Technische Universität München

Physik-Department
Experimental Physics with Cosmic Particles

and

TUM Institute for Advanced Study

The Origin of High-Energy Cosmic Particles: IceCube Neutrinos and the Blazar Case

Probing multi-messenger data by combining techniques from astrophysics, machine learning, and statistics

Theo Glauch

Vollständiger Abdruck der von der
Fakultät für Physik
der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender:	Prof. Dr. Alejandro Ibarra
Prüfende(r) der Dissertation:	1. Prof. Dr. Elisa Resconi
	2. Prof. Dr. Sherry Suyu

Die Dissertation wurde am 16.12.2020 bei der Technischen Universität München eingereicht und durch die Fakultät für Physik am 02.02.2021 angenommen.

Abstract

At the highest end of the energy spectrum, the universe is filled with three types of cosmic radiation[1] — cosmic rays, neutrinos, and gamma rays — of mostly unknown origin. We expect, however, that cosmic ray acceleration is tightly connected to neutrino and gamma-ray production. Hence, in this thesis, *multi-messenger* techniques are used to search for the long-sought sources of cosmic radiation.

In 2017, the IceCube Neutrino Observatory detected a neutrino[2] in proximity to a flaring gamma-ray blazar, TXS 0506+056. In this thesis, we probe their relationship in two ways: 1.) Using light curves of 3FGL sources, we estimate the chance coincidence probability to $\sim 0.1\%$ and 2.) Through a multi-messenger *dissection* of the region, we identify TXS 0506+056 as the only plausible neutrino emitter, making it the first compelling source of cosmic neutrinos. Based on this finding, three methods are developed to identify further candidates.

In the first step, IceCube analyses require the selection of suitable events based on their topology. This thesis, therefore, introduces a deep neural network (DNN) that — for the first time — classifies events into five categories directly based on their photon hits. The DNN is exemplarily used for a realtime selection of starting tracks that doubles event rates compared to its predecessor.

Complementary to single high-energy events, a 9-year sample of $\sim 650,000$ neutrino-induced tracks from the Northern Hemisphere is used to probe point-like astrophysical emission over the atmospheric background. We present an analysis that resolves previous biases by consistently modeling the likelihood function. A DNN further increases the energy resolution by $\sim 50\%$. Altogether, the discovery potential for point-like sources improves by $\sim 25\%$. Searching a catalog of potential neutrino candidates, we detect the starburst galaxy NGC 1068 as the *first* neutrino source with significance over 4σ (4.2σ).

Finally, we extend the multi-messenger *dissection* of TXS 0506+056 to 70 highly-energetic IceCube tracks, revealing a $3.23\,\sigma$ excess of IBL/HBL blazars with respect to random coincidence. For the future, an automated version of the dissection pipeline is triggered with every IceCube alert, enabling quasi-simultaneous follow-up measurements of potential counterparts. Through the combination of this pipeline and the DNN classifier, we identified the HBL 3HSP J095507.9+355101 as a promising neutrino counterpart. Overall, we find that IBL/HBL blazars are likely to be one population of astrophysical neutrinos and cosmic rays.

[1] often referred to as cosmic *messengers*

[2] IceCube-170922A

Zusammenfassung

Am oberen Ende des Energiespektrums findet man im Universum kosmische Strahlung, Neutrinos und Gammastrahlen mit ungeklärtem Ursprung. Es wird jedoch erwartet, dass die Produktion dieser drei *Messenger* in astrophysikalischen Quellen eng verbunden ist. Daher werden in dieser Arbeit *Multi-Messenger*-Methoden verwendet, um nach deren gemeinsamen Ursprung zu suchen.

Im Jahr 2017 entdeckte das IceCube Neutrino Observatory ein Neutrino[3] in der Nähe eines aktiven Blazars, TXS 0506+056. In dieser Arbeit evaluieren wir ihre Beziehung auf zwei Arten: 1.) Mit Hilfe von Lichtkurven schätzen wir die Zufallswahrscheinlichkeit auf $\sim$ 0.1% und 2.) Durch eine Multi-Messenger Analyse identifizieren wir TXS 0506+056 als den einzigen plausiblen Ursprung und damit als die erste Quelle kosmischer Neutrinos. Basierend darauf wurden drei Methoden zur Identifizierung weiterer Quellen entwickelt.

Im ersten Schritt erfordert jede IceCube Analyse die Auswahl geeigneter Ereignisse auf Basis ihrer Topologie. Diese Arbeit stellt daher ein Deep Neural Network (DNN) vor, das — zum ersten Mal — Ereignisse direkt auf Grundlage ihrer Erscheinung im Detektor in fünf Kategorien einordnet. Das DNN wird exemplarisch für eine Echtzeit-Selektion von startenden Spuren verwendet, die, verglichen zu ihrem Vorgänger, eine doppelt so hohe Ereignisrate hat.

Ergänzend zu einzelnen hochenergetischen Ereignissen werden $\sim$ 650.000 neutrino-induzierte Spuren aus der nördlichen Hemisphäre verwendet, um punktförmige astrophysikalische Emissionen über dem atmosphärischen Hintergrund zu identifizieren. Unsere Analyse behebt den Bias vorheriger Methoden durch eine konsistente Modellierung der Likelihood-Funktion. Ein DNN erhöht zudem die Energieauflösung um $\sim$ 50%. Insgesamt steigt das Entdeckungspotential für punktförmige Quellen um $\sim$ 25%. In einem Katalog potenzieller Neutrinoquellen haben wir die Starburst-Galaxie NGC 1068 als *erste* Quelle mit einer Signifikanz über 4σ ($4{,}2\sigma$) identifiziert.

Abschließend erweitern wir die Multi-Messenger Analyse von TXS 0506+056 auf 70 hochenergetische IceCube-Spuren, wobei wir einen $3,23\,\sigma$ Überschuss an IBL/HBL-Blazaren feststellen. Zukünftig wird mit jedem IceCube-Alert eine automatisierte Version der Analyse gestartet, um quasi-simultane elektromagnetische Beobachtungen von potenziellen Quellen zu ermöglichen. Durch die Kombination dieser Pipeline und der DNN Klassifizierung haben wir einen weiteren HBL[4] als potenzielle Neutrinoquelle identifiziert.

[3] IceCube-170922A

[4] 3HSP J095507.9+355101

Contents

List of Figures

List of Tables

1 Introduction

Seeing how we and our species fit in between the big bang and the closest science can take us to the end of time, is something that gives us deeply illuminating context. It allows us to see the human search for meaning and purpose in a different light, and with that to recognise that there is no ultimate answer hovering in the depths of space awaiting discovery. Instead, the context provided by the cosmological narrative frees us fully to develop our own, deeply and thoroughly subjective reasons for being.

— Brian Greene, Physicist [1]

To understand how little we know about the Universe we live in, we can take a journey to its high-energy end. There, invisible to the human eye, the Universe is filled with various types of cosmic radiation. Photons in the form of gamma rays, neutrinos, and charged particles fill the space and continuously penetrate the Earth's atmosphere without a clear trace of their origin and production mechanisms. It was in 1912 that Victor Hess discovered, for the first time, the high-energy radiation that we nowadays call *cosmic rays*.[2] With energies of up to 10^{20} eV, these charged nuclei are more than a million times more energetic than particles produced at human-made particle accelerators like the Large Hadron Collider (LHC) at CERN. Despite a century of intensive research, major questions about the high-energy cosmic radiation remain unanswered, including:

- What is the origin of high-energy cosmic radiation?
- How is high-energy cosmic radiation produced and what is the connection between the three *messengers* - gamma rays, neutrinos, and cosmic rays?
- What does high-energy cosmic radiation tell us about the cosmic evolution?

Today, around 100 hundred years after Victor Hess's balloon flights, we can measure cosmic radiation with the help of enormous observatories all around the globe. One of those, the IceCube Neutrino Observatory at the South Pole, plays a crucial role in the quest for cosmic-ray sources. In contrast to photons, which are produced in a large variety of electromagnetic processes, neutrinos only interact weakly. They are therefore closely connected to the acceleration and

---

[1] https://www.theguardian.com/science/2020/feb/06/brian-greene-theoretical-physicist-interview-until-the-end-of-time, accessed October 8th 2020

[2] Hess, "Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten".

interaction of cosmic rays in astrophysical environments. Unsurprisingly, the first detection of astrophysical neutrinos was thus recognized with the *Breakthrough of the Year* award of the American Physical Society in 2013[3]. Despite the cubic-kilometer volume of IceCube, the number of astrophysical neutrinos detected is still pretty low. Hence, it took another four years until the first compelling neutrino source candidate's observation, a flaring gamma-ray blazar [4] called TXS 0506+056[5,6,7].[8] While the observation of a high-energy neutrino, IceCube-170922A, played a crucial role as a trigger for follow-up observations, the association of the neutrino emission with TXS 0506+056 was only possible through an interplay of more than 17 telescopes that provided broad band measurements from radio, over optical and X-ray, up to the highest gamma-ray energies. The association, therefore, poses a milestone for *multi-messenger* astronomy.

This thesis combines innovative techniques from statistics, deep learning, and multi-messenger astronomy to provide a state of the art data analysis and explore the origin of cosmic neutrinos —and thereby cosmic rays.

The first three chapters comprise the introductory part of the thesis. **Chapter 2** discusses the current knowledge of astroparticle physics, including the production, propagation, and detection of cosmic radiation. Also, the chapter characterizes atmospheric air showers as the major background for searches of astrophysical neutrino sources. **Chapter 3** explains the fundamental detection principle of neutrinos, and summarizes the data acquisition principles and reconstruction algorithms of IceCube.

After the introductory part, **chapter 4** discusses the multi-messenger studies that led to the discovery of TXS 0506+056 as the compelling counterpart to the neutrino event IceCube-170922A. It includes calculating the chance coincidence probability of observing the neutrino coincident in space and time with an extragalactic flaring gamma-ray source,[9] as well as a comprehensive multi-frequency study of the sky region around it.[10] While this observation sheds first light on the origin of astrophysical neutrinos, further associations are required to complete the picture. Consequently, the following chapters present methods and analyses aiming to identify further source candidates.

A major challenge for any IceCube analysis is the selection of suitable events from the 2.6 kHz stream passing the hardware trigger of the detector.[11] In many cases, it can be advantageous to perform this selection based on the event's topology in the detector. **Chapter 5** therefore introduces a deep neural network (DNN) that can classify events directly based on their hits in the detector into five classes.[12] Despite the importance of the topology, IceCube currently does not have any comparable event classifier in place. Consequently, it is sometimes unclear which reconstruction algorithm is optimal for a given event, leading to mis-reconstructions. The DNN poses a solution for this dilemma, as it is sufficiently fast to be run very early in the data processing pipeline while also having a high prediction

---

[3] https://physicsworld.com/a/cosmic-neutrinos-named-physics-world-2013-breakthrough-of-the-year/, 09/2020

[4] Blazars are a rare type of elliptical galaxy hosting a supermassive black hole that produces a jet of particles in the direction of the Earth. See Urry and Padovani, "Unified schemes for radio-loud active galactic nuclei"

[5] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Samarai, et al., "Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert".

[6] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

[7] Padovani, Giommi, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source".

[8] Glauch, Padovani, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source".

[9] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Samarai, et al., "Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert".

[10] Padovani, Giommi, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source".

[11] M. Aartsen et al., "The IceCube Neutrino Observatory: Instrumentation and Online Systems".

[12] Kronmueller and Glauch, "Application of Deep Neural Networks to Event Type Classification in IceCube".

accuracy. As an exemplary application, **chapter 6** presents a new realtime selection of high-energy starting tracks of likely astrophysical origin. This new selection outperforms the previous version by a factor of more than two in rate while having an equal or better purity of astrophysical events.

Overall, IceCube measures neutrino events in a range between 100 GeV and several PeV. The largest fraction of these events originates from the interaction of cosmic rays in the atmosphere. Thus, astrophysical neutrinos are usually not identifiable on an event-to-event basis. However, emission from astrophysical neutrino point sources can still be identified through the clustering of events in space (and time). While this search strategy has a long tradition, previous implementations have suffered from biases caused by the mismodeling of the likelihood function. In contrast, the new analysis presented in **chapter 7** of this thesis follows a stringent statistical treatment to construct the point source likelihood from first principles. Instead of using analytical approximations, the underlying probability density functions are generated directly from Monte Carlo simulations, using kernel density estimation. Besides, a novel deep learning-based energy reconstruction algorithm improves the energy resolution by up to 50%. Overall, biases are resolved, the variance in the fit parameters is reduced, and the discovery potential is improved. After a detailed discussion of the methodology and the resulting improvements, 9 years of experimental data are analyzed using two search strategies: a sky scan of the Northern Hemisphere and a search in a gamma-ray based catalog of 110 potential source candidates.

From a multi-wavelength perspective, interactions of cosmic rays in astrophysical sources provide a natural connection between electromagnetic radiation and neutrinos through synchrotron emission of accelerated primary particles and the subsequent production of neutrinos and gamma rays through pion decays. Unambiguous identification of a neutrino source therefore, necessarily also requires a multi-wavelength counterpart. To facilitate the search for plausible counterparts, **chapter 8** presents a multi-frequency pipeline that dissects interesting sky regions, searching for blazar-like objects in a similar way as for the case of IceCube-170922A. Using tools from the United Nations *OpenUniverse* initiative, [13] the pipeline identifies counterpart candidates in radio, optical and X-ray, before running an automated gamma-ray analysis using data from the *Fermi*-LAT satellite. Finally, a machine-learning based algorithm analyses the multi-frequency emission measurements to classify the object automatically. In addition to the methodology, two applications are presented. Firstly, the multi-frequency analysis of chapter 4 is extended to all the well-reconstructed IceCube tracks to search for a statistical excess of blazars.[14] Secondly, the new analysis pipeline's full potential is discussed using the example of the event IceCube-200107A, for which a potential counterpart, 3HSP J095507.9+355101, was found.[15]

[13] https://openuniverse.asi.it/

[14] Giommi, Glauch, et al., "Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection".

[15] Giommi, Padovani, Oikonomou, et al., "3HSP J095507.9+355101: a flaring extreme blazar coincident in space and time with IceCube-200107A".

Finally, in **Chapter 9**, the results of the thesis are summarized and discussed in the light of future developments. Supplementary material can be found in the appendix.

2 Highly-Energetic Cosmic Particles

The quest for the origin and production of highly-energetic cosmic particles is one of the main outstanding challenges of modern astrophysics, which began with the first detection of cosmic rays in 1912.[1] Subsequent measurements of the sky have revealed that the high-energy cosmos is additionally filled with gamma rays[2] and neutrinos.[3] In fact, it is probable that there is a tight connection between the origin of these particles. The first evidence for this connection was the compelling multi-frequency detection of TXS 0506+056 as a common source of broad band electromagnetic radiation and neutrinos. While cosmic rays couldn't be observed directly, the detection of astrophysical neutrinos from a source provides clear evidence for the acceleration of protons and heavier nuclei. In general, the combination of cosmic rays, neutrinos, and the full electromagnetic spectrum lays the foundation for *multi-messenger astronomy*, a research area that strives for a more comprehensive understanding of the processes and particles that are shaping the universe. This chapter reviews the current knowledge and challenges in the field and introduces the most relevant observatories for this thesis.

2.1 Characterisation of the High-Energy Universe

While there is good reason to believe that there is a tight connection between neutrinos, gamma rays, and cosmic rays, each of their fluxes shows specific features connected to their underlying physical properties and production mechanisms. Before discussing the connection of the messenger in astrophysical source environments, we will, therefore, first discuss the production, propagation, and detection of the messengers separately.

2.1.1 Cosmic Rays

The term *cosmic rays* usually refers to any type of highly-energetic charged particles reaching the Earth's atmosphere from outer space. In a more narrow definition, it specifically describes the flux of hadronic particles, i.e., protons and heavier nuclei. In Figure 2.1 an overview over the most recent cosmic-ray flux measurements is shown. The

[1] Hess, "Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten".

[2] Morrison, "On gamma-ray astronomy".

[3] M. Aartsen et al., "Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector".

plot combines data from several observatories, thereby spanning several orders of magnitude, from around 40 GeV up to the most energetic cosmic particles at 300 EeV. While at the lower end of the spectrum balloon/satellite-borne experiments as the *Alpha Magnetic Spectrometer* (AMS) are used, event statistics quickly decrease to higher-energies requiring larger detection areas to perform precise measurements. For this reason, the largest cosmic-ray detector, the Pierre-Auger Observatory, covers an area of around $3000\,\mathrm{km}^2$ with 1660 surface detectors in the Pampa of Argentina. In contrast to space-borne experiments, ground-based observatories can not directly measure the cosmic-rays but rely on the atmosphere as a calorimeter. Consequently, these detectors can not directly address the composition of the cosmic rays, but use models of the induced air showers to deduce information about the primary particle.

One of the remarkable features of the cosmic-ray flux is its simplicity over a large energy range. Despite a softening at PeV energies (the *knee*) and a re-hardening at EeV energies (the *ankle*) it can be described by simple power laws. Overall, the shape of the spectrum is inconsistent with black body radiation, which leads to the conclusion that cosmic-rays are not produced in thermal processes but rather actively accelerated in astrophysical source environments. While the origin of knee and ankle is still a matter of debate, there is a wide consensus that the flux below the knee is dominated by cosmic rays produced in our Galaxy, most likely through diffuse shock acceleration in supernova remnants (SNRs).[4] Early works have shown, in fact, that from simple energetic considerations one supernova every 30 years would be sufficient to provide the required energy output of Galactic cosmic-rays.[5] More sophisticated simulations show that the combination of hadrons accelerated in the magnetic fields of SNRs with standard diffusion models of the Galactic plane can provide a good description of the cosmic-ray energy spectrum up to energies of $10^7 - 10^9$ GeV.[6] The major uncertainty in these models is the complex magnetohydrodynamic structure of the Galaxy itself. For a simple approximation one can consider the larmor radius of a charged particle moving in a perpendicular and uniform magnetic field. In convenient units it can be written as

$$r_L = 33.6\,\mathrm{km}\left(\frac{pc}{\mathrm{GeV}}\right)\left(\frac{1}{Z}\right)\left(\frac{\mathrm{G}}{B}\right) \tag{2.1}$$

with momentum p, charge Z and the magnetic field strength B. Eq. (2.1) can be further simplified by defining the *rigidity* R, i.e the resistance of a particle against deflection in a magnetic field via

$$R = \frac{pc}{eZ} \tag{2.2}$$

and hence

$$r_L = 33.6\,\mathrm{km}\left(\frac{R}{GV}\right)\left(\frac{\mathrm{G}}{B}\right) \tag{2.3}$$

with, for example, the rigidity of a 10 GeV proton being equivalent to 10 GV. Consequently, the larmor radius for a 3 PeV proton in a homogeneous magnetic field with the strength of the milky way ($6\,\mu\mathrm{G}$[7]) is

[4] Blasi, "The Origin of Galactic Cosmic Rays".

[5] Ginzburg and Syrovatskii, *The Origin of Cosmic Rays*.

[6] Ptuskin, Zirakashvili, and Seo, "Spectrum of Galactic Cosmic Rays Accelerated in Supernova Remnants".

[7] Haverkorn, "Magnetic Fields in the Milky Way".

given by 0.54 pc, which is approximately the thickness of the Galactic disc.

In the SNR scenario, the knee can therefore be interpreted in two ways: 1.) as the maximal energy that can be reached by Galactic accelerators and 2.) as the energy at which cosmic-rays start to escape the magnetic field of the Galaxy.[8] Both interpretations are currently equally viable, and in any case, a sequence of knee features is expected for the different constituents of the cosmic-ray flux; see the H4a model in Figure 2.1. Assuming the knee at 3 PeV to be dominated by protons, a similar feature at 100 PeV would be expected through the cut-off of heavier elements. Observational evidence for that has been reported by KASCADE-Grande.[9]

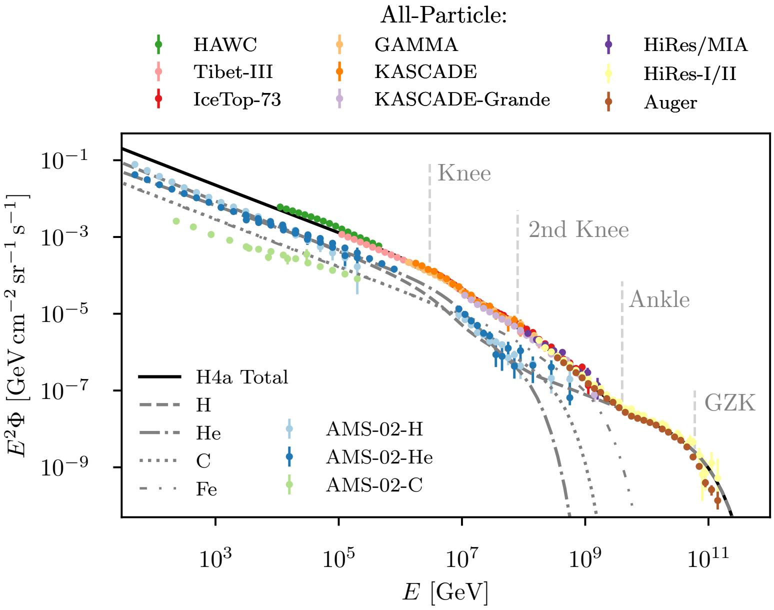

Figure 2.1: The energy flux density of cosmic rays between 40 GeV and several hundred EeV. Different measurements are shown in different colors. Overlayed, the H4a cosmic-ray model shows the energy-dependent contribution from the different nuclei. Important features of the spectrum are annotated. Data from `http://lpsc.in2p3.fr/crdb/` and the H4a model based on T. K. Gaisser, Stanev, and Tilav, “Cosmic Ray Energy Spectrum from Measurements of Air Showers”.

Observations beyond 100 PeV indicate a lightening of the composition, marking a likely transition to extragalactic cosmic rays. As seen from the H4a model[10] in Figure 2.1, the flux above this energy threshold can be purely described by hydrogen nuclei. Recent cross-correlation studies of those events have found evidence for a large scale anisotropy in the arrival direction of ultra-high energy cosmic rays (UHECR) consistent with the distribution of extragalactic matter as traced by the near-infrared 2MASS catalog.[11] In general, searches for the cosmic-ray sources are particularly challenging due to deflection in Galactic and extragalactic magnetic fields. Since the deflection decreases with growing rigidity, eq. (2.2), searches for extragalactic sources on small angular scales are usually limited to energies in the EeV range. Using magnetohydrodynamic models that incorporate local structures as well as intergalactic magnetic fields, it can be shown that the deflection angle for protons with energy $E_p = 4 \times 10^{19}$ eV is between 0.5 and 4 degrees depending on the propagation distance, see Figure 2.2.

[8] Alves Batista et al., “Open Questions in Cosmic-Ray Research at Ultrahigh Energies”.

[9] Apel et al., “Kneelike Structure in the Spectrum of the Heavy Component of Cosmic Rays Observed with KASCADE-Grande”.

[10] T. K. Gaisser, Stanev, and Tilav, “Cosmic Ray Energy Spectrum from Measurements of Air Showers”.

[11] Aab et al., “Observation of a Large-scale Anisotropy in the Arrival Directions of Cosmic Rays above 8×10^{18} eV”.

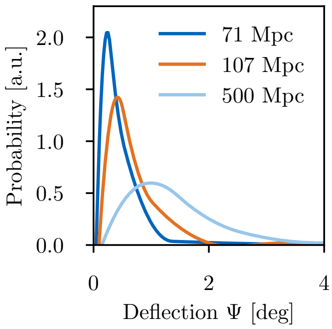

Figure 2.2: The distribution of deflection angles for cosmic rays with energy $E_p = 4 \times 10^{19}$ eV emitted at different distances as indicated in the plot. Models from Dolag et al., "Mapping deflections of ultrahigh energy cosmic rays in constrained simulations of extragalactic magnetic fields".

Above this energy threshold of $E_p = 4 \times 10^{19}$ eV, protons quickly loose their energy through interactions with photons of the cosmic microwave background (CMB) and the production of a delta resonance Δ^+

$$\gamma_{\text{CMB}} + p \rightarrow \Delta^+ \rightarrow p + \pi^0 \tag{2.4}$$

and

$$\gamma_{\text{CMB}} + p \rightarrow \Delta^+ \rightarrow n + \pi^+. \tag{2.5}$$

These processes thereby predict the GZK (Greisen–Zatsepin–Kuzmin) cutoff. At around 50 EeV in the cosmic-ray spectrum a cutoff is indeed observed, Figure 2.1, but an association with the GZK cutoff is not yet clearly established.[12] Despite the resonant pion production, another major energy loss, with lower energy threshold, is the production of electron-positron pairs via

$$\gamma_{\text{CMB}} + p \rightarrow p + e^+ + e^-. \tag{2.6}$$

Combining the two processes, the *cosmic-ray* horizon is defined as the distance up to which sources significantly contribute to the flux above a given energy threshold. As a useful measure, the attenuation length[13] can be used, see Figure 2.3. Note that, for cosmic rays the comparison of attenuation length to source distance is *optimistic* in a sense that the propagation length is in general longer than a source's luminosity distance due to random walk deflections in magnetic fields. Combining the cosmic-ray horizon in Figure 2.3 and the deflection in Figure 2.2 it is evident that cosmic ray-based astronomy is only possible up to a few hundred Megaparsec and therefore completely unfeasible for the majority of extragalactic objects in the universe. This poses a clear observational disadvantage against gamma rays and neutrinos as discussed in the following.

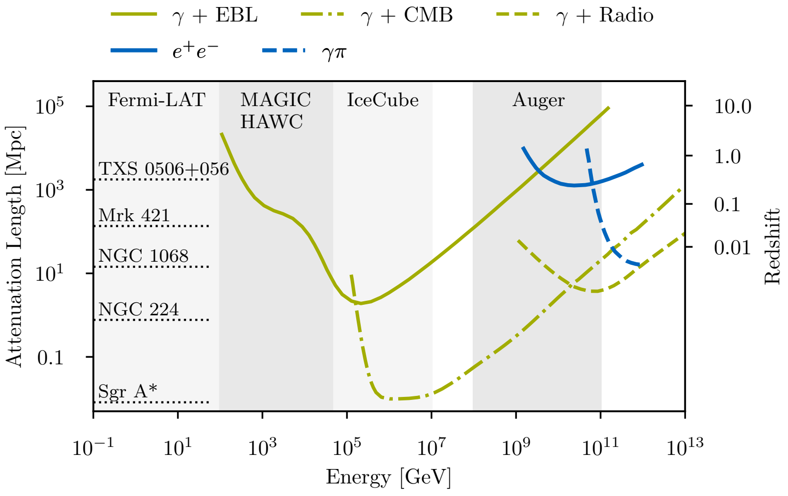

Figure 2.3: Attenuation length of cosmic rays (green) and gamma rays (blue as a function of their energy. Different line styles indicate different interaction processes as given in the legend. For comparison, the energy ranges of relevant observatories and the distance of well-known sources are shown. Model predictions from Harari, Mollerach, and Roulet, "On the ultrahigh energy cosmic ray horizon" and De Angelis, Galanti, and Roncadelli, "Transparency of the Universe to gamma rays".

[12] J. Abraham et al., "Observation of the suppression of the flux of cosmic rays above 4×10^{19}eV".

[13] The distance λ after which the initial flux has dropped to 1/e of its initial value. In general $N = N_0 \cdot e^{-x/\lambda}$.

2.1.2 *Gamma-Rays*

Gamma rays describe the part of the electromagnetic spectrum with photon energies above 100 keV. They play a crucial role in high-energy astrophysics and are produced in both, Galactic and extragalactic sources as seen in the *Fermi*-LAT skymap showing the arrival direction of photons in an energy range between 1 GeV and 300 GeV, Figure 2.4.

