shower particles being absorbed by the air, this leads to the shower front no longer increasing in density, and then gradually dying out [4]. The hadronic shower similarly dies out when its pions no longer have the energy to start cascades and instead only can decay.

Air shower particles are sometimes called secondary cosmic rays, as opposed to the primary which initiated the shower. They are the ionizing particles through which cosmic rays were discovered. Figure 1.2 provides an a sketch of the previously described chains of interaction and decay.

The relevant longitudinal scale for shower development is the slant depth, i.e. the atmospheric density integrated along the shower path in units of g/cm^2. The slant depth at which a shower reaches the maximum particle production (usually referred as X_{max}) gives an indication of the energy of the primary cosmic ray particle but also its mass number. Two detection techniques are mainly used to observe the development of the air shower.

- **Air shower arrays** are ground-based assemblies of distributed individual detector stations within which the shower particles are detected. The timing between stations then allows to reconstruct the incident direction, assuming a relativistic shower front, and the primary energy, assuming models for the shower. This detection principle becomes practical for energies $\gtrsim 10^{13}$ eV with e.g. the High Altitude Water Cherenkov Observatory (HAWC) [20] currently in operation. Extending the spacing of the array means detection requires a higher number of shower particles, and therefore a higher cosmic ray energy. However at the same time, the larger area spanned by the array corresponds to a higher acceptance for the cosmic rays at these higher energies as long as they arrive from high enough above the horizon, i.e. the **field of view** is $\approx 2\pi \cdot$ area. Therefore air shower arrays have measured the spectrum up to 10^{20} eV [6, 7], despite the small flux in that regime (see fig. 1.1). The stations can use different technologies, such as water Cherenkov (HAWC [20], IceTop [21], Pierre Auger Observatory [22], Telescope Array [23]), scintillators and ionization detectors. They may be buried in order to shield anything except for the shower's muon component. Since these detectors do not rely on external conditions, they have a high duty cycle, and can be efficiently combined in one station to capture the shower more completely.

- Another technique is based on the **electromagnetic radiation** emitted by air showers along their path, such as geomagnetic radio emission, Askaryan radiation, or fluorescence light emission due to the excitation of nitrogen in air. For the latter, the duty cycle is limited to clear moonless nights. All these telescopes are sensitive to the longitudinal profile of an air shower, which can also allow for direction and energy reconstruction, and is uniquely suited to indicate the mass number of the primary cosmic ray nucleus. Some observatories purely use one of these methods [24, 25], however a hybrid design is more common. For example in Auger, fluorescence telescopes [22] and radio antennas [26] augment a subset of events, which can also help calibrate the surface detector; in LOFAR, particle detectors help trigger the data acquisition (DAQ) of an extended antenna array [27].

1.1.3. Arrival directions of cosmic rays

Figure 1.3 shows the cosmic ray sky as seen by the Auger and Telescope Array. As can be seen by the general isotropy in this sky map, these arrival directions of cosmic rays are influenced by magnetic fields via the Lorentz force. Its relativistic calculation for a particle of momentum p and charge $q = Ze$ results in a turning radius for a magnetic field strength B of

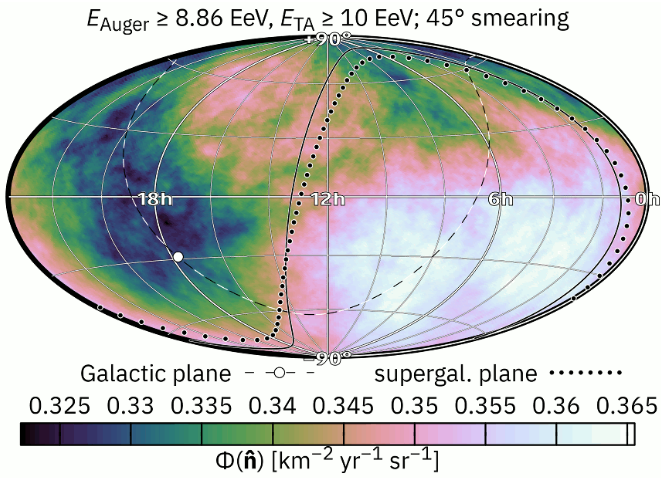

Figure 1.3.: Events from both Pierre Auger cosmic ray observatory ≥ 8.86 EeV, as well as Telescope Array ≥ 10 EeV, to cover the sky with their respective fields of view. Figure from [28].

$$r = p/(ZeB) \tag{1.2}$$

$$\approx \left(\frac{E}{10^{18}\,\text{eV}}\right)\left(\frac{1}{Z}\right)\left(\frac{\mu\text{G}}{B}\right)\text{kpc} \tag{1.3}$$

also known as the Larmor radius or gyroradius. The magnetic field within the Milky Way has a typical field strength of $B \approx 3\,\mu\text{G}$ [2], and therefore its regular component deflects Ultra-High Energy Cosmic Rays (UHECRs) [29] of energy E on their path of length scales $L{\approx}1$ kpc by

$$\delta \approx 3^\circ \times Z\left(\frac{6\times 10^{19}\,\text{eV}}{E}\right)\left(\frac{B}{3\,\mu\text{G}}\right)\left(\frac{L}{\text{kpc}}\right) \tag{1.4}$$

$$\approx 18^\circ \times Z\left(\frac{10^{19}\,\text{eV}}{E}\right) \tag{1.5}$$

The same kind of deflection can result from intergalactic magnetic fields, which must be weaker at $B \lesssim 10^{-9}$ G, but also act on paths of Mpc scale [30]. If the path of UHECR particles crosses an irregular magnetic field with a coherence length $\lambda_B \ll L$, they undergo a series of random deflections. These then accumulate to spread the cosmic ray directions by a root mean square (r.m.s.)

of $\delta_{\text{r.m.s.}} \propto \sqrt{L\lambda_B}$ [31, 30]. A cosmic ray's incidence direction therefore does not correspond to the direction of CR accelerators. However at high enough energies, inferences can be made. One success in this avenue is an analysis of the reconstructed directions of Auger events above 8 EeV which revealed an anisotropy with a 6.5% dipole excess at 5.2σ significance in a direction 125° away from the galactic centre [32]. The angular resolution of events at these energies is better than 1.6° [22, 33, 34] and the overall magnetic deflection of the primary nuclei with atomic number Z can be estimated as $30° \times Z \times (10\,\text{EeV}/E)$ [28]. Therefore a dominant galactic origin of these cosmic rays can be excluded.

If a particle's Larmor radius eq. (1.2) is smaller than the spatial extent of the responsible magnetic field, it will become trapped therein, although it may diffuse further than the Larmor radius. Approximating the Milky Way's magnetic field with a characteristic field strength of 3 µG [2], protons with energies of $\ll 10^{18}$ eV can not leave the galactic disk of thickness 300 pc and must be mostly galactic. This argument can be extended by assuming similar scales and magnetic field strengths for extragalactic acceleration sites.

Between these energies and the 8 EeV mentioned before one therefore expects the transition from galactic to extragalactic sources. The energy of this transition is still being disputed by orders of magnitude, since the underlying model of the cosmic ray spectrum is not determined. In the context of the question on which sources to target in observation it is therefore difficult to make an argument from energy. Nevertheless, from this point the dissertation gives priority to extragalactic cosmic rays because extragalactic source candidates in the form of blazars will be in the focus after chapter 5.

1.2. Cosmic ray acceleration

1.2.1. Fermi mechanism

Non-thermal processes are required to accelerate cosmic rays to their power law spectrum. The responsible mechanism is in principle unknown, but such spectra are predicted by a general principle called Fermi acceleration. Therein, the accelerated particle undergoes a cyclical process where at each cycle

1. its energy is increased by a constant factor $\alpha = \frac{\Delta E}{E}$
2. it has a constant probability P to interrupt the cycle

Injecting N_0 particles at energy E_0, after n such cycles there will be N_n particles with energy E_n left [2]:

$$N_n = N_0 \, (1 - P)^n \tag{1.6}$$
$$E_n = E_0 \, (1 + \alpha)^n \tag{1.7}$$
$$\Rightarrow n = \frac{\log(E/E_0)}{\log(1 + \alpha)} \tag{1.8}$$

Assuming that these will not lose energy, N_n is identical to the number of particles which will end the process, and ultimately escape, with an energy of E_n or higher. This "survival function" can be differentiated to determine the probability density function of emitted particle energy†:

$$N(\geq E) = N_0 \left(\frac{E}{E_0}\right)^{-\Gamma} \tag{1.9}$$
$$\text{where } \Gamma = -\frac{\log(1 - P)}{\log(1 + \alpha)} \tag{1.10}$$
$$\Rightarrow \frac{\mathrm{d}N_{\text{emitted}}}{\mathrm{d}E} \propto \left(\frac{E}{E_0}\right)^{-\gamma} \tag{1.11}$$
$$\text{where } \gamma = \Gamma + 1 \tag{1.12}$$
$$= -\frac{\log(1 - P)}{\log(1 + \alpha)} + 1 \tag{1.13}$$
$$\text{assuming } \alpha,\ P \ll 1 \tag{1.14}$$
$$\Rightarrow \gamma = \frac{P}{\alpha} + 1 \tag{1.15}$$

whereby the known power law is derived. The spectral index γ is not given, but can be determined by making assumptions about the physical mechanism responsible for P and α.

1.2.2. Diffusive shock acceleration

Diffusive shock acceleration [2, sec. 12.2] has become a baseline model. It assumes a planar shock wave of infinite size with ultra-sonic, but non-relativistic

†adapted from [2]

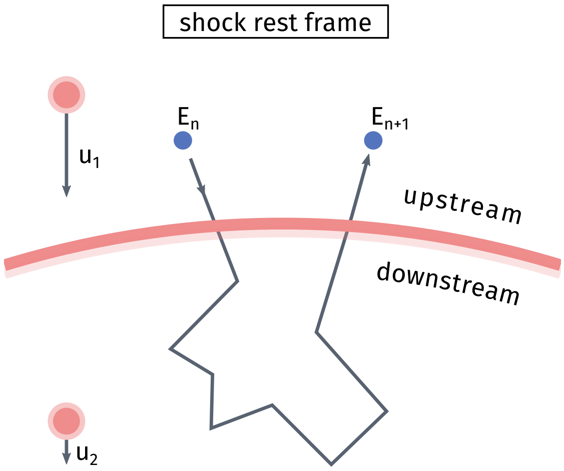

Figure 1.4.: A diagram of diffusive shock acceleration. The velocities involved are in the reference frame of the shock wave, which moves away from the energy source on the downstream side. The trajectory of a test particle, i.e. not interacting with any other particle population, is being shown in one cycle of Fermi acceleration as a solid line.

velocity βc in a fully ionized gas (a plasma). Its principles can be extended to more realistic geometries such as the cross section of a jet or a spherical shell [2, sec. 12.3]. A diagram of this process is shown in fig. 1.4. The shock is a moving discontinuity in the bulk velocity of the gas. As shown in the diagram, the shock approaches the unshocked, "upstream" plasma with speed u_1 in the latter's frame of reference. The shock leaves behind the shocked, "downstream" plasma, which in the shock reference frame recedes with speed u_2. The key idea behind shock acceleration is that any particle crossing the shock wave will always encounter a plasma flow moving towards it at a velocity $u_1 - u_2$ [35, sec. 1.1.2].

In this framework, the process necessary for Fermi acceleration consists of the CR particle first passing from one side of the shock to another, then being isotropically diffused without energy loss. The average energy gain in this mechanism can be calculated to be $\frac{\alpha}{2} = \frac{2}{3}\frac{u_1 - u_2}{c}$ per crossing in either direction [35]. This linear relation in $\beta = \frac{u_1 - u_2}{c}$ lends the mechanism the name of first-order Fermi acceleration, as compared to the second-order with $\alpha \propto \beta^2$ which failed to be sufficiently efficient.

In a full cycle, the particle passes from the shocked downstream region to the unshocked upstream, until the shock catches up with it and it passes back across. Then it has a chance of "escape" by way of convection, which means that the component of its velocity normal to the shock front does not suffice to catch up with it again. This probability is $P = \frac{4u_2}{c}$ [2]. On the other hand, as the particle diffuses into the upstream region, the shock will always catch up with it eventually.

One of the diffusive shock acceleration theory's successes is that it managed to give an estimate of the spectral index at the source. Since the universe is mostly made of hydrogen, the plasma and the shock therein can be described by the kinetic theory of monatomic gases. If the shock is by far supersonic, then $u_1 = 4u_2$‡. Inserting this relation into the result for α gives:

$$\alpha = \frac{4}{3}\frac{u_1 - u_2}{c} \tag{1.16}$$
$$= \frac{4u_2}{c} \tag{1.17}$$
$$\Rightarrow \gamma = \frac{P}{\alpha} + 1 \tag{1.18}$$
$$= 2 \tag{1.19}$$

Accounting for a shock that may be less super-sonic than assumed leads to the correction $u_1/u_2 = 1 + \frac{3}{1+\frac{4}{\mathcal{M}^2}}$ [2] where $\mathcal{M}$ is the Mach number, or u_1 relative to the speed of sound in the plasma. Since the only assumption considered is that of a monatomic gas, this prediction of the spectral index is in fact universal and independent of the specific source of acceleration.

1.2.3. Cosmic accelerators

Like the mechanism, the environments responsible for cosmic ray acceleration are also unknown, and so the subject needs to be resolved with observation of accelerator candidates.

But also for the environment, general requirements can be derived, particularly from the assumptions of first-order Fermi acceleration. The existence

‡using the ratio of specific heats, $c_p/c_v = 5/3$, and conservation of mass at the shock, $\frac{u_1}{u_2} \approx \frac{c_p/c_v+1}{c_p/c_v-1} = 4$ [2]

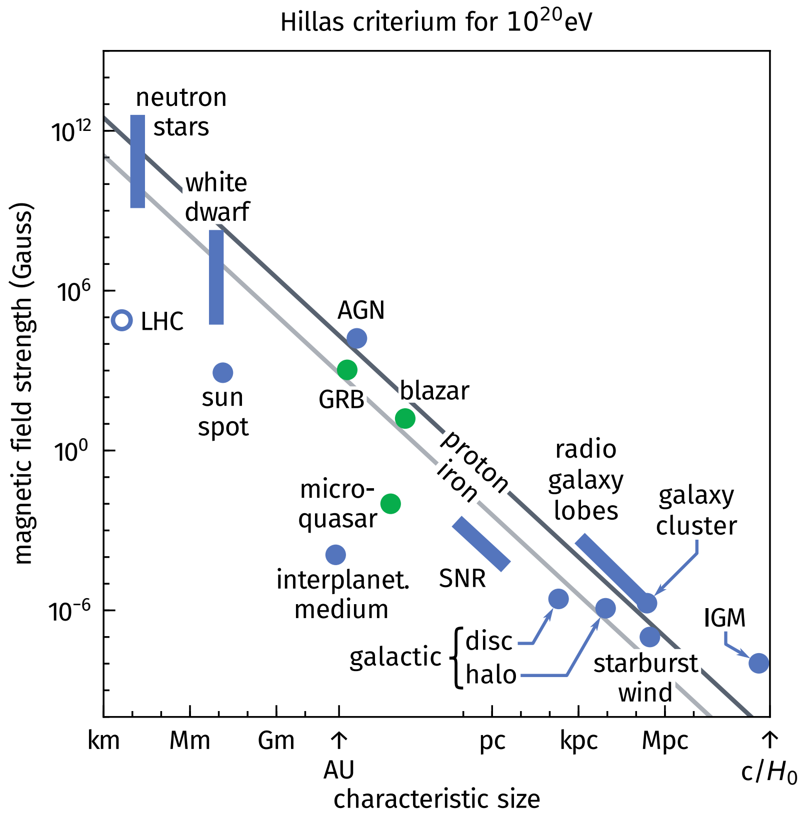

Figure 1.5.: The Hillas plot, adapted from [36], using [37]. The dark solid line shows the Hillas criterion to reach a 10^{20} eV proton, the lighter line the same for an iron nucleus. The actual cosmic ray composition at this energy will be in between iron and protons. The charm of this plot is that it can put wildly different source classes, shown as points and bars (green for jets), into a common context and, for some, reveals their surprising compatibility.

of shocks in an ionized medium is true for many objects such as supernova remnants and active galactic nuclei (AGN) [35] (see chapter 5). But beyond that, particles need to be contained by the environment magnetic field in order to have an opportunity to accumulate energy from the acceleration mechanism. The minimal requirement for a relativistic particle with momentum p and charge q in a magnetic field of magnitude B is that the Larmor radius $R = E/(q\ B)$ does not exceed the environment scale L. Rearranged, this translates into a maximum energy, in typical units [38]

$$E_{\max} < qBL \tag{1.20}$$

$$\lesssim Z \left(\frac{L}{\mathrm{kpc}}\right) \left(\frac{B}{10^{-6}\,\mathrm{G}}\right) 10^{18}\,\mathrm{eV} \tag{1.21}$$

also known as the Hillas criterion [39]. This allows to categorize astrophysical objects, and classes thereof, by their size and magnetic field strength in order to determine which can be considered as candidate accelerators for cosmic rays of a particular energy, such as in fig. 1.5. It is not a sufficient criterion for these candidates, since other properties of the environment also limit the maximum energy, such as the shock lifetime and speed. Due to conservation of energy, the power of emitted cosmic rays also can not exceed what a central engine provides to produce the shock wave.

1.3. Cosmic ray propagation

After being accelerated (section 1.2), cosmic rays can only be observed (section 1.1) if they propagate to the observer, which in the case of human astronomy is in the vicinity of the planet Earth in the Milky Way.

1.3.1. Scattering and diffusion for galactic sources

The galactic magnetic field changes at different scales [40, 29]. Scattering occurs via the Lorentz force on irregularities which have a similar length scale as the particle Larmor radius (see eq. (1.2) and fig. 1.6) [2, pp. 187]. This is the case for galactic cosmic rays of energies $\lesssim 10^{12}$ eV which satellite experiments like Pamela and AMS (see section 1.1.2) can observe and identify by nuclear species.

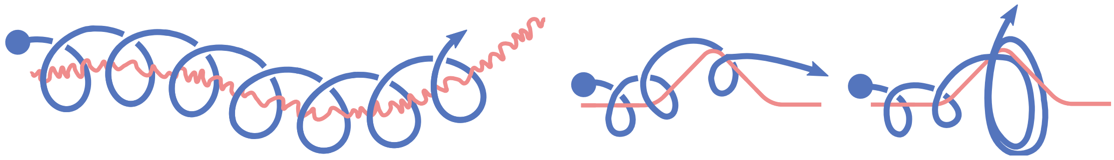

(a) The particle follows irregularities much larger than its Larmor radius and is unaffected by those much smaller.

(b) Two examples for scattering when the irregularities have a similar scale as the Larmor radius, adapted from [2, fig. 9.1].

Figure 1.6.: A charged particle (blue) in an irregular magnetic field (red).

The scattering leads to a random walk through the galactic disk, which can be characterized statistically with a simple diffusion model. However due to the high hydrogen density, cosmic ray nuclei are also subject to spallation, such as $\mathrm{C_{CR}} + p_{\mathrm{target}} \rightarrow \mathrm{B} + X$ where the energy per nucleon is approximately conserved. Spallation of carbon or oxygen is the dominant way in which boron is produced in the galactic disc. The production, diffusion, and further fragmentation of boron in the galactic disc, modelled as a homogeneous cylinder (for example $R = 30\,\mathrm{kpc}$, $h = 300\,\mathrm{pc}$) is summarily called the "leaky box model" [41, p. 20]. The amount of diffusion in the galaxy can be estimated by measuring the B/C ratio in the galactic cosmic ray composition vs. energy, under the assumption that the ratio is constant in time. The B/C ratio measurements clearly show that higher energetic particles diffuse away faster, while lower energetic particles persist for a longer time in the galaxy [42]. Furthermore, this diffusion dependence on energy is also a power law with a mean time to escape the leaky box $\tau_{\mathrm{escape}} \propto E^{-\delta}$, which for protons dominates over losses to interactions [41, p. 24]. The observed spectrum is therefore softer than that at the source, with $\gamma_{\mathrm{obs}} - \gamma_{\mathrm{src}} = \delta = 0.33$ [42]. At energies below $\approx 10\,\mathrm{GeV}$ meanwhile, energy loss effects influence the spectral index as well [2, pp. 193].

1.3.2. Deflection and energy loss for extragalactic sources

At higher energies, the Larmor radius is longer and cosmic rays experience the average magnetic field over the same scale, passing over (part of) the random irregularities in the Galactic magnetic field. As already described in section 1.1.3, this leads to a deflection, in combination with the effect of intergalactic magnetic fields.

Spallation, which has been discussed in section 1.3.1, requires interstellar hydrogen which is not abundant in the intergalactic space. On the other hand, cosmic rays do encounter background photons in their path. These originate from several sources of which the CMB is by far the most abundant, with $E_\gamma \approx 10^{-4}$ eV, followed by the extragalactic background light (EBL) at shorter wavelengths, as will be discussed in section 1.5. This causes cosmic rays to lose energy via $X_{\mathrm{CR}} - \gamma_{\mathrm{target}}$ interactions, however only for cosmic rays above a specific threshold energies which correspond primarily to when the reaction becomes possible with CMB photons.

Protons can undergo Bethe-Heitler pair production $p\gamma \to p\ e^+\ e^-$ if they exceed the corresponding energy threshold $\approx 2 \times 10^{17}$ eV [43] at which point this quickly becomes the dominant process for intergalactic energy loss.

Other than electron-positron pairs they can also produce pions such as in eq. (1.25), which will be discussed in more detail in section 1.4. The energy threshold for this is around $E_{\mathrm{th}} = 5 \times 10^{19}$ eV [44] , but the effect is amplified via the $\Delta^+(1232)$ resonance at $E_{\mathrm{res}} \sim 2 \times E_{\mathrm{th}} \approx 10^{20}$ eV §. This effect leads to the Greisen-Zatsepin-Kuzmin (GZK) limit on the energy of cosmic ray protons.

Cosmic rays also undergo a cosmological redshift due to the expansion of the universe, where the space-time metric gradually changes between the times of emission and observation. All cosmic rays regardless of energy lose energy this way, but it is only the dominant process for those below the threshold energy for Bethe-Heitler pair production [43, sec. 1.2].

Heavier nuclei meanwhile have low enough binding energy that they may disintegrate by interaction with the CMB in a process called photo-disintegration. Once the loss length is short enough, this effectively leads to a cutoff in energy, e.g. at 3×10^{20} eV for iron nuclei [35]. This super-imposes on the $E_{\max}$ of the accelerators themselves.

1.4. Cosmic ray reactions in the accelerator

Before cosmic rays start to propagate through interstellar or intergalactic space, they may encounter photon fields and matter in (or near) the acceleration environment. For the sake of simplicity, we focus on the reactions of protons in this section and forego the discussion of other accelerated nuclei.

§ $\frac{E_{\mathrm{res}}}{E_{\mathrm{th}}} = \frac{m_\Delta^2 - m_p^2}{m_\pi^2 + 2 m_\pi m_p} \approx 2.3$ [9]

Relativistic protons can react with matter, mostly hydrogen, in inelastic proton-proton collisions. This will produce secondary particles such as pions, gamma rays, electrons, positrons and neutrinos [45]. The energy threshold for pion production is relatively low ≈280 MeV in terms of kinetic energy of the proton [45]. In astrophysical environments where a radiation field of low-energy photons exceeds that of matter by orders of magnitude the p-γ interaction is more efficient [46, 47, 48]. This however has a higher energy threshold which can be expressed as a function of the photon energy in the same frame of reference E_γ:

$$E_{\mathrm{th}} = \frac{2m_{\mathrm{p}}m_\pi + m_\pi^2}{4E_\gamma} \tag{1.22}$$

$$\approx 7 \times 10^{16}\,\mathrm{eV}\frac{\mathrm{eV}}{E_\gamma} \tag{1.23}$$

$$\tag{1.24}$$

The target photons can for instance be synchrotron radiation by electrons accelerated in the same environment [43, p. 31], whose energies fall in the optical – X-ray band [49, ch. 8]. This process is therefore more relevant for high-energy extragalactic cosmic rays (see section 1.1.3). The following will focus on p-γ interactions although similar conclusions can be drawn from p-p interactions. Interactions of p-γ at threshold occur via the lightest Δ^+ resonance, which decays with a branching ratio of 99.4% [9] into [44, p. 42]:

$$\mathrm{p} + \gamma \rightarrow \Delta^+ \begin{cases} \xrightarrow{2/3} \mathrm{p} + \pi^0 \\ \xrightarrow{1/3} \mathrm{n} + \pi^+ \end{cases} \tag{1.25}$$

In both cases, the resulting pion is produced with an average energy of $\langle E_\pi \rangle \simeq E_p/5$ [50]. The pions will decay to stable particles before interacting, and assuming a low ambient photon density the same applies to the neutrons [43, p. 31]:

$$\mathrm{n} \xrightarrow{100\%} \mathrm{pe}^{-}\overline{\nu}_{\mathrm{e}} \tag{1.26}$$

$$\pi^{0} \xrightarrow{98.8\%} \gamma\gamma \tag{1.27}$$

$$\pi^{+} \xrightarrow{100.0\%} \mu^{+}\nu_{\mu} \xrightarrow{100\%} e^{+}\nu_{\mu}\overline{\nu}_{\mu}\nu_{\mathrm{e}} \tag{1.28}$$

$$\pi^{-} \xrightarrow{100.0\%} \mu^{-}\overline{\nu}_{\mu} \xrightarrow{100\%} e^{-}\overline{\nu}_{\mu}\nu_{\mu}\overline{\nu}_{\mathrm{e}} \tag{1.29}$$

As indicated by the branching ratios [9] above the arrows, these decay paths are almost exclusive. Due to the neutron lifetime of 880 s [9], its decay can also provide (secondary) cosmic ray protons from acceleration environments in which the primary protons might be magnetically confined [43, p. 31].

Therefore the same sources that can accelerate cosmic rays give rise to neutral messenger particles in the form of both (gamma-ray) photons and (high energy) neutrinos. These messengers and their relation will be discussed in the following sections.

1.5. Gamma rays as messengers from p-γ collisions

In the symmetric decay $\pi^0 \to \gamma\gamma$, each photon receives on average half the pion energy, and so $\langle E_\gamma \rangle = \langle E_\pi \rangle / 2 = E_p/10$ for p-γ interactions. This easily puts them into the open range >100 keV called gamma rays. The associated spectrum depends on the interaction cross-section with the ambient radiation, but in any case reproduces the non-thermal character of the beam proton spectrum [46].

Photons which reach the observer are red-shifted due to the expansion of the universe, depending on the distance of the source i.e. the age of emission. During their propagation, gamma-ray photons also encounter multiple effects mediated by the electromagnetic interaction.

Gamma rays are absorbed via pair production interactions $\gamma_{\mathrm{source}}\gamma_{\mathrm{target}} \to e^+e^-$ [52], where the target photons belong to the EBL and the CMB. Thereby the universe becomes opaque for gamma rays $\gtrsim 100$ GeV, called Very High Energy (VHE), with an energy-dependent horizon plotted in fig. 1.7.

Gamma rays can also be absorbed by clouds of dust, either in the path or more commonly near the source, or cascade down in optically thick environments, leading to obscured sources where only neutrinos could escape [53, 54].

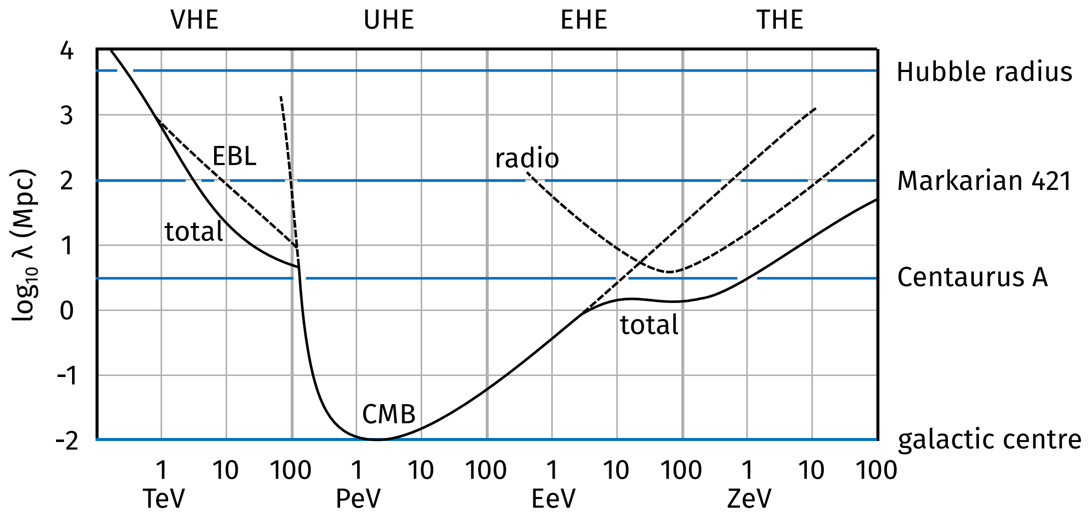

Figure 1.7.: The mean free path or horizon for gamma rays due to absorption by the EBL, CMB, and radio background vs. the gamma-ray energy. The individual contributions are drawn in dashed lines, the combined effect in a solid line. Adapted from [51].

1.6. Neutrinos as messengers from p-γ collisions

At this point we focus on the neutrinos resulting from the $\pi^{\pm}$ decay in eq. (1.26). Each lepton resulting from that decay receives approximately the same energy, leading to an average energy $\langle E_\nu \rangle = \langle E_\pi \rangle/4 = E_p/20$ relative to that of the beam proton.

The neutrino spectrum can be calculated the same way as the photon spectrum. The difference is that neutrinos do not interact electromagnetically, so neutrinos are essentially not absorbed and the the spectrum will be unchanged in propagation except via cosmological redshift. This makes neutrinos useful to look within even dense accelerator environments and sources beyond the EBL horizon (see section 1.5). However, due to neutrino flavour oscillations, the flavour composition will change. At the acceleration site there are neutrinos and anti-neutrinos in a ν_e:ν_μ:ν_τ=(1:2:0) proportion. For the extragalactic sources in question, the oscillation mixes this to a (1:1:1) composition within a typical baseline to the observer. Their detection at this point will be discussed in more detail in chapter 2. Until then, it is important to note that currently, flavour ratio measurements of the astrophysical neutrino flux can not exclude

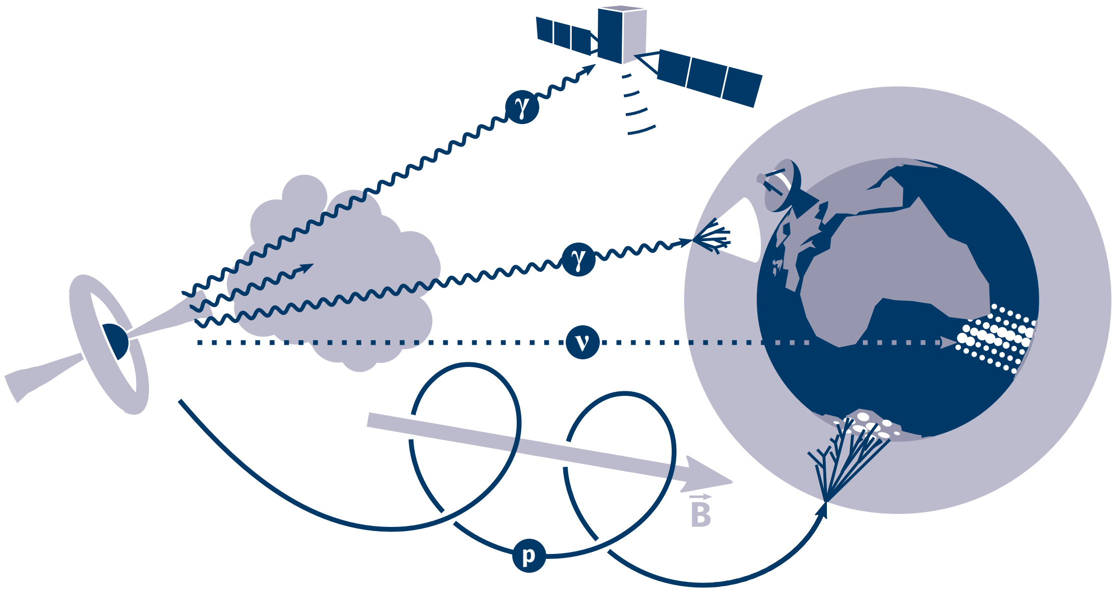

Figure 1.8.: A cosmic ray accelerator with gamma rays and neutrinos as messengers, detected at Earth. Not to scale.

any composition at its sources [55, 56].

1.7. Multimessenger principle

The identification of cosmic ray accelerators is challenged by the fact that they arrive deflected (see section 1.1.3), apart from the highest energy cosmic rays, for which not enough events can be collected. Other than ongoing efforts in tracing back to sources or source distributions through improved models of the galactic magnetic field [31, 40, 29] or future instruments which improve sensitivity to the highest cosmic ray energies by observing air showers from space [57, 58], the main hope lies in the messenger particles which can be produced as in section 1.4 and being electrically neutral will arrive undeflected at the observer. Figure 1.8 is a schematic of this principle.

Many sources of gamma rays have been identified, and to some extent classified [59]. However gamma rays can also be produced in leptonic scenarios such as inverse Compton scattering, rather than hadronic scenarios like the previously described π^0 decay. Additional observation ranging from radio waves to X-rays, called multi-wavelength (MWL) astronomy, provides a wealth of information about the source. It has historically been the most powerful tool to understand extreme astrophysics, however it is very challenging to distinguish between leptonic and hadronic gamma-radiation scenarios using MWL astronomy alone.

High-energy neutrinos on the other hand only arise in hadronic interactions and are therefore unambiguous evidence of cosmic ray acceleration. However their low interaction cross section makes them difficult to detect, so that while an approximately isotropic astrophysical neutrino flux has been discovered [60, 61], no sources could be identified with the current instruments [62]. Neutrino source searches therefore benefit from focusing on known astrophysical objects to reduce background. In addition, an eventually discovered neutrino source will need to be identified with a known object in order for its role in cosmic ray acceleration to be understood in more detail.

The connection between high-energy neutrinos and gamma rays as described in section 1.4 makes this a promising avenue. In this simple model, neutrinos and gamma rays would not only originate from the same direction, but due to the assumed decay chain also at the same time and in a fixed proportion, providing an even more powerful hypothesis.

Notable efforts which exploited this multimessenger principle were for example:

- Measurements of the Sun's electron-flavour neutrino emission, which confirmed that the Sun was undergoing stellar fusion. These measurements also showed a discrepancy with the solar model, itself derived from electromagnetic observations; this discrepancy ultimately led to the discovery of neutrino oscillations [63]. This can be considered the start of multimessenger astronomy, even if the target was still within the solar system.

- The first detection of neutrinos from a supernova SN 1987A [64], which laid foundations for multimessenger astronomy outside the solar system and advanced science on supernovae.

- The first observation of a multimessenger connection between gravitational waves and photons on the Gamma Ray Burst (GRB) 170817A coinciding with the gravitational wave event GW170817 [65]. Beyond the actual science, this was also instructive as the most extensive multimessenger campaign to date and proved the high promise inherent in gravitational wave observatories.

- The multimessenger and multi-wavelength follow-up campaign on a neutrino from the direction of the blazar TXS 0506+056, which will be discussed at great length in section 8.1.

1.8. Current multimessenger efforts

Since the discovery of cosmic rays, physicists realized that extensive collaboration was needed to answer the open questions about their accelerators. Particularly current efforts in multimessenger astronomy require a global, unhindered, and sometimes near-instantaneous flow of data and publications. Multimessenger astronomy has grown to involve all types of telescopes, which cover a broad spectrum of energies, and all messengers with gravitational waves as the most recent addition (see section 1.7). The necessary communication channels run through networks of both technical and organizational nature.

1.8.1. Roles of different telescopes

The pertinent distinction is between monitoring and pointed telescopes. Monitoring telescopes consistently capture a significant part of the sky with a high duty cycle. For neutrino telescopes, like IceCube, this is the case since at any moment an event can arrive in the detector from anywhere in the sky and its operation is not subject to the atmospheric conditions. Air shower detectors, such as those used for UHECR (see section 1.1.2) have a $\sim 2\pi$ aperture, but cover a wider field of view in the course of every 24h due to the Earth's rotation. Other instruments have a smaller aperture still, but still scan the sky within a defined time scale, either by virtue of their host satellite's rotation and orbit (like the gamma-ray telescope Fermi-LAT, see section 4.2), or systematically through preprogrammed schedules.

On the other hand, pointed telescopes are limited by the small part of the sky they can observe at a time, or Field of View (FoV). These telescopes usually take the source direction as the input in order to observe it, and in case of a transient event this input information sometimes has to be delivered within very short time scales.

1.8.2. Real-time

Alerts are produced when a monitoring telescope detects a transient with a certain significance versus the background. This might mean a gamma-ray brightness fluctuation (e.g. in Fermi-LAT's FAVA programme [66]), or in instruments with lower signal-to-background ratios, a particularly significant, interesting event or event cluster, as in the IceCube neutrino telescope's real-time programmes, described in section 2.5. The alert is then received by instruments which can perform a follow-up observation on the same direction and relevant time period. For other monitoring telescopes, the involved collaborations can perform follow-ups based on the recorded data, and therefore the time scales of human reaction are short enough to produce results. However for pointed telescopes, the follow-up needs to happen on the time scale of the transient itself, which might be near real time, before the observation opportunity has passed.

This strategy therefore requires technical infrastructure, as relying entirely on humans would introduce unpredictable latency. Networks that aggregate alerts from multiple telescopes make the communication more reliable.

- Since 1997, Astronomer's Telegram (ATel)¶ serves a wide community of professional astronomers, both individual and representing larger collaborations, to disseminate publicly accessible messages. These "telegrams" consist of the author's contact information, a free-text message, and a subject line consisting of multiple keywords which can specify: the type of object, from within the solar system to outside the galaxy; the observation of neutrinos, electromagnetic radiation at a wavelength from radio to VHE gamma rays, or gravitational lensing; a request for observation for an opportunity opening in the next 72 h. ATels concerning multimessenger follow-ups can have both messengers or wavelengths in the subject [67]. Within <1 s of reception, each "telegram" is posted to the website, and, if requested by the author for the purpose of transmitting new object or transient coordinates, forwarded by email to any readers subscribed to one of the telegram's keywords. Other telegrams arrive at the subscribers in a daily email digest.

- The SuperNova Early Warning System (SNEWS)‖ in contrast serves a spe-

¶ `http://www.astronomerstelegram.org/`
‖ `https://snews.bnl.gov/`

cific purpose, established 1998 as a consequence of SN1987A [68]. In a gravitational collapse supernova the $\mathcal{O}(10\,\mathrm{s})$ burst of neutrinos produced by inverse beta decay escapes — and thereby arrives — hours before any photons. A network of neutrino detectors can be the first to alert other telescopes of a sufficiently close supernova, which at the moment is limited to within the galaxy for the involved experiments of Super-K, LVD, IceCube, Borexino, KamLAND, Daya Bay and HALO. All these independently feed their respective supernova burst candidates into a central system, which automatically and blindly produces an alert only in the case of a coincidence within 10 s [69], which is classified by significance and transmitted via a public mailing list, as well as a dedicated service for amateur astronomers.

- The Gamma-Ray Coordinate Network (GCN) ** broadcasts coordinates of transients in the form of machine-readable messages, automatically generated by a number of approved senders. It also relays corresponding follow-up messages written as prose by a wider community. Its purpose originates from the discovery of optical counterparts to GRBs on archival astrographic plates and has since been extended to other astrophysical transients or candidates thereof. The network grew around existing infrastructure of the satellite-based gamma-ray telescope Batse when in May 1997 additional telescopes started becoming involved.
- The Astrophysical Multimessenger Observatory Network (Amon) [70] is designed to receive background-dominated event streams. It offers a platform for predefined, automatic real-time analyses which correlate these events in time and space. Coincidences might indicate a multimessenger connection, and can be ranked by their respective analysis likelihood in terms of significance or false-alarm rate. Amon then would send the most significant ones in a machine-readable alert.

  This allows for rapid follow-up observations. Amon has signed Memoranda of Understanding (MoUs) to share data with several observatories. For the role of triggering, it currently receives 1000 events per day according to filters decided by the observatories themselves, while others are intended to perform follow-up (see section 1.8.1), and some do both.

**`https://gcn.gsfc.nasa.gov/`

Together, these cover neutrinos, cosmic rays, and photons from radio to VHE gamma rays.

- Alert Management, Photometry and Evaluation of Lightcurves (AMPEL) [71] is an analysis framework and operational system developed to filter, distribute, (re)process and analyse alerts originating from the Zwicky Transient Facility (ZTF). The latter performs optical surveys of the Northern Sky, which are analysed for variable point sources, thereby compressing the high-throughput data stream to 20 transient detections per second.

1.8.3. Archival

Coincidence analyses can also search for a common signal in sub-threshold data, i.e. data in which no individual event can be attributed to a certain signal in light of a dominant background. During their typically high duty cycles, monitoring telescopes produce large archival sets of such data. Combining these data sets for a follow-up is a complementary approach to find a multimessenger connection. In contrast to real-time, these analyses are not limited to transient phenomena, but can include time-integrated emission hypotheses.

Observatories share their archival data in different ways. For Swift and Fermi, the raw event data is accessible publicly according to NASA policy, via widely used protocols like HTTP and FTP and updated daily. They also provide tools to analyse this data and produce derived data products like lightcurves and source catalogues on longer, but regular time scales.

The neutrino telescopes IceCube and ANTARES also provide public releases of their data. In addition, many collaborations share data privately according to MoUs with each other.

Examples of archival data multimessenger analyses can be found in this dissertation, particularly chapters 3 and 6 to 9.

"It's pretty clear we had no idea what we were doing, and so this was real research, right?"

Francis Halzen [72]

2 Neutrino Astronomy with IceCube

The previous chapter 1 has highlighted the importance of neutrinos as cosmic messengers. This chapter explains the principle of neutrino astronomy and how it is realized by the IceCube South Pole Neutrino Observatory. The description of the event selection is however deferred to chapter 3 which concludes with a neutrino data set, while the statistical method to analyse the latter is given in chapter 7.

IceCube is a neutrino observatory installed at the geographic South Pole. Its primary operational principle is to detect the Cherenkov light produced by the primary and secondary particles resulting from highly energetic interactions within and near a cubic kilometre volume of glacier ice. This is achieved using 5160 10-inch photomultiplier tubes (PMTs) integrated into autonomous data-taking units called digital optical modules (DOMs). 60 DOMs are installed at depths between 1450 m and 2450 m and connected by a vertical cable also called a **string**. The 86 strings converge on the IceCube Laboratory (ICL) at the surface. There, the per-DOM readout is bundled into events when the detector triggers on a likely physical interaction. These events are then further reconstructed, selected, stored and transmitted to data centres in the Northern hemisphere, where they are used in physics analyses. IceCube is the largest neutrino telescope on Earth. Since its completion in 2011 it has proven itself in reliability, stability and versatility, observing the entire sky for neutrino energies from 100 GeV to PeV with 97–98% clean uptime. By conservative estimates, the detector is expected to still perform with $97.4 \pm 0.3\%$ of its DOMs in 2030. Part of IceCube is the DeepCore infill array with 8 strings featuring a closer vertical and string spacing and high-efficiency PMTs, optimized for lower energies down to $\mathcal{O}(\mathrm{GeV})$. IceCube also works in tandem with the IceTop surface array consisting of two water tanks at the top of 81 strings, each with 2 PMTs. IceTop is built to detect cosmic-ray air showers which form the main background for IceCube.

This chapter gives an overview of IceCube, particularly the relevant points for the work presented in this thesis. For a more detailed description of: the detector see [73]; the alert system in section 2.5 see [74, 75, 76]; the studies for section 2.7 see [77].

2.1. Detection Principle

2.1.1. Neutrino interactions

Neutrinos are electrically neutral and do not participate in the strong interaction, so that on the microscopic level, they are only subject to the weak force. Therefore their interactions can be described by the exchange of a $W^{\pm}$ boson, in a charged current (CC) interaction (see fig. 2.1a), or a Z^0 boson, in a neutral current (NC) interaction (see fig. 2.1b). For energies $E_\nu \gtrsim 20\,\mathrm{GeV}$ neutrinos start to interact with the quarks inside hadrons dominantly via Deep Inelastic Scattering (DIS) [78], resolving the inner structure of nucleons in the target material.

The total neutrino cross-section (neutral and charged current) can be approximated up to 300 GeV as $\sigma \propto E_\nu$, with $\sigma / E_\nu = \mathcal{O}(10^{-38}\,\mathrm{cm}^2)$ [78].

At energies above 3 TeV the propagator term of the cross-section suppresses the cross section's linear growth with neutrino energy to $\sigma \propto E_\nu^{0.3}$ [79].

At $\sim 100\,\mathrm{TeV}$ the interaction length for neutrino-nucleus interaction is about 10^8 mwe. This is similar to the column depth integrated along a path which passes through the Earth core, as compared in fig. 2.1d. Therefore the high energy neutrino flux above these energies is suppressed along the nadir direction [80].

Neutrinos interact mostly with hadrons as the neutrino-electron cross-section is negligible because of its small target mass [78]. There is however one exception as the energy of the neutrino approaches the "Glashow resonance" at 6.3 PeV, where $\overline{\nu}_\mathrm{e} + e^- \rightarrow W^-$ raises the $\nu - e$ cross-section higher than that with nucleons [78]. At these energies the Earth is opaque to neutrinos, as seen in fig. 2.1d.

2.1.2. Event topologies

Once a neutrino crosses the Earth and arrives at the vicinity of the detector, it might interact with a nucleus of ice or bedrock. If the interaction vertex lies inside the detector medium, the nuclear remnant produces a shower of hadronic particles which for energies of 100 TeV typically propagate as far as $\mathcal{O}(100\,\mathrm{m})$ [82]. As shown in fig. 2.1, NC interactions produce a neutrino of the same flavour as the original, which escapes without interacting again, but for CC interactions, this is instead a charged lepton. For $\nu_\mathrm{e} \xrightarrow{\mathrm{CC}} e^-$, the electron cascades into an electromagnetic shower (see section 1.1.2), overlapping with the hadronic shower.

IceCube calls this overall event topology a **cascade**, see fig. 2.2a. A cascade is characterized by many particles travelling in a small region with a high dispersion in momentum.

In $\nu_\mu \xrightarrow{\mathrm{CC}} \mu^-$ the muon is less easily stopped and can propagate in ice or water (density $\approx 1\,\mathrm{g/cm^3}$) over $\mathcal{O}(10\,\mathrm{km})$, losing energy along its path to the medium and to secondary particles. IceCube calls this muon-dominated event a **track**, see fig. 2.2b. A track is characterized by a single particle travelling a long distance without noticeable deflection. In the ice with a density $0.922\,\mathrm{g/cm^3}$ [83], the muon will traverse a path length r of approximately

$$r(E_\mu^i, E_\mu^f) = 2.6\,\mathrm{km} \times \log\left[\frac{1 + E_\mu^i/(0.48\,\mathrm{TeV})}{1 + E_\mu^f/(0.48\,\mathrm{TeV})}\right], \tag{2.1}$$

to reach the energy E_μ^f from an initial energy E_μ^i [82]. This results from the energy loss along the path $-\frac{\mathrm{d}E_\mu}{\mathrm{d}x}$ which is $\approx 0.18\,\mathrm{GeV/m}$ in IceCube's lowest energy range, while a second term $\propto E_\mu$ dominates for $E_\mu \gtrsim$ TeV [82].

Muons will emerge at a **kinematic angle** away from neutrino direction, which is about $1°/\sqrt{E_\nu/\mathrm{TeV}}$ [82]. This makes them almost colinear for high-energy neutrinos, allowing for the identification of neutrino point sources.

Finally, for $\nu_\tau \xrightarrow{\mathrm{CC}} \tau^-$, the τ lepton decays, producing a second shower displaced from the interaction vertex. The tau decay path is of the order of

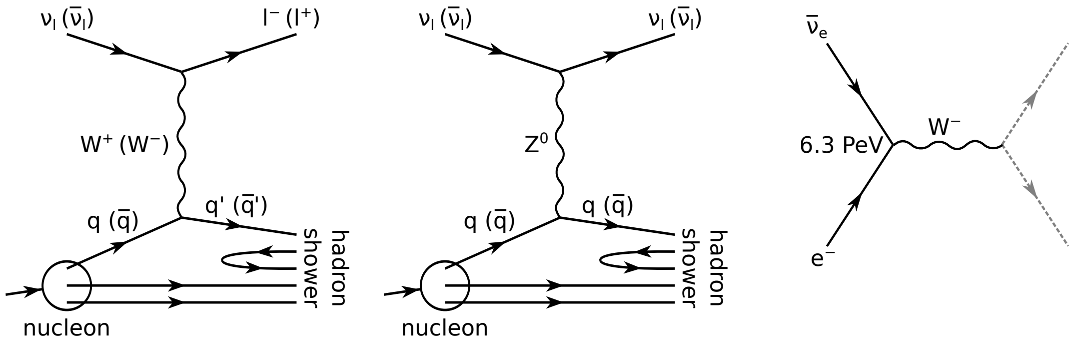

(a) Deep inelastic CC scattering on a nucleon, l $\in$ $\{e, \mu, \tau\}$.

(b) The same for NC.

(c) Glashow resonance.

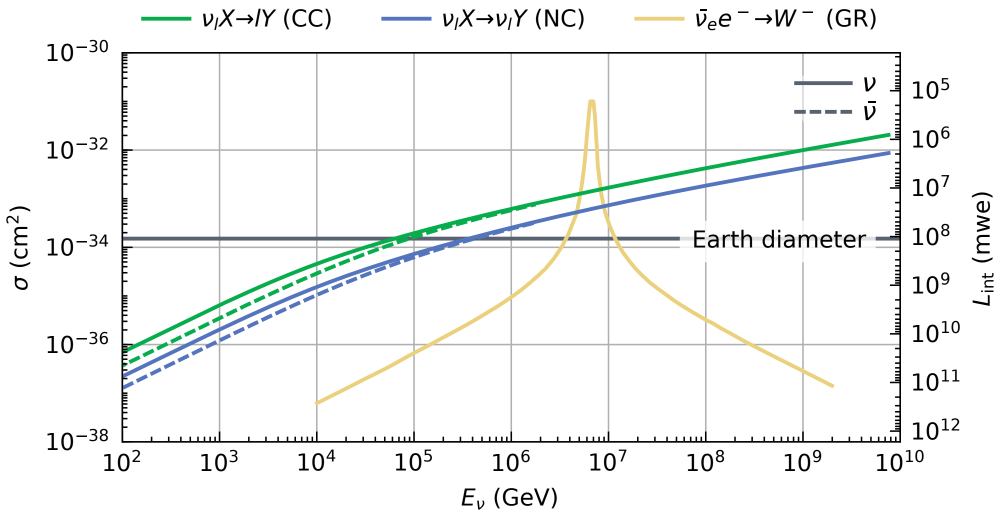

(d) Cross sections for the above interactions. The Earth diameter in meters water equivalent (m.w.e.) applies to the nucleon cross sections.

Figure 2.1.: Neutrino interactions, with first-order Feynman diagrams adapted from [81] with [35]. The corresponding cross sections are adapted from [38], using data and calculation from [80].

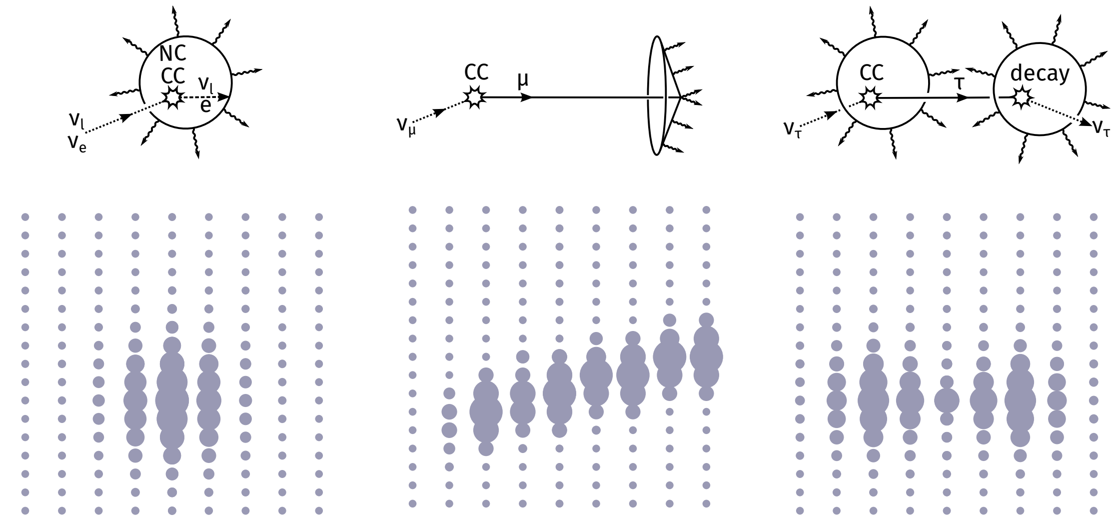

(a) cascade **(b)** track **(c)** double bang

Figure 2.2.: The three main event topologies.

$50\,\mathrm{m} \times E_\tau/\mathrm{PeV}$ [82] which means that below a few PeV the decay path is too short and the second shower will overlap with the initial one. These events are therefore identified as, and treated identically to, NC- or CC-ν_e-induced cascade events, except in specialized analyses which try to detect them [84, 85]. At higher energies, however, the tau track will extend enough to produce two distinctive showers. IceCube calls this event topology a **double bang**, see fig. 2.2c.

2.1.3. Cherenkov detectors

Cherenkov neutrino telescopes detect the charged particles described until now by means of Cherenkov light. This is a geometrical effect where the radiation of a medium excited by relativistic charged particle passing through it interferes constructively. The condition for this, as sketched in fig. 2.3, is

$$\frac{1}{\cos\theta} = \beta n \tag{2.2}$$

where θ is the angle of emission, β the particle speed relative to c and n the medium's refractive index. Given a particle with charge number z, the number of photons N is distributed over wavelength λ and trajectory length x according

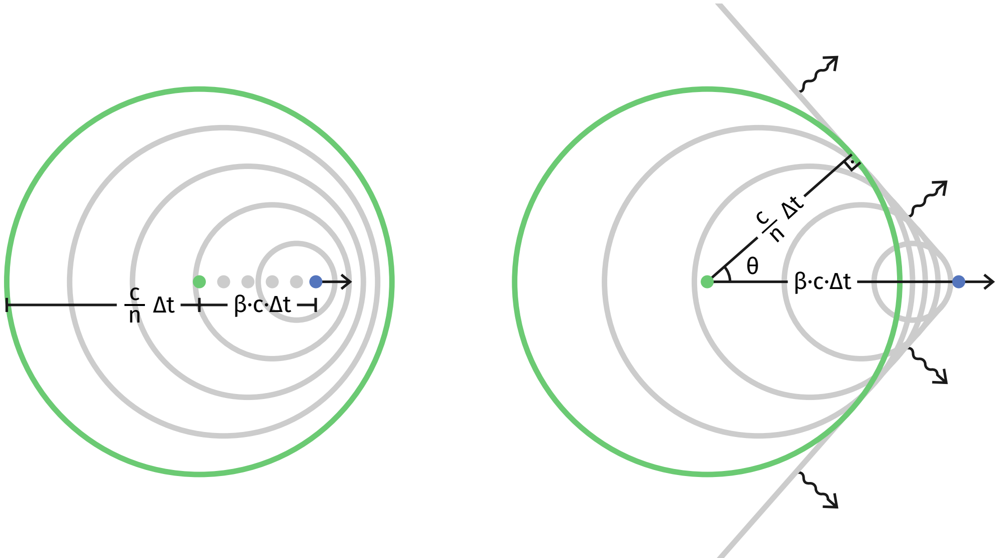

Figure 2.3.: The geometric condition for the Cherenkov effect. A particle with speed β is shown at one moment (green) as well as Δt later (blue). The emission from the excited medium of refractive index n propagates out from the particle location and, at the later moment, reaches the surrounding circle. In the second case where $\beta n > 1$, constructive interference along the Cherenkov cone is shown as arrows. Figure adapted from [86].

to the Frank-Tamm formula [87]:

$$\frac{\mathrm{d}^2 N}{\mathrm{d}x\,\mathrm{d}\lambda} = \frac{2\pi\alpha z^2}{\lambda^2}\left[1 - (\beta n_\lambda)^{-2}\right] \tag{2.3}$$

To effectively use the Cherenkov effect, the medium of the detector therefore needs a refractive index sufficiently larger than 1. It also needs to be transparent in order for the light to propagate to where it can be detected. Water is a prime candidate for this medium, due to its great natural abundance in liquid or solid form. To overcome the small neutrino cross-sections (see section 2.1.1), large instrumented detector volumes in the order of $1\,\mathrm{km}^3$ are required to detect the low expected fluxes of astrophysical neutrinos as first proposed by Markov in 1960 [88]. Cherenkov detectors therefore operate by collecting the Cherenkov light in a three-dimensional array of photomultipliers installed deep in a natural medium such as ice or water. The number of collected photons and their arrival times allow for the reconstruction of the neutrino energy and, in some cases, its direction.

2.1.4. Backgrounds

In addition to neutrinos of astrophysical origin, Cherenkov telescopes will also detect the atmospheric muons and neutrinos from cosmic-ray air showers (see section 1.1.2).

Atmospheric leptons are a product of cosmic-ray interactions with the upper shell of Earth's atmosphere. These interactions generate a flux of unstable mesons, such as pions and kaons, which in turn decay (see fig. 1.2), producing a flux of leptons aligned with the cosmic ray's incident direction. As such their angular distribution is decorrelated from the cosmic ray accelerators, as discussed in section 1.1.3, and so they form a background to the extraterrestrial neutrinos which neutrino telescopes seek to observe from the same astrophysical sources (see section 1.6).

Atmospheric neutrinos in particular are almost an irreducible source of background since the signature they produce in a Cherenkov telescope is indistinguishable from an astrophysical neutrino. Furthermore they point back to every region of the sky since they can originate both in the atmosphere above the telescope, as well as on the other side of the Earth.

However their energy spectrum is steeper than that usually expected for an extra-terrestrial origin. This is because for high energies ($\gtrsim 10\,\text{GeV}$), the decay length of the relativistic pions and kaons exceeds their interaction length. Higher-energy mesons are then increasingly likely to lose energy in interactions before they can decay and pass their energy on to neutrinos [89]. As a consequence, the atmospheric neutrino spectrum is one power steeper than that of primary cosmic rays ($\gamma \approx 2.7$ below the knee) [90].

The spectrum of atmospheric muons is closely related to that of atmospheric neutrinos, although at low energies energy losses and muon decay have to be taken into account. Muons are the most numerous charged particles at sea level and given their long decay length can reach underground detection facilities. One exception are those muons produced below the horizon, which are shielded by the Earth at all energies. An efficient way to reject the atmospheric muon background is therefore to select only track-like events with an upwards direction. Another strategy is to select tracks starting inside the instrumented volume. In both cases, the event must have been caused by a muon neutrino, passing through the Earth without being stopped, or though part of the detector volume without provoking Cherenkov light.

One final contribution from cosmic-ray air showers cascades are the so-called prompt fluxes. Prompt muons and neutrinos are the decay products of charmed hadrons. Charmed hadrons have far shorter lifetimes and can therefore dominate the lepton production at energies of 10 PeV despite their lower production cross-section [90]. The prompt neutrino component therefore has a harder spectrum ($\sim E_\nu^{-2.7}$) than the conventional one, reflecting the parent cosmic-ray spectrum. Prompt neutrinos have not yet been identified [91, 60] and upper limits set by IceCube currently rule out one of the more optimistic models [92].

2.2. Construction and geometry

When installing the IceCube array in the ice, the main challenge was to melt the 86 holes with a depth of 2500 m and width of 60 cm. The process involved three main phases. In the first phase, a conical drill head, heated through circulating hot water, was lowered down from a drilling tower located above the hole site, and melted through the top 50 m layer made of gradually compacted snow called “firn”. Then, a second drill head delivered a jet of 88 °C water as it was lowered the rest of the way through the ice shelf, which unlike the firn does not absorb water. In the third phase, the drill head was pulled back up, reaming the hole to its target diameter.

The IceCube string was then deployed from the same drilling tower, before the hole refroze to less than the DOM diameter 24 h later. One string consists of 2500 m of cable, with four 100-pound weights at the bottom, in total 60 DOMs with a spacing of 17 m between depths of 1450 m to 2450 m. As illustrated in fig. 2.4, these are mechanically attached to the cable by a harness, and electrically connected via a cable that breaks out of the string cable sheath.

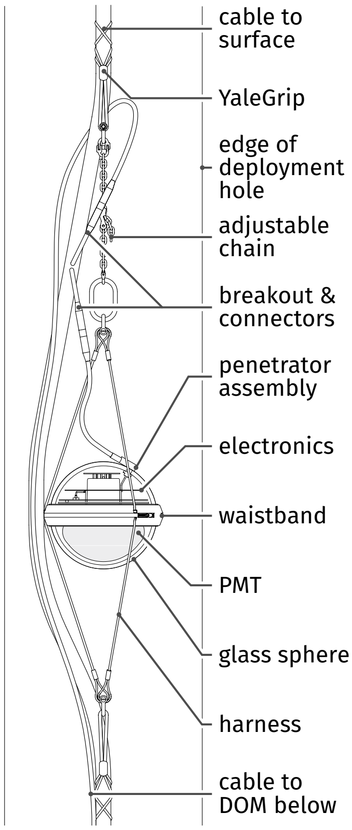

Figure 2.4.: A DOM as attached to its string in the ice. Adapted from [73].

The first string was deployed in the 2005 austral summer season. From there, IceCube's construction took a steadily increasing pace that peaked with 20 strings deployed between November 2009 and January 2010. Completion happened in January 2011 during the 7th austral summer of construction. Following each deployment period, the new cluster of strings was integrated into data-taking, typically in mid-May. The resulting head start to the development of data processing, physics analyses and detector studies was essential in light of the long construction timeline.

The resulting surface geometry (see fig. 2.5) is a hexagonal grid of in total 86 strings with an average 125 m spacing. In one of its corners, two intended strings could not be deployed due to debris from the previous South Pole station*. They were instead deployed closer to string 36 near the array centre, among 6 strings fully equipped with more efficient PMTs (see the following section 2.3). On these 8 strings of the DeepCore sub-detector, most DOMs are concentrated below the dust layer† in an exceptionally clear portion of the ice (see section 2.7). The smaller vertical and horizontal spacing of 7 m and 42 m to 72 m respectively make DeepCore sensitive to lower-energy neutrino interactions which are smaller in size (see section 2.1.2), becoming more

*K. J. Meagher, personal communication
†volcanic ash at depths of 1204 m to 1243 m, see section 2.7.3

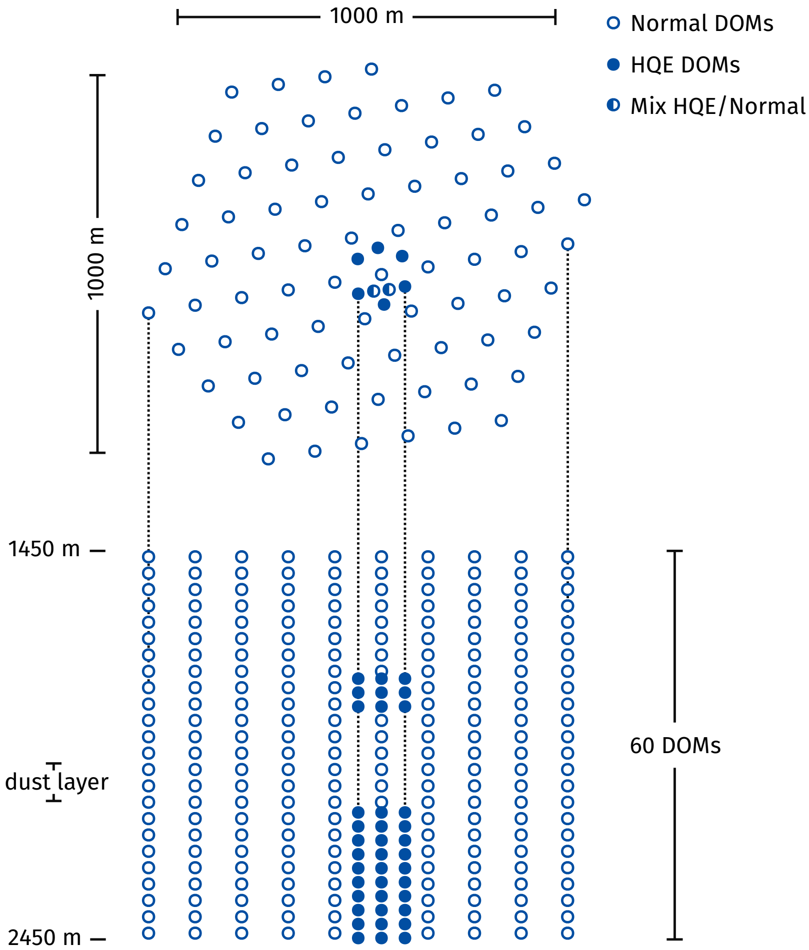

Figure 2.5.: An idealized view of the IceCube geometry. On the surface view, deviations from the hexagonal grid have been ignored. On the side view, the same is the case for deviations from horizontal planes. The IceCube strings show fewer DOMs than in reality, and the tightened spacing for DeepCore is not to scale. Adapted from [93].

sensitive than the rest of IceCube for $E_\nu \lesssim 100\,\mathrm{GeV}$ [93].