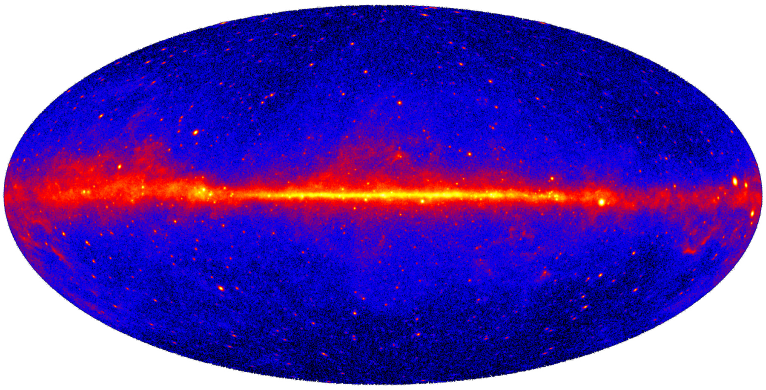

Figure 2.4: *Fermi*-LAT skymap above 1 GeV with 5 years of exposure in Galactic coordinates. Galactic gamma-ray emission is clearly visible as the horizontal band separating the two hemispheres. Outside of the galactic plane, small dots indicate the presence of extragalactic point sources. Plot taken from https://svs.gsfc.nasa.gov/11342; accessed December 12th 2020.

In contrast to cosmic rays, gamma rays are produced as secondary particles in a large variety of processes such as

- Inverse compton scattering: $e + \gamma \to e + \gamma *$, with the final state photon $\gamma *$ being more energetic than the original one.
- Bremsstrahlung: $e + N \to e + N + \gamma$ for some nucleus N
- Synchrotron radiation: $e \to e + \gamma$ in a magnetic field
- Meson Decay: $\pi^0 \to \gamma + \gamma$, where the pion is produced through proton-proton or proton-gamma interaction (see section 2.2 for details) .

Astrophysical source environments a rich of free electrons and charged ions, a physical state referred to as *plasma*. Through accretion, fusion, and acceleration of these particles, photons are additionally injected into the system. In the resulting sea of electromagnetically interacting particles, all of the previously mentioned processes are likely to occur. Unsurprisingly, the list of sources creating gamma rays, therefore, spans from Galactic objects as pulsars [14] and supernovas, over the diffuse emission from cosmic rays interacting with the interstellar medium, to extragalactic sources such as gamma-ray bursts and blazars (see section 2.3). Although the acceleration and interaction of cosmic rays is likely to produce a fraction of the detected gamma rays, most of their flux is produced in leptonic processes. In the following, we will briefly review the two dominant components of the gamma-ray sky

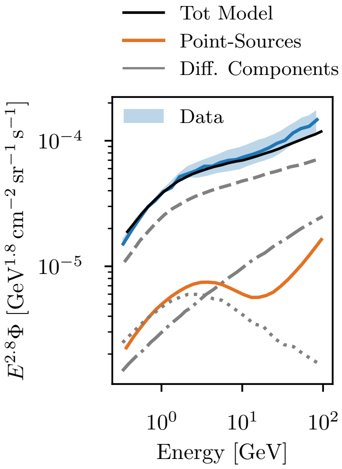

Figure 2.5: Different components of the Galactic gamma-ray flux according to the KRA$_\gamma$ model. The diffuse emission is shown in grey. Dashed, dashed-dotted, and dotted lines represent the emission from π^0 decay, inverse Compton, and bremsstrahlung, respectively. The solid, orange line indicates the contribution from gamma-ray point sources. Adapted from Gaggero, Urbano, et al., "Gamma-ray sky points to radial gradients in cosmic-ray transport".

[14] rapidly spinning neutron stars

- **Diffuse Galactic emission** is the product of cosmic rays interacting with the gas and radiation fields in the interstellar medium. Out of those processes, neutral (and charged) pions are created that further decay into gamma rays (and neutrinos). Besides, electrons created in the decay of charged pions can produce gamma rays through inverse Compton scattering and bremsstrahlung. Both components, the π^0 decay and the electromagnetic production of gamma rays, can be traced via the measurement of the Galactic gas density and the electron synchrotron emission, respectively. Combining those measurements with diffusion models of cosmic rays in the Galaxy, the diffuse Galactic gamma-ray emission can be well modeled[15].[16] One prominent example, the KRA_γ model, is shown in Figure 2.5. It can be seen that the dominant contribution of the Galactic flux originates from the decay of neutral pions, while point sources and other diffuse emission processes are subdominant. As neutral pions are inevitably produced together with charged pions that decay into neutrinos, the Galactic plane provides a guaranteed source of neutrino emission up to PeV energies.[17]

- The **extragalactic gamma-ray background** (EGB) is the integrated emission of all resolved and unresolved extragalactic GeV gamma-ray sources. The latest *Fermi* 4FGL catalog provides the largest sample of those objects so far. Of the 5064 sources detected in the entire sky, 3219 are of extragalactic origin.[18] Of these sources, blazars form by far the largest subclass with 3137 objects in total. Moreover, 42 radio galaxies and a hand full of galaxies have been detected. Unsurprisingly, blazars thereby account for the largest fraction of EGB. A quantitative study on that has been performed on the older 1FGL catalog by extrapolating the source count distribution of blazars below the *Fermi*-LAT detection threshold. It has been shown that blazars account for $50^{+12}_{-11}\%$ of the EGB photons, while $\sim$70% of the blazar contribution are already resolved.[19] Moreover, in the energy range above 100 GeV, blazars are responsible for almost the entire EGB emission, see Figure 2.6. The large number of gamma-ray blazars is specifically interesting as blazars have also been suggested as candidates for astrophysical neutrinos, and thereby cosmic rays, see chapter 2.2.

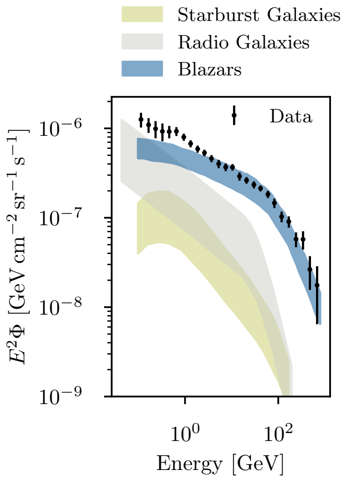

Figure 2.6: Energy-dependent contributions of different source types to the extragalactic gamma-ray background. Each source class is shown as a colored band. On top, the experimental measurement of the EGB is shown as black error bars. Plot adapted from Ajello et al., "The Origin of the Extragalactic Gamma-Ray Background and Implications for Dark-Matter Annihilation".

Beyond those two components, a smaller number of Galactic point sources has been observed. Specifically, the 4FGL catalog contains 239 pulsars, 40 supernova remnants, and 17 pulsar wind nebulae.[20]

Similarly to cosmic rays, gamma rays have a limited observational horizon, see Figure 2.3. Starting from around 100 GeV, at the upper end of *Fermi*-LAT's energy range, the interaction of gamma rays with photons from the extragalactic background light (EBL) [21] with subsequent production of electron-positron pairs becomes important. At 1 TeV, for example, EBL absorption limits the gamma-ray horizon to a redshift of around $z = 0.1$. This implies that the high-energy gamma-ray detection of objects that are further away requires very

[15] Ackermann et al., "Fermi-LAT Observations of the Diffuse γ-Ray Emission: Implications for Cosmic Rays and the Interstellar Medium".

[16] Gaggero, Urbano, et al., "Gamma-ray sky points to radial gradients in cosmic-ray transport".

[17] Gaggero, Grasso, et al., "The gamma-ray and neutrino sky: A consistent picture of Fermi-LAT, Milagro, and IceCube results".

[18] Abdollahi et al., "*Fermi* Large Area Telescope Fourth Source Catalog".

[19] Ajello et al., "The Origin of the Extragalactic Gamma-Ray Background and Implications for Dark-Matter Annihilation".

[20] Abdollahi et al., "*Fermi* Large Area Telescope Fourth Source Catalog".

[21] The EBL is the accumulated, diffuse radiation produced in star formation processes and other objects as AGNs.

long exposure times. The impact of EBL absorption increases up to a few hundred TeV where the interaction of gamma rays with photons from the Cosmic Microwave Background becomes dominant. The abundance of background photons in this energy range limits the gamma-ray horizon to a distance on the kiloparsec scale, making gamma-ray astronomy unfeasible.

2.1.3 Neutrinos

Neutrinos are uncharged fermions that are created through the weak force in one of three leptonic flavors: electron neutrinos ν_e, muon neutrinos ν_μ, and tau neutrinos ν_τ. In contrast to cosmic rays and gamma rays, neutrinos *only* interact through the weak force [22] and are thus neither deflected nor substantially absorbed. While this is advantageous to understand physical processes at the edge of the observable universe, neutrinos are extremely difficult to detect and event statistics are quite low. On the other hand, the limitation of neutrinos to hadronic processes guarantees that each astrophysical neutrino is directly tied to a cosmic ray acceleration site. Over the last ~ 10 years, the IceCube Neutrino Observatory (see chapter 3) has measured the flux of astrophysical neutrinos in various channels with high significance. A summary of the results are shown in Figure 2.7. The two most precise measurements come from the channel for through-going muons and cascades with power law spectral indices of $\gamma = 2.28^{+0.08}_{-0.09}$[23] and $\gamma = 2.53 \pm 0.07$,[24] respectively. Within the current statistics of these analyses, no more spectral complexity beyond the simple power law could be identified.

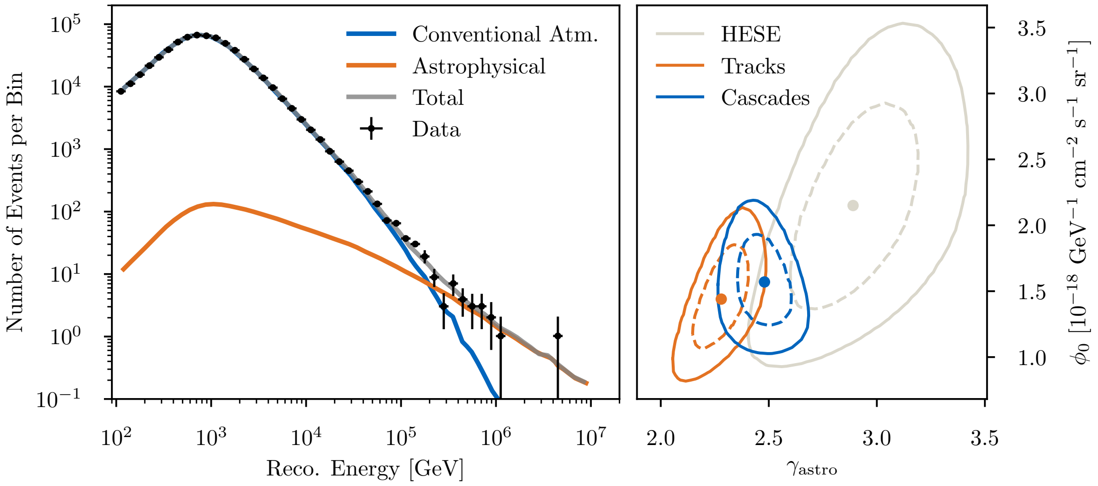

Figure 2.7: Different measurements of the astrophysical neutrino flux by IceCube. In the right plot the spectral index γ and the flux normalization ϕ_0 of a power law spectrum $\phi(E) = \phi_0 \cdot (E/100\,\mathrm{TeV})^{-\gamma}$ are shown for various detection channels. In the left plot, experimental data and best-fit model are shown for the sample of upgoing muon tracks. Adapted from Stettner, "Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data".

In summary, the astrophysical neutrino flux is well established in the energy range above $\sim 10\,\mathrm{TeV}$ up to a few PeV. While large efforts have been made, the bulk of the astrophysical neutrino flux's origin remains unknown. Searches for neutrino point sources, as well as the guaranteed diffuse galactic component [25], have yet been inconclusive [26]. However, it is expected that with growing statistics

[22] and gravitation

[23] Stettner, "Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data".

[24] M. Aartsen et al., "Characteristics of the diffuse astrophysical electron and tau neutrino flux with six years of IceCube high energy cascade data".

[25] see also section 2.1.2

[26] In chapter 7 we report on the first significant detection of a neutrino point source - the starburst galaxy NGC 1068

and improvements in the methodology, sources will start to appear over the diffuse background. While the case for point sources is complicated, as the underlying source count distributions and source spectra are speculative, recent IceCube analyses searching for the (guaranteed) Galactic diffuse component have already a sensitivity level of 58% of the KRA_γ baseline prediction.[27]

2.2 *The Multi-Messenger Connection*

The interactions of protons with ambient radiation fields or matter and the subsequent cascade of particles provides a natural connection between the three previously discussed *messengers*. Although this connection is theoretically on solid ground, a final proof can only be given through the experimental detection of neutrinos from a point source as a smoking gun. In this section, we will review and discuss the physical processes and particle spectra that are relevant in the scope of this thesis.

2.2.1 *Production of Gamma-Rays and Neutrinos from Pions*

In the multi-messenger scenario the most probable production of neutrinos happens through the decay of light mesons, specifically pions produced in hadronic interactions of cosmic rays. Being unstable particles, pions decay with a short lifetime of 2.6×10^{-8} s (π^0) and 8.4×10^{-17} s ($\pi^\pm$), respectively.[28] The major decay channels are

$$\pi^0 \rightarrow \gamma + \gamma \tag{2.7}$$
$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{2.8}$$
$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu \tag{2.9}$$

with the muons further decaying into

$$\mu^+ \rightarrow e^+ + \bar{\nu}_\mu + \nu_e \tag{2.10}$$
$$\mu^- \rightarrow \mathrm{e}^- + \nu_\mu + \bar{\nu}_e \tag{2.11}$$

resulting in a neutrino flavor ratio of $(\nu_\mu : \nu_e : \nu_\tau) = (2 : 1 : 0)$. Since the final state leptons (photons) are much lighter than the initial pion, they share around the same energy. Hence the two gamma-ray photons in eq. (2.7) carry one half of the pion energy, whereas each lepton from the decay chain of eq. (2.8) and eq. (2.10) carries around one fourth. There are two primary processes believed to produce these pions: photo-meson and proton-proton production.

2.2.2 *Photo-Meson Production*

In this production channel protons interact with ambient photons to create pions via

$$p + \gamma \rightarrow n + \pi^+ \tag{2.12}$$
$$p + \gamma \rightarrow p + \pi^0, \tag{2.13}$$

[27] M. Aartsen et al., "Search for Sources of Astrophysical Neutrinos Using Seven Years of IceCube Cascade Events".

[28] Tanabashi et al., "Review of Particle Physics".

with a ratio of $K = 1$ between charged and neutral pions. The required (lab frame [29]) proton energy E'_p for this process can be calculated through relativistic kinematics to

$$E'_p = \frac{m_p m_\pi c^4}{2E'_t}\left(1 + \frac{m_\pi}{2m_p}\right) \approx 10^{17}\left(\frac{1\,\mathrm{eV}}{E'_t}\right)\mathrm{eV} \tag{2.14}$$

with the proton and pion mass m_p and m_π and the (lab frame) photon energy E'_t, respectively. Hence, assuming exemplarily a target spectrum with photon energy of 2 eV (optical part of the spectrum) the required cosmic-ray energy is 5×10^7GeV, which is just above the knee of the cosmic-ray spectrum.

Assuming the pion production to be close to the threshold energy, the final state particles are at rest in the center of mass frame. Transforming back to the lab frame, one finds for the energy ratio

$$\kappa = \frac{m_\pi \gamma c^2}{m_p \gamma c^2} = \frac{m_\pi}{m_p} \approx 0.1, \tag{2.15}$$

which is a bit smaller than the average value of 0.2 found from numerical studies of the interaction phase space.[30] Hence, the proton loses around 20% of its energy to the pion. The previous assumption of the processes happening close to the threshold energy is, in fact, pretty realistic, as the cross section peaks sharply around the threshold and is basically constant for higher energies, see Figure 2.8.[31] In astrophysical environments, photo-meson production happens in emission regions that are boosted with Lorentz factor Γ to the observer. This implies two relevant production scenarios: photons being stationary or co-moving towards the proton emission region. In any case, the relation between the observed proton energy an the energy in the emission region is given by

$$E_p = \Gamma \cdot E'_p \tag{2.16}$$

and hence, by combining the branching ratios in section 2.2.1 with the fractional energy transfer from protons to pions, the neutrino energy is related to the proton energy via

$$E_\nu = \frac{E_p}{4 \cdot 5} = \frac{\Gamma \cdot E'_p}{20} = \Gamma \cdot 10^{15}\left(\frac{5\,\mathrm{eV}}{E'_t}\right)\mathrm{eV}. \tag{2.17}$$

The observed photon energy of the target photons is related to their energy in the emission frame via $E_t = \Gamma \cdot E'_t$ (co-moving photon field) and $E_t = \Gamma / E'_t$ (stationary/external photon field)[32]. Plugging this result into eq. (2.17), we find a relationship between the observed neutrinos and target photon energies. For the cases of the co-moving photon field, this yields

$$E_\nu = \Gamma^2 \cdot 10^{15}\left(\frac{5\,\mathrm{eV}}{E_t}\right)\mathrm{eV} \tag{2.18}$$

and for the case of the stationary photon field

$$E_\nu = 10^{15}\left(\frac{5\,\mathrm{eV}}{E_t}\right)\mathrm{eV}. \tag{2.19}$$

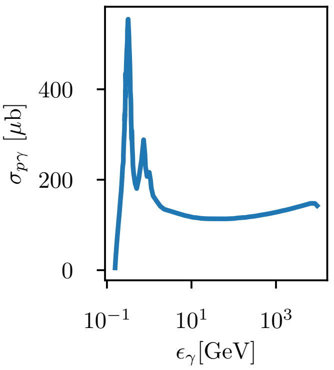

Figure 2.8: Interaction cross section for photo-hadronic interactions as a function of the required photon energy in the proton's rest frame. The cross section peaks strongly just above the pion rest mass and remains constants for higher energies. Adapted from Kelner and F. Aharonian, "Energy spectra of gamma-rays, electrons and neutrinos produced at interactions of relativistic protons with low energy radiation".

---

[29] for example the local emission region in an astrophysical jet.

[30] Kelner and F. Aharonian, "Energy spectra of gamma-rays, electrons and neutrinos produced at interactions of relativistic protons with low energy radiation".

[31] Ibid.

[32] assuming for the doppler factor δ that $\delta \approx \Gamma$, which is good assumption for high energies and small (boosted) angles

This implies that the photo-hadronic production of neutrinos requires much smaller (observed) photon energies if the field is stationary compared to a co-moving field.

In astrophysical source environments, the photo-meson production threshold manifests in flat neutrino spectra. Consider exemplarily a proton spectrum of $\mathrm{d}N_p/\mathrm{d}E_p \propto E_p^2$ as usually expected from Fermi acceleration[33] and a typical photon target spectrum of $\mathrm{d}N_\gamma/\mathrm{d}E_\gamma \propto E_\gamma^2$. In this case, the number of photons available above a given threshold E_p scales linear with the threshold. This implies that the optical depth, i.e., the probability for a proton to interact, increases linearly in E_p. Consequently, the power law spectrum of the resulting neutrino flux is harder than the initial proton flux with $\gamma_\nu = \gamma_p + 1$.[34]

2.2.3 *Proton-Proton Interaction*

In this production channel meson are produced through the inelastic collision of protons with ambient matter via processes like

$$p + p \rightarrow p + p + \pi^0 \tag{2.20}$$

$$p + p \rightarrow p + n + \pi^+ \tag{2.21}$$

$$p + p \rightarrow p + p + \pi^+ + \pi^- \tag{2.22}$$

with a usual ratio of $K = 2$ between charged and neutral pions.[35] Similarly as for photo-meson production, the energy threshold for pion production through proton-proton interaction can be calculated from the kinematics of the collision to

$$E_p = m_p c^2 \left(1 + \frac{m_\pi^2 + 4 m_\pi m_p}{2 m_p^2}\right). \tag{2.23}$$

Numerically, this results in a required proton kinetic energy in the lab frame [36] of

$$E_{kin} = E_p - m_p c^2 \approx 280\,\mathrm{MeV} \tag{2.24}$$

and thereby much smaller than for photo-meson production in (2.14). Thus, the pion and neutrino spectra from proton-proton interactions follows the initial proton spectrum. This has been verified using detailed numerical simulations, also showing that the average pion carries away around 17% of the proton's kinetic energy.[37]

2.2.4 *Multi-Messenger Constraints*

The previous sections have highlighted the natural connection between the three astrophysical messenger particles: cosmic rays, neutrinos, and gamma rays, and the energetics of their production. It has been shown that independently of the production process, the pion carries about 20% of the initial proton energy. In the subsequent processes, charged pions (neutral pions) decay into four leptons (two gamma-ray photons), each getting an equal share of the energy. Thus, each neutrino carries around 5% of the initial proton

[33] Bell, "The acceleration of cosmic rays in shock fronts – I".

[34] Waxman and J. N. Bahcall, "High-energy neutrinos from astrophysical sources: An Upper bound".

[35] Kelner, F. A. Aharonian, and Bugayov, "Energy spectra of gamma-rays, electrons and neutrinos produced at proton-proton interactions in the very high energy regime".

[36] where the second proton is at rest

[37] Kelner, F. A. Aharonian, and Bugayov, "Energy spectra of gamma-rays, electrons and neutrinos produced at proton-proton interactions in the very high energy regime".

energy ($E_\nu = 0.05E_p$), while each gamma-ray photon carries around 10% ($E_\gamma = 2E_\nu = 0.1E_p$). Although the specific realization of the neutrino production strongly depends on the astrophysical source environments, some model-independent constraints can be derived.

- The **Waxman-Bahcall Bound** was the first attempt to calculate a calorimetric limit on the possible *extragalactic* (muon) neutrino flux from the observed cosmic ray density[38] at PeV energies. The calculation assumes a (constant) injection spectrum of protons $\mathrm{d}\dot{N}_p/\mathrm{d}E_p = N_0 E_p^{-2}$ and an energy production rate of $\dot{E}_p \sim 5 \times 10^{44}\,\mathrm{erg\,Mpc^{-3}\,yr^{-1}}$ between 10^{19}eV and 10^{21}eV.[39] Under these assumptions the integrated energy density is given by

$$\dot{E}_p = \int_{10^{19}}^{10^{21}} E_p \times N_0 E_p^{-2}\,\mathrm{d}E_p = \int_{10^{19}}^{10^{21}} N_0 E_p^{-1}\,\mathrm{d}E = N_0 \times \log\left(\frac{10^{21}}{10^{19}}\right) \tag{2.25}$$

and hence

$$E_p^2 \frac{\mathrm{d}\dot{N}_p}{\mathrm{d}E_p} = N_0 = \frac{\dot{E}_p}{\log(10^{21}/10^{19})} = 5 \times 10^{44}\,\mathrm{erg\,Mpc^{-3}\,yr^{-1}}. \tag{2.26}$$

Assuming that the cosmic rays loose a fraction $f_{pp/p\gamma} < 1$ of their energy before leaving the source, the present neutrino energy density is given by

$$E_\nu^2 \frac{\mathrm{d}N_\nu}{\mathrm{d}E_\nu} \approx \zeta \cdot f_{pp/p\gamma} \cdot t_H \cdot \frac{\mathrm{d}\dot{N}_p}{\mathrm{d}E_p} E_p^2 \tag{2.27}$$

$$\approx \zeta \cdot f_{pp/p\gamma} \cdot t_H \cdot 5 \times 10^{44}\,\mathrm{erg\,Mpc^{-3}\,yr^{-1}} \tag{2.28}$$

with hubble time t_H and $E_\nu = 0.05E_p$ as discussed above. The factor ζ describes the process-dependent fractional energy transferred to neutrinos. In general

$$\zeta = \frac{3}{4}\frac{K}{K+1} \tag{2.29}$$

with the average fraction of charged to neutral pions K, i.e. $K \approx 1$ for photo-meson and $K \approx 2$ for proton-proton interaction. The ratio $K/(K+1)$ represents the fractional energy in charged pions, while the factor $3/4$ describes the ratio of the number of neutrinos to the total number of particles in the pion decay chain. Assuming that neutrinos are isotropic and move with the speed of light, multiplying eq. (2.27) by c and dividing by 4π gives the expected all-flavor neutrino flux[40]

$$E_\nu^2 \phi_\nu(E_\nu) \approx \zeta \cdot E_p^2 \frac{\mathrm{d}\dot{N}_p}{\mathrm{d}E_p} \cdot f_{pp/p\gamma} \cdot t_H \cdot \frac{c}{4\pi} \tag{2.30}$$

$$\approx \zeta \cdot 6 \times 10^{-8}\,\mathrm{GeV\,cm^{-2}\,s^{-1}\,sr^{-1}} \tag{2.31}$$

which is well consistent with the experimental result as illustrated in Figure 2.9 for the limit case of $f_{pp/p\gamma} = 1$ [41].

[38] Waxman and J. N. Bahcall, "High-energy neutrinos from astrophysical sources: An Upper bound".

[39] Waxman, "Cosmological origin for cosmic rays above 10**19-eV".

[40] Waxman and J. N. Bahcall, "High-energy neutrinos from astrophysical sources: An Upper bound".

[41] Note that the original derivation only considers the muon neutrino flux and hence uses a ζ of 0.25.

- **Gamma Ray - Neutrino connection**: With the production of pions and their subsequent decays, there is a natural connection between neutrinos and gamma rays. For photo-hadronic and proton-proton interactions the relationship of their fluxes after propagation on astrophysical distances is

$$\left.\frac{\partial \phi_\nu}{\partial E_\nu}\right|_{E_\nu=\frac{E_\gamma}{2}} = \eta \left.\frac{\partial \phi_\gamma}{\partial E_\gamma}\right|_{E_\gamma} . \tag{2.32}$$

with $\eta = 0.5$ for proton-proton interactions and $\eta = 1/3$ (neutrinos) or $\eta = 1/6$ (anti-neutrinos) in photo-meson production, respectively. The propagation on astrophysical distances assures that the (anti-)neutrino flavor ratio oscillates to $(\nu_\mu : \nu_e : \nu_\tau) = (1 : 1 : 1)$ allowing to treat all neutrino flavors in eq. (2.32) equally. In many theoretical considerations it can be useful to not consider the fluxes, but rather the energy budget. Using the production rates and energies of neutrinos and gamma rays from pion decays, [42], the relationship between the total luminosity per logarithmic interval, $E_X L_{E_X} = \mathrm{d}L_{E_X}/\mathrm{d}\log E_X$, can be written as

$$E_\gamma L_{E_\gamma} \approx \frac{4}{3K} \left.(E_\nu L_{E_\nu})\right|_{E_\nu=E_\gamma/2} . \tag{2.33}$$

with the fraction of charged to neutral pions, K [43]. Unfortunately, those relations are only of limited usage for experimental purposes, as 1.) gamma rays are also produced in many other, purely leptonic, processes (see section 2.1.2), 2.) gamma rays can interact with ambient photons at the production side and cascade down to lower energies and 3.) gamma rays in the energy range of IceCube are quickly attenuated through interaction with the CMB and EBL (see Figure 2.3) and are recovered at sub-TeV energies.[44] While 1.) simply poses an additional experimental background 2.) and 3.) require further theoretical considerations on how the hadronic gamma rays can be detected. The major interaction process driving the intra-source cascade is the electron-positron pair production in two-photon interaction, $\gamma + \gamma \to e^- + e^+$. From kinematics, the threshold for this interaction can be calculated to

$$E_1 E_2 = m_e^2 c^4 \tag{2.34}$$

with electron (positron) mass m_e and the energies E_1 and E_2 of the incident photons. Dedicated measurements show that the cross section peaks at $E_1 E_2 \approx 4 m_e^2 c^4$ or

$$E_t = E_2 \approx \left(\frac{1\,\mathrm{TeV}}{E_2}\right) \mathrm{eV}. \tag{2.35}$$

Hence a 1 TeV photon interacts most efficiently with an infrared target photon of $E_t = 1\,\mathrm{eV}$ energy. This puts us into a dilemma, as optical photons are expected in many source environments and are also a viable target photon field for the production of IceCube neutrinos, see (2.14). The cross section for pair production is, in

[42] For pion ratio $\pi^0/\pi^\pm = 1$ the energies of neutrinos and gamma rays relate as $E_\nu^{tot} = 3 \cdot E_\nu = 3/2 \cdot E_\gamma = 3/2 \cdot E_\gamma^{tot}/2 = 3/4 \cdot E_\gamma^{tot}$

[43] $K = 2$ for proton-proton interaction and $K = 1$ for photo-meson production

[44] Berezinsky and Smirnov, "Cosmic neutrinos of ultra-high energies and detection possibility".

fact, $\sim 10^4$ times larger than the cross section for photo-meson production.[45] In addition to the electromagnetic cascades caused by EBL and CMB interactions, it is therefore expected that gamma rays produced in π^0 decays will initiate source internal electromagnetic cascades, which become visible at lower energies. Dedicated studies indicate that sources have to be at least partially opaque to gamma rays to not violate the isotropic gamma-ray background as it would be, for example, the case when extending the *HESE 6yr* flux in Figure 2.9 to *Fermi*-LAT energies.[46] In this case, it could be promising to search for sources in the X-ray and MeV gamma-ray range.

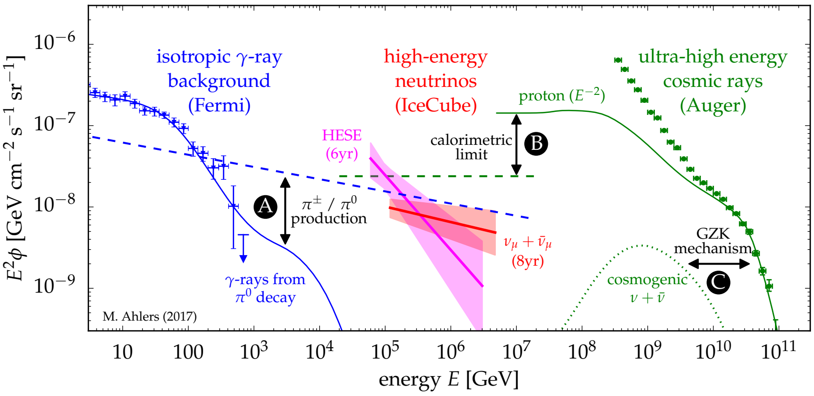

Figure 2.9: The connection between gamma rays, neutrinos, and cosmic rays. While the calorimetric limit provides an upper bound to the possible neutrino flux (B), the extrapolated neutrino flux should not exceed the gamma-ray flux if sources are assumed to be gamma-ray transparent (A). The interaction of cosmic rays with CMB photons and the subsequent decay of pions theoretically provides a further component to the neutrino flux. This *cosmogenic* neutrinos are, however, at energies above the IceCube energy range and have not been detected, yet. Courtesy of Markus Ahlers.

2.3 A Review of Cosmic Accelerators

To accelerate charged particles to the ultra-high energies of up to 10^{11} GeV, the particles must be confined in a small acceleration region through strong magnetic fields. The Hillas criteria gives a simple approximation for possible regions. It states that the maximum energy E_{max} that can be reached by an accelerator is given by the energy at which the gyro-radius r_L of a particle with charge Ze exceeds the size of the emission region R, allowing the particle to escape. Quantitatively, the criteria reads

$$E_{max} = ZeBR \approx Z \left(\frac{R}{\mathrm{pc}}\right) \left(\frac{B}{\mathrm{G}}\right) 10^{21}\,\mathrm{eV}. \tag{2.36}$$

Based on this equation, the Hillas plots, Figure 2.10, shows all astrophysical source classes that are theoretically able to accelerate charged particles to ultra-high energies. In fact, there is only a hand full of good candidates, three of which — AGNs, Starburst Galaxies. and SNRs — will be discussed in the following. While this list is certainly incomplete, it represents the most common Galactic and extragalactic production mechanism of cosmic rays. While the Hillas criterion gives a model-independent constraint on the possible source classes, it does not make a statement about the nature of acceleration

[45] Waxman and J. N. Bahcall, "High-energy neutrinos from astrophysical sources: An Upper bound".

[46] Murase, Guetta, and Ahlers, "Hidden Cosmic-Ray Accelerators as an Origin of TeV-PeV Cosmic Neutrinos".

mechanism. Although theories are very speculative, diffuse acceleration mechanisms — e.g., First and Second-Order Fermi acceleration[47] — are promising as they naturally produce a $\phi \propto E^{-2}$ cosmic-ray source spectrum that is (after propagation) consistent with our observations.[48]

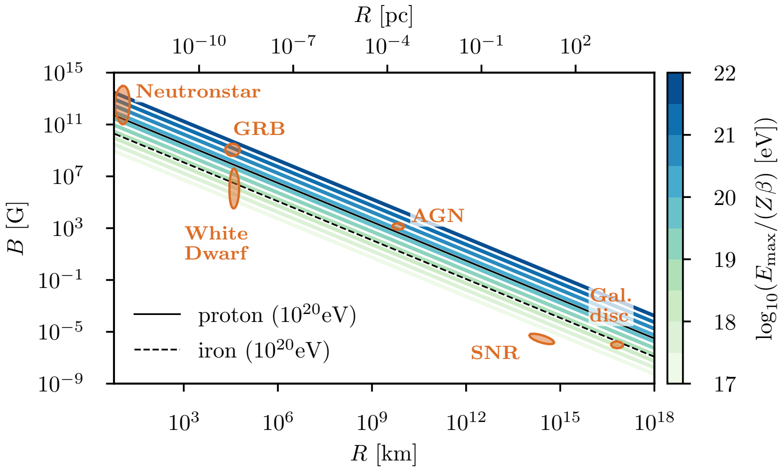

Figure 2.10: The Hillas plot showing Galactic and extragalactic sources classes that have sufficiently small emission regions and sufficiently strong magnetic fields to produce high-energy cosmic rays. The colored diagonal lines show the maximum energy that can be reached by a nucleus with total charge Ze for a specific configuration of the emission region. The solid and dashed black line show the exemplary cases for a 10^{20} eV proton and iron nuclei, respectively. Adapted from T. K. Gaisser, Engel, and Resconi, *Cosmic Rays and Particle Physics*, p.256.

2.3.1 Supernova Remnants

When stars run out of material to keep up their fusion process, the outward radiation pressure stops, and the star collapses. This happens particularly for very heavy stars ($M \geq 8M_\odot$) that burn all the way through until the core is made out of iron. During the collapse, the infalling matter compresses the star until some form of pressure leads to a new equilibrium. Depending on the star's mass, the new final state object is a white dwarf, a neutron star, or a black hole. Finally, through the abrupt ending of the collapse, the infalling matter gets reflected, and the star explodes.[49]

For one subclass, the *shell-type* explosions, the ejected material produces a shell around the central objects that lasts for thousands of years and provides a good environment for the acceleration of particles. This picture is, for example, supported by the observation of radio and X-ray emission from highly-energetic electrons in the evolving shock wave of SN1006.[50] Besides, direct evidence for the presence of highly-energy particles comes from the observation of gamma rays going up to at least 100 TeV.[51] These observations imply high magnetic fields and, therefore, the possibility of also accelerating nuclei to high energies. If there are sufficiently dense clouds of ambient gas, they might interact with protons from the shock wave in proton-proton interactions. Through the subsequent decay of mesons (section 2.2.1) also neutrinos and gamma rays are produced with a spectrum $\phi(E) \propto E^{\gamma} \cdot \exp -E/E_C$ that follows approximately

[47] Fermi, "On the Origin of the Cosmic Radiation".

[48] T. K. Gaisser, Engel, and Resconi, *Cosmic Rays and Particle Physics*.

[49] Reynolds, "Supernova remnants at high energy."

[50] Koyama et al., "Evidence for shock acceleration of high-energy electrons in the supernova remnant SN1006".

[51] Abdalla et al., "H.E.S.S. observations of RX J1713.7−3946 with improved angular and spectral resolution: Evidence for gamma-ray emission extending beyond the X-ray emitting shell".

the initial cosmic ray spectrum with $\gamma \gtrsim 2$. The cut off energy E_C depends on the specific acceleration mechanism. In general, supernova remnants are believed to accelerate protons up to several PeV.[52]

If the gravitational pressure during the collapse is large enough, shell electrons are pressed into the atom's core leading to the process $p + e^- \rightarrow n + \nu_e$. As a result, a *neutron star* is formed. Since the angular momentum and the magnetic field are conserved during the collapse, the resulting neutron star quickly rotates, producing a strong magnetic field that can accelerate particles in the surrounding gas clouds up to relativistic energies. Such objects are called *pulsar wind nebula* and have also been speculated to be a source of cosmic rays and neutrinos. A recent IceCube analysis has, however, found no evidence for the production of neutrinos in PWN.[53]

2.3.2 *Starburst Galaxies*

Starburst galaxies are a special type of galaxies undergoing an episode with an extraordinary high star-formation rate of up to 100 times higher than the one in the Milky Way.[54] Surveys in the infrared show that the evolution of these objects increases up to redshifts of $z \approx 2$ where they have the highest density before decreasing to higher redshifts.[55]

While the class of starburst galaxies itself is very diverse and difficult to categorize, there is solid evidence that they are often triggered by galaxy merger events that channel fresh gas into the system, which allows the formation of massive stars.[56] Those young and extremely bright stars burn their fuel extremely quickly and finally explode as supernovas. Importantly, the dynamic timescale of the concentration of gas is comparable to the lifetime of those massive stars.[57] Hence, the exploding supernovas and the accompanying shock waves eject relativistic protons and electrons into an environment of dense gas. Additional relativistic particle might be injected through outflows of a central black hole.[58] Similar to the processes leading to the diffuse gamma ray (and neutrino) emission from the Milky Way (chapter 2.1.2), protons in starburst galaxies can collide with the interstellar matter in inelastic scattering processes. The efficiency of these processes depends very much on the gas densities and the confinement of the cosmic rays in the galaxy's magnetic field, something that is not well understood.[59]

In contrast to other extragalactic sources classes, that are more likely to produce neutrinos via photo-hadronic interactions, the proton-proton collisions in starburst galaxies produce softer neutrino spectra with $\phi_\nu(E) \propto E_\nu^\gamma$ and $\gamma \gtrsim 2$, following approximately the cosmic-ray injection spectrum. In recent years, multi-wavelength studies have shown, however, that the contribution of starburst galaxies to the astrophysical neutrino flux at 100 TeV can not exceed more than $\sim 30\%$.[60] The argument is based on the fact, that the production of neutrinos is inevitably accompanied by electromagnetic cascades, and therefore gamma rays, that are visible in the *Fermi*-LAT energy

[52] Villante and Vissani, "How precisely neutrino emission from supernova remnants can be constrained by gamma ray observations?"

[53] M. Aartsen et al., "IceCube Search for High-Energy Neutrino Emission from TeV Pulsar Wind Nebulae".

[54] Tamborra, Ando, and Murase, "Star-forming galaxies as the origin of diffuse high-energy backgrounds: Gamma-ray and neutrino connections, and implications for starburst history".

[55] Gruppioni et al., "The Herschel PEP/HerMES Luminosity Function. I: Probing the Evolution of PACS selected Galaxies to z~4".

[56] Tamborra, Ando, and Murase, "Star-forming galaxies as the origin of diffuse high-energy backgrounds: Gamma-ray and neutrino connections, and implications for starburst history".

[57] Loeb and Waxman, "The Cumulative background of high energy neutrinos from starburst galaxies".

[58] Desai et al., "PAH Emission from Ultraluminous Infrared Galaxies".

[59] Loeb and Waxman, "The Cumulative background of high energy neutrinos from starburst galaxies".

[60] Bechtol et al., "Evidence against star-forming galaxies as the dominant source of IceCube neutrinos".

range. It has been shown, however, that a large fraction of the extragalactic *Fermi*-LAT detected gamma rays are already resolved in blazars[61] (section 2.1.2 and specifically Figure 2.6). The question of whether or not the gamma-ray flux related to a potential TeV neutrino production from starburst galaxies is compatible with the gamma-ray constraints depends very much on the assumed spectral indices[62].[63] It is also possible to loosen gamma-ray constraints if starburst galaxies are (partially) opaque to gamma rays due to the production of electron-positron pairs in the intense radiation fields of the galaxy itself. However, it has been argued that this is not expected to happen up to multi-TeV energies and thereby not affecting the arguments laid out above.[64] Overall starburst galaxies are a possible scenario for TeV neutrino production, and we will see in chapter 7 that, in fact, the current most significant neutrino point source — NGC 1068 — belongs to this source class.

2.3.3 *Active Galactic Nuclei and Blazars*

Active Galactic Nuclei (AGN) are compact regions in the middle of galaxies with high-luminosity broad band electromagnetic emission that can not be associated with stars. Since their first detection in the middle of the 20th century, they have been studied in a large range of frequency bands. Despite their versatile appearance many AGNs show similar features as[65]

- Strongly Doppler-broadened emission lines (high velocities)
- High variability
- Strong non-thermal emission / Extended radio jets
- Compact radio cores
- X-ray, γ-ray and TeV-emission.

With (bolometric) luminosities going up to $L_{bol} \approx 10^{48}\,\mathrm{erg\,s^{-1}}$, AGNs are further the most powerful non-explosive sources in the universe.[66] As a consequence, AGNs, and specifically the rare [67] subclass of *blazars*, (AGNs with a jet pointing towards us, i.e. viewing angle $\theta \lesssim 15^\circ - 20^\circ$) belong to the most promising neutrino emitters and have been extensively discussed in the literature[68,69,70].[71]

The taxonomy of AGNs has developed historically and is often confusing and overlapping. Most fundamentally, however, the underlying objects are believed to be pretty similar. It is now well established that the emission of AGNs is driven by the presence of a supermassive black hole ($> 10^6\,M_\odot$) that is accreting hot gas and plasma which is forming a thin disk,[72] see Figure 2.11. Based on this picture, the maximum luminosity of the object can be estimated using the Eddington criterion, i.e., requiring outward radiative pressure and inward gravity to be in equilibrium

$$L_{edd} = 1.4 \times 10^{38}\,\mathrm{erg\,s^{-1}}\,\frac{M}{M_\odot}. \tag{2.37}$$

[61] Ajello et al., "The Origin of the Extragalactic Gamma-Ray Background and Implications for Dark-Matter Annihilation".

[62] Bechtol et al., "Evidence against star-forming galaxies as the dominant source of IceCube neutrinos".

[63] Palladino et al., "IceCube Neutrinos from Hadronically Powered Gamma-Ray Galaxies".

[64] Bechtol et al., "Evidence against star-forming galaxies as the dominant source of IceCube neutrinos".

[65] Padovani et al., "Active galactic nuclei: what's in a name?"

[66] Ibid.

[67] A few thousand known objects in the entire universe

[68] Mannheim, "High-energy neutrinos from extragalactic jets".

[69] Halzen and Zas, "Neutrino Fluxes from Active Galaxies: A Model-Independent Estimate".

[70] Mücke et al., "BL Lac objects in the synchrotron proton blazar model".

[71] Padovani and Resconi, "Are both BL Lacs and pulsar wind nebulae the astrophysical counterparts of IceCube neutrino events?"

[72] Padovani et al., "Active galactic nuclei: what's in a name?"

When conditions are suitable, a fraction of the plasma is (magnetically) deflected and emitted perpendicular to the disk forming two jets of accelerated particles. Surrounding the central engine, the broad- and narrow-line region, BLR and NRL, consist of clouds of hot gas. While the emission lines in the BLR are wide and indicate supersonic velocities with temperatures of $\sim 10^4$ K, the narrow line region is further away from the central disk and has lower velocities and densities. Finally, the central region of the AGN is surrounded by a torus of hot gas and dust. Depending on the luminosity ratio between the central region to the galaxy, AGNs are called *Quasars* if the central region strongly outshines the rest of the galaxy [73] and *Seyfert* Galaxies otherwise. AGNs with extended ratio emission are called *Jetted AGN*. Over the years, different categorizations of AGNs have been established, mainly driven by phenomenological observations. Some of them are shown in Figure 2.11. For example, Seyfert I and Seyfert II galaxies are separated by the width of the emission lines [74], while radio-loud and radio-quiet AGN distinguish the objects based on their radio flux. Unification schemes haven shown, however, that things are much simpler, with the major difference between the classes coming from a) The viewing angle (see Figure 2.11) b) the accretion rate and c) the presence or absence of strong jets[75].[76] In this context, Seyfert I and Seyfert II are intrinsically the same type of objects, with the BLR being visible or not-visible, respectively. Using the unified schemes one can, in contrast to the historical, phenomenological classifications, define physically driven classes. Two important examples are:

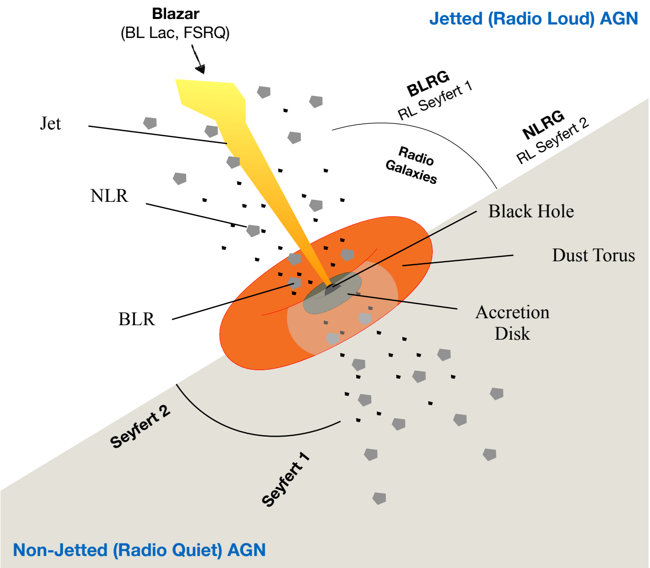

Figure 2.11: A schematic view of the different classes of AGNs depending on the viewing angle and the presence or absence of a jet. In addition to the most important morphological features, also the source classifications are shown. Abbreviations are given as follows: NLR - narrow line region, BLR - broad line region, BLRG - broad line radio galaxy, NLRG - narrow line radio galaxy. Adapted from Urry and Padovani, "Unified schemes for radio-loud active galactic nuclei".

- HEG/LEG AGNs: This classification separates AGNs into two classes, objects with high-excitation emission lines (HEG) and low-excitation emission lines (LEG). Most importantly, the difference between the two can be related to the accretion efficiency. While

[73] Which lets them looks like a quasi-stellar object (quasar)

[74] larger or smaller a full-width half-maximum of a few times 100 km s^{-1}

[75] Urry and Padovani, "Unified schemes for radio-loud active galactic nuclei".

[76] Padovani et al., "Active galactic nuclei: what's in a name?"

HEGs accrete in a radiative efficient way with $0.01 < L/L_{edd} < 1$, particles in LEGs lose less energy before crossing the event horizon of the black hole resulting in $0.01 < L/L_{edd}$.[77]

- FRI and FRII AGNs: This definition distinguishes jetted AGNs on the basis of R_{FR}, the ratio of the distance between the regions of highest surface brightness to the total extent of the source. AGNs are labeled as FRI if $R_{FR} < 0.5$ and as FRII otherwise. It is generally believed that the distinction between the two classes is caused by differences in the jets' morphology. While the jets of FRI galaxies are often distorted and plume-like, the jets of FRII galaxies are straight and dominated by highly supersonic flows.[78]

There is, in fact, a large overlap between the two classes. Almost all FRI galaxies are associated with LEGs, while FR II galaxies are HEGs.

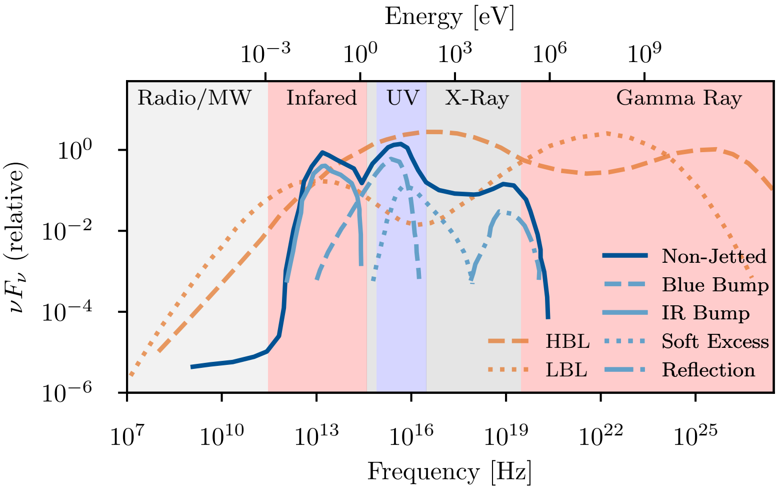

Figure 2.12: Three schematic spectral energy distributions of AGNs. Orange dashed and dotted lines show the typical double bump structure of blazars, here shown for a HSP and a LSP case, respectively. For comparison, an exemplary SED of a non-jetted AGN is shown in blue, with the different contributions as purple lines (see legend for details). Shaded bands in the background show the electromagnetic spectrum. The small band between infrared and UV is the optical spectrum. Plot adapted from Padovani, "Active Galactic Nuclei at All Wavelengths and from All Angles".

As noted before, AGNs emit radiation in a broad wavelength band. In Figure 2.12 three spectral energy distributions (SEDs) are illustrated: one of a typical non-jetted AGN, and two of blazars. A major difference is the presence or absence of galaxy features from the infrared to X-ray. While in the case of blazars the SED is very much dominated by the non-thermal emission of the jet, the non-jetted AGN shows clear features of the accretion and the galaxy itself. The former manifest as the *blue bump* at UV energies and is the result of superimposed thermal black body spectra radiated during the accretion. The *IR bump*, in contrast, is the result of thermal emission from heated dust in the torus.[79] Since the host galaxies of Quasars are nearly all elliptical with luminosity around 10^{44} erg/s,[80] the detection of the galactic features in the SED can be used to estimate the source's redshift.[81]

The SEDs of blazars significantly differs from other AGNs. It is dominated by non-thermal emission and shows a characteristic double bump structure. While the first bump is believed to be the result

[77] Padovani et al., "Active galactic nuclei: what's in a name?"

[78] Fanaroff and Riley, "The morphology of extragalactic radio sources of high and low luminosity".

[79] Padovani, "Active Galactic Nuclei at All Wavelengths and from All Angles".

[80] Ibid.

[81] Vanden Berk et al., "Composite Quasar Spectra from the Sloan Digital Sky Survey".

of the acceleration and subsequent synchrotron emission of charged particles, the second bump is associated with inverse Compton scattering or, potentially, hadronic production of gamma rays through π^0 decays.[82] The common presence of gamma rays makes blazars unique among AGNs and places them on the list of the most promising candidate sources for neutrinos. As blazars are of special importance for this thesis (see chapter 4 and 8), they are in the following discussed in more detail.

In Figure 2.13 a lepto-hadronic model of a typical blazar SED is shown on the example of 3HSP J095507.9+355101, a blazar associated with IceCube-200107A, see section 8.5. Separating the total flux into its different components shows that the electromagnetic emission consists of a complex combination of leptonic and hadronic processes. Although these lepto-hadronic models provide decent agreement with the actual data, they usually need to make strong assumptions on parameters as the geometry of the emission region, the spectra of the primary particles, and the magnetic field.[83] Despite those uncertainties, lepto-hadronic models give a comprehensive inside into the processed that *could* happen in neutrino sources.

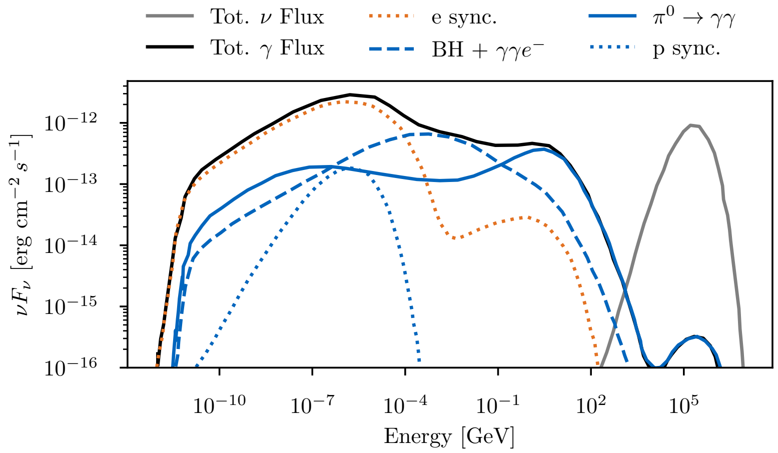

Figure 2.13: SED model for 3HSP J095507.9+355101. The blue and orange lines show the various hadronic and leptonic components as given in the legend. The resulting total electromagnetic flux is shown as black solid line. In addition, the predicted hard-spectrum neutrino flux is shown in grey. Plot adapted from Petropoulou, Oikonomou, et al., "Comprehensive Multimessenger Modeling of the Extreme Blazar 3HSP J095507.9+355101 and Predictions for IceCube".

.

While the synchrotron emission of the electrons dominates the first hump in Figure 2.13, the second hump is a combination of leptonic and (potentially) hadronic processes. In this specific example, the gamma-ray emission is assumed to be dominated by the electromagnetic cascades resulting from the photons produced in photo-hadronic interactions of protons (see section 2.2.2). Another purely leptonic contribution comes from the inverse component scattering of electrons with photons. A potential probe for the existence of protons in a blazar jet could come from Bethe-Heitler (BH) pair production, i.e., $p + \gamma \to p + e^+e^-$, which manifests as an electromagnetic cascade in the MeV or X-ray part of the spectrum. Finally, the same photo-hadronic interactions that produce the gamma rays also produce a hard-spectrum neutrino flux in an energy range of high TeV

---

[82] Padovani, "Active Galactic Nuclei at All Wavelengths and from All Angles".

[83] Petropoulou, Oikonomou, et al., "Comprehensive Multimessenger Modeling of the Extreme Blazar 3HSP J095507.9+355101 and Predictions for IceCube".

to PeV energies. This is a common feature of photo-meson production as discussed in section 2.2.2. Proton-proton interactions that can produce neutrinos with softer spectra/lower energies are believed to be irrelevant in blazars as the target matter densities are insufficient to facilitate the required interaction rates.