2.3. The DOM

An IceCube DOM is contained by a glass sphere of 13 inch diameter and 0.5 inch thickness, able to withstand the hydrostatic pressure at depths down to 2.6 km as well as pressure spikes of 690 bar. It is assembled from two halves joined with an aluminium waistband by which it is also mounted to a steel cable harness (see fig. 2.4). Within, a PMT is optically coupled to the glass with a silicon gel and protected from the ambient South Pole magnetic field with a mu-metal cage.

The PMT is a 10 inch diameter Hamamatsu R7081-02. Its functional principle can be summarized as photons hitting a cathode, each with a 25% chance (called the **quantum efficiency**) [73] of freeing a photoelectron (PE) via the photoelectric effect. The PE is then amplified by a factor of 10^7 in a series of dynodes, and ultimately collected as a **pulse** of an amplitude from 1 mV to over 2 V.

The total photocathode area of IceCube is limited by the cost of constructing DOMs and drilling holes (see section 2.2). It allows IceCube to detect approximately 70 photons per TeV energy deposited in the ice [94], with a timing accuracy of $\mathcal{O}(\mathrm{ns})$.

The DOM has a degree of autonomy as a data-taking unit, thanks to its onboard electronics which are spread over five circuit boards.

When the PMT produces a pulse, the signal passes a number of these components:

1. A toroid transformer couples it to the main board.
2. A discriminator detects waveforms exceeding a charge equivalent to 0.25 PE (one quarter of the pulse induced by a single photon hitting the PMT). This launches the readout of the DOM and signals neighbouring DOMs of that fact.
3. A coincidence unit determines whether any of the nearest and next-to-nearest neighbouring DOMs on the same string launched their readout within $\pm 1\,\mu\mathrm{s}$. It indicates that the signal in this hard local coincidence (HLC) hit is due to **bright light** produced by a source closer to the DOM, with $\gg 1$ photons/DOM, and undergoing less scattering. The full detail

of the waveform is relevant for the physical event, and so the output of all available digitizers is saved. Without a local coincidence, the signal can belong to three additional classes. One is **dim light** from a source far in the detector, attenuated to below 1 photon/DOM and arriving at times strongly smeared by scattering; its intensity is still relevant for the physical event, but not the timing, and so the charge summary and time of threshold crossing are saved. The second are **afterpulses**, one of which can follow on average 6 µs after another pulse due to the ionized gas left between the dynodes. The third is **dark noise** without light entering the PMT at all, which is thus prevented thanks to the local coincidence distinction from overwhelming the data transmission with its rate of 560 Hz, and can still be partially separated from dim light during later data processing.

4. While this decision is being made, the signal passes through a delay line in order to arrive 75 ns later at the preamplifiers, and then the digitizers.

5. There are seven digitizers in the DOM which turn the analogue waveform into a series of digital samples.
   - The flash Analogue-Digital-Converter (fADC) which samples continuously, but with a rough time resolution. With no local coincidence, this is the only digitizer needed.
   - Six faster sampling digitizers which record the waveform at higher time resolution, but can only capture 427 ns, followed by a deadtime, so that a pair of them operates in tandem. There are three such sets receiving signal preamplified with a low, medium, and high gain respectively, to increase the DOM's dynamic range.

All relevant information is then combined in a **hit record** and cached for up to typically 1 s until it is transmitted by request of the DAQ on the surface. Communication to the surface is achieved via a twisted cable pair shared by two DOMs. These cables also supply them with 96 VDC.

Furthermore the DOM fulfils other functions, among others:

- High voltage power supply to the PMT.
- Timestamping hits with a built-in clock, whose latency and offset is calibrated from the surface.

- Self-calibration of the electronics.
- An array of high-intensity pulsed LEDs to help calibrate the detector geometry, ice properties, and energy reconstruction algorithms (see section 2.7.2).

2.4. Data processing

2.4.1. On-site hardware

The cables of each in-ice string are routed via a junction box, through trenches in the snow, to the ICL, a two-story building which sits on stilts near the centre of the array surface. Inside the ICL, each surface cable is split up and connected via a patch panel to a DOMHub, a rack-mounted computer with specialized components such as:

- Eight custom readout cards, to which the cables connect, over which the DOMHubs have direct access to the hit records taken by each DOM.
- Two 48V power supplies, which supply the DOMs with power over the same cables.
- A fanout card, which supplies the DOM clocks with the time received from a master clock, which itself is synchronized via a GPS antenna mounted on the roof of the ICL.

DOMHubs watch their string for trigger conditions (see section 2.4.2), compile monitoring quantities, and calibrate the signalling latency to each DOM. They connect by ethernet to the rest of the South Pole System (SPS), a computing cluster that fulfils tasks of detector control, data processing and data handling.

2.4.2. Data acquisition

The DOMHub queries each DOM once a second for hits. By calibrating the signalling latency to each DOM, the hit record time stamps are transferred onto a common time axis synchronized between all hubs/strings. This is called DAQ time. The hits are cached locally, and a summary forwarded a summary to the trigger system.

The trigger has the task of deciding when there was a physical event in the detector, upon which the entire detector should be read out. This happens in software. The basis of the trigger definitions are HLC hits in temporal and spatial coincidence, ignoring all other hits without local coincidence to their neighbouring DOMs, called soft local coincidence (SLC) hits (see section 2.3). This is in order to suppress the majority of dark noise. As some HLC hits due to dark noise remain, the multiplicity of this coincidence needs to be set high enough within a sliding time window to ensure triggering only on light in the detector. IceCube can use several patterns of spatial coincidence, but for the brightest events, the detector can be triggered by counting the number of hits within an entire sub-detector e.g. IceCube or DeepCore. For IceCube (without DeepCore), this Simple Multiplicity Trigger (SMT) is configured to consider 5 μs and require 8 HLC hits, giving it the name SMT8.

The trigger time of the SMT is defined as the beginning of the sliding time window. For SMT8, all hits from the in-ice DOMs are read out 4 μs before and 11 μs after the trigger time (6 μs past the end of the sliding time window) while the IceTop DOMs are read out within ±10 μs of it, as sketched in fig. 2.6. These time windows are based on generous estimates for the duration of an event caused by a relativistic particle moving through the detector. An event builder bundles all the hits from overlapping readout windows to an event and saves them to disk.

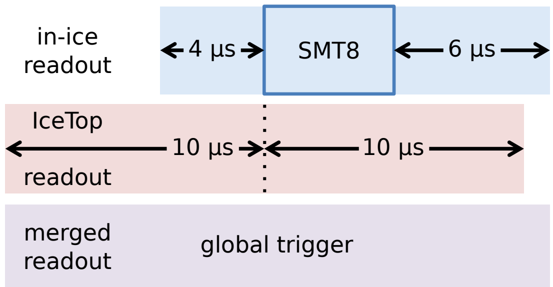

Figure 2.6.: Readout windows given an SMT8 trigger, adapted from [73].

A Command and Control (CnC) system launches the detector's data acquisition according to a configuration file that specifies among other things the status of the SPS, the definitions of the triggers, and which DOMs will be involved. The latter most importantly allows to temporarily exclude individual malfunctioning DOMs in order to avoid losing uptime until they return to normal operation. The configuration and detector status are prepended to the data files, so that further processing (see the following sections 2.4.3 and 2.4.4) can distinguish DOMs that received no light from those that were not operational. CnC monitors the data-taking operation, and if it ends prematurely will record the time up to which the detector was stable and the data can be used nevertheless. Otherwise, 8 h later the "run number" associated with the data will change, which can happen once more for a total of 24 h i.e. 3 runs until data-taking is completely restarted, including the option to adapt the configuration.

2.4.3. Online processing and filtering

The DAQ produces around 1 TB of raw data per day. Around 20 servers of the SPS compact this data so that it can be archived long-term (170 GB/day), and select a subset of 90 GB/day which is transmitted to the North via satellite.

Processing the digitized DOM waveforms starts by converting them to physical charge units. The necessary calibration constants for each DOM, each digitizer within that DOM, and within the ATWDs even each channel, are measured at the beginning of each run as they depend among other things on temperature.

Each waveform with no saturation can be represented as a linear combination of pulses from photons arriving at different times. The times and amplitudes of these pulses are extracted by a deconvolution algorithm, significantly reducing the storage space. The list of times and amplitudes is the basis for the long-term archive of all triggered events, which begins as a crate of hard disk drives shipped from the South Pole each season.

Once waveforms are deconvolved, a series of algorithms can handle the entire event thanks to this compacted representation. These algorithms first take care of separating pulses related to different coincident events, or dark noise. The simplest of these exploits the existing local coincidence between (next-to-)neighbouring DOMs (see section 2.3) and a general time correlation. Because Cherenkov light is produced exclusively by relativistic particles, the size of IceCube limits the duration of an event. Consequently, the set of HLC hit pulses within a 6μs time window will have a significantly reduced dark noise contamination, while also rejecting some desired pulses ($\gtrsim 30\%$ for tracks). Therefore this heuristic is only complementary to more sophisticated ones, which are described in section 3.1.

Variables that characterize the event can also be computed. This includes reconstructions of the interacting neutrino's kinematic parameters. Using all these techniques (which are described further in section 3.2), events are selected and divided into data streams called filters. These generally focus on a particular physical case and apply complementary methods to reject background.

The reduced data volume from the filters can be transmitted to the North via a satellite link which is available during ∼12 h per sidereal day. There, they are further processed, as described shortly in the following section 2.4.4.

2.4.4. Offline processing and filtering

To fully exploit the scientific potential of IceCube data requires more processing, and therefore more computing capacity than can be practically hosted at the South Pole. But these additional tasks do not need to run concurrent with data-taking, or "online". Instead, IceCube uses "offline" systems installed on other continents, with the largest such system maintained by the Wisconsin IceCube Particle Astrophysics Centre[‡].

Offline processing proceeds in multiple levels, building on the set of triggered events called level 1.

- Level 2 processing is characterized by repeating the filter selections from the Pole, but using improved, more CPU-intensive reconstructions. Additional quality cuts, noise removal, and splitting of coincident events are also involved. This results in data that is as before sorted into (overlapping) filter streams, which are still background-dominated but with a lower rate so that more specialized processing is feasible.
- Level 3 selections are again a combination of processing and selection, but usually combine several level 2 filters as their input in order to collect events relevant to a common physics purpose in order to serve many analyses. There is e.g. a cascade level 3 and a muon level 3.
- From here on, individual analysis-level data samples are prepared, without necessarily a formal "level" nomenclature.

One such sample is the subject of chapter 3, along with the on- and offline processing that produces it. Therefore no specific levels 2, 3 or further will be described in this chapter.

[‡] `https://wipac.wisc.edu/`

2.5. Real-time alerts

2.5.1. Infrastructure

Although a complete discrimination against the backgrounds described in section 2.1.4 cannot been done in an event-by-event basis, individual events have a higher probability of originating from an astrophysical source. A small number of such events can already be singled out by the online processing and filtering at the South Pole, and handled by a dedicated real-time (low-latency) system. This composes **alerts** which describe the events based on the preliminary reconstructions and variables available at this point. Due to their small storage size and number, alerts can be transmitted via a constantly available, but low-bandwidth satellite system, instead of the less available high-throughput link used for filtered full event data (see sections 2.4.3 to 2.4.4). A possible message queue on this link adds a typical latency of 13 s (but always <300 s) before the event is received in the Northern hemisphere.

There, they are accepted by a group of real-time clients which run in parallel to:

- publish the alert in a different format;
- start a follow-up analysis;
- correlate it with a database of past alerts or source candidates;
- and automatically check the detector health around the alert time, in lieu of the manual monitoring checks applied to every run to certify its suitability for offline processing.

This higher-priority treatment helps pointing telescopes make follow-up observations of a potential transient neutrino source, a principle already discussed in section 1.8.2.

Another type of message that follows the same path is the **compact data record** of an alert-producing event. This contains mainly the extracted pulses, the minimal amount of data to let a real-time client in the North launch a follow-up reconstruction. Within a few hours, this computational effort improves the alert's angular resolution to typically $\leq 1°$ and provides a better estimate of its uncertainty, including e.g. ice model systematics (see section 2.7).

These are again useful for the vast majority of telescopes whose field of view does not cover the original alert error circle. In addition, more elaborate study of the event might either confirm its assumed character as a signal candidate, or reveal it to be background or a detector artefact. If the alert was published, then so is the result of this follow-up, i.e. either the updated direction and uncertainty, or a retraction.
As of 2019, IceCube has four real-time alert streams belonging to two types.

2.5.2. Singlet alerts for the greater astronomical community

The **Gold** and **Bronze** singlet alert streams both originate from a set of three filters:

1. Gamma-Ray Follow-Up (GFU) candidates (see section 2.5.3), a sample of well-reconstructed throughgoing track events similar in purity to an analysis-level data set, selected by a boosted decision tree (BDT) [76].
2. High Energy Starting Events (HESE), which uses the outer layers of the detector as a veto in order to strongly reject muons. Among these mostly cascade events left by the filter, an additional cut selects well-reconstructed tracks [95, 96].
3. Extremely High Energy (EHE) are the brightest events above a threshold of total charge measured in the detector. This filter as well is followed by a quality cut, leaving mostly throughgoing tracks similar to the GFU [97].

The reconstructed direction (RA, δ) (see section 3.3.2) and energy proxy E of these events are used to calculate the **signalness** which is defined as the fraction of expected signal events in the event's respective region of the parameter space:

$$\text{signalness}(E,\delta) = \frac{N_{\text{signal}}(E,\delta)}{N_{\text{signal}}(E,\delta) + N_{\text{background}}(E,\delta)} \tag{2.4}$$

where the numbers N are the expectation values based on simulation (see section 2.6). Energy has an important role since it can distinguish the $\propto E^{-2.19}$ signal spectrum from the softer atmospheric backgrounds (see section 2.1.4). Since signal and background are uniformly distributed with respect to right ascension this parameter is ignored. However they do change with declination, particularly rapidly in the case of the background, making the dependency on this parameter crucial.

Even though the three filters use different reconstructions, the signalness allows the comparison of their events. Events with a signalness above 50% pass into the **Gold** alert stream, and those above 30% into **Bronze** (unless only selected by the EHE filter). After typically ≈20 s, for each such event one single **alert** message is composed. It is based on the best reconstructions available for this event and consists of time, direction, angular uncertainty, signalness and false alert rate, as well as the most likely neutrino energy given the signal spectrum as a prior.

In total, this results in an average of ∼ 10 Gold alerts per year, of which ∼ 52% are expected to have an astrophysical origin, in addition to ∼ 20 Bronze alerts containing ∼ 16% of astrophysical events [75]. The responsible real-time clients in the North disseminate these alerts to the astronomical community via AMON and GCN (see section 1.8.2).

2.5.3. Multiplet alerts for the gamma-ray / optical / X-ray follow-up programmes

Gold and Bronze alerts simply exploit the overall distribution of the background and signal in declination and especially energy in order to find a small number of high-energy events which reach the required significance for real-time follow-up on a possible transient source. A complementary method additionally makes the distinction that background events are not correlated in time and space, while signal events from a transient source are. Therefore, the required significance can also be attained by a **multiplet** of several events from a larger sample. This has long been the principle of offline point source searches (see e.g. chapter 7) but IceCube also applies it in real time.

Starting from events selected by the Muon filter, two stages of additional online filtering and processing result in a $\mathcal{O}$(mHz) stream of events that is similarly sensitive to point sources as the traditional offline event selections, thanks to similar purity and similarly advanced reconstructions, but compromising on the angular uncertainty estimator for some events. The selection uses the score from a BDT, after which they are transmitted to the North as alerts followed by compact data records, as described in section 2.5.2. There they are combined into multiplets by both the optical follow-up (OFU) and gamma-ray follow-up (GFU) programmes described in the following. The latter only uses alerts passing a tightened BDT cut, which are also those entering into the previously described selection of high-energy singlets.

IceCube carries out its OFU programme with the help of a real-time follow-up client which is triggered by incoming alerts of OFU stream events. It clusters events of the last 100 s which are closer than 3.5°. All multiplets of 3 or more events, as well as doublets satisfying an additional tightness criterion, lead to an alert of the multiplet time and direction sent immediately to the observing partners, which at this time are Palomar, Swift-XRT, and Master. These telescopes will look for a possible transient event as the origin of the neutrino multiplet.

The GFU programme meanwhile uses a more sophisticated maximum likelihood method to find and characterize multiplets. Its time clustering test statistic compares directional information as well as reconstructed energies to a source location and spectrum, while considering only GFU events within the most recent time window of ΔT. With each alert of the GFU event stream that arrives in the Northern real-time system, the analysis is updated for a range of ΔT up to ≈ 0.5 year. This clustering time is chosen to encompass multiplets like that found near TXS 0506+056 (see section 8.1.3).

Part of the GFU programme is an all-sky time-clustering search, analogous to the OFU programme. In this case, the targeted source location is the nearest point on a HealPix grid [98] that evenly divides the sky with 0.9° spacing. The location and parameters of a candidate multiplet are iteratively optimized for the maximum test statistic (TS). A threshold is set on the latter leading to on average one public alert per year due to background [99, ch. 6] sent via Amon/GCN.

In addition, the GFU also focuses on a catalogue of gamma-ray sources selected for the three imaging atmospheric Cherenkov telescopes (IACTs) Magic, Veritas, and H.E.S.S. which are sensitive to VHE gamma rays. The total $\mathcal{O}(400)$ sources do not cover the sky in neutrino observation to the extent optical or X-ray sources would, making this mode of operation unique to GFU. In contrast to the all-sky search, the location of a candidate multiplet is fixed to the source coordinates. For each applicable IACT, when the maximum test statistic surpasses a respective significance of 3σ the telescope is alerted via an automated, private email. These $\mathcal{O}(10)$ alerts per year [99, ch. 6] (the actual rate depends on the catalogue) consist of the source direction, multiplet time, and significance or false alert rate for the source. Their rate is chosen somewhat higher than in the all-sky clustering to still achieve one follow-up observation a year despite the ~ 10% duty cycle of any IACT.

2.6. Simulation

Event simulation is to leverage knowledge about the underlying physics and detector into higher-level predictions concerning event selections, reconstructions, and analyses, without making use of a calibration signal. Simulation uses a chain of software that replicates the stages of a physical event one after the other: from the primary cosmic particle, to the point where software processing and filtering is identical to that of experimental data.

The chain to simulate atmospheric background events commonly starts with an air shower simulated by Corsika [100]. This was used to validate the event selections of chapter 3, but plays no further role in the context of this dissertation, where the background can instead be sufficiently characterized through its abundance in the experimental data itself. The same is however not the case for signal neutrino events.

A high-energy (anti-)neutrino begins as a randomly generated momentum and starting point on the Earth's surface. From there it propagates through the Earth, taking into account all standard model processes but replacing absorption (such as in CC interactions) with a statistical weight to reflect the survival probability rather than needing to restart the propagation. Once the neutrino has reached the predefined detector volume, it is forced to interact, resulting in a set of secondary particles originating from the interaction vertex (see section 2.1.1) each with their own randomly generated momentum. The interaction probability, highly suppressed by the cross-section, adds another factor to the event weight, as do the initial generation and the detector volume. Thanks to this combined weight, the artificially distributed neutrino events can be translated into a physical flux. Responsible up to this point is the IceCube-internal software NuGen based on the public Anis [101].

The secondary particles are then propagated from the vertex through the detector volume. In case of a muon neutrino CC interaction, the software Muon Monte Carlo (MMC) [102] is responsible for the muon trajectory. It also simulates a pattern of the energy lost by the muon and the light emitted by the medium. MMC is now succeeded by Proposal [103, 104].

This light is then propagated through the ice as individual photons, which can be absorbed or scattered according to the simulated ice model (see section 2.7). Some photons will meet a DOM according to the simulated geometry, and these enter the next stage of the simulation. Two compatible pieces of software can handle this ray-tracing, PPC [105] or clsim [106]. Alternatively, the light propagation can be performed independently for a tabulated set of idealized in-ice light sources. After accumulating enough statistics, these tables can translate every event hypothesis into the expected number of photons per arrival time per DOM; numerically sampling these photon statistics has an equivalent outcome to the full light propagation. The software that implements these **photon tables** is called Photospline [107].

Once a photon is determined to be absorbed by a DOM given the angular acceptance and quantum efficiency of the PMT, the signal path of section 2.3 is simulated in detail: from the PMT, to the toroid, to the amplifiers, to the digitizers. Afterpulses are also added, as is uncorrelated dark noise from a Poissonian process with a DOM-dependent rate. More detailed models which also reproduce the observed bursts of correlated noise are used for DeepCore simulation [108, ch. 5]. With all the available pulses, hit records are created, also taking into account the local coincidence between neighbouring DOMs. These then trigger a simulated readout of the detector according to the same criteria as in real data, and software processing proceeds in an identical fashion.

2.7. Ice properties

2.7.1. The Antarctic ice shelf

The ice shelf near the geographic South Pole formed during the past 120 kyr [109] by microscopic ice crystals slowly accumulating on the surface as snow [110], gradually filling in the valleys in the bedrock. Small air bubbles remain trapped within the snow as it compacts under its own weight, first into porous firn, and then solid ice. With increasing depth and hence pressure, these bubbles compress, but continue to scatter light, until they finally dissolve into the local ice, converting it into a clathrate hydrate which has almost the same refractive index as ice [109]. Below 1400 m [111], this transition is complete for all bubbles, which makes the ice shelf the most transparent natural solid on Earth for UV to visible light [109]. This medium in its natural abundance makes a neutrino telescope on the scale of IceCube feasible, and in fact was found to permit a high-energy extension with wider sensor spacing [112].

2.7.2. Optical calibration data

The optical properties of the ice are studied with the following methods:

- Extracting ice cores, which can be studied in great detail but with relatively high effort. At the South Pole only a depth of 1751 m has been reached [110].
- Deploying special-purpose laser dust loggers (where a laser beam is scattered back by the surrounding dust into an optically isolated PMT) during the drilling process. This method covers the full depth but was only applied to 8 of the drill holes [109]. The advantage of this method is a vertical resolution $\mathcal{O}(\mathrm{mm})$.
- Illuminating the detector with the DOM's own 400 nm wavelength calibration LEDs, called **flashers**. Their pulses, down to a FWHM of 6 ns, are practically point-like light sources in time and space. In a special data-taking mode, the shape of the pulse is recorded by the DOM's own digitizers, and can therefore be correlated through the DAQ time with the PMT waveforms of other DOMs to which the light propagates. On its path, the light's amplitude, angular shape and residual time distribution collect information on the optical ice properties. In addition, the orientation of the DOM which is a priori unknown, can be determined this way. Current general-purpose flasher data sets flash on 85 strings. This method directly measures the light propagation across the entire detector volume, and can be repeated whenever necessary with little additional effort.

Since light can enter the array also from above and below, the calibration of IceCube's predecessor Amanda [113] and dust logger data help extend or extrapolate measurements to shallower depths. The EPICA Dronning Maud Land deep ice core [114] does the same for deeper ice [77].

2.7.3. Ice model features

The ice model is fit to this data by likelihood [115] and unfolding [77, sec. 5.2.5.1] methods. Its primary features are [77, sec. 5.2.5]:

Scattering and absorption lengths: The scattering length b^{-1} is the mean free path that light travels before scattering and has an explicit dependence on depth due to the amounts of impurities, colloquially called **dust**, included in different layers which form a record of the conditions at the time they were deposited. At 400 nm this dust also dominates the absorption length a^{-1} which is the mean free path of light absorption, as can be seen in fig. 2.7.

Both these lengths are modelled in layers of 10 m. For each layer, the scattering and absorption lengths are determined from flasher data at the wavelength of 400 nm. As seen in fig. 2.7, for the deepest part of IceCube below 2100 m the absorption length never goes under 70 m, nor the scattering length under 20 m.

Wavelength- and temperature dependence: IceCube detects Cherenkov light primarily with wavelengths 300 nm to 600 nm. Since the main flashers can only probe scattering and absorption at 400 nm, 16 DOMs have LEDs which emit at different wavelengths of 340 nm to 500 nm [73]. A global parameterization is fit to this data, which scales the scattering and dust absorption coefficients from their 400 nm values with a respective global power law. The ice-intrinsic absorption becomes more important towards infrared wavelengths according to an exponential function [83] and dominates above 500 nm [113]. The latter also weakly depends on temperature, for which the model adopts a coefficient previously determined by AMANDA [116] as a fixed parameter in the likelihood. The necessary temperature profile is based on measurements in depths of 800 m to 2345 m [117] which largely overlap with IceCube and can be extrapolated down to the bedrock.

In summary, the absorption and scattering coefficients a and b are modelled to depend on depth z, wavelength λ and temperature T as follows:

$$a = a_{\text{dust}}(z; 400\,\text{nm}) \cdot \left(\frac{400\,\text{nm}}{\lambda}\right)^{\kappa} + A_{\text{ice}} \mathrm{e}^{-\lambda_0/\lambda} \cdot (1 + 0.01(T - T_0)) \tag{2.5}$$

$$b = b_{\text{dust}}(z; 400\,\text{nm}) \cdot \left(\frac{400\,\text{nm}}{\lambda}\right)^{\alpha} \tag{2.6}$$

where T_0 is an arbitrary reference temperature, and the other parameters are determined in the ice model fit.

Scattering function: As photons diffuse through the detector to arrive at a DOM, the scattering function affects both their spatial pattern as well as the

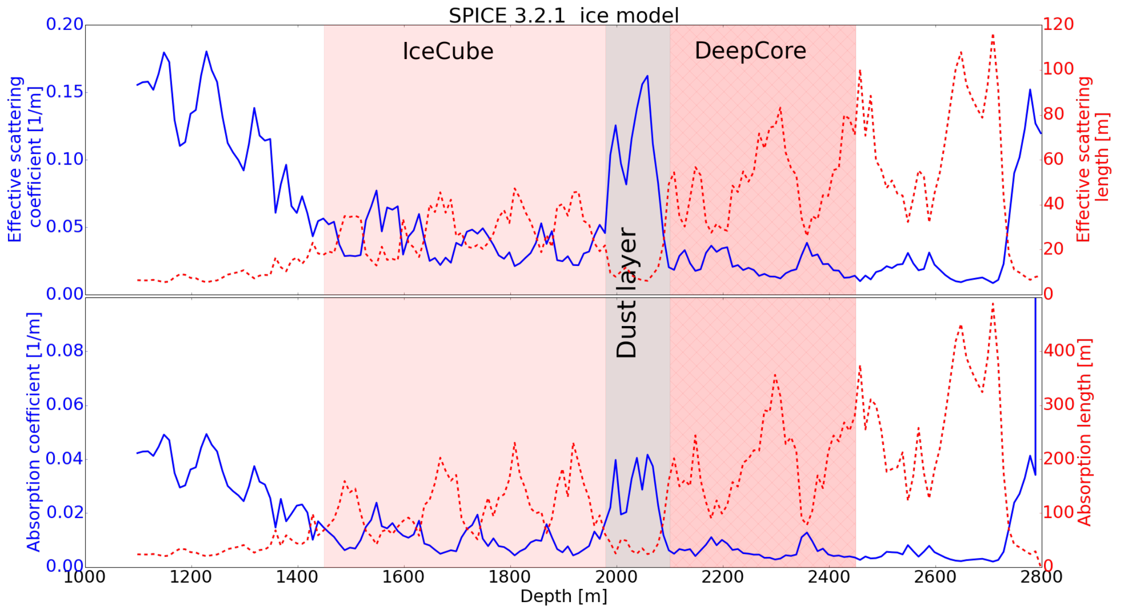

Figure 2.7.: Scattering and absorption coefficients and lengths vs. depth [77].

arrival time distribution. Shape parameters of the scattering function are fit in the flasher data likelihood, thereby avoiding related biases when analysing the flasher data as hitherto described.

Layer topography: Although the surface of Antarctica is flat, at the depths of IceCube the ice layers formed during the same epoch still follow the topography of the bedrock with variations of 20 m to 52 m in elevation, as shown in fig. 2.8. Since the flashers can not simulate light sources in between DOM layers, this could only be accessed by a dedicated **dust logger**. As it is deployed down a freshly-drilled hole, mm-size layers are illuminated with a laser beam, and their dust content scatters it back into a PMT [109]. The layer topography in the ice model is defined by fitting a depth-dependent parameterization to the dust logger data matched between 8 holes [109], as shown in fig. 2.8. This takes advantage of the expected smoothness and vertical continuity, and is extrapolated to regions outside of the octagonal prism enveloped by the dust-logged holes.

The high-resolution stratigraphy also allows to attribute thin, opaque layers to ash deposited by volcanic eruptions. For IceCube, the most pertinent of these are eruptions in the Southern Andes 24.5 kyr and 26.1 kyr ago, which are consistent

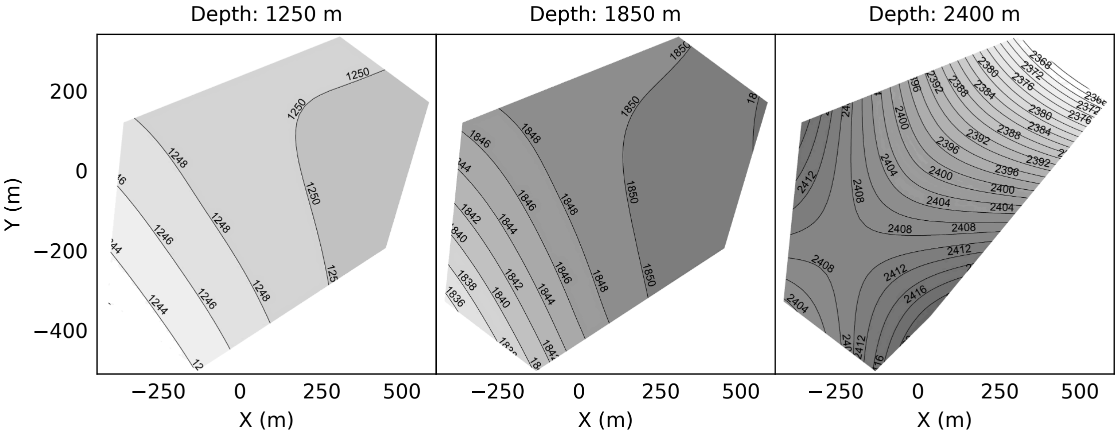

Figure 2.8.: The ice layer topography modelled on dust layer data. Shown as isochrons which intersect different depths, adapted from [109].

with such layers at depths 1204 m to 1243 m [109]. The light absorption in this **dust layer** is visible both in the flasher data (see fig. 2.7), as well as physics data acquisition.

Anisotropy: Given isotropic light emission, flasher data show that for light propagating through the horizontal plane from one string to its nearest neighbour, half as much light reaches a DOM on the axis of tilt as on the perpendicular axis [118]. This anisotropy is modelled with a modified scattering function independent of depth, following the principle of its first implementation [119]. This leads to a reinterpretation of the observed scattering in terms of the effective coefficient b_e, rather than the purely geometric length between scattering sites b^{-1}.

The underlying physical effect is not fully understood. More recently, birefringence on ice mono-crystals deformed by the glacial flow has been explored [118], instead of the impurities which normally dominate scattering.

Angular acceptance: Each DOM is embedded in the refrozen ice of a drill hole which presents a different optical medium than the undisturbed ice shelf that surrounds it. A remote-controlled system of two cameras at the bottom of one string [73] found a cloudy column of 16 cm diameter. This is consistent with the expectation that, as the freshly drilled hole refroze from the outside in, the advancing ice pushed air bubbles and impurities ahead of itself into the centre of the hole. This column scatters away a portion of the light that would meet the DOM vertically from below, obscuring part of the downward-facing PMT where it would otherwise present its maximum photocathode area. It can also scatter light from above the DOM "equator" into the PMT. A possible bubble column, along with the rest of the hole ice and the shadow of the string cable all change the angular acceptance of the DOM from its intrinsic laboratory-measured curve [120], that is the relative efficiency to detect light coming from different angles enclosed with the vertical axis, often considered in a far field approximation where the light source is broad and distant. Although there are attempts to access the hole ice properties via dedicated measurements where a DOM is illuminated by its own flashers [121], they are currently not commonly used as an explicit part of the ice model. Instead, the splined angular acceptance curve due to all effects is treated as a nuisance parameter individually for each DOM and determined by an unfolding of the general-purpose flasher data [77, sec. 9.2.3].

2.7.4. Ice-related systematics

Knowledge of the ice properties is extremely important, both for simulation (see section 2.6) and reconstruction (see section 3.2), and among the dominant sources of systematic uncertainty for IceCube analyses [122].

The magnitude of this uncertainty in the result of an analysis can be estimated with the help of dedicated simulation data sets. Discrete variations in the ice model parameters are propagated to the result by simulating events under this updated assumption, replacing the baseline simulation data used in the entire analysis chain (such as chapter 7). This approach for instance leads to uncertainties of ±5.3% on a point source flux sensitivity [122].

To track progress in the ice model development, IceCube has defined a **model error** which measures a model's quality independent of how its parameters were defined. Given one DOM being flashed by another, it divides the measured charge by the charge predicted through simulations of the ice model. Taking many such transmitter-receiver pairs results in a distribution of charge ratios. The latter's r.m.s. on a logarithmic scale defines the model error [77, sec. 5.2.7].

The ice model has evolved through numerous iterations. What has been described in this section is an ice model released in 2016 with a model error of 10% [77, sec. 5.2.8] currently considered a standard in IceCube. In the analyses of this dissertation, simulation and data processing used two previous versions of the ice model for historical reasons; one from 2010 with a model error 30% [83] in older samples, and one from 2012 with a model error 20% [123] in the more recent ones.

“Don’t put your Condor logs in /data/user.”
IceCube proverb

3
Neutrino Data

The neutrino data used in this thesis comprises a set of IceCube events starting in 2008. The exact data set is the same as used for the observation of TXS 0506+056 in [124], except for a period of 122 days at the beginning with no overlap with the Fermi-LAT data-taking period. The use of a single, standard data set shared by other IceCube analyses benefits reproducability. It also makes it easier to directly compare the analyses of chapter 8 and chapter 9, as has been done in chapter 10. In this chapter, the IceCube data flow from the previous chapter 2 is further described and continued up to the specific input variables given to the likelihood in chapter 7.

Author’s contribution: One variable in part of the data uses a novel calibration method introduced by me, see section 3.4.2.

Data sample in IceCube are demarcated by changes in the detector configuration and geometry or by the processing software used. The first available data sample was taken during the construction phase of IceCube when the detector consisted only of 40 strings (IC40). With each additional set of strings software updates continued in parallel to provide new reconstruction methods and improved data selection (IC50 – IC79). Acquisition with the full 86 string configuration (IC86a) started three seasons later, while additional software development continued leading to another data sample spanning the following three years (IC86b). Finally, a continuously updated sample was added which

Sample	Start	End	Livetime (days)	Events
IC40	2008-08-06	2009-09-20	263.80	25 353
IC59	2009-05-20	2010-05-31	353.58	107 011
IC79	2010-05-31	2011-05-13	316.05	93 133
IC86a	2011-05-13	2012-05-16	332.96	136 244
IC86b	2012-05-16	2015-05-18	1058.48	338 590
IC86c	2015-05-18	2017-11-01	886.33	509 756

Table 3.1.: Summary of data samples in the neutrino data set.

extends it through October 2017, in order to cover as much as possible of the multimessenger observations described in [125] and chapter 8 (IC86c).

The sum of data-taking periods which contribute to a sample is called its livetime. This differs from the period from the sample's start to its end by the time lost between data-taking runs (see section 2.4.2), and the runs excluded from the sample based on data quality. Systems at the ICL determine during data-taking when each run begins and ends, whether the detector was operating normally, and in what configuration, and record this information into a database (see section 2.4.2).

In order to perform neutrino astronomy on point sources, we prefer track events for their better suitability to reconstruct the incident neutrino direction (see section 2.1.2). The event selections (see section 3.5) therefore start with track-focused filters (see section 2.4.3). For each run, the event rate at the filter level is calculated. An annual modulation from atmospheric effects [126, 127] is visible, but shorter scale deviations are due to runs which used only a partial configuration of the detector or had other issues. Runs which deviate in rate by more than ±5% from the running median are removed from the final selection. The final clean uptime surpassed 90% in 2011 and reached 98% in 2015 [73].

Chapter 3 shows an overview of the data samples used with their start date, stop date, and total livetime. For a time-dependent analysis of the data, this information is not sufficient and the start- and stop- times of the individual runs are also required, which are hence also part of the data set. For the runs included in this data set, fig. 3.1 shows their event rates, as well as the fraction of clean uptime to which they contribute.

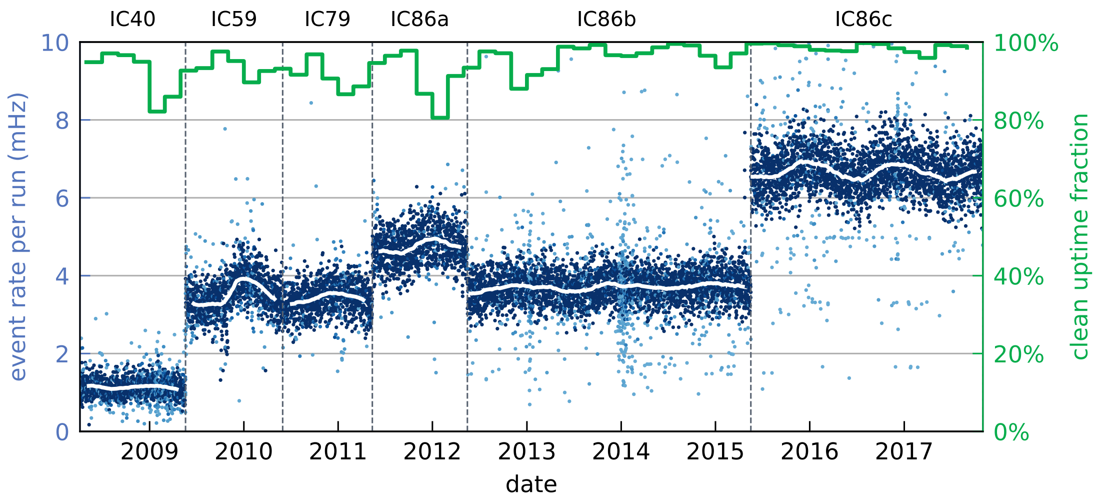

Figure 3.1.: Left axis, blue: all-sky event rates during runs in the data set. Runs shorter than 8 h vary more widely in rate and are de-emphasized through a lighter colour. At the transition between samples, shown as dashed vertical lines, the differences in detector geometry and event selection are apparent. Within some samples, seasonal variations of downgoing events are apparent in the 2-month running average of these rates, shown as a white curve. Right axis, green: the clean uptime which the data set comprises as a fraction of 2-month intervals. Comparing earlier to later data-taking seasons shows an improvement in detector stability and data quality.

3.1. Noise cleaning and event splitting

The pulses which make up the digitized DOM waveforms (see section 2.4.3) can have origins other than light in the ice from the interaction which triggered the event. For instance all events have dark noise which is uncorrelated to light in the ice, and “afterpulses” which happen on average 6 μs after the original pulse. Since the rate of incident atmospheric muons which can trigger the detector is 3 kHz and a common readout window lasts 15 μs [73, tab. 8], for $\mathcal{O}(5\%)$ of triggers in IC86, the DAQ captures the track from a coincident atmospheric muon after the original trigger, or before it if the first muon did not trigger the detector.

In order to analyse and reconstruct an event, separating pulses of different origins is very important. This happens through a combination of local coincidence and its extension to larger scale through general causality criteria, aided by preliminary track reconstructions. Various techniques exist specifically to clean away noise or to split coincident events, and some more general techniques achieve both. Different ones are used at different stages for the samples in this data set. This section describes the most common ones.

Noise hits are causally disconnected in time and space from the rest of the event, filling the whole detector volume and acquired time. This stands in contrast to signal hits, which cluster in time and space due to the causal connection to in-ice interactions. On this basis, **time window cleaning** finds the biggest cluster of HLC hits in a 6 μs time window. Another method called seeded RT (SRT) cleaning requires each hit to lie within $R = 150\,\mathrm{m}$ and $T = 1\,\mu\mathrm{s}$ of an HLC hit. The SLC hits which are added by this criterion help reject muons enter the detector.

Algorithms such as TopologicalSplitter cluster hits which are likely to be causally connected to each other through a cut in a cylindrical coordinate system and time difference. Once an event is split, reconstructions and event selections are repeated on each part. This helps to reject still more background since in coincident events, i.e. triggered events that actually contained more than a single event signature, the angular reconstruction often fails. This in turn misleads background rejection cuts which rely on reconstructed upgoing events never being atmospheric muons. Therefore, after splitting additional background can be rejected. Reconstructions are also generally improved through the splitting away of noise hits and after-pulses. The more advanced HiveSplit

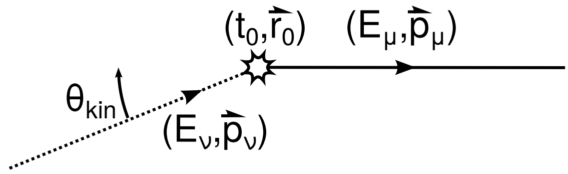

Figure 3.2.: The kinematics of a track event along with the parameters involved in its reconstruction.

algorithm [128] also searches for causal connections to a hypothetical track, and therefore enables the recovery of signal hits from a faint track which were lost e.g. to noise cleaning. For a downgoing track this might also identify an event as a likely atmospheric muon which enters the detector, rather than a ν_μ induced track that starts inside of it.

3.2. Reconstruction

Once the waveforms are unfolded into a series of pulses (equivalent to incident photons, see section 2.4.3) with an associated time and location, the event reconstructions try to find the emitting source for these photons, be it a muon track or an hadronic or electromagnetic cascade. The reconstructions described here explicitly or implicitly assume an idealized hypothesis on the pattern of light emission where the track is a line segment of uniform emission, and the cascade at its vertex is a point of isotropic emission. The final result is an albeit imperfect inference on the kinematic parameters of the particles in the detector, as shown in fig. 3.2. These would be for instance the neutrino interaction vertex and the momentum of the muon, which in turn will be correlated to the momentum of a possible interacting ν_μ.

Following are the reconstruction procedures most relevant to the data set, either in the final events or during the selection, generally focusing on muon track reconstruction. From sample to sample, their combination and configurations change.

In the IceCube data pipeline, more and more computing-intensive event reconstructions are applied as the events are selected into thinner and thinner data streams. At an early stage of the acquisition only very fast reconstructions can be used, especially for the filters which run at Pole.

3.2.1. Angular reconstruction methods

The fastest reconstruction for track directions is a line fit. This reconstruction minimizes the distance between the locations $\vec{r}_i$ of hit DOMs and the hypothesis of a straight trajectory at constant velocity,

$$\vec{r}(t) - \vec{r}_0 = \vec{v} \cdot (t_i - t_0), \tag{3.1}$$

where t_i is the time of the hit, i.e. the time of the first pulse for each DOM. This pulse ideally corresponds to the photon that has undergone the least scattering, subsequent photons detected by the same DOM would be worse-correlated with the track. The fit parameters are the velocity $\vec{v}$ and the vertex $\vec{r}_0$, through which the track passes at time t_0. The distance measure is a modified χ^2, which smoothly transitions from quadratic to linear at a distance of 153 m in order to suppress hits in regions which are not dominated by the track anymore. Like an unmodified χ^2 fit, this so-called Huber fit can be solved analytically and is statistically robust, without the requirement for a first guess [129]. It is therefore used as a seed for the reconstruction method which is applied next.

Most other reconstructions use a maximum likelihood method. These can reflect more detailed physical hypotheses, such as the pattern of light emission, the geometry of its propagation and its absorption and scattering in the ice. The latter introduces statistical deviations in the propagation direction of light, which diffuses the Cherenkov cone of a muon. This is also reflected in the delay with which photons arrive at a DOM versus the time expected from geometry, called the time residual $t_{\text{res}} = t_{\text{pulse}} - t_{\text{geo}}$, with

$$t_{\text{geo}} = t_0 + \frac{1}{c}\left(|\vec{p} \cdot (\vec{r}_i - \vec{r}_0)| + d \cdot \tan\theta_C\right). \tag{3.2}$$

The probability density function (PDF) $p(t_{\text{res}})$ describes the time of photons arriving at a DOM. In principle it depends on the specific DOM and the vector from it to the light source, but in practice can be approximated using fewer inputs.

IceCube commonly uses an approximation called the **Pandel** function. This is an analytical model which interpolates between an unscattered Cherenkov cone at short distances, to a diffusive model at longer distances, assuming a homogeneous medium. A Gaussian kernel convolution remains analytically computable and is used to extend the PDF to cover PMT jitter and left-over noise pulses [130].

The single photo-electron (SPE) likelihood also uses the time of the first pulse in each hit DOM as in the LineFit method, and is defined as

$$\mathcal{L}_{\mathrm{SPE}}(\vec{x}|\vec{a}) = \prod_i p(t_{\mathrm{res},i}|i,\vec{a}), \tag{3.3}$$

where i are the indices of hit DOMs and $t_{\mathrm{res},i}$ the time residual for the first pulse in DOM i. An estimate for the track parameters $\vec{a}$ results from a numerical maximization, starting at a track derived from LineFit.

This likelihood formulation is valid in the edge case where each DOM only detected one photon or none. However this is not true in general, particularly for brighter events. If a DOM i has detected n_i photons in total, the reconstruction can be improved by taking into account that they are all sampled from the same $p(t_{\mathrm{res}})$. This leads to the definition of the multi-photo-electron (MPE) likelihood as

$$\mathcal{L}_{\mathrm{MPE}}(\vec{x}|\vec{a}) = \prod_i n_i p(t_{\mathrm{res},i}|i,\vec{a}) \cdot \left(\int_{t_{\mathrm{res},i}}^{\infty} \mathrm{d}t \; p(t|i,\vec{a}) \right)^{n_i-1} \tag{3.4}$$

Maximizing this likelihood with respect to the track parameters results in the PandelMPE reconstruction.

Up to this point, using the Pandel function has assumed that the ice is homogeneous. Calibration measurements performed in the ice show that instead, the absorption and scattering depend on the depth and even angle of the incident light (see section 2.7). A full ray-tracing simulation of photons can take all of this into account, but needs to be cached for efficiency. The possible track hypotheses are represented by a grid of their parameters $\vec{a}$, which are collapsed down to a subset by exploiting the azimuthal and translational symmetries of the IceCube array (see section 2.2). For each such track, the simulated time residuals for each DOM are histogrammed in time. The resulting multi-dimensional histogram is stored as a spline which can be efficiently retrieved, interpolated for track parameters between grid points, and smoothly evaluated for times between histogram bins.

The combination of MPE likelihood and spline PDF yields the SplineMPE reconstruction. This provides the neutrino arrival angles for the samples IC79 onward. The samples IC40 and IC59 used unsplined histograms [131].

3.2.2. Energy reconstruction methods

In a ν_μ charged-current interaction, the neutrino energy is divided between a hadronic cascade at the vertex and the escaping muon. The inelasticity $1 - E_\mu/E_\nu$ is distributed around mean values of 0.5–0.3 in the energy range 10^2 GeV to 10^6 GeV [132]. The energies involved also mean that often, the muon enters IceCube from the outside, leaving behind the cascade and having lost an unknown fraction of its own energy. For these reasons, any reconstructed muon energy only has a purpose when applied to a set of events, for example to statistically separate a harder from a softer neutrino spectrum.

The fastest estimator for the energy deposited in the detector is the number of DOMs which had hits, a quantity often called N_{channel}. This estimator relies on an approximately even DOM density, and therefore the denser DeepCore infill is often disregarded, or weighted down. However this estimator loses relevance even for starting events, as they may leave the detector again. To be more accurate and less dependent on the event topology, a likelihood method is required. These make use of a closer approximation of the physical event hypothesis, which can be supplemented with an an existing angular reconstruction either to give a first guess for the direction of the momentum, or to fix it entirely and leave only the normalization as a free parameter.

The muon energy loss $\frac{\mathrm{d}E}{\mathrm{d}x}$ consists of both a constant minimally ionizing term, and stochastic energy losses in interactions with the medium (see section 2.1.2). As the latter scale with energy, they dominate for the events in the data set. By approximating the true energy loss pattern with a continuous one, the energy-normalized expected light yield Λ as a function of distance and angle can be described by a form of the Pandel function. Adding an expectation of noise hits ρ, this results in a Poissonian expectation for the number of photons n_i detected by DOM i. Combining these terms into a likelihood yields

$$\mathcal{L} = \prod_i \frac{\lambda_i^{n_i}}{n_i} \cdot \mathrm{e}^{-\lambda_j} \tag{3.5}$$

$$\text{where } \lambda_i = E \cdot \Lambda_i + \rho$$

This is again numerically maximized to solve for E to produce the `MuE` energy reconstruction. It assumes that the track emits light regularly, however this neglects the stochastic losses which are responsible for the $\frac{\mathrm{d}E}{\mathrm{d}x}$ behaviour it exploits. Each Poissonian distribution is further convoluted in λ with a kernel G in order to cover the irregularity of stochastic losses. This more correct version is called `MuEX`.

3.2.3. Angular uncertainty estimate

The angular reconstruction has an intrinsic uncertainty, which is unique to each event, where longer and brighter tracks are easier to reconstruct. Characterizing this uncertainty is crucial to neutrino point source searches, and several methods are employed across the samples:

Paraboloid sigma: In this method, the likelihood of the given reconstruction is evaluated on a grid of directions which surround the global best fit direction, while being maximized in all other parameters. This profile of the local log-likelihood maximum is fit to a paraboloid function, with major and minor axes σ_1 and σ_2. A combination of these axes is used as an angular uncertainty estimator for the reconstruction:

$$\sigma_{\text{para}} = 0.57\sqrt{\sigma_1\sigma_2}\left(\frac{\sigma_1}{\sigma_2} + \frac{\sigma_2}{\sigma_1}\right) \tag{3.6}$$

which scales with the radial median of the modelled point spread function (PSF) irrespective of $\frac{\sigma_1}{\sigma_2}$ [133]. This is the most popular estimator, and the sole one used up to IC86b.

There are two additional estimators available in IC86c:

Bootstrapping: A purely statistical method can also be used, which is independent of the reconstruction method and choice of numerical minimizer. In this method, the photons which serve as an input to the reconstruction are resampled to produce a new set of pulses with the same total charge as the original, representing a statistical variation of an event equivalent to the one analysed. The corresponding statistical variation from reconstructing many such simulated events is then used as the angular reconstruction uncertainty [131, sec. 3.1.3].

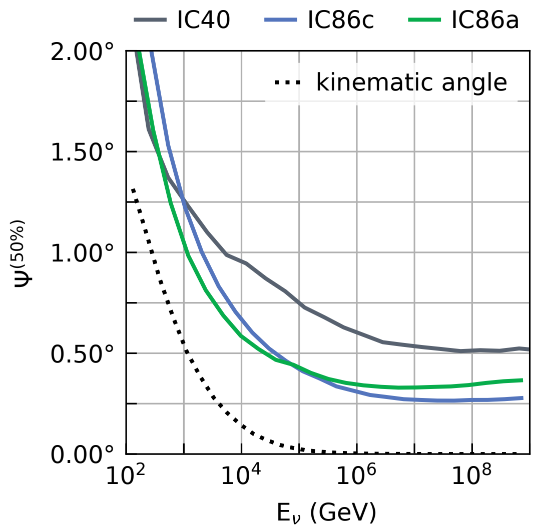

Figure 3.3.: The angular resolution of three samples in the data set used as examples: IC40 as the oldest, IC86a as the first using the full detector, and IC86c as the most recent. The median kinematic angle, here according to [136], contributes at lower energies.

These two methods require substantial CPU time proportional to the light yield in the event.

Cramér-Rao: An analytical method where the first and second derivatives of the likelihood are evaluated through knowledge of the PDFs. These then define the Fisher information matrix, which due to the Cramér-Rao bound sets limits on the covariance matrix of the maximum likelihood estimate [134, sec. 5.1] [135, sec. 4.2]. This less precise but fast estimator can be computed for all events in IC86c. This makes it available as a fallback for events in which the other two methods are not fast enough to keep up with the sample's online processing (see section 3.5).

3.2.4. Resolution

The resolution of the reconstruction methods is evaluated by comparing their results with the true parameters known in simulated data (see section 2.6). Several sources of randomness contribute to it. On one hand it is the unavoidable stochastic development of the event itself which affects the light emission of the involved particles for the same initial kinematic parameters, through effects such as the kinematic angle (see later in this subsection and fig. 3.2), stochastic energy losses, and the interaction vertex possibly at an unknown distance outside the array. But also the light emission is sparsely sampled by the array, so that limited photon statistics hinder the reconstruction as several patterns of light emission could have produced the same set of recorded waveforms. The severity of the latter can be judged based on properties of the event itself and of its reconstruction, so that ultimately the resolution can also be improved by sacrificing effective area to a tighter event selection.

Angular: Every reconstructed direction ideally approaches the momentum of the muon track, and the deviation therefrom represents the reconstruction precision. However due to the kinematics of a charged-current interaction, the muon itself deviates from the incident ν_μ by the **kinematic angle** θ_{kin}. Due to its energy dependence, it becomes a significant effect for $E_{\mathrm{true}} \lesssim 1\,\mathrm{TeV}$ as can be seen in fig. 3.3.

The reconstruction's angular resolution depends on the energy, as does the kinematic angle, and so both effects can be summarized as the median angle $\Psi^{(50\%)}(E_\nu)$ between the reconstructed and the incident ν_μ direction, as shown in fig. 3.3. When attempting to resolve a neutrino point source, the observed Ψ follows a distribution called the PSF, which then depends on the source spectrum and is usually approximated with a two-dimensional normal or Kent distribution.

Energy: Since E_{reco} is not a physical quantity, its absolute scale is irrelevant. All cases where there is a monotonous correspondence to E_{true} are equally ideal. Because of the reconstruction uncertainty, the dispersion distribution $P(E_{\mathrm{reco}}|E_{\mathrm{true}})$ deviates from such an ideal case, and is commonly summarized as the relative uncertainty

$$\frac{\sigma(E_{\mathrm{reco}})}{\langle E_{\mathrm{reco}}\rangle}, \tag{3.7}$$

which depends on the reconstruction method, track event selection, and true neutrino energy. The relative uncertainty of $\log_{10} E_{\mathrm{reco}}$ takes values in the range of 5%–25%. Using common reconstruction methods to infer the true energy of a throughgoing muon results in 30%–50% uncertainties on this quantity [137].

The previously mentioned resolutions help characterize how the reconstruction relates to true kinematic parameters. They could be used in an appropriate parameterization which, convoluted with a chosen neutrino flux, predicts the observable distributions. This is however not necessary thanks to a sufficiently large sample of Monte Carlo simulation (see section 2.6). The desired convolution is achieved through weighting the latter, without losing details to the choice of parameterization.

3.3. Coordinate systems

Reconstructed event directions in IceCube are provided in a given coordinate system. Space and time are at first defined in the local coordinate system of IceCube. They can then be perfectly transformed into coordinate systems which are in line with astronomical conventions, enabling multimessenger observations.

3.3.1. Time axes

IceCube synchronizes its time with the world by receiving GPS signals, and produces a local time for the DAQ which is compatible with Coordinated Universal Time (UTC) by manually adding leap seconds. Together with the Gregorian date this makes up an unambiguous time stamp for any IceCube event.

Astronomers are not concerned with the leap days of the Gregorian calendar. The Coordinated Universal Time is combined with the Julian day to make a real-valued time axis for all astronomical observations of human history. For easier use from the last century onward this is commonly zeroed at midnight on 1858-11-17, to define the Modified Julian Date (MJD). IceCube DAQ timestamps can be converted to MJD.

3.3.2. Spatial coordinates

The IceCube coordinate system has its origin near the detector centre of gravity, between strings 86 and 36 at a depth of $\approx$2000 m below the surface, or 1948.07 m below the survey point for string 21's deployment [73]. It is a right-handed Cartesian coordinate system where the z-axis points upwards, and the y-axis points along the 0-degree of longitude, i.e. towards the Royal Observatory in Greenwich. In spherical coordinates, the direction of a particle's origin, opposite to the direction of its momentum, is expressed by azimuth and zenith θ in the same system. An event's angular reconstruction is expressed in these two angles. Events with $\theta < \pi/2$ are (reconstructed) downgoing, and $\theta > \pi/2$ (reconstructed) upgoing. Celestial objects are assigned a position by virtually projecting a line from the centre of the Earth that meets them at a particular time, called the epoch, which everywhere in this thesis unless specified otherwise is J2000.0 i.e. 12:00 Terrestrial Time on 2000-01-01 [138, p. 23]. This direction is expressed in equatorial coordinates, in which the angle enclosed by the projected line and the equator is the declination δ, so that a point above the North Pole would be $\delta = +90°$ and a point above the South Pole at $\delta = -90°$. The degree of longitude crossed by the projection is the right ascension (R.A. or α). The latter is commonly written in units of hours where $24\text{h} \equiv 360°$.

When an event is reconstructed in local coordinates, these need to be transformed to the equatorial coordinates of its direction of origin at the instant of its detection. This calculation includes multiple effects that make up the evolution of the Earth's rotation. Excluded is the parallax motion of the axis due to the Earth's orbit around the Sun, since this depends on the source's distance; its magnitude is negligible for any source targeted by IceCube, with the notable exception of the Sun itself. Nowadays this calculation is achieved with an IceCube-internal wrapper around the positional astronomy library `PAL` [139].

3.4. Calibration

The observables in the data set are reconstructed direction, estimated angular uncertainty, reconstructed energy, and the time of the event. The latter is determined with $\mathcal{O}(\text{ns})$ accuracy (see section 2.4.2) and can be reliably used for analyses at all timescales common in astronomy without any further calibration. Furthermore, there is no need to calibrate the energy reconstruction to a true energy scale. Instead it can be used directly as an observable by the likelihood in chapter 7 since it is forward-folding, made possible by simulated data. The situation is different for the remaining observables, which need to be calibrated.

3.4.1. Moon shadow

A significant systematic offset of the reconstructed muon direction, for instance related to detector geometry or DAQ timing, would be fatal to neutrino astronomy. Such a bias could be shown with a known ν_μ or μ point source which is sufficiently bright and energetic. While IceCube observes no such natural or artificial sources, there are known sinks in the form of the Sun and Moon, which absorb cosmic rays leading to a deficit in atmospheric muons and neutrinos (see section 2.1) from their direction. Their angular size is small enough compared to the point spread function (see section 3.2.4) so that a study of the Moon shadow based on a large event sample from IC40 and IC59 dominated by atmospheric muons was able to limit the size of an angular reconstruction bias, i.e. the pointing accuracy, to $\lesssim 0.2°$ [140]. An updated study of the Moon and Sun shadows using IC79 and IC86 did not analyse them to produce limits on the pointing accuracy for these detector configurations [141, 142] but unpublished results from these detector configurations are compatible with the limits for IC40/59 [143, sec. 6.5.1].

3.4.2. Pull correction

Another observable that needs calibration is the angular uncertainty estimator. A bias here would lead to worse sensitivity in neutrino astronomy, since the likelihood will not describe the true point spread of a source (see section 7.1.2). In simple terms, an overestimation would lead to larger background contamination, and an underestimation to a loss of signal efficiency. In theory the reconstructed direction, for events with an estimated angular uncertainty of σ, should follow a 2-dimensional normal distribution centred at the true neutrino direction and with the same standard deviation of σ. For most estimators however this is not the case, as can be shown by calculating the so-called **pull distribution**, by dividing the angular difference between truth and reconstruction Ψ by the estimated uncertainty, Ψ/σ. In the ideal case the pull would follow a Rayleigh distribution, but fig. 3.4 shows this to not be the case, as the percentile $(\Psi/\sigma)^{(39.3\%)} \neq 1$ and the median $(\Psi/\sigma)^{(50\%)} \neq 1.18$ in the form of a bias which depends on (reconstructed) energy. This is partially due to the biased estimator σ itself, and partially due to the kinematic angle (see earlier section 3.2.4).

This bias can be approximately removed by applying a correction factor to σ

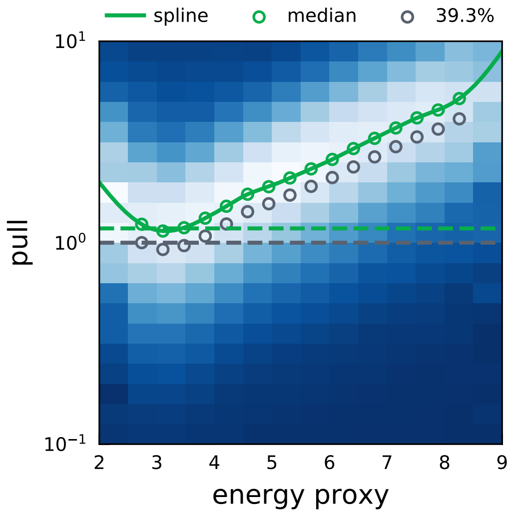

Figure 3.4.: The pull vs. $\log_{10} E_{\text{reconstructed}}$ for the uncorrected paraboloid sigma of IC86a. The energy-dependent bias can be seen both on the median (green) and 39.3-percentile (grey), which differ from 1.1774 and 1 respectively (dashed lines). The solid line shows a spline to correct the pull, although this method wasn't used on this particular sample.

depending on reconstructed energy, which is expected to be a proxy for different statistical behaviour of an estimator and the kinematic angle. The true bias may also depend on the sample composition, which often changes with zenith, and so this correction may be done separately in zenith bands corresponding to the event selection. Finally, where σ is aggregated from multiple estimators, these need to be corrected separately and then recombined.

The energy dependence is accessed by an interpolation of the binned $(\Psi/\sigma)^{(50\%)}(E)$. The interpolation functions have historically been polynomials of 3rd–5th order which are fit to the median data points with a χ^2 minimization. In later samples, increased Monte Carlo statistics showed that these functions could not represent $(\Psi/\sigma)^{(50\%)}(E)$ in sufficient detail. This could be achieved by adding orders to the polynomial, however this makes the extrapolation outside of the reconstructed energy bins increasingly uncontrollable. To mitigate the problem, I introduced the use of cubic splines, which are segmented cubic polynomials. In the fitting method, implemented by `SciPy` [144], the interpolation in one segment of the x-axis is constrained by the contained data points, and does not influence the interpolation of another segment. This method was used for IC86c.

Any extrapolation may lead to extreme values for the corrected σ, in particular among simulation events which span a wider reconstructed energy range then the data. The various samples of the data set remove these events as a quality cut, or replace unrealistically small σ, either due to the extrapolation or the estimators themselves, with a value considered to be a lower bound.

3.5. Event selection

The event selection primarily seeks to distinguish the tracks due to ν_μ interactions from tracks induced by atmospheric μ. Cascades from atmospheric ν also need to be rejected, even if they form a background only at a far lower level. Tracks from atmospheric ν_μ are an irreducible background. The usefulness of a sample depends on how much background is rejected, how much signal retained, and how well-reconstructed it is. The event selection progresses in so-called levels, where each reduces the event rate and therefore makes it feasible to compute more advanced cut variables. Since muons are screened by Earth, the background differs between the Northern hemisphere (upgoing events) and the Southern hemisphere (downgoing events). The event selections at every level take this difference into account.

I will only outline the principles of the event selections and characterize them briefly. These event selections are applied both on the experimental data, and on the signal simulation for the matching detector configuration.

As already mentioned in sections 2.4.3 to 2.4.4, all event selections start with the same online processing, called level 1, followed by filters running at Pole, which are then rerun in the North as level 2. The Muon filter (35 Hz in IC86) is geared towards finding tracks with a minimum quality and brightness. This cut hardens towards vertically downgoing events in order to reject muons based on their lower average energy compared to events from a E^{-2} neutrino spectrum. For upgoing events, the data is also still dominated by misreconstructed atmospheric muons, so the focus is explicitly to reject unreliably reconstructed events. The cut is applied once with LineFit zenith and N_{channel} for both regions, and then again based on SPE zenith, its likelihood i.e. how well it fits a track hypothesis for upgoing, and the total charge for downgoing, where it is also smoothly parameterized to depend on zenith.

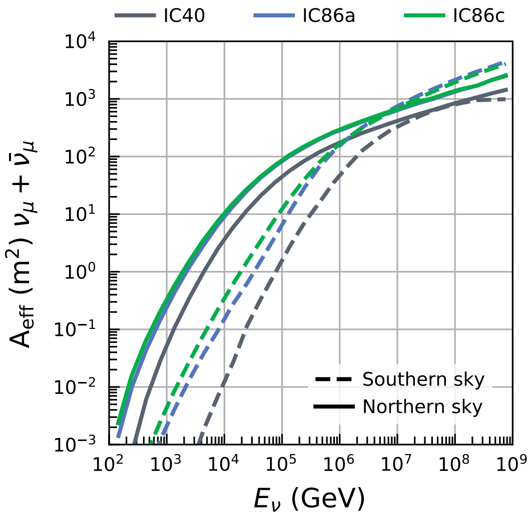

Figure 3.5.: The effective areas for three samples chosen as examples, separated for hypothetical sources in the Northern and Southern hemisphere, with the horizon at $\delta = -5°$.

3.5.1. Point Source samples

The samples IC40 – IC86b also include events from the EHE filter (1 Hz, see section 2.5.2) which selects extremely bright events, although these events are almost all contained in the Muon filter. They then undergo a level 3 offline selection (2 Hz in IC86) whose cuts follow the same principles but exploit the more accurate track reconstructions, for instance to tighten the existing cuts or to find **direct** i.e. minimally scattered hits and cut based on these. The final analysis level generally also uses the upgoing and downgoing events each as input to train a respective multivariate classifier called a BDT. The BDT scores make it possible to distinguish between data (as background) and simulation (as signal) in each of these regions and are used as the final cut.

3.5.2. Gamma-ray Follow-Up sample

The sample IC86c/GFU is special in the sense that its selection (6 mHz) happens entirely online, as described previously in section 2.5.3. It is reprocessed offline with better reconstructions and angular uncertainty estimators. It includes cuts similar to the Muon L3, but using variables optimized for online processing such as an empirical blend between Pandel-SPE and -MPE reconstruction. Like Muon

L3, the reconstruction is also SplineMPE but with an enhanced configuration that provides better speed and angular resolution. It proceeds similarly to the Point Source samples with a BDT cut for up- and downgoing events respectively [75]. This then yields a stream of events with a purity optimized for point source searches on $\lesssim$ 180 day time scales, as it is used for real-time analyses (see section 2.5.3).

The number of neutrino events N_ν in a sample caused by a point source flux $\frac{\mathrm{d}\Phi}{\mathrm{d}E_\nu}$ from declination δ is described by the effective area: The effective areas, representing signal acceptance for a point source flux, of three samples in the data set are shown in fig. 3.5.

These past two chapters have set the stage for neutrino astronomy, a crucial component of multimessenger astronomy which was motivated in chapter 1. We have thus arrived at the starting point for many previous analyses of IceCube data that seek to find sources for the astrophysical neutrino flux. From this point, the work in this thesis takes its own direction towards studying the neutrino emission during blazar flares, which will be introduced in chapter 5. Before these methods are developed and characterized in chapters 7 to 9, they need additional input. This takes us on a detour to gamma-ray astronomy, with the instrument in the following chapter 4 and its higher-level data in chapter 6.

"During the period of the present analysis 127 events occurred which could be gamma rays."

W. L. Kraushaar and G. W. Clark [145]

4 Gamma-ray Astronomy with Fermi-LAT

Chapter 1 established photons as cosmic messengers. This chapter discusses their extreme spectral range known as gamma rays, mainly how these are detected by the Fermi-LAT instrument, and result in the data which forms the basis of the lightcurves in chapter 6. After a brief overview of the Fermi-LAT sky at the end of this chapter, chapter 5 will go on to describe the source class of blazars, including their gamma-ray emission.

4.1. Gamma-ray astronomy

Gamma rays are photons of energy $\gtrsim 100\,\mathrm{keV}$, a range hardly reached by black body radiation. They were first discovered in radioactive nuclear decays, other natural sources include thunderstorms [146] and solar flares [147].

Astrophysical gamma rays have been assumed to be produced by the decay of π^0 produced in hadronic collisions, as well as Bremsstrahlung and inverse Compton scattering involving relativistic electrons [8]. They therefore trace energetic environments in the universe. Since they are deeply penetrating they can reach observers across intergalactic distances. These considerations motivated the first experiment to detect astrophysical gamma rays which was carried by

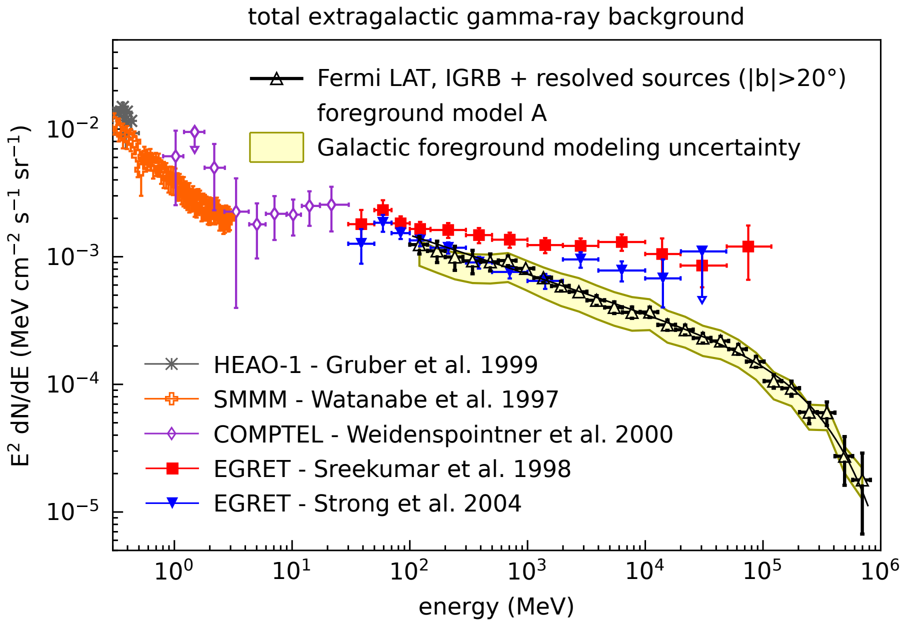

Figure 4.1.: The extragalactic gamma-ray background from 300 keV to 1 TeV, adapted from [150].

the satellite Explorer XI in 1961 [148]. The field of gamma-ray astronomy thus born has since grown to span a range of energies best divided into bands [149]:

300 keV		30 MeV		30 GeV		30 TeV
	low energy		High Energy (HE)		VHE	

Due to the origin of gamma rays, their spectrum follows a non-thermal distribution, as seen in fig. 4.1, where fluxes decrease with energy approximately following a broken power law.

Gamma rays are subject to the propagation effects already discussed in section 1.5 before they reach Earth's atmosphere.

Other than for visible light and radio waves, the atmosphere is far from transparent to gamma rays. Above 30 MeV, this is dominantly due to e^+e^- pair production [149]. For VHE gamma rays, this results in an electromagnetic particle shower (see section 1.1.2) bright enough to be detected on Earth. This detection is mostly achieved either by IACTs when they capture the Cherenkov light (see section 2.1.3) produced by the shower in their field of view, or by water Cherenkov detectors installed on mountaintops when a shower's charged particles reach into their effective area.

HE gamma rays do not cause showers extensive enough to be imaged or detected over large distances. However above the atmosphere, or above a large part of it, they are abundant enough that detectors can have smaller areas, and hence low enough masses to be carried there on a balloon or, more commonly, a spacecraft. Since gamma rays are more penetrating than optical light, such detectors also operate with a high duty cycle, and in the case of an Earth-orbiting satellite with an independent power supply, over the course of years. One such detector currently in operation is Fermi-LAT.

4.2. The Fermi-LAT instrument

Where not explicitly cited, this section refers to the technical papers [151, 152, 153] or the review paper [154].

Figure 4.2a shows the Fermi spacecraft, which was launched on 2008-06-11*, and since then orbits Earth every 96 minutes.

It carries two instruments, the Fermi Gamma Burst Monitor (GBM) and the Fermi Large Angle Telescope (LAT). The latter is a high-energy gamma-ray telescope, sensitive to an energy range of 20 MeV to 300 GeV.

It has three detector subsystems, described in the following sections 4.2.1 to 4.2.3 and sketched in fig. 4.2b. Photons pair-convert in the **tracker**, which measures the trajectory of the resulting $e^{\pm}$ pair. This then deposits its energy in the **calorimeter** which lies underneath. The tracker is constructed with an aspect ratio of $\frac{\text{height}}{\text{width}} = 0.4$ which opens a large part of the sky for which photons and converted $e^{\pm}$ have all tracker layers in their way before reaching the calorimeter. The latter is therefore constructed with an aspect ratio of $\frac{\text{height}}{\text{width}} = 0.4$ which opens a large part of the sky up for observation at any instant, more accurately characterized in section 4.3.3.

Fermi-LAT can observe the rest of the sky thanks to its ±35° rocking motion, which provides an almost uniform exposure after two orbits, or ≈3 h. This **scanning mode** lets Fermi-LAT monitor sources across the entire gamma-ray sky down to $\mathcal{O}(\mathrm{d})$ time scales (see section 4.4).

Another component to Fermi-LAT's ongoing pursuit of its science goals has been its reliability. Relevant in this regard are its radiation hardness, micrometeoroid shield, redundant electronic components, heat management, detectors which do not require consumables such as gas, an on-board power source in the

*then called Glast, `https://www.nasa.gov/mission_pages/GLAST/launch/index.html`

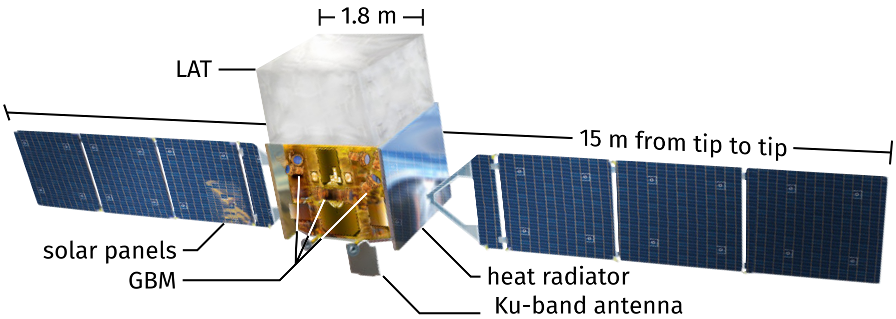

(a) The Fermi spacecraft, 3d model render courtesy of NASA/JPL-Caltech.

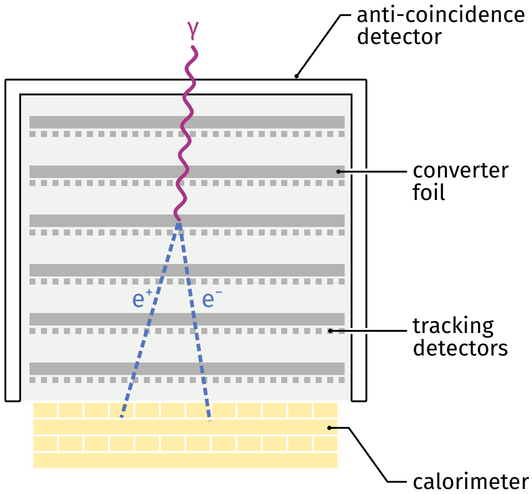

(b) Schematic view of a photon event in Fermi-LAT and the three detector subsystems, adapted from [151], not to scale.

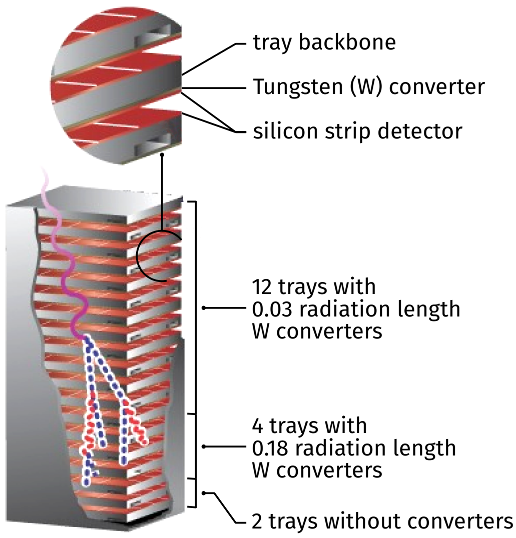

(c) Three-dimensional view of one tracker module and its internal structure, adapted from [155].

Figure 4.2.: Fermi-LAT and its components.

form of its photovoltaic panels, and remote control from the ground to recover from failures. As such, the LAT duty cycle of 88% is primarily limited by the time its orbit crosses the **South Atlantic Anomaly**.

4.2.1. Converter-tracker

The converter-tracker has an approximately square footprint of ~148 cm×148 cm, a height of ~85 cm, and is made up of a 4×4 array of 16 modules. Figure 4.2c shows one such module in a cross section. Its layer structure contains 16 planes of high-Z material, namely 93% tungsten, where incoming photons can pair convert. The more abundant, lower-energy parts of a typical gamma-ray spectrum will result in a significant number of such conversions in the frontmost 12 planes, which have a thickness of 0.03 radiation lengths in order to limit the effect of Coulomb scattering and Bremsstrahlung. However in order to convert enough in the upper energy bands, the 4 planes at the back are six times thicker. Each region ultimately contributes similarly to point source sensitivity.

The silicon strip detectors (SSDs) track the passage of a (minimum ionizing) charged particle like the electron and/or positron through their fiducial volume with $> 99\%$ efficiency. Bonding together 16 of them results in 1536 amplifier channels connected to readout strips 0.2 mm wide and 35.8 cm long, of which 35 cm are active. Two such layers are mounted ~3 cm apart and rotated 90° against one another to form an **x-y plane**. Each tracker module contains 18 such planes in total. The converter material is interrupted above/beneath the inactive portions of the SSDs in order to reduce conversions that are detected only one plane down.

All components are mounted on a light-weight, heat-conductive support structure. Each SSD strip's readout passes through an amplifier-discriminator which provides the information of "hit" or "no hit", with the required electronics mounted around the sides of each module, out of the way of incoming photons.

4.2.2. Electromagnetic calorimeter

The 1.8 t calorimeter uses Thallium activated Caesium Iodide scintillating crystals which produce light proportional to the energy deposited by the cascade of particles produced when an $e^{\pm}$ pair hits a crystal. It is divided into 16 independent modules, aligned with the tracker modules sitting above. Its footprint fills almost the same active area: twelve 32.6 cm × 2.7 cm × 2.0 cm crystals are lined up to form an approximately square, 2 cm thick layer. Alternating layers are rotated by 90°, and the 8 total layers combine to a depth of 8.6 radiation lengths, the latter representing a length scale for shower development. This depth contributes to the energy range, as does the segmentation (see section 4.3.1) which means that each crystal is optically isolated. On both of its opposing small faces, a pair of photo diodes are mounted to detect the scintillation light. One of them is large, the other small, so that they combine to a dynamic range corresponding to 2 MeV–70 GeV deposited per crystal. After a preamplifier and shaper, the analogue signal is multiplexed according to which region of energy deposition is most relevant, and ends at one of 3072 digitizer channels which provide pulse heights. The energy deposited in a crystal is measured via the sum of light collected at both its ends. Taking the ratio instead relates to where between these two faces the energy was deposited, on average.

4.2.3. Anti-coincidence detector

In Fermi-LAT, charged cosmic particles cause a dominant rate of background events which need to be identified. The entire tracker subsystem is therefore "shielded" by an anti-coincidence detector (ACD), which detects > 99.97% of singly charged particles which enter the LAT from its field of view. It consists of plastic scintillator tiles, plus plastic scintillator ribbons underneath the gaps around the edges. The ribbons, and the wavelength-shifting optical fibres threaded through the tiles, couple the scintillation light to photomultiplier tubes mounted at the bottom of the assembly. Each PMT has a redundant spare to ensure the ACD's reliability, which is crucial for Fermi-LAT's continued operation. For the same reason, light must never enter the ACD scintillator from the outside, which are therefore both wrapped, and the wrapping protected by a micrometeoroid shield.

From each PMT, the pulse height is digitized and can be compared with a set threshold to consider the corresponding ACD element as hit.

4.2.4. Trigger and data acquisition

The signals from each calorimeter and tracker module are buffered so they can be correlated into trigger conditions. In order to trigger on photon events, these involve a sliding time window containing a coincidence between a specific combination of **trigger primitives** from different modules. For trackers, the primitive is a coincidence between three planes. For calorimeters, it is when any crystal's signal exceeds a threshold. The ACD checks which tower triggered the event and then requires the absence of a signal in those tiles or ribbons through which the observed particles in the LAT could have entered.

All three are required to trigger the LAT on low-energy photon candidates, while higher-energy events are rare and interesting enough that the ACD doesn't need to be used at this stage.

The spacecraft has a GPS-synchronized 20 MHz clock accurate to ±1.5 μs, which is distributed to the modules/subsystems so that the hit and pulse height data read out from the entire detector upon a global trigger can be built into a single event, which is therefore itself timed with an accuracy <10 μs.

These events are sent at a rate of ~2.5 kHz to onboard processing and filtering units, which apply mainly a preliminary filter which rejects charged-particle background and low-quality events while preserving most photon-induced events. It consists of a series of cuts and processing steps, culminating in a preliminary, fast track reconstruction. This allows to use the ACD as a more precise veto for higher-energy events. Events above 20 GeV deposited energy meanwhile are rare and interesting enough that the filter keeps all of them [154]. The remaining 400 Hz from all filters can be accommodated by the average 1 Mb/s data downlink. However background still dominates this data, necessitating further processing and filtering on the ground, described in the following section 4.3.

4.3. Analysis-level data

As Fermi-LAT's event data arrives on the ground it is processed further. Individual event streams are selected that build into data sets of observables. Following revisions of the involved software chain, all historical data is reprocessed. This section will focus on the event stream optimized in purity for analysing gamma-ray point sources on long time scales, in the software revision "Pass 8, release 2, version 6". This data was used for the lightcurves in chapter 6.

4.3.1. Reconstructions

The ideal event signature is a photon pair-converting in the tracker, the electron and positron leaving distinct tracks through the layers, and ranging out in the calorimeter to deposit all their remaining energy. This ideal is rarely realized. Instead, a photon pair-conversion event may have any of these features:

- at higher energies, the forward boost of the tracks does not let them be separated;
- and the calorimeter may be too small to contain the entire cascade;
- at lower energies, one of the tracks might be too faint to be continuously reconstructed;
- and can be stopped already within the tracker, never reaching the calorimeter.

In light of this, even the advanced Fermi-LAT reconstruction which is used for gamma-ray astronomy at the analysis level is designed as a robust system of fallback methods. The following gives an overview, divided by subsystem, disregarding the order in which they accept each other's output as seed values. The final step is a multivariate classifier tree (CT) which selects the methods which for the specific event are most likely to have yielded the best estimate for the incident photon direction and energy, respectively. Using the known true parameters in Monte Carlo data, CTs also gauge the reconstruction quality.

Calorimeter: For each photodiode, its collected charge is calculated from the digitizer output via the known pedestals and gain. The sum of charges from either end of a crystal is a proxy for the energy deposited therein, and calibrated with Monte Carlo data. Each energy deposition has associated spatial coordinates, on one hand thanks to the segmentation of the calorimeter which lets no light pass between crystals. On the other hand, the balance of the charge between either side corresponds to a mean longitudinal coordinate which can be determined with $O(\text{mm})$ accuracy. This information is then used to find spatial clusters of energy depositions, separated by minimum distances characteristic to the deposited energy. Using a cluster's energy deposition pattern and associated trigger signals, it can then be classified as part of a gamma-ray event, or for example as a **ghost signal**. In Fermi-LAT this describes a phenomenon also called **pile-up** in accelerator experiments, where part of the signal from a previous event, particular one with higher energy, can linger on into the current trigger's readout window.

When the shower is completely contained in the calorimeter, the sum of all energy depositions corresponds to the total energy of the incident particle(s), and their direction can be estimated as colinear to their dipole moment. To determine energies above $\mathcal{O}(\mathrm{GeV})$ however, the three-dimensional profile of the shower needs to be fit as it develops until reaching the boundaries of the calorimeter module. This method also manages to take into account the saturation of the readout, extending the energy range to 3 TeV. The energy reconstructions need to be calibrated to Monte Carlo.

Tracker: Tracks are not generally straight since especially at low energies they can undergo multiple Coulomb scattering within the tracker. Furthermore they can cause showers of their own. The first step in reconstructing the direction of a track is therefore to find the hits which belong to it, for instance with a Kalman filter which also solves for the initial momentum as the r.m.s. deflection from repeated Coulomb scattering is $\propto 1/p_{e^\pm}$ [10, sec. 34.3]. This iterative search can begin at a hit consistent with a shower in the calorimeter, or in lack of such systematically evaluate combinations of hits in the first three tracker layers.

The set of hit tracker strips allows multiple track hypotheses, although these may only share their first hit, and so preference is given to the longest and straightest which are then interpreted as the electron and/or positron trajectories. Their direction in the frontmost hit layers is least affected by scattering, and thanks to the ≈0.2 mm width of the silicon strips contributes most to reconstructing the incident direction of the converted photon. If two tracks can be separated by the tracker reconstruction, they need to be shown consistent with the hypothesis of a common origin vertex by requiring a projected point of approach consistent with the tracker uncertainties.

The energy lost by the electron-positron pairs and their secondaries in the tracker is not negligible for lower energies. It can be reconstructed as a sum of energy losses from layer to layer, due to the known material burden.

Anti-coincidence detector: Similar to the calorimeter, the anti-coincidence tiles and ribbons indirectly measure the energy deposited in them. The ACD also suffers from ghost signals, which are similarly mitigated by checking the timing of its trigger.

A track may not be considered consistent with a electron-positron pair converted inside the tracker if it propagates back to a hit tile or ribbon, defined as exceeding a threshold for deposited energy. The calorimeter direction is more robust in most events and used in the same way, although it requires an event-based uncertainty estimator for angular scale. Both together help reject charged particles entering Fermi-LAT from the outside.

4.3.2. Selection and classification

The higher-level reconstructions which were described in the previous subsection are used to find events as candidates for gamma-ray events. First, events are required to have a track and >5 MeV in the calorimeter, and no hit ACD segments consistent with the track [153]. These events are then scored by classification trees which use the reconstructed observables to judge their suitability as gamma-ray events as well as the quality of the reconstruction. Finally, events are selected with an energy-dependent cut on either of these scores [153]. The Source event class is intended for the study of point sources, and was used in this way for the 3FGL catalogue [156] and their lightcurves in chapter 6. It is a subset of the Transient class, which allows to study transient sources with a larger signal acceptance where the background contamination is limited by the transient duration. It is also a superset of event classes with higher signal purity, which is required e.g. to analyse the diffuse gamma-ray sky [154].

These events may be further selected for the purpose of specific analyses. One example is to require a minimum angle from the Sun, whose activity can contaminate the given region of interest (ROI). During parts of an orbit, the Earth's albedo may also be a source of contamination in the form of events arriving from below, which is thus avoided by requiring a maximum zenith angle [157, p. 4.5.1].

4.3.3. Performance

The LAT instrument, the reconstruction of its events, and their selection described in this and the previous section all determine how well the final data set performs in an analysis. Part of this performance is how it reflects a given gamma-ray flux. Fermi-LAT describes this with instrument response functions (IRFs) which depend on e.g. gamma-ray energy and incident angle. These will be the focus of this subsection.

Effective area and acceptance: Fermi-LAT's tracker intercepts a flux Φ of gamma rays of an energy E from a direction Ω with its geometric area. However only a fraction of these gamma rays leaves an event in the data set, since they may not be correctly identified by the detector and processing, or not pair-converted in the first place. This fraction is determined, in the absence of a calibration source, with specialized Monte Carlo simulations where Fermi-LAT is illuminated by N_{gen} photons from a sphere of $6\,\mathrm{m}^2$ cross section that envelops the entire LAT [152]. The density function of N_{gen} with respect to direction Ω and $\log E$ is known, and its counterpart for the number of events in the data sample n can be estimated by binning the simulation results. This defines the **effective area**:

$$A_{\mathrm{eff}}(E,\Omega) = 6\,\mathrm{m}^2 \left(\frac{\mathrm{d}^2 n}{\mathrm{d}\log E\,\mathrm{d}\Omega}\right)\left(\frac{\mathrm{d}^2 N_{\mathrm{gen}}}{\mathrm{d}\log E\,\mathrm{d}\Omega}\right)^{-1} \tag{4.1}$$

Taking Ω as the direction in the instrument frame of reference shows a strong dependence on the incident zenith angle shown in fig. 4.3a because of geometric effects. The flux from any fixed source in the sky meanwhile will be exposed at different angles, which leads to the definition of **acceptance** $\mathcal{A}(E) = \int_\Omega \mathrm{d}\Omega A_{\mathrm{eff}}(E,\Omega)$. The remaining dependency on energy is graphed in fig. 4.3b and reveals the energy range of Fermi-LAT.

As this acceptance is distributed over the sky, it is possible to define its characteristic extent, called the FoV, as

$$\mathrm{FoV} = \frac{\mathcal{A}(E)}{A_{\mathrm{eff}}(E,\theta=0)} \tag{4.2}$$

(4.3)

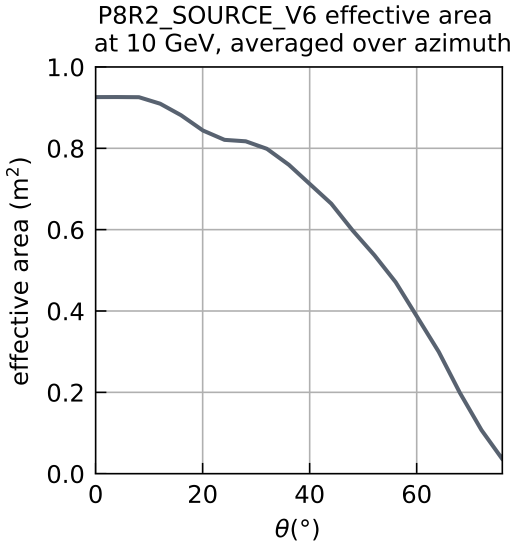

(a) Effective area versus zenith angle for 10 GeV photons.

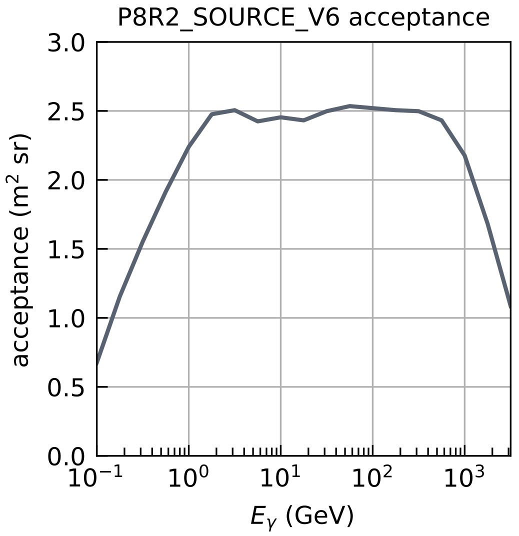

(b) Acceptance versus energy.

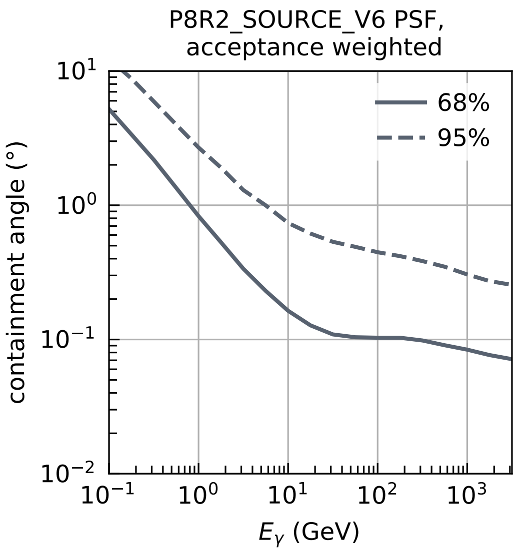

(c) Angular resolution in terms of the angular radius containing 68% or 95% of the PSF in an energy bin.

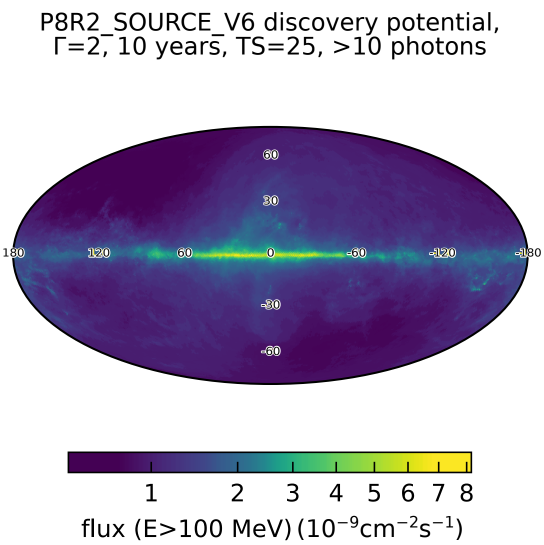

(d) Power law flux required to discover a point source at any point in the sky, assuming known diffuse backgrounds but no other sources.

Figure 4.3.: Plots of the IRFs described in section 4.3.3, for the P8R2_SOURCE_V6 data set, digitized and/or adapted from the originals.

which peaks around 2.6 sr for energies 1 GeV to 100 GeV.

Angular resolution: The effective area or acceptance determine how many gamma rays from a point source end up in the data set. How many of them can be associated to the source depends, besides other factors, on how widely their reconstructed directions scatter around the source location. Among gamma rays above 1 GeV, at least 68% are reconstructed within 1° of their true direction. As shown in fig. 4.3c, this **angular resolution** improves with higher energy as the total deflection by Coulomb scattering in the converter foils (see section 4.3.1) diminishes [152]. For energies >10 GeV it asymptotically approaches $\mathcal{O}(0.1°)$, which is the angular distance corresponding to the width of one SSD strip after traversing a characteristic number of tracker planes [152].

Sensitivity: Given a point source, whether Fermi-LAT will significantly detect it depends on the IRFs described so far, as well the observing profile of that location in the sky and the diffuse background in its vicinity. The instrument performance can be summarized in the flux normalization required for the 5σ observation of a $\propto E^{-2}$ spectrum after 10 years sky scanning, defined for each point in the sky as a test statistic of 25 given at least 10 contributing photons. The map fig. 4.3d shows that this **discovery potential** does not exceed $4 \times 10^{-9}\,\mathrm{cm^{-2}s^{-1}}$ in regions away from the galactic plane.

4.4. The Fermi-LAT sky

Fermi-LAT sees the high-energy gamma-ray sky (see fig. 4.4) primarily as diffuse emission [158, 159]. This is brightest in the Galactic plane, which Fermi-LAT can attribute to the interaction of high-energy cosmic ray nucleons and leptons with the interstellar medium and radiation field [158]. In addition, Fermi-LAT has discovered two structures extending from the Galactic centre to 55° on either side of it. These may be related to past outflow or jets from the black hole at Sagittarius A*, and have been dubbed the **Fermi bubbles** [160]. The remaining diffuse emission when not associated to the Milky Way is isotropic. While this can contain residual Galactic foregrounds, spectral analysis points to unresolved sources as a significant contribution [150].

The other 20% of detected photons belong to sources bright enough to be

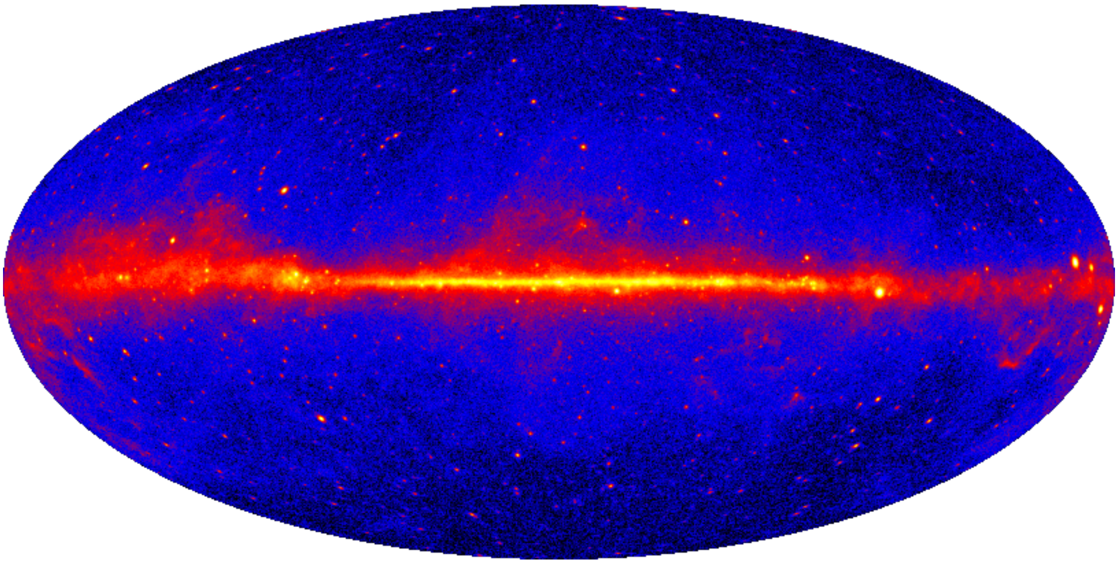

Figure 4.4.: The gamma-ray sky above 1 GeV, seen in 5 years of Fermi-LAT observation, galactic coordinates. Image credit: NASA/DOE/Fermi LAT Collaboration.

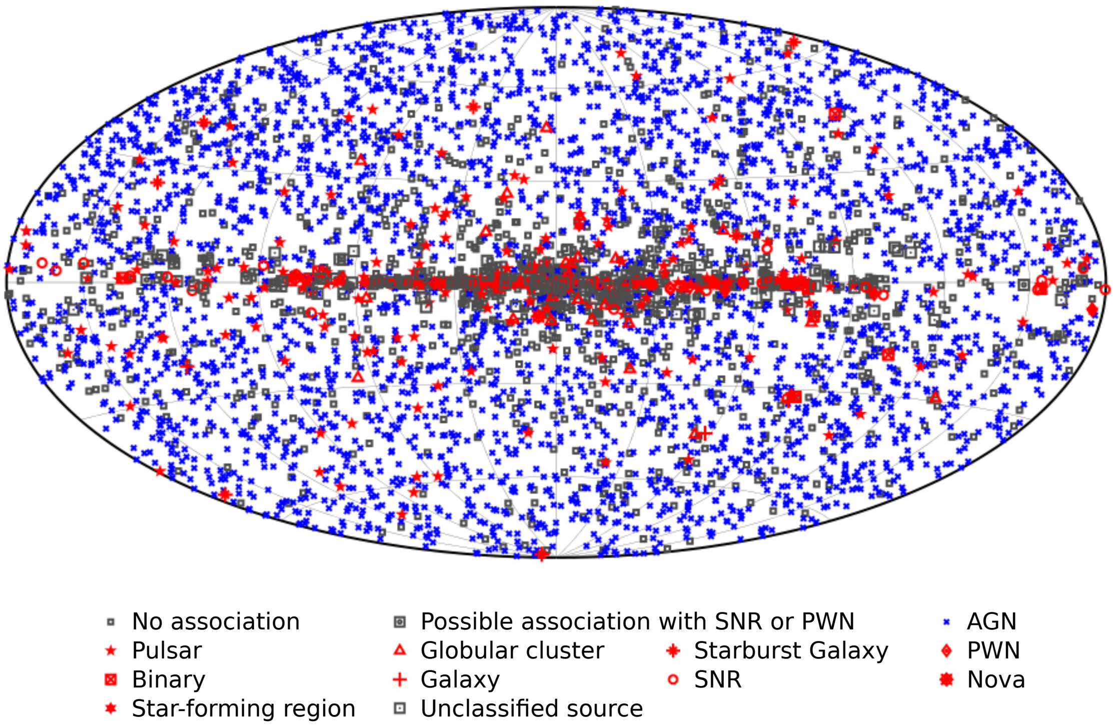

Figure 4.5.: A sky map of sources in the fourth Fermi-LAT source catalogue. The class is indicated for sources that have been associated or identified [59].

resolved and localized. Some of these are transients, namely GRBs discovered e.g. by Fermi-GBM [161], and novae [162]. Fermi-LAT also detects gamma rays (or misidentified charged particles) from the Earth limb, the Moon, and the Sun [59], including in the form of solar flares [163].

A further 5064 sources are revealed by the integrated observation over the course of 8 years [59] (see fig. 4.5). Of these, 74% have been identified or associated to sources known to astronomy at lower energies. They represent a broad set of galactic and extragalactic classes of extended and point sources, such as pulsars, pulsar wind nebulae, supernova remnants, globular clusters, binaries and starburst galaxies. Among the number of Fermi-LAT sources of their combined flux, a majority is however due to AGN, and in particular blazars. They are the subject of dedicated catalogues [164], as well as the following chapter.

“Everything that flares in gamma rays should be treated like a blazar.”

Kevin J. Meagher

5 Blazars

Blazars dominate the gamma-ray sky opened up in chapter 4. This chapter provides an insight first into their taxonomy and morphology, which are introduced in sections 5.1 to 5.3. It goes on to describe the high-energy processes suspected in blazars in section 5.4, and their classification in section 5.5. Finally, section 5.6 introduces the possibility of neutrino production, which ties them to the multimessenger principle presented in chapter 1, and leads to the work described in chapters 7 to 10.

Unless otherwise cited, the source for the following sections 5.1 to 5.3 is [165].

5.1. Active and inactive galaxies

The observable universe's luminous matter is largely bound in galaxies, containing stars. These emit most of their electromagnetic radiation due to the temperature of their surface [138, p. 227], with temperatures of 3×10^3 K to 3×10^4 K [138, p. 230] leading to black-body spectra which therefore cover wavelengths of infra-red, visible, and ultra-violet light. Those combine to make up the spectrum of an **inactive galaxy**. **Active galaxies** meanwhile radiate significantly more than would be expected from their stars, and cover the entire electromagnetic spectrum from radio waves to gamma rays. Figure 5.1 shows

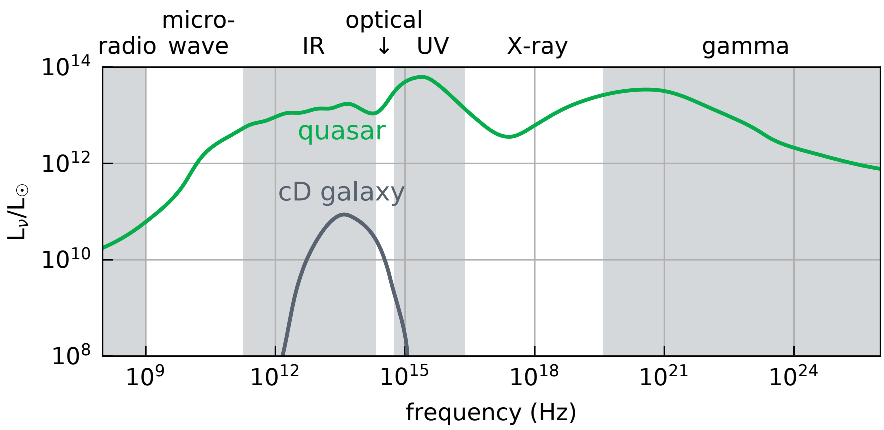

Figure 5.1.: The spectra in terms of solar luminosities for a quasar and a cD-type galaxy, respective examples of an active and inactive galaxy. Figure adapted from [166, fig. 3.4].

this difference in terms of the monochromatic luminosity* L_ν. This additional radiation has long been observed to originate from the active galaxy's central region [168], called the **nucleus**.

5.2. Central engine

Is is now understood that these AGN contain a supermassive rotating black hole [169, pp. 37] which acts as an engine to convert the potential energy of accreted material [169, sec. 3.2], and possibly its own angular momentum [169, sec. 3.7], into the observed non-thermal spectrum via a network of mechanisms.

Around the black hole, the accreting material forms a disk [167, p. 125]. This process gives the accretion disk a total luminosity L equivalent to $\lesssim$ 10% the rate at which matter is accreted, in terms of mass energy [170]. This luminosity scatters on the accreting material, assumed to be a fully ionized gas, and thereby imparts a force. For a unit volume at a distance r from the source of radiation, this force is $\propto \frac{L}{4\pi r^2}\sigma_T N_e$, with N_e the electron density and σ_T the Thomson cross section. Accretion can only continue as long as the resulting outward radiation pressure does not overcome the gravity of the black hole. The same unit volume has a mass $N_e \mu m_p$ with μ the number of nucleons per electron and m_p the proton

*the energy radiated per unit time and unit frequency [167, ch. 1] mass. It therefore experiences an inward gravitational force $\propto \frac{GM}{r^2} N_e \mu m_p$. These effects result in a maximum luminosity L called the **Eddington limit** [171, p. 40]:

$$L_{\text{Eddington}} = \frac{4\pi c G M \mu m_p}{\sigma_T} \tag{5.1}$$

$$\simeq 1.5 \times 10^{38} \frac{M}{M_{\text{sun}}} \text{erg/s}, \tag{5.2}$$

where the characteristic value is given for a μ equivalent to that of the Sun. For an AGN with luminosity $L = 10^{47}$ erg/s this corresponds to a minimum black hole mass

$$M_{\text{Eddington}} \geq \left(\frac{L}{1.5 \times 10^{38}\,\text{erg/s}} \right) M_{\text{sun}} \tag{5.3}$$

$$= 7 \times 10^{8} M_{\text{sun}}, \tag{5.4}$$

In practice, this limit only holds strictly for spherical symmetry and steady states, while the accreting matter will form a disk, and the radiation is variable (see section 5.4.4) and anisotropic (see section 5.3). But even with some orders of magnitude tolerance, the actual masses 10^6–$10^{10} M_{\text{sun}}$ [172] fall firmly in the range of supermassive black holes (SMBHs), just like those in the core of many galaxies.

5.3. Multi-wavelength overview of AGN

5.3.1. Centaurus A and basic unification morphology

The MWL image of Centaurus A, Earth's nearest AGN, is shown in figs. 5.2a to 5.2e. The complex morphology means that an observer sees only the spectral energy distribution (SED)† according to the angle at which they observe the AGN. In Centaurus A we find many common features of an AGN. As close as Cen A is, other AGN do not allow that kind of resolution and therefore their classification has relied on the SED in different bands of wavelengths. In principle this produces many possible combinations to define an AGN class,

† $E^2\,\mathrm{d}N/(\mathrm{d}A\,\mathrm{d}t\,\mathrm{d}E)$ with number of photons N, energy E, time t and area A, units (erg cm^{-2} s^{-1})

(a) Radio **(b)** Infrared

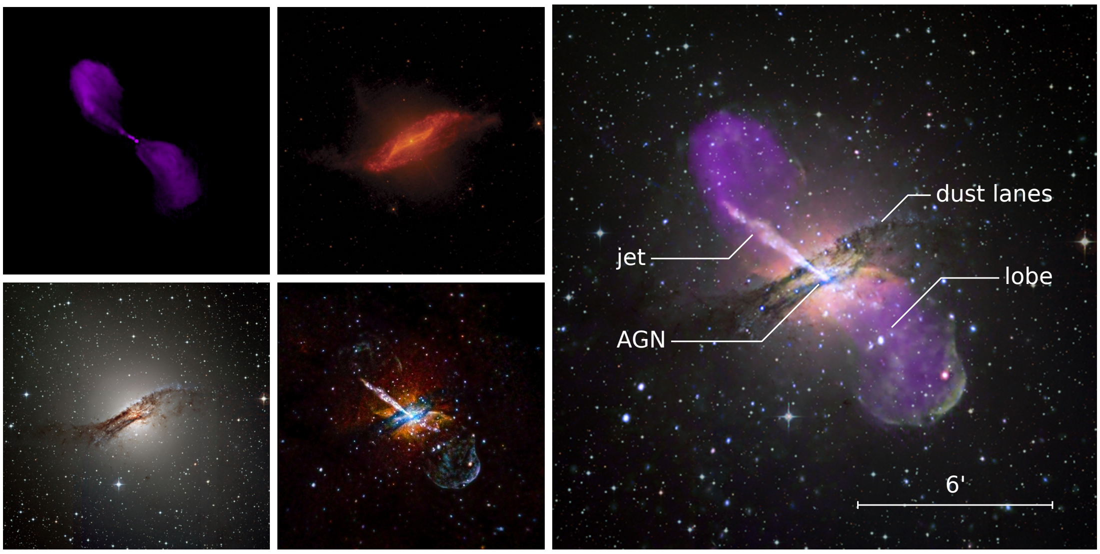

(c) Optical **(d)** X-rays **(e)** Composite of radio, optical and X-ray

Figure 5.2.: The jetted AGN Centaurus A in multiple wavelength bands. Image credit, radio: NSF/VLA/Univ.Hertfordshire/M.Hardcastle, infrared: NASA/JPL-Caltech/J.Keene(SSC), optical: ESO/WFI/M.Rejkuba et al., X-ray: NASA/CXC/U.Birmingham/M.Burke et al., composite, scale, labels: NASA/CXC/CfA/R.Kraft et al.

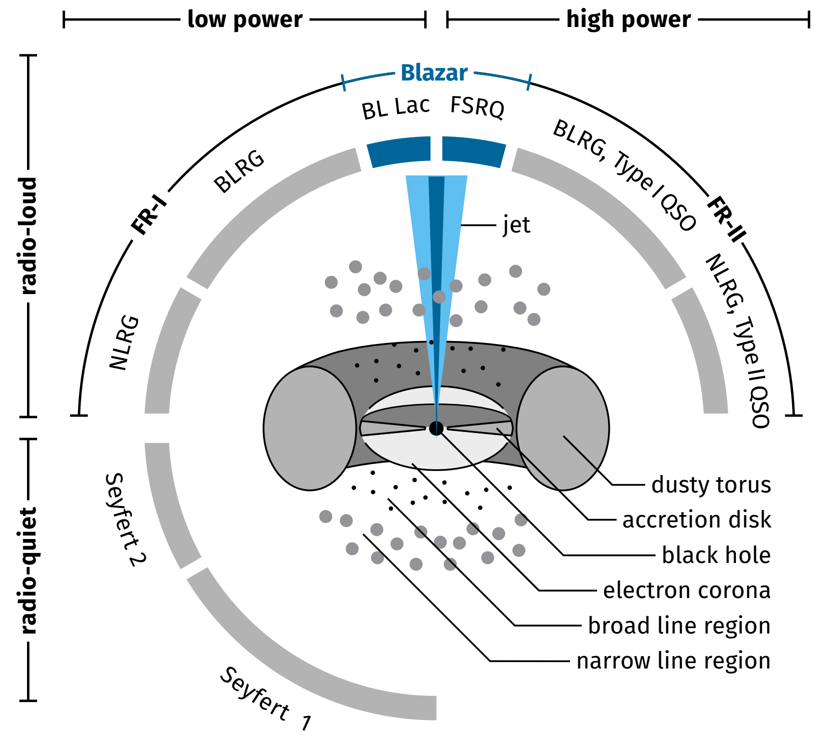

Figure 5.3.: Classification of AGN, adapted from [169, fig. 1.1], original graphic by Marie-Luise Menzel, inspired by [173]. The grey arc indicates different viewing angles. Note that while one side of the diagram shows no jet to represent the radio-quiet AGN, most radio-loud AGN in fact feature symmetrical jets.

and an observation in one band does not help to understand physics that would be reflected in another. By now, however, the research has culminated in the unified AGN model [173, 174] where many of these differences are understood as different viewing angles on a morphology made up of a standard set of components. This is sketched in fig. 5.3. These are [167, sec. 8.3]:

- The supermassive black hole (10^6–$10^{10} M_{\mathrm{sun}}$) which acts as the **central engine** by accreting matter.
- The latter is thus shaped into an **accretion disk**.
- This is accompanied by a **corona** of hot electrons, which radiate in X-rays.
- A **torus** made of dust and gas, which surrounds this central region and can obscure its radiation, but also re-emit it in the infrared.
- Further gas is found either orbiting or as an in-/outflow [167, p. 157] in clouds further from the accretion disk, where they as well intercept a portion of its radiation. The lines which they re-emit are Doppler-broadened, lending the name broad line region (BLR) to these closer, faster-moving clouds. Meanwhile the slower clouds which orbit farther from the accretion disk emit narrower lines, and are thus called the narrow line region (NLR).
- Cen A, as $\sim 10\%$ of AGN, also has relativistic outflows called **jets** projecting from the central region perpendicular to the accretion disk or colinear with the black hole rotation axis. These are particularly visible in the radio band due to synchrotron radiation.
- These more-or-less collimated jets end in broader **radio lobes** which move more slowly [169, sec. 4.5].

5.3.2. Optical/UV

Optical spectra are characterized by their emission lines from the recombination of nuclei. All lines have a width $\Delta\lambda$ because of the Doppler shift of the emitting atoms:

$$\frac{\Delta\lambda}{\lambda_0} = \frac{\delta v}{c} \tag{5.5}$$

where λ_0 is the unbroadened wavelength and δv the width of the velocity distribution. Observations show that these relative velocities are higher than could be explained by the thermal motion of an un-ionized gas [165], leading to the conclusion that these atoms must be for example orbiting the central region. The unification scheme explains the emission lines of type 1 AGN as originating from a BLR where the emitting gas moves at a higher velocity, located close to the accretion disk. Therefore the dusty torus obscures the BLR for a range of viewing angles - in which case one observes a type 2 AGN instead, where the optical spectra are characterized by emission from the more distant, slower NLR. Optical spectra which deviate from these major groups are ascribed to "type 0" AGN, which include those without strong emission or absorption lines [173].

5.3.3. Radio

One distinction independent of the unification is that between radio-quiet and radio-loud AGN, where the latter exhibit synchrotron radiation from accelerated electrons in magnetic fields of the turbulent plasma [165]. They can be identified spectrally by comparing the monochromatic radio and optical flux as $F(5\,\mathrm{GHz}) > 10 \times F(5000\,\text{Å})$ [174], which applies to $\approx 10\%$ of AGN, including Cen A. Morphologically, they have **jets** projecting from the centre out to distances of $\mathcal{O}(10\,\mathrm{kpc})$–$\mathcal{O}(1\,\mathrm{Mpc})$ [174], generally greater than the parent galaxy radius, ending in **radio lobes**. Their emission at many wavelengths is dominated by the jet [175].

The spectrum of most such AGN in the radio band depends on the frequency ν with a power law $F(\nu) \propto \nu^{-\alpha}$ [171]. The lobes have a steeper spectrum, which is observable when seen from the side. But when the jet's beamed emission points towards the observer, the core (see section 5.4.3) dominates with its flat spectrum [174] [171, p. 232]. The index α lends itself to an according division into **flat** ($\alpha < 0.5$) and **steep** ($\alpha > 0.5$).

Next to optical, radio telescopes deliver the highest-resolution images of AGN jets, especially using very-long-baseline interferometry (VLBI). This reveals a dichotomy of two kinds of jets, for which examples are shown in fig. 5.4. One is called Fanaroff-Riley type 1 (FR-I). Here, the intensity of the jets diminishes with the distance from the centre [173]. For Fanaroff-Riley type 2 (FR-II) meanwhile, the more collimated jets [173] end in sharp radio lobes whose respective brightest spots are separated more than half the total size of the source, called "edge-

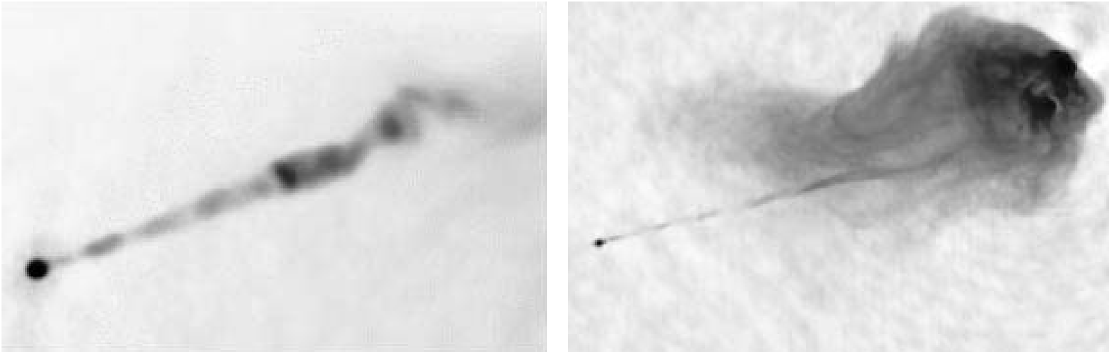

(a) The FR-I M 87 [169, fig. 3.25]. **(b)** The FR-II Cygnus A [169, fig. 4.9].

Figure 5.4.: Examples for the Fanaroff-Riley types mapped in the radio with VLBI by the Karl G. Jansky Very Large Array (VLA). The respective host galaxy can not be seen in the radio, and so a scale for the relative jet extension is absent.

brightened". The latter are less numerous, but more luminous [171, pp. 127].

5.3.4. X-ray

The core regions of AGN emit X-rays. Consequently this is particularly visible where the core region is exposed to the observer in type 1 AGN. X-rays are produced when lower-energy photons, which are emitted from the accretion disk, meet a surrounding corona of hot electrons, and are Comptonized up to X-ray energies. They mainly follow a power-law spectrum of index 0.9–1.0 which cuts off at 40 keV to 300 keV, but also feature fluorescent lines, and a portion of the X-rays scattered back from the accretion disk [167, sec. 8.7]. Often, this X-ray emission from the core also extends into the jet, as fig. 5.2d shows.

5.3.5. Infrared

As can be seen in fig. 5.2b, the infrared emission traces the dusty torus which glows at 10 K to 10^3 K. The responsible heat source is the accretion process which radiates as previously described into the surrounding torus. This way, even heavily obscured AGN can reveal the existence of their central engine.

5.3.6. Gamma ray

Some AGN also exhibit gamma-ray emission, in which case they are also jetted [171, p. 11]. This has been detected in GeV energies for many AGN [164], and in TeV energies for some [176, 177]. As described in the following section 5.4, this is a combined result of relativistic particles and relativistic Doppler beaming.

5.3.7. Unifying picture

The unification of AGN [173, 174] accommodates classes that were historically distinguished observationally (see [169, tab. 4.4]). A selection of them is indicated in fig. 5.3, along with distinctions not based on orientation, and the following gives an overview of their differences in optical and radio. The viewing angle determines the type 1 or type 2. Those without a jet are accordingly called Seyfert-1 and Seyfert-2. Jetted AGN meanwhile are further divided into "weak jet" FR-I and "strong jet" FR-II. Also for these, a viewing angle perpendicular to the jet axis explains the dominance of narrow optical lines. Moving closer to the jet axis, the dusty torus reveals the broad line region. Depending on geometry and optical depth of the torus, the radio spectrum at this point may still be steep, dominated by the lobes [174]. For viewing angles aligned with the jet axis to $\lesssim \mathcal{O}(1°)$ [171, p. 131], the observer is within the jet's beamed emission. This leads to a flat radio spectrum and is assumed to characterize blazars. They are this chapter's subject from here on out, and further classified in section 5.5.

5.4. High-energy emission from blazars

Figure 5.5 shows an example of a blazar SED, which are thought to be dominated by the jet itself and exhibit two typical humps. The first (at radio–X-ray) originates with electrons and positrons in the relativistic jet plasma emitting synchrotron radiation [45], as indicated by its polarization [169, sec. 4.5]. For that reason, it is often called the **synchrotron hump**.

The second hump extends to high energies, with TeV gamma rays having been observed. The same figure also shows a range between low and high fluxes (see section 5.4.4) observed over the course of several years. Before unification, blazars were identified by observational characteristics, one such characteristic being this intense and highly variable gamma-ray emission [171, p. 130].

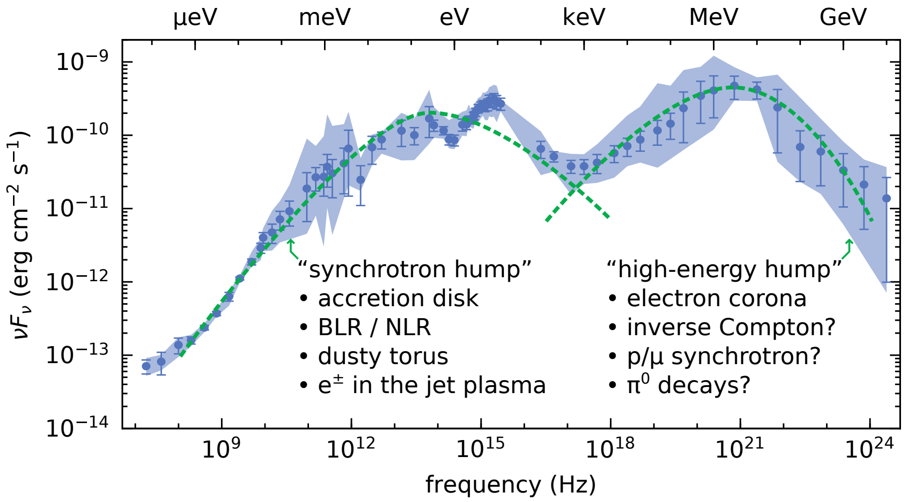

Figure 5.5.: The SED of the blazar 3C 273 (as in fig. 5.1), combined from 4–44 years of observations at different wavelength. The points show average values with standard deviations, and the shaded band the range in the observation. Figure adapted from [178].

There are multiple hypotheses on where it is produced, by which particles and mechanisms, and how the latter are injected.

5.4.1. Hadronic and leptonic

One distinguishes two main scenarios [179, 45].

Leptonic: The same electrons responsible for the synchrotron hump up-scatter softer photons in the inverse Compton effect. The softer photons could either

- be the synchrotron radiation from the same electrons, called synchrotron-self Compton (SSC);
- originate outside the jet, for example from the BLR, i.e. external Compton (EC) [167].

Hadronic: Protons and possibly other nuclei are accelerated along with the known relativistic electrons.

- These protons can themselves emit synchrotron radiation. This proton synchrotron radiation can possibly reach higher energy than that from electrons in the same environment. This is because the latter have a shorter synchrotron cooling time, and therefore have a lower maximum energy than said protons (see [167, sec. 4.3], [2, sec 12.3.2]).
- Photo-meson interactions (p-p or p-γ, see section 1.4), wherein $\pi^0 \rightarrow \gamma\gamma$. Softer target photons are available in abundance e.g. from the synchrotron hump [180].
- Another by-product of hadronic interactions are muons, which would synchrotron-radiate as well [179].
- Finally, these reactions also produce electrons (e.g. via Bethe-Heitler pair production), which could then produce the low-energy "synchrotron" hump as well as contribute to the "inverse Compton" hump. This scenario is called **purely hadronic** [45].

For any particular blazar, either of these scenarios may contribute to the entire network of particle populations and radiation fields. For some, it has been shown that a purely leptonic scenario can not supply the entire luminosity [169, pp. 232]. A unique feature of an hadronic scenario is neutrino production (see section 1.6). This is particularly interesting since the responsible relativistic protons would also contribute to the extragalactic cosmic ray flux, whose sources have not been determined (see section 1.1).

5.4.2. Shocks

Either of these scenarios requires particle acceleration. Assuming this happens as discussed in section 1.2, it requires a plasma, which is given by the jet and/or accretion disk, and ultra-sonic shocks.

The jet is launched close to the black hole [171] in the so-called **blazar zone**, closer than a parsec to the black hole [181]. Fitting observed blazar SEDs with SSC favours a scenario where the jet energy is mostly carried by particles (at its launch, or soon after), which are then further accelerated at mildly relativistic shocks within the relativistic bulk motion of the jet [181].

Particles accelerated in the blazar zone are similarly boosted and collimated as the overall jet, and thus flow further out to kpc scales [181]. For radio-loud AGN which are not blazars, these jets can be mapped, revealing a lumpy structure (see fig. 5.4) called **knots**. These could mark a trail of shocks left behind by plasma blobs travelling from the central region [169, sec. 3.7.2]. The knots may also be the blobs themselves [169, sec. 3.7.2]. In addition, there are FR-I jets that feature an X-ray component compatible with models which feature synchrotron emission from electrons of individual Lorentz factors $\gamma \sim 10^7$. These can not travel far from the acceleration site before cooling down. That would imply that acceleration does not occur at individual sites, but continuously along the jet [181].

Other candidate shock sites are

- **Mach disks**, where a shock across the jet cross-section arises due to its own dynamics under internal and external pressure [169, sec. 4.5], and
- the **termination shock** where a "cocoon" around the jet, part of which are the radio lobes, expands into the intergalactic medium [2, sec. 12.6] [182, 183]. This is the largest scale of shock considered for AGN.

If thus the jet plasma contains shocks (see section 5.4.2) and protons (see section 5.4.1) which run into them, they will be accelerated. This forms the theoretical basis for blazars as cosmic ray accelerators.

5.4.3. Relativistic beaming

We assume the previously described emission originates from regions which move along with the relativistic jet [181], or at least move relativistically within it [184]. Their emission is therefore subject to the relativistic Doppler effect which blue-shifts the emitted frequency ν' to the observed frequency $\nu = \delta\nu'$ by the Doppler factor [167, p. 43]:

$$\delta = \frac{1}{\Gamma(1 - \beta \cos\theta)} \tag{5.6}$$

$$\xrightarrow{\theta\approx 0} \delta = \Gamma(1 + \beta) \tag{5.7}$$

$$\approx 2\Gamma, \tag{5.8}$$

where Γ and β are the Lorentz boost and velocity of the emitting region moving at an angle θ with respect to the line of sight. An example is the blazar

Markarian 421, where modelling several multi-wavelength flares with SSC in a spherical region results in an estimated Doppler factor $\delta \gtrsim 50$ [185]. Considering special relativity for the entire scenario, one finds the radiation is both collimated, intensified and blueshifted in the observer frame so that

$$L_{\text{obs}} = \delta^4 L_{\text{src}}, \tag{5.9}$$

at least for the simplified assumption of a single, spherical emission region moving at a uniform velocity [167, p. 45].

As a result of beaming, blazars dominate the extragalactic sky of persistent gamma-ray point sources [175]. This is true both above 100 MeV with Fermi-LAT (see section 4.4), and above 100 GeV with various IACTs [186, 187]. Beaming is also necessary to verify whether radio luminosity functions of different AGN classes are consistent with the unification paradigm [174].