From a SED perspective blazars are usually classified based on the frequency of their synchrotron peak, ν^S_{peak}, as a measure for the acceleration power[84]

$$\text{Blazar Class} = \begin{cases} \text{LSP, for } \nu^S_{\text{peak}} < 10^{14} \\ \text{ISP, for } 10^{14} \leq \nu^S_{\text{peak}} < 10^{15} \\ \text{HSP, for } 10^{15} \leq \nu^S_{\text{peak}}. \end{cases} \tag{2.38}$$

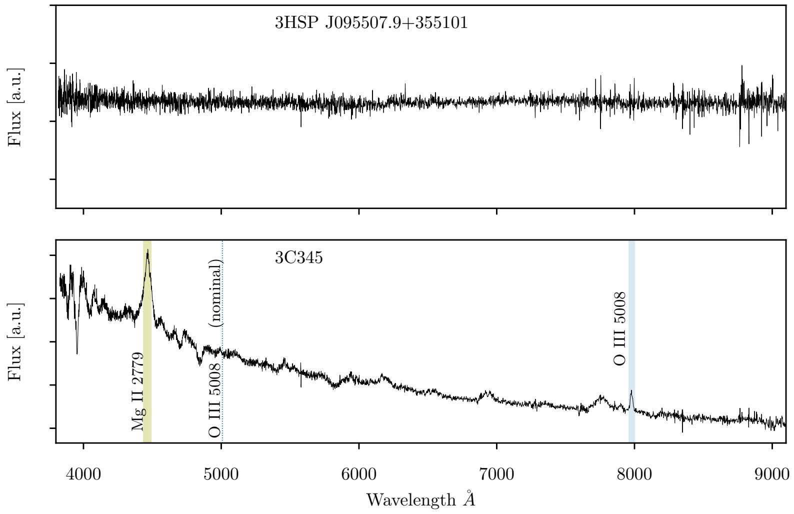

Figure 2.14: The optical spectrum of a typical BL Lac (top) and FSRQ (bottom) in arbitrary flux units. While the BL Lac spectrum is mostly flat, the FSRQ shows clear emission lines. Highlighted are the emission line of MgII and OIII (nominal and after accounting for the redshift). In addition to the blazar categorization, the spectrum can hence be used to determine the source's redshift, via $z = \lambda/\lambda_0 - 1$, with nominal and shifted wavelength λ_0 and λ, respectively. Data from the SDSS Science Archive Server; https://dr15.sdss.org/optical/spectrum/search.

Further fundamental differences are found in the optical spectrum of blazars, Figure 2.14. While one subclass of sources, flat-spectrum radio quasars (FSRQ), shows strong emission lines, the other class, BL Lacs, is nearly featureless. Quantitatively, the two classes are separated by the equivalent width (EW) of the emission lines and the existence/absence of the Ca H&K break, a stellar feature frequently found in elliptical galaxies.[85] Catalogs of blazars showed that the majority of sources *classified* as FSRQs have a small $\nu^S_{\text{peak}}$and a large luminosity in the *Fermi*-LAT energy range between 100 MeV to 100 GeV.[86] This observation triggered a long-standing debate of whether this is caused by a *blazar sequence* of the form, FSRQ → LBL → IBL → HBL [87],[88] or whether it is an effect of the emission lines of high-ν^S_{peak}FSRQs being strongly diluted from the non-thermal emission of the jets.[89] While both approaches *statistically* reproduce the observations, the detailed study of the optical emission from single high-ν^S_{peak}blazars have revealed objects with strong emission lines that are intrinsic to FSRQs[90].[91] This group of hard to identity, high ν^S_{peak}FSRQs is called *masquaraded* BL Lacs. They complete the list of blazars by adding the high-ν^S_{peak}analogous

[84] Padovani and Giommi, "The Connection between X-Ray– and Radio-selected BL Lacertae Objects".

[85] Giommi, Padovani, and Polenta, "A simplified view of blazars: the gamma-ray case".

[86] Ghisellini et al., "The Fermi blazar sequence".

[87] LBL = LSP BL Lac, IBL = ISP BL Lac, HBL = HSP BL Lac

[88] Ghisellini et al., "The Fermi blazar sequence".

[89] Giommi, Padovani, and Polenta, "A simplified view of blazars: the gamma-ray case".

[90] Padovani, Oikonomou, et al., "TXS 0506+056, the first cosmic neutrino source, is not a BL Lac".

[91] Keenan et al., "The Relativistic Jet Dichotomy and the End of the Blazar Sequence".

	BL Lacs	FSRQ	Masquerading BL Lacs
accretion	inefficient	efficient	efficient
EW	$< 5\text{Å}$	$> 5\text{Å}$	$> 5\text{Å}$
L/L_{edd}	$\lesssim 0.01$	$\gtrsim 0.01$	$\gtrsim 0.01$
ν^S_{peak}	any	$\lesssim 10^{14}$	$\gtrsim 10^{14}$
Associated with	LEG/FR I	HEG/FR II	

Table 2.1: Classification of blazars into BL Lacs, FSRQs, and Masquerading BL Lacs (FSRQs with strongly diluted emission liens). In contrast to BL Lacs, the latter two have strong emission lines and efficient accretion. The association with HEGs/FRII and LEGs/FRI given in the bottom row holds true for most of the objects. Table based on Giommi, Padovani, and Polenta, "A simplified view of blazars: the gamma-ray case".

to (standard) FSRQs. In this scheme, there is mainly one fundamental difference between FSRQs and BL Lacs. Specifically, FSRQs are nearly exclusively connected to radiatively efficient HEG galaxies, while BL Lacs are connected to radiatively inefficient LEG galaxies. A summary of the different blazar classes is shown in Table 2.1.

2.4 *Neutrino Oscillations*

The history of neutrino flavor oscillations dates back to the middle of the 20th century when physicist were struggling with the *solar neutrino problem*, i.e. the discovery that the electron neutrino flux at Earth is only about one-third of is expected from energetic considerations of the fusion processes in the sun[92].[93] As neutrinos are only detected in small numbers and the calculations of the sun required confidence in the Standard Solar Model, many physicist at this time believed, however, that the experiment was wrong. It was only in 1968 when Bruno Pontecorvo suggested a solution for this problem by introducing neutrino oscillations and it took another 44 years until the solar neutrino problem was finally solved in favor of neutrino oscillations by the Sudbury Neutrino Observatory in 2002.[94] In the paradigm of neutrino oscillations, the flavor eigenstates, which are relevant for neutrino interactions mediated by the weak force, are a superposition of the mass eigenstates, that are relevant for propagation in vacuum. In a modern way this can be written as

$$\begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} = U \begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} . \tag{2.39}$$

and since U is unitary ($U^{-1} = U^\dagger$)

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U^\dagger \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{2.40}$$

The matrix U is called the *Pontecorvo–Maki–Nakagawa–Sakata* matrix (PMNS matrix) and has three free mixing angles (θ_{12}, θ_{23}, and θ_{13}) and one complex phase (δ_{CP}).[95] A full parameterization is given

[92] Davis, Harmer, and Hoffman, "Search for Neutrinos from the Sun".

[93] J. N. Bahcall, N. A. Bahcall, and Shaviv, "Present Status of the Theoretical Predictions for the ^{37}Cl Solar-Neutrino Experiment".

[94] Ahmad et al., "Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory".

[95] Maki, Nakagawa, and Sakata, "Remarks on the unified model of elementary particles".

by

$$U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \tag{2.41}$$

$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{\mathrm{CP}}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{\mathrm{CP}}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{\mathrm{CP}}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{\mathrm{CP}}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{\mathrm{CP}}} & c_{23}c_{13} \end{pmatrix} \tag{2.42}$$

where s_{ij} and c_{ij} are used as simplified notation for $\sin\theta_{ij}$ and $\cos\theta_{ij}$, respectively. The current (July 2020) best-fit values in degrees [°] from a global analysis of all experimental data are given as[96] [97]

- $\theta_{sol} \approx \theta_{12} = 33.82^{+0.78}_{-0.76}$
- $\theta_{atm} \approx \theta_{23} = 48.3^{+1.1}_{-1.9}$
- $\theta_{13} = 8.61^{+0.13}_{-0.13}$
- $\delta_{\mathrm{CP}} = 222^{+38}_{-28}$.

While two of these angles, θ_{12} and θ_{23}, are found to be large, θ_{13} is much smaller, but unequal zero. The two large mixing angles play a dominant role in the oscillation of solar neutrinos (θ_{12}, conversion of ν_e to ν_τ) and atmospheric neutrinos (θ_{23}, conversion of ν_μ to ν_τ).

Most importantly, the detection of neutrino oscillations implies that neutrino masses are non-zero and unequal between the mass eigenstates, see eq. (2.43). Latest result by the KATRIN collaboration place an upper limit to the electron neutrino mass to 1.1 eV[98] and cosmological constraints require the sum of the neutrino masses to be smaller than 0.23 eV.[99]

After production, the neutrino mass eigenstate propagated freely according to the vacuum Hamiltonian. The propagation can be described in terms of a plane quantum-mechanical wave with the frequencies depending on the mass of the neutrinos, i.e., heavier neutrinos oscillate faster than lighter ones. The interference of these waves at the point of observation allows the neutrino of a given initial flavor to be observed in a different one. Using relativistic approximation in the full 3-flavor model, the transition probabilities are given in natural units by[100]

$$P_{\alpha\to\beta} = \delta_{\alpha\beta} - 4\sum_{i>j} \mathrm{Re}\left(U^*_{\alpha i}U_{\beta i}U_{\alpha j}U^*_{\beta j}\right)\sin^2\left(\frac{\Delta m^2_{ij}L}{4E}\right) \tag{2.43}$$

$$+ 2\sum_{i>j} \mathrm{Im}\left(U^*_{\alpha i}U_{\beta i}U_{\alpha j}U^*_{\beta j}\right)\sin\left(\frac{\Delta m^2_{ij}L}{2E}\right) \tag{2.44}$$

for a relativistic neutrino with energy E travelling a distance of L. From this equation it is evident that the mixing angles define the *amplitude* of the oscillation while the energy, propagation length and mass difference define the *frequency*. For interpretability it is useful

[96] Esteban et al., "Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23}, δ_{CP}, and the mass ordering".

[97] http://www.nu-fit.org/?q=node/211

[98] Aker et al., "Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN".

[99] Ade et al., "Planck 2013 results. XVI. Cosmological parameters".

[100] T. K. Gaisser, Engel, and Resconi, *Cosmic Rays and Particle Physics*, p.154.

to convert the argument of the $\sin^2$ in eq. (2.43) to astronomical units, i.e

$$\frac{\Delta m^2 L}{4E} \rightarrow 3.09 \times 10^7 \frac{\Delta m^2}{\mathrm{eV}^2} \frac{L}{\mathrm{pc}} \frac{\mathrm{PeV}}{E}. \tag{2.45}$$

Although the masses of the single eigenstates are not known, the mass differences are well measured to $7.54 \times 10^{-5}\,\mathrm{eV}^2$ and $2.43 \times 10^{-3}\,\mathrm{eV}^2$ for Δm_{12}^2 and Δm_{23}^2, respectively.[101] Concerning the signs of the mass differences it is unknown whether m_3 is larger or smaller than m_1, which is usually referred to as normal and inverted hierarchy. In order to calculate the relevant oscillation length for astrophysical neutrinos we consider a neutrino with energy of $E_\nu = 1\,\mathrm{PeV}$ and assume $\Delta m^2 = 10^{-4}\,\mathrm{eV}^2$. The resulting propagation length ΔL that is needed to shift the phase angle by 2π is then

$$\Delta L \approx \frac{2\pi}{3.09 \times 10^7 \cdot 10^{-4}}\,\mathrm{pc} = 2\,\mathrm{mpc}. \tag{2.46}$$

With an oscillation length on the milliparsec scale, neutrino oscillations are hence expected to play a significant role on astrophysical distances. In Figure 2.15 the flux ratio of the neutrino flavors is shown against the oscillation length L/E for an initial flavor ratio of $(\nu_\mu : \nu_e : \nu_\tau) = (2:1:0)$. It can be seen that the ratio quickly oscillates towards $(\nu_\mu : \nu_e : \nu_\tau) = (1:1:1)$ on astrophysical distances as expected from eq. (2.46).

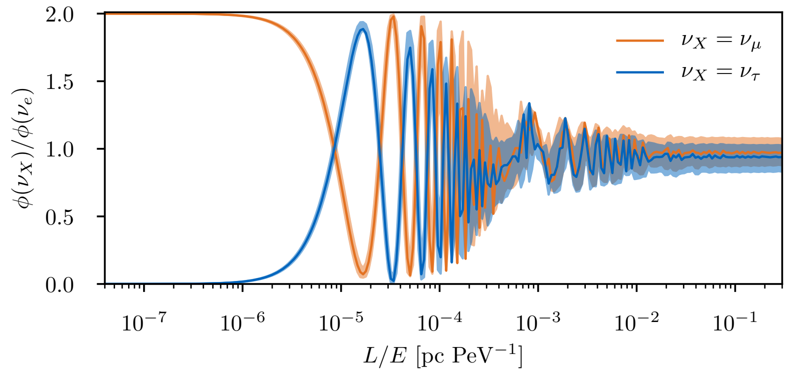

Figure 2.15: Neutrino flavor oscillations as a function of distance over energy in astrophysical units. On the y-axis the flux of tau (blue) and muon (orange) neutrinos is shown with respect to electron neutrinos. Starting from an initial ratio of $(\nu_\mu : \nu_e : \nu_\tau) = (2:1:0)$ the flux ratio converges quickly (on astrophysical distances) against $(\nu_\mu : \nu_e : \nu_\tau) = (1:1:1)$.

Consequently, it is usually assumed that the astrophysical neutrino flux reaching the Earth has an equal flavor distribution, independent of the initial production ratio. As neutrino oscillations happen on the milliparsec scale this statement is not only valid for the diffuse flux, where many sources at various distances are integrated, but also for a single neutrino source with an emission region not much smaller than the oscillation length.

2.5 *Atmospheric Air Showers*

When cosmic-ray nuclei reach the Earth' atmosphere, they undergo deep-inelastic hadronic interactions with the gas molecules, producing similar cascades of particles as those discussed for astrophysical

---

[101] Esteban et al., "Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23}, δ_{CP}, and the mass ordering".

source environments in section 2.2.3. In this sense, the atmosphere is a natural fixed target for a cosmic-ray beam dump. The cascade of particles produced in these interactions is called an *atmospheric air shower*. Analogous to the astrophysical case, *(atmospheric) muons* and *(atmospheric) neutrinos* are produced through the decay of mesons and thereby providing a background to the search for astrophysical neutrinos. Depending on their lifetime and energy, particles in the air shower undergo different interactions and energy losses while propagating through the atmosphere. The entire sequence of interactions and decays, as well as the (re-)productions of particles can be mathematically described through a set of coupled differential equations (cascade equations) for each particle type[102].[103]

In the following, we discuss some semi-analytical calculations to illustrate the origin and behavior of the most important flux components for neutrino telescopes. The complex hadronic cascade that is triggered by the interaction of the cosmic-ray nuclei with molecules in the atmosphere is composed of secondary protons and neutrons, as well as pions, kaons and heavier mesons. Neutral pions quickly decay into a pair of gamma rays, initiating a *electromagnetic cascade* of subsequent electron-positron pairs and bremsstrahlung. The other mesons decay further into

$$\pi^{\pm} \rightarrow \mu^{\pm} + \nu_{\mu}\,[+\bar{\nu}_{\mu}] \qquad (99.99\%), \tag{2.47}$$

$$K^{\pm} \rightarrow \mu^{\pm} + \nu_{\mu}\,[+\bar{\nu}_{\mu}] \qquad (64\%), \tag{2.48}$$

$$K^0_L \rightarrow \pi^{\mp} + \mu^{\pm} + \nu_{\mu}\,[+\bar{\nu}_{\mu}] \qquad (27\%), \tag{2.49}$$

$$K^0_L \rightarrow \pi^{\mp} + e^{\pm} + \nu_{e}\,[+\bar{\nu}_{e}] \qquad (41\%). \tag{2.50}$$

with the given branching ratios and the respective neutrino or antineutrino.[104] If the velocity of a muon is not sufficiently high, they further decay via

$$\mu^{\pm} \rightarrow e^{\pm} + \bar{\nu}_{\mu} + \nu_{e}\,[+\nu_{\mu} + \bar{\nu}_{e}], \tag{2.51}$$

In general, the development of the shower depends strongly on the energy of the particles involved. Muons with energies below $\approx 2\,\mathrm{GeV}$ have a decay length that is shorter than their average height of production (15 km) and hence decay via eq. (2.51) into electrons and neutrinos. In this energy regime neutrinos are thus produced with a flavor ratio of $(\nu_\mu : \nu_e : \nu_\tau) = (2 : 1 : 0)$. When the energy increases, the muon lifetime of $2.19 \times 10^{-6}\,\mathrm{s}$[105] is sufficient for them to reach the surface of the Earth before decaying. Hence, with growing energy, the muon decay and thereby the flux of electron neutrinos is strongly suppressed. Consequently, at 100 GeV the flux of electron neutrinos is already by a factor of 10 smaller than the muon neutrino flux.[106]

The energy loss pattern is another important characteristic of muons in the atmosphere. In general, the energy loss for long-lived particles can be described by

$$\frac{\mathrm{d}E}{\mathrm{d}X} = -a - E/b. \tag{2.52}$$

[102] T. K. Gaisser, Engel, and Resconi, *Cosmic Rays and Particle Physics*.

[103] Fedynitch et al., "Calculation of conventional and prompt lepton fluxes at very high energy".

[104] Tanabashi et al., "Review of Particle Physics".

[105] Ibid.

[106] T. K. Gaisser, Engel, and Resconi, *Cosmic Rays and Particle Physics*.

with critical energy E_C describing the energy at which the radiation losses (b) is equal to ionization losses (a). While E_C is around 87 MeV for electrons, the mass of the muons suppresses radiative processes, pushing the critical energy up to around $E_C = 500\,\mathrm{GeV}$.[107] Consequently, muons with energies below E_C can travel large distances before significantly loosing energy. The transition between ionization and radiation losses manifests in Figure 2.16 as a softening of the muon spectrum.

The ensemble of neutrinos produced through the decay of pions and kaons is called *conventional* atmospheric neutrinos. An energy and inclination dependent parameterization for the neutrino production from the decay of those two mesons in the atmosphere is given by[108]

$$\frac{\mathrm{d}\Phi_\nu}{\mathrm{d}E_\nu}(E,\theta) = \frac{\mathrm{d}\Phi_{CR}}{\mathrm{d}E_{CR}}(E,\theta) \times \left(\frac{A_{\pi\nu}}{1 + B_{\pi\nu} \cdot \frac{E_\nu \cos\theta}{\epsilon_\pi}} + \frac{A_{K\nu}}{1 + B_{K\nu} \cdot \frac{E_\nu \cos\theta}{\epsilon_K}} \right). \tag{2.53}$$

Here $A_{X\nu}$ and $B_{X\nu}$ are parameters describing the pion and kaon production and decay probabilities, respectively. Importantly, ϵ_K and ϵ_π denote the critical energies of kaons and pions, i.e. the energies above which re-interaction becomes dominant over decays. The critical energy of kaons, $\epsilon_K = 850\,\mathrm{GeV}$, is significantly larger than the one of pions with $\epsilon_\pi = 115\,\mathrm{GeV}$[109] . Plugging these values into eq. (2.53) the conventional atmospheric neutrino flux can be roughly divided into three energy regions

- $E_\nu < \epsilon_\pi = 115\,\mathrm{GeV}$: Below the pion critical energy, the denominators in eq. (2.53) are close to 1, and the neutrino flux spectrum follows the primary cosmic ray spectrum with a power law index of $\gamma = 2.7$.

- $\epsilon_\pi < E_\nu < \epsilon_K$: In this regime, pions lose significant energy, making the production of neutrinos less efficient. Therefore the production of neutrinos from kaon decays starts to dominate.

- $E_\nu > \epsilon_K$: Both kaons and pions significantly loose energy before decaying. As a result, the spectrum gets softer with a spectral index of around $\gamma = 3.7$.

All these effects are visible in the numerical solution of the full cascade equations, see Figure 2.16.

Despite kaons and pions, a minor fraction of charmed D and Λ mesons are produced in the cosmic rays' primary interaction with the atmosphere. The decay times of these heavy mesons are extremely short, producing a *prompt* neutrino flux[110] with a spectral index of $\gamma = 2.7$, following the initial cosmic ray spectral index. Depending on the neutrino flavor, the prompt component starts to dominate the total atmospheric neutrino flux between $\sim 20\,\mathrm{TeV}$ and several hundred TeV (Figure 2.16). In addition to muon and electron neutrinos,

[107] Ibid.

[108] Ibid.

[109] Ibid.

[110] Ibid.

the prompt flux is expected to produce a minimal number of tau neutrinos. Up to today, the prompt flux has not been detected by any analysis.[111]

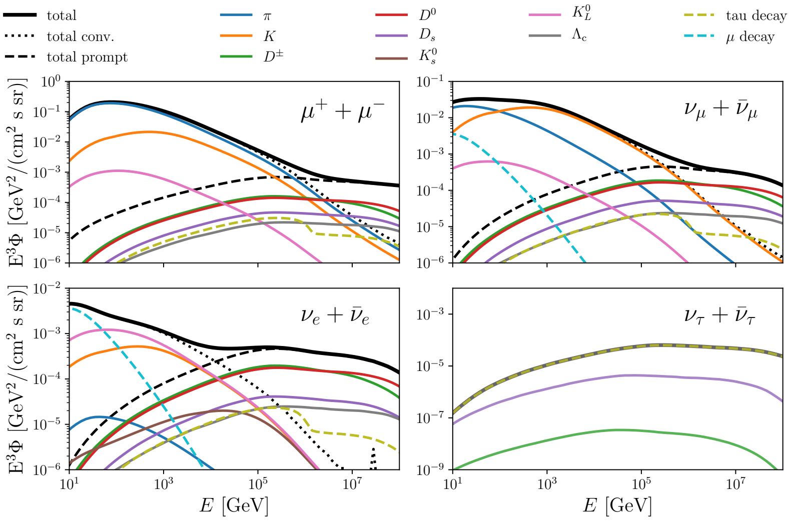

Figure 2.16: The spectrum of atmospheric muons (top left) and neutrinos (others) as calculated with MCEq. Conventional and prompt contributions are shown in different colors as indicated in the legend. Plot adapted from Fedynitch et al., "Calculation of conventional and prompt lepton fluxes at very high energy"

2.6 *Important Astronomical Observatories*

At the end of this chapter, we briefly introduce three astronomical observatories that play an important role in multi-messenger astrophysics and the rest of this thesis. A detailed discussion of the IceCube Neutrino Observatory can be found in chapter 3.

2.6.1 *Fermi-LAT*

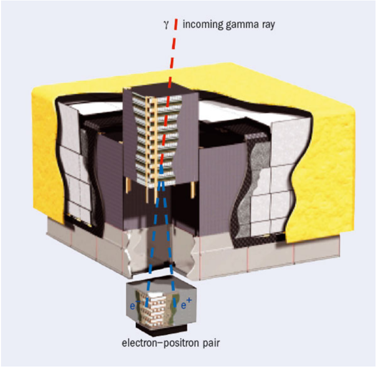

Figure 2.17: The detection principle of the Large Area Telescope (LAT) on board the Fermi Gamma-ray Space Telescope. An incoming gamma ray produces a pair of electrons and positrons which are reconstructed by silicon trackers and a calorimeter. Graphic taken from Atwood et al., "The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission" .

The Fermi Gamma-ray Space Telescope is a pair conversion telescope, measuring gamma rays from a low orbit. Its main instrument is the Larger Area Telescope (LAT), which is used to perform an all-sky survey of gamma-ray emission in the energy range between 100 MeV and 100 GeV every 3 hours.[112]

The detection principle of the LAT is illustrated in Figure 2.17. An incident gamma ray interacts with a heavy nuclei, tungsten, in this case, to create an electron-positron pair. The subsequent electromagnetic shower is tracked through a few layers of silicon before its energy is measured in a calorimeter made out of the scintillator material cesium iodid. As the initial gamma rays' energy is much higher than the rest mass of the electron and positron, the shower is strongly boosted, enabling a good directional reconstruction. To avoid the large background of cosmic rays the LAT is surrounded by an anti-coincidence detector (ACD) made out of scintillator tiles.

[111] Stettner, "Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data".

[112] Atwood et al., "The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission".

2.6.2 *OVRO*

The OVRO 40 meter Telescope is a radio telescope at the Owens Valley Radio Observatory in California, US. It gained major importance as an accompanying radio telescope for the *Fermi*-LAT satellite. In its current observation campaign around 1800 *Fermi*-LAT detected blazars are observed at 15 GHz about twice a week [113].[114] It thereby provides an unprecedented opportunity to study the multi-frequency behavior of blazars with time simultaneous gamma ray and radio data.

2.6.3 *Swift*

The Neil Gehrels Swift Observatory is a NASA-operated satellite dedicated to the multi-frequency observation of gamma-ray burst (GRBs). Swift carries three different instruments: 1.) the Burst Alert Telescope (BAT) is a large-field of view telescope in the energy range between 15 keV to 150 keV that is used to quickly locate GRBs. For precise follow-up observations, the 2.) X-ray Telescope (XRT) allows to take X-ray spectra in an energy range between 0.3 keV and 10 keV and the 3.) UV/Optical Telescope (UVOT) takes spectral measurements in the optical range between 170 nm and 600 nm.[115] With this multi-frequency instrumentation, Swift is well equipped for analyzing the regions around IceCube high-energy neutrinos, as shown in chapter 8.

---

[113] https://www.astro.caltech.edu/ovroblazars/

[114] Richards et al., “Blazars in the Fermi Era: The OVRO 40 m Telescope Monitoring Program”.

[115] Burrows et al., “The Swift X-ray Telescope”.

3 The IceCube Neutrino Observatory

The IceCube Neutrino Observatory is the world's first cubic-kilometer-scale neutrino detector, primarily built to detect and characterize the flux of high-energy neutrinos ranging between 100 GeV and several PeV. Like its predecessor AMANDA, it is located deep in the clear ice at the geographic South Pole, close to the Amundsen–Scott Station.[1] IceCube's measurements are based on the detection of Cherenkov photons produced by charged particles propagating in ice. Complementary observatories in the Northern Hemisphere — ANTARES,[2] KM3NeT[3] (both in the Mediterranean Sea), and Baikal-GVD[4] (at Lake Baikal) — use water as the target medium instead of ice. In addition, there are plans to construct another detector, P-ONE,[5] at the Cascadia Basin off the coast of Vancouver, Canada. This chapter starts with a review of the detection principle and the requirements for neutrino telescopes before discussing the design, event detection, reconstruction, and data processing of the IceCube Observatory in detail.

3.1 Detection Principle of Neutrinos

Neutrinos only interact through the weak force and are therefore not directly visible through any electromagnetic detection principle. An *indirect* detection and reconstruction is, however, possible by observing the light emission of secondary charged particles produced in the inelastic scattering of neutrinos with nuclei. As a result of the interaction, hadronic and electromagnetic cascades are produced. A convenient detection principle that facilitates large target volumes, is based on the production of Cherenkov radiation of charged particles in dense, but transparent, target media. The production of Cherenkov light is thereby a direct consequence of Maxwell's wave equations for a charged point-particle moving along a given axis with velocity v larger than the phase velocity of light in the medium. The calculations show that as long as the particle moves with velocities smaller than the local phase velocity, the field energy of the moving particle is deposited close to its path. Once the velocity exceeds the phase velocity, i.e. $v > c_0/n$ [6], the radiation escapes and is observable as *Cherenkov radiation*.[7] The wavelength-dependent deposition

[1] M. Aartsen et al., "The IceCube Neutrino Observatory: Instrumentation and Online Systems".

[2] Ageron et al., "ANTARES: The first undersea neutrino telescope".

[3] Margiotta, "The KM3NeT deep-sea neutrino telescope".

[4] Avrorin et al., "Baikal-GVD: status and prospects".

[5] Agostini et al., "The Pacific Ocean Neutrino Experiment".

[6] with refractive index n and vacuum speed of light c_0

[7] Jackson, *Classical Electrodynamics, 3rd Edition*.

of Cherenkov radiation is described through the Frank–Tamm equation

$$\frac{dN}{d\lambda dx} = \frac{2\pi\alpha}{\lambda^2} \sin^2 \theta_c \tag{3.1}$$

with fine-structure constant $\alpha = e^2/4\pi$. Following eq. (3.1) the emission rate is inversely proportional to the wavelength squared, which explains why Cherenkov radiation appears blue in the optical spectrum. The Cherenkov opening angle θ_c is defined as

$$\cos \theta_c = \frac{1}{n(\lambda)\beta}. \tag{3.2}$$

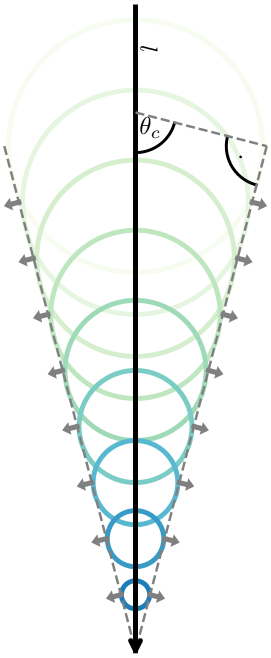

Figure 3.1: Geometry of the Cherenkov cone. A charged particle moving with a velocity above the local phase velocity of light in the medium emits a sequence of circular waves. The superposition of these waves produces a cone-like wavefront with an opening angle of θ_c. Courtesy of M. Huber.