5.4.4. Blazar variability

One observational characteristic of blazars is their rapid variability. When the flux in question changes on a time scale Δt in the observer reference frame, this can be explained as a region of size R changing its emission activity. Neglecting the time it takes for energetic particles to be injected into the region and then radiate, the light crossing time of the region is a lower limit on the variability time scale Δt [180]. Therefore when given Δt, R is limited by [188]

$$R \leq \Delta t \cdot c \cdot \delta / (1 + z), \tag{5.10}$$

with c the speed of light, δ the relativistic Doppler factor (see eq. (5.6)), and z the redshift.

The high-energy emission from blazars is often understood as a superposition of **quiescent** and **flaring** states. The former are responsible for periods of approximately steady emission. Then during flares, the gamma-ray flux increases, sometimes by several orders of magnitude and with $\mathcal{O}$(minute) doubling time scales [188, 189]. This fact suggests high-energy emission regions in a suitably small region of the jet:

$$\Delta t = 300\,\mathrm{s} \tag{5.11}$$

$$\Rightarrow R \leq 10^{11}\,\mathrm{m} \cdot \left(\frac{\delta}{1+z}\right) \tag{5.12}$$

$$= 10^{-4}\,\mathrm{pc} \cdot (\delta/50) \tag{5.13}$$

assuming values for the blazar 3C 279, whose gamma-ray variability is described in [188]. The same reference goes on to assume that a $x = 10$–100 times larger region supplies the energy to where it is radiated. In the conical jet with opening angle $\theta \ll 1$, this is the characteristic distance found at a radius $r \approx R/\theta$ near the black hole, and finally estimates a scale of $\approx 100 R_S$ measured in a typical SMBH Schwarzschild radius, still within the BLR. A similar estimate was found for the central emitting region of the FR-I radio galaxy M 87 [181].

This principle can be observed when resolving "blobs" in the AGN jet (see section 5.4.2) and isolating their radio emission. The latter is found to fluctuate on longer time scales than the other emission from the same AGN [190].

5.5. Blazar classification

5.5.1. FSRQ vs. BL Lac

The optical spectrum of some blazars shows broad lines on top of the non-thermally produced continuum. These are called Flat Spectrum Radio Quasars (FSRQs), and those without such lines BL Lacertae objects (BL Lacs). A precise criterion has been adopted using the **equivalent width**‡ of these lines; >5 Å indicates an FSRQ, ≤5 Å a BL Lac [173]. This implies that blazars with an especially bright optical continuum that hides such lines may be (mis-)classified as BL Lacs [192]. The unification identifies

- BL Lacs with radio-loud FR-I AGN,
- FSRQs with radio-loud FR-II AGN,

when observed near the jet axis, by virtue of multi-wavelength astronomy of the AGN and their host galaxies [173, sec. 5]. This unification also agrees with the

‡the integrated flux density of the line, divided by the surrounding continuum flux density [138, fig. 5.6]

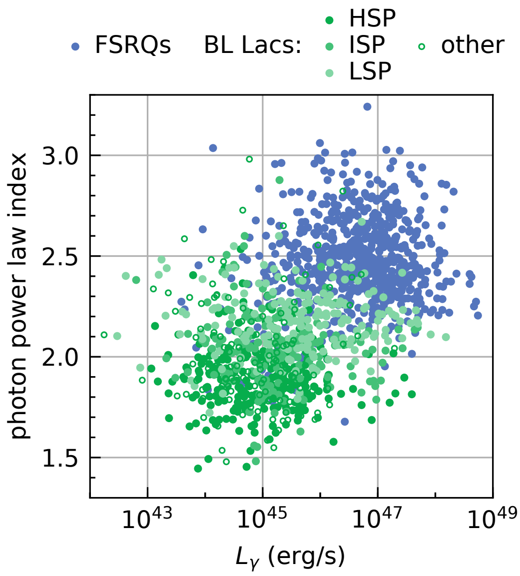

Figure 5.6.: The gamma-ray luminosity and spectral index of blazars in the Fermi 4LAC-DR2 catalogue [164, 191], divided into FSRQs (blue) and BL Lacs (green). The latter are subdivided further (see the following section 5.5.2).

observed statistical distributions of radio luminosity [173, sec. 6]. Exceptions to this identity have however been found [193]. Although the Fanaroff-Riley classification uses the radio jet morphology (see section 5.3.3), the FR-II also exhibit stronger optical emission lines than FR-I [173, sec. 5.2]. As the emission lines are produced when the BLR is illuminated by the accretion disk, there is an argument that that FSRQs have a higher disk luminosity, particularly in the ionizing UV band [194]. This suggests a link to the accretion mechanism itself, which in turn may explain the difference in the jet launched near the central engine [171, p. 140]. Among the blazars in the Fermi sky, FSRQs have higher luminosities and higher redshifts [164]. If a higher luminosity is intrinsic to FSRQs, then their population as observed in the flux-limited Fermi catalogue would stretch to higher redshifts (see section 9.2.1). The separation of Fermi-BL Lac and -FSRQ luminosity becomes cleaner when adding the gamma-ray spectral index as another dimension [170, 164], as shown in fig. 5.6.

5.5.2. Synchrotron peak

The frequency ν^S_{peak} at which the synchrotron hump of the SED reaches its peak is found to differ between blazars. This is interesting since according to our

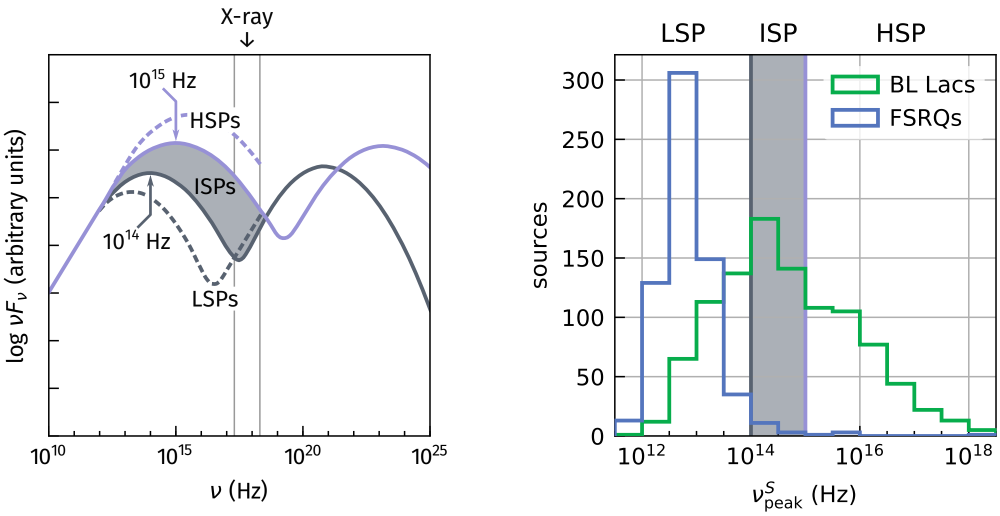

(a) Idealized example SEDs for the synchrotron peak classification, adapted from [175].

(b) The synchrotron peak frequencies of blazars in the Fermi 4LAC-DR2 catalogue [164, 191].

Figure 5.7.

understanding of blazar emission (see section 5.3), it relates to the Lorentz factor γ of the corresponding relativistic electrons accelerated in the jet [175]:

$$\nu^S_{\text{peak}} = 3.2 \times 10^6 \frac{(\gamma^S_{\text{peak}})^2 B\delta}{1+z} \tag{5.14}$$

with magnetic field intensity B, Doppler factor δ (see section 5.4.3) and redshift z. However, ν^S_{peak} is not determined by γ as a degeneracy with the product $B\delta$ remains [175]. It is thus possibly related to high-energy emission from the jet. Therefore a classification based on the synchrotron peak frequency was introduced [175]:

- Low synchrotron peaked (LSP) blazars first peak at low energy with $\nu^S_{\text{peak}} \lesssim 10^{14}$ Hz.
- Intermediate synchrotron peaked (ISP): 10^{14} Hz $\lesssim \nu^S_{\text{peak}} \lesssim 10^{15}$ Hz.
- High synchrotron peaked (HSP): to reach $\nu^S_{\text{peak}} \gtrsim 10^{14}$ Hz, the jet requires more energetic particles.

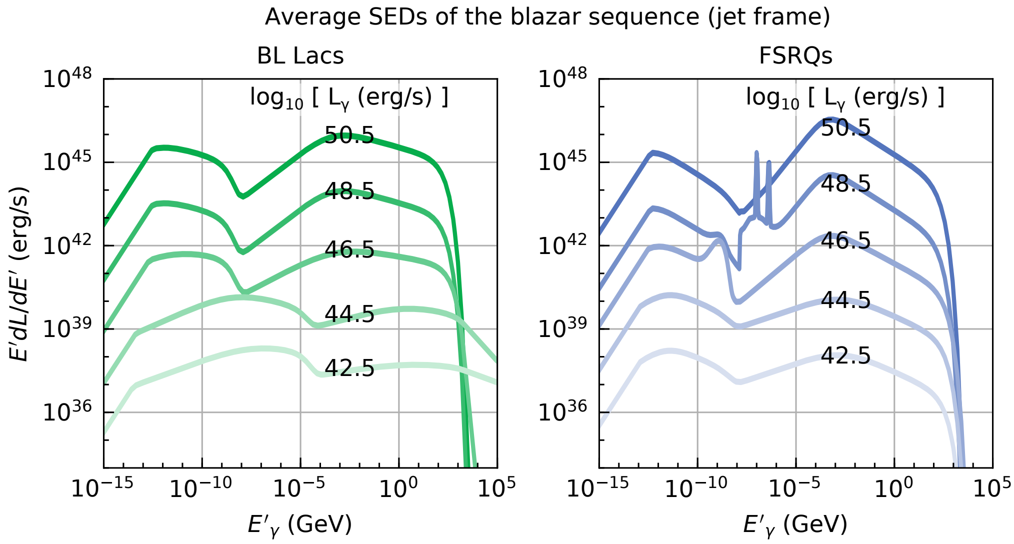

Figure 5.8.: The average SEDs predicted by the blazar sequence, depending on energy in the jet frame, given a K-corrected gamma-ray luminosity (labelled as $\log_{10}\,[L(\mathrm{erg/s})]$). Optical lines for FSRQs are added on top, and largely overlap for the relevant spectra. Figure adapted from [195].

More immediately, the peak frequency also affects observations in the X-ray band, as shown in fig. 5.7a.

Among the blazars seen and classified by Fermi-LAT (see section 4.4), the vast majority of FSRQs are LSP, while BL Lacs are broadly distributed between the classes. This distribution is shown in fig. 5.7b.

5.5.3. Blazar sequence

As seen in the previous subsection, blazar SEDs vary in shape. The **blazar sequence** describes the finding that this is correlated to their intrinsic luminosity. A recent version [190] parameterizes it separately for BL Lacs and FSRQs, in dependence of the K-corrected luminosity§ in the 0.1 GeV to 100 GeV band and fits the parameters to Fermi blazars. The resulting evolution is shown in fig. 5.8.

The previously mentioned BL Lacs with a larger ν^S_{peak} are rather found at lower luminosities, while FSRQs do not shift their synchrotron peak like this. For both classes, the second hump of all blazar SEDs grows faster than the first,

§the emission within a chosen energy band in the source rest frame, calculated from the redshifted spectrum in the observer rest frame [196] but to a lesser extent in BL Lacs [190].

Note that the blazar sequence only describes average SEDs within certain blazar populations. Extending it to all unresolved blazars may create tensions with the observed diffuse gamma-ray background [197].

5.6. Neutrino production

Multi-wavelength astronomy has been successful in studying certain aspects of blazars such as their morphology. Multimessenger astronomy on the other hand is required to clear up the particle physics happening inside, and with that the question whether they are acceleration sites for primary cosmic rays (see chapter 1). Indications for this fact have already been given in the form of the presence of shocks (section 5.4.2), ionized plasma (section 5.3.3), and high-energy photons (section 5.3.6), but the ultimate piece of evidence would be the neutrinos produced when these primary cosmic rays react with the known matter and photon fields in this environment.

As later described in section 9.2.1, the predicted neutrino emission depends on the low-energy photon fields and the amount of protons injected. Different geometries of the emission region have also been proposed, where the simpler models work with a spherical region somewhere in the jet [195], and more sophisticated ones structure the jet as a whole [198, 199].

Finally, more advanced models actually account for variability, which is a key feature of blazars, and describe the dynamics involved. This results in a more detailed prediction of the neutrino arrival times [200]. Without assuming such a model, the simplest approximation is that the cause of blazar variability propagates into all relevant particle populations and photon fields simultaneously, so that neutrino arrival times are proportional to the electromagnetic high-energy emission.

The previous chapter 4 showed how Fermi-LAT can monitor source variability. The present chapter introduced the fact that particularly blazars, the most common source in the Fermi-LAT sky, exhibit flares. The following chapter 6 therefore discusses how these blazar flares can be observed in Fermi-LAT **lightcurves**. Combining these with the physical principles which make blazars candidates for multimessenger observation then leads to the analysis formalism in the later chapter 7.

"The greyhound shuffle is not just a
dance, kid! It's a state of unrest."

Dave Chambers and Bob McGlynn [201]

6
Gamma-ray Lightcurves

A lightcurve refers to a measurement of a (photon) flux at multiple periods in time. This type of measurement can reveal the features in sources' time variability. The Fermi Large Area Telescope (see chapter 4) is able to scan the entire sky every 3h and can provide lightcurve measurements for many gamma-ray sources.

This chapter describes the lightcurves used as the input to the formalism of section 7.1.2. The fundamental method by which they are derived from Fermi data (see chapter 4) is sketched in section 6.2. The likelihood method for lightcurve calculation presented here differs from the method used in previous analyses which is explained in section 6.1. The way lightcurves are post-processed and smoothed is detailed in section 6.3.

Author's contribution: The latter section describes an optimization of the smoothing method which was my main contribution to this chapter. I also produced and studied lightcurves of TXS 0506+056 according to the old method, as described in the following.

6.1. Lightcurves from aperture photometry

Previous IceCube analyses looking for correlation with Fermi-LAT lightcurves [202, 157, 203, 204, 205] used lightcurves obtained by a method called aperture photometry, which is essentially a cut-and-count analysis. The procedure goes as follows:

1. Events are selected within a circular aperture of e.g. 2° around a given source location.
2. An energy threshold $E_{\min}$ of e.g. 100 MeV is applied to the data, to improve the angular resolution.
3. The remaining photon events are counted N within each time bin of e.g. 1 day.
4. The detector response, source location, observation time, energy range and an assumed spectral index of e.g. 2.4 are combined to calculate an exposure e per time bin.
5. The flux estimate is then given by $\Phi = N/e$, and its uncertainty $\sigma = \sqrt{N}/e$ from Poissonian statistics.

There is however a basic limitation of aperture photometry which is that the flat circular aperture can not distinguish between photons from the source of interest (SOI), a neighbouring source, or the diffuse background. This leads to photon contamination that depends on the PSF of the detector (see section 4.3.3). This limitation of the method became evident in the case of TXS 0506+056 and the neighbour source PKS 0502+049, located 1.2° away [206] i.e. within the standard aperture radius. The latter had flares in 2014-15, making it brighter than TXS 0506+056 itself. These can be seen contaminating the TXS 0506+056 lightcurve in fig. 6.1.

Any attempt to reduce the contamination problem induces additional complications, because of the common power-law photon spectrum and how the PSF widens towards lower energy with $\Psi^{(68\%)}(E_\gamma = 100\,\mathrm{MeV}) \approx 5°$ (see section 4.3.3). For example narrowing the aperture (e.g. to 1°) can only remove contamination to a limited extent while it also removes photons from the SOI. On the other hand, increasing the energy threshold (e.g. to 800 MeV) to improve the angular resolution is more effective, but it reduces the photon statistics significantly leading to uncertainties on the flux estimates. Even then, some contamination remains, as shown in fig. 6.2, and its level is not trivial to estimate.

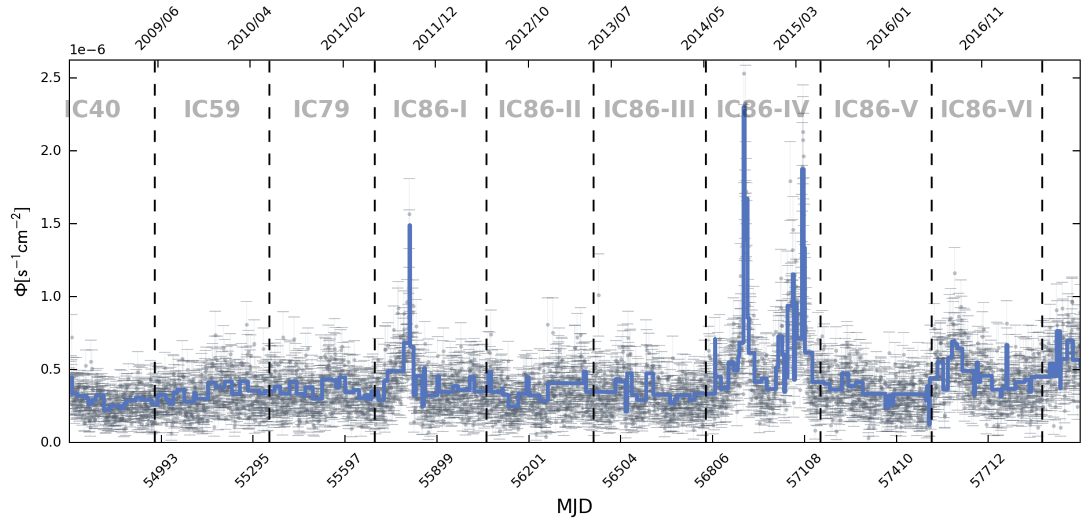

Figure 6.1.: Photometric Fermi-LAT lightcurve of TXS 0506+056 for the energy range of 100 MeV to 800 MeV assuming a photon spectral index of $\gamma = -2.2$. During the IC86-IV season, two flares from PKS 0502+049 contaminate this lightcurve and dwarf the actual biggest flare of TXS 0506+056 starting after IC86-VI, which consequently can not be seen clearly in this plot.

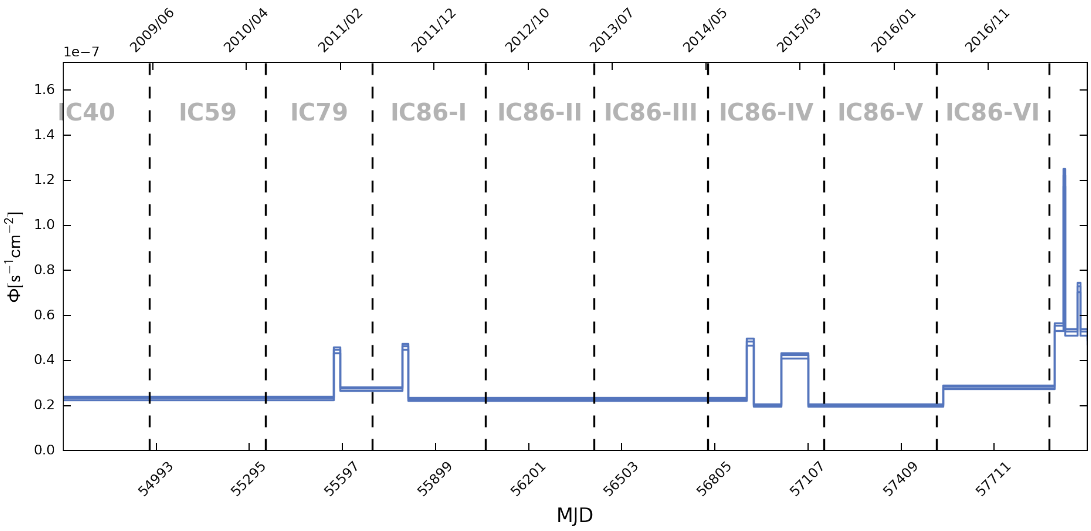

Figure 6.2.: Photometric Fermi-LAT lightcurve of TXS 0506+056 in the energy range of 0.8 GeV to 300 GeV and restricting the aperture to 1° to exclude PKS 0502+049. The relative magnitude of the contaminating flares is reduced, but still exceeds that of the 2017 flare. The drastically decreased photon statistics are reflected in the broader Bayesian blocks (see section 6.3).

6.2. Lightcurves from likelihood fits

Compared to aperture photometry, likelihood fits are a method to obtain lightcurves which is more sophisticated, sensitive and robust against contamination.

6.2.1. Method

These lightcurves are the same as in [125]. They are based on the same data as selected in section 4.3 but with an energy threshold applied.

The analysis for a particular SOI takes into account a $10° \times 10°$ ROI centred around it. The model of gamma-ray emission in this ROI consists of several components:

1. A diffuse background due to the Galaxy, unresolved sources, and misclassified charged particles, all described by standard templates*.
2. The SOI.
3. Other point sources. These sources are first obtained from the most recent Fermi catalogue 3FGL [156] (2017 at the time). However since the available data already covered more than 8 years, compared to the catalogue's 4 years, a search within the ROI is performed for additional sources that now surpass the detection threshold.

A binned likelihood, common to Fermi-LAT analyses, is defined with a spatial binning of 0.1° and 8 logarithmic bins per decade of energy, and evaluated to compute a TS which can be numerically maximized.

As a first step the full data set is fit for the SEDs of the point sources. This uses the 3FGL/3FHL catalogue's choice for the source's SED parameterization of power law, power law with exponential cutoff, or log-parabola.

Then the data is divided into time bins. Within each, the variable parameters are the flux normalization of the SOI and those sources within 3° of it. This radius is close enough for the PSFs to overlap in some cases as ~2.2° is the average distance between two 3FGL sources for $|b| > 10°$ [156]. If a neighbouring source flares, the likelihood statistically separates these photons with the help of their distribution in space and energy. The neighbouring source's flux fit

*https://fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html then absorbs the flare and ideally no gamma-ray excess remains which would influence the flux fit of the SOI. This way, the method becomes robust against contamination.

6.2.2. Parameters

For each time bin this procedure results in a maximal TS, at a corresponding best fit flux normalization Φ of the SOI. Due to the limited number of detected photons, this has a statistical uncertainty σ which is computed using the `MINOS` procedure in the `Minuit` package [207, 208].

The energy threshold and time binning need to match the available computing resources. In this context, the energy threshold controls the number of photons which have to be divided into energy and space bins. The time binning controls the number of likelihood maximizations that need to be performed. Therefore, to feasibly compute the lightcurves of all 2,254 Fermi sources in chapter 9, a 28-day time binning and 1 GeV energy threshold were chosen. For TXS 0506+056, the lightcurve in chapter 8 instead used a 7-day binning and a 300 MeV energy threshold. These parameters also affect properties of the lightcurve. Those are

- the **time resolution**; a 7-day flare averaged over a 28-day bin would be underestimated in brightness and overestimated in duration. Due to statistical uncertainties, the best possible time resolution depends on the brightness of the source, which itself is variable. Therefore adaptive binning techniques exist [209], though these lightcurves do not use them to select the bins subject to the likelihood fits. We instead chose to apply a variable-width smoothing afterwards, as described in the following section 6.3.

- the **purity** due to the narrower PSF at higher energies; one might raise the energy threshold from the 100 MeV commonly used for point source analyses up to 300 MeV in order to improve the likelihood's source discrimination.

- the **uncertainties** σ due to the photon statistics per time bin; these pull the parameters in the opposite direction of the previous two criteria. Depending on the source, with a 1 GeV energy threshold the fit in a 7-day bin

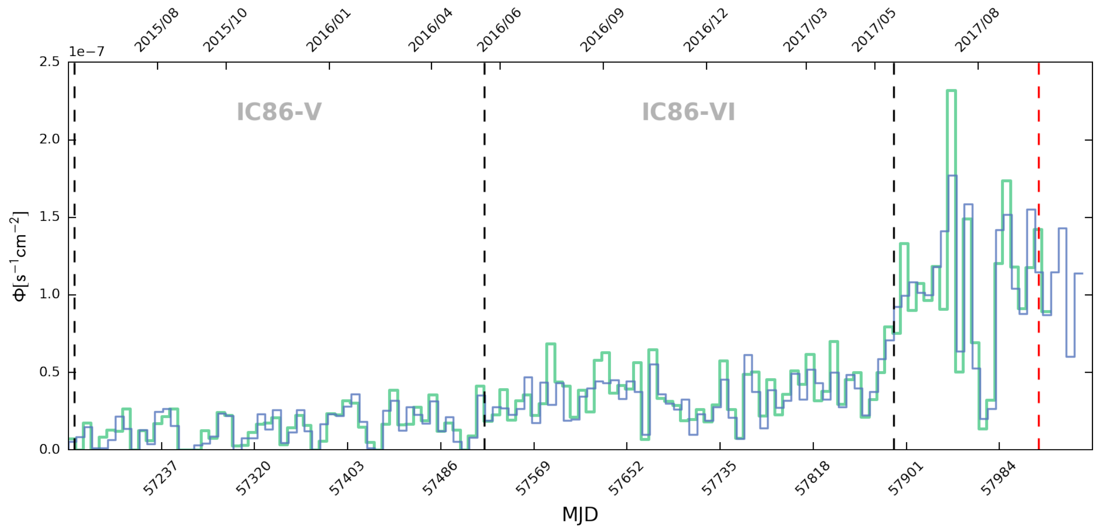

Figure 6.3.: Two LLH Fermi-LAT lightcurves of TXS 0506+056 during the last 2.5 years of data-taking, analysing the same photon data with the same method, but in different time bins. Compared to the green, the blue is extended by four weeks and 1 day out of phase.

> would have too large statistical uncertainties to be useful without additional smoothing. To make better use the lightcurve as-is, one could then choose 28-day bins or a 300 MeV threshold.

For quality reasons, only the fits of the 1 GeV lightcurves with TS > 0.1 are kept from this point on. Those monthly bins which fail this criterion are treated as gaps in the normal gamma-ray observation.

6.3. Bayesian block smoothing

6.3.1. Motivation

Figure 6.3 shows two lightcurves taken of TXS 0506+056, both with $E_\gamma > 300$ MeV and in 7-day bins. However one was extended by a month, and coincidentally starts 1 day out of phase from the other. The large difference between each pair of bins suggests that there are large statistical fluctuations. This is of course supported by the uncertainties estimated in the likelihood fit.

We do not wish to transfer these fluctuations onto the signal hypothesis, and therefore need to insert a smoothing step before using the lightcurve in the likelihood method.

Since blazars can in principle also exhibit physical variability on such short time-scales as those where the lightcurve is dominated by statistics (see section 5.4.4), there are two soft requirements on the smoothing.

1. If this variability is weak, it will be entirely hidden. In order to exclude it as a source of systematic uncertainty on the limits, the time PDF should remain flat and only change on the existing bin boundaries, i.e be a rebinning of the original lightcurve.

2. If it is strong enough, it will be apparent and the smoothing should not lose this information. This is not possible with rebinning at a constant scale, and so it needs to be adaptive to the magnitude of fluctuations relative to the uncertainties.

The Bayesian block algorithm [210] fulfils both these requirements, and has long been used in previous, similar analyses [204]. In this work I use the implementation from the `astropy` software library [211, 212].

6.3.2. Method

In one of its variants the algorithm operates on a lightcurve as in section 6.2, consisting of bins with each a flux measurement and associated uncertainty.

It partitions this lightcurve into a series of so-called blocks. It evaluates a likelihood function which for each block has a factor of the Gaussian likelihood that the bins which it comprises are statistically consistent with a constant flux at the level of their (weighted) average. Each block also adds a penalty term e^{-p}, and so p is a parameter which can be used to optimize the rebinning strength via the total number of blocks.

The algorithm starts with a block which contains only the first bin of the lightcurve. It then iteratively constructs the blocks by

1. extending the most recent block with the next bin;

2. evaluating the likelihood given this extended block
    - as a whole;

- split into two blocks at any of its bin edges;

3. making the choice to grow or split which maximizes the likelihood.

For the bins contained in each block I compute the mean of their fluxes Φ, without taking their uncertainties σ into account. This is because the latter can be approximated with a Poissonian model $\sigma^2 \propto \Phi$ (see fig. 6.4), which then predicts that a mean weighted by $\frac{1}{\sigma^2} \propto \frac{1}{\Phi}$ would be biased towards a lower value than the true mean.

The resulting block lightcurve strikes a balance between resolving variability during bright parts of the lightcurve, and smoothing out statistical fluctuations during its faint parts. This makes them particularly suited for the typical blazar lightcurve anatomy of quiescence and flares (see section 5.4.4), without requiring any model of their variability.

6.3.3. Optimization

As already mentioned in section 6.3.2, the Bayesian blocks have a tunable parameter, p, which describes the rebinning strength. This parameter was optimized separately for the 7-day TXS 0506+056 lightcurve (see section 6.2.2) as well as the monthly blazar lightcurves, however for the sake of simplicity this subsection will only use the former as an example.

The optimization proceeded according to a predefined criterion, based on the concept of a type-I error. In this scenario,

- the null hypothesis is that the source in question emits a steady flux,
- the test statistic, q, is the duration (relative to the entire lightcurve) for which the blocks are inconsistent with this flux,
- which is defined as lying outside the central 68% interval of the lightcurve bin fluxes.

The threshold above which we define the block lightcurve to contain a "false flare" is $q_{\mathrm{thres}} = 1\%$. While $q_{\mathrm{thres}} = 0\%$ was considered, in this case the probability of a false flare at any time increases with the lightcurve duration.

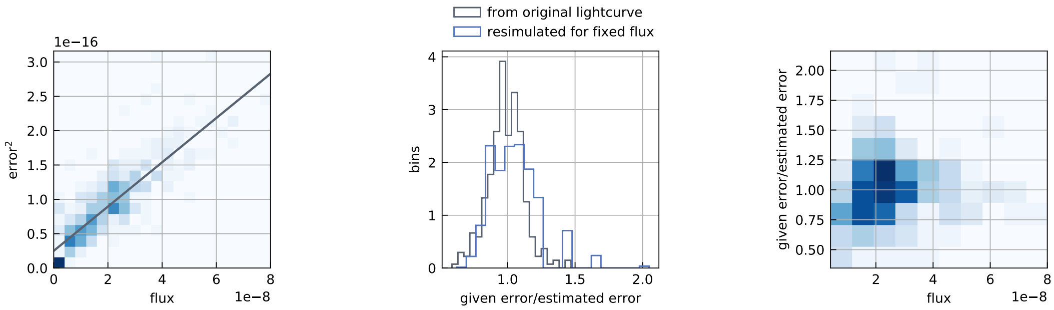

(a) A two-dimensional histogram of flux fits and corresponding squared errors, which show a linear correlation represented by the linear regression in grey.

(b) The given error relative to its estimation from the linear regression, in grey for the same set of bins and in blue for a resimulated lightcurve.

(c) The same relative error vs. flux, for which the correlation is weaker than in a), correlation coefficient 0.1.

Figure 6.4.: Statistics to characterize the behaviour of the likelihood fit from its outcome on the 482 bins of the example lightcurve.

The optimal p is chosen so that if the null hypothesis is true, a lightcurve has a probability of 32%[†] to contain a false flare. We call this type-I error the false flare rate (FFR).

To determine the FFR, we generate an ensemble of pseudo-lightcurves according to the null hypothesis using a toy Monte Carlo (MC). The latter is constructed to match the characteristics of the lightcurve to which the optimized Bayesian block rebinning will be ultimately applied. For this purpose, the lightcurve is examined as a set of best fit fluxes Φ and associated uncertainties σ. Figure 6.4a exhibits a linear correlation of

$$\sigma^2 = a + b\Phi, \tag{6.1}$$

approximately as would be expected for a rescaled Poissonian or from the central limit theorem [134, sec. 4.1]. A linear regression estimates a and b. In order to simulate the lightcurve corresponding to a true flux Φ_{true} we use this linear dependence to predict the error σ_{true}, thereby sampling a random Φ_{sim} for each bin. We invert the relation from aperture photometry between flux Φ, uncertainty σ, photon count n and exposure e

† $32\% = 1 - 68\%$ where the latter is the central 1σ interval of a normal distribution

$$\Phi = \frac{n}{e} \tag{6.2}$$

$$\sigma = \frac{\sqrt{n}}{e} \tag{6.3}$$

in order to calculate an expected photon count λ and an exposure e_{true}:

$$\lambda = \frac{\Phi_{\text{true}}^2}{\sigma_{\text{true}}^2} \tag{6.4}$$

$$e_{\text{true}} = \frac{\Phi_{\text{true}}}{\sigma_{\text{true}}^2} \tag{6.5}$$

A trial n_{sim} according to the Poissonian $\mathcal{P}(n_{\text{sim}}; \lambda)$ is converted into $\Phi_{\text{sim}} = n_{\text{sim}}/e_{\text{true}}$, while $\sigma_{\text{sim}} \equiv \sigma_{\text{true}}$.

The toy MC method is only valid so far as the spread in fig. 6.4a can largely be attributed to the simulated statistical variations, of a magnitude according to the linear regression. In order to test this, $\frac{\sigma}{\sqrt{a+b\Phi}}$ is computed both for the TXS 0506+056 lightcurve and a simulated lightcurve. Comparing the two distributions in fig. 6.4b shows that the former does not have a wider spread than the latter, and correlation with the flux is reduced (see fig. 6.4c). The method is therefore accepted.

Further parameters for the MC are lightcurve duration and binning, which we choose identical to the TXS 0506+056 lightcurve, and the level of the simulated steady flux, which we set at the 25th percentile of its 7-day fluxes. Using a grid scan of p and interpolating, we find the FFR=32% for $p = 3.05$, see fig. 6.5a. Fixing this, we separately also vary the lightcurve duration and steady flux from their baseline values to examine whether they are stable. Figure 6.5b and Figure 6.5c show that this is approximately the case, and particularly so for the duration when compared to the alternative 0%-definition of q. The final rebinned lightcurve is shown in fig. 6.6. Applying the same smoothing procedure to the other, phase-shifted lightcurve of fig. 6.3 reveals the statistical fluctuations averaging to the same flux during longer blocks. The fluctuations can affect the placement of block boundaries, which is noticable during more variable periods. However, the overall shapes are compatible within the unbinned lightcurve's uncertainties and show the same set of features.

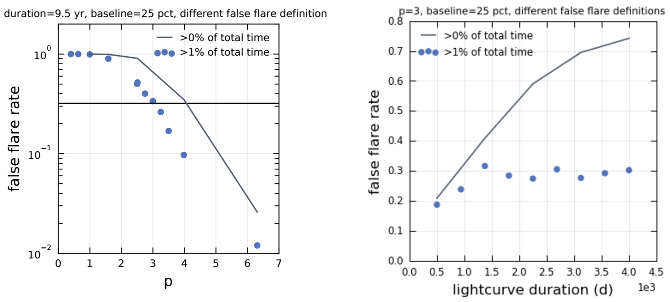

(a) Varying the rebinning parameter p itself, whose optimized value is thereby determined.

(b) Varying the lightcurve duration as a check of robustness, which the tests using the older FFR definition fail.

(c) Varying the baseline flux level as a similar robustness check, which in this case mostly tests the lightcurve simulation itself.

Figure 6.5.: Blue circles: the false flare rates (y-axis) among sets of simulated block lightcurves which depend on three parameters (x-axes). Grey lines: an older definition of the false flare rate, retained in the plot to motivate the current one.

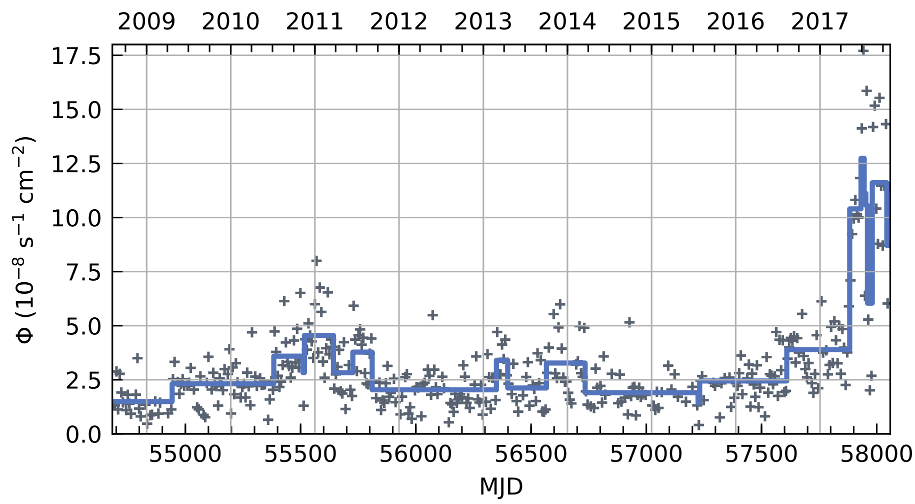

Figure 6.6.: The result of applying Bayesian blocks to the example lightcurve with the optimized value of 3.05 for the rebinning strength parameter. The weekly bins are shown with their uncertainties as grey crosses while the Bayesian block representation is shown as a blue curve.

Two further methods to select p have been described in [210, sec. 2.7], one of them based on a similar principle but making use of an analytical calculation.

"If minimizers don't work it's not your fault."
Jonas Verhellen

7 Likelihood Method

When neutrinos are emitted by astrophysical objects and measured in IceCube, their arrival directions in equatorial coordinates will cluster around the direction of the source. The background due to atmospheric muons and neutrinos on the other hand (see section 2.1)covers the entire sky. Due to the detectors response, which includes the Earth absorption of muons and the event selection, the background varies smoothly with the zenith angle. Due to the Earth's rotation, it is uniform in right ascension. In this thesis we search for a time-dependent neutrino signal and hence another dimension is added to the standard, time-integrated point-source analysis [213, 214].

This chapter describes the method used for the time-dependent searches. The likelihood description is given in section 7.1 together with the signal and background probability density functions. This is inverted when generating a data set which may contain signal from a particular hypothesis, according to the methods of section 7.2. Finally, the application of the method is a likelihood ratio test described in section 7.3, where the likelihood and generation are also combined to compute sensitivities and limits.

Author's contribution: For this chapter, I first re-implemented the time-dependent likelihood method in a new framework, starting from an existing re-implementation of the time PDF. I implemented the time MC according to an established method, but produced its input in a faster way enabled by current IceCube DAQ metadata.

Rather than only using these simulated times for background as previously, also my signal injection reflects IceCube data-taking. That allowed me to extract MC integrals of the time PDF, which let me normalize it over livetime. Therefore I could already develop and characterize the analysis of chapter 8 when the data set of chapter 3 was not yet complete. For chapter 9, I extended the time-dependent likelihood to stacking, implementing this novel method in the same analysis framework which was already capable of time-integrated stacking. As a particular aspect of this analysis, I needed to define the common threshold parameter. To apply the method in feasible time, I composed a numerical maximizing procedure and parallelized the existing method to calculate limits and sensitivities. To improve the discovery potential calculation, I introduced a new extrapolation method for the background TS distribution.

7.1. Likelihood fit

In this work we use an unbinned maximum likelihood method to search for a neutrino signal from a set of flaring blazars. The method used here is an extension of the likelihood method applied to point sources in IceCube [213, 214, 62]. It represents data as a set of N events with observables $\{\vec{x}_i\}$. The probability that one event follows a given hypothesis is given by the PDF $\mathcal{P}(\vec{x}_i;\vec{a})$, determined by model parameters $\vec{a}$. The likelihood that the data is distributed according to the hypothesis is then given by:

$$\mathcal{L} = \prod_{i=1}^{N} \mathcal{P}(\vec{x}_i;\vec{a}). \tag{7.1}$$

Maximizing the likelihood gives an estimate $\widehat{\vec{a}}$ that better describes the data. This general method is applicable to any estimation problem, and in contrast to binned or cut-and-count methods uses all available information [134].

The signal hypothesis H_1 can be described with a sum of both a background and a signal PDF. The latter is weighted to reflect the signal strength, which is expressed as the expected number of signal events n_S in the data set of size N, where $n_S \ll N$:

$$\mathcal{P}(\vec{x}_i,\vec{a}) = \frac{n_S}{N}\mathcal{S}(\vec{x}_i,\vec{b}) + \frac{N-n_S}{N}\mathcal{B}(\vec{x}_i). \tag{7.2}$$

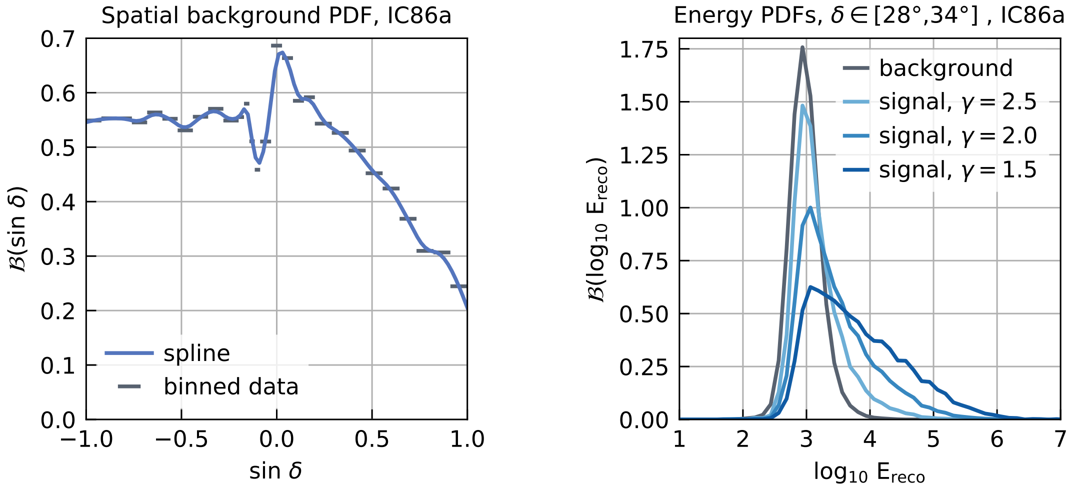

(a) The background binned in $\sin\delta$ (grey lines) and the spline evaluated as the PDF (blue).

(b) The distribution of $\log_{10} E_{\mathrm{reco}}$ in one δ bin, convoluted with a smoothing kernel but not yet splined. Grey: background. Blue: signal at different spectral indices.

Figure 7.1.: Background PDFs for the IC86a sample.

Here, $\vec{b}$ are the parameters apart from n_S which characterize the unknown signal. The background is known and so its PDF has no parameters.

In the case of point source searches, the measurements $\vec{x}_i$ are the observables of chapter 3 which include the reconstructed event direction $\vec{r}_i = (\mathrm{RA}_i, \delta_i)$ and estimated angular uncertainty σ_i, its reconstructed energy E_i, and arrival time t_i. At $n_S = 0$, one speaks of the null hypothesis H_0. In the search for neutrino sources, the hypothesis space is divided between acceptance and rejection of H_0 primarily along the parameter of signal strength.

7.1.1. Background PDF

The background PDF $\mathcal{B}$ should describe the atmospheric muons and neutrinos observed by the detector. Since data contains only a negligible fraction of signal*, it is appropriate to approximate the background PDF with the empirical distribution of experimental data. We use a definition that factorizes into a spatial, energy and time PDF:

*see e.g. [61] which studied ν_μ in IceCube at similar energies as the data set of chapter 3

$$\mathcal{B}(\vec{r}, E, t) = \mathcal{B}_r(\delta) \cdot \mathcal{B}_E(E; \delta) \cdot \mathcal{B}_t \tag{7.3}$$

IceCube is located at the geographic South Pole and rotates with the Earth. Therefore, this definition reflects the assumption that the backgrounds do not depend on right ascension, and the spatial PDF $\mathcal{B}_r$ is taken from the declination distribution of data. The distribution of events is stored in a histogram where bins increase in density near the horizon to better describe the rapid change in rate due to the transition from atmospheric muons to an atmospheric neutrino dominated background. Statistical fluctuations are smoothed by splines, as shown in fig. 7.1a.

The same procedure produces the background energy PDF $\mathcal{B}_E$ (fig. 7.1b). Since the detector acceptance depends on the declination, the energy PDF consists of a two-dimensional map with reconstructed energy in one axis and reconstructed declination in another. This map is normalized for each declination bin, convoluted with a smoothing kernel, and also splined. The energy bins span the range of the sample, so that some of them are empty due to statistical fluctuations.

The background time PDF $\mathcal{B}_t$ is flat, which exploits the uniform detector performance during the selected good uptime (see chapter 3), but also neglects the seasonal modulation, apparent in fig. 3.1 (see also section 7.2.2). The PDF can either be normalized over the sample's livetime (as in chapter 8) or the time span of the sample (as in chapter 9), to match the choice for the signal time PDF described in the following subsection.

7.1.2. Signal PDF

The signal PDF $\mathcal{S}$ is described by the source position, spectral shape and time distribution. Here, we again use a definition which factorizes into spatial, energy, and time PDFs:

$$\mathcal{S}(\vec{r}, E, t; \vec{b}) = \mathcal{S}_r(\vec{r}, \sigma) \cdot \mathcal{S}_E(E; \delta) \cdot \mathcal{S}_t(t) \tag{7.4}$$

Events coming from a point source of neutrinos will cluster around the source direction following IceCube's PSF, which is correlated with energy. Instead of explicitly characterizing the energy dependency, the likelihood uses the event-wise angular uncertainty estimator σ and assumes symmetry so that it only

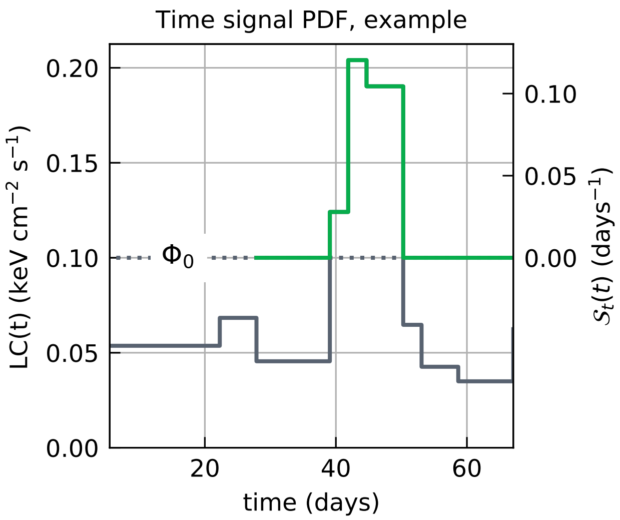

Figure 7.2.: An example lightcurve (left y-axis, grey) along with the corresponding signal time PDF (right y-axis, green).

depends on the angular distance between the source and the event $\Delta\Psi_i = |\vec{r}_i - \vec{r}_s|$:

$$\mathcal{S}_r(\vec{r}_i; \vec{r}_s) = \frac{1}{2\pi\sigma_i^2} e^{-\frac{\Delta\Psi_i^2}{2\sigma_i^2}} \tag{7.5}$$

Signal neutrinos are further distinguished from background by their harder energy spectrum, as illustrated in fig. 7.1b. The procedure to build the energy PDF is identical as for background but using simulation events weighted according to a power-law spectrum. PDFs are computed for a set of spectral indices, and locally interpolated in-between the selected values.

For the time domain, our hypothesis is that the timing of neutrino emission traces the flaring state of the blazar and that the quiescent gamma-ray emission is steady at a level of Φ_0 or **threshold**. The time PDF for neutrino emission from a blazar is therefore derived from the lightcurve LC according to the definition:

$$\mathcal{S}_t(t_i, \Phi_0) = \frac{\max(0, LC(t_i) - \Phi_0)}{\int \mathrm{d}t \max(0, LC(t_i) - \Phi_0)} \tag{7.6}$$

This is illustrated in fig. 7.2. The time PDFs for chapter 9 are normalized over the time span of the sample. The non-stacking analysis of chapter 8 more correctly normalizes its time PDF taking into account the detector dead-times due to run transitions and runs excluded from the sample. The method of section 7.2.1 provides random times to calculate the necessary integral with a Monte Carlo method.

It is very difficult to identify the level of a lightcurve above which the flaring state begins and neutrinos are emitted. In principle neutrino emission could happen during only the most intense flares, or on the other hand, follow the full lightcurve profile which also includes the quiescent part. The latter case corresponds to $\Phi_0 = 0$, implying that all gamma-ray emission of the source is due to a hadronic mechanism which yields a proportional neutrino flux. Therefore the threshold is used as a free parameter in this analysis, similar to γ, able to enhance the contrast between the signal and background PDFs.

7.1.3. Combining samples

The likelihood combines different samples j by multiplying their respective likelihood terms. Each likelihood receives the same signal parameters γ and Φ_0 as these are common to all periods. The signal strength n_S, however, is split between the samples, according to their exposure to the signal:

$$\mathcal{L}(n_S, \vec{r}_s, \gamma, \Phi_0) = \prod_j^{\text{samples}} \mathcal{L}_j \left(n_S^j, \vec{r}_s, \gamma, \Phi_0 \right) \tag{7.7}$$

$$n_S^j = n_S \times w_j \left(\vec{r}_s, \gamma, \Phi_0 \right) \tag{7.8}$$

where $w_j = \omega_j / \sum_j \omega_j$, and ω_j is the **exposure** of the hypothetical source neutrino flux $\frac{\mathrm{d}\Phi}{\mathrm{d}E}(t, E)$ to the sample effective area and livetime (see fig. 3.5, chapter 3). This is defined as an integral over true neutrino energy E (not reconstructed energy, as in the likelihood), and over the time spanned by the sample $[t_j^-, t_j^+]$:

$$\omega_j(\vec{r}_s, \gamma, \Phi_0) = \int_{t_j^-}^{t_j^+} \mathrm{d}t \frac{\mathrm{dlivetime}_j}{\mathrm{d}t} \int \mathrm{d}E A_{\mathrm{eff},j}(E, \vec{r}_s) \frac{\mathrm{d}\Phi}{\mathrm{d}E}(t, E; \vec{b}). \tag{7.9}$$

Following the hypothesis that the flux factorizes into a time PDF $P(t)$ and a spectrum $\frac{\mathrm{d}\Phi}{\mathrm{d}E}(E)$, so does the integral:

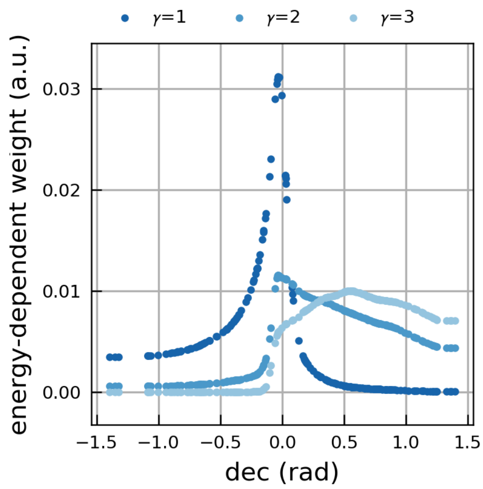

Figure 7.3.: The energy integral within ω, see eq. (7.10), normalized over a set of sources (points on the declination axis), given different power law indices (colours).

$$\omega_j(\vec{r}_s, \gamma, \Phi_0) = \left(\int_{t_j^-}^{t_j^+} \mathrm{d}t \frac{\mathrm{livetime}_j}{t_j^+ - t_j^-} P(t; \Phi_0) \right) \times \left(\int_{E_j^-}^{E_j^+} \mathrm{d}E A_{\mathrm{eff},j}(E, \vec{r}_s) \, \frac{\mathrm{d}\Phi}{\mathrm{d}E}(E; \gamma) \right). \tag{7.10}$$

Simulation events and their weights (see section 2.6) are used to compute the energy integral in a Monte Carlo method. The results of this calculation in fig. 7.3 show that the power-law index γ changes the contribution of the signal originating from different hemispheres. The time integral on the other hand is evaluated analytically. This assumes that, within each lightcurve block during the sample, the term $\frac{\mathrm{dlivetime}}{\mathrm{d}t}$ averages to the same value.

7.1.4. Stacking

A stacking analysis looks for the cumulative signal distributed between multiple sources, k. The stacking likelihood for the sample, j, therefore sums their respective signal PDFs:

$$\mathcal{S}_j(\vec{b}) = \sum_k w_{jk}(\vec{b})\mathcal{S}_{jk}(\vec{b}) \tag{7.11}$$

where $w_{jk}(\vec{b}) = \omega_{jk}(\vec{b}) / \sum_k \omega_{jk}(\vec{b})$ is proportional to the same exposures ω as when combining samples. Now the exposure for different sources, and thus w_{jk}, depends not only on detector properties, but also on the assumption on the brightness of the sources. This is often derived from a theoretical model, expressed as a **weighting scheme**, which thus becomes part of the signal hypothesis for which the likelihood is designed and against which it is tested. As before with $\mathcal{S}_j$, the likelihood parameters also propagate to the normalized PDFs $\mathcal{S}_{jk}$, as well as to w_{jk}.

7.1.5. Common threshold parameter

The threshold parameter Φ_0 differs from the spectral index γ in so far as that there is no argument as to why it should not be specific to a blazar. This raises the question whether in the case of the stacking analysis the signal hypothesis should have each blazar's threshold as an individual variable parameter. The disadvantage of this direct approach is that the hypothesis space becomes highly multi-dimensional and practically intractable. First, the resulting multitude of limits would be difficult to interpret, even assuming the available time suffices to produce enough signal trials (see section 7.3.4). Second, a likelihood with this many degrees of freedom has more ways to fit background, which raises the discovery potential (see section 7.3.5). Finally, the numerical maximization (see section 7.3.2) would either take too much time, or have a too low success rate.

The stacking likelihood instead focuses on an attempt to decompose lightcurves into quiescence and flares. Similarly to [179], it does this on the base of the parameters q and σ_q, the estimated level of quiescent flux and its statistical or physical fluctuations respectively. The threshold for source k is defined as:

$$\Phi_{0,k} = \max(0, q_k + \tau\sigma_{q,k}) \tag{7.12}$$

We assume henceforth that the parameter τ has a common value for all blazars. Because no threshold is allowed to be negative, in the case of $\tau = -\infty$ we obtain

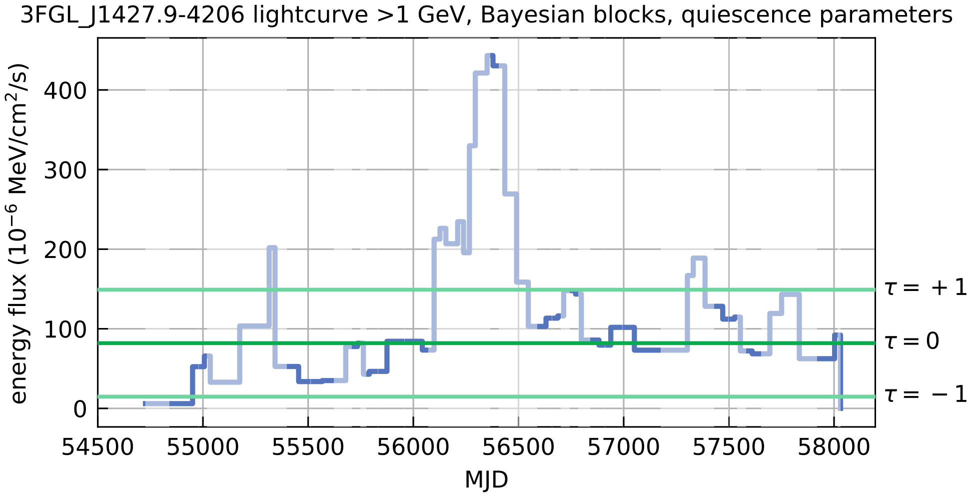

Figure 7.4.: The τ parameter and underlying quiescence parameters illustrated on a lightcurve (light blue) where the general definition of the latter applies, not requiring any edge cases. The quiescent state selected for this purpose is emphasized in darker blue. Its weighted mean and r.m.s. $q \pm \sigma_q$ are the green lines.

$\Phi_{0,k} = 0$ for all k.

In order to estimate q and σ_q from the Bayesian block lightcurve, we proceed as follows:

1. For each transition between two blocks, connect their centre points to calculate the derivative. Low values belong to steady stretches.

2. Assuming each transition takes half of the blocks surrounding it, find the median derivative value for which the lightcurve is steadier for half its duration. We consider the block halves adjacent to these transitions as part of the blazar's quiescent state.

3. Average the flux over this time period to obtain q.

4. Taking the same time period into account, calculate the r.m.s. flux, which then is σ_q.

An example for the result of these parameters for a particular lightcurve is shown in fig. 7.4.

7.2. Data set generation

Testing the likelihood performance requires a statistical ensemble of simulated data sets corresponding to a specific hypothesis. This will necessarily contain the background, which as in the likelihood is represented by the data once its observables have been partially decorrelated or **scrambled** to hide possible signals. To this, an **injector** adds a specific level of signal. Here, the injected signal follows the same hypothesis as in the likelihood. In the following, I explain how this signal injection and background scrambling are implemented.

7.2.1. Time Monte Carlo

Generated data sets use times sampled from the sample livetimes, i.e. the detector runs included in the data set, in order to simulate actual data-taking. To this end I devised a MC to produce times corresponding to a variable flux with constant spectrum. For each data sample, it proceeds as follows:

1. Construct the curve of cumulative livetime at a certain time, using the start time, stop time, and duration of each run in the sample.
2. Uniformly sample $N = \mathcal{O}(10^6)$ random values from the interval [0, livetime].
3. Map them to times $\{t_i\}$ using the inverse of the curve, and store the result.
4. Use the flux time PDF $P(t)$ to calculate weights $w_i = P(t_i)$ for the stored times, and make an accordingly weighted random choice from among them.

This numerical method is applicable to any definition of the time PDF, not requiring any analytical integrals or piece-wise defined functions/integrals. It is also meant to be more efficient than **rejection sampling**, another general method. Furthermore it lends itself to a Monte Carlo integral of a flux $\Phi(t)$ over the set of livetime:

$$\int_{\{\text{livetime}\}} \mathrm{d}t\ \Phi(t) = \frac{\text{livetime}}{N} \sum_{i=1}^{N} \Phi(t_i) \tag{7.13}$$

For a steady flux this exposure integral simplifies to a factor $\propto$ livetime.

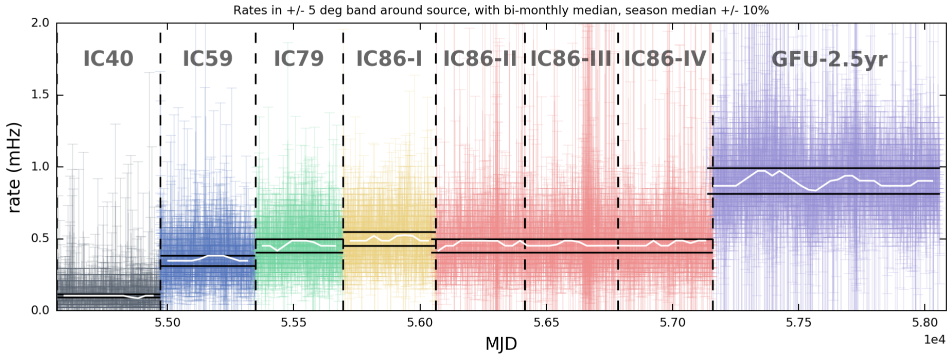

Figure 7.5.: Event rates in a declination band ±5° within TXS 0506+056 per run of the data set (coloured markers). Their bi-monthly running median (white line) reveals a seasonal modulation. The bounds of ±10% within the sample average are shown in black.

The method's downsides are that it is only accurate for time dependencies of scales $\gtrsim \frac{\text{livetime}}{N} \approx O(100\,\text{s})$, and that storing the weights requires additional space in memory for every additional time PDF.

7.2.2. Background scrambling

For background, the primary time-dependence is a seasonal variation. It affects both atmospheric muons and neutrinos, depends on declination, and its magnitude is limited to < 10% [127, 126]. While the analysis-level data set in principle reflects the combined $\mu + \nu$ seasonal modulation depending on declination, it does not provide enough statistics to derive an accurate empirical model (see fig. 7.5). This makes deriving a general model for the seasonal modulation at any source declination less than trivial. In light of the modulation's limited magnitude, we therefore assume that the background times follow a uniform distribution. As a corollary, this distribution does not depend on declination and so a new time is generated for each event with the same time Monte Carlo regardless of its other observables. This also replaces any signal-related time distribution which the data may have had before.

Each event keeps its local coordinates, i.e. zenith and azimuth (see section 3.3.2), while being assigned the new time. The approximation is made that up to this time, the detector has been rotating precisely and uniformly around its Z axis, back to the epoch (J2000). At that epoch, the local coordinates then point back to new equatorial coordinates, where due to the approximation the declination is unchanged. The right ascension meanwhile is affected by the time scrambling. Any space-time clusters that may have been present in the data are thus scattered.

The scrambling changes none of the other observables, which are declination, energy, and angular uncertainty. This way, their distributions and correlations are preserved, so that no further assumptions on the background need to be added to the experimental data. In particular, the background energy PDF near the source declination is represented by random subsets of the same energies.

7.2.3. Injector

To a scrambled background realization, the injector adds signal events from simulation. As the characteristics of the sample change with zenith, it selects simulation events in a declination band around the source, with a typical width being $\sin(1°)$ on a $\sin(\delta)$ scale. This produces $\approx \frac{N_{\mathrm{total}} \Delta \sin\delta}{2}$ candidate events for injection.

As their number, true energy and interaction probability (see section 2.6) are known, these candidates can be reweighted to a power law flux $E^{-\gamma_{\mathrm{inj}}}$ and sampled with these weights, up to a number n_{inj}. By rotating their true directions onto the source, their reconstructed directions then form a sampling of the PSF (see section 3.2.4) for this source and spectrum. Finally, the time MC assigns times according to the time PDF as described in section 7.2.1.

The number n_{inj} is a Poissonian variation around a mean μ_{inj} which corresponds to a normalization of $\frac{\mathrm{d}\Phi}{\mathrm{d}E}(t, E)$, exposed to the sample livetime and effective area. The integral over time and energy is computed entirely with Monte Carlo methods (see eqs. (7.9) and (7.13)).

Injecting a signal from multiple sources (stacking) at once requires keeping track of the individual sources' normalization as in the likelihood. From this point, due to the additive property of the Poissonian distribution, it is enough to combine all candidates and their weights together, and $\mu_{\mathrm{inj}} = \sum_k \mu_{\mathrm{inj},k}$.

The samples in the data set do not overlap, in the sense that their livetimes or event selections prevent them from sharing events. Therefore they combine in exactly the same way as the stacking, using the respective sample's simulation.