Importantly, the refraction index n is a function of λ, which ensures that the integrated energy deposition of eq. (3.1) is finite. As usual, β is the ratio between the particle's velocity v and the vacuum speed of light, i.e. $\beta = v/c_0$. A visual explanation of the Cherenkov angle is given in Figure 3.1. At each time step, the moving particle emits a spherical wave that is propagating with velocity c_0/n. Like the Mach cone for supersonic movements, the superposition of all the spherical waves results in a Cherenkov cone. Consequently, the Cherenkov angle can be directly deduced from the ratio of the adjacent to the hypotenuse.

The relation between a particle's energy and the resulting number of Cherenkov photons is of large importance for the reconstruction of events in neutrino detectors. Through the interaction of the neutrino with a target nuclei and subsequent energy losses of secondary particles, a large number of hadronic and electromagnetic cascades is produced. In dedicated Geant4[8] simulations of electromagnetic and hadronic cascades in ice, it can be shown that the number of particles produced in these processes scales linearly with the energy of the initial primary. As each of these particles further radiates Cherenkov photons also the total photon count is linear to the energy of the primary.[9] Small corrections are due to heavier nuclei that do not have sufficient energy to reach the Cherenkov threshold.

An optimal medium for a Cherenkov detector is sufficiently dense to facilitate the production of Cherenkov photons while being transparent enough to keep scattering and absorption of the photons as low as possible. Two media that fulfill this criterion — while being available in large quantities – are ice and water. With a depth-independent refractive index of $n \sim 1.31$ South Pole ice, for example, has a Cherenkov angle of $\sim 41^\circ$ assuming $\beta \approx 1$. At a wavelength of $\sim 200\,\text{nm}$ the absorption of photons rises sharply, an effect known as *Urbach-Kante*.[10] For this reason the optimal wavelength for Cherenkov detectors is at $\sim 300\,\text{nm}$ where sufficient photons are produced (see eq. (3.1)) and absorption is low. The same reasoning applies when considering water as the target material.

[8] Agostinelli et al., "GEANT4–a simulation toolkit".

[9] Radel and Wiebusch, "Calculation of the Cherenkov light yield from electromagnetic cascades in ice with Geant4".

[10] Urbach, "The Long-Wavelength Edge of Photographic Sensitivity and of the Electronic Absorption of Solids".

3.2 *Requirements for Neutrino Telescopes*

For most energies the neutrino flux reaching the Earth is dominated by conventional atmospheric neutrinos. It is only above several hundreds of TeV that the diffuse astrophysical flux dominates (see Figure 2.7). Assuming the astrophysical flux to be dominated by point sources, this energy threshold might be lower locally, but nevertheless neutrino astronomy becomes only feasible above a few TeV of neutrino energy. At these energies neutrino interactions with nuclei are dominated by deep-inelastic scattering. Depending on the nature of the mediating gauge boson, there are two different interactions

- *Neutral-current* (NC) interactions

$$\nu_l + X \xrightarrow{Z^0} \nu_l + Y \tag{3.3}$$

- *Charged-current* (CC) interactions

$$\nu_l + X \xrightarrow{W^\pm} l^\pm + Y. \tag{3.4}$$

In either case, the initial nucleus X is destroyed, and a hadronic cascade Y is created. A major difference comes from the fact that neutral current interactions are mediated by the uncharged Z^0 boson, and hence only energy and momentum are transferred. On the other hand, in charged current processes, the charged $W^\pm$ also converts the primary neutrino into an (anti-)lepton of the same flavor.

Despite the interaction with a nucleus, neutrinos can also interact with the bound electrons of the target atoms. This process is, however, strongly suppressed in most parts of the energy range. An exception is the *Glashow resonance* of electron antineutrinos and bound electrons interacting with a center of mass energy close to the mass of the $W^\pm$ boson. The required neutrino energy is $m_W^2/(2m_e) \sim 6.3\,\mathrm{PeV}$. Formally, the interaction can be written as

$$\bar{\nu}_e + e^- \rightarrow W^- \rightarrow X. \tag{3.5}$$

Due to its short lifetime of only $\sim 3 \times 10^{-25}\,\mathrm{s}$[11] the W^- boson quickly decays into other particles. The branching ratios for the decay are $\sim 67.6\%$ for hadronic and $\sim 10.8\%$ for each leptonic channel ($W^- \rightarrow l^- \bar{\nu}_l$), respectively.[12] A summary plot of all cross sections is shown in Figure 3.2. It can be seen that the CC and NC cross section increase linearly until $\sim 10\,\mathrm{TeV}$. Above this energy the momentum transfer k becomes relevant in the boson propagator and the cross sections flattens to

$$\sigma_{CC/NC} \propto \left(\frac{1}{k^2 + m_{W/Z}^2}\right)^2 \times E \tag{3.6}$$

resulting in an effective cross section dependence of $\sigma \propto E^{0.36}$.[13] The ratio between charged current and neutral current cross sections, $\sigma_{CC}/\sigma_{NC} \approx 2.4$, is constant over the entire energy range.[14]

---

[11] Tanabashi et al., "Review of Particle Physics".

[12] Ibid.

[13] Formaggio and Zeller, "From eV to EeV: Neutrino Cross Sections Across Energy Scales".

[14] Ibid.

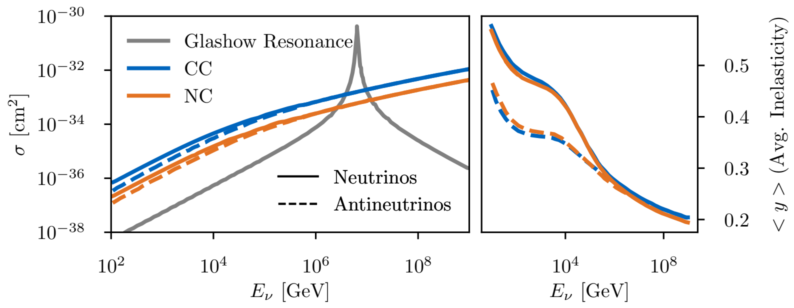

Figure 3.2: Deep-inelastic neutrino cross sections (left) and average inelasticities (right) for neutrino interactions with matter. Different colors represent different processes. Solid and dashed lines indicate neutrinos and antineutrinos, respectively. Data from Cooper-Sarkar, Mertsch, and Sarkar, "The high energy neutrino cross-section in the Standard Model and its uncertainty".

Finally, the discrepancy in neutrino and antineutrino cross section is the consequence of the $\mathcal{CP}$ asymmetry of the valence quarks in the target nuclei. The asymmetry is resolved above $\sim 10\,\mathrm{TeV}$ where vacuum fluctuations of quark/antiquark pairs become visible in large numbers. Consequently, the neutrino and antineutrino cross sections converge for high energies.[15]

According to eq. (3.3) and eq. (3.4) the deep-inelastic interaction of neutrinos results in a hadronic cascade and a neutrino or charged lepton, respectively. The energy shared between the initial neutrino and the nuclear cascades is called the inelasticity y with

$$y = \frac{E_{had}}{E_\nu}. \tag{3.7}$$

It can be seen from Figure 3.2 that the inelasticity depends on particles involved, as well as their energy. As a general trend, the inelasticity decreases toward higher energies, from $\sim 0.5/\sim 0.4$ (neutrinos/antineutrinos) at 100 GeV to ~ 0.25 at 1 PeV. This implies that for NC interactions, only a minor fraction of the energy is visible. In contrast for the reconstruction of CC events, the energy of the leptonic energy losses is most important.

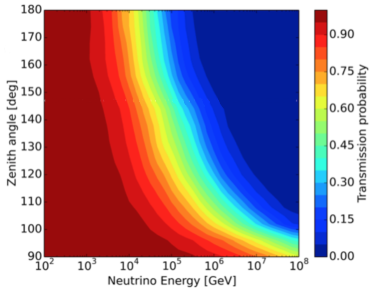

Figure 3.3: Energy-dependent transmission probability of neutrinos passing through the Earth. The zenith angle on the y-axis is measured in IceCube coordinates, i.e., starting from the South Pole. Plot taken from M. G. Aartsen et al., "Measurements using the inelasticity distribution of multi-TeV neutrino interactions in IceCube" .

Despite the background of conventional atmospheric neutrinos, neutrino telescopes are swamped by a large background of atmospheric muons. In order to suppress this background, the Earth can be used as a natural shield. In contrast to the muons, neutrinos are less affected when propagating through the Earth with absorption effects only being relevant towards higher energies. Formally, this can be written in terms of the mean free path length

$$L(E, \vec{x}) = \frac{1}{\sigma(E)\rho(\vec{x})} \tag{3.8}$$

with the total energy dependent neutrino cross section $\sigma(E)$ and the positional dependent density of the Earth $\rho(\vec{x})$. The probability that a neutrino is absorbed can hence be calculated as

$$P(E,\theta) = \int L(E, \vec{x})\, d\,|\vec{x}| = \int \frac{1}{\sigma(E)\rho(\vec{x})}\, d\,|\vec{x}| \tag{3.9}$$

with the integral going along the trajectory of the neutrino. From numerical calculations the energy and zenith dependent survival probability can be calculated using realistic Earth models. The result is

[15] Formaggio and Zeller, "From eV to EeV: Neutrino Cross Sections Across Energy Scales".

shown in Figure 3.3. Importantly, it can be seen that absorption of high-energy ($>$ 100 TeV) neutrinos becomes relevant at large declinations when the neutrinos are crossing the dense core of the Earth. As a result, high-energy events are usually observed from directions close to the horizon.

Historically, the size of a neutrino telescope was estimated using the Waxmann-Bahcall bound,[16] which, in fact, turns out to be quite close to recent measurements (also compare section 2.2.1). Following the historic argument, the number of neutrino-induced muons is about 20 to 50 per km^2 per year.[17] Consequently, cubic kilometer-scale detectors are needed to detect a sufficient amount of event to perform *neutrino astronomy*.

3.3 *Event Topologies*

When neutrinos interact through charged current or neutral current interactions, (eq. (3.3) and eq. (3.4)), secondary charged particles are produced. During propagation those particle further loose energy, creating subsequent electromagnetic and hadronic cascades. The sum of Cherenkov photons produced from all of these particles forms the *event signature*. Depending on the flavor of the initial neutrino, different secondaries, and most importantly, different leptons are produced. As the three leptons have different energy loss patterns and lifetimes, they produce unique event topologies. In the following the three major event signatures — tracks, cascades and double bangs — are discussed. Exemplary photon emission patterns for each of these signatures are shown in Figure 3.5, Figure 3.7 and Figure 3.8, respectively.

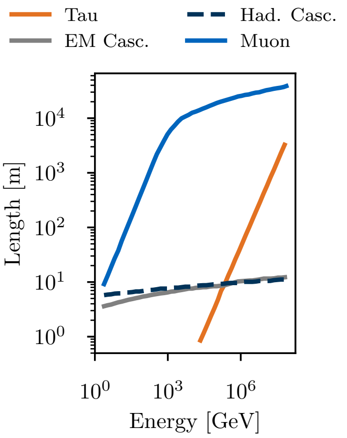

Figure 3.4: Mean propagation length for lepton tracks and cascades. The plot is based on water as target medium, but the behaviour in ice is very similar. Adapted from Kopper, "Performance Studies for theKM3NeT Neutrino Telescope".

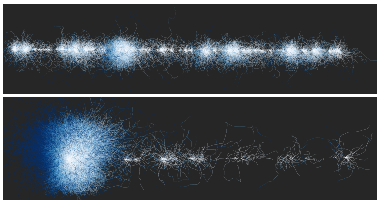

Figure 3.5: Photon trajectories for two $\sim$ 100 TeV muon tracks with the initial hadronic cascade being visible (bottom) and invisible (top). The trajectories are based on an IceCube simulation with the ice model described in section 3.6. The color gradient represents the time from white (early) to blue (late). The starting track on the bottom shows that photons are boosted into the event direction, therefore appearing later on the 'backside' of the cascade.

- **Tracks** are created in processes with a high-energy muon in the final state, e.g., ν_μ charged current interactions or muons from atmospheric showers. With a lifetime of 2.2×10^{-6} s,[18] muons with velocity close to the speed of light can travel several kilometres before decaying, see Figure 3.4. During the propagation through dense media, high-energy muons ($\gtrsim$ 100 GeV) loss their energy through four difference processes. *Ionization losses* occur through the interactions of the muon with the shell electron of the target nuclei. It is the dominant process below around 2 TeV and is basically independent of the muon energy. In contrast, there are three major energy-dependent radiative losses: muon-induced *pair production*, *photo-nuclear* interactions and *bremsstrahlung* in the magnetic fields of the target nuclei. Three exemplary Feynman diagrams for these processes are shown in Figure 3.6. Radiative losses are stochastic by nature and increase linearly in energy, forming a sequence of electromagnetic cascades. Figure 3.5 (top) exemplarily illustrates the Cherenkov emission for a 100 TeV muon track. It is clearly visible that the light pattern is created by a sequence of stochastic energy losses. Due to their elongated shape, tracks

[16] Waxman and J. N. Bahcall, "High-energy neutrinos from astrophysical sources: An Upper bound".

[17] T. Gaisser, "Neutrino astronomy: Physics goals, detector parameters".

[18] Tanabashi et al., "Review of Particle Physics".

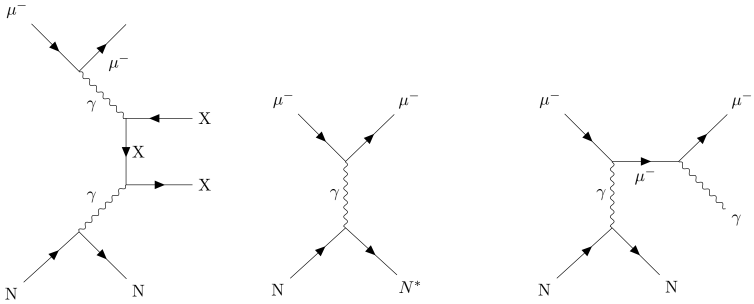

Figure 3.6: Feynman diagrams of muon radiative losses. From left to right one exemplary diagram is shown for pair production, photo-nuclear interaction and bremsstrahlung, respectively. N represents the nucleus and N^* its exited state. X represents any possible pair of electromagnetically interacting fermions, mostly electrons. Processes are shown for muons, but are also valid for anti-muons under CPT symmetry.

have a long lever arm, enabling a good directional reconstruction. In neutrino detectors, tracks usually appear in two types, through-going, and starting, with the primary neutrino interaction vertex outside and inside the fiducial volume, respectively. While the majority of tracks originates from atmospheric muons and neutrino induced charged current interactions, a minor fraction is produced through the decay of *tau*-leptons, or W bosons produced in Glashow resonances $W^- \rightarrow \bar{\nu}_\mu \mu^-$.

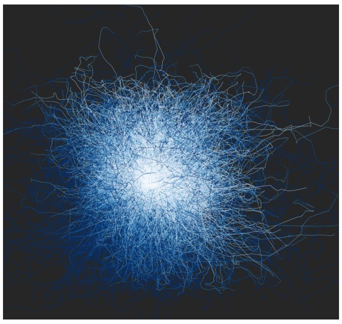

Figure 3.7: Exemplary photon trajectories for a $\sim 100\,\mathrm{TeV}$ cascade. The typical spherical emission profile is clearly visible. See Figure 3.5 for more details about the simulation.

- **Cascades** are produced in charged current ν_e and any kind of neutral current interaction. The characteristic feature of a cascade is its nearly spherical emission profile, which is the consequence of the short propagation length of electromagnetic and hadronic cascades, see Figure 3.4. In Figure 3.7 the Cherenkov photons produced in a $100\,\mathrm{TeV}$ simulated cascade are shown. It is clearly visible how the scattering of the Cherenkov photons diffuses the shower direction. As a result, the angular resolution of cascades is limited to $\sim 10^\circ$.[19] While this is disadvantageous for the identification of spatial clustering, the confinement of cascades in the detector enables the accurate reconstruction of the deposited energy. Depending on the neutrino interaction, there is a different relation between deposited and initial neutrino energy. For charged-current interactions, the entire neutrino energy is converted into the hadronic and electromagnetic cascades. In the case of neutral-current interactions, on the other hand, the final state neutrino still carries a fraction of its initial energy, and hence only a fraction of the total neutrino energy is visible; see Figure 3.2 for details.

- Similar to the muon also the tau lepton is unstable, but with a much smaller lifetime of only $2.9 \times 10^{-13}\,\mathrm{s}$.[20] As a result, charged current interactions of tau neutrinos produce two cascades, one at the interaction vertex and one during decay, connected by a short track. This topology is usually referred to as **double bang**. The separation length between the two cascades can be approximated using the energy E_τ of the tau lepton as

$$l_\tau \approx 50\,\mathrm{m} \times \left(\frac{E_\tau}{\mathrm{PeV}}\right). \tag{3.10}$$

[19] M. Aartsen et al., "Search for Sources of Astrophysical Neutrinos Using Seven Years of IceCube Cascade Events".

[20] Tanabashi et al., "Review of Particle Physics".

Consequently, low energy double bangs are basically indistinguishable from cascades. In general, there are two relevant decay channels for this topology: hadronically into pions and electromagnetically via $\tau \rightarrow \nu_\tau e \nu_e$ [21]. The hadronic decay is dominant with a branching ratio of 64.79% compared to the electromagnetic one with 17.82%.[22] In the remaining 17.39% of cases, the tau lepton decays leptonically into a muon, not producing a double bang but a track-like signature. While expected event numbers are low, the detection of double-bangs poses a clear trace of astrophysical neutrinos, as the production of tau neutrinos in atmospheric showers is strongly suppressed. Dedicated IceCube analysis have recently found the first two double bang candidates, being consistent with an astrophysical neutrino flavor ratio of $(\nu_\mu : \nu_e : \nu_\tau) = (1:1:1)$.[23]

Conveniently, the different topologies enable a broad range of physics analyses taking advantage of their respective properties. Specifically, the topologies help to study better the different components of the neutrino flux, e.g., flavors or the underlying particle physics. Technically, this requires, however, a clear identification of those topologies in the detector. An efficient approach for that is the deep learning-based classification algorithm presented in chapter 5. The classifier is trained to rapidly distinguish between the major topologies observed in IceCube, facilitating the application to large event samples very early in the data processing pipeline.

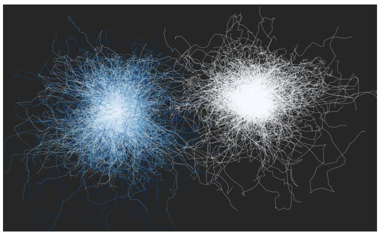

Figure 3.8: Photon trajectories for a double bang. While the first cascade is early (white), the second one happens significantly later (blue). See Figure 3.5 for more details about the simulation.

3.4 *The Design of the IceCube Neutrino Observatory*

IceCube consists of 86 strings with 5160 optical modules (DOMs), effectively instrumenting $\sim 1\,\mathrm{km}^3$ of clear ice in a depth of 1450 m to 2500 m close to the geographic South Pole. The detector's overall uptime is more than 99% and events are observed from the full sky.[24] IceCube consists of several sub-detectors:

- The main **In-Ice Array** contains 79 of the 86 strings with 60 DOMs on each string. The strings are arranged in an approximately hexagonal shape with a horizontal spacing of $\sim 125\,\mathrm{m}$ and a vertical spacing of 17 m. The large intra-string spacing is the major driver of IceCube's large effective area. As a result, the In-Ice Array is only triggered by events with energy larger than $\sim 100\,\mathrm{GeV}$.[25]

- **Deep-Core** is a region in the center of the detector with 6 additional strings and a denser string spacing of only 72 m (horizontal) and 7-10 m (vertical). Of the 60 DOMs, 50 are located in the clearest part of the detector below the dust layer (see section 3.6). An additional 10 are located above the dust layer as veto cap against atmospheric muons. In contrast to the In-Ice Array, six Deep-Core strings are equipped with high quantum efficiency DOMs.

---

[21] with the respective charges and chiralities

[22] Tanabashi et al., "Review of Particle Physics".

[23] Stachurska, "First Double Cascade Tau Neutrino Candidates in IceCube and a New Measurement of the Flavor Composition".

[24] M. Aartsen et al., "The IceCube Neutrino Observatory: Instrumentation and Online Systems".

[25] Ibid.

All together, the energy threshold can be lowered to as small as $\sim 10\,\mathrm{GeV}$, facilitating the measurements of oscillation patterns of atmospheric neutrinos.[26]

- On the surface, an array of 82 Cherenkov tanks called **Ice Top** measures cosmic-ray air showers. Each tank consists of two DOMs. The main goal is to measure and characterize atmospheric air showers between the first and second knee (section 2.1.1), as well as to veto atmospheric muons and atmospheric neutrinos for the In-Ice Array.[27]

A sketch of the detector design is shown in Figure 3.9 (left) together with a close-up illustration of a single DOM (right). The DOMs are the major measurement devices of the detector. Inside a spherical glass pressure vessel, they are housing a 10" Hamamatsu R7081 Photomultiplier Tube (PMT), high voltage generation, waveform digitization, calibration LEDs, digital communication hardware, and all the required peripheral electronics. The glass vessel itself has a diameter of 33 cm and is the combination of two symmetrical half-spheres. Each DOM is connected to three wires in the string using a penetrator. Two of these cables go directly to the IceCube Lab at the surface, transferring power and bidirectional digital communication streams. The remaining cable connects to neighboring DOMs, enabling fast detection of coincident hits.[28]

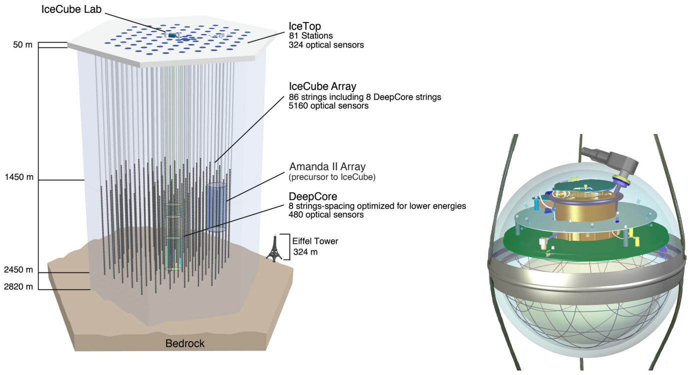

Figure 3.9: Schematic view on the IceCube Detector (left) and a single In-Ice DOM (right). The 86 IceCube strings reach from the surface — with the IceCube laboratory (ICL) — down to a depth of around 2500m, close to the Antarctic bedrock. The instrumented volume with its 5160 DOMs is indicated through circles on the strings. Images from M. Aartsen et al., "The IceCube Neutrino Observatory: Instrumentation and Online Systems"

The 25 cm PMTs are the photon detection unit of the DOMs. A sketch of a PMTs operating principle is shown in Figure 3.10. When an incoming photon hits the photocathode, it can release a photoelectron. Through an array of dynodes, further electrons are released and accelerated, causing an amplified cascade of electrons that produce a measurable current at the PMTs anode. To release electrons through the photoelectric effect, the initial photons need a wavelength smaller than $\sim 700\,\mathrm{nm}$. For smaller wavelength, the sensitivity increases until at around $\sim 300\,\mathrm{nm}$ the absorption of the surrounding PMT glass starts dominating. In the glass vessel, the PMTs are facing downwards, being optimized for direct light from events from the Northern Hemisphere, i.e., declinations larger than $\delta > 0^\circ$. They are shielded from magnetic fields that would influence the electrons' trajectories through a mu-metal grid.[29]

[26] M. Aartsen et al., "The IceCube Neutrino Observatory: Instrumentation and Online Systems".

[27] Ibid.

[28] Ibid.

[29] M. Aartsen et al., "The IceCube Neutrino Observatory: Instrumentation and Online Systems".

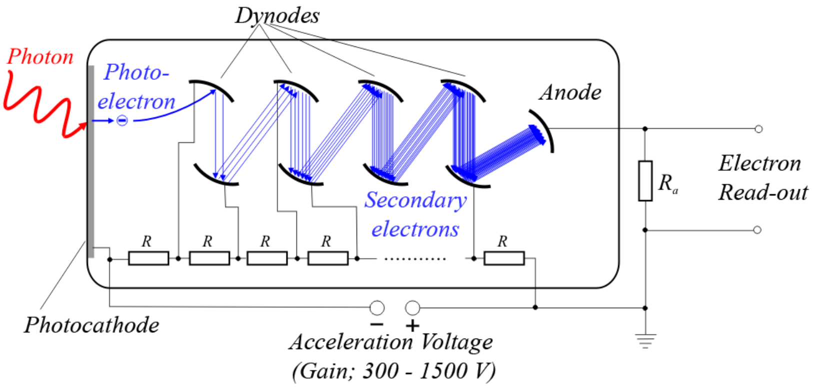

Figure 3.10: Scheme of a photomultiplier (PMT). An incident photon (left) releases electrons from the dynodes which produce a cascade of secondary electrons (blue). The electrons are accelerated through an acceleration voltage and measured on the right side of the PMT. Adapted from Wikimedia Commons, *Scheme of a Photomultiplier Tube (PMT)*, https://upload.wikimedia.org/wikipedia/commons/5/55/Photomultiplier_schema_en.png; accessed December 12th 2020.

For calibration purposes, each DOM holds a *flasher board* that consist of 12 LEDs with an output wavelength of $\sim 400\,\mathrm{nm}$. The LEDs are arranged in six evenly distributed pairs. One LED for each of these pairs faces ∼ 10 degrees downwards, which, after refraction through the DOM glass results in a horizontal emission. The other LED is tilted slightly upward at an angle of ∼ 50 degrees, close to the Cherenkov angle in ice. Calibration task that make use of the flasher board are[30]

- Timing response of the DOMs throughout the analysis software chain,
- Measuring the position of the DOMs in the ice,
- Measuring the optical properties of the ice.

3.5 IceCube Data Acquisition System

Cherenkov photons hitting the photocathode of a PMT release a cascade of electrons that get further amplified and collected, resulting in a waveform describing a current over time. Once the waveform passes a discriminator threshold of 0.25 photoelectrons, the DOM is *launched*, i.e., starts the recording and digitization of the waveform. The waveform digitization is thereby based on custom integrated circuits (ATWDs) and continuously sampling fast ADC (fADCs). While the ATWDs are optimized to digitize the waveform at a quick sampling rate of $3\,\mathrm{ns}$ ($427\,\mathrm{ns}$ in total) fast enough for light produced in the close vicinity of the DOM, the fADCs samples at a rate of $25\,\mathrm{ns}$ ($6.4\,\mu\mathrm{s}$ in total). To reduce dead time, two sets of ATWD chips are used in turns in each PMT.[31]

Even without an incoming event, the detector measures a few launches due to dark noise. These unwanted artifacts are suppressed by only sending the collected waveforms if the next or next-to-next DOM is also launched within a time window of $\pm 1\,\mu\mathrm{s}$. This requirement is called *hard local coincidence* (HLC). To identify incoming events, the surface computers continuously evaluate all the incoming HLC hits for correlations. If the signature of a likely event is

---

[30] Ibid.

[31] Ibid.

observed, the detector is *triggered* and all the photon hits recorded in a time window of $\Delta t \in [-4\,\mu\mathrm{s}, 6\,\mu\mathrm{s}]$ are combined into an *event*. IceCube's major trigger is the *simple multiplicity trigger*, which requires at least N HLC hits within a sliding time window. Specifically, for the InIce-Array 8 HLC hits are required within $5\,\mu\mathrm{s}$ (SMT8) and for Deep Core 3 HLC hits within $2.5\,\mu\mathrm{s}$ (SMT3). If trigger conditions are met multiple times, the time window is extended to account for all triggers. IceCube's SMT8 trigger rate is around 2.8 kHz, which is mainly caused by atmospheric muons penetrating the detector from above. Seasonal variations of this rate are caused by changes in the atmosphere's pressure, which affects the interaction rate of pions and kaons in atmospheric showers.[32] An overview of the event rates at the trigger level is shown in Figure 3.11.

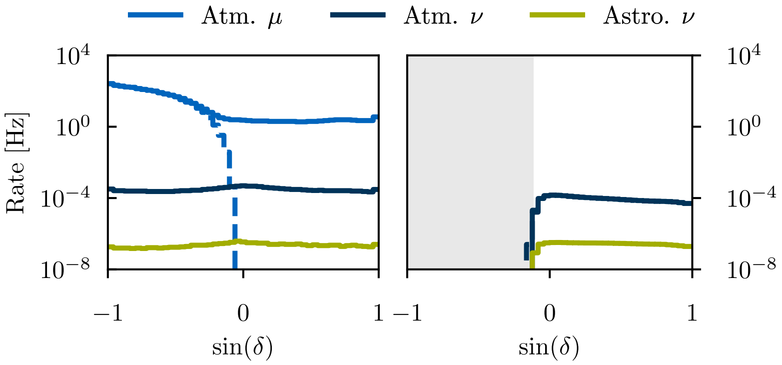

Figure 3.11: Comparison of IceCube event rates at trigger level (left) to a final level sample of well-reconstructable up-going muon tracks (right). Different contributions are shown in different colors, with solid lines indicating reconstructed quantities and dashed line the corresponding truth. The final level sample is optimized for tracks from the Northern Hemisphere, in order to remove the large contamination of atmospheric muons.

Events passing the SMT8 trigger are transferred to the Processing and Filtering (PnF) system. During the PnF processing, events are calibrated, reconstructed, and filtered to provide a first selection of potentially interesting physics events. The pre-selection is specifically important as the Iridium Satellite bandwidth at the South Pole is limited to only around $\sim$ 15% of the total events. In the first step, the waveforms are calibrated using DOM-specific calibration data, which provide information about the digitizer response function. Subsequently, the waveforms are unfolded into a sequence of discrete charges and times, where each tuple of charge and time is called a *pulse*. An example of a waveform unfolding is shown in Figure 3.12. The timing resolution achieved through this procedure is around 1 ns for ATWD recordings and 8 ns if only the fADC waveform is available. The sequence of pulses on a single DOM is usually called a *pulsemap* and is the basis for any IceCube reconstruction.[33]

Since South Pole computing facilities are limited, only inexpensive reconstruction algorithms are applied at the pole. While they are less accurate, they can still be used to filter potentially interesting events. In the following, we discuss three of the major filters that are relevant for this thesis:

[32] M. Aartsen et al., "The IceCube Neutrino Observatory: Instrumentation and Online Systems".

[33] Ibid.

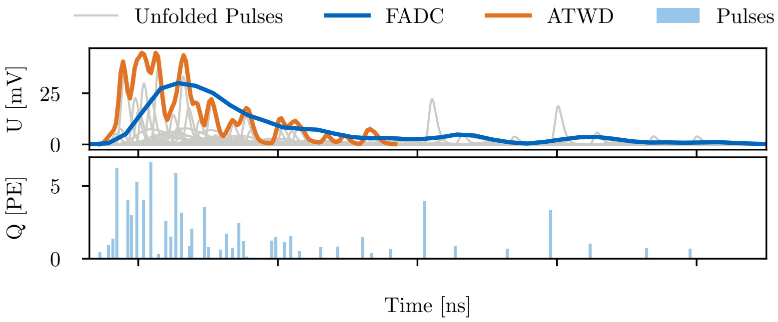

Figure 3.12: An exemplary transformation from recorded waveforms (top) to pulses (bottom). The ATWD and FADC waveform information is unfolded into pulses using pulse templates (thin, grey lines in the top panel). The resulting pulses are a set of energy and time tuples, modelling the photon detection.

- The **muon filter** aims to select events that have a track-like signature and could be produced through neutrino interactions. Filter cuts are applied on the direction and fit-quality of a simple directional reconstruction, *Improved Linefit*, and the total charge collected. In the up-going region (declination $\delta > -2.5°$), the atmospheric muon background is suppressed through the shielding of the Earth. Hence, we only need to apply an additional reconstruction quality cut to select clear tracks and reduce the rate of misreconstructed atmospheric muons. In the down-going region, the flux at the detector is largely dominated by atmospheric muons. Thus, a charge and zenith dependent cut is applied. Charge, as a simple estimator for an event's energy, can be used to suppress the steeply falling spectrum of background muons. The resulting total event rate of the muon filter is around 34 Hz. For events that pass the muon filter, a fast version of the Spline MPE reconstruction (section 3.8.1) is used to identify well reconstructable muon tracks. This **OnlineL2 filter** reduces the rate of muons down to 6 Hz. Events that pass the OnlineL2 filter are the basis for the realtime selection of through-going tracks.

- The **cascade filter** is designed to select events with cascade-like topology. Like the muon filter above, the sky is separated into up- and down-going, depending on the Linefit reconstruction direction. For all events in the up-going region (declination $\delta > -2.5°$), a single cut on the cascade likelihood is placed. For the down-going region, additional cuts on the reconstructed velocity of Linefit and the tensor of inertia of the deposited charge are used to minimize track-like contributions from the muon background. The resulting event rate is approximately 30 Hz.

- The **HESE filter** selects highly energetic starting tracks with a deposited charge of more than 6000 photoelectrons(p.e.). The starting criterion is verified by requiring that less than 3 p.e. of charge are deposited in a veto region in the first 250 ns of the event. Finally, an atmospheric self-veto[34,35] is used to discard events that are accompanied by an atmospheric muon. The HESE Filter has a very low rate of only a few events per year and is of major importance for the realtime pipeline, see chapter 6.

[34] Schonert et al., "Vetoing atmospheric neutrinos in a high energy neutrino telescope".

[35] T. K. Gaisser, Jero, et al., "Generalized self-veto probability for atmospheric neutrinos".

While filters are the first step to the data selection, they have limited performance and accuracy. Instead of making a clean selection, their major goal is to not miss any interesting events. In fact, there is a large overlap between the muon and the cascade filter, although the underlying topologies are fundamentally different. For the future, deep learning-based classification and reconstruction algorithms — as presented in chapter 5 — can significantly improvement the selection efficiency and purity, as they combine high speed with the performance of the best likelihood reconstructions. This could be especially interesting for future detectors — as IceCube-Gen2 — that significantly increase the data rate on trigger level.[36]

Events that pass any of the filters are transferred North to the IceCube computing facilities for further processing. In the first step, more sophisticated reconstruction algorithms are applied. This stage of the processing is referred to as *Level 2*. Starting from there, the event stream is split up into two branches, the Muon Level 3 and the Cascade Level 3 branch, consisting of well-reconstructed events of the respective topology. At this stage, the total event rate is lowered to around 2 Hz. The Level 3 branches are the starting point for the selection of events for final level analysis.

3.6 Properties of the Antarctic Ice

The accurate modeling of the ice properties in and around IceCube is an important factor for the calculation of the detector response to photons produced by traversing particles. In general, the ice at the South Pole has been formed gradually in layers through the accumulation and compression of the annual snow. The age of the ice in the detector volume is between ≈ 40 kyr (at 1500 m depth) and ≈ 100 kyr (at 2500 m depth).[37] Due to the high pressure at this depth, the ice is nearly free of air bubbles and thereby extremely transparent. The absorption length lies between 50 m to 350 m and is therefore much larger than in water.[38] While absorption decreases the number of photons detected, *scattering* of the photons on dust particles in the ice changes their direction and arrival times. With the scattering coefficient being a factor of 2 - 2.5 larger than the absorption coefficient, scattering is the major factor influencing the light propagation, see Figure 3.13. In fact, with an horizontal inter-string spacing of around 125 meters, most of the photons are expected to scatter several times before being detected.

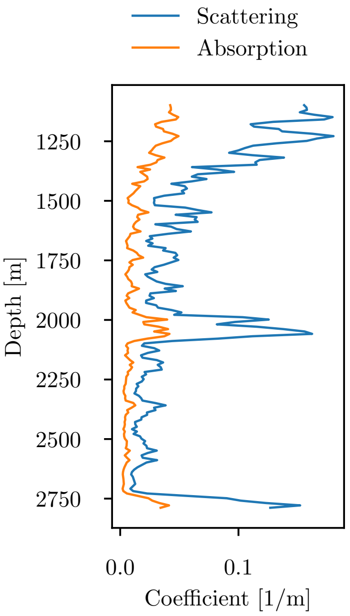

Figure 3.13: Depth-dependent scattering and absorption coefficient of photons in the deep Antarctic ice. Evidently, scattering effect are dominating over absorption. Three basic layers are visible: clear ice on the bottom, a dust layer in a depth of ∼ 2000 m, and clear ice up to the top of IceCube at 1450 m. Data from the Spice 3.2.1 ice model.

To account for the changes of the optical properties with depth, the ice is modeled in 100 layers with a thickness of 10 m each.[39] In consistency with topological studies of the bedrock, the single layers are tilted up to 70 m over one kilometer. The absorption and scattering parameters are then calibrated independently for each of the resulting layers using the on-board LED flashers of the DOMs.[40] In Figure 3.13 the measured absorption and scattering length are shown for the latest ice model *Spice 3.2.1*. It can be seen that the detector is effectively separated into three parts: Ice of good quality at the top, a

[36] M. Aartsen et al., "IceCube-Gen2: The Window to the Extreme Universe".

[37] "South Pole glacial climate reconstruction from multi-borehole laser particulate stratigraphy".

[38] Agostini et al., "The Pacific Ocean Neutrino Experiment".

[39] M. G. Aartsen et al., "Measurement of South Pole ice transparency with the IceCube LED calibration system".

[40] "South Pole glacial climate reconstruction from multi-borehole laser particulate stratigraphy".

dust layer with lower transparency and high scattering in the middle, and ultra-pure ice on the bottom. Finally, comparisons of the model predictions and flasher data show azimuthal intensity modulations. While the origin of this anisotropy is still under investigation, it can be corrected for in the ice model[41].[42]

3.7 *Energy Reconstructions*

The reconstructed energy is an important observable used in many IceCube analyses to distinguish between atmospheric and astrophysical events and characterize the underlying spectra. Different energy reconstructions are used depending on the respective event topologies. While for cascade-like topologies usually the total deposited energy in the detector is measured, through-going events are better reconstructed by their energy on entry as it is independent of the track length in the detector. Although other algorithms exist, this section only reviews one algorithm called *truncated energy*,[43] which is used in many high-level analyses searching for astrophysical neutrinos with through-going muon tracks. In section 7.2 we will show, however, that the energy resolution can be significantly improved through a deep learning-based energy reconstruction algorithm.

3.7.1 *Truncated Energy*

Muon energy losses are the composition of ionization losses and radiative losses. While the former are continuous and dominate below $\sim 1\,\mathrm{TeV}$, radiative losses produce a stochastic sequence of cascades along the track. Figure 3.14 shows a summary of the relevant loss contributions at different energies.

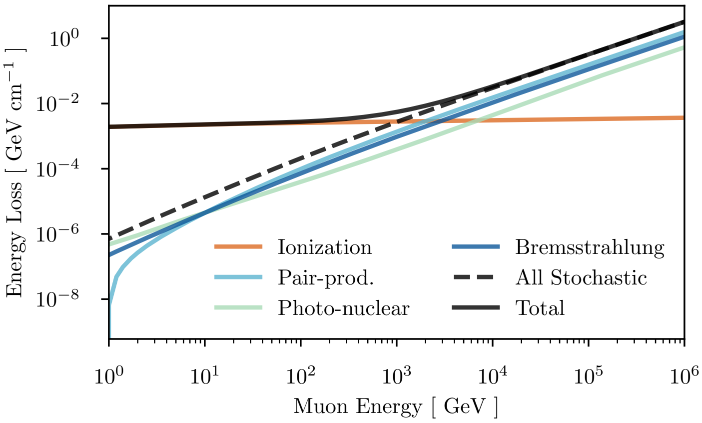

Figure 3.14: The energy losses of muons in ice. The total loss (black, solid line) is split up into the different contributions (colored lines as in the legend). Stochastic losses start dominating above $\sim 2\,\mathrm{TeV}$ of muon energy. Data from Tanabashi et al., "Review of Particle Physics".

Analytically, the total average loss per propagation length $\mathrm{d}x$ can be written as

$$-<\frac{\mathrm{d}E}{\mathrm{d}x}>=a+bE \tag{3.11}$$

with constants a and b.[44] Eq. (3.11) thereby provides a direct connec-

---

[41] Rongen, "Calibration of the IceCube Neutrino Observatory".

[42] Chirkin and Rongen, "Light diffusion in birefringent polycrystals and the IceCube ice anisotropy".

[43] Abbasi et al., "An improved method for measuring muon energy using the truncated mean of dE/dx".

[44] Tanabashi et al., "Review of Particle Physics".

tion between the muon energy and the observed losses. In contrast to the sub-TeV range where energy losses are continuous, higher energies are intrinsically more challenging to reconstruct, as the stochasticity of the radiative losses can produce varying energy loss patterns for the same muon energy. This holds specifically true in the transition region between continuous and stochastic losses as further discussed in section 7.2. In order to be stable against stochastic effects a truncation strategy is employed, which excludes the largest energy losses to be insensitive against strong stochastic variations.[45] Using a reconstructed track direction, truncated energy uses the expected light yield (total charge) to calculate the $\mathrm{d}E/\mathrm{d}x$ at each single DOM. The calculation is based on tabulated photon data for a muon with constant energy loss of $1\,\mathrm{GeV\,m^{-1}}$. Since the relation between the photon yield and the energy loss is linear, this can be used to infer the energy loss at a single DOM via

$$\left.\frac{\mathrm{d}E}{\mathrm{d}x}\right|_{DOM} = \frac{q_{obs}}{\hat{q}_{1GeV}} 1\,\mathrm{GeV\,m^{-1}} \tag{3.12}$$

with observed and expected charge, q_{obs} and $\hat{q}_{1GeV}$, respectively. Finally, the average loss $< \mathrm{d}E/\mathrm{d}x >$ is determined as the mean of the smallest 50% energy losses per DOM. Using eq. (3.11) this can be converted into an energy on entry assuming that the muon energy is approximately constant along the track [46].

3.8 Directional Reconstructions

The reconstructed event direction is an important observable for any IceCube analysis because fluxes are inherently anisotropic due to several reasons such as atmospheric effects, the propagation through the Earth, and the detector response. For astrophysical neutrino point source searches, specifically, the direction is the most important observable to identify spatial clustering of events above the background. Depending on the event topology and the requirements towards speed and accuracy, different algorithms are used. In this section we discuss IceCube's most important reconstruction methods for muon tracks - *Spline MPE* and *Millipede*. The underlying likelihoods are called single photoelectron (SPE) likelihood and multiphotoelectron (MPE) likelihood, depending of the treatment of the measured photon pulses.

3.8.1 Spline MPE

Spline MPE reconstructs track-like events based on the hypothesis of a straight line moving through the detector at the speed of light with continuous energy losses. Stochastic losses, as shown in Figure 3.5, are neglected. While this leads to a mismodeling of the photon arrival times, the impact on the reconstruction quality remains minor, as it is mainly dominated by the long lever arm of the track. An

[45] Abbasi et al., "An improved method for measuring muon energy using the truncated mean of dE/dx".

[46] which is valid for tracks with energy $E_\mu >\sim 1\,\mathrm{TeV}$, i.e. track length longer than the extension of the detector.

improved method that accounts for stochastic effects is called *Segmented SplineMPE*. While it improves the angular resolution slightly at muon energies above 100 TeV, it is still in the test phase and computationally much more expensive.[47]

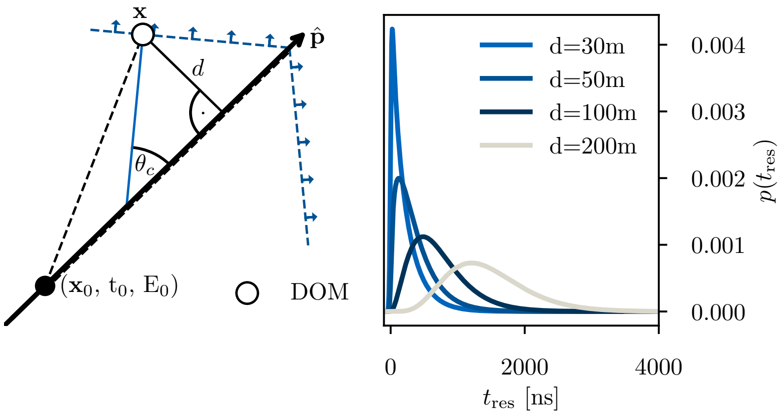

Figure 3.15: Geometry and timing properties of IceCube's directional reconstruction. The left panel shows the geometry of the photons observed from a track that is passing a DOM at position **x**. The right panels shows the residual time of the photons arriving at the DOM as a function of the distance d assuming the Pandel function. See text for more details. Courtesy of M. Huber.

The hypothesis of Spline MPE incorporates 5 parameters, two for the direction $\mathbf{v} = (\phi, \theta)$ and three for the position $\mathbf{r}_0 = (x, y, z)$ at arbitrary time t_0. ϕ and θ are thereby the azimuth and zenith angle on a celestial sphere. Figure 3.15 illustrates the geometry of a muon track passing a single DOM. Based on this, the expected geometrical arrival time t_{geo} of an unscattered Cherenkov photon is

$$t_{geo} = t_0 + \frac{\hat{\mathbf{p}} \cdot (\mathbf{x} - \mathbf{x}_0) + d \tan \theta_c}{c_0}. \tag{3.13}$$

with the DOM position $\mathbf{x}$, the perpendicular distance between track and DOM d, and the vacuum speed of light c_0. The difference between the arrival time of an observed photon t_{obs} and the geometrical time t_{geo} is called the residual time $t_{res} = t_{obs} - t_{geo}$. Over the years different methods have been developed in order to model the time residuals distribution, which is the sum of several effects: the scattering of the photons during propagation, the detection effects of the PMTs, and the dark noise of the detector. An approximate, analytical distribution is given through the *Pandel* function[48]

$$p(t_{res}) \propto \frac{\tau^{-d/\lambda} \cdot t_{res}^{d/\lambda - 1}}{\Gamma(d/\lambda)} \times e^{-t_{res}(1/\tau + \hat{c}/\lambda_a) - d/\lambda_a} \tag{3.14}$$

with $\hat{c} = c_0/n$, absorption length λ_a, and two free parameters λ and τ, which are estimated from Monte Carlo simulation. In Figure 3.15 the residual time $p_{res}(t)$ is illustrated for different distances between DOM and track. As expected, the time residual distribution shifts and broadens for larger DOM distances. While the Pandel function is a decent approximation and fast to evaluate, it does not account for the variations of the ice properties throughout the detector, as explained in section 3.6. As a more accurate alternative, smooth Monte Carlo-based spline tables can be used to calculate the

[47] Bradascio and Glüsenkamp, "Improving the muon track reconstruction of IceCube and IceCube-Gen2".

[48] J. Ahrens et al., "Muon track reconstruction and data selection techniques in AMANDA".

expected photon arrival time distributions. The generation of the tables is based on a full photon simulation chain for a large number of light sources placed at different positions and orientations in the detector volume. The results are stored in high-dimensional penalized B-spline,[49] which is reasonably fast in its evaluation. Tables are generated for different event hypotheses as infinite through-going muons and single electromagnetic cascades. Note that, although the spline table approach improves significantly over the analytical approach in eq. (3.14), it still assumes lateral translational invariance and therefore does not account for effects as the ice anisotropy and tilt. Methods to overcome this limitation are currently developed using Generative Adversarial Networks.[50]

Based on the distribution of photon arrival times the single photoelectron (SPE) likelihood is defined as

$$\mathcal{L}_{SPEall} = \prod_{i}^{N_{DOMs}} \prod_{j}^{N_{pulses,\,i}} p(t^{j}_{res,\,i}\,|\,\mathbf{x}_i, H). \tag{3.15}$$

The product runs over all detected pulses with time residual t^{j}_{res} on a DOM with distance $\mathbf{x}_i$ to the track. While this parameterization of the likelihood includes the maximum of information, it has been shown to be unstable, since effects such as afterpulses are not modeled accurately in the arrival distributions. In fact, the reconstruction quality improves when only using the arrival time of the first photon, which is the one that is least affected by scattering and thereby carries most of the information. As a trade-off between *only* using the arrival time of the first photon and using *all* arrival times, the MPE likelihood is constructed via

$$\mathcal{L}_{MPE} = \prod_{i}^{N_{DOMs}} p(t^{1}_{res,i}\,|\,\mathbf{x}_i, H) \left(\int_{t^{1}_{res,i}}^{\infty} p(t\,|\,\mathbf{x}_i, H)\mathrm{d}t \right)^{N_{pulses,i}-1} . \tag{3.16}$$

In this construction the arrival time of the first photon provides the dominant timing information, while also being set into relation to the probability of observing $N^{i}_{pulses} - 1$ hits at a later time. The best-fit direction is found by minimizing the negative logarithm of eq. (3.16). As the likelihood spaces are quite complex with potential local minima, Spline MPE is usually seeded with a different angular reconstruction method that has been run earlier in the data processing pipeline. In Figure 3.16 the performance of Spline MPE is shown in comparison to other algorithms for the track sample used in the point source analysis presented in chapter 7.

[49] Whitehorn, Santen, and Lafebre, "Penalized Splines for Smooth Representation of High-dimensional Monte Carlo Datasets".

[50] Huennefeld, "Reconstruction Techniques in IceCube using Convolutional and Generative Neural Networks".

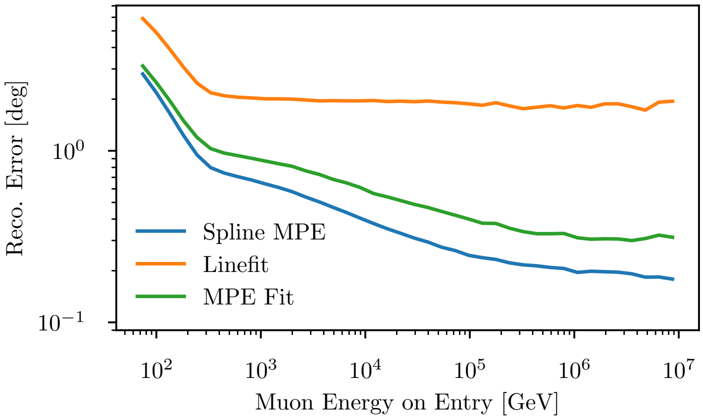

Figure 3.16: Comparison of the angular resolution of different IceCube track reconstructions. While simple approximated algorithms like Linefit only reconstruct within a few degrees, the Spline MPE reconstruction reaches around 0.2° at 100 TeV.

3.8.2 Estimating the Angular Uncertainty - Paraboloid

The estimation of the directional uncertainty is of major importance to calculate the compatibility of an event with a specific location in the sky. Thus, after the determination of the best-fit direction by maximizing eq. (3.16), the likelihood region around the minimum is analyzed, as shown exemplarily in Figure 3.17.

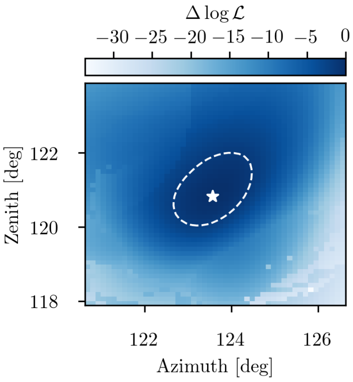

Figure 3.17: An exemplary Spline MPE likelihood scan. The white star represents the best-fit direction and the white, dashed contour the 95% confidence region.

For the large sample limit, the confidence intervals of a likelihood estimator can be approximated through the likelihood function with[51]

$$\Delta \log \mathcal{L} = \log \mathcal{L}(\hat{\theta}^{+\sigma}_{-\sigma}) - \log \mathcal{L}_{max} = -\frac{N^2}{2}, \tag{3.17}$$

where N corresponds to the desired confidence level quantile of the Gaussian distribution (N=1 for a 68.3% confidence interval). Hence the one-sigma error region for a parameter σ is defined by the contour at which the $\Delta \log \mathcal{L}$ is equal to -1/2. In IceCube this method is implemented in a package called *paraboloid*,[52] which scans the directional likelihood landscape at 24 distinct points in three rings around the minimum. For each direction, the vertex position and time are re-optimized. Finally, a 2-dimensional parabola is used to interpolate between the points and to calculate the confidence region σ of the direction.

Conveniently, the estimated *confidence region* σ can be related to the *variance* of directional reconstruction assuming the Cramer-Rao bound that defines a lower limit for the variance of an estimator $\hat{\theta}$ via

$$var(\hat{\theta}) \geq \frac{1}{I(\theta)}. \tag{3.18}$$

$I(\theta)$ is the *Fischer Information*, which for a twice differentiable function can be written as

$$I(\theta) = -E\left[\left.\frac{\partial^2 \mathcal{L}(x|\theta)}{\partial \theta^2}\right|_{\hat{\theta}}\right], \tag{3.19}$$

with E indicating the expectation value and $\hat{\theta}$ the maximum likelihood estimate. Visually, eq. (3.19) represent the expectation value of the likelihood's curvature around the maximum. In many cases

[51] Cowan et al., "Asymptotic formulae for likelihood-based tests of new physics".

[52] Neunhoffer, "Estimating the angular resolution of tracks in neutrino telescopes based on a likelihood analysis".

it is, however, unfeasible to calculate the expectation, as it requires re-sampling of the same true hypothesis. Fortunately, the problem simplifies significantly for the case of photon statistics going to infinity. In this case, the Cramer-Rao bound in equation (3.18) becomes exact and the expectation value in eq. (3.19) can be estimated through a single measurement. Hence the variance is given by

$$var(\hat{\theta}) = \frac{1}{-\frac{\partial^2 \mathcal{L}(x|\theta)}{\partial \theta^2}}. \tag{3.20}$$

In order to determine the right side of this equation we return to the estimation of the confidence region σ, for which we had assumed that the likelihood function can be locally described by a parabola [53]

$$\mathcal{L}(x|\theta) = -a \cdot (\theta - \hat{\theta})^2 + \mathcal{L}_{max}, \ a > 0 \tag{3.21}$$

with second derivative

$$-\frac{\partial^2 \mathcal{L}(x|\theta)}{\partial \theta^2} = 2a. \tag{3.22}$$

By comparing eq. (3.20) and eq. (3.22) we find that the variance can thus be simplified to

$$var(\hat{\theta}) = \frac{1}{2a}. \tag{3.23}$$

On the other hand, we also used the parameterization in eq. (3.22) and the Cramer-Rao bound in eq. (3.17) to construct σ in a way that

$$\mathcal{L}(x|\hat{\theta}^{+\sigma}_{-\sigma}) - \mathcal{L}_{max} = -\frac{1}{2} = -a \cdot \sigma^2 \tag{3.24}$$

and hence

$$\frac{1}{\sigma^2} = 2a. \tag{3.25}$$

Plugging this into equation (3.23) the result reads

$$var(\hat{\theta}) = \sigma^2.$$

The size of the confidence region can therefore be used as a measure for the variance of the estimator. In a final step the normality of maximum likelihood estimators, eq. (7.5), guarantees that the deviation of the estimated parameters $\hat{\theta}$ from the truth θ_0 can be modeled by a normal distribution N, i.e.

$$\sqrt{n}(\hat{\theta} - \theta_0) \rightarrow N\left(0, \frac{1}{I(\theta)} = \sigma^2\right) \tag{3.26}$$

leading to the Gaussian approximation of the spatial term in the current point source likelihood, eq. (7.26). We will see throughout chapter 7 that this approximation is not valid over the entire parameter space of the analysis. While the discrepancy is mainly caused by the kinematic angle between neutrino and muon, in some cases the *paraboloid* minimization fails and thereby produces inaccurate results.