7.2.4. Neutrino flux normalization

A steady signal flux $\Phi = \frac{\mathrm{d}N}{\mathrm{d}E\,\mathrm{d}t\,\mathrm{d}A}$ of a source at declination δ_{src} can be written as

$$\Phi_{\mathrm{steady}}(t, E) \coloneqq f \times \left(\frac{E}{E_0}\right)^{-\gamma} \tag{7.14}$$

where f is the flux normalization at a given energy, E_0, and has units

$$[f] = \frac{1}{\mathrm{cm}^2\mathrm{GeVs}} \tag{7.15}$$

The expected number of events for any flux is

$$\mu = \sum_{j\in\{\mathrm{samples}\}} \int_{\mathrm{livetime}_j} \mathrm{d}t \int \mathrm{d}E A_{\mathrm{eff},j}(E, \delta_{\mathrm{src}})\Phi(t, E) \tag{7.16}$$

$$\tag{7.17}$$

given the detector livetime and effective area. In the steady case, this is

$$\mu_{\mathrm{steady}} = f \times \sum_{j\in\{\mathrm{samples}\}} \mathrm{livetime}_j \times \left(\int \mathrm{d}E A_{\mathrm{eff},j}(E, \delta_{\mathrm{src}}) \left(\frac{E}{E_0}\right)^{-\gamma}\right). \tag{7.18}$$

If meanwhile the flux is time-dependent, it can be defined using a time PDF $P(t)$ as

$$\Phi(t, E) \coloneqq g \times P(t) \times \left(\frac{E}{E_0}\right)^{-\gamma} \tag{7.19}$$

where now the normalization g has units

$$[g] = \frac{1}{\mathrm{cm}^2\mathrm{GeV}}. \tag{7.20}$$

The expected number of events is then

$$\mu = g \times \sum_{j\in\{\text{samples}\}} \left[\left(\int_{\text{livetime}_j} \mathrm{d}t P(t) \right) \times \left(\int \mathrm{d}E A_{\text{eff},j}(E, \delta_{\text{src}}) \left(\frac{E}{E_0}\right)^{-\gamma} \right) \right]. \tag{7.21}$$

The terms are the global normalization factor g, and for each sample j the energy and time integral which the injector needed to compute (see section 7.2.3). Given a certain μ, this calculation makes it possible to calculate the flux, with the normalization

$$g = \mu \times \left[\sum_{j\in\{\text{samples}\}} \left(\int_{\text{livetime}_j} \mathrm{d}t P(t) \right) \times \left(\int \mathrm{d}E A_{\text{eff},j}(E, \delta_{\text{src}}) \left(\frac{E}{E_0}\right)^{-\gamma} \right) \right]^{-1} \tag{7.22}$$

In the stacking hypothesis, the signal hypothesis is

$$\Phi(t, E) \to \sum_k \Phi_k(t, E) = g \sum_k w_k P_k(t) \left(\frac{E}{E_0}\right)^{-\gamma}. \tag{7.23}$$

With the detector response depending on the source declination, this means a normalization factor

$$g = \mu \times \left[\sum_{k\in\{\text{sources}\}} \sum_{j\in\{\text{samples}\}} \left(\int_{\text{livetime}_j} \mathrm{d}t P_k(t) \right) \times \left(\int \mathrm{d}E A_{\text{eff},j}(E, \delta_k) \left(\frac{E}{E_0}\right)^{-\gamma} \right) \right]^{-1} \tag{7.24}$$

7.3. Likelihood ratio hypothesis test

7.3.1. Test statistic

The logarithm is computed of the ratio between the likelihood for a set of given parameters, and the likelihood given $n_S = 0$, where it is also degenerate with respect to all other signal parameters. This is maximized with respect to $\vec{a}$, resulting in an estimate $\hat{a}$. This is used to define the Wilks test statistic (TS):

$$\mathrm{TS} = 2\,\mathrm{sgn}(\hat{n}_S)\log\frac{\mathcal{L}(\hat{a})}{\mathcal{L}(n_S = 0)} \tag{7.25}$$

- If the data contained an over-fluctuation over background then $\mathrm{TS} > 0$, and the larger the value the better the fit to the signal hypothesis.
- If the best-fit $\hat{n}_S = 0$ then the likelihoods in the ratio are identical and $\mathrm{TS} = 0$.
- An under-fluctuation, where the examined portion of the data undershoots the background expectation, will result in $\hat{n}_S < 0$ and therefore $\mathrm{TS} < 0$, however in our case we are not interested in under-fluctuations and clip them to $\hat{n}_S = 0, \mathrm{TS} = 0$.

The hypothesis test consists in performing this maximization using the events of chapter 3 and comparing the resulting TS with a predefined threshold. Up to this point, the likelihood is developed blindly by using data scrambled as in section 7.2.2 in order to avoid biases or hidden trials.

7.3.2. Maximizing

Due to the complicated form of the likelihood, the maximization needs to happen numerically. Therefore the performance of the analysis hinges on the performance of the method used therefor. It is assisted by bounds on the parameters:

- $n_S \in [0, 1000]$ since we do exclusively look for over-fluctuations. We also do not need to fit for very strong signals as these would have been found in previous analyses of the same data, regardless of the specific signal hypothesis.
- $\gamma \in [0, 4]$ which is a wider range than the values expected from diffusive shock acceleration (see section 1.2.2).
- $\Phi_0 \in [0, \Phi_{\mathrm{max}}]$ the maximum of the considered lightcurve, due to the parameter's definition. This is in the case of a single source. The τ parameter used in stacking has an upper bound where the threshold surpasses the maximum block among all the lightcurves.

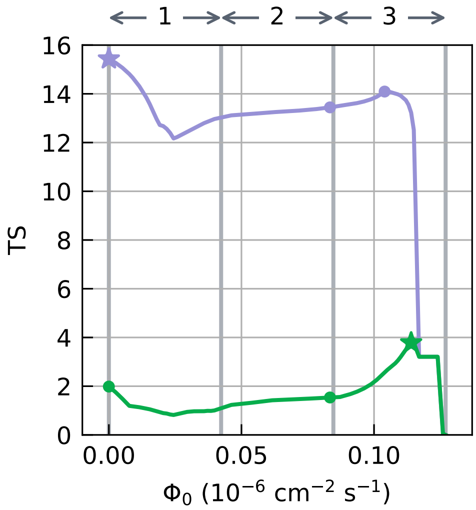

Figure 7.6.: Two examples for a profile of the TS landscape obtained by scanning in Φ_0 and numerically maximizing TS with respect to its other parameters. The star marks the true global maximum, while the circles are the local maxima within the three threshold sub-intervals.

Due to the definition of the time PDF in section 7.1.2, the test statistic is not continuously differentiable where it crosses over the level of a block in the lightcurve, and equivalently for τ. Therefore these analyses require minimizers which do not explicitly use the derivative. Furthermore the TS landscape is often bimodal in Φ_0, as seen in fig. 7.6, or multimodal in τ and a minimizer might converge on a local instead of a global maximum. If this happens often enough that the unblinding fit is also concerned, this directly affects the significance of the result and the limits set therefrom. No single minimizer managed to satisfactorily solve this issue, and so I devised a composite, multi-stage minimization scheme:

1. The Φ_0 or τ axis is divided into equal intervals. This way, each interval contains fewer, ideally one, local maxima. Each subsequent fit is limited to one interval.

2. A global-style minimizer such as `DIRECT` is run with a large tolerance, in order to robustly and quickly reach the vicinity of the maximum within the interval.

3. A local-style minimizer such as `MIGRAD` is run to precisely find the maximum, with a starting estimate provided by the global-style minimizer.
4. In some cases, this fails to converge, and a second attempt is made with another local-style minimizer like `COBYLA`.

The best fit between all intervals and stages is used as the result. This scheme was not completely applied to the stacking analysis, where the τ axis was still split but minimization limited to one stage of `MIGRAD`. This represents a compromise given the greater CPU time required to evaluate the test statistic, while still respecting its multimodality.

7.3.3. p-value

A trial is a fit on a data realization, according to section 7.2. The TS distribution of trials under a single hypothesis help interpret the TS of the unblinding fit. More specifically, its significance is defined as the probability to get a larger TS under the null hypothesis, i.e. with background trials where the data is only scrambled (see section 7.2.2) without adding signal events. This is also called the p-value. In a fraction $1 - \eta$ it is degenerate $p = \eta$ since under-fluctuations accumulate with TS = 0. From there, it decreases approximately exponentially with TS. Therefore it is handy to express it as the one-sided quantile of a standard normal distribution, e.g.

$$\begin{aligned} p &= 2.3\% && (7.26) \\ &= \frac{1}{\sqrt{2\pi}} \int_2^\infty \mathrm{d}x\, \mathrm{e}^{-x^2} && (7.27) \\ &\equiv 2\sigma \end{aligned}$$

7.3.4. Limits and sensitivities

Now trials are made where signal is added according to a specific hypothesis (spectral index, time PDF(s)) by the injector. The resulting TS distribution depends on the signal strength expectation μ. After an unblinding fit, the hypothesis test might decide that the result is background and accept the null hypothesis. In that case, a limit on μ given that particular hypothesis can be set at β = 90% confidence level (C.L.) by finding the $\mu^{(\beta)}$ which, if present in the data, would have produced a higher TS in β of the trials. Figure 7.7 is a graphical representation of this **Neyman construction**, originally introduced to calculate confidence intervals [215].

Before unblinding, the entire analysis can be characterized with its median sensitivity which is nothing but the median of those limits set when performing trials with realizations of pure background. This means replacing TS → $\mathrm{TS}_0^{(50\%)}$, the α = 50% quantile of the background TS distribution.

Since producing signal trials specifically for both a number of signal hypotheses and a range of μ is computing-intensive, I use a weighted method commonly used in IceCube. In short, signal trials are collected together with the information n_{inj} on how many injected signal events they contain. The TS distribution is then derived by reweighting this set of trials so that its n_{inj} follow the Poissonian distribution for a given μ. Trials are accumulated until the statistical uncertainty passes below a predefined tolerance, with efficient sampling techniques to help reach this point in as few trials as possible. Especially slow likelihood evaluations such as in chapter 9 parallelize this procedure, which makes it feasible to obtain results even if some of the efficiency is lost.

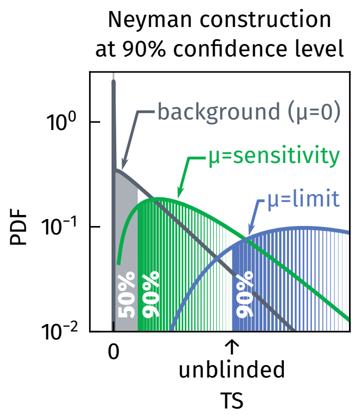

Figure 7.7.: A diagram of the Neyman construction, using test statistic PDFs given hypotheses parameterized by μ.

7.3.5. Discovery potentials

Another useful characteristic are discovery potentials. Here, β = 0.5, and e.g. $\alpha = 2.87 \times 10^{-7}$, i.e. the one-sided 5σ quantile of a standard normal distribution, which is a significance commonly used as the threshold in a hypothesis test for considering a result a discovery. This can be interpreted as finding the signal strength for a given hypothesis where half of the trials would qualify as a discovery. The numerical method is identical to the limits and sensitivities. However to find $\mathrm{TS}_0^{(5\sigma)}$, producing $\mathcal{O}(10^8)$ background trials would be too computing-intensive, and so the distribution needs to be extrapolated, for which $\mathcal{O}(10^5)$ trials are sufficient.

This distribution can be parameterized as

$$(1-\eta)\,\delta + \eta\chi^2(\mathrm{TS}; n_{\mathrm{dof}}), \tag{7.28}$$

given the fraction of overfluctuations, η, and a number of degrees of freedom, n_{dof}. If the likelihood perfectly described the data, and the minimization was unbounded, Wilks' theorem predicts $\eta = \frac{1}{2}$, $n_{\mathrm{dof}} = 2$. This means that half the trials are under-fluctuations, and the rest exponentially decline according to a χ^2 distribution with two degrees of freedom. In reality, the fraction of under-fluctuations $(1-\eta)$ will be different, and the rest of the distribution will follow a χ^2 distribution with a different number of degrees of freedom, which can be determined from a fit in the $\mathrm{TS} > 0$ region.

One of several previously used fit methods maximized an unbinned likelihood on the background trials. This was dominated by the shape of the distribution at the lower range of TS, where the most trials are. For this thesis, I developed a new extrapolation method which seeks to more accurately describe the distribution in its tails, which are more relevant to the extrapolation. It is based on a binned curve fit, with three peculiarities:

1. When fitting the shape of the PDF, discrepancies can accumulate unnoticed, with a big impact on the survival function at the tail, which however is the desired quantity of the distribution. I therefore directly bin the survival function $\mathrm{SF}(\mathrm{TS}) = \int_{\mathrm{TS}}^{\infty} \mathrm{dTS}'\, p(\mathrm{TS}')$.

2. Still, the curve subject to fit traverses several orders of magnitude on its ordinate. To ensure these regions of the abscissa have approximately equal influence on the fit, the fit curve is $\log \mathrm{SF}$.

3. Instead of weighting the bins according statistical uncertainties, a cutoff is made at the point where 400 trials are still in the tail in order to avoid sensitivity to fluctuations.

This method is shown in comparison in fig. 7.8. The resulting survival function matches the background trials better than the old method.

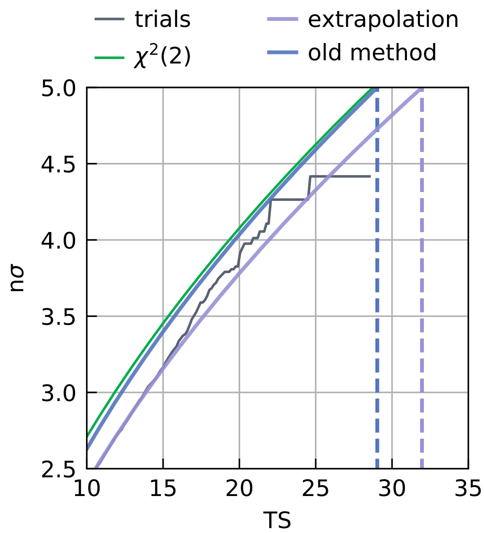

Figure 7.8.: A comparison between different methods to extrapolate the background TS distribution, specifically in its tail. Rather than the PDF, the y-axis is the survival function (1 - CDF) in terms of $n\sigma$. Grey: binned background trials. Green: $\delta + \chi^2$ with fixed $\eta = 0.5$ and $n_{\mathrm{dof}} = 2$. Blue and violet: old and new extrapolation methods, where dashed vertical lines mark the respective $\mathrm{TS}^{(5\sigma)}$.

“They might be giants, and what are we going to do unless they are?”

John Flansburgh and John Linnell [216]

“He charged at Rocinante’s fullest gallop [. . .]”

Miguel de Cervantes

8 TXS 0506+056

This chapter motivates and describes a lightcurve correlation analysis as outlined in chapter 7 targeting TXS 0506+056. The input is the data set described in chapter 3 and the source’s lightcurve mentioned in chapter 6.

First, sections 8.1 to 8.2 describe the historical context from which this analysis arose and recapitulate its motivation. Then, section 8.3 describes properties of the analysis, and finally section 8.4 related cross-checks. The analysis results follow in chapter 10, and a perspective on possible future analyses in chapter 11.

Author’s contribution: To support the unblinding of IceCube data for this analysis, I studied its discovery potentials, sensitivities and biases, and performed the aforementioned cross-checks.

8.1. Context and motivation

8.1.1. IceCube-170922A

On 2017-09-22 at 20:54:30.43 UTC, IceCube observed a track event registering a total charge of 5.8 kPE (see fig. 8.1). This event consequently passed the Extremely High Energy (EHE) filter (see section 2.5.2), and caused an automated real-time alert [217] which informed the astronomical community about its preliminary reconstructed direction and energy 43 seconds later (see section 2.5 for

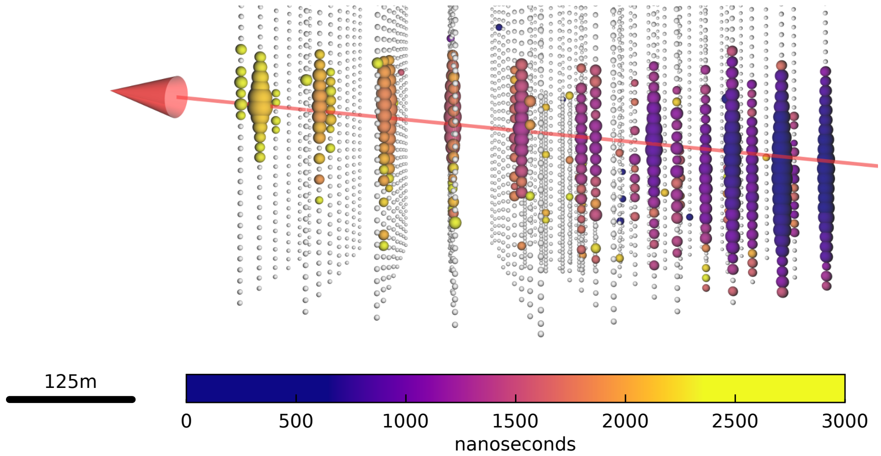

Figure 8.1.: The EHE event 170922-A. Each sphere represents an IceCube DOM, its size the measured charge, and its colour the time it was hit. The arrow shows an angular reconstruction of the event, which enters the detector from below the horizon and whose upper portion is partially obscured by the dust layer.

a description of this process in its current form).

With the event automatically transmitted to the North via satellite, manual inspection of the event and detector status excluded the possibility of a misreconstructed coincident background event, and confirmed that the detector was operating normally. More sophisticated reconstruction algorithms were consequently applied (see section 2.5). This led to a second alert 4 hours later [218]. This included the updated direction RA $= 77.4^{\circ}{}^{+1.0}_{-0.7}$, $\delta = 5.7^{\circ}{}^{+0.5}_{-0.3}$ (J2000) with 90% confidence intervals. The latter are particularly important for pointing telescopes which can not expose the entire angular surface of IceCube's PSF at once.

Using dedicated neutrino event simulations, a prior of an $E^{-2.13}$ spectrum is convoluted with variations on the ice characteristics, detector response, event statistics, and the distance of the interaction vertex from the detector. By comparing the event's deposited energy of (23.7 ± 2.8) TeV with these simulations, the most likely primary neutrino energy is 290 TeV, with a 90% C.L. lower limit of 183 TeV.

Already the brightness and declination provided by the first alert translate into a **signalness** of 56.5% [217], i.e. a relatively high probability of being astrophysical in nature based on the rates from the known fluxes. The improved reconstructions confirmed this; the arrival direction below the horizon almost completely excludes muons, and the energy significantly above 100 TeV strongly limits the rate of atmospheric neutrinos. As such, this event was treated as a candidate for an astrophysical neutrino and follow-up observations by other instruments which could reveal its source were strongly encouraged.

8.1.2. Multi-wavelength observations

The event IceCube-170922A arrived within 0.1° of the blazar TXS 0506+056. Six days later, Fermi-LAT reacted to the neutrino alert as part of its multimessenger programme, reporting TXS 0506+056 to be in a flaring state which had started 5 months prior in April 2017 and was ~ 6 times brighter than the average flux during the 3FGL catalogue (August 2008 – July 2012) [67]. The IceCube alert and the Fermi coincidence together triggered an extensive campaign of MWL follow-up observations:

1. In VHE gamma rays by the IACTs H.E.S.S., Veritas, and finally Magic. Their observations directly following the alert suffered from bad conditions and could not detect the source. However MAGIC resumed observations on September 28 as a reaction to the flare reported by Fermi-LAT, and with 13h observation time accumulated until October 4 could detect it at 6.2σ in the band of $[80, 400]$ GeV. Magic also found a $p = 1.35\%$ hint towards day-scale variability. The Magic observations are consistent with the limits from the other IACTs as well as a non-detection by the water Cherenkov observatory Hawc.

2. In HE gamma rays by Fermi-LAT, which observed the full sky since 2008, allowing it to make a continuous lightcurve as well as a SED in the band of $[100\,\text{MeV}, 100\,\text{GeV}]$. Additionally, AGILE measured the elevated flux $\gtrsim 100$ MeV during 10–23 September.

3. In X-rays by Swift-XRT, MAXI GSC, NuSTAR and Integral. Among these it was detected by NuSTAR and Swift-XRT, which together covered the photon energy range 0.3 keV to 79 keV. They identified 9 X-ray sources within 2.1 deg of the alert direction, as well as finding spectral variability which correlates with the hint of variability seen by MAGIC.

4. In the optical, including infrared, by seven telescopes. Their combined limit on the TXS 0506+056 redshift implied the source to have high lumi-

nosity. Later observations using the Gran Telescopio Canarias were the first establish the redshift, setting it to $z = 0.3365 \pm 0.0010$ [219].

5. In radio by OVRO at 15 GHz as part of its continued monitoring of the source, and by VLA at 2 GHz to 12 GHz starting two weeks after the alert. Both telescopes detected it and found variability in its flux.

The lightcurves in fig. 8.2 show an overview of some of these observations. Compressing instead the time axis, they are also summarized as a mostly complete, contemporaneous MWL SED in fig. 8.3. This rich data was crucial for theorists to build models of the blazar's gamma-ray and possible neutrino production.

The gravitational wave telescopes LIGO and VIRGO were not observing at the time of the alert, LIGO having ended its second observing run a month prior on 2017-08-25 [220].

TXS 0506+056 is a known blazar first discovered and localized by the Texas Interferometer, and published in the Texas Survey of Radio Sources in 1996 [221]. It has remained in the Fermi source catalogue since 1FGL [222], based on the first 11 months of Fermi-LAT data. Its synchrotron peak frequency $\lesssim 10^{15}$ Hz (see fig. 8.3) lies on the upper end of the IBL range (see section 5.5.2). Its average gamma-ray luminosity during the Fermi-LAT observation period up to and including October 2017 is 2.8×10^{47} erg s^{-1} in the band of $[0.1, 100]$ GeV [125].

Recent studies find that this luminosity combined with its synchrotron peak frequency makes TXS 0506+056 an outlier in the blazar sequence (see section 5.5.3) [192]. While it has historically been considered a BL Lac because of the lack of evidence for broad optical lines, it has also been argued [192] that based on multiple observational aspects, it might belong to a special class of FSRQs where these lines are outshone by the enhanced optical light from the jet.

8.1.3. Neutrino follow-up

IceCube possesses a fast response analysis (FRA) [223] which performs a predefined point source search within the analysis-level GFU data stream that the IceCube real-time system transmits to the North via satellite. It accepts a source direction, extension, and time interval for the signal, and returns a list of coincident events, best-fit test statistic, and either an estimate or an upper limit for the signal flux. In the current context, this analysis was applied twice [124]:

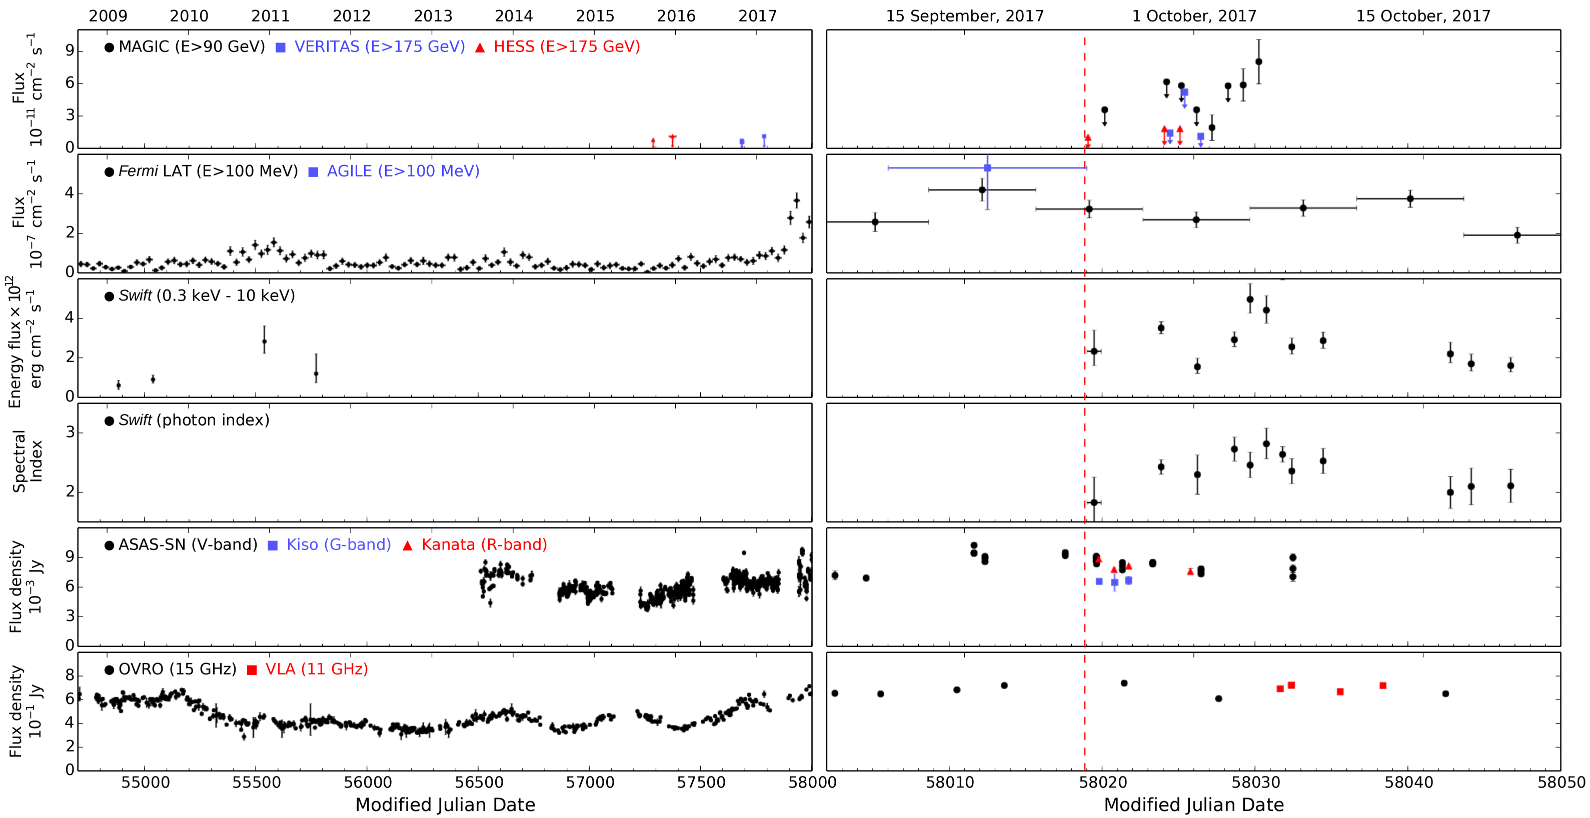

Figure 8.2.: The multi-wavelength lightcurves of TXS 0506+056 ranging from VHE gamma rays (A) to radio (F). The panels on the right-hand side show the period around IceCube 170922-A (dashed red line) in more detail. Note that not all instruments provide the same coverage. The VHE observation by MAGIC is consistent with the limits set by H.E.S.S. and VERITAS.

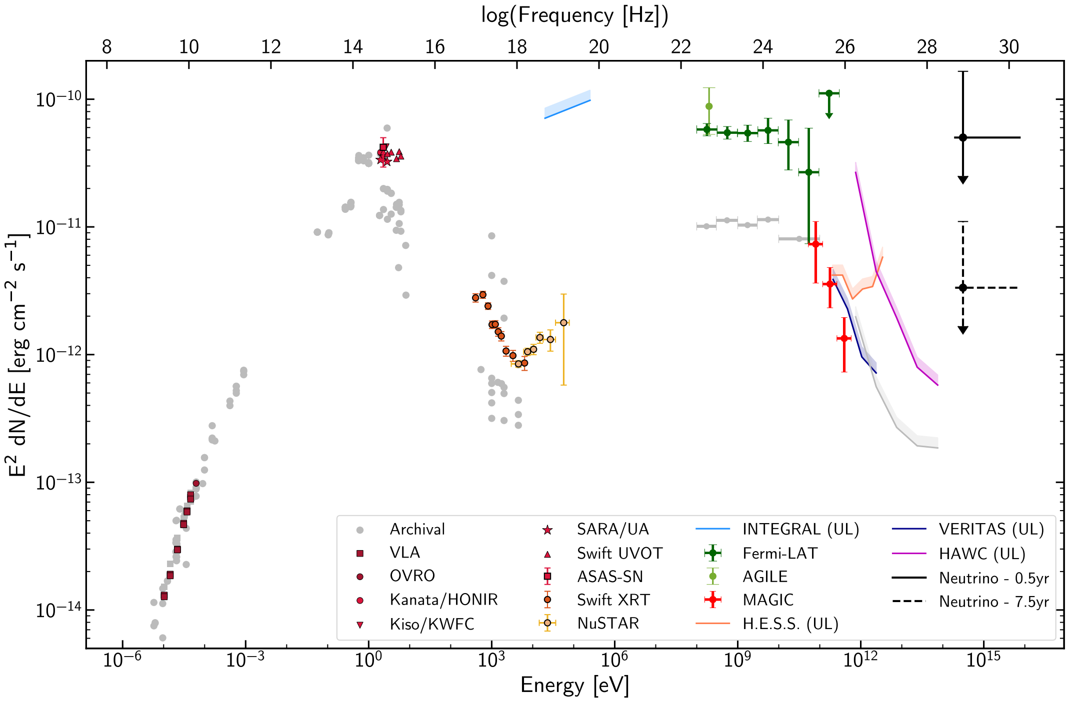

Figure 8.3.: The multi-wavelength electromagnetic SED of TXS 0506+056. Thanks to the involvement of many instruments, it covers a broad band of 10^{-5} eV to 10^{11} eV, revealing the two humps characteristic to a blazar spectrum (see section 5.4). Shaded bands represent 95%C.L. upper limits, and markers show observations. Colour is semi-contemporaneous in the sense of falling within 14 days of IceCube 170922-A while grey shows archival data. The observation of a single Extremely High Energy (EHE) neutrino can also be interpreted as an upper limit on the $\nu_\mu + \overline{\nu}_\mu$ flux [125]. The SED displays this assuming two different emission periods.

- As a response to the EHE neutrino alert. The target was the reconstructed neutrino direction and used an extension of 0.8° to represent the localization uncertainty. The search period was an interval of ±1 day.
- As a response to Fermi's report of the TXS 0506+056 flare. The target was the direction of the same, with no extension. The search period was ±7 days.

Both analyses removed IceCube 170922-A from the data sample to avoid bias from the choice of observation target. Consequently, neither analysis found an overfluctuation (see section 7.3). The additional events revealed in spatial and temporal coincidence by the second analysis are consistent with background. For a steady flux within the search window which follows an E^{-2} spectrum, the resulting limits on the fluence* at 90% C.L. are 3.52×10^{-5} TeV/cm^2 for the first, and 4.63×10^{-5} TeV/cm^2 for the second analysis [124].

Another follow-up of neutrino observations by ANTARES in December 2017 within ±1 day of the alert resulted in no candidate events, and limits that are weaker than the corresponding IceCube analysis by two orders of magnitude [224].

The event stream of EHE alerts constitutes only a small part of IceCube data. Assuming a power-law spectrum and extrapolating to lower energies means that IceCube should observe additional events. Analogously, the FRA examined only a two-week period. This leads to archival point-source searches which use analysis-level (atmospheric neutrino-dominated) data samples reaching back to IC40, i.e. the data set of chapter 3, to examine TXS 0506+056. The IceCube collaboration launched three archival searches in the month following the alert. They can be ordered by an decreasing number of restrictions on the timing aspect of the signal hypothesis:

1. A lightcurve-correlated search for neutrinos which are observed during the higher parts of the gamma-ray lightcurve, and therefore particularly during its highest flare in 2017.
2. An untriggered time cluster search, that looks for a flare of neutrinos of any width at any time using a box or Gaussian profile.

*fluence=$\iint F(E,t)\,\mathrm{d}t\,\mathrm{d}E = \frac{\mathrm{d}E}{\mathrm{d}A}$

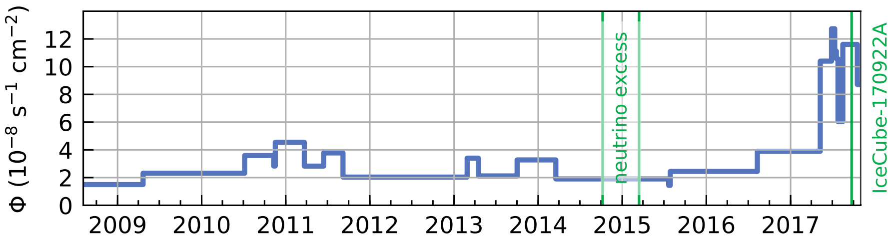

Figure 8.4.: The lightcurve from fig. 6.6 shown in context with the neutrino observations of section 8.1.

3. A time-integrated search, for any excess of neutrinos above the background, regardless of their times.

Search 2 [124] found a neutrino excess during a window of 158 days (box function) or 110 days (Gaussian) centred on December 2014, however this reaches only 3.5σ significance. Search 3 is dominated by the same neutrino excess [124], which brings it to 2.1σ in the first 7 years, i.e. excluding the data sample which contains IceCube 170922-A. During the time of the excess, the TXS 0506+056 lightcurve was in a quiescent state. The searches 2 and 3 found no cluster during the 2017 flare when disregarding IceCube 170922-A. Therefore these results are not relevant to search 1, which this chapter concerns.

8.2. Argument for a lightcurve correlation analysis

The general arguments for this analysis method from chapter 7 will not be repeated here, neither those for targeting blazars in general of chapter 5. Instead, this section makes the case for applying the analysis to the specific case of TXS 0506+056.

When choosing to interpret the coincidence of a blazar flare and a single EHE neutrino as causal in nature, this implies a neutrino flux which is correlated to the HE gamma-ray lightcurve, which fig. 8.4 shows in more detail. While its magnitude is subject to the FRA limit and could be further suppressed by assuming a large population of similar sources [225], its timing provides interesting points of attack for neutrino observation:

1. The flare was longer than the two weeks analysed by the FRA, beginning

already in April 2017 and reaching the flux level observed on September 22nd for the first time three months prior.

2. The flare also shows substructure, primarily in the form of two peaks separated by a sharp dimming around August.
3. Integrating over the entire Fermi-LAT observation time, most of the source's luminosity falls in the period preceding the 2017 flare.

The lightcurve correlation analysis (see chapter 7) is specifically constructed for such a signal. It was used before on other blazars, but not on TXS 0506+056 since it is not a member of the Fermi Monitored Source List (MSL), on which previous lightcurve correlation neutrino searches focused. The criterion for the MSL is that the photon flux exceeds $10^{-6}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ for 1 day. However as will be discussed in section 8.4.1 this biases it against harder sources which actually have a lower photon flux, but higher energy flux. To some extent this can be considered an unphysical criterion, since the energy flux is what represents the luminosity, and is conserved during the reprocessing of gamma-ray photons in the source environment and on the path to the observer.

The other proposed archival searches would not be as likely to discover a lightcurve-correlated neutrino signal. The integrated search can not exploit the potential variability of the source, giving equal weight to neutrinos during the 2017 flare as to the rest. The untriggered search on the other hand might be sensitive to additional neutrinos during the 2017 flare, but this is diluted by the number of trials due to the moving search window. Furthermore it can not integrate neutrinos from every period of the data set.

8.3. Analysis

8.3.1. Discovery potentials and sensitivities

Figure 8.5 shows the discovery potentials for an energy spectrum with $\gamma = 2$ and a range of thresholds. They are presented both as the pure number of signal events, μ, as well as the **fluence**. The latter is defined given the signal flux, Φ, calculated from μ (see section 7.2.4), as

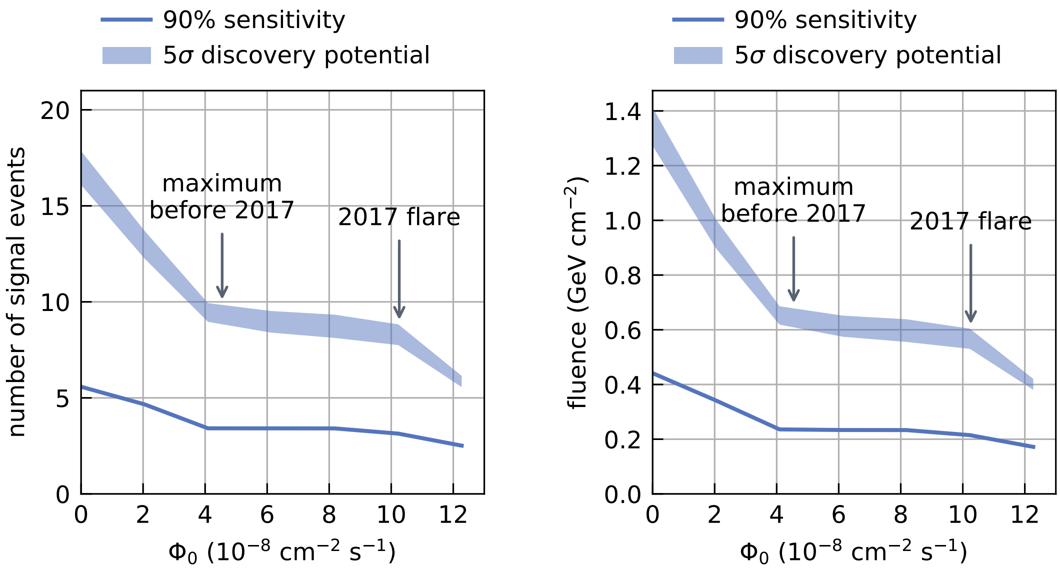

(a) Number of signal events. **(b)** Fluence.

Figure 8.5.: The median 5σ discovery potentials and 90% sensitivities for this analysis, assuming a signal with spectral index $\gamma = 2$ and a range of thresholds Φ_0 on the x-axis. The shaded region covers the range of $\mathrm{TS}_{5\sigma}$ extrapolations (see fig. 7.8).

$$\mathcal{F} := \sum_j \int_{\text{livetime}_j} \mathrm{d}t \int_{E_j^-}^{E_j^+} \mathrm{d}E E \times \Phi(t, E) \tag{8.1}$$

where $[E_j^-, E_j^+]$ is the energy interval containing 90% of the observed signal in sample j, as in the fluence definition presented in [203]. This integral is done analytically, but the bounds determined based on IceCube simulation.

Using either quantity, the discovery potentials reveal three domains of the threshold, which reflect the three domains of the source flux level (compare with fig. 8.4):

1. Low, $\Phi_0 < 4 \times 10^{-8}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$: the analysis needs to search for signal neutrinos distributed across the entire lightcurve. Raising the threshold removes periods of the lightcurve, restricting the search space, giving more weight to events induced by the signal flux. This improves the discovery potential as the log-likelihood terms of fewer events sum up to the same $\mathrm{TS}_{5\sigma}$.
2. Middle, $4 \times 10^{-8}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1} < \Phi_0 < 1 \times 10^{-7}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$: the threshold has passed over the pre-2017 lightcurve but remains below most of the 2017 flare. The time distribution of the signal hypothesis only changes slowly as low blocks are weighted down more and high blocks are weighted up more, and so the curve reaches an almost-plateau.
3. High, $\Phi_0 > 1 \times 10^{-7}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$: the threshold starts to eat into the 2017 flare and the sensitivity improves again, until the maximum Φ_0.

8.3.2. Bias

The result of the likelihood maximization is a point in the likelihood parameter space. When injecting a signal according to certain parameters, we expect the likelihood maximization to recover them as best fit parameters once the signal injection is strong enough. To study the dependence on signal strength, this experiment was performed with a set of n_{inj}, i.e. without Poissonian variations.

Figure 8.6 shows the results of these experiments. The mean spectral index converges with n_{inj}, but with an asymptotic bias that increases towards harder spectra (see fig. 8.6b).

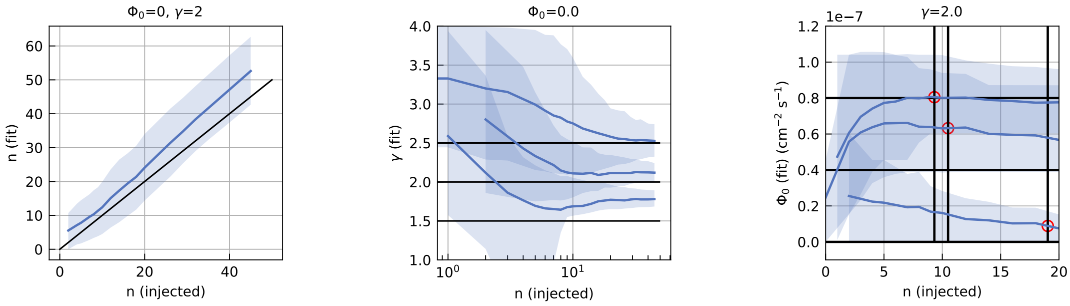

(a) Injecting with $\Phi_0 = 0$ and $\gamma = 2$.

(b) The injected γ assumes values of $\{1.5, 2.0, 2.5\}$ and $\Phi_0 = 0$.

(c) With $\gamma = 2$, Φ_0 assumes a range of values. Red circles mark for each injected threshold what would be fit in case of a discovery-strength signal.

Figure 8.6.: Plots showing the bias of best fit parameters versus their injected values. Starting from a set of signal trials for each individual injected hypothesis and number of signal events n_{inj} (shown on the x-axis), the blue lines show the median, while the shaded areas show the central 1σ interval. Black lines are the injected values.

The speed at which the biases converge given different signal hypotheses is not fairly measured by comparing them all at the same signal strength. After unblinding the analysis, we would treat the best fit parameters as physical measurements if the hypothesis test had yielded a discovery-level test statistic. The signal strengths at which to compare are thus the respective discovery potentials. For the threshold fit, these have been marked in fig. 8.6c as vertical lines and red circles. We see that the likelihood, if challenged with discovery-level observations of $\gamma = 2$, tends to distinguish a lower threshold from a higher. However for the upper two thresholds, there is a large overlap of the central 1σ intervals. This corresponds to how in the middle domain of the lightcurve flux level, the time PDF shape changes the least (see section 8.3.1).

8.4. Systematics and checks

All analyses suffer from systematic effects that can affect the interpretation of results. In this section we quantize some of these systematic effects, and explore

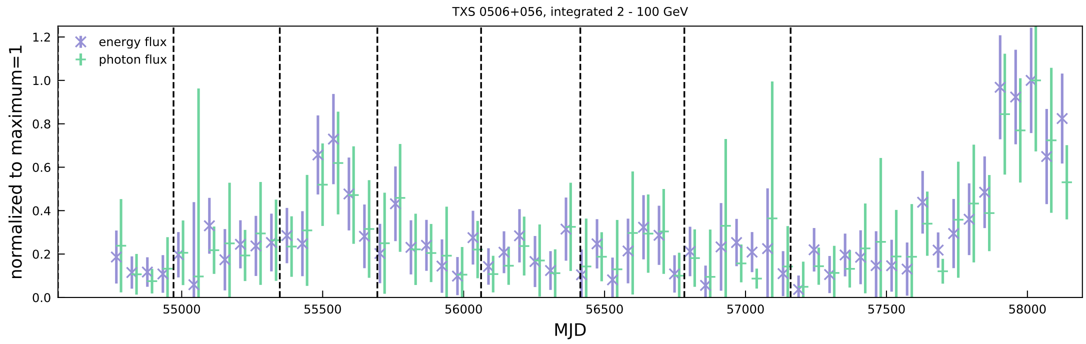

Figure 8.7.: Lightcurves of energy flux and photon flux between 2 GeV to 100 GeV. In an extension of the basic lightcurve fit procedure described in section 6.2, here the gamma-ray spectral index is fit independently within each 55-day bin. Having been divided by their peak value, these lightcurves have a largely compatible shape.

how the limits can be extended into larger regions of the hypothesis space.

8.4.1. Systematics with energy and photon flux

One source of systematics in this analysis is the choice of using a lightcurve in terms of either energy flux or photon flux. For lightcurves where the energy spectrum is constant in time, i.e. of the form $\frac{\mathrm{d}\Phi}{\mathrm{d}E}(t, E) = \Phi(t)\left(\frac{E}{E_0}\right)^{-\gamma}$, photon- and energy flux are defined respectively as

$$\mathrm{PF}(t) = \int_{E_-}^{E_+} \mathrm{d}E \frac{\mathrm{d}\Phi}{\mathrm{d}E}(t,E) \tag{8.2}$$

$$= \Phi(t) \int_{E_-}^{E_+} \mathrm{d}E \left(\frac{E}{E_0}\right)^{-\gamma}$$

$$\mathrm{EF}(t) = \int_{E_-}^{E_+} \mathrm{d}E E \frac{\mathrm{d}\Phi}{\mathrm{d}E}(t,E) \tag{8.3}$$

$$= \Phi(t) \int_{E_-}^{E_+} \mathrm{d}E E \left(\frac{E}{E_0}\right)^{-\gamma}.$$

The time-dependent term factorizes so that $\mathrm{EF}(t) = \alpha \mathrm{PF}(t)$, where α will differ between sources of different spectra, but for each source is constant in time. It also depends on the integral bounds $[E_-, E_+]$. However by extending the bin width to 55 days (for a lightcurve above 1 GeV), it becomes possible to fit the spectral index within each bin. Thereby, α is no longer constant in time, and the choice between the differently-shaped energy flux or photon flux lightcurve will affect the time PDF of the analysis which uses it.

Two tests have been undertaken to estimate this potential impact. The available data consists of a 55-day lightcurve above 2 GeV from [206], plotted in both variants in fig. 8.7 given integral bounds of $[2, 100]$ GeV. The incompatible shape with the fixed-spectrum 7-day lightcurve leads to the main complication in the comparison, and therefore comparisons are only made perpendicular to the intended systematic axis i.e. internally between photon flux and energy flux.

1. Applying the likelihood according to the original 7-day lightcurve, but injecting according to the 55-day photon flux or energy flux lightcurve.

2. Also switching the likelihood over to the alternative lightcurves, and comparing the matching with the mismatched configuration.

Evaluating these test cases in terms of sensitivities and discovery potentials, the differences are $\lesssim 7\%$, a consequence of how little the shape changes in

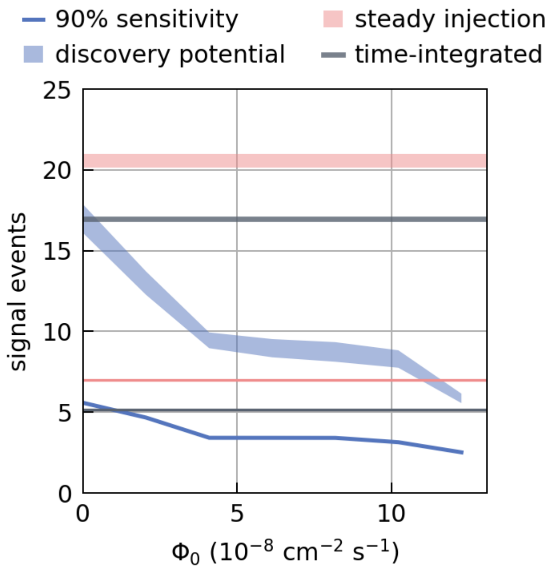

Figure 8.8.: Lightcurve systematic check when injecting a steady lightcurve.

fig. 8.7. The analysis is more sensitive to the overall shape of the lightcurve than to local fluctuations at small scale relative to the lightcurve maximum. Figure 8.7 contains, albeit not noticeable, one specific example of this in the form of a spectral hardening in one bin in 2014. This was the subject of [206], and prompted this study. It could be promoted to affect the energy flux lightcurve and thus possibly the analysis to a greater extent by extending the energy integral limits. However this would require extrapolating the SED beyond the range actually covered by the Fermi-LAT measurement.

In general, the photon flux is more sensitive to a variable spectral index, whereas energy flux compensates the degeneracy in the data between normalization and hardness.

8.4.2. Injecting a steady flux

I also computed discovery potentials (sensitivities) when applying the lightcurve correlation analysis, but injecting according to the hypothesis of a steady emission. Figure 8.8 shows the results. The steady injection requires 15% (25%) more signal events compared to when events are distributed according to the lightcurve at $\Phi_0 = 0$. Also shown for comparison are the sensitivity and discovery potential resulting from a time-integrated analysis, which performs better by 5% (7%).

8.4.3. Reinserting the EHE event

According to section 8.1, the event IceCube 170922-A triggered the observation of TXS 0506+056 in general and the chosen method of this analysis in particular. Also, IceCube 170922-A is a member of the chosen data sample in its own right. Therefore the result of this analysis carries a bias, a hidden trial factor for all the sources not observed and analyses not performed.

The motivation of the analysis is not to recast the observed coincidence, but rather to find additional events. Therefore I remove IceCube 170922-A from the data set to achieve the results to report.

If the analysis was functioning properly, it should however be sensitive to IceCube 170922-A. Therefore a test is undertaken by reinserting the event and comparing the unblinded result. The outcome of this is positive, with an unblinded TS = 15.44, $p = 3.2\sigma$.

8.4.4. Local scan p-value map

As discussed in section 3.4.1, observing the Moon shadow in atmospheric muons can show that the angular reconstruction is not noticeably biased. However these observations are integrated only over those declinations crossed by the Moon's trajectory when it stands above the South Pole horizon. For the declination of TXS 0506+056 below the horizon this result can only be transferred based on secondary arguments as there is no equivalent experimental result. Therefore the unblinded result of this chapter's analysis might be affected by such a bias, where a discovery-level excess of neutrinos exists but is offset from the TXS 0506+056 position, thus failing to reach the required significance in the specified hypothesis test. This is avoided with a check which also unblinds the fit for hypothetical source positions near the actual one, in a grid covering $[-0.3°, +0.3°]^2$.

Results of the analysis introduced in this chapter will be shown in chapter 10.

"One is an exception and two is a population."
Pieter van Dokkum et al. [226]

9
Blazar Flare Stacking

This analysis uses the time-dependent stacking method defined in chapter 7, the neutrino data of chapter 3, and lightcurves of chapter 6. Its results follow in chapter 10. In this chapter, the method's application to blazar flares is first motivated and then described.

Author's contribution: Starting from an existing set of Fermi sources and their lightcurves, I used the latter to select particularly variable blazars. I defined one of their three weighting schemes based on a model which was not used in previous stacking analyses. This also required me to find source redshifts not available in the Fermi catalogues. I then implemented the trivial method to combine and trial-correct the p-values and characterized the analysis according to sensitivities, discovery potentials, biases and cross-checks. I developed a bootstrapping method to estimate the discovery potential's uncertainty.

The astrophysical neutrino flux arrives from all regions of the sky and is consistent with being isotropic [227, 60]. Since the supposed mechanisms of neutrino production require compact environments (see section 1.4), in contrast to the cosmogenic neutrino flux, this implies that there must be extragalactic point sources in sufficient number to cover the sky when blurred by the IceCube PSF. The large number of trials from searching any direction in the sky can be traded for the assumption of specific known sources. Single-source lightcurve-correlated searches have also yielded null-results [204], which implies that any individual source is too weak to be discovered in the accrued IceCube livetime.

One possible way forward is to search explicitly for the sum of the signal from multiple sources, a method called **stacking**. The sensitivity thus gained to weaker individual sources which would not individually be discernible from background is again bought by adding an assumption on not only the source list (see section 9.1), but also how the signal is distributed within it (see section 9.2.1).

Blazars have also been the focus of time-integrated stacking analyses [228, 229, 230]. Generic population-based arguments can also be used to limit the contribution of sources with typical blazar luminosity and density distribution based on single-source searches [231] or generic all-sky methods [232]. In each of these cases, a time-dependent signal could still be discoverable below these blazar limits by means of a time-dependent search.

9.1. Source list

Since we use Fermi lightcurves the source list necessarily starts from Fermi catalogues, but followed by a selection for relevance and merit to the blazar flare stacking. One possible starting point, the Monitored Source List, is unsuitable according to the arguments laid out in section 8.2. Instead for practical purposes, the starting point is the broader selection of 2,254 extragalactic Fermi sources which were deemed relevant for the coincidence probability calculation in [125]. At the time of source list selection, 3FGL (4 years, >100 MeV) was the most recent general-purpose catalogue published by the Fermi collaboration. This is supplemented with the smaller companion catalogue 3FHL (7 years, >1 GeV), however at the final selection no sources remain that are only in 3FHL, but not in 3FGL.

9.1.1. Source class

Fermi data releases indicate the class of a source as long as it could be associated with an object from a range of existing astronomical catalogues. Most commonly, this will be the candidate object whose direction is closest to the Fermi source. There are also stronger cases of identification where, in addition, the variability is correlated to other wavelengths.

The selection of 2,254 extragalactic objects in [125] already uses only sources

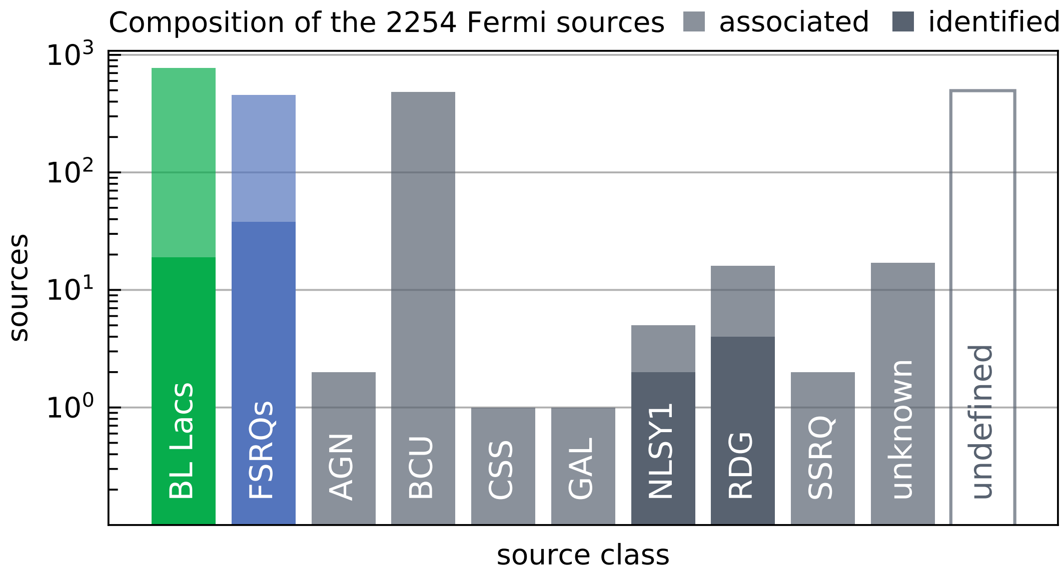

Figure 9.1.: The source classes present in the starting list of 2,254 sources, according to the relevant Fermi catalogue.

which either lie more than 5° from the galactic plane or have been classified as extragalactic objects, e.g. no pulsar wind nebulae. An overview of the remaining source classes in this set is shown in fig. 9.1 together with their acronyms according to [156, tab. 6], which are capitalized in case of a strong identification. In order to exclusively analyse Fermi blazars, from here on I restrict the selection to sources classified as a BL Lac, FSRQ, or blazar of unknown type (BCU).

The analysis does not expect all sources to have the same intrinsic neutrino flux. Instead, it explores different theoretically motivated hypotheses for their relative brightness, (see in the later section 9.2.1). One of those depends on the source redshifts. In the starting source list provided to us, some redshifts were already available from the underlying Fermi catalogues. For those which were not, we automatically searched the NED* for fuzzy matches to the source name, as formatting and nomenclature choices prevented an exact match in some cases. Finally, redshifts from TeVCat [186]† were added, replacing any existing one for the same source. The latter catalogue also provided the redshift

*NASA/IPAC Extragalactic Database `https://ned.ipac.caltech.edu`, accessed July 2018
†`http://tevcat2.uchicago.edu`, accessed July 2018

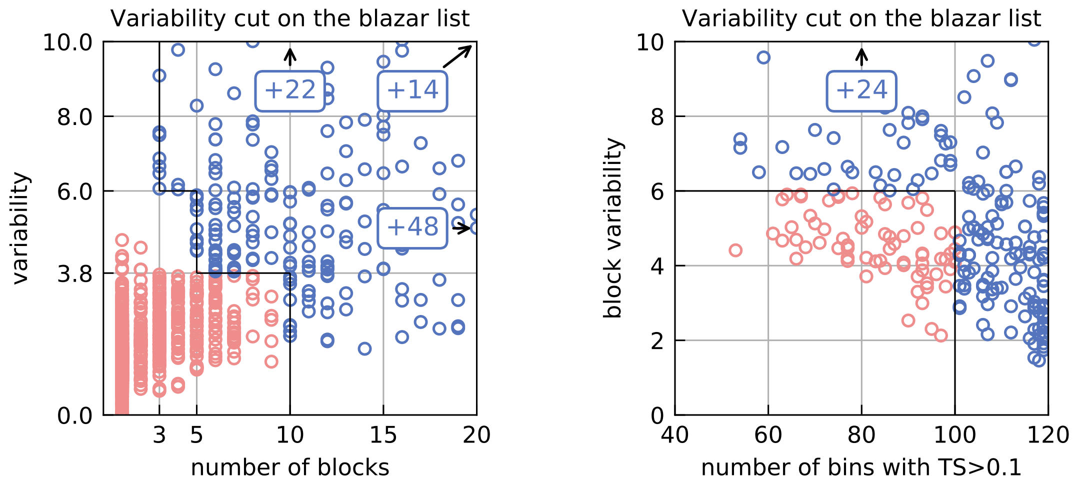

(a) The first variability cut.

(b) The second variability cut.

Figure 9.2.: Sources accepted (blue) and rejected (red) by the two-dimensional variability and quality cuts (black lines). The axis ranges have been chosen to show the cuts, and the number of sources falling outside them is indicated.

for TXS 0506+056, first measured in January 2018 [219]. The remaining fraction of 42% missing redshifts means the full source list is not fairly represented in the case of removing these sources, which needs to be respected when extrapolating to the source list before cut in order to compare results from both cases. This is similar to how the selection of Fermi sources does not evenly cover all redshifts, since farther sources will appear fainter. The latter has to be respected when extrapolating back to the entire population. These concerns are assuaged by the fact that this number sinks to 11% after the next level of source list selection, which favours brighter sources. As such, the bias from source selection largely overlaps with the bias from requiring redshift.

9.1.2. Lightcurve quality and variability

The analysis benefits from sources which show the most variability in the sense of the shortest flares with the highest amplitude relative to their quiescent state. This is a vague notion which must be implemented with a suitable variable. In order for the latter to be reliable or even defined, the underlying lightcurve also needs to fulfil quality criteria, represented by another variable.

First, I filter the lightcurves to bins where the gamma-ray likelihood fit (see section 6.2) has a test statistic TS > 4. With the remaining monthly bins I calculate the variability score v_{bin}[‡]:

$$v_{\text{bin}} = \frac{\Phi^{\max} - \Phi^{10\%}}{< \Phi >_{(\text{central } 80\%)}}, \tag{9.1}$$

where the average and minimum of the (energy) flux Φ take only the central 80% of bins into account. In order to define these central 80%, I need to require that at least 10 bins survive the filter TS > 4.

Then I rebin the lightcurve (see section 6.3), filtered to TS>0.1, and count the number of Bayesian blocks N_{block}. The first variability cut is then

$$N_{\text{block}} \geq 10 \ \lor \ (N_{\text{block}} \geq 5 \ \land \ v_{\text{bin}} > 3.8) \ \lor \ (N_{\text{block}} \geq 3 \ \land \ v_{\text{bin}} > 6).$$

In the next cut the variability score is calculated using the block representation, and so for this v_{block} to be defined, the first cut is designed such that in all cases $N_{\text{block}} \geq 3$. Those lightcurves with $N_{\text{block}} < 10$ are on average not as bright. However if their variability is high enough, this indicates that at least a small number of their bins capture a flare, and so they are kept by a v_{bin} cut which loosens in two steps with growing N_{block}. For $N_{\text{block}} \geq 10$, I consider the block lightcurve detailed enough so that a variability cut with v_{block} is preferable and so none needs to be made with v_{bin}.

Figure 9.2a represents this cut by the 256 sources selected and rejected on the two cut axes.

When computing v_{block}, outliers will already be smoothed away so all blocks are used for the minimum and the duration-weighted average instead of the central 80%. The quality of the lightcurve is related to the number N_{bin} of bins in the filtered lightcurve, i.e. the ones with TS>0.1 used for the block lightcurve. The second variability and quality cut is

$$v_{\text{block}} > 6 \ \lor \ N_{\text{bin}} > 100 \text{ (out of 119)}$$

where again, a trade-off between brightness and variability is struck. Figure 9.2b shows the blazars thus rejected or accepted on the cut axes. The latter are 65 BL Lacs, 114 FSRQs, and 8 FSRQs of unknown type, in total 187. Figure 9.3 shows

[‡]developed by Asen Christov

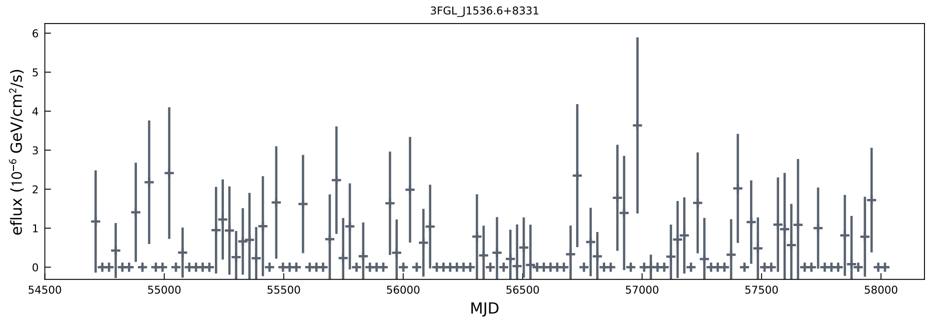

(a) Lightcurve failing the requirement of ≥ 10 bins with TS > 4.

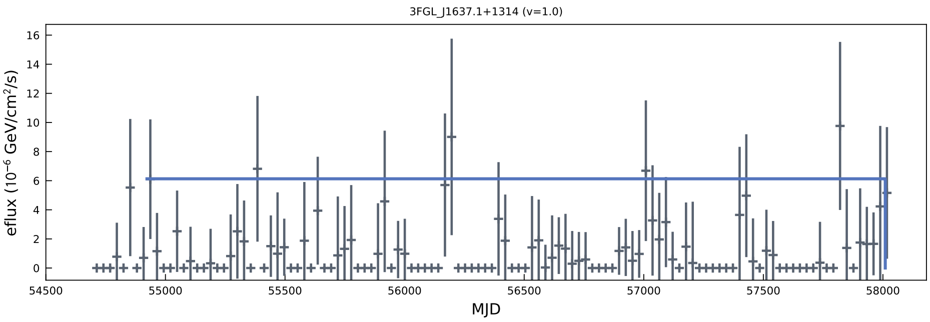

(b) Lightcurve failing the first cut.

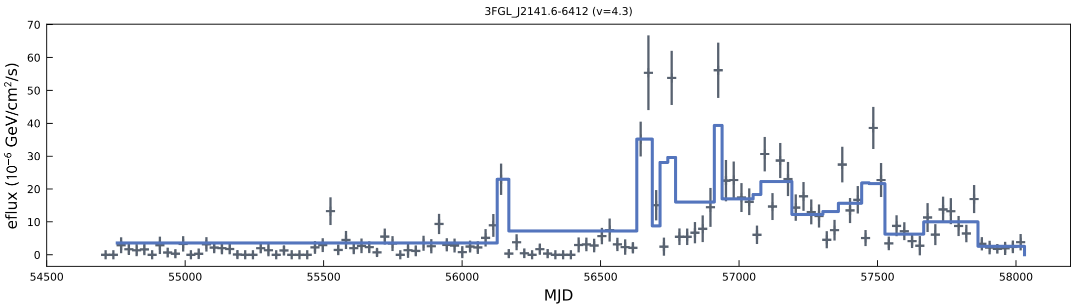

(c) Lightcurve failing the second cut.

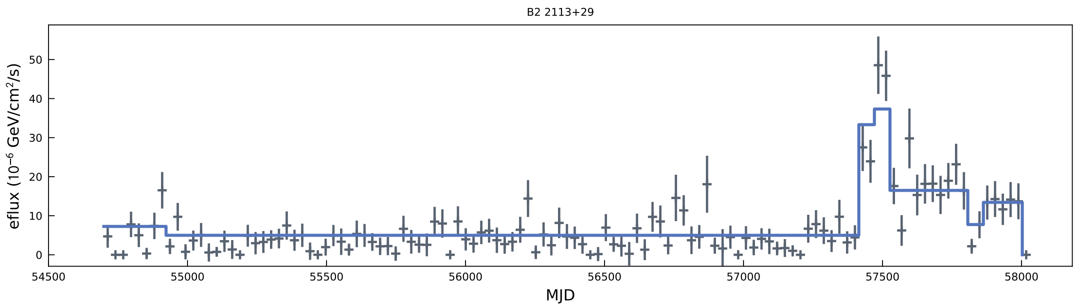

(d) Lightcurve passing all cuts.

Figure 9.3.: Examples lightcurves chosen to represent the source list selection based on them.