[53] which can always be reached through Taylor expansion around the minimum

3.8.3 Millipede Reconstruction

In section 3.8.1 Spline MPE has been introduced as an angular reconstruction algorithm for through-going muon tracks. While it works well for a large range of events, it becomes unreliable for events that deviate from the assumption of an infinite track with continuous losses, e.g., starting tracks and through-going tracks with large stochasticity. Those two topologies play, however, a crucial role in the searches for (realtime) counterparts of the most promising astrophysical event candidates, as discussed in chapter 4 and chapter 8. An alternative reconstruction algorithm that is stable for those topologies and furthermore considers the timing of *all* the photons is called *Millipede*. The reconstruction principle is illustrated in Figure 3.18. In contrast to Spline MPE, the general strategy of Millipede is the reconstruction of single energy losses in small segments along the track. The hypothesis H is therefore defined by a vector of energy losses $\mathbf{E} = (E_1, \dots, E_N)^T$ along a track with given direction. The expected charge at DOM i is then given by the matrix product

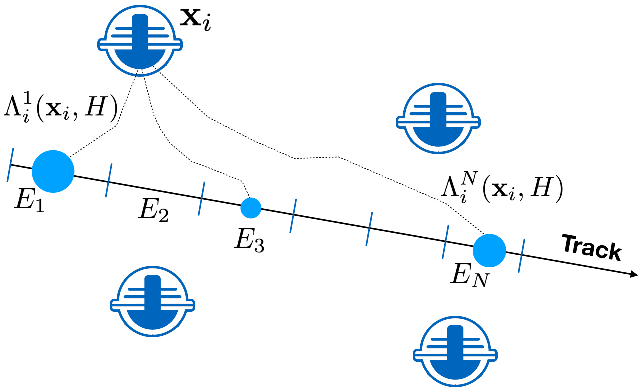

Figure 3.18: Sketch of the Millipede reconstruction. The track is segmented into chunks with energy losses $E_1, \dots, E_N$. The expected charge at a DOM with coordinates $\mathbf{x}_i$ is given as the superposition of the expectations from the single losses.

$$\hat{n}_i = \rho_i + \Lambda_i(\mathbf{x}_i, H) \cdot \mathbf{E}, \tag{3.27}$$

where $\Lambda_i(\mathbf{x}_i, H)$ is the i's row of a response matrix describing the expected hits at the DOM position $\mathbf{x}_i$ for N 1 GeV cascades along a given track. In addition, ρ_i adds a model for the DOM-specific noise. Using the Poisson likelihood of observing k hits for an expectation value of $\hat{n}_i$,

$$\mathcal{L}_{poisson} = \frac{\hat{n}_i}{k!} e^{-\hat{n}_i}, \tag{3.28}$$

the full likelihood is given by

$$\log \mathcal{L} = \sum_i^N (k_i \log(\Lambda_i \cdot \mathbf{E} + \rho_i) - (\Lambda_i \mathbf{E} + \rho_i) - \log k_i!). \tag{3.29}$$

In order to find the best-fit direction (and deposited energy) a grid scan in azimuth and zenith direction is performed. For each direction the vector of energy losses is adjusted and the best-fit loss pattern determined by minimizing the millipede likelihood in eq. (3.29).

Unsurprisingly, this procedure is computationally extremely expensive and can only be performed on a small number of interesting events.

3.8.4 Event Resimulation and Error Contours

As a result, the millipede *scan* described in section 3.8.3 produces a fine grid of likelihood values. Despite the identification of the best-fit direction, the scan can also be used to calculate an error contour. This is especially relevant for multi-frequency follow-up searches of the IceCube high-energy events. As telescope time is limited, an accurate determination of the error contour helps to efficiently identify promising candidates and discard background sources outside the contour. Unfortunately, however, the contours derived assuming a simple χ^2 distribution of the $\Delta \log \mathcal{L}$ between best-fit and true direction [54] do not show a good coverage. This has been confirmed with a small sample of events, which have been resimulated with events of similar energy, direction, and varying ice-systematics.[55] For some extraordinary experimental events, a dedicated resimulation effort has been performed to derive the event-specific $\Delta \log \mathcal{L}$ distribution. For example, this is the case for the high-energy track leading to the identification of the first compelling neutrino source TXS 0506+056 as discussed in chapter 4. In fact, the resimulation of those few events provide the underlying test-statistic distribution that is used to derive the error contours for the current IceCube realtime stream. While this is expected to have better coverage than the simple χ^2 assumption, the validation and characterization of this technique on a larger sample is still pending.

[54] As expected from Wilk's theorem

[55] Rädel, "Measurement of High-Energy Muon Neutrinos with the IceCube Neutrino Observatory".

4 TXS 0506+056 - The First Source of Astrophysical Neutrinos

On the 22nd of September 2017, the IceCube Neutrino Observatory detected a highly-energetic neutrino — IceCube-170922A. From the energy estimation, the event has a high probability of being astrophysical. Furthermore, a bright gamma-ray blazar — TXS 0506+056 — was found close to the best-fit direction of the event.[1] Follow-up observations revealed that the source was undergoing a strong gamma-ray flare, and searches in archival neutrino data found an additional accumulation of neutrinos from a direction consistent with the source.[2] Based on a multi-wavelength study of the region, we find that TXS 0506+056 is indeed the only energetically convincing counterpart for this emission.[3] This chapter will review and connect the different components that made TXS 0506+056 the first compelling, non-stellar source of astrophysical neutrinos. Discussions on the multi-frequency dissection follow closely our previously published results.[4]

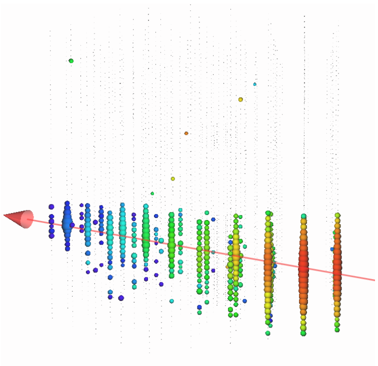

Figure 4.1: Event view of IceCube-170922A. The size of the bubbles is proportional to the charge and the timing is encoded in the color (red to blue).

4.1 Detection of the Neutrino Event IceCube-170922A

Historically, IceCube's realtime pipeline consists of two event streams, the EHE stream for extremely-high-energy through-going tracks and the HESE event stream for high-energy starting tracks. While the former covers the energy range between 500 TeV to 10 PeV with a rate of 4 to 6 events per year at a purity [5] of >50%, the latter has a rate of around 5 events per year and purity of 29%.[6]

IceCube-170922A was identified as a through-going track through the EHE alert stream at MJD 58018.87. An event view can be seen in Figure 4.1. After applying the millipede reconstruction the best-fit direction is given as right ascension (RA) $77.43^{+0.95}_{-0.65}$ and declination (Dec) $+5.72^{+0.50}_{-0.30}$ (deg, J2000, 90 per cent containment region)[7] and thereby only 0.06° away from the flaring gamma-ray source TXS 0506+056, see Figure 4.4 and 4.9. As this is a very rare coincidence, a re-simulation of the event was performed to determine an event-specific confidence contour as discussed in section 3.8.4. Being clearly up-going, i.e., coming from the Northern Hemisphere, an atmospheric muon origin can be excluded as they are shielded by the Earth. Using the reconstructed direction and energy, one can further estimate the energy of the initial neutrino. For this calculation, EHE

[1] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

[2] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Samarai, et al., "Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert".

[3] Padovani, Giommi, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source".

[4] Ibid.

[5] Fraction of signal events among the total event rate

[6] M. Aartsen et al., "The IceCube Realtime Alert System".

[7] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

Monte Carlo events with similar reconstructed energy and direction are selected, resulting in a distribution of true energies as seen in Figure 4.2.[8] The resulting energy distribution depends on the assumed astrophysical spectrum since the vertex, and thereby the length of the track before entering the detector, is unknown. The most-likely neutrino energy lies in a range between 259 TeV and 311 TeV for spectral indices between $\gamma = 2.0$ and $\gamma = 2.5$.[9] However, note that photo-hadronic models of neutrino production in blazars usually predict harder spectral indices with cut off. In the case of harder spectra, the estimated energy can be significantly higher. Finally, assuming an astrophysical spectrum with a power law index $\gamma = 2.13$, the signalness of the event, i.e., the fraction of signal events among the total event rate at this declination and energy is 56.5%.[10] Albeit this number is not particularly high, the discussion of the connection between IceCube-170922A and TXS 0506+056 in the next section will show that the chance coincidence probability for a *spatial* coincidence of the event with such a source is extremely small.

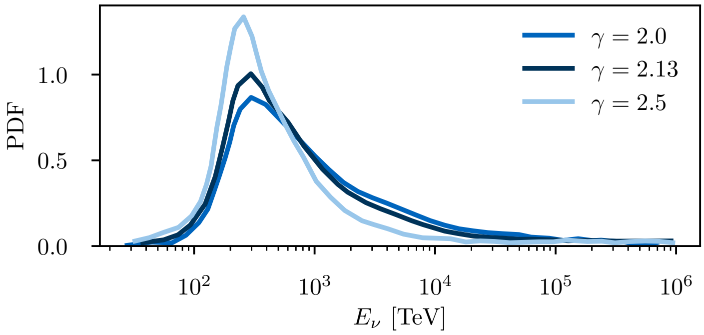

Figure 4.2: Distribution of the neutrino energies leading to the observed direction and muon energy of IceCube-170922A. As the position of the vertex and thereby the track length before reaching the detector is unknown, the distribution depends on the assumed neutrino spectrum. Plot from IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

4.2 *Chance Coincidence Probability*

As discussed in section 3.8, IceCube uses likelihood-based algorithms in order to reconstruct the event direction along with the uncertainty. In this *frequentist approach* the error contours are constructed in a way that they contain the true direction in 90% of the experimental outcomes [11]. The resulting contour does not, however, make any statement on how probable a certain direction $\mathbf{x}$ is the true direction, i.e., it does not make a statement on $P(\mathbf{x}|\text{data})$, which is a strictly Bayesian concept. Similarly to the point source analysis in chapter 7.1.3, we can, however, construct a frequentist test that compares the signal hypothesis of an event with direction $\mathbf{x}$ to originate from a nearby source with source flux ϕ_s to the null hypothesis of the event being of atmospheric or diffuse astrophysical origin. The corresponding signal likelihood can be written as

$$\mathcal{L}(H_1 \,|\, \mathbf{O}) = p(\mathbf{O}|H_1) = p(\mathbf{O}|\alpha_s, \omega) = p(\mathbf{x}|\alpha_s) \cdot p(\phi_s|\omega) \tag{4.1}$$

where the observable space $\mathbf{O}$ consist of the event direction $\mathbf{x}$ and the source flux ϕ_s. The hypothesis H_1 is the composition of a source position $\alpha_s = (\phi_s, \theta_s)$ and some flux weighting ω. The formalism

---

[8] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

[9] Ibid.

[10] Ibid.

[11] Assuming that the likelihood is correctly modeled and takes into account all the relevant systematics.

in eq. (4.1) can then simply be expanded to the case of a catalog of sources by summing the single contributions to the total likelihood

$$\mathcal{L}(H_1 \,|\, \mathbf{O}) = \sum_s p(\mathbf{x}|\alpha_s) \cdot w_s. \tag{4.2}$$

with the single source flux weight $w_s = p(\phi_s \,|\, \omega)$. The point spread function $p(\mathbf{x}|\alpha_s)$ can be approximated by a Gaussian

$$p(\mathbf{x}|\alpha_s) = \frac{1}{2\pi\sigma^2} \times \exp\left(-\frac{1}{2}\frac{|\mathbf{x}-\alpha_s|^2}{\sigma^2}\right) \tag{4.3}$$

with the angular error estimator σ being calculated on an event-to-event basis. In case of this follow-up analysis we use the Cramér-Rao lower bound, which becomes a reasonably good estimator for the variance of the directional reconstruction at high photon statistics, i.e., high energies.[12]

As the production mechanisms of astrophysical neutrinos are unknown, there is no obvious choice of the weighting term w_s in eq. (4.2). As both neutrinos and gamma rays, are, however, expected to be produced with a similar luminosity, it is natural to assume a simple model of the neutrino energy flux at Earth being linearly proportional to the gamma-ray energy flux, i.e., $S_\nu \propto S_\gamma$. Consequently, we define the neutrino and gamma-ray energy flux over the energy range between 1 GeV to 100 GeV as

$$S_\nu \propto S_\gamma = \int_{1\,\mathrm{GeV}}^{100\,\mathrm{GeV}} E^\gamma \frac{\mathrm{d}\Phi_\gamma}{\mathrm{d}E_\gamma} \mathrm{d}E_\gamma. \tag{4.4}$$

To calculate the expected neutrino flux at the detector, one needs to additionally account for absorption effects of neutrinos in the Earth as well as the detector efficiency. The final weight is therefore the normalized product of the time-dependent flux weight and zenith-dependent acceptance

$$w_s(t, \theta_s) = \frac{S_{\gamma,s}(t) \cdot w_{acc}(\theta_s)}{\sum_{\tilde{s},\tilde{t}} S_{\gamma,\tilde{s}}(\tilde{t}) \cdot w_{acc}(\theta_{\tilde{s}})} \tag{4.5}$$

with $\tilde{s}$ going over all the 2257 extragalactic *Fermi*-LAT 3FGL sources known at that time.[13] Note that in order to account for the flaring of the source, the analysis includes an additional time series parameter t which is based on a sample of 28-day 1 GeV gamma-ray light curves for all the sources.[14] The final signal likelihood is then written as

$$\mathcal{S} = \mathcal{L}(H_1 \,|\, \mathbf{O}) = \sum_s \frac{1}{2\pi\sigma^2} \times \exp\left(-\frac{1}{2}\frac{|\mathbf{x}-\alpha_s|^2}{\sigma^2}\right) \cdot w_s. \tag{4.6}$$

For the background hypothesis we assume a conventional atmospheric plus best-fit diffuse astrophysical spectrum leading to

$$\mathcal{B} = \mathcal{L}(H_0 \,|\, \mathbf{O}) = \frac{p(\theta|H_0)}{2\pi \cdot N_t \cdot N_s}, \tag{4.7}$$

[12] M. Aartsen et al., "Very High-Energy Gamma-Ray Follow-Up Program Using Neutrino Triggers from IceCube".

[13] Acero et al., "Fermi Large Area Telescope Third Source Catalog".

[14] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

assuming a uniform arrival direction in right ascension due to the rotation of the detector at the South Pole; $N_t = 154$ and $N_s = 2257$ are the number of time bins and extragalactic 3FGL sources, respectively. Finally, the test statistic is constructed over

$$\mathcal{TS} = 2 \times \log \frac{\mathcal{S}}{\mathcal{B}}. \tag{4.8}$$

In order to determine the background probability for the space and time coincidence of TXS 0506+056 and IceCube-170922A, a background test statistic distribution is derived using Monte Carlo-based pseudo-experiments. Comparing to the experimental test statistic, the resulting pre-trial p-value is $p_{local} = 2.1 \times 10^{-5}$ or 4.1σ in a one-sided test. In order to account for the rate of background events, we further correct for the 51 archival events that passed the realtime selection prior to the observation of IceCube-170922A. For independent trials, the binomial probability of observing at least one trial with $p_{local} \leq 2.1 \times 10^{-5}$ under the background hypothesis in N trials is

$$p_{global} = 1 - (1 - p_{local})^N. \tag{4.9}$$

leading to a post-trial significance of 3.0σ.

Note that *a posteriori* calculations — as the one presented here — should only be taken as a hint that something *might* be interesting since the result is susceptible to two biases: 1.) the *a posteriori* choice of the (signal) hypothesis and 2.) the *a posteriori* choice of the data analysis method. To illustrate this, one can think of any other reasonable method to calculate the chance coincidence probability of IceCube-170922A and TXS 0506+056. For example, one could, instead of assuming a Gaussian point spread function, uses the published millipede contours and conservatively assign the same spatial weight for every point inside the contour (which is done in similar publications[15,16]). Additionally, instead of taking the gamma-ray flux, one could also argue that TXS 0506+056 is very special in its radio flux (there are only 12 brighter radio sources in the 3HSP blazar catalog[17]). As a result the probability of seeing a source that is as bright as TXS 0506+056 in radio (gamma rays) and included in an error contour with a size of $\Omega_{\text{event}} = 0.97\,\text{sqdeg}$ (90% error contour of IceCube-170922A[18]) is

$$P = \frac{\Omega_{\text{event}}}{41252.96\,\text{sqdeg}} \times N(\Phi > \Phi_{\text{TXS 0506+056}}), \tag{4.10}$$

with the denominator giving the total solid angle of a sphere and assuming that the number of sources with a flux higher than TXS 0506+056 is small enough to avoid any higher-order binomial effects. The resulting chance coincidence probabilities are 0.028% based on the radio flux and 0.44% based on the gamma-ray flux (TXS 0506+056 is among the 189 brightest objects in the *Fermi* 4FGL catalog[19]). The entire issue about the *a posteriori* chance coincidence calculation is a strong motivation to move beyond single-event analysis and towards the multi-frequency correlation studies of larger event samples as shown in chapter 8.2.

[15] Stein et al., "A high-energy neutrino coincident with a tidal disruption event".

[16] Giommi, Padovani, Oikonomou, et al., "3HSP J095507.9+355101: a flaring extreme blazar coincident in space and time with IceCube-200107A".

[17] Y.-L. Chang, Arsioli, et al., "The 3HSP catalogue of extreme and high-synchrotron peaked blazars".

[18] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

[19] Abdollahi et al., "Fermi Large Area Telescope Fourth Source Catalog".

4.3 *Dissecting the Region around IceCube-170922A*

The previous section has shown that a coincidence between a high-energy neutrino and a gamma-ray flaring source, as bright as TXS 0506+056, is extremely rare. Despite this being interesting on its own, a search in archival IceCube data has revealed a flare of neutrino events from a location consistent with the position of TXS 0506+056, Figure 4.3. The flare duration for a Gaussian time window is estimated to 110^{+35}_{-24} days between MJD 56949 and 57059 with a post-trial significance of 3.5σ.[20] The corresponding ν_μ fluence at 100 TeV is $2.1^{+0.9}_{-0.7} \times 10^{-4}\,\mathrm{TeV\,cm^{-2}}$ for a best-fit power law spectral index of $\gamma = 2.1 \pm 0.2$ in a sensitive energy range (68 per cent) between 32 TeV and 3.6 PeV[21] [22]. To connect all the observations in the region around IceCube-170922A and to exclude other potential candidate sources despite TXS 0506+056, we have developed a pipeline that performs a *dissection* of the entire region of interest (ROI) from a multi-frequency, i.e., radio to gamma ray, perspective.

4.3.1 *Search for Multi-Frequency Counterparts*

In astrophysical source environments, neutrinos are usually expected to be created through the decay of charged pions produced in inelastic proton-proton and proton-gamma interactions. In the same processes, neutral pions are produced that further decay into two photons ($\pi^0 \to 2\gamma$). The energy of protons ϵ_p is related to the energy of the neutrinos (ϵ_ν) and gamma rays (ϵ_γ) as

$$\epsilon_\nu \approx 0.5\epsilon_\gamma \approx 0.05\epsilon_p. \tag{4.11}$$

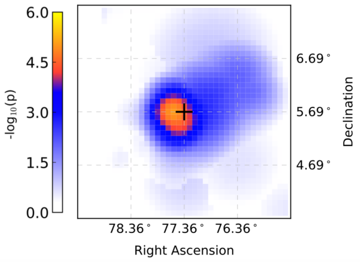

Figure 4.3: The local p-value map showing the time-integrated emission around TXS 0506+056. It can be seen that the emission is well localized and consistent with the position of TXS 0506+056. Plot taken from IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Samarai, et al., "Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert".

The production of astrophysical neutrinos is therefore accompanied by the production of gamma rays at similar energies and rates [23]. However, the detection of those gamma rays is not always guaranteed as they can be a) absorbed by the EBL or b) deflected in gamma ray opaque sources due to two-photon annihilation processes. Consequently, it is advantageous to extend searches for possible neutrino counterparts towards radio, optical and X-ray energies, where accelerators can be identified through their characteristical non-thermal synchrotron spectrum. The pipeline presented here is based on two components 1.) a search for blazar-like counterpart candidates using the VOU-Blazar tool[24] and 2.) a dedicated analysis of the available multi-frequency data, including an analysis of the *Fermi*-LAT data for the region and all the previously identified counterpart candidates. This work is the foundation of the automated follow-up pipeline presented in chapter 8.4.

In Figure 4.4 the radio and X-ray emission in the region around IceCube-170922A is shown as red dots and blue circles, respectively. The data is provided through the VOU-Blazar tool that queries 34 catalogs through the *conesearch pipeline* [25].[26] The 34 catalogs include basically all the publicly available radio and X-ray data, as well as additional type-specific catalogs of known sources to resolve ambiguities. In a second step, counterpart candidates are identified as

[20] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Samarai, et al., "Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert".

[21] Ibid.

[22] Note that the point source analysis method applied here has an intrinsic bias in the spectral index and flux normalization as shown in chapter 7.5. Although this effect is small for an E^{-2} spectrum close to the horizon, it could still impact the results by $\sim 10\%$.

[23] see section 2.2 for a detailed review

[24] Y.-L. Chang, Brandt, and Giommi, "The Open Universe VOU-Blazars tool".

[25] https://github.com/chbrandt/eada/blob/master/eada/vo/conesearch.py

[26] Y.-L. Chang, Brandt, and Giommi, "The Open Universe VOU-Blazars tool".

positions in the sky where there is both, radio *and* X-ray emission, at a flux ratio consistent with that of a blazar. A convenient definition for this ratio is given by[27]

$$\alpha_{rx} = -\frac{\log(f_{\nu_1}/f_{\nu_2})}{\log(\nu_1/\nu_2)} \tag{4.12}$$

where $\nu_1 = 1.4\,\text{GHz}$ and $\nu_2 = 1\,\text{keV}$, with respective flux values f_{ν_1} and f_{ν_2}. Different classes of blazars have different emission profiles and hence the α_{rx} can be used to classify the blazars via[28,29]

$$\text{Class} = \begin{cases} \text{HSP, for } 0.43 < \alpha_{rx} \leq 0.78\ \&\ \nu^S_{\text{peak}} \geq 10^{15}, \\ \text{ISP, for } 0.43 < \alpha_{rx} \leq 0.78\ \&\ \nu^S_{\text{peak}} < 10^{15}, \\ \text{LSP, for } 0.78 < \alpha_{rx} < 0.95, \\ \text{non-jetted AGN candidate, for } \alpha_{rx} < 0.43, \\ \text{Unknown, for } \alpha_{rx} > 0.95. \end{cases} \tag{4.13}$$

Considering the region around IceCube-170922A, this leads to the source candidates shown in the right panel of Figure 4.4. There are three known objects, TXS 0506+056 (an IBL/HBL at redshift z = 0.3365,[30] source no. 5, distance $\Delta\psi = 0.06°$), PKS 0502+049 (an LBL/FSRQ at z = 0.954, source no. 1, distance $\Delta\psi = 1.2°$), and 3HSP J050833.4+053109 [31] (source no.3, distance $\Delta\psi = 0.35°$), as well as four additional blazar candidates. The first two objects, TXS 0506+056 and PKS 0502+049, are bright gamma-ray emitters and therefore included in the 4FGL catalog, both being among the 3.5% brightest objects in photon flux. Visual inspection of the other sources's SEDs show that source no. 4 is a HSP candidate, source no. 7 is likely a cluster of galaxies (due to its extended X-ray emission), source 6 is a steep radio spectrum object, and source 2 is a nearby elliptical galaxy showing low luminosity X-ray emission that could be due to a jet or even to non-nuclear sources.[32]

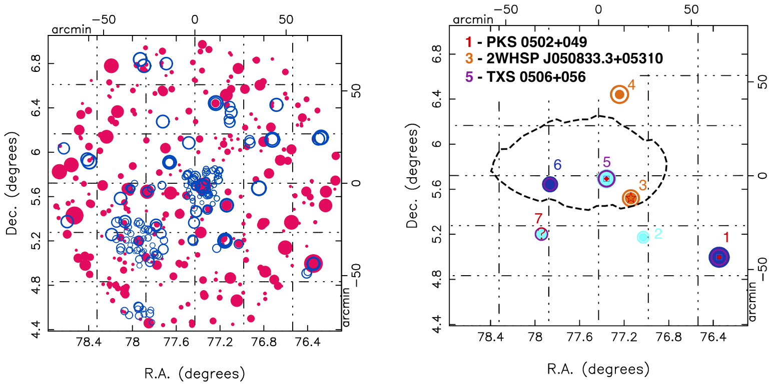

Figure 4.4: Result of the multi-frequency dissection of the region around IceCube-170922A using the OpenUniverse VOU-Blazar tool. The left plot shows all the known radio (red) and X-ray emission (blue). Cross-matching these emissions, only a few blazar-like counterpart candidates remain (right plot). Dark blue circles represent LBL, cyan symbols IBL, and orange symbols HBL candidates. Known blazars are marked by red diamonds if they are part of the BZCAT catalogue or by a star if they are part of the 2WHSP sample.

4.3.2 *Analysis of the Fermi-LAT Gamma-Ray Data*

In order to further address the emission spectra of the candidates, *Fermi*-LAT gamma-ray data collected between August 4, 2008 to February 10, 2018 are analysed [33]. In Figure 4.5 and Figure 4.6 the gamma-

[27] Padovani and Giommi, "The Connection between X-Ray– and Radio-selected BL Lacertae Objects".

[28] Padovani, Perlman, et al., "What types of jets does nature make: A New population of radio quasars".

[29] Padovani and Giommi, "The Connection between X-Ray– and Radio-selected BL Lacertae Objects".

[30] Paiano, Falomo, Treves, et al., "The redshift of the BL Lac object TXS 0506+056".

[31] the source was named 2WHSP J050833.3+05310 in the previous catalog of HSP sources, see Figure 4.4

[32] Padovani, Giommi, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source".

[33] see Padovani, Giommi, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source" for technical details

ray test statistic maps are shown for various photon energy thresholds around the arrival of IceCube-170922A and during the neutrino flare, respectively. First of all, it can be seen that the gamma-ray sky looks very much different in the two time windows. While PKS 0502+049 shows strong emission during the neutrino flare, TXS 0506+056 clearly dominates at the observation time of IceCube-170922A. The strong emission of PKS 0502+049 during the neutrino flare has, in fact, lead people to suspect that something could be wrong with the pointing of the IceCube point source analysis.