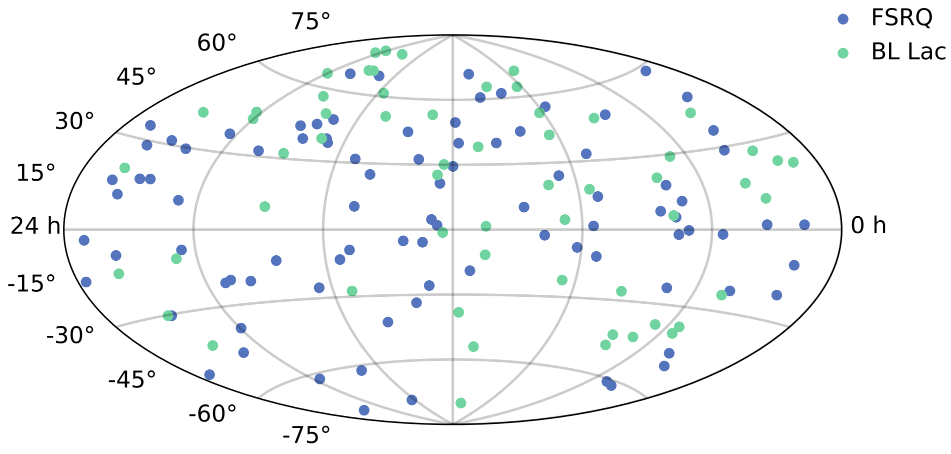

Figure 9.4.: The BL Lacs and FSRQs in the final source list, on a sky map in equatorial coordinates (J2000).

examples of lightcurves removed by or surviving this succession of cuts.

Among sources passing the variability and quality cuts, the average energy flux above 1 GeV is a factor 3.8 larger than among the sources going into the cuts. They can therefore be said to favour brighter sources. This is a direct result of the higher photon statistics in their lightcurves, which mean a flare can be more clearly distinguished from the quiescent state and more of the gamma-ray fits converge to a high TS value. The cuts also favour FSRQ type objects, whose fraction among blazars in the source list increases from 27% to 61%. The origin of this effect is less clear and might partially be an enhancement of existing catalogue biases.

9.1.3. Splitting into BL Lacs and FSRQs

Blazars are divided observationally into the classes of BL Lacertae objects (BL Lacs) and Flat Spectrum Radio Quasars (FSRQs), which are supposed to be physically different (see section 5.5.1). This includes the possibility that their potential neutrino production physics differs as well. In this case, limits produced separately for the two source classes will be easier to interpret. Additionally, if one class turns out to produce the majority of the neutrino signal in the analysis, the significance would be diluted by combining both.

Considering these points, I split the source list and henceforth analyse FSRQs and BL Lacs separately. In order to limit cross-contamination, this also means removing blazars of unknown class (see section 9.1.1). Figure 9.4 maps these two source lists in equatorial coordinates.

9.1.4. Removing and reinserting TXS 0506+056

This analysis was conceived before the observations of section 8.1, however I chose the data sample and source list for the eventual unblinding afterwards. Therefore these choices are biased from the prior knowledge of the data sample and the blazar TXS 0506+056. Because I already examined this source on its own in an analogous fashion (see chapter 8), I prefer to remove it from the source list in order to prevent the bias, rather than e.g. removing the event IceCube 170922-A from the data set.

Analogous to section 8.4.3, post-unblinding tests which reinsert the source will show whether the analysis is in fact sensitive to such a coincidence. Because of the doubts raised on the TXS 0506+056 classification in [192] (see section 8.1.2) it is reinserted either into the BL Lac or the FSRQ source list. This test is done purely as a check to help interpret the results, and does not contribute to the final reported p-value (see section 9.2.2).

9.2. Analysis method

This section describes points specific to the analysis. The general method it uses was already laid out in chapter 7 and specifically section 7.1.4.

9.2.1. Weighting schemes

Three weighting schemes cover part of an uncountable hypothesis space. Figure 9.6 presents the weight of each source in the analysis at $\tau = 0$, and a simplified view of the τ-dependence resulting from each weighting scheme for one source is shown in fig. 9.5.

Gamma-ray energy flux weighting: Based on $\pi^0 \rightarrow \gamma\gamma$ decays dominating the gamma-ray spectrum, energy flux weighting is the continuation of the time

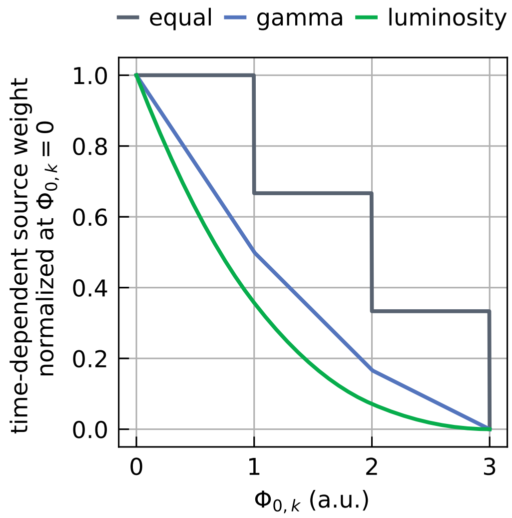

Figure 9.5.: An artificial lightcurve gets gradually covered by the threshold (x-axis). The y-axis shows the resulting factor to the source weight, before normalizing between sources and samples. This depends on the choice of weighting scheme: gamma-ray energy flux weighting (blue), luminosity-squared weighting (green), or equal weighting (grey).

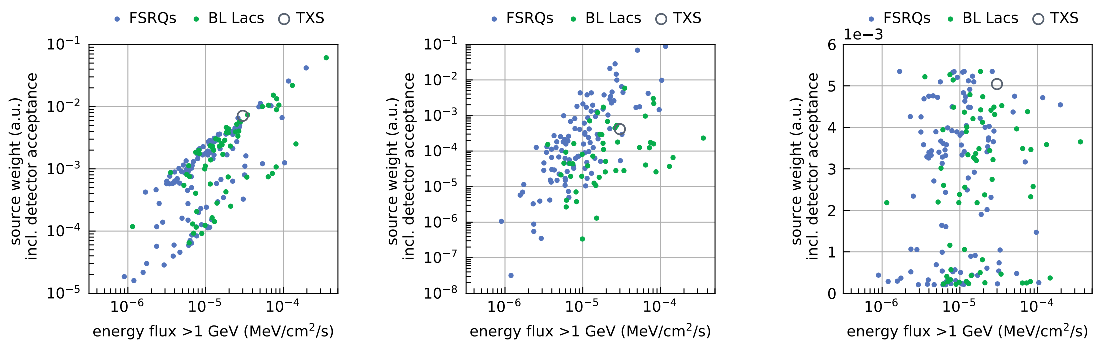

(a) Gamma-ray energy flux **(b)** Luminosity-squared **(c)** Equal

Figure 9.6.: The weighting schemes at $\tau = -\infty$ vs. the Fermi energy flux. The y-axis (arbitrary units) includes both the source weight as well as the detector acceptance at the respective source declination assuming a spectral index $\gamma = 2$, see fig. 7.3.

PDF definition into the source flux normalization. It consists of the time integral of the energy flux lightcurve. Another way to write it is in term of luminosity L:

$$L_\nu \propto L_\gamma, \tag{9.2}$$

where however L_γ is integrated only over the Fermi range and L_ν over the IceCube range. To assume that they are equivalent is a strong assumption. A more sophisticated treatment would be source-specific and involve the known spectral shapes and redshifts. The τ-dependence of the weighting reflects the thresholds at which the lightcurves are truncated by analytically integrating the truncated block lightcurves. These integrals are however rescaled to reproduce the time-integrated best-fit energy fluxes given $\Phi_0 = 0$.

Luminosity-squared weighting: A general one-zone $p - \gamma$ production model such as in [195] can be characterized by the luminosity L_γ and the injected proton flux = $\xi \times L_\gamma$ where ξ is called the baryonic loading. The emitted neutrino flux relative to the proton flux is called the neutrino efficiency ϵ_ν and will be $\propto L_\gamma$ if the size of the neutrino production zone is constant, and thus the photon field density increases with luminosity. Since the two humps of blazar SEDs are observed to be correlated (see section 5.5.3), the flux of lower-energy photons which are the targets in said $p-\gamma$ reactions can be approximated as proportional to the luminosity L_γ in the Fermi-LAT range of 0.1 GeV to 100 GeV [195].

If furthermore ξ is constant, this results in a relation between neutrino and gamma-ray luminosity:

$$L_\nu \propto L_\gamma^2. \tag{9.3}$$

In this case we use the redshift z explicitly to derive the respective energy flux normalisations from the luminosity in the form of

$$EF_{\nu/\gamma} = \frac{L_{\nu/\gamma}}{4\pi d_L(z)^2} \tag{9.4}$$

where d_L is the luminosity distance based on the Planck2015 cosmology as implemented in the `astropy.cosmology.Planck15` package [212]. The τ-dependence is implemented the same way as for the energy flux, again making use of the time-integrated gamma-ray fit but with the analytical integral instead over the square of truncated lightcurve.

Equal weighting: The weighting schemes so far have covered very specific areas of the physical hypothesis space. In stacking analyses, it is traditional to add another "equal" weighting scheme which covers a much larger fraction of the same, i.e provides sensitivity towards a relatively broad range of neutrino emission models as long as they sufficiently overlap with the chosen source list. In this time-dependent analysis the chosen "equal" weighting scheme still takes into account the time each lightcurve spends above the threshold.

At $\tau = -\infty$ most lightcurves have the same time range, although some are shorter due to the TS filter (see section 9.1.2). For higher τ the weights diverge further, for instance cutting out those sources which are entirely below the threshold.

It can be argued that a fairer weighting scheme would instead assume an equal neutrino luminosity for each blazar, so that the weights would be $\propto d_L(z)^{-2}$. In contrast, the "equal" weighting scheme as defined here is biased in favour of distant sources with the same luminosity.

9.2.2. Combining p-values

We desire to combine the hypothesis tests for the three weighting schemes and two source lists, and finally report only one p-value. The combination happens in two steps, once for the weighting schemes and then again for the source lists.

The three p-values for different weighting schemes are represented by the smallest (most significant) of them, called p_{min}. If the three hypothesis tests were identical, so would the p-values and p_{min} would be the significance of the combined test. However since they are different, they count as additional trials. This is reflected in how p_{min} will be smaller on average and needs to be "trial-corrected" upwards. If the three trials were completely uncorrelated this could be done with a simple analytical calculation, but in this case the correlation between the hypothesis tests, and the resulting correlation between their background p-value distributions, prevents this.

Instead I use scrambling to sample the p_{min} distribution under the background hypothesis, and calculate a trial-corrected p-value p_{post} according to it. Since the three unblinded hypothesis tests get the same data, they should also get the same scramblings within each particular background trial. I achieve this by controlling the seed of the pseudorandom number generator (PRNG). For each individual p_i, the largest possible value is the fraction of over-fluctuations η_i

(see section 7.3.3). Defining $\eta_1 < \eta_2 < \eta_3$, the largest possible value of p_{min} is $\eta_{\mathrm{min}} = \eta_1$. It applies to all trials where $\mathrm{TS}_1 = 0$, $p_2 > \eta_1$ and $p_3 > \eta_1$, even if $\mathrm{TS}_2 > 0$ and $\mathrm{TS}_3 > 0$ are not underfluctuations themselves.

The outcome of this procedure can be seen in the following chapter, alongside the unblinded p-values in fig. 10.5c.

The two trial-corrected p-values from the different source lists can be considered as independent since the source lists do not overlap. In this case, the smallest of them $p_{\mathrm{post,min}}$ can be trial-corrected with an analytical calculation:

$$p_{\mathrm{post,post}} = 1 - \left(1 - p_{\mathrm{post,min}}\right)^2 . \tag{9.5}$$

This can be derived from the fact that obtaining any less extreme result of $> p_{\mathrm{post,min}}$ from p-values p_k is identical to the case where all $p_k > p_{\mathrm{post,min}}$. If there are N independent such p-values, the probability of this outcome under the null hypothesis is $\left(1 - p_{\mathrm{post,min}}\right)^N$. A more extreme result $\leq p_{\mathrm{post,min}}$ is the complement of this outcome and its probability under the null hypothesis is therefore the p-value $1 - \left(1 - p_{\mathrm{post,min}}\right)^N$.

9.2.3. Sensitivities and discovery potentials

The sensitivities and discovery potentials for the six different hypothesis tests of the blazar flare stacking analysis are summarized in fig. 9.7. Similar to section 8.3.1 the hypothetical signal depends on the spectral index, which I chose to be $\gamma = 2$ as a benchmark, and on the common threshold parameter τ. For the latter I inject a range of values, where the main benchmark values are $\tau = 0.0$, where all thresholds lie at the estimated quiescent level, and $\tau = -\infty$ where all thresholds are at zero. The rest of the range is scaled relative to the greatest defined τ value for the respective source list. Similar to what was observed in section 8.3.1, the discovery potentials and sensitivities decrease with this parameter. The FSRQs in the luminosity weighting are an exception, which is explained in section 10.2.5.

The flux calculation of section 7.2.4 in the stacked case is a sum over all sources, and thus also for the fluence derived in section 8.3.1. This can then be understood as a certain fraction of the entire source population, and accordingly extrapolated.

Since the stacking likelihood evaluation requires more CPU time, the number

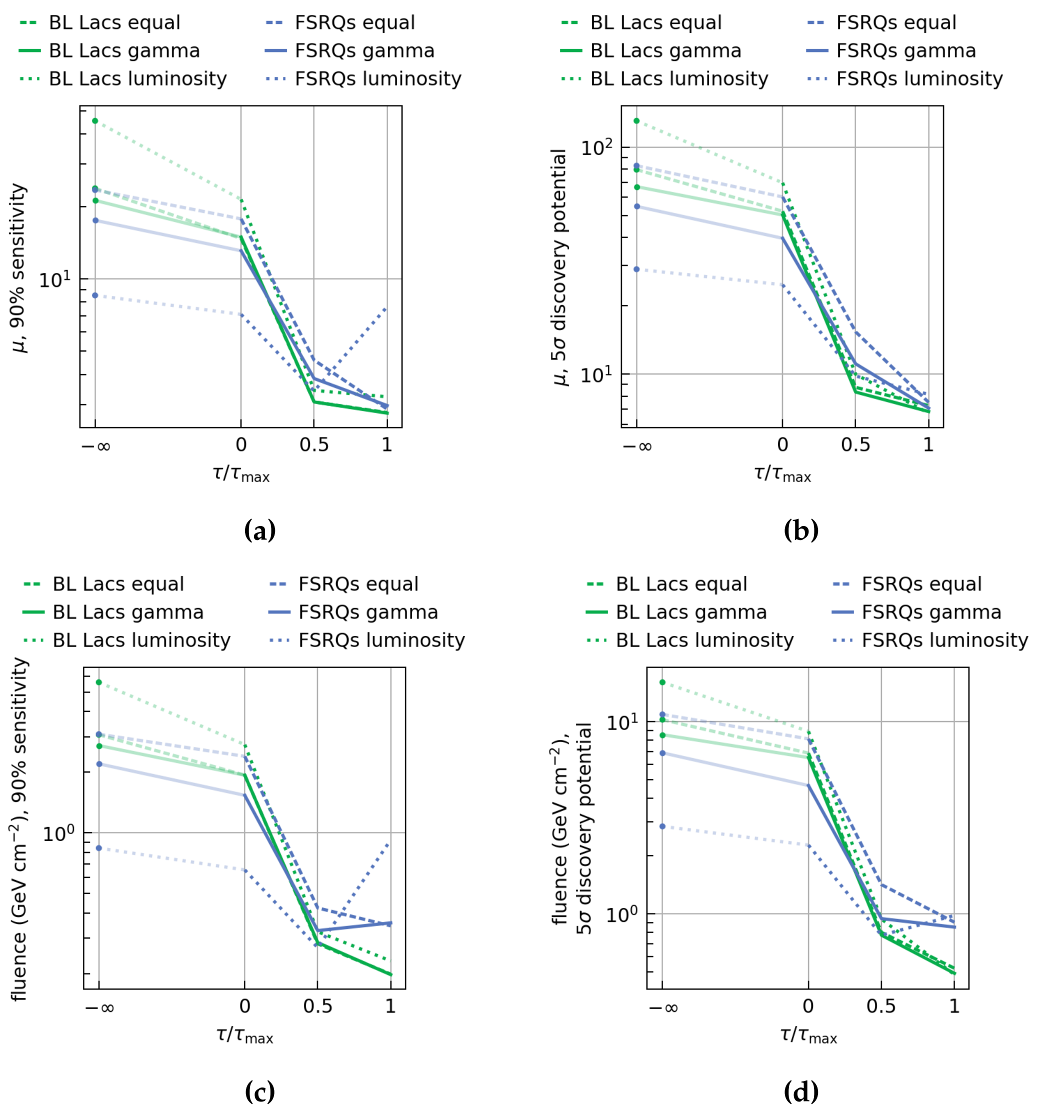

Figure 9.7.: The performance of the analysis, expressed in the number of events expected due to the signal flux as well as fluence (see sections 7.2.4 and 8.3.1).

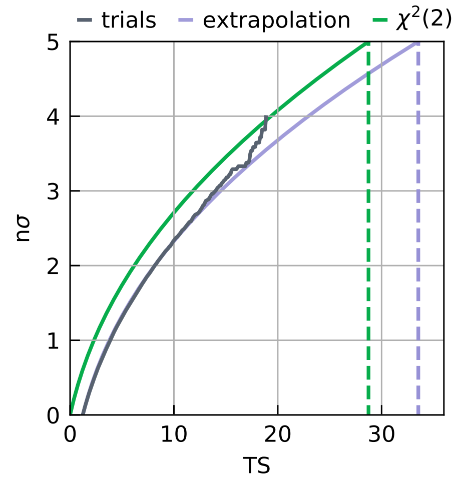

(a) A distribution of 30k background trials (grey), compared with the fit extrapolation (violet) and a χ^2 with $n_{\mathrm{dof}} = 2$ (green). The respective $\mathrm{TS}_{5\sigma}$ are vertical lines.

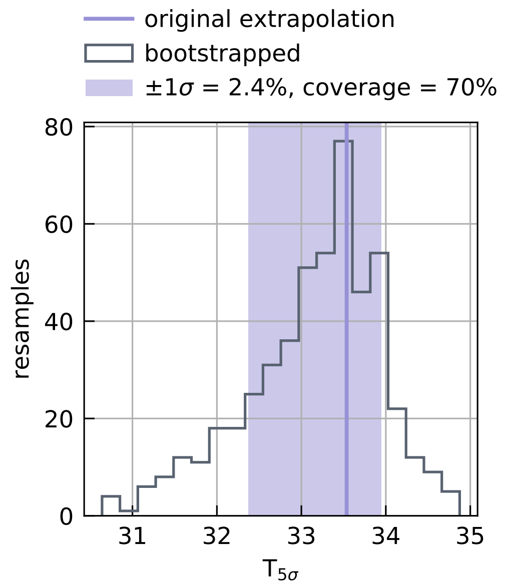

(b) Again the extrapolated $\mathrm{TS}_{5\sigma}$ (vertical line), the distribution resulting from bootstrapping the underlying background trials (grey), and the latter's arithmetic mean ± r.m.s. (shaded violet region).

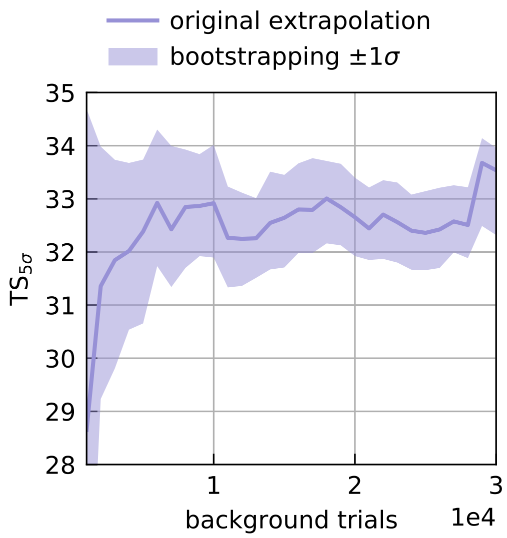

(c) The extrapolation and range from the previous plot (line and shaded region) depending on the number of underlying background trials (x-axis).

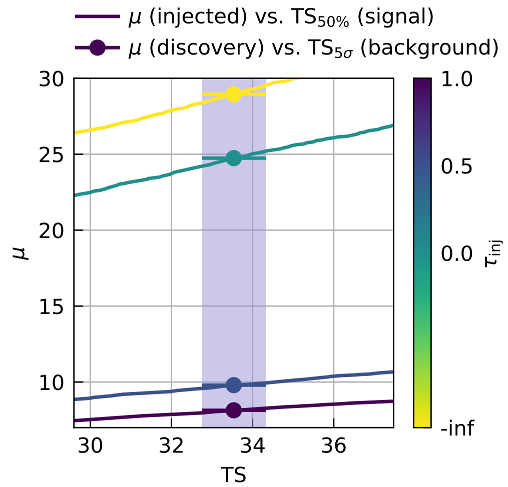

(d) The $\mathrm{TS}_{5\sigma}$ uncertainty band intersecting with the approximately linear curves of signal strength vs. median TS, for a range of injected τ_{inj} (colours).

Figure 9.8.: The entire path of bootstrapping the contribution of limited background trials to the discovery potential uncertainty, given an example likelihood configuration.

of available background trials is limited. This introduces a statistical uncertainty into the extrapolation of the test statistic distribution to the 5σ level, which can be translated into a relative uncertainty on the discovery potential. To estimate the magnitude of this uncertainty, I developed a bootstrapping method. Here, additional realizations of the background TS distribution and corresponding extrapolation (see fig. 9.8a) are created by resampling the same with the same statistics, thus covering the expected statistical variance in the latter. This is then translated into a a range of statistical variation in the derived extrapolation function and consequently $\mathrm{TS}_{5\sigma}$, as in fig. 9.8b. As the final step, we find the injected signal strength, μ, where the curve of median $\mathrm{TS}_{(50\%)}(\mu)$ crosses through this range of $\mathrm{TS}_{5\sigma}$, shown in fig. 9.8d. This then is the statistical uncertainty of the discovery potential, which here is of the order 3–4%. The calculation is also robust for smaller uncertainties because this curve can be linearized around the discovery potential.

As a general rule, the uncertainty becomes more narrow the more background trials are available (see fig. 9.8c). Here, when including a progressive number of background trials up to 3×10^4, both the extrapolated $\mathrm{TS}_{5\sigma}$ and its bootstrapped statistical uncertainty undergo an evolution that points to new features revealed in the tail of the distribution which the extrapolation does not yet fully reflect. In this case, and others like it, the apparent break would correct the extrapolation to smaller discovery potentials. Therefore by being unable to take it into account, our discovery potentials remain conservative.

9.3. Systematics and checks

9.3.1. Cross-check with integrated stacking

The analysis is a general variant of the time-integrated stacking such as in [228, 229, 230] where the signal time PDF is uniform i.e. cancels out with the background time PDF. Given that the integrated analysis is well-established and -checked, we use it to check our likelihood method's definition and implementation.

The test case consists of 16 sources at random, but fixed locations from both hemispheres. Their lightcurves are identical to a parameterized lightcurve which is steady at a baseline of 1 in arbitrary units, except for a 3-month flare

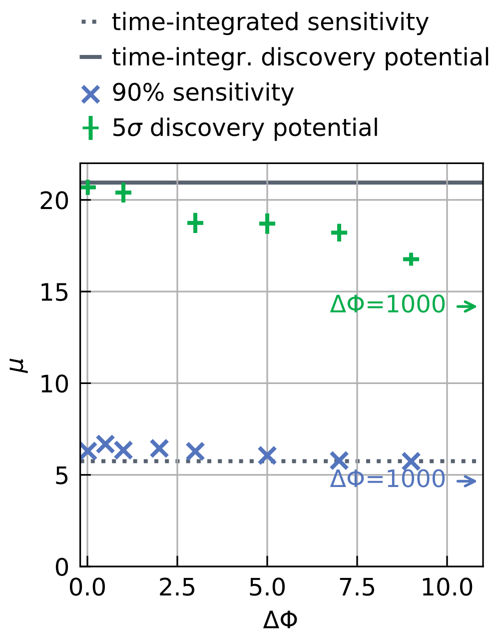

Figure 9.10.: Discovery potentials (green error bars) and sensitivities (blue x-es) for the cross-check. These are respectively compared to their counterparts for a time-integrated stacking of the same sources in the same data sample (horizontal lines). The parameter $\Delta\Phi$ describes the time-variability of the signal hypothesis.

of magnitude $\Delta\Phi$ (sketched in fig. 9.9). The data set is restricted to the IC86a sample in order to save computing time.

Used as indicators are the sensitivity and discovery potential, which for the flare stacking analysis depend on $\Delta\Phi$. The results in fig. 9.10 show the discovery potentials, considering their uncertainty (derived by bootstrapping, see section 9.2.3) are compatible with their time-integrated equivalent as $\Delta\Phi \to 0$. This is no longer the case for $\Delta\Phi > 2.5$, where the discovery potentials begin to decrease significantly. The sensitivities meanwhile show a weaker dependence on $\Delta\Phi$ and remain more or less consistent with the time-integrated sensitivity over the entire range. This example underlines the choice of analysis strategy with discovery as the objective.

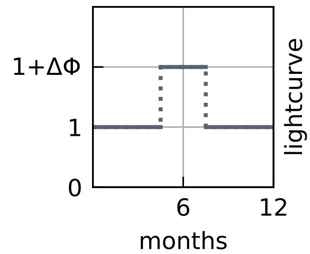

Figure 9.9.: The constructed lightcurve for the cross-check.

9.3.2. Bias

Figure 9.11 shows a study of the bias for this analysis using a similar method as section 8.3.1, repeated for each combination of the two source lists and three weighting schemes. Here, the calculation was done not for a set of n_{inj}, but for expectation values μ_{inj} by means of reweighting the signal trials to the corresponding Poissonian distribution. The injection happens at $\tau = 0$ and $\gamma = 2$. For the fit results, the figure shows their mean value and variance, represented here by the central $1-\sigma$ interval. The check confirms that for a large multiplicity of injected events, the mean fit results $\langle\hat{n}_S\rangle$, $\langle\hat{\gamma}\rangle$, $\langle\hat{\tau}\rangle$ approach μ_{inj}, γ_{inj}, τ_{inj}, respectively. The BL Lac source list in the luminosity weighting is an exception, where from $\mu_{\text{inj}} = 30$, the fit $\hat{n}_S$ is inconsistent with the injection.

Results of the analysis introduced in this chapter will be shown in chapter 10, particularly section 10.2.

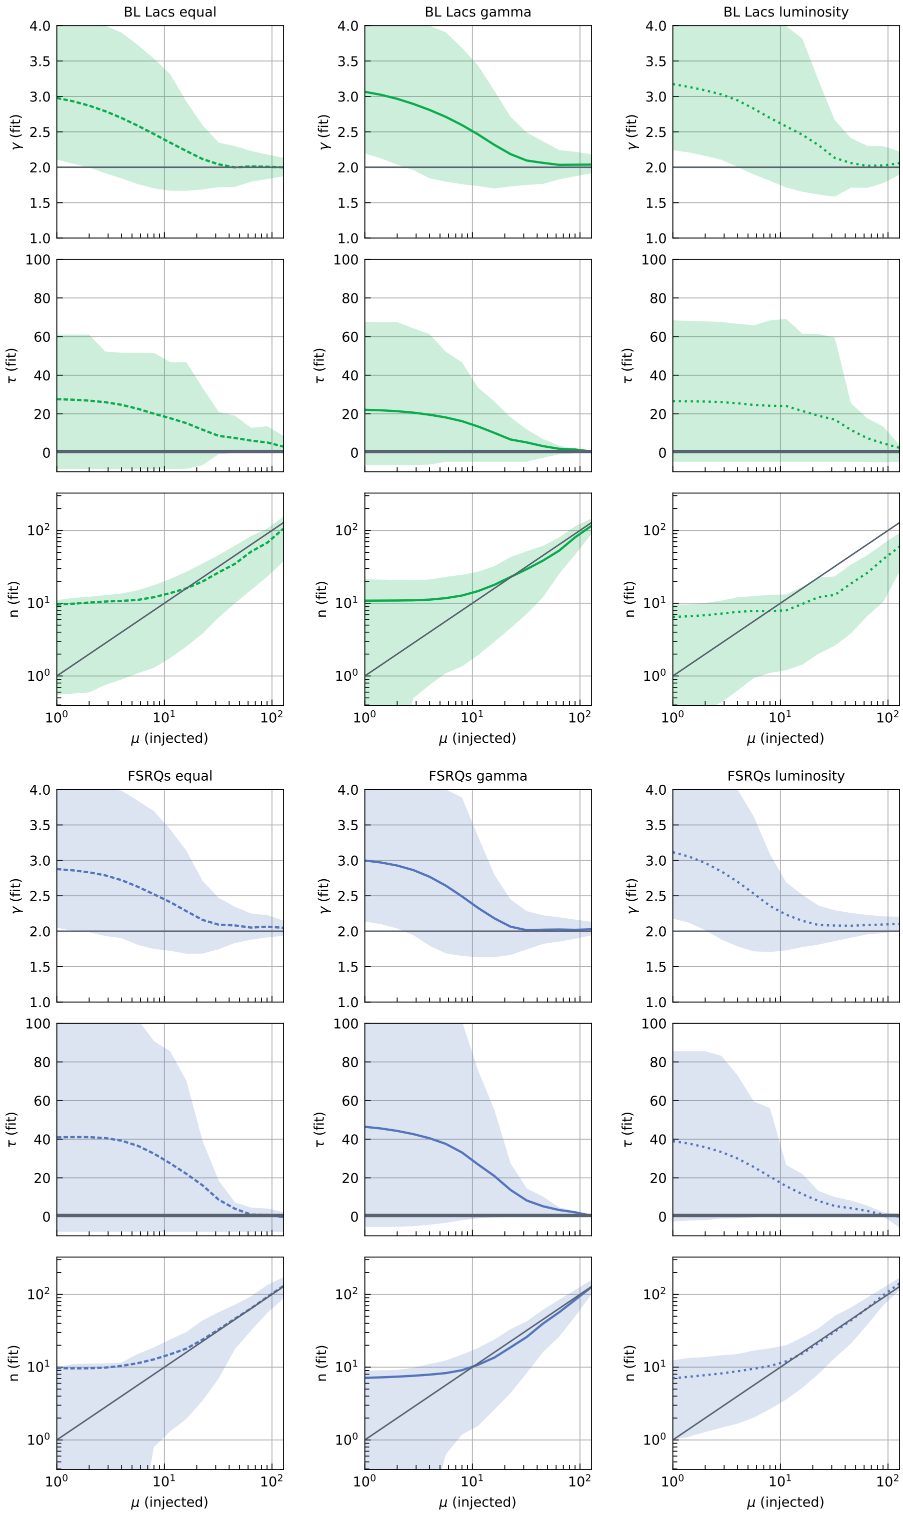

Figure 9.11.: Bias of fitting signal parameters given the six different likelihood configurations, and a benchmark signal hypothesis of $\gamma = 2$, $\tau = 0$.

"Sometimes you just have to muddle through."
Vera C. Rubin [233]

10 Results

This chapter contains the results of the analyses presented in chapter 8 and chapter 9. Its final section 10.3 discusses their physical context.

Author's contribution: I produced all results in this chapter, except for the other analyses shown for comparison in the final section. The limits are calculated according to an established method with a finite number of signal trials. I validated the methods used to sample these trials by defining their effective number and estimating the associated uncertainty of the limits.

10.1. TXS 0506+056

10.1.1. Unblinding

The analysis of chapter 8 is unblinded while removing the event IceCube-170922A from the data sample. This reveals the true times when events arrived, which are shown in relation to the lightcurve in fig. 10.1. The resulting best fit test statistic is 4.0, which compared to the background test statistic distribution results in a p-value $p = 12.6\%$ (1.1σ). This distribution is plotted in terms of its survival function in fig. 10.2a. This result does not show a significant correlation between the arrival times of neutrino events in IceCube and the HE gamma-ray variability of TXS 0506+056. After adding the alerting event

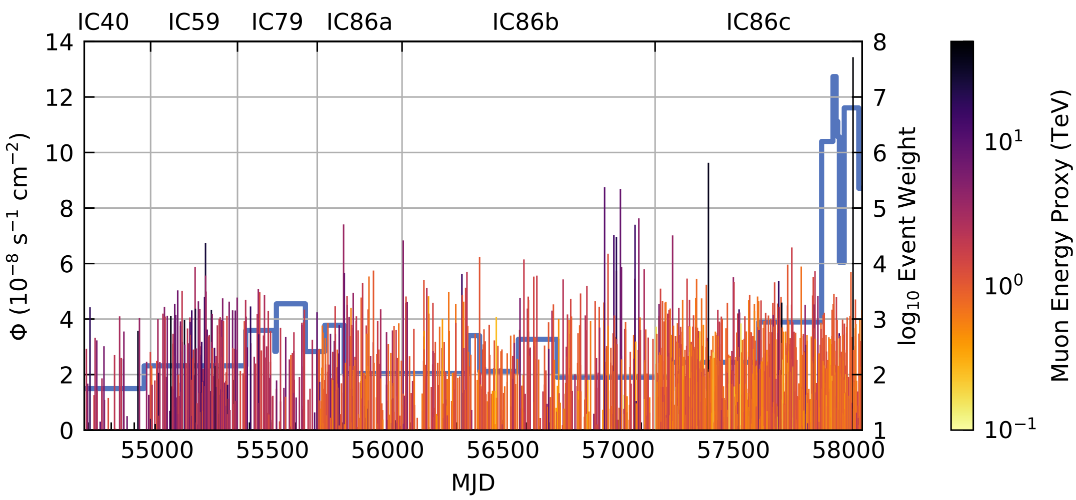

Figure 10.1.: Unblinded signal-like events arriving near TXS 0506+056, plus IceCube 170922-A. Their times (x-axis) can be compared with the lightcurve (blue), as can the ratio of signal and background $\mathcal{S}/\mathcal{B} > 1$ sans time PDF (see section 7.1) (right-hand side y-axis). Shown are only events whose reconstructed muon energy $E_{\text{reco}} > 100\,\text{GeV}$ (colour bar). This figure assumes a spectral index $\gamma = 2.06$ instead of the unblinded best fit $\gamma = 4$, the latter case being shown in fig. A.4.

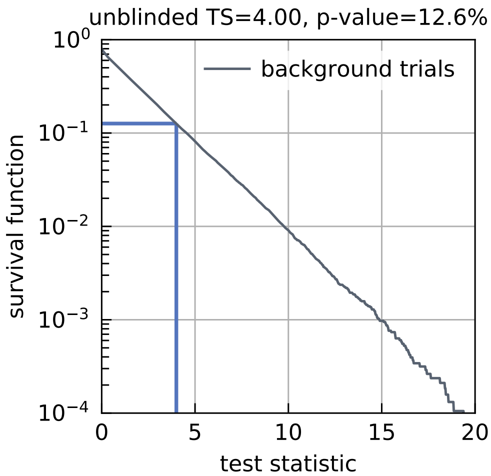

(a) The baseline analysis of TXS 0506+056. The blue line is the unblinded $\mathrm{TS}_{\mathrm{ub}}$, and the grey curve is the survival function $p(\mathrm{TS}) = \int_{\mathrm{TS}} \mathrm{dts} \frac{\mathrm{d}n_{\mathrm{BG}}}{\mathrm{dts}}$.

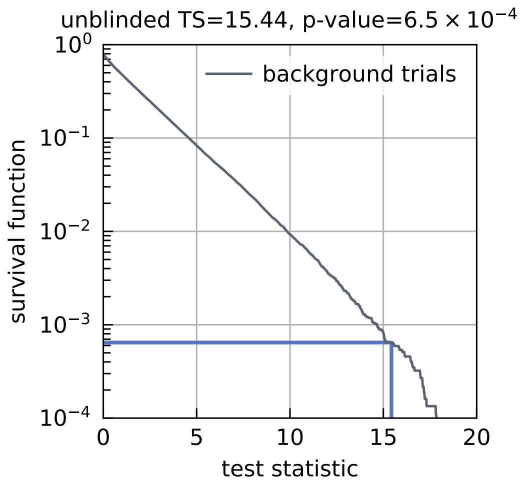

(b) A post-unblinding check where IceCube-170922A is added back into the sample. The p-value p is now expressed in terms of $n\sigma$, where $p = \int_{n\sigma} \mathrm{dts}\ \mathrm{norm}(\mathrm{ts})$ using the normal distribution.

Figure 10.2.: The test statistic TS of the unblinded fit and its translation into a p-value via the background test statistic distribution $\frac{\mathrm{d}n_{\mathrm{BG}}}{\mathrm{dTS}}$.

		without IceCube 170922-A	with IceCube 170922-A
	$\hat{n}_S$	7.02	10.15
	$\hat{\gamma}$	4.00 (parameter bound)	2.06
	$\hat{\Phi}_0$	$1.14 \times 10^{-7}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$	$0\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$
$\rightarrow$	TS	4.00	15.44
$\rightarrow$	p-value	12.6% (1.1σ)	6.5×10^{-4} (3.2σ)

Table 10.1.: Results (best-fit parameters, test-statistic, and p-value) of the TXS 0506+056 analysis.

IceCube-170922A back in, these background trials change little, but the unblinding fit does, due to this alert's space-time-coincidence with the 2017 flare of TXS 0506+056. The latter can now also be seen in fig. 10.1, where the event has been added and signal-over-background weights calculated assuming the resulting best-fit spectral index $\gamma = 2.06$. The test statistic of 15.44 is further in the tail of the background distribution in fig. 10.2b, reaching $p = 6.5 \times 10^{-4}$ (3.2σ). This also falls short of the discovery-level significance 5σ, but is similar to the significance calculated in [125]; there, the same coincidence is analysed in the context of similar alerts and more Fermi sources.

The values of the best-fit likelihood parameters are summarized in table 10.1. It is interesting that when including IceCube 170922-A, the fit prefers $\Phi_0 = 0\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, allowing additional, lower-energy events to contribute to the likelihood.

10.1.2. Local sky scan

The local scan announced in section 8.4.4 is also performed on the unblinded data sans IceCube-170922A. The resulting p-value map fig. 10.3 contains no discovery-level result within ±0.3° of the originally unblinded source direction. Thereby, an angular reconstruction which may be biased up to this scale is taken into account when reporting the null result of this analysis.

10.1.3. Limits

According to section 7.3.4, limits are calculated for a neutrino signal hypothesis following an E^{-2} spectrum and a time distribution corresponding to a range of thresholds applied to the lightcurve. In terms of number of neutrino events,

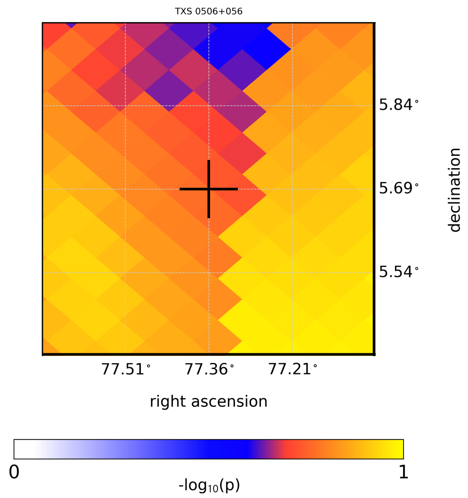

Figure 10.3.: A grid scan of alternative source hypothesis directions within ±0.3° of TXS 0506+056. Each pixel corresponds to a p-value, assuming the same background test statistic distribution.

their threshold dependence is shown in fig. 10.4a and follows the same evolution as sensitivities and discovery potentials (see section 8.3.1). Limits which include IceCube-170922A are higher, corresponding to the unblinded test statistic in this case.

Other analyses, data samples and physical hypotheses can not be compared based on μ. This is no longer the case for $\Phi(t, E)$, which can be calculated (see section 7.2.4) but the possible comparisons at the same t and E remain limited. Therefore, this signal hypothesis is summarized in different ways. Based on the livetime-integrated flux

$$\Phi_{\text{int}}(E) := \int_{\text{livetime}} \mathrm{d}t \Phi(t, E) \tag{10.1}$$

one can calculate the **average flux**,

$$\langle \Phi \rangle (E) := \frac{\Phi_{\text{int}}(E)}{\text{livetime}}, \tag{10.2}$$

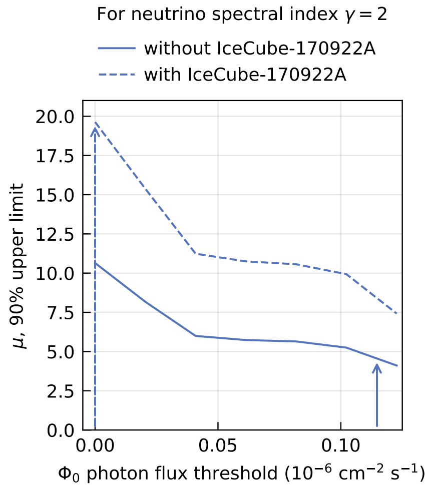

(a) Number of signal events in the sample.

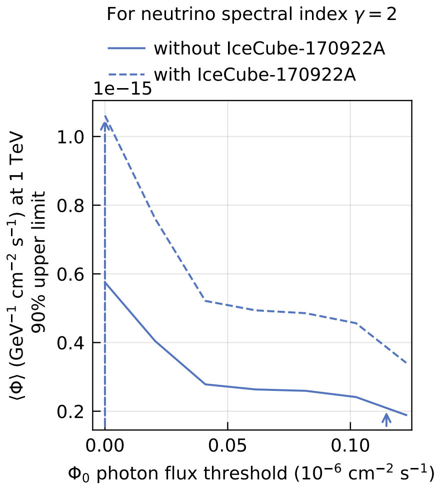

(b) Average flux during the livetime.

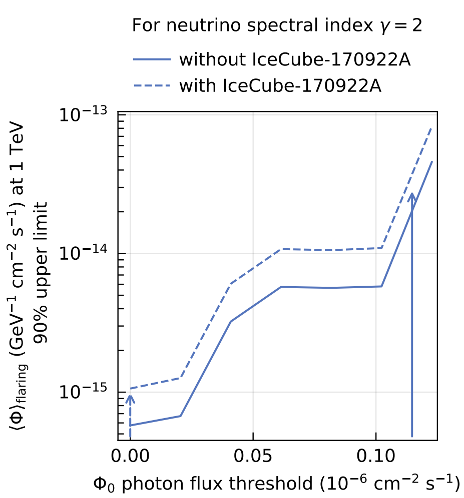

(c) Average flux during the livetime when the source is on.

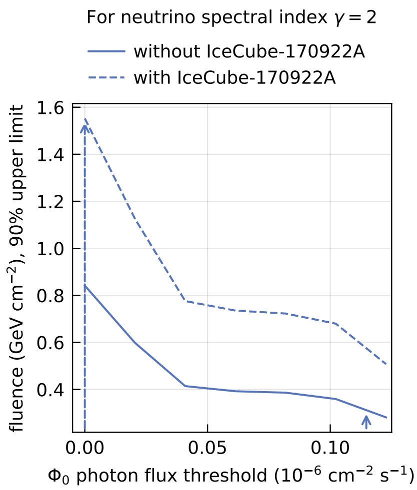

(d) Fluence within the central 90% energy interval of the source's signal in each sample.

Figure 10.4.: 90% CL upper limits for the TXS 0506+056 lightcurve correlation analysis, assuming a spectrum of $\gamma = 2$. The upper limit is shown as a function of the photon flux threshold, ranging from $\Phi_0 = 0$ where signal neutrinos correlate to the whole lightcurve, up to thresholds where they arrive only during the highest parts of the 2017 flare. Vertical lines show the best fit thresholds: without IceCube 170922-A (solid), and with it (dashed). The choices of y-axis are described in in section 10.1.3.

whose limits are presented in fig. 10.4b. One further quantity is the **flaring flux**, $\langle\Phi\rangle_{\text{flaring}}(E)$, analogous to presenting the flux Φ_{on} for a source that is steadily $\Phi = \Phi_{\text{on}}$ within a time interval, and $\Phi = 0$ for all other times:

$$\langle\Phi\rangle_{\text{flaring}}(E) := \frac{\Phi_{\text{int}}(E)}{\int_{\text{flaring}} \mathrm{d}t} \tag{10.3}$$

$$\text{where } \{\text{flaring}\} = \{t : t \in \text{livetime} \wedge \Phi(t, E) > 0\}. \tag{10.4}$$

Due to this normalization, the resulting limits in fig. 10.4c now increase with threshold. The general principle of integrating over times relevant to the observed signal is also applied to the energy integral when defining the **fluence**, already introduced (see section 8.3.1):

$$\mathcal{F} := \sum_j \int_{\text{livetime}_j} \mathrm{d}t \int_{E_j^-}^{E_j^+} \mathrm{d}E E \times \Phi(t, E). \tag{10.5}$$

Figure 10.4d shows the corresponding limits.

10.2. Blazar flare stacking

10.2.1. Unblinding

The complete unblinding of the analysis in chapter 9 comprises a total of six different fits, for the two source lists and three weighting schemes. All resulting parameters are listed in table 10.2. The highest test statistic value is 1.6, or 9.7 when adding TXS 0506+056.

The test statistic values are then translated into individual p-values via the corresponding background trials (see section 7.2.2), seen in figs. 10.5a and 10.5b. Their statistical uncertainty due to the available number of background trials N is calculated with an approximation, valid when p is far both from 0 or 1 [234]:

$$\sigma_p = \sqrt{\frac{p(1-p)}{N}} \tag{10.6}$$

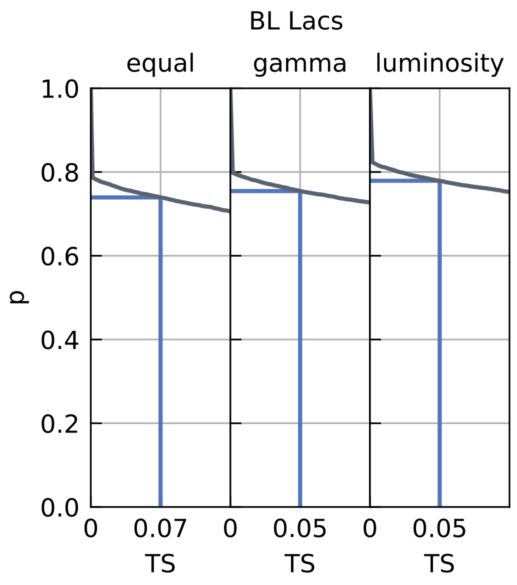

(a) The p-value calculation $p(\mathrm{TS})$ for BL Lacs in their three weighting schemes (grey), with the unblinded results (blue).

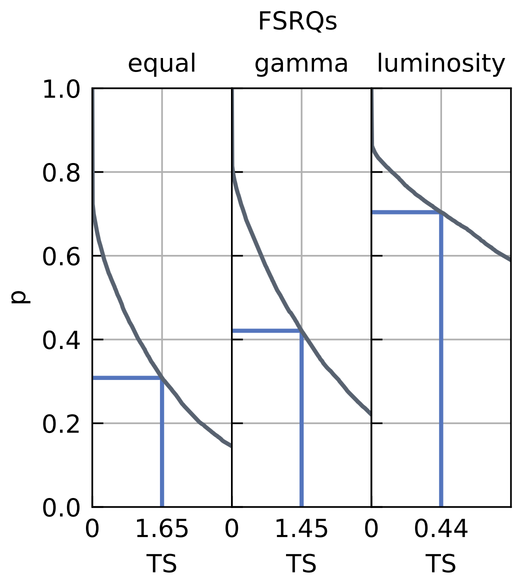

(b) The same as (a) for FSRQs.

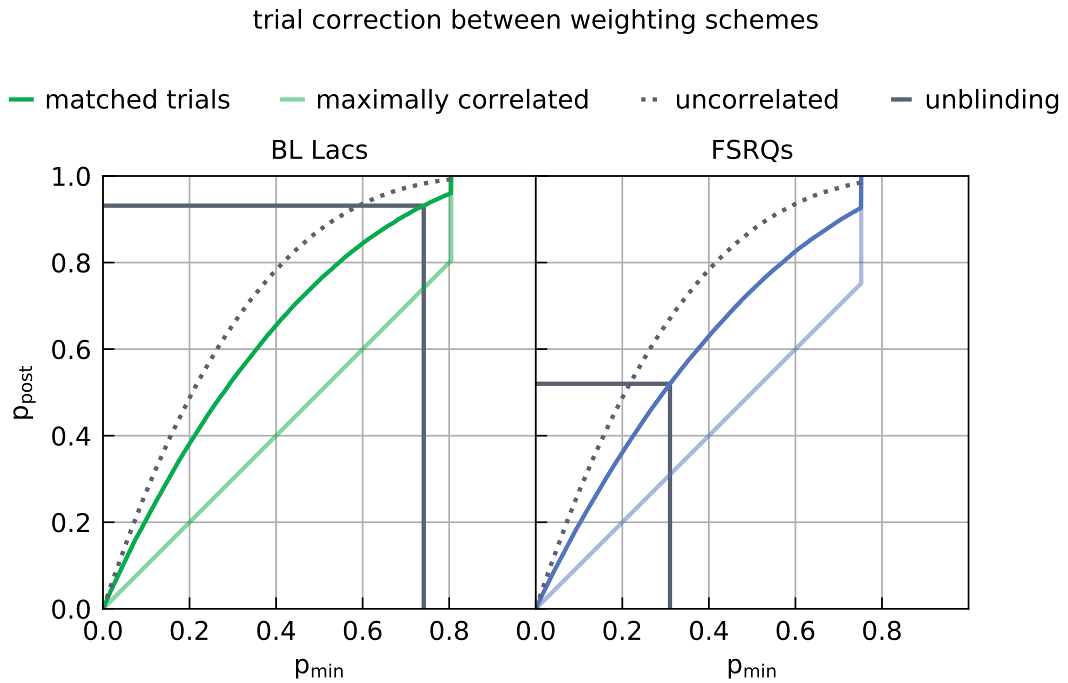

(c) After matching where the same background data realization shown above has been fit by each weighting scheme, given the smallest p-value p_{min} out of those which result, these plots show the trial correction $p_{\mathrm{post}}(p_{\mathrm{min}})$ as solid coloured lines. The lighter solid lines represent an example where the p-values are maximally correlated, i.e. the weighting schemes are equivalent. The dotted grey line is the other extreme, where fitting the same data realization results in completely uncorrelated results for each weighting scheme. The grey solid lines show the unblinded result.

Figure 10.5.

source class	weighting	$\hat{n}_S$	$\hat{\gamma}$	$\hat{\tau}$	TS	p	
	gamma	2.07	1.58	6.99	9.07	1.0%	
BL Lacs + TXS	luminosity	1.38	1.31	7.10	5.90	5.5%	1.9%
	equal	2.94	1.67	7.04	9.70	0.8%	
	gamma	0.44	4.00	74.41	0.05	75.4%	
BL Lacs	luminosity	0.44	4.00	74.49	0.05	77.9%	93.1%
	equal	0.55	4.00	68.64	0.07	73.9%	
	gamma	3.36	3.13	103.01	1.45	42.1%	
FSRQs	luminosity	1.78	2.77	131.12	0.44	70.4%	54.3%
	equal	3.37	3.07	101.97	1.65	30.8%	
	gamma	3.80	1.84	−4.42	4.36	11.2%	
FSRQs + TXS	luminosity	1.78	2.77	131.12	0.44	70.4%	4.6%
	equal	3.57	1.76	13.10	7.62	2.0%	

Table 10.2.: Best-fit likelihood parameters and resulting test statistic of the unblinding (section 10.2.1) and post-unblinding check (section 10.2.2). The resulting p-values and their trial corrections are also shown. The smallest value of τ where all lightcurves are below threshold is 75.2 for the BL Lacs, 133.6 for FSRQs (both with and without TXS 0506+056).

Combined via the trial correction defined in section 9.2.2 and visualized in fig. 10.5c, they result in $p_{\text{BL Lacs}} = (93.1 \pm 0.1)\%$, $p_{\text{FSRQs}} = (54.3 \pm 0.3)\%$ (where the uncertainty is recalculated the same way). These two are now independent and are analytically combined into the ultimate p-value $p = (79.1 \pm 0.3)\%$.

10.2.2. Adding TXS 0506+056

As post-unblinding check, the unblinding is repeated while adding TXS 0506+056 back into either source list (see section 9.1.4). The resulting fit parameters are also listed in table 10.2, with the higher test statistic resulting in higher significance. After trial-correcting between weighting schemes, this is $p_{\text{BL Lacs+TXS}} = (1.9 \pm 0.1)\%$, $p_{\text{FSRQs+TXS}} = (4.6 \pm 0.1)\%$. The higher significance indicates that indeed, the flare stacking analysis is able to identify a coincidence similar to that of TXS 0506+056 and IceCube-170922A. One exception to this is the FSRQ source list in the luminosity-squared weighting scheme, whose test statistic (and p-value) is unchanged by adding TXS 0506+056. This is because the best-fit τ, unlike in the other cases, does not shift downwards to include this source (see table 10.2). This is a property of the likelihood landscape which was explored for this specific reason, and not a result of the numerical minimizer failing to find the true optimum.

10.2.3. Contributing events and sources

The blazar flare stacking TS (TS > 0, see eq. (7.25)) is the sum of many terms comprising different events in the sample:

$$\mathrm{TS} = 2\,\mathrm{sgn}\,n_S \sum_{\text{events}} \Delta\log\mathcal{L}. \tag{10.7}$$

However, these terms have a strongly skewed distribution, so that only a small number of events dominate the unblinded test statistic. We choose to select those by first restricting the sum to those $\Delta\log\mathcal{L} > 0$, and then finding the fewest events possible to make up half that sum. Between the source lists and weighting schemes, we find 7 different events, shown next to their closest source in fig. A.1.

There are 4 such sources between the different unblinding fits: OM 484, OT 081, PKS 0507+17, and TXS 0506+056. Their relative contributions to the sum $\sum_{>0} \Delta\log\mathcal{L}$ are shown in fig. 10.6.

As expected, when adding TXS 0506+056 to the source list the corresponding IceCube-170922A dominates the likelihood, except in one case. The other events which dominate are also close to the source and on-time with a flare, in accordance with the likelihood definition. However interesting these coincidences are, none of them are similarly significant as that of TXS 0506+056.

As already mentioned in the previous subsection, all but one of the unblinded fits which include TXS 0506+056 feature a lower fit τ. On the other hand, those without TXS 0506+056 prefer a τ near its maximum (see table 10.2), thereby selecting a smaller number of blocks from the entirety of lightcurves. This can explain the skew in the distribution of per-source contributions to $\sum_{>0} \log\mathcal{L}$, which is greater in the latter case (see fig. 10.6).

10.2.4. Limits

Each of the six unblinded test statistic values is translated into limits on the corresponding signal hypothesis, according to the method of section 7.3.4. Analo-

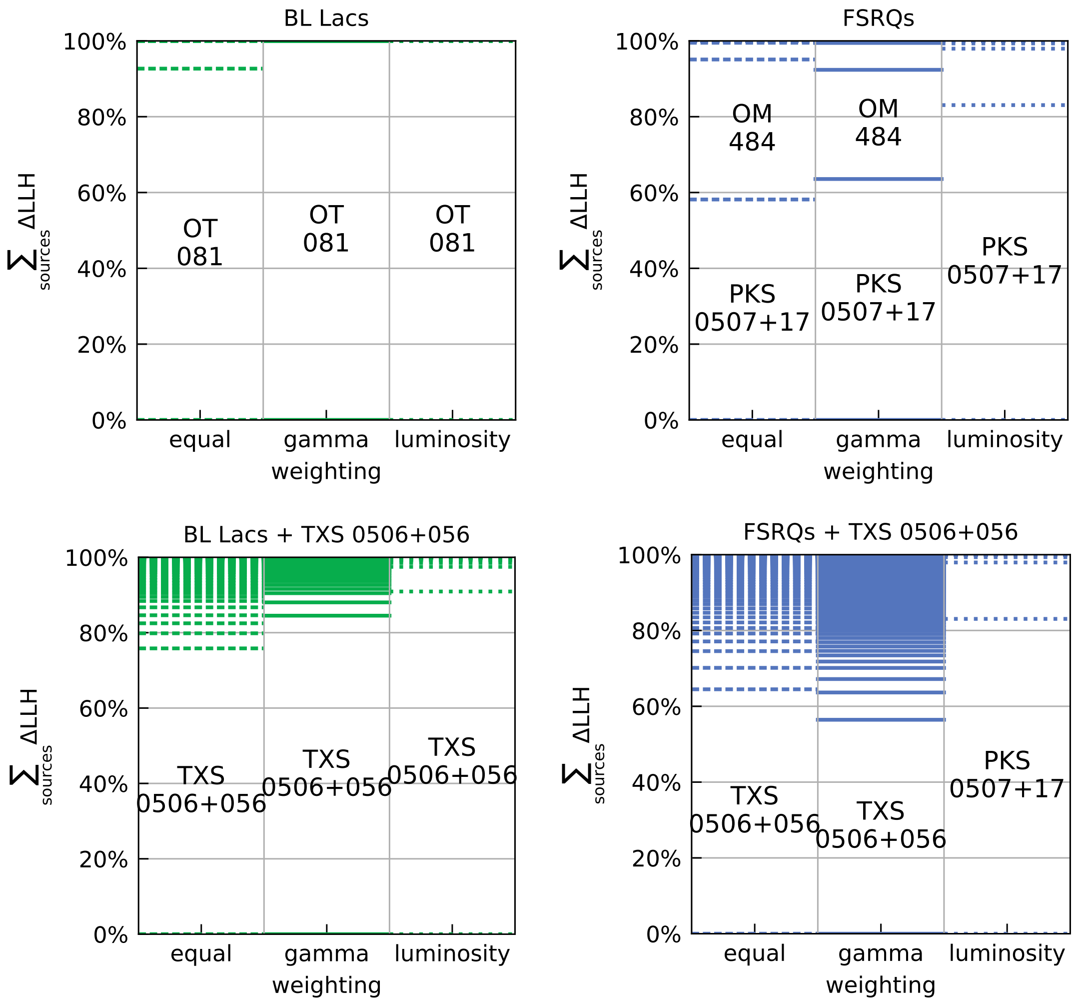

Figure 10.6.: The positive log-likelihood terms, broken down by the closest source. The y-axis is the cumulative fraction of the total sum.

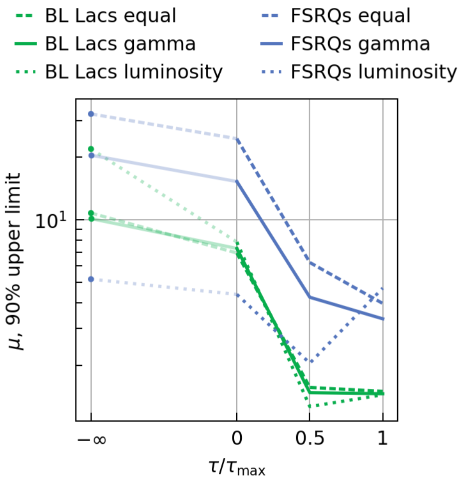

(a) Number of signal events in the sample.

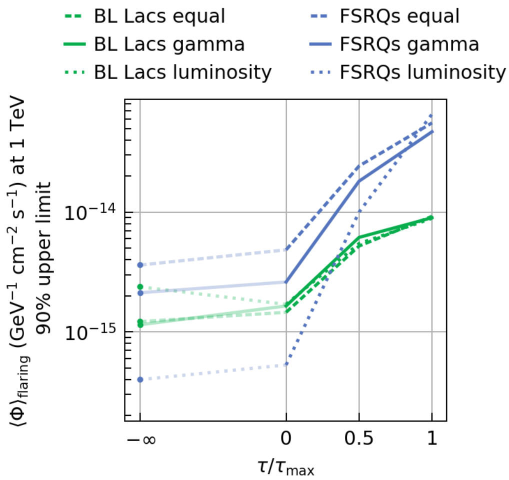

(b) Average flux during the livetime when the source is on.

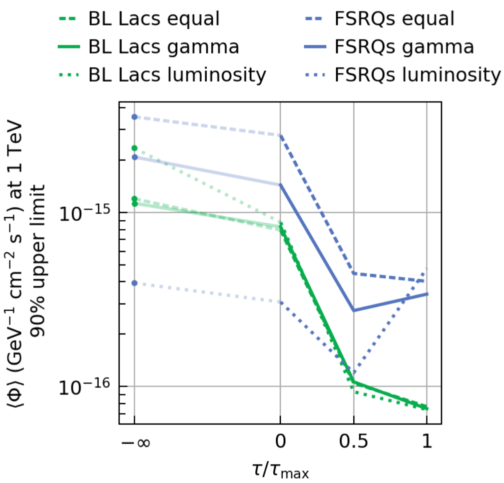

(c) Average flux during the livetime.

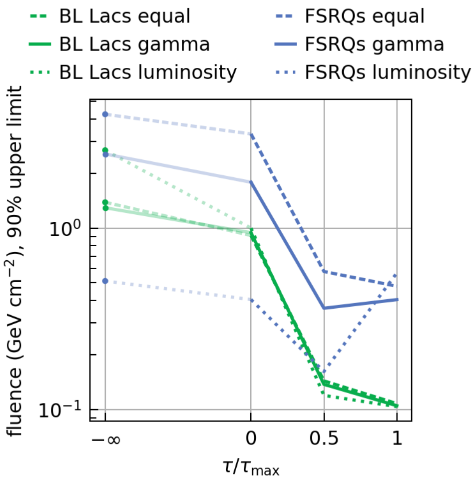

(d) Fluence within the central 90% energy interval of the source's signal in each sample.

Figure 10.7.: Limits of the blazar flare stacking analysis, for each source list and weighting scheme. They are represented by a choice of quantities (see section 7.2.4).

gous to section 10.1.3, this is actually a hypothesis space spanned by τ, which is sampled at certain points:

- $\tau = -\infty$, which corresponds to all thresholds $\Phi_{0,k} = 0$.
- $\tau = 0$, or all thresholds at the quiescent baseline, q_k (see section 7.1.5).
- $\tau = 0.5\tau_{\max}$ and $\tau = \tau_{\max}$, where $\tau_{\max}$ is the highest τ that still leaves one lightcurve block of the source list uncovered; this is 75.2 for the BL Lacs, 133.6 for FSRQs.

All these limits are shown in fig. 10.7. Here in the stacking case, we redefine the limit quantities of section 10.1.3 as the sum over all involved sources. Figure 10.8 also illustrates how they are derived, previously described in section 7.3.4: a signal is injected at different strengths μ. At any signal strength, the test statistic exceeds that of the unblinded fit TS_0 for a fraction β of signal trials, called the power. A stronger signal leads to higher TS and therefore higher $\beta(\mu)$, which leads to a monotonously increasing curve in the figure. The limit at 90% C.L. is defined such that $\beta(\mu^{(90)}) = 90\%$. The statistical uncertainty on β is quantified by approximating that the number of signal trials $> \mathrm{TS}_0$ per injected number of signal events is Poisson-distributed. Projecting this onto $\beta(\mu)$, it can be translated into a proxy for the uncertainty of $\mu^{(90)}$, as shown in the figure by vertical lines. Figure A.3 shows that all these uncertainties relative to the respective limits fall in the range of 7%–26%.

10.2.5. Limit discussion

Generally the limits on μ decrease with the injected τ, which reflects a likelihood that gives higher weight to a smaller number of events which occur while the blazar lightcurves are high.

The definitions of these three weighting schemes are degenerate at $\tau = \tau_{\max}$, which is reflected in the signal injection, and in the likelihood to the extent that the fit τ manages to approach the injected. Overall, this explains how the limits approach each other at $\tau_{\max}$. At $0.5 \times \tau_{\max}$, there is already a minority of sources which exceed their threshold at any time (3/64 BL Lacs, 9/114 FSRQs). Here, the luminosity weighting scheme breaks the pattern of the τ dependency, especially for FSRQs.

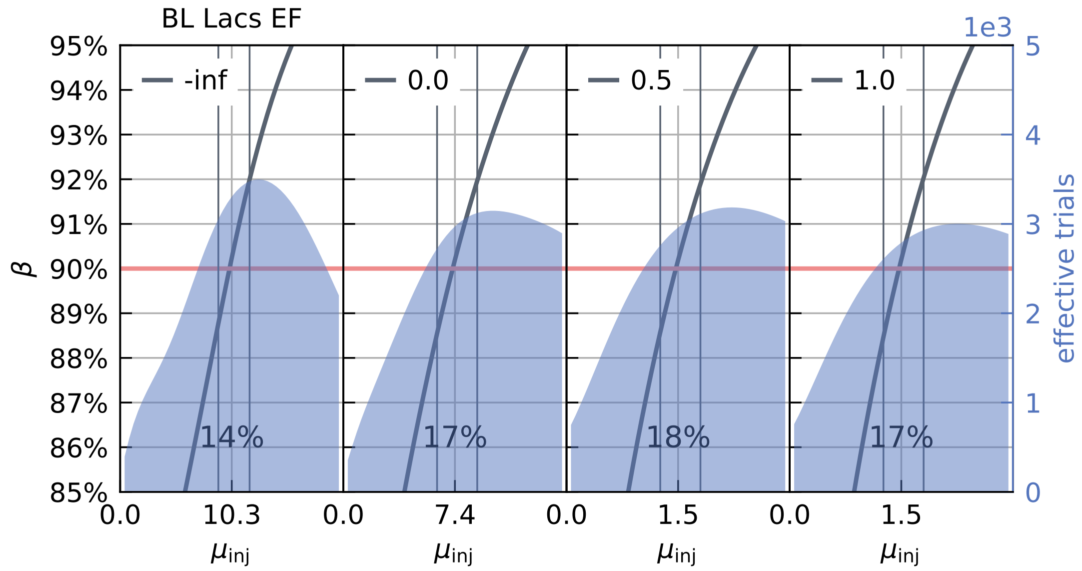

Figure 10.8.: The origin of the limits of section 10.2.4 as the crossing of the $\beta(\mu)$ curve (grey) with 90% (red). Also shown is the effective number of signal trials (blue) whose reweighting produces the curve.

An attempt was made to understand this by comparing the neutrino samples that contribute to the various limits. Figure A.2 captures the relative importance of these samples in terms of the number of signal events injected, n_{inj}. At low τ, this strongly reflects the effective area and livetime of the sample, as IC86b and IC86c covered 3 and 2.5 seasons respectively instead of one, like the other samples. Given the FSRQ source list at $\tau/\tau_{\text{max}} = 0.5$ (9 sources), we see n_{inj} affected by the particular set of lightcurves which still exceed their thresholds. Here, the luminosity weighting scheme is dominated by IC86c, whereas the other weighting schemes have a stronger contribution from IC86b. In comparison, the luminosity weighting also shows an absence of injected events from samples IC79 and earlier. Those used an incomplete detector, and hence feature worse sensitivity. This again comes through when continuing to the extreme case of $\tau/\tau_{\text{max}} = 1.0$ with one single lightcurve block surviving in all weighting schemes, and IC79 the only sample injected.

The efficient sampling methods required are validated in two ways. First, we can check how well the region of $\mu^{(90)}$ is covered by the signal trials with injected number of events $\{n_{\text{inj},i}\}$. For this, the number of effective trials is defined as

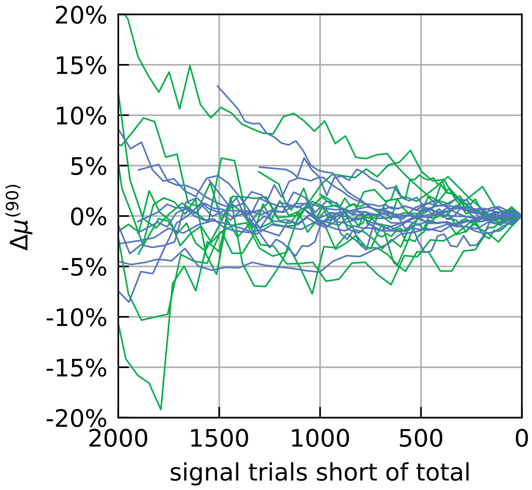

Figure 10.9.: For all calculated limits, their relative change is shown which would result from removing up to 2000 of the signal trials. No limit changed more than 5% in the last 347 trials, or more than 9.4% in the last 1000.

$$N_{\text{eff}}(\mu) := \frac{\sum_k N_k P_k(k;\mu)}{\sum_k P_k(k;\mu)^2} \tag{10.8}$$

where P is the Poisson distribution, μ its expectation value, and N_k the number of signal trials with k injected events. This way, a Poissonian at μ itself sampled N times contributes N to the effective number, for sufficiently high N. Additional samples more spread out than that Poissonian contribute less, and those far enough outside its tails are virtually disregarded. Figure 10.8 shows this effective number for the signal trials used in the limit calculation for BL Lacs under the gamma-ray energy flux weighting. There, and for the rest in fig. A.3, we find that it does not decay by more than half within the uncertainties estimated for $\mu^{(90)}$ (see section 10.2.4).

The second validation is based on examining the convergence of $\mu^{(90)}$ when adding additional signal trials. This can be seen in fig. 10.9, where the scale of fluctuations becomes increasingly suppressed.

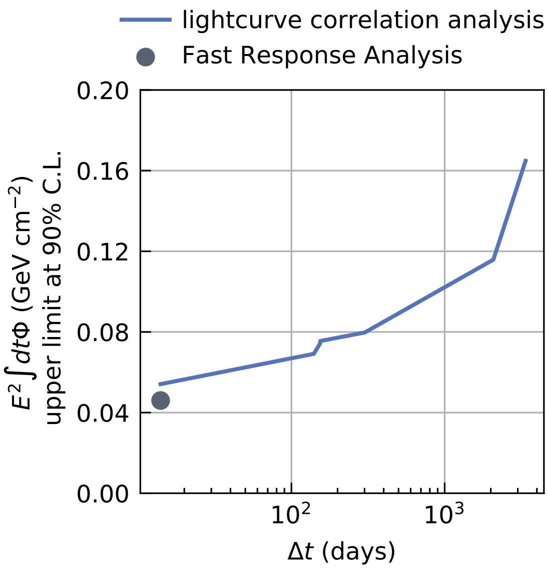

Figure 10.10.: The limits from two time-dependent analysis focusing on the 2017 flare of TXS 0506+056. One is the FRA, which selected a window of $\delta t = \pm 7$ days around IceCube 170922-A, the other is the lightcurve correlation analysis. There, δt is the duration the lightcurve spends above the injected threshold.

10.3. Discussion

10.3.1. TXS 0506+056 in the context of its other neutrino observations

The upper limit on the time-dependent hypothesis in this analysis can be compared to the other neutrino observations of the same source, reported in section 8.1.

The neutrino flare in 2014, which is also reflected in the time-integrated analysis, can only be seen as complementary. The events which give it 3.5σ significance arrived during a low-emitting state of the blazar. In the present analysis results, they thus receive low weights and contribute little. In the specific case of unblinding without IceCube 170922-A, the best-fit threshold Φ_0 exceeds that part of the lightcurve and so the neutrino flare has no impact on the resulting limit.

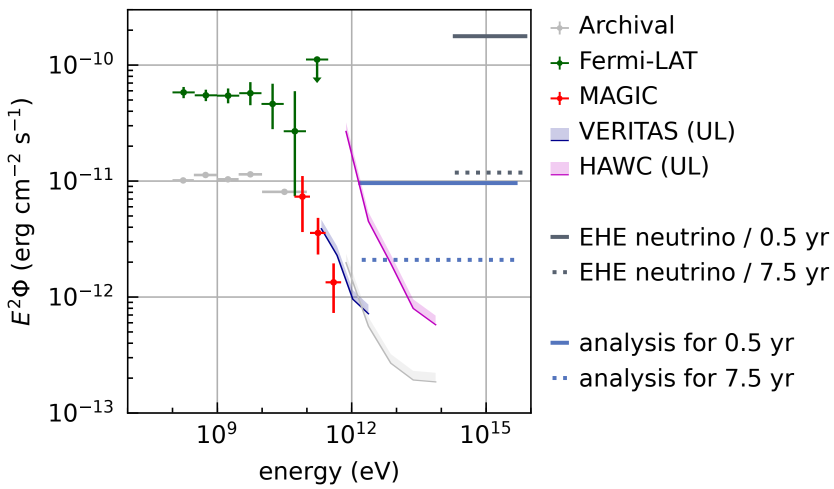

Figure 10.11.: Neutrino limits assigned to a 90% energy range, in the context of the high-energy electromagnetic SED of TXS 0506+056, both within 14 days of IceCube 170922-A and archival. The latter is taken from [125], along with an upper limit interpretation of the EHE event-flare coincidence.

The FRA resembles the present analysis, using a similar data sample at an earlier stage of processing and excluding IceCube 170922-A. Instead of using a threshold to isolate the high part of the lightcurve, it manually restricts itself to a time window of ±7 days within IceCube 170922-A. The single highest block of the lightcurve has a width of 14 days, and so the limit given the maximum $\Phi_0 = 1.2 \times 10^{-7}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ can be compared with that from the FRA, as both analyses report limits on a E^{-2} spectrum. Presented in terms of $E^2 \int \mathrm{d}t\Phi$, the limit in our analysis is 17% larger, as seen in fig. 10.10. This is comparable, despite specific differences. Principally, the FRA likelihood only had the signal strength as a free parameter, whereas the lightcurve analysis had to fit the injected spectral index and the threshold parameter to the time PDF. The x-axis of fig. 10.10 shows at once the injected duration of the lightcurve above threshold, and for the FRA the length of the predefined time window where the search was performed.

The coincidence with IceCube 170922-A is also similar to the signal sought in the present analysis, but restricted to the higher end of the energy range. We can compare it to the results of the search for correlation with the entire lightcurve, in the case where IceCube 170922-A remains in the data sample (see section 10.1.1).

Upper limits on the flux to produce one similar EHE event during 7.5 or 0.5 years of data-taking are presented [125]. This time span corresponds to signal hypotheses assuming a threshold of $\Phi_0 = 1.8 \times 10^{-8}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ or $\Phi_0 = 4.5 \times 10^{-8}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ respectively. The limits on these hypotheses are chosen for the comparison in fig. 10.11. Of the quantities defined in section 10.1.3, the **flaring flux** corresponds to the principle of the event-based estimate, based on which we compute the SED $E^2\Phi_{\mathrm{flaring}}$ as in the original MWL plot from [125]. We find limits a factor of 6–18 stronger. The same figure shows that the source's gamma-ray SED during the flare is comparable to our limits.

This suggests that the EHE neutrino itself was not representative of a lightcurve-correlated E^{-2} neutrino flux at EHE energies from TXS 0506+056. Besides the possibility of of being a pure background fluctuation, the following hypotheses can be invoked to make the two compatible:

- The spectrum does not follow E^{-2} across the entire considered range, but is skewed towards higher energies. Importantly, the two sets of limits in fig. 10.11 do not apply to the same energy range. For IceCube 170922-A, it shows the central 90% interval of the event's estimated energy given a E^{-2} signal spectrum prior. For the lightcurve analysis limits, there's the central 90% interval of the events making up a E^{-2} signal in IceCube. Although the latter range largely covers the former, the differential limit would be less stringent at the higher energies, easing the tension.

- The neutrino production follows another time PDF than the total gamma-ray variability, for any subtracted threshold. An example of this are the models designed to predict neutrino emission during the quiescent states of this blazar.

- TXS 0506+056 can be considered as a member of a population of similar sources, such as blazars. In this case, the detection of a single (EHE alert) event could result from a Poissonian expectation of ~ 1 in the entire population, and thus lower for TXS 0506+056 individually. This effect is known in astronomy as the **Eddington bias**, and also described in neutrino astronomy in terms of few-event observations [225]. This possibility was already considered in the plot from which fig. 10.11 is adapted [125]. The

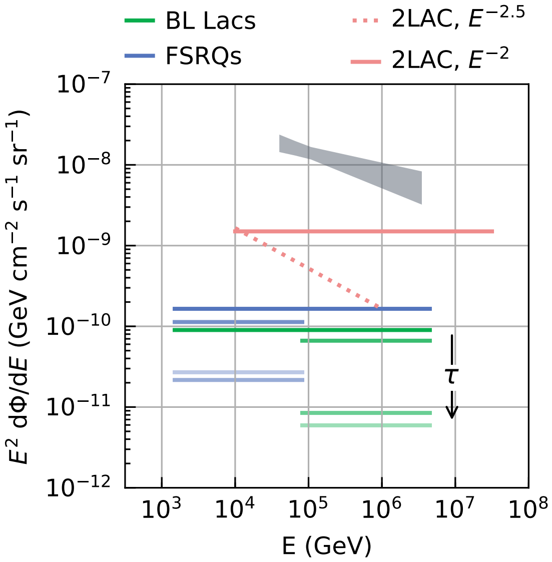

Figure 10.12.: The blazar flare stacking limits in the gamma weighting scheme, as the time-averaged SED per solid angle. The BL Lacs (green) and FSRQs (blue) have limits for $\tau_{\text{inj}} = -\infty$ (horizontal line) and $0 < \tau_{\text{inj}} < \tau_{\text{max}}$ (gradient, shown only in half of the energy range). The astrophysical $\nu_\mu + \overline{\nu}_\mu$ flux (grey) lies 2–3 orders of magnitude higher [61]. Limits from a time-integrated stacking [228] (red) are included for an E^{-2} spectrum and $E^{-2.5}$ from the astrophysical fit [235].

blazar flare stacking analysis then addresses the case of a neutrino flux shared between a larger population of sources.

10.3.2. Blazar flare stacking in the context of the astrophysical neutrino flux

Time-integrated stacking neutrino source searches have reported an upper limit on the average flux per solid angle in the region of the sky covered by their respective source list [228]. This quantity can be compared to the observed astrophysical $\nu_\mu + \overline{\nu}_\mu$ flux, compatible with being isotropic in the Northern [61] and Southern hemisphere [60]. In the case of this analysis, the BL Lacs and FSRQs it comprises were selected from the entire sky, corresponding to a solid angle $\Omega = 4\pi$.

Referring to section 7.2.4, the limit on the time-averaged flux is the right choice to answer the question how much the analysed sets of blazars contribute to the neutrino sky over the course of a year, not during particular flares. The definition we have adopted for the stacking case, summing the contribution of each source's flux in the combined hypothesis, is also the result of an integral over a sky which contains these point sources.

As the search focused on track events (see chapter 3), the appropriate astrophysical flux to compare with is that of muon (anti-)neutrinos [235, 61, 236]. The time-integrated stacking includes a **gamma** source weighting scheme [228] which is equivalent to ours at $\tau = -\infty$, except that it uses Fermi-LAT gamma rays $> 0.1\,\text{GeV}$ instead of $> 1\,\text{GeV}$. The range of values for $0 < \tau < \tau_{\text{max}}$ is represented by a gradient in fig. 10.12. The time-integrated stacking sets limits on a $E^{-2.5}$ spectrum like one astrophysical fit [235], as well as E^{-2} like the present analysis. The latter spectral index is only a test case which can be motivated from acceleration physics (see eq. (1.16)), and for the astrophysical neutrino flux as a whole is disfavoured at 3.8σ [235]. Our limits can thus not be interpreted as a fraction of the astrophysical flux, and neither was this the intention.

These limits are at least an order of magnitude lower than those set by the time-integrated stacking. However, this does not compare the respective analysis methods' strengths, as they apply to different signal hypotheses. Either can be expressed as an average flux, but the blazar flare stacking concerns emission from a smaller number of sources, assuming specific time-dependence where the time-integrated makes no assumptions. This fact is also illustrated by how the limits are even lower when assuming a higher threshold, which is a choice of signal hypothesis within the same analysis. The analysis was designed with discovery in mind, not to extrapolate its limits onto a general physical hypothesis involving the entire class of resolved and un-resolved blazars.

10.3.3. Comparing TXS 0506+056 to the other stacked Fermi blazars

The post-unblinding check of the blazar flare stacking analysis (see section 10.2.2) has shown that it is sensitive to multimessenger coincidences of the type exhibited by TXS 0506+056. The upper limits obtained can be divided by the corresponding number of sources (64 BL Lacs, 114 FSRQs, with redshift 49 and 110 respectively). This would be the upper limit on the flux of an individual, average blazar in the list, assuming that they all contribute to the neutrino emission as expected by the stacked signal hypothesis.

On the other hand there is the upper limit on lightcurve-correlated emission of the blazar excluded from the list, TXS 0506+056. In the case of including the event IceCube 170922-A, that limit also reflects said coincidence.

For a comparison where all lightcurves are untruncated, we choose a signal hypothesis of $\tau = -\infty$, $\Phi_0 = 0$. This gives

$$\langle \Phi^{(90\%)}_{\text{BL Lac}} \rangle = 1.8 \times 10^{-17} \text{ — } 4.9 \times 10^{-17}\,\text{GeV}^{-1}\ \text{cm}^{-2}\ \text{s}^{-1} \tag{10.9}$$

$$\langle \Phi^{(90\%)}_{\text{FSRQ}} \rangle = 3.6 \times 10^{-18} \text{ — } 3.2 \times 10^{-17}\,\text{GeV}^{-1}\ \text{cm}^{-2}\ \text{s}^{-1} \tag{10.10}$$

$$\langle \Phi^{(90\%)}_{\text{TXS}} \rangle = 1.1 \times 10^{-15}\,\text{GeV}^{-1}\ \text{cm}^{-2}\ \text{s}^{-1} \tag{10.11}$$

at an energy of 1 TeV, with the range in the stacking corresponding to different weighting schemes.

The ratio of the TXS 0506+056 limit vs. BL Lacs (22–60) and vs. FSRQs (34–297) can be understood to say that the average stacked blazar is less correlated to neutrino arrival times than TXS 0506+056 was, making it exceptional among bright, variable Fermi blazars. This assumption could already be made **a priori** since TXS 0506+056 was selected for analysis because of the same coincidence, while the stacked blazars were only selected based on gamma-ray observations (see section 9.1). Still in the sense of comparing limits, the source lists of the blazar flare stacking (depending on the respective choice of weighting scheme) can contain up to the equivalent of approximately 1.5–5.7 TXS 0506+056-level coincidences. These must be spread between multiple sources, and may therefore be the subject of future blazar flare stacking analyses.

"You should never compare yourself to other people."

Catherine de Clercq

11 Conclusions & Outlook

11.1. Conclusions

This thesis concerns two analyses. Both use IceCube to search for neutrino production in blazars, correlated to their gamma-ray emission as observed in Fermi-LAT lightcurves, smoothed with Bayesian blocks. The first analysis was motivated by an example of such a coincidence, the event IceCube 170922-A arriving during a flare of the blazar TXS 0506+056. The search reimplemented a prior likelihood analysis whose hypothesis allows for the blazar emission to be distinct between flares and a steady, quiescent flux, with the two separated by a threshold which is a likelihood parameter. Finding no significant excess this way besides the original IceCube 170922-A, we could set limits at 90% C.L. that $0.05\,\mathrm{GeV/cm^2}$ to $0.17\,\mathrm{GeV/cm^2}$ of lightcurve-correlated neutrino emission from TXS 0506+056 extends down to $\mathcal{O}(10\,\mathrm{TeV})$ energy. The limits depend on the assumed threshold in the signal hypothesis, but can still be interpreted as complementary to the other observations of this blazar. Those notably included a 3.5σ cluster of neutrinos during a quiescent period of the blazar, a signal to which the analysis is not sensitive.

The second analysis, begun before the first but refined afterwards was motivated by the past analyses which similarly failed to discover a neutrino signal. Under the assumption that at least a portion of the blazars in the Fermi-LAT sky are physically similar, the individually undetectable neutrino emission from their flares could still be discoverable when combined. This analysis technique is called stacking. A common parameter τ was used to assign each source a threshold, using parameters calculated on its lightcurve. In order to benefit this analysis, the selected blazars were required to show well-resolved variability, which led to source lists of 65 BL Lacs and 114 FSRQs, a selection skewing to brighter sources. However, TXS 0506+056 was not included as it was already unblinded in the first analysis. Combining either of these lists with a choice of three weighting schemes led to six unblinding fits, none of which significant. The combined trial-corrected p-value was $p = (78.5 \pm 1.0)\%$. A set of six limits for a range of hypotheses depending on τ can also be calculated. However it should be noted that these are primarily limits on a specific hypothesis applied to a limited set of sources. Unlike time-integrated stacking analyses which have managed to constrain the contribution of blazars to the observed astrophysical neutrino flux, this one had discovery as its purpose.

This was the first analysis in a new approach, and its design involved many choices. Future analyses may reconsider these choices, and better exploit the blazar flare stacking principle thanks to several possible improvements.

11.2. More analysis targets

Markarian 421: This is a nearby blazar, the third brightest source (above 100 MeV) in the Fermi 4LAC-DR2 [191], prominent in the TeV gamma-ray sky [186]. It is also highly variable in bands from X-rays to VHE gamma rays [237], as well as optical and radio [238]. The brightness and MWL variability make Markarian 421 a promising candidate for time-dependent neutrino searches [239]. Its exceptional brightness allows HAWC, a monitoring telescope, to provide lightcurves in the 0.5 TeV to 100 TeV band which reveal $\mathcal{O}(\mathrm{d})$ variability [240]. This enables a similar analysis as those we've performed using Fermi-LAT lightcurves ($\leq$ 100 GeV), but in an energy band closer to that of IceCube neutrinos.

Antiflares: There are models which calculate how a hadronic AGN jet would meet nearby dust or gas, undergo proton-proton (or proton-nucleus) interactions, and thereby produce neutrinos [54]. This would be proportional to the density of the dust or gas, integrated over the proton path. This in turn can be measured via the attenuation of X-rays, which are known to be produced in the AGN central region. In addition, any hadronically produced gamma rays are also attenuated. Beyond selecting AGN where such obscuration is found averaging over time [54], some AGN also show X-ray attenuation which varies on the scale of days. Also there, the cause of this X-ray **antiflare** is interpreted as dust or gas, for example belonging to the BLR [241]. As already noted in [54], this motivates a search for neutrino production during times where X-ray or gamma-ray emission dims. The searches in this thesis have instead focused on the orthogonal hypothesis of brightening gamma rays, but the mathematical and technical methods are almost entirely transferable.

11.3. Other data

X-ray lightcurves: Since the 3.5σ cluster of neutrinos during a low state of TXS 0506+056, there has been increased interest in models where neutrino production in blazars is correlated more to their X-ray emission than to that of gamma rays. This may be because these X-rays are understood as proton-synchrotron emission [242] or as the target photons for p-γ interactions [243, 244]. Other proposed models see the multimessenger connection between blazar X-rays and neutrinos as a key to understanding the underlying processes and geometry [245, 246], as a two-zone rather than a single-zone model may be required to accommodate the available data. Even when only assuming X-rays as part of the possibly hadronic high-energy blazar hump, the aforementioned Markarian 421 is an example of exhibiting higher variability in X-rays than in gamma rays [238], where thus an X-ray lightcurve would provide a stronger separation between quiescence and flares. Such lightcurves are now available for $\mathcal{O}(100)$ blazars thanks to Swift-BAT and other instruments [247]. X-rays can therefore be used for blazar flare stacking analyses.

IceCube-Gen2: The current IceCube detector has not yet discovered a neutrino point source, and after almost 10 years of data-taking in the full detector configuration (see chapter 3) the expected remaining $1/\sqrt{\text{time}}$ improvement on sensitivity has diminished. This can be understood as the astrophysical neutrino flux, discovered with IceCube itself, only exceeds that of atmospheric neutrinos above a few 100 TeV [61]. An array with a wider spacing would shift more of its sensitivity into this energy range*, and given the same number of modules could be larger, and thus detect neutrinos at a higher rate. A further consequence is a better angular resolution for muon tracks that are captured over a longer distance. Hence the planned 8 km^3 extension IceCube-Gen2 was designed to provide, after 8 years of construction and deployment, a 5 times better discovery potential (see section 7.3.5) towards neutrino point sources. This includes improved prospects for the detection of neutrinos from blazar flares [248].

*having a larger effective area to compensate for the lower flux, and still resolving and reconstructing the brighter events

A
Supplementary Figures

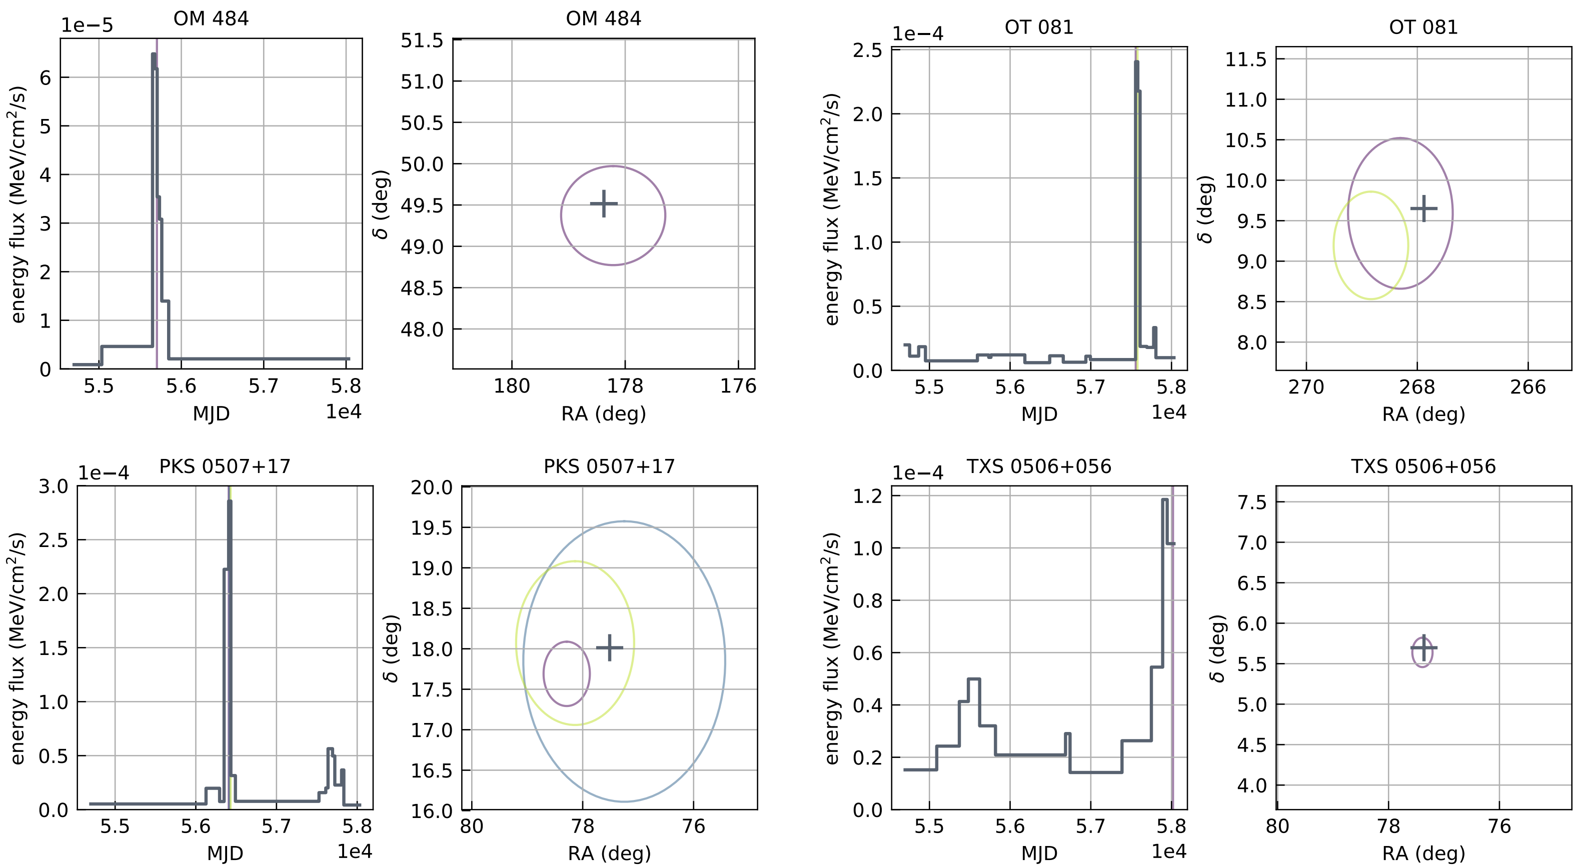

Figure A.1.: The seven events which together dominate the various unblinded fits, shown in the context of their closest source. For each source, the left panel is the lightcurve with vertical lines at the event time, coloured on a time scale between the earliest and latest event. The right panel is a map of right ascension and declination, using the same colours for error circles to show the event reconstructed direction and reconstruction uncertainty. The source direction is in the grey cross in the centre.

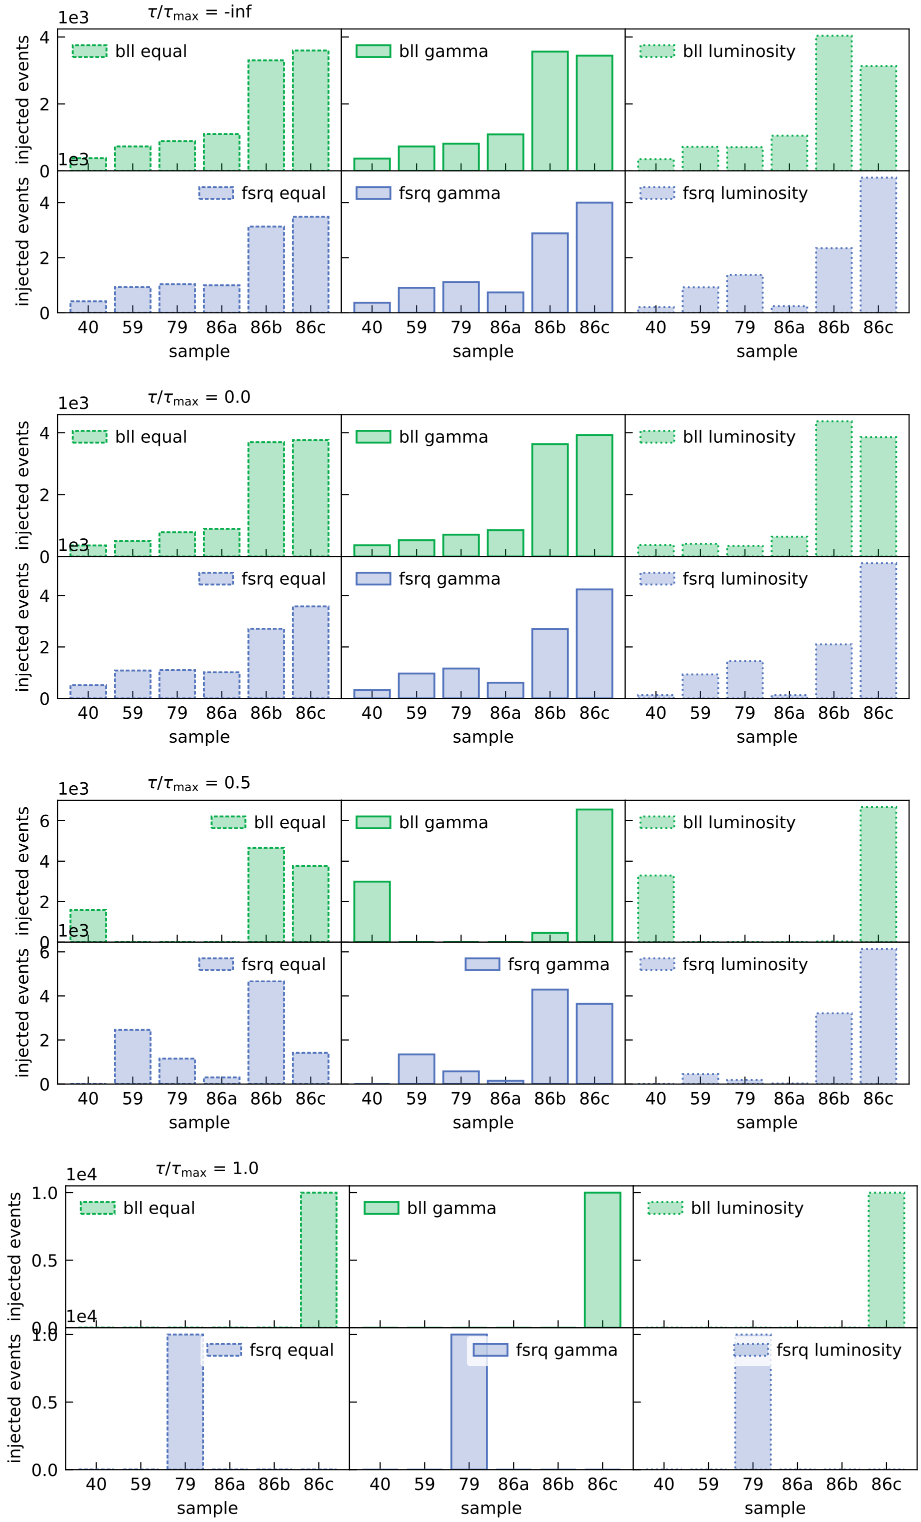

Figure A.2.: The distribution of 10^4 injected events between the neutrino event samples for each signal hypothesis used in the limit calculation.

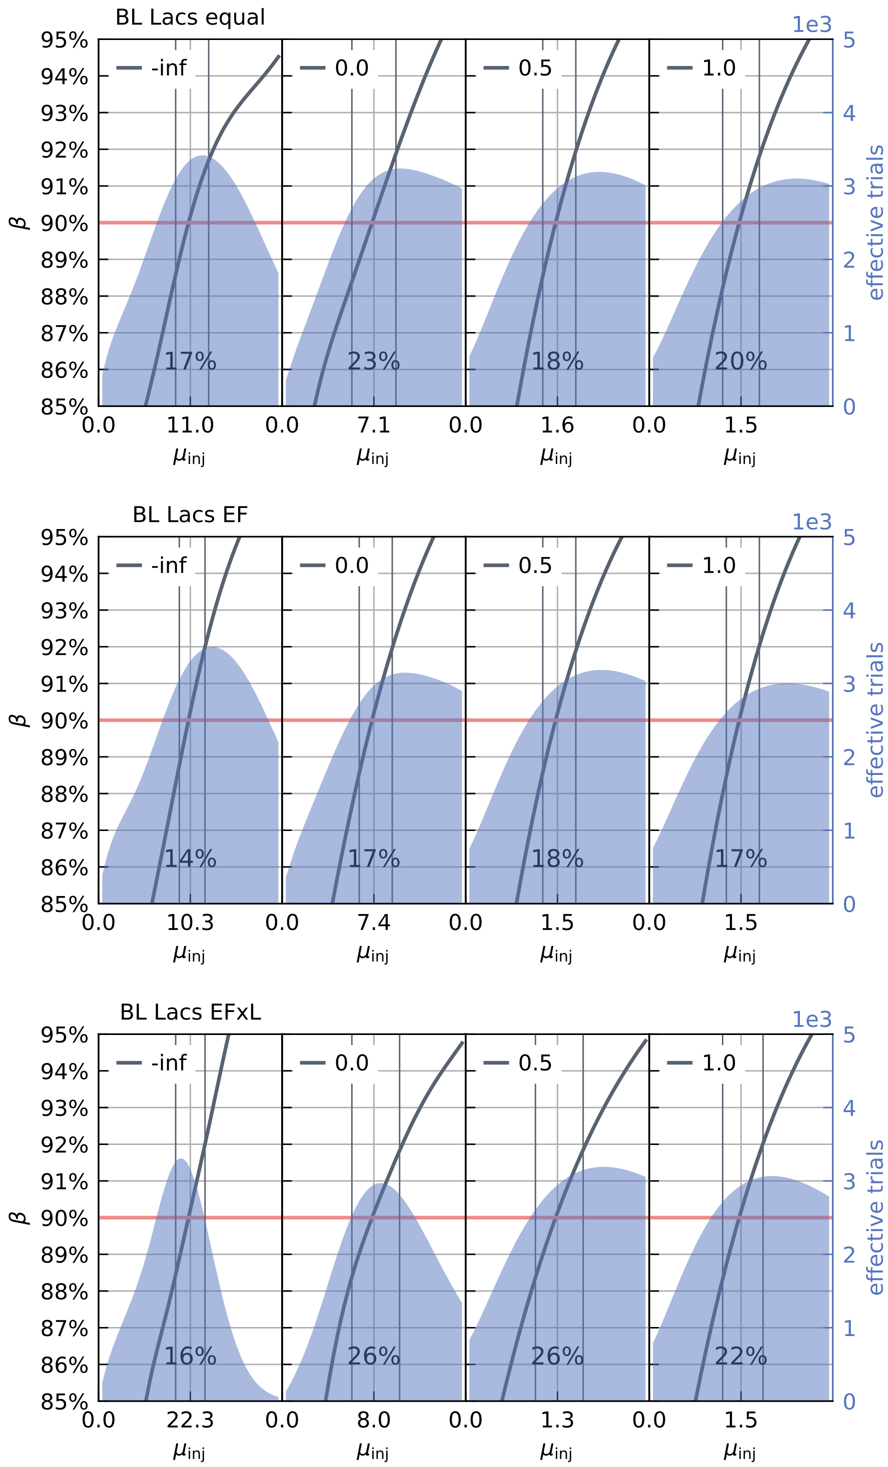

(a) BL Lacs

Figure A.3.: The origin of the limits of section 10.2.4 as the crossing of the $\beta(\mu)$ curve (grey) with 90% (red). Also shown is the effective number of signal trials (blue) whose reweighting produces the curve.

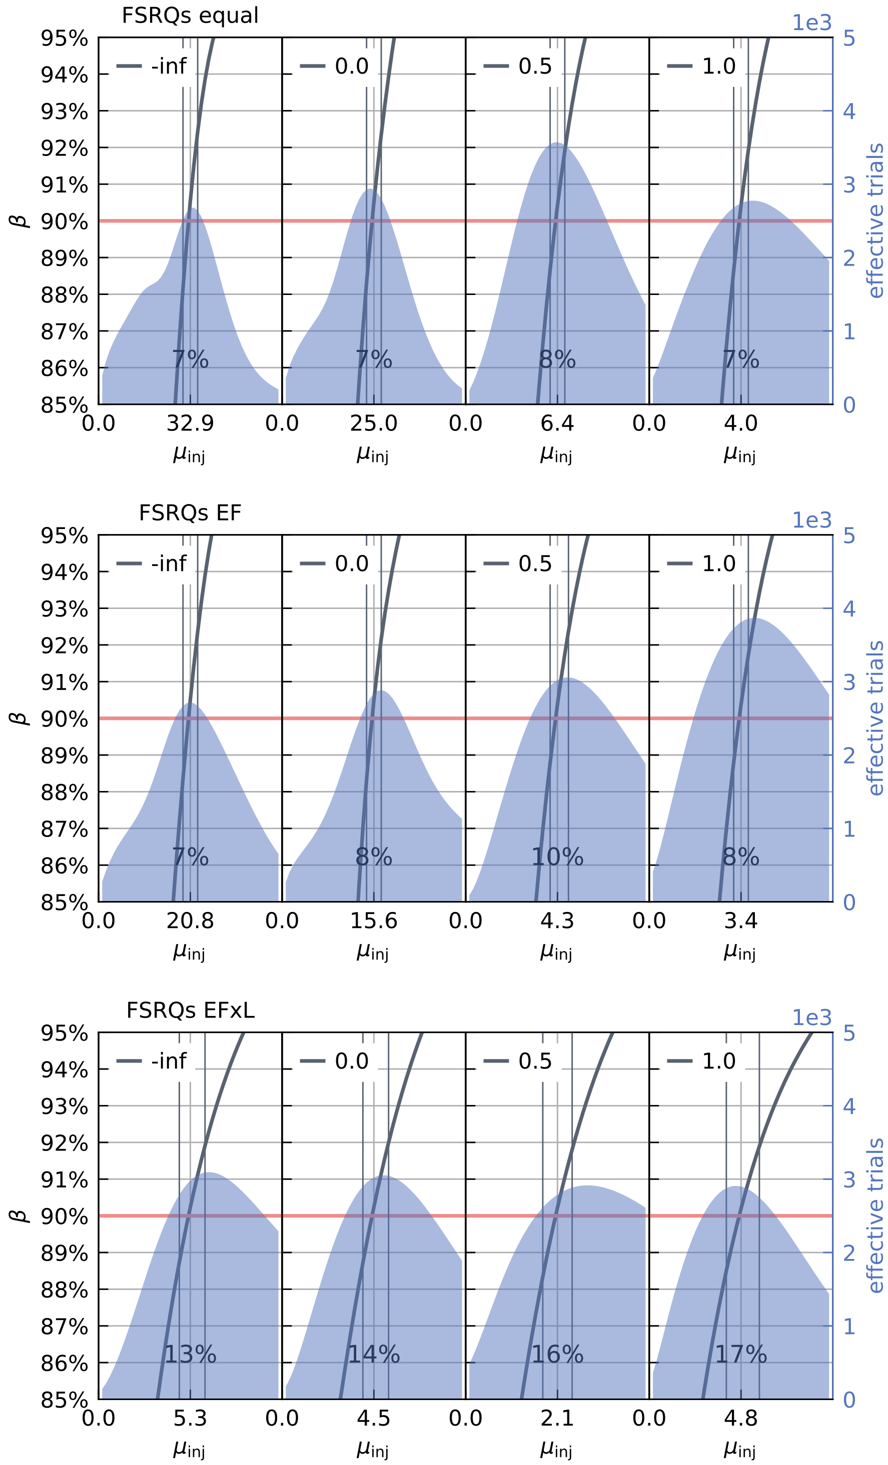
FSRQs equal
FSRQs EF
FSRQs EFxL
-inf
0.0
0.5
1.0
95%
94%
93%
92%
91%
90%
89%
88%
87%
86%
85%
β
1e3
effective trials
7%
7%
8%
7%
32.9
25.0
6.4
4.0
7%
8%
10%
8%
20.8
15.6
4.3
3.4
13%
14%
16%
17%
5.3
4.5
2.1
4.8
μ_{inj}

(b) FSRQs

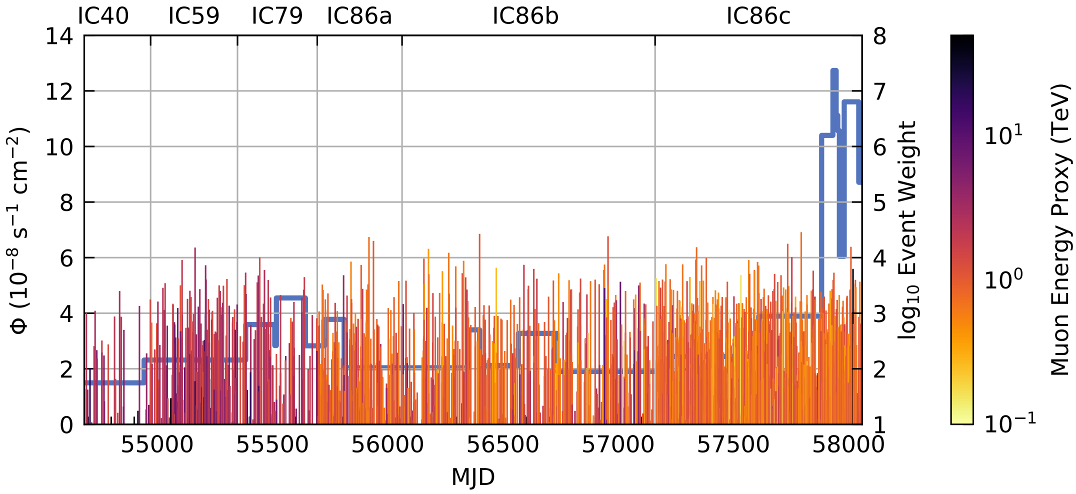

Figure A.4.: A version of fig. 10.1 where weights are calculated assuming a spectral index of $\gamma = 4$, as in the unblinding fit.

B
Source Lists

These tables state properties of the stacked sources. Each source list is sorted by the declination measured by Fermi-LAT in equatorial coordinates (see section 3.3.2). The source name belongs to the object associated by Fermi catalogues at the time the analysis was prepared. Redshifts are also provided by the Fermi catalogues as well as others (see section 9.1.1 for a detailed description). The Fermi energy flux is integrated in the band $[1, 300]$GeV and averaged over 9.5 years of observation.

B.1. BL Lacs

Table B.1.

Fermi-LAT J2000		associated source	energy flux	redshift (reference)	
RA (deg)	δ (deg)	**name**	**F** (MeV cm^{-2} s^{-1})	**z**	
164.64	−80.06	PKS 1057-79	1.54×10^{-3}	0.569	(Fermi)
165.97	−53.98	PKS 1101-536	2.24×10^{-3}	—	—
83.00	−48.46	PMN J0531-4827	3.27×10^{-3}	—	—
84.70	−44.08	PKS 0537-441	1.60×10^{-2}	0.892	(Fermi)
72.36	−43.84	PKS 0447-439	1.81×10^{-2}	0.343	(TeVCat)
324.86	−42.59	MH 2136-428	5.27×10^{-3}	—	—
53.55	−40.14	PKS 0332-403	2.47×10^{-3}	1.445	(Fermi)
176.76	−38.20	PKS 1144-379	1.76×10^{-3}	1.048	(Fermi)
67.17	−37.94	PKS 0426-380	2.02×10^{-2}	1.110	(Fermi)
53.56	−37.43	PMN J0334-3725	3.04×10^{-3}	—	—
259.39	−33.70	TXS 1714-336	1.90×10^{-3}	—	—
329.71	−30.23	PKS 2155-304	3.64×10^{-2}	0.116	(TeVCat)

(continued on next page)

Table B.1 (continued)

RA (deg)	δ (deg)	**name**	**F** (MeV cm^{-2} s^{-1})	**z**	
230.66	−27.51	PKS 1519-273	1.71×10^{-3}	1.297	(Fermi)
95.59	−26.10	PMN J0622-2605	3.51×10^{-3}	0.414	(Fermi)
45.86	−24.12	PKS 0301-243	8.51×10^{-3}	0.266	(TeVCat)
126.50	−22.50	PKS 0823-223	4.58×10^{-3}	0.910	(Fermi)
97.35	−20.00	PKS 0627-199	1.93×10^{-3}	1.724	(Fermi)
339.14	−14.55	PKS 2233-148	4.90×10^{-3}	0.325	(Fermi)
164.81	−11.57	PKS B1056-113	2.88×10^{-3}	—	—
309.75	−10.78	TXS 2036-109	1.87×10^{-3}	—	—
323.57	−1.89	PKS 2131-021	8.93×10^{-4}	1.283	(Fermi)
184.65	−1.32	PKS 1216-010	1.87×10^{-3}	0.415	(NED)
164.63	1.57	4C +01.28	4.63×10^{-3}	0.890	(Fermi)
127.97	4.48	PKS 0829+046	2.06×10^{-3}	0.174	(Fermi)
267.88	9.65	OT 081	3.39×10^{-3}	0.322	(TeVCat)
32.82	10.86	MG1 J021114+1051	4.95×10^{-3}	0.200	(Fermi)
39.67	16.62	AO 0235+164	8.70×10^{-3}	0.940	(Fermi)
114.54	17.71	PKS 0735+17	6.89×10^{-3}	0.424	(Fermi)
259.81	17.76	PKS 1717+177	2.09×10^{-3}	0.137	(Fermi)
133.71	20.11	OJ 287	4.65×10^{-3}	0.306	(TeVCat)
340.99	20.35	RGB J2243+203	6.31×10^{-3}	—	—
80.44	21.21	TXS 0518+211	1.87×10^{-2}	0.108	(Fermi)
11.34	21.46	GB6 J0045+2127	2.88×10^{-3}	—	—
18.03	22.75	S2 0109+22	7.05×10^{-3}	0.265	(Fermi)
187.56	25.30	ON 246	3.70×10^{-3}	0.135	(Fermi)
26.15	27.09	TXS 0141+268	4.91×10^{-3}	—	—
68.41	29.10	MG2 J043337+2905	3.08×10^{-3}	0.970	(Fermi)
184.48	30.12	1ES 1215+303	1.35×10^{-2}	0.131	(TeVCat)
30.93	30.71	NVSS J020344+304238	1.52×10^{-3}	0.761	(Fermi)
268.55	32.20	RX J1754.1+3212	3.03×10^{-3}	—	—
166.12	38.21	Mkn 421	8.92×10^{-2}	0.030	(Fermi)
253.47	39.76	Mkn 501	3.27×10^{-2}	0.033	(Fermi)
330.70	42.28	BL Lacertae	1.61×10^{-2}	0.069	(TeVCat)
124.56	42.38	S4 0814+42	4.42×10^{-3}	0.530	(Fermi)

(continued on next page)

Table B.1 (continued)

RA (deg)	δ (deg)	**name**	**F** (MeV cm^{-2} s^{-1})	**z**	
35.67	43.03	3C 66A	2.06×10^{-2}	0.340	(TeVCat)
300.30	43.88	MG4 J200112+4352	6.09×10^{-3}	—	—
303.02	46.49	7C 2010+4619	4.02×10^{-3}	—	—
91.85	47.66	TXS 0603+476	3.19×10^{-3}	—	—
262.08	50.23	I Zw 187	3.78×10^{-3}	0.055	(Fermi)
223.63	51.41	TXS 1452+516	3.91×10^{-3}	1.083	(NED)
122.46	52.31	1ES 0806+524	6.70×10^{-3}	0.138	(TeVCat)
193.31	53.02	S4 1250+53	4.09×10^{-3}	0.178	(NED)
276.04	56.85	4C +56.27	1.44×10^{-3}	0.663	(Fermi)
235.75	61.50	GB6 J1542+6129	6.53×10^{-3}	0.507	(Fermi)
123.67	64.50	GB6 J0814+6431	1.55×10^{-3}	0.239	(Fermi)
300.01	65.15	1ES 1959+650	2.20×10^{-2}	0.048	(TeVCat)
149.70	65.56	S4 0954+65	2.91×10^{-3}	0.367	(Fermi)
271.74	69.82	3C 371	2.49×10^{-3}	0.050	(Fermi)
267.16	70.09	S4 1749+70	5.23×10^{-3}	0.770	(Fermi)
110.49	71.34	S5 0716+71	2.05×10^{-2}	0.127	(TeVCat)
305.64	76.20	S5 2023+760	1.32×10^{-3}	0.594	(Fermi)
301.42	77.88	S5 2007+77	2.44×10^{-3}	0.342	(Fermi)
270.17	78.47	S5 1803+784	4.18×10^{-3}	0.680	(Fermi)
58.70	80.18	S5 0346+80	2.89×10^{-4}	—	—

B.2. FSRQs

Table B.2.

Fermi-LAT J2000		associated source	energy flux	redshift (reference)	
RA (deg)	δ (deg)	**name**	**F** (MeV cm^{-2} s^{-1})	**z**	
244.45	−77.31	PKS 1610-77	9.63×10^{-4}	1.710	(Fermi)
326.82	−75.61	PKS 2142-75	2.84×10^{-3}	1.138	(Fermi)
255.92	−62.22	MRC 1659-621	1.66×10^{-3}	1.755	(Fermi)

(continued on next page)

Table B.2 (continued)

RA (deg)	δ (deg)	**name**	**F** (MeV cm^{-2} s^{-1})	**z**	
295.34	−62.18	PKS 1936-623	1.52×10^{-3}	—	—
39.16	−61.62	PKS 0235-618	7.35×10^{-4}	0.467	(Fermi)
47.48	−60.97	PKS 0308-611	1.17×10^{-3}	1.479	(Fermi)
32.69	−51.02	PKS 0208-512	2.69×10^{-3}	1.003	(Fermi)
352.33	−49.93	PKS 2326-502	7.65×10^{-3}	0.518	(Fermi)
81.58	−48.51	PKS 0524-485	1.93×10^{-3}	1.300	(Fermi)
314.07	−47.23	PKS 2052-47	2.66×10^{-3}	1.489	(Fermi)
41.51	−46.85	PKS 0244-470	1.58×10^{-3}	1.385	(Fermi)
216.99	−42.10	PKS B1424-418	2.60×10^{-2}	1.522	(Fermi)
270.67	−39.67	PMN J1802-3940	3.01×10^{-3}	1.319	(Fermi)
299.50	−38.76	PKS 1954-388	1.67×10^{-3}	0.630	(Fermi)
223.59	−37.76	PKS 1451-375	3.01×10^{-4}	0.314	(NED)
37.37	−36.73	PKS 0227-369	3.95×10^{-4}	2.115	(NED)
224.36	−35.66	PKS 1454-354	2.22×10^{-3}	1.424	(Fermi)
326.27	−33.95	PMN J2145-3357	9.87×10^{-4}	1.360	(Fermi)
199.02	−33.64	PKS 1313-333	1.91×10^{-3}	1.210	(Fermi)
327.99	−30.46	PKS 2149-306	4.39×10^{-4}	2.345	(NED)
191.68	−25.80	PKS 1244-255	4.25×10^{-3}	0.635	(Fermi)
55.80	−25.50	PKS 0341-256	2.26×10^{-4}	1.419	(NED)
246.45	−25.46	PKS 1622-253	3.60×10^{-3}	0.786	(Fermi)
74.26	−23.41	PKS 0454-234	1.35×10^{-2}	1.003	(Fermi)
43.20	−22.32	PKS 0250-225	2.57×10^{-3}	1.419	(Fermi)
352.74	−21.77	PMN J2331-2148	5.89×10^{-4}	0.563	(NED)
19.73	−21.68	PKS 0116-219	2.04×10^{-3}	1.165	(Fermi)
290.87	−21.08	TXS 1920-211	1.49×10^{-3}	0.874	(Fermi)
278.41	−21.06	PKS 1830-211	7.94×10^{-3}	2.507	(Fermi)
287.80	−20.12	PKS B1908-201	2.18×10^{-3}	1.119	(Fermi)
171.76	−18.96	PKS 1124-186	5.60×10^{-3}	1.048	(Fermi)
23.17	−16.91	PKS 0130-17	1.25×10^{-3}	1.020	(Fermi)
356.30	−15.92	PMN J2345-1555	8.03×10^{-3}	0.621	(Fermi)
233.20	−13.33	TXS 1530-131	3.47×10^{-3}	—	—
263.25	−13.09	PKS 1730-13	3.95×10^{-3}	0.902	(Fermi)

(continued on next page)

Table B.2 (continued)

RA (deg)	δ (deg)	**name**	**F** (MeV cm^{-2} s^{-1})	**z**	
132.54	−12.21	PMN J0850-1213	1.49×10^{-3}	0.566	(Fermi)
112.57	−11.69	PKS 0727-11	8.27×10^{-3}	1.589	(Fermi)
19.01	−11.61	PKS 0113-118	7.05×10^{-4}	0.670	(NED)
207.37	−11.51	PKS 1346-112	5.90×10^{-4}	0.340	(NED)
228.21	−9.11	PKS 1510-08	2.41×10^{-2}	0.360	(TeVCat)
214.86	−8.64	NVSS J141922-083830	1.66×10^{-3}	—	—
337.41	−8.54	PKS 2227-08	1.47×10^{-3}	1.560	(NED)
122.07	−7.85	PKS 0805-07	4.72×10^{-3}	1.837	(Fermi)
306.42	−7.59	PKS 2023-07	5.62×10^{-3}	1.388	(Fermi)
194.05	−5.79	3C 279	1.74×10^{-2}	0.536	(TeVCat)
203.00	−5.16	PKS 1329-049	1.43×10^{-3}	2.150	(Fermi)
350.89	−3.30	PKS 2320-035	2.79×10^{-3}	1.393	(Fermi)
137.44	−2.52	PKS 0907-023	1.41×10^{-3}	0.957	(Fermi)
75.31	−1.98	S3 0458-02	2.33×10^{-3}	2.291	(Fermi)
54.88	−1.77	PKS 0336-01	3.65×10^{-3}	0.850	(Fermi)
70.67	−0.29	PKS 0440-00	2.04×10^{-3}	0.449	(Fermi)
48.21	1.56	PKS 0310+013	4.23×10^{-4}	0.664	(NED)
17.17	1.58	4C +01.02	6.54×10^{-3}	2.099	(Fermi)
114.82	1.62	PKS 0736+01	2.86×10^{-3}	0.189	(TeVCat)
34.46	1.74	PKS 0215+015	2.12×10^{-3}	1.715	(Fermi)
187.26	2.05	3C 273	2.73×10^{-3}	0.158	(Fermi)
189.89	4.73	MG1 J123931+0443	3.24×10^{-3}	1.761	(Fermi)
76.34	5.00	PKS 0502+049	3.80×10^{-3}	0.954	(Fermi)
83.17	7.55	OG 050	3.37×10^{-3}	1.254	(Fermi)
146.65	10.29	TXS 0943+105	1.14×10^{-3}	1.007	(Fermi)
47.27	10.49	PKS 0306+102	1.00×10^{-3}	0.863	(Fermi)
226.11	10.49	PKS 1502+106	1.27×10^{-2}	1.839	(Fermi)
308.85	10.93	PKS 2032+107	2.97×10^{-3}	0.601	(Fermi)
72.28	11.34	PKS 0446+11	1.30×10^{-3}	1.207	(NED)
338.16	11.73	CTA 102	2.92×10^{-2}	1.037	(Fermi)
238.40	12.96	PKS 1551+130	8.70×10^{-4}	1.290	(NED)
111.33	14.42	4C +14.23	2.91×10^{-3}	1.038	(Fermi)

(continued on next page)

Table B.2 (continued)

RA (deg)	δ (deg)	**name**	**F** (MeV cm^{-2} s^{-1})	**z**	
343.49	16.15	3C 454.3	4.93×10^{-2}	0.859	(Fermi)
330.88	17.43	PKS 2201+171	2.00×10^{-3}	1.076	(Fermi)
325.90	17.73	OX 169	3.26×10^{-3}	0.211	(Fermi)
77.51	18.01	PKS 0507+17	3.42×10^{-3}	0.416	(Fermi)
186.23	21.38	4C +21.35	1.21×10^{-2}	0.432	(TeVCat)
117.18	24.02	S3 0745+24	5.79×10^{-4}	0.410	(NED)
127.69	24.17	OJ 248	7.96×10^{-4}	0.941	(NED)
220.98	25.04	PKS 1441+25	4.58×10^{-3}	0.939	(TeVCat)
339.10	28.49	B2 2234+28A	3.04×10^{-3}	0.790	(Fermi)
39.48	28.81	4C +28.07	5.99×10^{-3}	1.206	(Fermi)
179.88	29.25	Ton 599	5.30×10^{-3}	0.729	(TeVCat)
318.89	29.55	B2 2113+29	1.75×10^{-3}	1.514	(Fermi)
329.39	31.45	B2 2155+31	1.99×10^{-3}	1.486	(Fermi)
230.55	31.74	B2 1520+31	6.89×10^{-3}	1.489	(Fermi)
282.09	32.27	B2 1846+32A	1.24×10^{-3}	0.798	(Fermi)
197.66	32.39	OP 313	2.30×10^{-3}	0.997	(Fermi)
194.48	32.47	ON 393	7.93×10^{-4}	0.806	(NED)
109.84	33.13	B2 0716+33	2.67×10^{-3}	0.779	(Fermi)
347.77	34.42	B2 2308+34	2.53×10^{-3}	1.817	(Fermi)
35.28	35.93	B0218+357	5.89×10^{-3}	0.944	(Fermi)
303.91	37.18	MG2 J201534+3710	5.84×10^{-3}	—	—
248.82	38.14	4C +38.41	6.77×10^{-3}	1.814	(Fermi)
263.60	38.97	B2 1732+38A	1.52×10^{-3}	0.976	(Fermi)
155.79	39.81	4C +40.25	8.59×10^{-4}	1.254	(NED)
250.73	39.82	3C 345	1.69×10^{-3}	0.593	(Fermi)
176.74	39.97	S4 1144+40	2.63×10^{-3}	1.090	(NED)
341.06	40.96	TXS 2241+406	5.73×10^{-3}	1.171	(Fermi)
257.43	43.31	B3 1708+433	1.62×10^{-3}	1.027	(Fermi)
270.38	44.06	S4 1800+44	1.08×10^{-3}	0.663	(NED)
140.23	44.70	S4 0917+44	2.86×10^{-3}	2.186	(Fermi)
206.40	44.88	B3 1343+451	5.61×10^{-3}	2.534	(Fermi)
103.60	45.24	B3 0650+453	1.35×10^{-3}	0.928	(Fermi)

(continued on next page)

Table B.2 (continued)

RA (deg)	δ (deg)	**name**	**F** (MeV cm^{-2} s^{-1})	**z**	
261.86	45.51	S4 1726+45	1.01×10^{-3}	0.717	(NED)
24.26	47.86	OC 457	2.68×10^{-3}	0.859	(Fermi)
254.44	48.14	4C +48.41	1.07×10^{-3}	1.669	(Fermi)
83.33	48.38	TXS 0529+483	1.65×10^{-3}	1.160	(Fermi)
178.38	49.52	OM 484	9.47×10^{-4}	0.334	(Fermi)
18.39	49.82	S4 0110+49	1.50×10^{-3}	0.389	(Fermi)
103.61	50.71	GB6 J0654+5042	9.97×10^{-4}	1.253	(Fermi)
115.66	54.73	GB6 J0742+5444	1.49×10^{-3}	0.723	(Fermi)
15.71	58.42	TXS 0059+581	3.85×10^{-3}	0.644	(Fermi)
158.49	60.86	S4 1030+61	3.22×10^{-3}	1.401	(Fermi)
140.42	62.26	OK 630	1.16×10^{-3}	1.446	(NED)
282.32	67.09	S4 1849+67	3.12×10^{-3}	0.657	(Fermi)
255.02	68.50	TXS 1700+685	1.64×10^{-3}	0.301	(Fermi)
130.29	70.90	S5 0836+71	6.34×10^{-4}	2.172	(NED)
162.10	71.72	S5 1044+71	6.32×10^{-3}	1.150	(Fermi)

B.3. TXS 0506+056

Table B.3.

Fermi-LAT J2000		associated source	energy flux	redshift (reference)	
RA (deg)	δ (deg)	**name**	**F** (MeV cm^{-2} s^{-1})	**z**	
77.36	5.70	TXS 0506+056	7.54×10^{-3}	0.337	(TeVCat)

Acronyms

fADC flash Analogue-Digital-Converter 38

FFR false flare rate 115, 117

FoV Field of View 22, 83

FR-I Fanaroff-Riley type 1 95–97, 100, 102, 103

FR-II Fanaroff-Riley type 2 95–97, 102, 103

FRA fast response analysis 142, 145, 146, 188, 189

FSRQ Flat Spectrum Radio Quasar 102, 142

GBM Gamma Burst Monitor 75, 87

GCN Gamma-Ray Coordinate Network 24, 45, 46

GFU gamma-ray follow-up 45, 46

GRB Gamma Ray Burst 21, 24, 87

GZK Greisen-Zatsepin-Kuzmin 16

HE High Energy 74, 141, 146, 173

HLC hard local coincidence 37, 40, 41, 58

HSP high synchrotron peaked 104

IACT imaging atmospheric Cherenkov telescope 46, 47, 74, 101, 141

ICL IceCube Laboratory 27, 39

IRF instrument response function 83–85

ISP intermediate synchrotron peaked 104

LAT Large Angle Telescope 75

LSP low synchrotron peaked 104, 105

m.w.e. meters water equivalent 30

MC Monte Carlo 115, 116, 119, 120, 128, 130

MJD Modified Julian Date 66

MMC Muon Monte Carlo 47

MoU Memorandum of Understanding 24, 25

MPE multi-photo-electron 61

MWL multi-wavelength 20, 21, 91, 141, 142, 190, 196

NC neutral current 28–30

NLR narrow line region 94, 95

OFU optical follow-up 45, 46

PDF probability density function 60, 61, 64, 113, 119–121, 130, 131, 150, 152, 164, 169, 174, 189, 190

PE photoelectron 37, 139

PMT photomultiplier tube 27, 28, 35, 37, 38, 48, 49, 51, 53, 78

PRNG pseudorandom number generator 165

PSF point spread function 63, 65, 84, 108, 110, 111, 122, 130, 140, 155

r.m.s. root mean square 8, 53, 81, 127, 168

ROI region of interest 82, 110

SED spectral energy distribution 91, 97, 99, 105, 110, 141, 142, 144, 153, 189–191

SLC soft local coincidence 40

SMBH supermassive black hole 91, 102

SMT Simple Multiplicity Trigger 40

SOI source of interest 108, 110, 111

SPE single photo-electron 61, 70

SPS South Pole System 39–41

SRT seeded RT 58

SSC synchrotron-self Compton 98, 99, 101

SSD silicon strip detector 77, 85

TS test statistic 46, 110, 111, 120, 133, 182

UHECR Ultra-High Energy Cosmic Ray 8, 22

UTC Coordinated Universal Time 66, 139

VHE Very High Energy 18, 23, 25, 46, 74, 141, 143, 196

VLA Karl G. Jansky Very Large Array 96

VLBI very-long-baseline interferometry 95, 96

ZTF Zwicky Transient Facility 25

“What is naive and dangerous is to expect science to thrive when information exchange is hindered.”
Alexandra Elbakyan

Bibliography

[1] N. Sammon, director. **Northern Star**. In collab. with C. Moloney and J. Bell Burnell. 2007. URL: `https://www.bbc.co.uk/programmes/b007qf14`.

[2] T. K. Gaisser, R. Engel, and E. Resconi. **Cosmic Rays and Particle Physics**. 2nd ed. 2016. ISBN: 978-1-139-19219-4. DOI: `10.1017/CBO9781139192194`.

[3] R. Engel. “Very High Energy Cosmic Rays and Their Interactions”. In: **Nuclear Physics B Proceedings Supplements** 151 (2006), pp. 437–461. DOI: `10.1016/j.nuclphysbps.2005.07.079`.

[4] F. G. Schröder. “Radio Detection of Cosmic-Ray Air Showers and High-Energy Neutrinos”. In: **Progress in Particle and Nuclear Physics** 93 (2017), pp. 1–68. DOI: `10.1016/j.ppnp.2016.12.002`.

[5] P. Lipari. “Spectral Shapes of the Fluxes of Electrons and Positrons and the Average Residence Time of Cosmic Rays in the Galaxy”. In: **Physical Review D** 99 (2019), p. 043005. DOI: `10.1103/PhysRevD.99.043005`.

[6] V. Verzi. “Measurement of the Energy Spectrum of Ultra-High Energy Cosmic Rays Using the Pierre Auger Observatory”. In: **PoS**. 36th International Cosmic Ray Conference (ICRC2019). Vol. ICRC2019. 2019, p. 450. DOI: `10.22323/1.358.0450`.

[7] D. Ivanov. “Energy Spectrum Measured by the Telescope Array”. In: **PoS**. 36th International Cosmic Ray Conference (ICRC2019). Vol. ICRC2019. 2019, p. 298. DOI: `10.22323/1.358.0298`.

[8] A. De Angelis and M. Pimenta. “Messengers from the High-Energy Universe”. In: **Introduction to Particle and Astroparticle Physics: Multimessenger Astronomy and Its Particle Physics Foundations**. 2018, pp. 575–681. ISBN: 978-3-319-78181-5.

[9] M. Tanabashi et al. “Review of Particle Physics”. In: **Phys. Rev.** D98.3 (2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001.

[10] P. Zyla et al. “Review of Particle Physics”. In: **PTEP** 2020.8 (2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.

[11] Y. S. Yoon et al. “Cosmic-Ray Proton and Helium Spectra from the First CREAM Flight”. In: **Astrophys. J.** 728 (2011), p. 122. arXiv: 1102.2575 [astro-ph.HE].

[12] K.-H. Kampert and M. Unger. “Measurements of the Cosmic Ray Composition with Air Shower Experiments”. In: **Astropart. Phys.** 35 (2012), pp. 660–678. arXiv: 1201.0018 [astro-ph.HE].

[13] C. Grupen. **Astroparticle Physics**. 2005. ISBN: 978-3-540-25312-9.

[14] R. Aloisio, V. Berezinsky, and A. Gazizov. “Transition from Galactic to Extragalactic Cosmic Rays”. In: **Astropart. Phys.** 39-40 (2012), pp. 129–143. arXiv: 1211.0494 [astro-ph.HE].

[15] P. Zuccon and AMS Tracker Collaboration. “The AMS Silicon Tracker: Construction and Performance”. In: **Nuclear Instruments and Methods in Physics Research A** 596 (2008), pp. 74–78. DOI: 10.1016/j.nima.2008.07.116.

[16] K. Lübelsmeyer et al. “Upgrade of the Alpha Magnetic Spectrometer (AMS-02) for Long Term Operation on the International Space Station (ISS)”. In: **Nuclear Instruments and Methods in Physics Research A** 654 (2011), pp. 639–648. DOI: 10.1016/j.nima.2011.06.051.

[17] O. Adriani et al. “The PAMELA Mission: Heralding a New Era in Precision Cosmic Ray Physics”. In: **Physics Reports** 544 (2014), pp. 323–370. DOI: 10.1016/j.physrep.2014.06.003.

[18] E. S. Seo and ISS-CREAM Collaboration. "Cosmic Ray Energetics And Mass: From Balloons to the ISS". In: **PoS**. 34th International Cosmic Ray Conference. Vol. ICRC2015. 2015, p. 574. DOI: 10.22323/1.236.0574.

[19] A. D. Panov et al. "Energy Spectra of Abundant Nuclei of Primary Cosmic Rays from the Data of ATIC-2 Experiment: Final Results". In: **Bulletin of the Russian Academy of Sciences, Physics** 73 (2009), pp. 564–567. DOI: 10.3103/S1062873809050098.

[20] HAWC Collaboration et al. "All-Particle Cosmic Ray Energy Spectrum Measured by the HAWC Experiment from 10 to 500 TeV". In: **Phys. Rev. D** 96.12 (2017), p. 122001. arXiv: 1710.00890.

[21] R. Abbasi et al. "IceTop: The Surface Component of IceCube". In: **Nucl. Instrum. Meth. A** 700 (2013), pp. 188–220. arXiv: 1207.6326 [astro-ph.IM].

[22] Pierre Auger Collaboration. "The Pierre Auger Cosmic Ray Observatory". In: **Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment** A798.FERMILAB-PUB-15-034-AD-AE-CD-TD (2015), pp. 172–213. arXiv: 1502.01323 [astro-ph.IM].

[23] H. Kawai et al. "Telescope Array Experiment". In: **Nuclear Physics B Proceedings Supplements** 175 (2008), pp. 221–226. DOI: 10.1016/j.nuclphysbps.2007.11.002.

[24] R. U. Abbasi et al. "First Observation of the Greisen-Zatsepin-Kuzmin Suppression". In: **Physical Review Letters** 100 (2008), p. 101101. DOI: 10.1103/PhysRevLett.100.101101.

[25] S. Hoover et al. "Observation of Ultrahigh-Energy Cosmic Rays with the ANITA Balloon-Borne Radio Interferometer". In: **Physical Review Letters** 105 (2010), p. 151101. DOI: 10.1103/PhysRevLett.105.151101.

[26] T. Huege and Pierre Auger Collaboration. "Radio Detection of Cosmic Rays with the Auger Engineering Radio Array". In: **EPJ Web of Conferences**. UHECR2018. Vol. 210. 2019, p. 05011. arXiv: 1905.04986.

[27] S. Thoudam et al. "LORA: A Scintillator Array for LOFAR to Measure Extensive Air Showers". In: **Nuclear Instruments and Methods in Physics Research A** 767 (2014), pp. 339–346. DOI: 10.1016/j.nima.2014.08.021.

[28] A. di Matteo et al. "Full-Sky Searches for Anisotropies in UHECR Arrival Directions with the Pierre Auger Observatory and the Telescope Array". In: **PoS**. Vol. ICRC2019. 2019, p. 439. arXiv: 2001.01864 [astro-ph.HE].

[29] M. S. Pshirkov et al. "Deriving the Global Structure of the Galactic Magnetic Field from Faraday Rotation Measures of Extragalactic Sources". In: **The Astrophysical Journal** 738 (2011), p. 192. DOI: 10.1088/0004-637X/738/2/192.

[30] R. Durrer and A. Neronov. "Cosmological Magnetic Fields: Their Generation, Evolution and Observation". In: **Astron Astrophys Rev** 21 (2013), p. 62. arXiv: 1303.7121.

[31] Pierre Auger Collaboration et al. "Large-Scale Cosmic-Ray Anisotropies above 4 EeV Measured by the Pierre Auger Observatory". In: **ApJ** 868.FERMILAB-PUB-18-393 (1 2018), p. 4. arXiv: 1808.03579.

[32] Pierre Auger Collaboration. "Observation of a Large-Scale Anisotropy in the Arrival Directions of Cosmic Rays above 8×10^{18} eV". In: **Science** 357.6357 (2017), pp. 1266–1270. arXiv: 1709.07321.

[33] Pierre Auger Collaboration et al. "The Pierre Auger Observatory: Contributions to the 36th International Cosmic Ray Conference (ICRC 2019)". In: **PoS**. Vol. ICRC2019. 2019. arXiv: 1909.09073.

[34] Pierre Auger Collaboration. "Reconstruction of Inclined Air Showers Detected with the Pierre Auger Observatory". In: **JCAP** 1408.FERMILAB-PUB-14-234-AD-AE-CD-TD (2014), p. 019. arXiv: 1407.3214 [astro-ph.HE].

[35] T. Carver. “Time Integrated Searches for Astrophysical Neutrino Sources Using the IceCube Detector and Gender in Physics Studies for the Genera Project”. PhD thesis. Université de Genève, 2019. URL: https://nbn-resolving.org/urn:nbn:ch:unige-1209245.

[36] M. Aartsen et al. “Astrophysical Neutrinos and Cosmic Rays Observed by IceCube”. In: **Adv. Space Res.** 62 (2018), pp. 2902–2930. arXiv: 1701.03731 [astro-ph.HE].

[37] Education, Communications and Outreach Group. **LHC the Guide**. 2017. URL: https://home.cern/resources/brochure/cern/lhc-guide.

[38] S. Coenders. “High-Energy Cosmic Ray Accelerators: Searches with IceCube Neutrinos”. PhD thesis. Technische Universität München, 2016. URL: http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss-20161028-1327578-1-4.

[39] A. M. Hillas. “The Origin of Ultrahigh-Energy Cosmic Rays”. In: **Ann. Rev. Astron. Astrophys.** 22 (1984), pp. 425–444. DOI: 10.1146/annurev.aa.22.090184.002233.

[40] R. Jansson and G. R. Farrar. “A New Model of the Galactic Magnetic Field”. In: **The Astrophysical Journal** 757 (2012), p. 14. DOI: 10.1088/0004-637X/757/1/14.

[41] M. Kachelriess. **Lecture Notes on High Energy Cosmic Rays**. 2008. arXiv: 0801.4376 [astro-ph].

[42] M. Aguilar et al. “Precision Measurement of the Boron to Carbon Flux Ratio in Cosmic Rays from 1.9 GV to 2.6 TV with the Alpha Magnetic Spectrometer on the International Space Station”. In: **Phys. Rev. Lett.** 117.23 (2016), p. 231102. DOI: 10.1103/PhysRevLett.117.231102.

[43] M. Ahlers. “Cosmic Neutrinos as a Probe of TeV-Scale Physics”. PhD thesis. Universität Hamburg, 2006. URL: https://www-library.desy.de/preparch/desy/thesis/desy-thesis-07-002.pdf.

[44] J. K. Becker. "High-Energy Neutrinos in the Context of Multimessenger Astrophysics". In: **Physics Reports** 458.4-5 (2008), pp. 173–246. DOI: 10.1016/j.physrep.2007.10.006.

[45] C. Tchernin et al. **An Exploration of Hadronic Interactions in Blazars Using IceCube**. 2013. arXiv: 1305.3524 [astro-ph.HE].

[46] S. R. Kelner and F. A. Aharonian. "Energy Spectra of Gamma-Rays, Electrons and Neutrinos Produced at Interactions of Relativistic Protons with Low Energy Radiation". In: **Phys. Rev. D** 78.3 (2008), p. 034013. arXiv: 0803.0688.

[47] S. R. Kelner, F. A. Aharonian, and V. V. Bugayov. "Energy Spectra of Gamma-Rays, Electrons and Neutrinos Produced at Proton-Proton Interactions in the Very High Energy Regime". In: **Phys. Rev. D** 79.3 (2009), p. 039901. arXiv: astro-ph/0606058.

[48] F. Stecker et al. "High-Energy Neutrinos from Active Galactic Nuclei". In: **Phys. Rev. Lett.** 66.NASA-HEAPTH-91-007 (1991), pp. 2697–2700. DOI: 10.1103/PhysRevLett.66.2697.

[49] M. S. Longair. **High Energy Astrophysics**. 3rd ed. 2011. 861 pp. ISBN: 978-0-521-75618-1.

[50] F. Halzen and A. Kheirandish. "High Energy Neutrinos from Recent Blazar Flares". In: **ApJ** 831.1 (2016), p. 12. arXiv: 1605.06119.

[51] Zyxwv99. **CC BY-SA 4.0**. 2014. URL: https://en.wikipedia.org/wiki/File:VHE-UHE-EHE-THE_gamma-ray_attenuation_intergalactic.svg.

[52] Fermi-LAT Collaboration et al. "A Gamma-Ray Determination of the Universe's Star-Formation History". In: **Science** 362.6418 (2018), pp. 1031–1034. arXiv: 1812.01031.

[53] E. Waxman and J. N. Bahcall. “High-Energy Neutrinos from Astrophysical Sources: An Upper Bound”. In: **Phys. Rev. D** 59.IASSNS-AST-98-38 (1999), p. 023002. arXiv: `hep-ph/9807282`.

[54] G. Maggi et al. “Obscured Flat Spectrum Radio Active Galactic Nuclei as Sources of High-Energy Neutrinos”. In: **Physical Review D** 94 (2016), p. 103007. DOI: `10.1103/PhysRevD.94.103007`.

[55] J. Stachurska. “First Double Cascade Tau Neutrino Candidates in IceCube and a New Measurement of the Flavor Composition”. In: **PoS**. Vol. ICRC2019. 2019. arXiv: `1908.05506`.

[56] M. Aartsen et al. “Flavor Ratio of Astrophysical Neutrinos above 35 TeV in IceCube”. In: **Phys. Rev. Lett.** 114.17 (2015), p. 171102. arXiv: `1502.03376 [astro-ph.HE]`.

[57] Y. Takahashi and JEM-EUSO Collaboration. “The JEM-EUSO Mission”. In: **New Journal of Physics** 11 (2009), p. 065009. arXiv: `0910.4187`.

[58] A. Olinto et al. **POEMMA: Probe of Extreme Multi-Messenger Astrophysics**. 2019. arXiv: `1907.06217`.

[59] Fermi-LAT Collaboration. “Fermi Large Area Telescope Fourth Source Catalog”. In: **ApJS** 247.1 (2020), p. 33. arXiv: `1902.10045`.

[60] R. Abbasi et al. **The IceCube High-Energy Starting Event Sample: Description and Flux Characterization with 7.5 Years of Data**. 2020. arXiv: `2011.03545 [astro-ph.HE]`.

[61] M. G. Aartsen et al. “Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere Using Six Years of IceCube Data”. In: **Astrophys. J.** 833.1 (2016), p. 3. arXiv: `1607.08006 [astro-ph.HE]`.

[62] M. G. Aartsen et al. "Time-Integrated Neutrino Source Searches with 10 Years of IceCube Data". In: **Phys. Rev. Lett.** 124.5 (2020), p. 051103. arXiv: 1910.08488.

[63] B. T. Cleveland et al. "Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector". In: **The Astrophysical Journal** 496.1 (1998), pp. 505–526. DOI: 10.1086/305343.

[64] K. Hirata et al. "Observation of a Neutrino Burst from the Supernova SN1987A". In: **Phys. Rev. Lett.** 58.14 (1987), pp. 1490–1493. DOI: 10.1103/PhysRevLett.58.1490.

[65] B. P. Abbott et al. "Multi-Messenger Observations of a Binary Neutron Star Merger". In: **The Astrophysical Journal** 848.2 (2017), p. L12. DOI: 10.3847/2041-8213/aa91c9.

[66] M. Ackermann et al. "The Fermi All-Sky Variability Analysis: A List of Flaring Gamma-Ray Sources and the Search for Transients in Our Galaxy". In: **Astrophys. J.** 771 (2013), p. 57. arXiv: 1304.6082 [astro-ph.HE].

[67] Y. T. Tanaka, S. Buson, and D. Kocevski. "Fermi-LAT Detection of Increased Gamma-Ray Activity of TXS 0506+056, Located inside the IceCube-170922A Error Region." In: **The Astronomer's Telegram** 10791 (2017). URL: http://www.astronomerstelegram.org/?read=10791.

[68] A. Habig. "SNEWS: A Neutrino Early Warning System for Galactic SN II". In: **Cosmic Explosions. Proceedings, 10th Astrophysics Conference, College Park, USA, October 11-13, 1999**. Vol. 522. AIP Conf. Proc. 1. 2000, pp. 169–172. arXiv: astro-ph/9912293.

[69] K. Scholberg. "The SuperNova Early Warning System". In: **Astronomische Nachrichten** 329.3 (2008), pp. 337–339. DOI: 10.1002/asna.200710934.

[70] H. A. Ayala Solares et al. "The Astrophysical Multimessenger Observatory Network (AMON): Performance and Science Program". In: **Astropart. Phys.** 114 (2020), pp. 68–76. arXiv: 1903.08714 [astro-ph.IM].

[71] J. Nordin et al. “Transient Processing and Analysis Using AMPEL: Alert Management, Photometry, and Evaluation of Light Curves”. In: **Astronomy & Astrophysics** 631 (2019), A147. DOI: 10.1051/0004-6361/201935634.

[72] M. Bowen. **The Telescope in the Ice: Inventing a New Astronomy at the South Pole**. 2017. ISBN: 978-1-4668-7898-3. URL: https://telescopeintheice.com/.

[73] M. Aartsen et al. “The IceCube Neutrino Observatory: Instrumentation and Online Systems”. In: **JINST** 12.03 (2017), P03012. arXiv: 1612.05093 [astro-ph.IM].

[74] M. Aartsen et al. “The IceCube Realtime Alert System”. In: **Astropart. Phys.** 92 (2017), pp. 30–41. arXiv: 1612.06028 [astro-ph.HE].

[75] E. Blaufuss et al. “The next Generation of IceCube Realtime Neutrino Alerts”. In: **PoS**. Vol. ICRC2019. 2020, p. 1021. arXiv: 1908.04884 [astro-ph.HE].

[76] IceCube Collaboration, MAGIC Collaboration, and VERITAS Collaboration. “Very High-Energy Gamma-Ray Follow-up Program Using Neutrino Triggers from IceCube”. In: **Journal of Instrumentation** 11.11 (2016), P11009. arXiv: 1610.01814 [hep-ex].

[77] M. Rongen. “Calibration of the IceCube Neutrino Observatory”. PhD thesis. RWTH Aachen, 2019. arXiv: 1911.02016 [astro-ph.IM].

[78] J. A. Formaggio and G. P. Zeller. “From eV to EeV: Neutrino Cross Sections Across Energy Scales”. In: **Rev. Mod. Phys.** 84.3 (2012), pp. 1307–1341. arXiv: 1305.7513.

[79] M. Aartsen et al. “Measurement of the Multi-TeV Neutrino Cross Section with IceCube Using Earth Absorption”. In: **Nature** 551 (2017), pp. 596–600. arXiv: 1711.08119 [hep-ex].

[80] R. Gandhi et al. "Ultrahigh-Energy Neutrino Interactions". In: **Astroparticle Physics** 5.2 (1996), pp. 81–110. arXiv: hep-ph/9512364.

[81] J. M. Conrad, M. H. Shaevitz, and T. Bolton. "Precision Measurements with High Energy Neutrino Beams". In: **Rev. Mod. Phys.** 70.4 (1998), pp. 1341–1392. arXiv: hep-ex/9707015.

[82] J. F. Beacom et al. "Measuring Flavor Ratios of High-Energy Astrophysical Neutrinos". In: **Phys. Rev. D** 68.FERMILAB-PUB-03-180-A, MADPH-03-1336 (2003), p. 093005. arXiv: hep-ph/0307025.

[83] M. Aartsen et al. "Measurement of South Pole Ice Transparency with the IceCube LED Calibration System". In: **Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment** 711 (2013), pp. 73–89. arXiv: 1301.5361.

[84] M. Meier and J. Soedingrekso. "Search for Astrophysical Tau Neutrinos with an Improved Double Pulse Method". In: **PoS**. Vol. ICRC2019. 2019, p. 960. arXiv: 1909.05127 [astro-ph.HE].

[85] L. Wille and D. Xu. "Astrophysical Tau Neutrino Identification with IceCube Waveforms". In: **PoS**. Vol. ICRC2019. 2019, p. 1036. arXiv: 1909.05162 [astro-ph.HE].

[86] D. Freiherr Heereman von Zuydtwyck. "HitSpooling: An Improvement for the Supernova Neutrino Detection System in IceCube". PhD thesis. Université libre de Bruxelles, 2015. URL: https://difusion.ulb.ac.be/vufind/Record/ULB-DIPOT:oai:dipot.ulb.ac.be:2013/209179/Details.

[87] J. P. Dumm. "Searches for Point-like Sources of Neutrinos with the 40-String IceCube Detector". PhD thesis. University of Wisconsin - Madison, 2011. URL: https://inspirehep.net/files/6f9a098d25a2884262c842536387f808.

[88] M. Markov and I. Zheleznykh. "On High Energy Neutrino Physics in Cosmic Rays". In: **Nucl. Phys.** 27 (1961), pp. 385–394. DOI: 10.1016/0029-5582(61)90331-5.

[89] A. G. Rosso et al. **Introduction to Neutrino Astronomy**. 2018. arXiv: 1806.06339 [astro-ph].

[90] T. K. Gaisser. **Atmospheric Neutrinos**. 2019. arXiv: 1910.08851 [astro-ph].

[91] C. Haack, C. Wiebusch, and IceCube Collaboration. "A Measurement of the Diffuse Astrophysical Muon Neutrino Flux Using Eight Years of IceCube Data." In: **PoS**. Vol. ICRC2017. 2017, p. 1005. DOI: 10.22323/1.301.1005.

[92] M. G. Aartsen et al. "Search for a Diffuse Flux of Astrophysical Muon Neutrinos with the IceCube 59-String Configuration". In: **Phys. Rev. D** 89.6 (2014), p. 062007. arXiv: 1311.7048 [astro-ph.HE].

[93] IceCube Collaboration. "The Design and Performance of IceCube Deep-Core". In: **Astroparticle Physics** 35.10 (2012), pp. 615–624. arXiv: 1109.6096.

[94] M. G. Aartsen et al. "First Observation of PeV-Energy Neutrinos with IceCube". In: **Phys. Rev. Lett.** 111 (2013), p. 021103. arXiv: 1304.5356 [astro-ph.HE].

[95] M. G. Aartsen et al. "Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector". In: **Science** 342 (2013), p. 1242856. arXiv: 1311.5238 [astro-ph.HE].

[96] M. G. Aartsen et al. "Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data". In: **Phys. Rev. Lett.** 113.10 (2014), p. 101101. arXiv: 1405.5303.

[97] M. G. Aartsen et al. “First Observation of PeV-Energy Neutrinos with IceCube”. In: **Phys. Rev. Lett.** 111.2 (2013), p. 021103. DOI: `10.1103/PhysRevLett.111.021103`.

[98] K. M. Górski et al. “HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere”. In: **The Astrophysical Journal** 622 (2005), pp. 759–771. DOI: `10.1086/427976`.

[99] T. Kintscher. “Rapid Response to Extraordinary Events: Transient Neutrino Sources with the IceCube Experiment”. PhD thesis. Humboldt-Universität zu Berlin, 2020. DOI: `10.18452/21948`.

[100] D. Heck et al. **CORSIKA: A Monte Carlo Code to Simulate Extensive Air Showers.** FZKA-6019. 1998. DOI: `10.5445/IR/270043064`.

[101] A. Gazizov and M. P. Kowalski. “ANIS: High Energy Neutrino Generator for Neutrino Telescopes”. In: **Comput. Phys. Commun.** 172.DESY-04-101 (2005), pp. 203–213. arXiv: `astro-ph/0406439`.

[102] D. Chirkin and W. Rhode. **Muon Monte Carlo: A High-Precision Tool for Muon Propagation through Matter**. 2004. arXiv: `hep-ph/0407075`.

[103] J.-H. Koehne et al. “PROPOSAL: A Tool for Propagation of Charged Leptons”. In: **Computer Physics Communications** 184.9 (2013), pp. 2070–2090. DOI: `10.1016/j.cpc.2013.04.001`.

[104] M. Dunsch et al. “Recent Improvements for the Lepton Propagator PROPOSAL”. In: **Comput. Phys. Commun.** 242 (2019), pp. 132–144. arXiv: `1809.07740 [hep-ph]`.

[105] D. Chirkin. “Photon Tracking with GPUs in IceCube”. In: **Nucl. Instrum. Meth. A** 725 (2013). Ed. by G. Anton et al., pp. 141–143. DOI: `10.1016/j.nima.2012.11.170`.

[106] C. Kopper et al. **Clsim: An OpenCL-Based Photon-Tracking Simulation Using a (Source-Based) Ray Tracing Algorithm Modeling Scattering**

and Absorption of Light in the Deep Glacial Ice at the South Pole or Mediterranean Sea Water. URL: https://github.com/claudiok/clsim.

[107] N. Whitehorn, J. van Santen, and S. Lafebre. "Penalized Splines for Smooth Representation of High-Dimensional Monte Carlo Datasets". In: **Comput. Phys. Commun.** 184 (2013), pp. 2214–2220. arXiv: 1301.2184 [physics.data-an].

[108] M. J. Larson. "A Search for Tau Neutrino Appearance with IceCube-DeepCore". PhD thesis. Niels Bohr Institute, Københavns Universitet, 2018. URL: https://www.nbi.ku.dk/english/theses/phd-theses/phd_theses_2018/michael_james_larson/.

[109] IceCube Collaboration. "South Pole Glacial Climate Reconstruction from Multi-Borehole Laser Particulate Stratigraphy". In: **J. Glaciol.** 59.218 (2013), pp. 1117–1128. DOI: 10.3189/2013JoG13J068.

[110] J. Fegyveresi et al. **Visual Observations and Stratigraphy of the South Pole Ice Core (SPICEcore): A Preliminary Holocene (10.2 Ka) Accumulation Record and Depth-Age Chronology**. Engineer Research and Development Center (U.S.), 2019. DOI: 10.21079/11681/33378.

[111] P. B. Price. "Kinetics of Conversion of Air Bubbles to Air Hydrate Crystals in Antarctic Ice". In: **Science** 267 (1995), p. 1802. arXiv: astro-ph/9501073.

[112] M. Aartsen et al. **IceCube-Gen2: A Vision for the Future of Neutrino Astronomy in Antarctica**. 2014. arXiv: 1412.5106 [astro-ph.HE].

[113] M. Ackermann et al. "Optical Properties of Deep Glacial Ice at the South Pole". In: **Journal of Geophysical Research: Atmospheres** 111.D13 (2006). DOI: 10.1029/2005JD006687.

[114] S. H. Faria, S. Kipfstuhl, and A. Lambrecht. **The EPICA-DML Deep Ice Core: A Visual Record**. Ed. by S. H. Faria, S. Kipfstuhl, and A. Lambrecht.

Frontiers in Earth Sciences. 2018. ISBN: 978-3-662-55306-0. DOI: 10.1007/978-3-662-55308-4.

[115] D. Chirkin. **Likelihood Description for Comparing Data with Simulation of Limited Statistics**. 2013. arXiv: 1304.0735 [astro-ph.IM].

[116] K. Woschnagg and P. B. Price. "Temperature Dependence of Absorption in Ice at 532 Nm". In: **Appl. Opt.** 40.15 (2001), pp. 2496–2500. DOI: 10.1364/AO.40.002496.

[117] P. B. Price et al. "Temperature Profile for Glacial Ice at the South Pole: Implications for Life in a Nearby Subglacial Lake". In: **Proceedings of the National Academy of Sciences** 99.12 (2002), pp. 7844–7847. DOI: 10.1073/pnas.082238999.

[118] D. Chirkin and M. Rongen. "Light Diffusion in Birefringent Polycrystals and the IceCube Ice Anisotropy". In: **PoS**. Vol. ICRC2019. 2020, p. 854. arXiv: 1908.07608 [astro-ph.HE].

[119] D. Chirkin and IceCube Collaboration. "Evidence of Optical Anisotropy of the South Pole Ice". In: **PoS**. 33rd International Cosmic Ray Conference. Vol. ICRC2013. 2013, p. 0580. URL: https://ui.adsabs.harvard.edu/abs/2013ICRC...33.3338C.

[120] IceCube Collaboration et al. "Calibration and Characterization of the IceCube Photomultiplier Tube". In: **Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment** 618.1-3 (2010), pp. 139–152. arXiv: 1002.2442.

[121] M. Rongen. "Measuring the Optical Properties of IceCube Drill Holes". In: **EPJ Web of Conferences**. Ed. by A. Capone et al. Vol. 116. 2016, p. 06011. DOI: 10.1051/epjconf/201611606011.

[122] M. G. Aartsen et al. “Search for Steady Point-like Sources in the Astrophysical Muon Neutrino Flux with 8 Years of IceCube Data”. In: **Eur. Phys. J. C** 79.3 (2019), p. 234. arXiv: 1811.07979.

[123] D. Williams. “Light Propagation in the South Pole Ice”. In: **AIP Conf. Proc.** 1630.1 (2015). Ed. by C. Finley et al., pp. 146–149. DOI: 10.1063/1.4902793.

[124] M. G. Aartsen et al. “Neutrino Emission from the Direction of the Blazar TXS 0506+056 Prior to the IceCube-170922A Alert”. In: **Science** 361.6398 (2018), pp. 147–151. arXiv: 1807.08794 [astro-ph.HE].

[125] M. G. Aartsen et al. “Multimessenger Observations of a Flaring Blazar Coincident with High-Energy Neutrino IceCube-170922A”. In: **Science** 361.6398 (2018), eaat1378. arXiv: 1807.08816 [astro-ph.HE].

[126] P. Heix et al. “Seasonal Variation of Atmospheric Neutrinos in IceCube”. In: **PoS**. Vol. ICRC2019. 2020, p. 465. arXiv: 1909.02036.

[127] S. Tilav et al. “Seasonal Variation of Atmospheric Muons in IceCube”. In: **PoS**. Vol. ICRC2019. 2020, p. 894. arXiv: 1909.01406 [astro-ph.HE].

[128] M. C. R. Zoll. “Preparations for the next Solar WIMP Analysis with IceCube”. Licentiate thesis. Stockholms universitet, 2014. URL: http://www.diva-portal.org/smash/get/diva2:699510/FULLTEXT01.pdf.

[129] M. Aartsen et al. “Improvement in Fast Particle Track Reconstruction with Robust Statistics”. In: **Nucl. Instrum. Meth. A** 736 (2014), pp. 143–149. arXiv: 1308.5501 [astro-ph.IM].

[130] N. van Eijndhoven, O. Fadiran, and G. Japaridze. “Implementation of a Gauss Convoluted Pandel PDF for Track Reconstruction in Neutrino Telescopes”. In: **Astropart. Phys.** 28 (2007), pp. 456–462. arXiv: 0704.1706 [astro-ph].

[131] K. Schatto. "Stacked Searches for High-Energy Neutrinos from Blazars with IceCube". PhD thesis. Johannes Gutenberg-Universität Mainz, 2014. URL: https://inspirehep.net/files/7925173613df3ce30d1cf9b8c5fa96d8.

[132] IceCube Collaboration et al. "Measurements Using the Inelasticity Distribution of Multi-TeV Neutrino Interactions in IceCube". In: **Phys. Rev. D** 99.3 (2019), p. 032004. arXiv: 1808.07629 [hep-ex].

[133] T. Neunhöffer. "Estimating the Angular Resolution of Tracks in Neutrino Telescopes Based on a Likelihood Analysis". In: **Astropart. Phys.** 25 (2006), pp. 220–225. arXiv: astro-ph/0403367.

[134] R. Barlow. **Statistics. A Guide to the Use of Statistical Methods in the Physical Sciences**. 1989. ISBN: 0-471-92295-1.

[135] K. Wiebe. "All-Flavor Based Searches for Solar Dark Matter with the IceCube Neutrino Observatory". PhD thesis. Johannes Gutenberg-Universität Mainz, 2017. URL: https://publications.ub.uni-mainz.de/theses/volltexte/2017/100000952/pdf/100000952.pdf.

[136] E. Pinat. "The IceCube Neutrino Observatory: Search for Extended Sources of Neutrinos and Preliminary Study of a Communication Protocol for Its Future Upgrade". PhD thesis. Université Libre de Bruxelles, 2017. URL: https://difusion.ulb.ac.be/vufind/Record/ULB-DIPOT:oai:dipot.ulb.ac.be:2013/253046/Details.

[137] M. Aartsen et al. "Energy Reconstruction Methods in the IceCube Neutrino Telescope". In: **JINST** 9 (2014), P03009. arXiv: 1311.4767 [physics.ins-det].

[138] H. Karttunen et al. **Fundamental Astronomy**. 2017. ISBN: 978-3-662-53045-0.

[139] T. Jenness and D. S. Berry. "PAL: A Positional Astronomy Library". In: **Astronomical Data Analysis Software and Systems XXII**. Ed. by D. N. Friedel. Vol. 475. Astronomical Society of the Pacific Conference Series.

2013, p. 307. URL: https://ui.adsabs.harvard.edu/abs/2013ASPC..475..307J.

[140] IceCube Collaboration et al. "Observation of the Cosmic-Ray Shadow of the Moon with IceCube". In: **Phys. Rev. D** 89.10 (2014), p. 102004. arXiv: 1305.6811 [astro-ph.HE].

[141] M. G. Aartsen et al. "Detection of the Temporal Variation of the Sun's Cosmic Ray Shadow with the IceCube Detector". In: **Astrophys. J.** 872.2 (2019), p. 133. arXiv: 1811.02015 [astro-ph.HE].

[142] M. G. Aartsen et al. **Measurements of the Time-Dependent Cosmic-Ray Sun Shadow with Seven Years of IceCube Data – Comparison with the Solar Cycle and Magnetic Field Models**. 2020. arXiv: 2006.16298 [astro-ph.HE].

[143] F. Tenholt. "Studying the Temporal Variation of the Cosmic-Ray Sun Shadow". PhD thesis. Ruhr-Universität Bochum, 2020. URL: https://hss-opus.ub.ruhr-uni-bochum.de/opus4/frontdoor/index/index/docId/7021.

[144] P. Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: **Nature Methods** 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[145] W. L. Kraushaar and G. W. Clark. "Search for Primary Cosmic Gamma Rays with the Satellite Explorer XI". In: **Physical Review Letters** 8 (1962), pp. 106–109. DOI: 10.1103/PhysRevLett.8.106.

[146] S. A. Cummer et al. "The Lightning-TGF Relationship on Microsecond Timescales". In: **Geophysical Research Letters** 38.14 (2011). DOI: 10.1029/2011GL048099.

[147] M. Pesce-Rollins et al. "Fermi Large Area Telescope Observations of the Sun: The First Ten Years". In: **PoS** IFS2017 (2017), p. 173. DOI: 10.22323/1.312.0173.

[148] W. Kraushaar et al. "Explorer XI Experiment on Cosmic Gamma Rays". In: **The Astrophysical Journal** 141 (1965), p. 845. DOI: 10.1086/148179.

[149] A. De Angelis and M. J. M. Pimenta. "Particle Detection". In: **Introduction to Particle and Astroparticle Physics: Questions to the Universe**. 2015, pp. 101–187. ISBN: 978-88-470-2688-9.

[150] Fermi-LAT Collaboration et al. "The Spectrum of Isotropic Diffuse Gamma-Ray Emission between 100 MeV and 820 GeV". In: **ApJ** 799.1 (2015), p. 86. arXiv: 1410.3696.

[151] W. B. Atwood et al. "The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission". In: **ApJ** 697.2 (2009), pp. 1071–1102. DOI: 10.1088/0004-637X/697/2/1071.

[152] Fermi-LAT Collaboration. "The Fermi Large Area Telescope On Orbit: Event Classification, Instrument Response Functions, and Calibration". In: **ApJS** 203.1 (2012), p. 4. arXiv: 1206.1896.

[153] W. Atwood et al. "Pass 8: Toward the Full Realization of the Fermi-LAT Scientific Potential". In: 2013. arXiv: 1303.3514 [astro-ph].

[154] M. Tinivella. **A Review of Cosmic-Ray Electrons and Fermi-LAT**. 2016. arXiv: 1610.03672 [astro-ph].

[155] M. Pesce-Rollins. "The Cosmic Ray Electron Spectrum Measured by the Fermi LAT". PhD thesis. Università Degli Studi di Siena, 2017. URL: http://www.infn.it/thesis/PDF/getfile.php?filename=4734-Pesce-rollins-dottorato.pdf.

[156] F. Acero et al. "Fermi Large Area Telescope Third Source Catalog". In: **Astrophys. J. Suppl.** 218.2 (2015), p. 23. arXiv: 1501.02003 [astro-ph.HE].

[157] A. Christov. "Multi-Messenger Studies with the IceCube Detector". PhD thesis. Université de Genève, 2016. URL: https://nbn-resolving.org/urn:nbn:ch:unige-917668.

[158] L. Tibaldo. **Fermi LAT Measurements of Diffuse Gamma-Ray Emission: Results at the First-Year Milestone**. Version 1. 2010. arXiv: 1002.1576 [astro-ph].

[159] Fermi-LAT Collaboration. **Galactic Interstellar Emission Model for the 4FGL Catalog Analysis**. 2019. URL: https://fermi.gsfc.nasa.gov/ssc/data/analysis/software/aux/4fgl/Galactic_Diffuse_Emission_Model_for_the_4FGL_Catalog_Analysis.pdf.

[160] Fermi-LAT Collaboration. "The Spectrum and Morphology of the Fermi Bubbles". In: **ApJ** 793.1 (2014), p. 64. arXiv: 1407.7905.

[161] M. Ajello et al. "A Decade of Gamma-Ray Bursts Observed by Fermi-LAT: The Second GRB Catalog". In: **Astrophys. J.** 878.1 (2019), p. 52. arXiv: 1906.11403 [astro-ph.HE].

[162] M. Ackermann et al. "Fermi Establishes Classical Novae as a Distinct Class of Gamma-Ray Sources". In: **Science** 345 (2014), pp. 554–558. pmid: 25082700 (astro-ph.HE).

[163] M. Ackermann et al. "High-Energy Gamma-Ray Emission from Solar Flares: Summary of Fermi Large Area Telescope Detections and Analysis of Two M-Class Flares". In: **Astrophys. J.** 787 (2014), p. 15. arXiv: 1304.3749 [astro-ph.HE].

[164] Fermi-LAT Collaboration. **The Fourth Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope**. 2020. arXiv: 1905.10771 [astro-ph].

[165] R. G. Ruiz. "Search for Populations of Unresolved Sources of High Energy Neutrinos with the ANTARES Neutrino Telescope". PhD thesis. Université Paris Diderot, 2016. 205 pp. URL: http://www.theses.fr/2016USPCC150/document.

[166] P. Schneider. **Extragalactic Astronomy and Cosmology**. 2nd ed. 2015. ISBN: 978-3-642-54083-7. DOI: 10.1007/978-3-642-54083-7.

[167] G. Ghisellini. **Radiative Processes in High Energy Astrophysics**. 2013. arXiv: 1202.5949 [astro-ph].

[168] C. K. Seyfert. "Nuclear Emission in Spiral Nebulae." In: **The Astrophysical Journal** 97 (1943), p. 28. DOI: 10.1086/144488.

[169] V. Beckmann and C. Shrader. **Active Galactic Nuclei**. 2012. ISBN: 978-3-527-66682-9.

[170] G. Ghisellini, L. Maraschi, and F. Tavecchio. "The Fermi Blazars' Divide". In: **Monthly Notices of the Royal Astronomical Society** 396 (2009), pp. L105–L109. DOI: 10.1111/j.1745-3933.2009.00673.x.

[171] H. Netzer. **The Physics and Evolution of Active Galactic Nuclei**. 2013. ISBN: 978-1-107-02151-8.

[172] P. N. Best et al. "The Host Galaxies of Radio-Loud Active Galactic Nuclei: Mass Dependences, Gas Cooling and Active Galactic Nuclei Feedback". In: **Monthly Notices of the Royal Astronomical Society** 362 (2005), pp. 25–40. DOI: 10.1111/j.1365-2966.2005.09192.x.

[173] C. M. Urry and P. Padovani. "Unified Schemes for Radio-Loud Active Galactic Nuclei". In: **PASP** 107 (1995), p. 803. arXiv: astro-ph/9506063.

[174] C. M. Urry. "AGN Unification: An Update". In: **ASP Conf. Ser.** 311 (2004). Ed. by G. T. Richards and P. B. Hall, p. 49. arXiv: astro-ph/0312545.

[175] A. A. Abdo et al. "The Spectral Energy Distribution of Fermi Bright Blazars". In: **The Astrophysical Journal** 716 (2010), pp. 30–70. DOI: 10.1088/0004-637X/716/1/30.

[176] A. Albert et al. "A Survey of Active Galaxies at TeV Photon Energies with the HAWC Gamma-Ray Observatory". In: **Astrophys. J.** 907.2 (2021), p. 67. arXiv: 2009.09039 [astro-ph.HE].

[177] W. Benbow and the VERITAS Collaboration. **Highlights from the VERITAS AGN Observation Program**. 2017. arXiv: 1708.02374 [astro-ph].

[178] S. Soldi et al. “The Multiwavelength Variability of 3C 273”. In: **A&A** 486.2 (2008), pp. 411–425. DOI: 10.1051/0004-6361:200809947.

[179] E. Resconi et al. “On the Classification of Flaring States of Blazar”. In: **A&A** 502.2 (2009), pp. 499–504. arXiv: 0904.1371.

[180] A. Mücke and R. J. Protheroe. “A Proton Synchrotron Blazar Model for Flaring in Markarian 501”. In: **Astroparticle Physics** 15.1 (2001), pp. 121–136. DOI: 10.1016/S0927-6505(00)00141-9.

[181] A. Bykov et al. “Particle Acceleration in Relativistic Outflows”. In: **Space Science Reviews** 173 (2012), pp. 309–339. DOI: 10.1007/s11214-012-9896-y.

[182] M. C. Begelman and D. F. Cioffi. “Overpressured Cocoons in Extragalactic Radio Sources”. In: **The Astrophysical Journal Letters** 345 (1989), pp. L21–L24. DOI: 10.1086/185542.

[183] E. G. Berezhko. “Cosmic Rays from Active Galactic Nuclei”. In: **ApJ** 684.2 (2008), pp. L69–L71. arXiv: 0809.0734.

[184] M. Petropoulou et al. **The Many Faces of Blazar Emission in the Context of Hadronic Models**. 2016. arXiv: 1601.06010 [astro-ph].

[185] H. Krawczynski et al. “Simultaneous X-Ray and TeV Gamma-Ray Observation of the TeV Blazar Markarian 421 during 2000 February and May”. In: **The Astrophysical Journal** 559 (2001), pp. 187–195. DOI: 10.1086/322364.

[186] S. P. Wakely and D. Horan. “TeVCat: An Online Catalog for Very High Energy Gamma-Ray Astronomy”. In: International Cosmic Ray Conference. Vol. 3. 2008, pp. 1341–1344. URL: http://tevcat2.uchicago.edu/.

[187] A. De Angelis and M. Mallamaci. “Gamma-Ray Astrophysics”. In: **European Physical Journal Plus** 133 (2018), p. 324. DOI: 10.1140/epjp/i2018-12181-0.

[188] Fermi-LAT Collaboration. "Minute-Timescale >100 MeV Gamma-Ray Variability during the Giant Outburst of Quasar 3C 279 Observed by Fermi-LAT in 2015 June". In: **ApJ** 824.2 (2016), p. L20. arXiv: `1605.05324 [astro-ph.HE]`.

[189] F. Aharonian. "An Exceptional VHE Gamma-Ray Flare of PKS 2155-304". In: **ApJ** 664.2 (2007), pp. L71–L74. arXiv: `0706.0797`.

[190] G. Ghisellini et al. "The Fermi Blazar Sequence". In: **Monthly Notices of the Royal Astronomical Society** 469 (2017), pp. 255–266. arXiv: `1702.02571 [astro-ph.HE]`.

[191] B. Lott, D. Gasparrini, and S. Ciprini. **The Fourth Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope – Data Release 2**. 2020. arXiv: `2010.08406 [astro-ph]`.

[192] P. Padovani et al. "TXS 0506+056, the First Cosmic Neutrino Source, Is Not a BL Lac". In: **Monthly Notices of the Royal Astronomical Society: Letters** 484.1 (2019), pp. L104–L108. arXiv: `1901.06998`.

[193] E. T. Meyer et al. "From the Blazar Sequence to the Blazar Envelope: Revisiting the Relativistic Jet Dichotomy in Radio-Loud Active Galactic Nuclei". In: **The Astrophysical Journal** 740 (2011), p. 98. DOI: `10.1088/0004-637X/740/2/98`.

[194] G. Ghisellini et al. "The Transition between BL Lac Objects and Flat Spectrum Radio Quasars". In: **Monthly Notices of the Royal Astronomical Society** 414.3 (2011), pp. 2674–2689. arXiv: `1012.0308 [astro-ph.CO]`.

[195] A. Palladino et al. "Interpretation of the Diffuse Astrophysical Neutrino Flux in Terms of the Blazar Sequence". In: **ApJ** 871.1 (2019), p. 41. arXiv: `1806.04769`.

[196] D. W. Hogg et al. **The K Correction**. 2002. arXiv: `astro-ph/0210394`.

[197] P. Giommi and P. Padovani. "A Simplified View of Blazars: Contribution to the X-Ray and Gamma-Ray Cosmic Backgrounds". In: **Monthly Notices of the Royal Astronomical Society** 450.3 (2015), pp. 2404–2409. arXiv: 1504.01978.

[198] C. Righi, F. Tavecchio, and D. Guetta. "High-Energy Emitting BL Lacs and High-Energy Neutrinos - Prospects for the Direct Association with IceCube and KM3NeT". In: **A&A** 598 (2017), A36. arXiv: 1607.08061.

[199] S. Ansoldi et al. "The Blazar TXS 0506+056 Associated with a High-Energy Neutrino: Insights into Extragalactic Jets and Cosmic Ray Acceleration". In: **Astrophys. J. Lett.** (2018). arXiv: 1807.04300 [astro-ph.HE].

[200] S. Gao et al. "Modelling the Coincident Observation of a High-Energy Neutrino and a Bright Blazar Flare". In: **Nat Astron** 3.1 (2019), pp. 88–92. arXiv: 1807.04275.

[201] Written by The Rubberbandits. Published by Lovely Men, Ltd. 2011.

[202] M. F. Baker. "Time-Dependent Searches for Neutrino Point Sources with the IceCube Observatory". PhD thesis. University of Wisconsin - Madison, 2011. URL: https://inspirehep.net/files/51884b083f79ce79aa31c63eff1ae291.

[203] R. Abbasi et al. "Time-Dependent Searches for Point Sources of Neutrinos with the 40-String and 22-String Configurations of IceCube". In: **Astrophys. J.** 744.1 (2012), p. 1. arXiv: 1104.0075.

[204] M. G. Aartsen et al. "Searches for Time Dependent Neutrino Sources with IceCube Data from 2008 to 2012". In: **Astrophys. J.** 807.1 (2015), p. 46. arXiv: 1503.00598 [astro-ph.HE].

[205] R. Abbasi et al. **A Search for Time-Dependent Astrophysical Neutrino Emission with IceCube Data from 2012 to 2017**. 2020. arXiv: 2012.01079 [astro-ph.HE].

[206] P. Padovani et al. "Dissecting the Region around IceCube-170922A: The Blazar TXS 0506+056 as the First Cosmic Neutrino Source". In: **Monthly Notices of the Royal Astronomical Society** 480.1 (2018), pp. 192–203. DOI: `10.1093/mnras/sty1852`.

[207] **ROOT Reference Guide: Minuit2**. URL: `https://root.cern.ch/doc/master/Minuit2Page.html`.

[208] F. James and M. Roos. "Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations". In: **Comput. Phys. Commun.** 10.CERN-DD-75-20 (1975), pp. 343–367. DOI: `10.1016/0010-4655(75)90039-9`.

[209] B. Lott et al. "An Adaptive-Binning Method for Generating Constant-Uncertainty/Constant-Significance Light Curves with Fermi-LAT Data". In: **Astron. Astrophys.** 544 (2012), A6. arXiv: `1201.4851 [astro-ph.HE]`.

[210] J. D. Scargle et al. "Studies in Astronomical Time Series Analysis. VI. Bayesian Block Representations". In: **Astrophys. J.** 764.2 (2013), pp. 26–. arXiv: `1207.5578`.

[211] Astropy Collaboration et al. **Astropy: A Community Python Package for Astronomy**. 2013. arXiv: `1307.6212 [astro-ph.IM]`.

[212] Astropy Collaboration et al. "The Astropy Project: Building an Inclusive, Open-Science Project and Status of the v2.0 Core Package". In: **AJ** 156.3, 123 (2018), p. 123. arXiv: `1801.02634 [astro-ph.IM]`.

[213] J. Braun et al. "Methods for Point Source Analysis in High Energy Neutrino Telescopes". In: **Astropart. Phys.** 29 (2008), pp. 299–305. arXiv: `0801.1604 [astro-ph]`.

[214] M. G. Aartsen et al. "All-Sky Search for Time-Integrated Neutrino Emission from Astrophysical Sources with 7 Years of IceCube Data". In: **Astrophys. J.** 835.2 (2017), p. 151. DOI: `10.3847/1538-4357/835/2/151`.

[215] H. B. Prosper. "Practical Statistics for Particle Physicists". In: **CERN Yellow Rep. School Proc.** 6 (2019). Ed. by M. Mulders and C. Duhr, pp. 261–292. DOI: 10.23730/CYRSP-2019-006.261.

[216] Written by They Might Be Giants. Published by TMBG Music. 1990.

[217] IceCube Collaboration. **GRB Coordinates Network/AMON Notice 50579430 130033**. 2017. URL: https://gcn.gsfc.nasa.gov/notices_amon/50579430_130033.amon.

[218] IceCube Collaboration. **GCN Circular 21916**. 2017. URL: https://gcn.gsfc.nasa.gov/gcn/gcn3/21916.gcn3.

[219] S. Paiano et al. "The Redshift of the BL Lac Object TXS 0506+056". In: **The Astrophysical Journal** 854.2 (2018), p. L32. DOI: 10.3847/2041-8213/aaad5e.

[220] B. P. Abbott et al. "GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs". In: **Physical Review X** 9.3 (2019), p. 031040. arXiv: 1811.12907 [astro-ph.HE].

[221] J. N. Douglas et al. "The Texas Survey of Radio Sources Covering -35.5°<δ<71.5° at 365 MHz". In: **Astronomical Journal** 111 (1996), p. 1945. DOI: 10.1086/117932.

[222] A. A. Abdo et al. "Fermi Large Area Telescope First Source Catalog". In: **Astrophys. J. Suppl.** 188 (2010), pp. 405–436. arXiv: 1002.2280 [astro-ph.HE].

[223] K. Meagher. "IceCube as a Neutrino Follow-up Observatory for Astronomical Transients". In: **PoS**. 35th International Cosmic Ray Conference. Vol. ICRC2017. 2017, p. 1007. DOI: 10.22323/1.301.1007.

[224] A. Coleiro and D. Dornic. "Search for Neutrinos from TXS 0506+056 with the ANTARES Neutrino Telescope". In: **The Astronomer's Telegram**

12274 (2018). URL: http://www.astronomerstelegram.org/?read=12274.

[225] N. L. Strotjohann, M. Kowalski, and A. Franckowiak. “Eddington Bias for Cosmic Neutrino Sources”. In: **Astronomy & Astrophysics** 622 (2019), p. L9. DOI: 10.1051/0004-6361/201834750.

[226] P. van Dokkum et al. “A Second Galaxy Missing Dark Matter in the NGC1052 Group”. In: **ApJ** 874.1 (2019), p. L5. arXiv: 1901.05973.

[227] C. Kopper and IceCube Collaboration. “Observation of Astrophysical Neutrinos in Six Years of IceCube Data”. In: **PoS**. 35th International Cosmic Ray Conference. Vol. ICRC2017. 2017, p. 981. DOI: 10.22323/1.301.0981.

[228] M. Aartsen et al. “The Contribution of Fermi-2LAC Blazars to the Diffuse TeV-PeV Neutrino Flux”. In: **Astrophys. J.** 835.1 (2017), p. 45. arXiv: 1611.03874 [astro-ph.HE].

[229] M. Huber and K. Krings. “Results of IceCube Searches for Neutrinos from Blazars Using Seven Years of Through-Going Muon Data”. In: **PoS** ICRC2017 (2018), p. 994. DOI: 10.22323/1.301.0994.

[230] M. Huber. “Searches for Steady Neutrino Emission from 3FHL Blazars Using Eight Years of IceCube Data from the Northern Hemisphere”. In: **PoS**. Vol. ICRC2019. 2020, p. 916. arXiv: 1908.08458 [astro-ph.HE].

[231] P. Mertsch, M. Rameez, and I. Tamborra. “Detection Prospects for High Energy Neutrino Sources from the Anisotropic Matter Distribution in the Local Universe”. In: **JCAP** 03 (2017), p. 011. arXiv: 1612.07311 [astro-ph.HE].

[232] IceCube Collaboration et al. **A Search for Neutrino Point-Source Populations in 7 Years of IceCube Data with Neutrino-Count Statistics**. 2019. arXiv: 1909.08623 [astro-ph, physics:hep-ex].

[233] C. M. Urry. “Vera Cooper Rubin (1928—2016)”. In: **Science** 355 (2017), pp. 462–462. DOI: 10.1126/science.aam7920.

[234] M. Paterno. **Calculating Efficiencies and Their Uncertainties**. FERMILAB-TM-2286-CD, 15017262. 2004. DOI: 10.2172/15017262.

[235] M. G. Aartsen et al. “A Combined Maximum-Likelihood Analysis of the High-Energy Astrophysical Neutrino Flux Measured with IceCube”. In: **The Astrophysical Journal** 809 (2015), p. 98. DOI: 10.1088/0004-637X/809/1/98.

[236] J. Stettner. **Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data**. 2019. arXiv: 1908.09551 [astro-ph].

[237] V. A. Acciari et al. “TeV and Multi-Wavelength Observations of Mrk 421 in 2006-2008”. In: **The Astrophysical Journal** 738 (2011), p. 25. DOI: 10.1088/0004-637X/738/1/25.

[238] J. Aleksić et al. “The 2009 Multiwavelength Campaign on Mrk 421: Variability and Correlation Studies”. In: **Astronomy and Astrophysics** 576 (2015), A126. DOI: 10.1051/0004-6361/201424216.

[239] M. Petropoulou, S. Coenders, and S. Dimitrakoudis. “Time-Dependent Neutrino Emission from Mrk 421 during Flares and Predictions for IceCube”. In: **Astroparticle Physics** 80 (2016), pp. 115–130. arXiv: 1603.06954.

[240] A. U. Abeysekara et al. “Daily Monitoring of TeV Gamma-Ray Emission from Mrk 421, Mrk 501, and the Crab Nebula with HAWC”. In: **ApJ** 841.2 (2017), p. 100. arXiv: 1703.06968.

[241] S. Bianchi, R. Maiolino, and G. Risaliti. “AGN Obscuration and the Unified Model”. In: **Advances in Astronomy** 2012 (2012), pp. 1–17. DOI: 10.1155/2012/782030.

[242] A. Mastichiadis and M. Petropoulou. “Hadronic X-Ray Flares from Blazars”. In: **ApJ** 906.2 (2021), p. 131. arXiv: `2009.12158`.

[243] R. Xue et al. “A Two-Zone Blazar Radiation Model for "Orphan" Neutrino Flares”. In: **ApJ** 906.1 (2021), p. 51. arXiv: `2011.03681`.

[244] A. Reimer, M. Böttcher, and S. Buson. “Cascading Constraints from Neutrino-Emitting Blazars: The Case of TXS 0506+056”. In: **The Astrophysical Journal** 881 (2019), p. 46. DOI: `10.3847/1538-4357/ab2bff`.

[245] A. Keivani et al. “A Multimessenger Picture of the Flaring Blazar TXS 0506+056: Implications for High-Energy Neutrino Emission and Cosmic-Ray Acceleration”. In: **The Astrophysical Journal** 864 (2018), p. 84. DOI: `10.3847/1538-4357/aad59a`.

[246] K. Murase, F. Oikonomou, and M. Petropoulou. “Blazar Flares as an Origin of High-Energy Cosmic Neutrinos?” In: **ApJ** 865.2 (2018), p. 124. arXiv: `1807.04748`.

[247] K. Oh et al. “The 105-Month Swift-BAT All-Sky Hard X-Ray Survey”. In: **The Astrophysical Journal Supplement Series** 235 (2018), p. 4. DOI: `10.3847/1538-4365/aaa7fd`.

[248] M. G. Aartsen et al. **IceCube-Gen2: The Window to the Extreme Universe**. 2020. arXiv: `2008.04323 [astro-ph]`.

Acknowledgements

Even a concise attempt to acknowledge the support I've received during my PhD and while writing this thesis could never be brief. Wherever names are listed, they appear in alphabetical order.

Naturally the first to thank is my advisor, Juan Antonio Aguilar Sánchez, for his dedicated and compassionate mentorship. In this quality, as well as your versatile research skills and attention to quality science communication, you've set an example which I have tried (and will continue trying) to emulate. Importantly, whenever I got lost in details and detours, you saw a way forward. But you also had the foresight to let me embark on the TXS 0506+056 detour, which eventually improved the whole thesis. Eventually writing the text would not have been possible without your many helpful comments which he patiently provided during two years of uneven progress (see fig. 1) until it was ready for my jury. Markus Ahlers, Krijn De Vries, Laurent Favart, and Simona Toscano, I am deeply grateful that you agreed to serve this important role, and agreed to a very optimistic timeline. Your comments have significantly improved this thesis, and I'm deeply Diego Beghin, Pablo Correa and Nadège Iovine helped me with their respective mother tongues in the abstracts; Paul Coppin, Simon De Kockere and Giovanni Renzi took on the tedious task to proofread the bibliography. I really appreciate their help, and hope to return it one day. For ensuring the writing progress, I'd also like to recognize the help of my comité d'accompagnement, Ioana Mariș with our bet, and the peer writing group. This thesis would never have seen the light of day without financial support during the writing phase from Fondation Jaumotte-Demoulin and the Fonds David et Alice Van Buuren, ONEM, and Angelika Raab, for which I'm deeply grateful. And finally, everyone who told me to "start writing now": I'm sorry, you were right.

I sincerely thank Markus Ahlers, Krijn De Vries, Xavier Rodrigues and the IIHE LOFAR group for crucial discussions in the later stage of my analysis and for making me feel like part of a scientific community. Many aspects of the analyses were discussed, challenged and refined in meetings with the IIHE IceCube group, who never tired to decipher my plots. This is even more true for the meetings with only Juanan's group. I believe such discussions are the best way to learn physics. At the IIHE, I also enjoyed excellent support of the administration and computing teams.

I further want to acknowledge that this work builds on that of my predecessors in IceCube, especially Asen Christov who did the previous iteration of lightcurve analyses, and my advisor who came up with the basic idea of time-dependent stacking together with Mohamed Rameez. The analysis could not have been conducted without the framework and data set provided by Stephanie Bron, Tessa Carver, Stefan Coenders, Thomas Kintscher, Elisa Pinat, and Josh Wood; thank you. Of course this goes along with sincere thanks to Sara Buson, Anna Franckowiak, and Yasuyuki Tanaka who provided the likelihood lightcurves. In particular, I'm deeply indebted to Anna Franckowiak for her incredible patience answering all my Fermi questions, her valuable advice, and all the discussions we had, both during the turbulent TXS 0506+056 months as well as in 2018. Of course my analyses only got unblinded thanks to the constructive criticism from my reviewers Timo Karg, Thorben Menne, Mohamed Rameez, and Steve Sclafani, the Neutrino/Point Sources working group, and IceCube in general.

My PhD was all the more enjoyable for feeling connected and welcomed in the global IceCube collaboration. I would particularly like to highlight two groups: Uppsala who taught me my first steps, and Karlsruhe who let me gather hardware experience. From that visit, I also appreciated Michael Korntheuer's insights into electronics. Not to be forgotten are my colleagues. You've encouraged and guided me countless times, both regarding life and academia, in many cases even after defending your thesis, changing jobs, leaving Brussels, or retiring. I am especially grateful to those who showed me in 2019 how it was possible to finish my PhD. I regret that I can not list every name. Thank you for making me feel welcome at the IIHE, I leave with many good memories and hope to see you all again.

These acknowledgements must absolutely credit the enduring support from my family and friends. Mama, Lissi, besonderen Dank für eure Unterstützung in den letzten Schreibmonaten, und auch gerade diese Tage mit dem Auszug. All the Brussels gals and lads, Solvay girls, Café Latte, meine GenossX, F21B1B, USG and everyone I had the fortune of meeting in the past years. Seeing you root for me was a big boost, and knowing you'll always stand by me helped through the rough times. I'd like to close with a look back at those who helped me on the path to the PhD: Carlos Pérez de los Heros, who supervised me in Uppsala, Elisa Resconi who gave me the impulse to go there, Florian Haas who

gave me the first taste of research, and Wolfgang Heine who won me back for mathematics. Und natürlich meine Eltern, die mich von früh gefördert und immer für mich gekämpft haben.

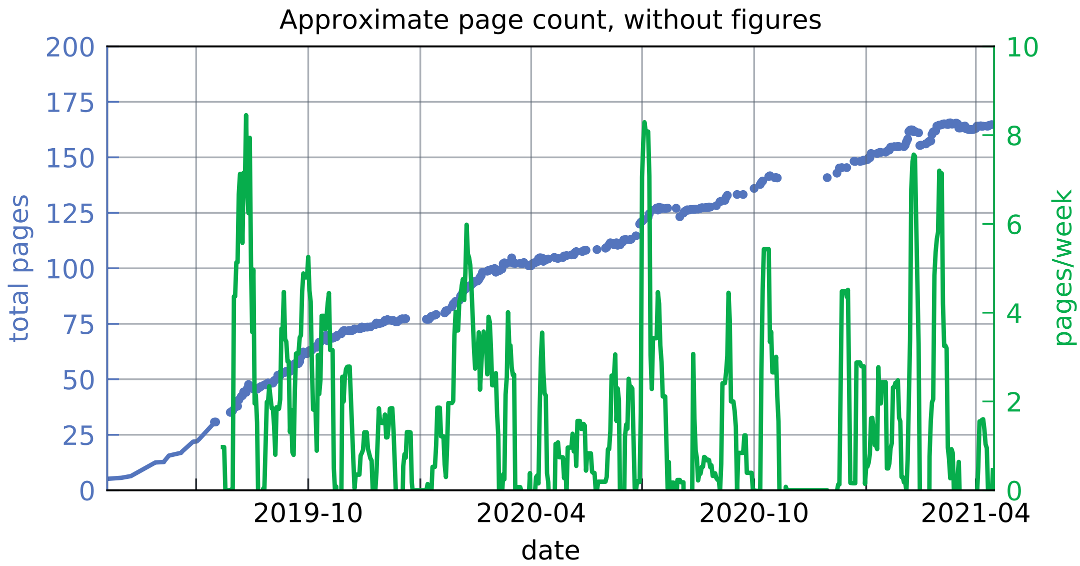

Figure 1.: Writing progress by date (year-month) until submission to the jury, represented by the approximate page count (left y-axis, blue, assuming 285 words per page), and the change during a sliding 7-day window (right y-axis, green).