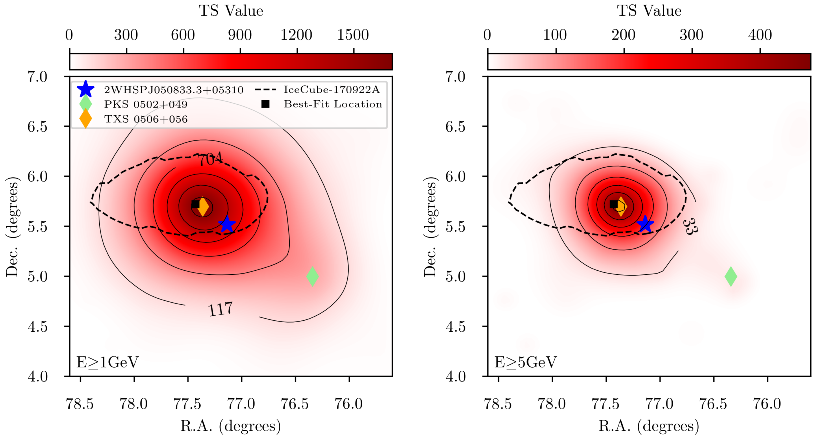

Figure 4.5: Test statistic map between MJD 57908 and 58018, i.e., in a time window 100 days before the arrival of IceCube-170922A. Colored symbols represent the sources of interest as shown above. The best-fit direction and 90% contour of IceCube-170922A are shown as black square and black dashed line, respectively. According to Wilk's theorem (theorem 7.1.2) the significance Σ is related to the TS value as $\Sigma = \sqrt{TS}$.

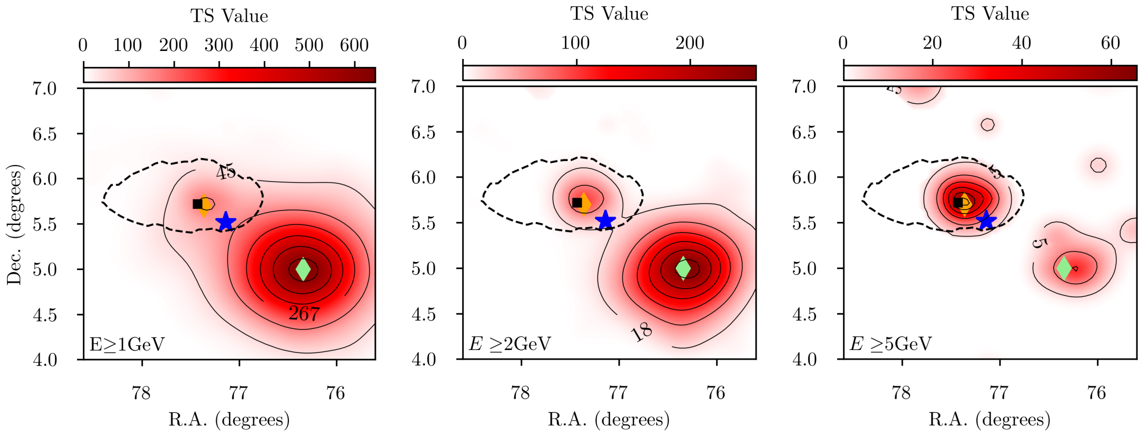

Figure 4.6: Test statistic map during the IceCube neutrino flare (MJD 56949 to 57059). See Figure 4.5 for more details.

Looking closer at the time profile of the emission shows, however, that the flare of PKS 0502+049 happens directly before and after, but not during the neutrino flare, see the light curve in Figure 4.7. Also, the dominance of PKS 0502+049 results from a bulk of low energy photons connected to the comparably soft spectrum of the source. Moving up to higher energies, TXS 0506+056 remains the strongest emitter in the region (Figure 4.6). Hence, both the spatial distances and the emission profile, point clearly against the neutrino emission being related to PKS 0502+049.

From Figure 4.5 and Figure 4.6, no other gamma-ray emitters are visible despite TXS 0506+056 and PKS 0502+049. As one of the other counterpart candidates, 3HSPJ050833.4+053109, lays right in between

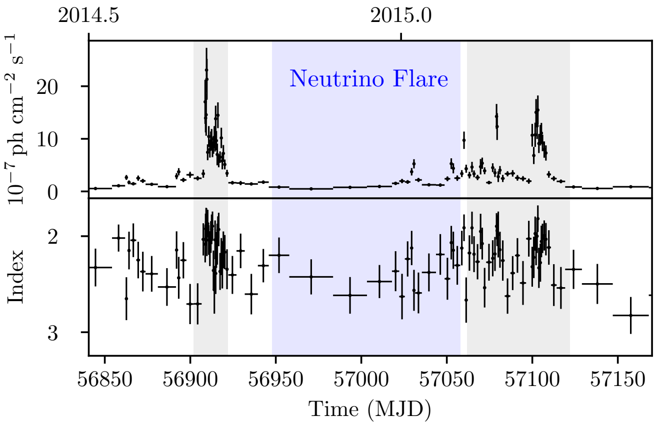

Figure 4.7: The 214 MeV to 100 GeV gamma-ray light curve of PKS 0502+049 around the time of the neutrino flare (blue). The width of the bins is chosen based on the adaptive-method by Lott et al., "An adaptive-binning method for generating constant-uncertainty/constant-significance light curves with Fermi-LAT data". Two periods of high flux can be identified (grey). Most importantly, however, the periods of high emission do *not* overlap with the time window of the neutrino flare.

the two, a dedicated gamma-ray analysis has been performed with a lower energy threshold of 100 MeV. The result shows, however, no clear sign of emission above the 3σ level. In fact, the fitted spectrum is extremely soft and most likely dominated by source confusion with TXS 0506+056.

Overall, both the analysis of the IceCube neutrino data and the multi-frequency dissection point to TXS 0506+056 as the only plausible (blazar) candidate in the region. In the following, the source's emission is analyzed in more detail to complete the phenomenological observations and search for a coherent picture. In Figure 4.8 the hybrid-SEDs, combining electromagnetic and neutrino emission[34],[35] for TXS 0506+056 are shown for the two relevant time windows. Evidently, the source is in very different states during those two emission periods. During the neutrino flare, the gamma-ray emission is faint at low energies but indicates a hardening towards higher energies. This is particularly visible when fitting the lower part of the spectrum (100 MeV-30 GeV) with a log-parabola and adding a second power law component starting at higher energies (>2 GeV). On the contrary, at the arrival time of IceCube-170922A, the gamma-ray flux is in its highest state while being well consistent with a simple power law model.[36]

The change of the spectral shape and the integrated photon flux are also visible in the 2 GeV - 300 GeV gamma-ray light curve, Figure 4.9. Although the states during the two neutrino emission periods are very different, they are both rare compared to the rest of the light curve. This can be estimated by comparing to the distribution of *normal states* given by the sum of Gaussians with respective mean and standard deviation over all light curve bins without associated neutrino emission (Figure 4.9, upper panel). As a result, the probability of observing a gamma-ray spectral index as hard as it is reached during the neutrino flare turns out to be only $\sim 2\%$. Integrating over the entire period of 110 days, the average spectral index above 2 GeV is 1.62 ± 0.20. Additionally, a detailed analysis of the highest-energy part of the spectrum has been done with an adaptive binning method above 10 GeV (Figure 4.9, right panel). Two emission periods

[34] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

[35] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Samarai, et al., "Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert".

[36] Glauch, Padovani, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source".

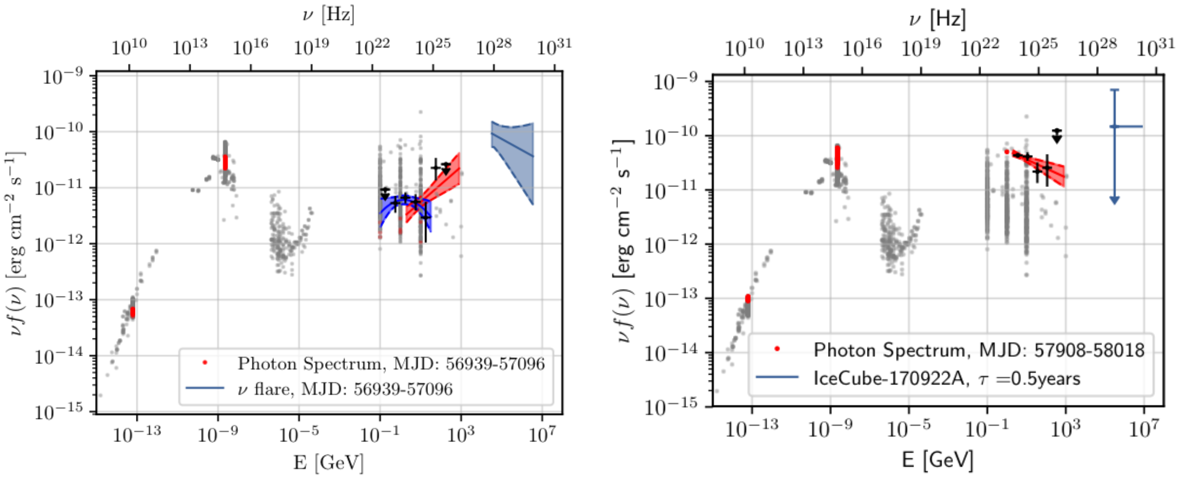

Figure 4.8: Hybrid-SEDs of TXS 0506+056 during the neutrino flare (left) and around the arrival time of IceCube-170922A (right). Archival data points are shown in grey and time simultaneous data in color. In addition to the neutrino and gamma-ray spectra, all the publicly available multi-frequency data are shown. The median (upper limit) neutrino flux in the right SED equals the flux required to observe on average one event (not more than one event at 90% C.L.) in 0.5 years.

of high energy gamma rays are identified with photon energies going up to 52.6 GeV at more than 99% confidence level. Similarly to the hard state during the neutrino flare, we find that the strong emission during the arrival of IceCube-170922A is pretty rare, with a chance probability of $\sim 1\%$. Most importantly, both gamma-ray states are also consistent with the respective neutrino emission flux levels at that time. This is an important requirement, as lepto-hadronic models predict a clear connection between the observed gamma-ray and neutrino luminosity, i.e.,

$$L_\nu = Y_{\nu\gamma} L_\gamma \tag{4.14}$$

with $Y_{\nu\gamma}$ being on the order of 1 for a photon energy range range between 1 GeV and 1 TeV.[37] For the case of TXS 0506+056 neutrino flare we find an all-flavor neutrino luminosity of $L_\nu = 1.4^{+0.6}_{-0.5} \times 10^{47}\,\mathrm{erg\,s^{-1}}$ between 32 TeV and 3.6 PeV and a simultaneous (extrapolated) gamma ray luminosity of $L_\gamma \sim 3 \times 10^{46}\,\mathrm{erg\,s^{-1}}$.[38] With those numbers being on a similar scale, we conclude that TXS 0506+056 is also energetically a consistent candidate.

In summary, all the spatial, timing, and energetic results, from both the IceCube neutrino analysis and the multi-wavelength dissection point to TXS 0506+056 as the first compelling non-stellar source of astrophysical neutrinos.

4.4 The Nature of TXS 0506+056 and its Multi-Frequency Emission

With the compelling detection of TXS 0506+056 as the first non-stellar neutrino point source, large interest sparked around the nature and the underlying mechanisms of the neutrino emission. Although many questions are still unresolved, some general conclusions and constraints can be derived.

[37] Petropoulou, Dimitrakoudis, et al., "Photohadronic origin of γ-ray BL Lac emission: implications for IceCube neutrinos".

[38] Padovani, Giommi, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source".

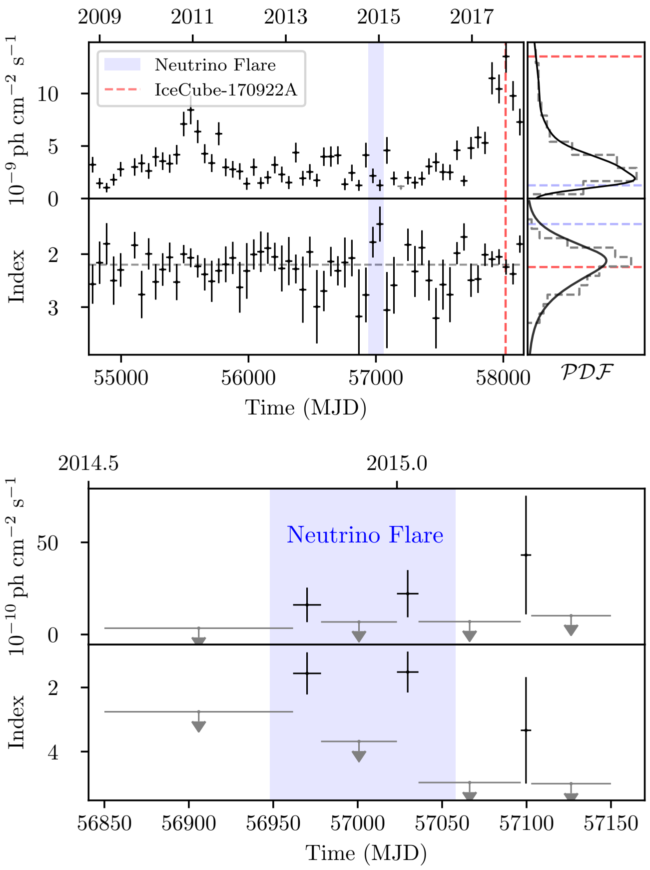

Figure 4.9: Gamma-ray light curves of TXS 0506+056. The upper panel shows the light curve over the full mission using a binning of half the neutrino flare time (55 days) and the lower panel shows the high energy emission (above 10 GeV) during the neutrino flare. The neutrino flare period and arrival time of IceCube-170922A are indicated in blue and red, respectively. In addition to the photon fluxes on the top, the spectral indices are shown on the bottom of each plot. The right panel in the full-emission light curve additionally shows the distribution of neutrino-quiescent states of the source.

First, there is the challenge of classifying TXS 0506+056 among the classes of blazars. As optical emission lines are strongly diluted by the jet, early works on the source assumed it is a BL Lac. A dedicated analysis of the optical spectra showed, however, the presence of OII and OIII emission lines with luminosities typical for high-excitation galaxies (HEGs, section 2.3.3).[39] From a synchrotron peak perspective, the source is further classified in the transition region between IBLs and HBLs and, hence, falls into the category of masquerading BL Lacs, i.e., FSRQs that appear as BL Lacs due to the strong non-thermal jet emission.[40] The classification of TXS 0506+056 as an FSRQ has important implications as the narrow line region might provide a stationary photon field for photo-meson production of neutrinos in the jet. Following eq. (2.19) these target photons are required to be in the ultraviolet band to produce neutrinos at a few hundred TeV. For an FSRQ such target photons naturally arise from the accretion disk and the narrow line region. In contrast, the scenario of neutrino production in a co-moving photon field would require target photons in the keV (X-ray) range. Multi-wavelength measurements around the arrival time of IceCube 170922A have found, however, a strong dip in the X-ray range, and hence the amount of target photons in this scenario is too low to produce neutrinos in reasonable numbers without requiring unreasonable proton luminosities.[41]

[39] Paiano, Falomo, Treves, et al., "The redshift of the BL Lac object TXS 0506+056".

[40] Padovani, Oikonomou, et al., "TXS 0506+056, the first cosmic neutrino source, is not a BL Lac".

[41] Reimer, Boettcher, and Buson, "Cascading Constraints from Neutrino Emitting Blazars: The case of TXS 0506+056".

Several authors have used numerical lepto-hadronic models to perform a detailed analysis of the spectral energy distribution of TXS 0506+056 around the two relevant neutrino emission time windows[42,43,44].[45] Despite some differences in the assumptions, there is a consensus that major constraints for the proton, and thereby the neutrino luminosity, come from the X-ray dip. Specifically, the proton luminosity is constrained by the Bethe-Heitler emission that should not exceed the measured X-ray flux. While X-ray measurements are available around the arrival time of IceCube-170922A, there are no measurements during the neutrino flare, leaving more space for interpretation. In summary, most works agree that consistent modeling of the IceCube-170922A flux is possible in standard one-zone models. For the neutrino flare, in contrast, things are more complicated as the overall gamma-ray flux is lower. A possible explanation could be the existence of several emission zones or the presence of a large X-ray flux.[46] In any case, it will need further neutrinos observations to unambiguously identify the processes that can explain the neutrino production from TXS 0506+056. Finally, note that the observation of neutrinos with an energy of a few hundred TeV to 1 PeV only requires proton energies of $\sim 10\,\mathrm{PeV} (\sim 10^{16}\,\mathrm{eV})$ and hence does not necessarily implicate the production of the highest energies of cosmic rays.

[42] Keivani et al., "A Multimessenger Picture of the Flaring Blazar TXS 0506+056: implications for High-Energy Neutrino Emission and Cosmic Ray Acceleration".

[43] Cerruti et al., "Leptohadronic single-zone models for the electromagnetic and neutrino emission of TXS 0506+056".

[44] Liu et al., "Hadronuclear interpretation of a high-energy neutrino event coincident with a blazar flare".

[45] Rodrigues et al., "Leptohadronic Blazar Models Applied to the 2014–2015 Flare of TXS 0506+056".

[46] Ibid.

5 Deep Learning Event-Type Classification in IceCube

Although the idea of artificial neural networks (ANNs) has been around for decades, it was only recently that improvements in GPU power and new network architectures have paved the way for widespread applications.[1] Since then, their development has accelerated tremendously and, nowadays, ANNs are reaching human performance levels in various tasks such as pattern recognition,[2] medical diagnosis,[3] and playing highly-complex board games such as *Go*.[4] Notably, a readily trained neural network is fast in the application, with prediction times on the order of milliseconds (see Figure 5.14). Of special interest in the context of ANNs are the so-called *deep neural networks* (DNNs), which describe ANNs where the input is connected to multiple *hidden layers* before reaching an output node. In this way, modern DNNs can handle up to several million free parameters, which the network freely adjusts based on a set of training samples. This procedure is sometimes referred to as *learning*, in analogy to human brains.

Physics research is another large field of application for DNNs. With data acquisition rates rapidly increasing, there is a growing desire for fast and powerful algorithms to process and filter interesting measurements as early as possible in the data processing pipeline. Consequently, several groups within the IceCube Collaboration have worked on different applications such as track and cascade reconstructions using convolutional (CNNs)[5] and recurrent neural networks (RNNs), as well as muon bundle identifications using graph neural networks (GNNs).[6] Complementary, the identification of different event topologies is an important part of the IceCube processing pipeline. Therefore, this chapter presents a deep neural network that is trained to identify the five major event topologies observed in the IceCube In-Ice Array at energies above $\sim 100\,\mathrm{GeV}$. As a practical application example, the subsequent chapter 6 presents a new selection of realtime events that combines the DNN classification with the RNN muon reconstruction. This new selection significantly improves event rates over its predecessor while having a comparable purity.

[1] LeCun, Bengio, and G. Hinton, "Deep learning".

[2] He, X. Zhang, et al., "Deep residual learning for image recognition".

[3] Litjens et al., "A survey on deep learning in medical image analysis".

[4] Silver et al., "Mastering the game of Go with deep neural networks and tree search".

[5] Huennefeld, "Deep Learning in Physics exemplified by the Reconstruction of Muon-Neutrino Events in IceCube".

[6] Choma et al., "Graph Neural Networks for IceCube Signal Classification".

This chapter begins with a short review of the concepts behind neural networks before explaining the CNN-based event classifier and evaluating its performance and consistency on different datasets. An earlier version of the classifier was presented at the ICRC 2019.[7]

5.1 *Artificial Neural Networks and Deep Learning*

5.1.1 *The Neuron and Simple Neural Networks*

The idea of modelling the human brain through computational systems has inspired researches and philosophers for decades. Starting from the early ideas of computer models[8] and perceptrons[9] in the 1940s and 1950s, the concepts have become more and more abstract, with modern types of artificial neural networks (ANNs) only being vaguely connected to their biological counterparts. From a modern perspective, deep learning describes a subset of machine learning methods that are based on a system of artificial neurons. A sketch of such a unit is shown in Figure 5.1. The neuron takes a vector of inputs $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$ which is multiplied by a set of weights $\mathbf{W} = (w_1, w_2, \ldots, w_n) \in \mathbb{R}^n$ to produce an output y_1

$$y_1 = \mathbf{W}^T \cdot \mathbf{x} + b_1, \tag{5.1}$$

where the b_1 adds an offset or bias. Eq. 5.1 describes a simple linear model. Based on this prototype neuron, more complex architectures can be constructed. An extended example is shown in Figure 5.2.

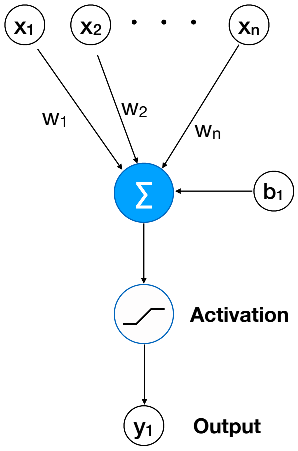

Figure 5.1: Sketch of a simple neuron, with inputs $\mathbf{x}$, a set of weights $\mathbf{W}$ and a bias b_1. The inputs are combined according to eq. (5.1). Finally, an activation function can be used, eq. (5.5).

There are two major differences to the previous model: 1.) the outputs of the first layer are used as inputs for the second layer (*hidden layer*) and 2.) several outputs are produced at the same time. Using the previous definition of the neuron the hereby constructed network can be written as

$$\mathbf{y} = \mathbf{W}_2(\mathbf{W}_1\mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2 \tag{5.2}$$
$$= \mathbf{W}_2\mathbf{W}_1\mathbf{x} + \mathbf{W}_2\mathbf{b}_1 + \mathbf{b}_2 \tag{5.3}$$
$$= \mathbf{W}_{1,2}\mathbf{x} + \mathbf{b}_{1,2} \tag{5.4}$$

where $\mathbf{W}_1 \in \mathbb{R}^{m\times n}$ and $\mathbf{W}_2 \in \mathbb{R}^{l\times m}$ are now matrices, with m and l nodes in the first and second layer, respectively. Although the model seems to have more complexity than the previous one, it can be seen from eq. (5.4) that this linear combination of neurons can be simplified and is itself again a linear model. In order to overcome this limitation and construct a network that can model non-linear mappings, a (non-linear) *activation function* σ can be used. Eq. (5.1) changes thereby to

$$y_1 = \sigma(\mathbf{W}^T \cdot \mathbf{x} + b_1). \tag{5.5}$$

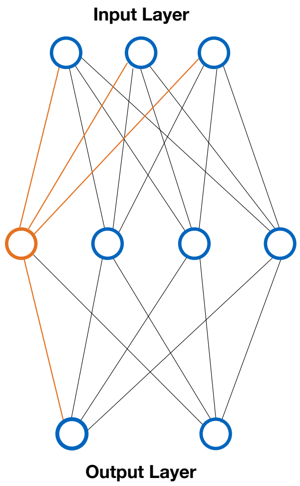

Figure 5.2: Sketch of a simple neural network with input, hidden and output layer. As all nodes between subsequent layers are connected, it is called a *fully-connected* network. A single neuron in the network is highlighted.

Common choices of activation functions are depicted in Figure 5.3. A network which passes data in only one direction, as the example in Figure 5.2, is called a *feed-forward network*. Although such networks

---

[7] Kronmueller and Glauch, "Application of Deep Neural Networks to Event Type Classification in IceCube"; Kronmüller, "Application of Deep Neural Networks on Event Type Classification in IceCube".

[8] Shannon and McCarthy, *Automata Studies. (AM-34), Volume 34*.

[9] Rosenblatt, "The perceptron: a probabilistic model for information storage and organization in the brain."

with only one hidden layer are not used in practice, they are the basis for one of the major results in the mathematical theory of artificial neural networks, the universal approximation theorem.

Theorem 5.1.1 (Universal Approximation Theorem) [10] *Let σ be a non-constant, bounded, and monotone-increasing continuous function. Let I_{m_0} denote the m_0-dimensional unit hypercube $[0,1]^{m_0}$. The space of continuous functions on I_{m_0} is denoted as $C(I_{m_0})$. Then given any function $f \in C(I_{m_0})$ and $\epsilon > 0$ there exists a set of real constants a_i, b_i and w_{ij}, where $i = 1, \ldots, m_1$ and $j = 1, \ldots, m_0$ such that we may define*

$$F(x_1, \ldots, x_{m_0}) = \sum_{i=1}^{m_1} a_i \sigma \left(\sum_{j=1}^{m_0} w_{ij} x_j + b_i \right) \tag{5.6}$$

as an approximate realization of the function f; that is,

$$|F(x_1, \ldots, x_{m_0}) - f(x_1, \ldots, x_{m_0})| < \epsilon \tag{5.7}$$

for all $x_1, x_2, \ldots, x_{m_0}$ that lie in the input space.

This means that every well-behaved function can theoretically be approximated through a simple feed-forward network with one hidden layer. While this is a powerful statement, the theorem contains no information on the network's size and the procedure to choose the optimal weights. As we will see later, modern applications of neural networks tend to be deeper, i.e., have many hidden layers instead of having many parameters in one layer, which is why we frequently talk about *deep* neural networks.

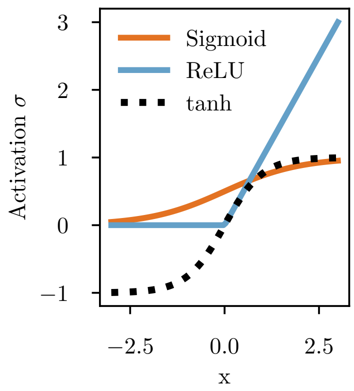

Figure 5.3: Common activation functions used in artificial neural networks.

5.1.2 *Learning Strategies*

In general, machine learning methods can be used for a wide range of problems, which are usually split up into

- **Supervised Learning**: Given input data x and output values y, the network is a model $\hat{y} : x \rightarrow y(x)$ that approximates the dependency of y from x using a set of weights θ. Typical examples are regression and classification problems.
- **Unsupervised Learning**: Given some input data x, the algorithm identifies underlying patterns. An example is the estimation of a probability density function $f(x)$ using kernel density estimation, see section 7.4.
- **Reinforcement Learning**: Some environments, such as games, are highly dynamical. In this case, the model tries to optimize a reward function to perform the best action.

Physics applications usually have large-scale Monte Carlo simulations that connect a set of observables x to an underlying truth y. Hence we can apply supervised learning strategies to find an approximation $\hat{y}(x)$ for the unknown mapping between x and y. After

---

[10] Haykin, *Neural Networks: A Comprehensive Foundation*.

defining the neural network architecture, sets of training samples - batches - are propagated through the network and evaluated against the truth. To measure the quality of the predictions for a given set of weights θ, a loss or objective function $J(\theta)$ is defined. It measures the distance between the truth and the prediction under a given metric. Based on this measure, the weights between the input neuron i and output neuron j are adjusted via *backpropagation*, i.e.

$$w_{ij}^{t+1} = w_{ij}^{t} - \alpha \frac{\mathrm{d}J(\theta)}{\mathrm{d}w_{ij}}. \tag{5.8}$$

with the weights w_{ij} of the t and $t+1$ iteration and a *learning rate* parameter α. Note that the calculation of the derivative in eq. (5.8) requires the chain rule to include recursively all the derivatives up to the output node. The general idea is to do this procedure, starting from the output node and then going all the way to the input node to reuse recurring derivatives. The entire procedure of repeatedly feeding data through the network and adjusting the weights through backpropagation is called *gradient descent*.

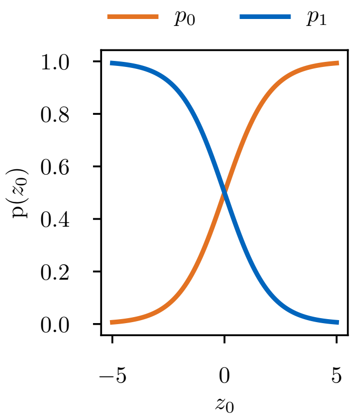

Figure 5.4: Sketch of the softmax activation function for the binary case of two output nodes z_0 and z_1. z_1 is fixed to zero and z_0 varied, leading to the output probabilities p_1 and p_2 as shown.

5.1.3 Loss Functions

The choice of the loss function depends on the specific task and the activation function in the last layer. Typical combinations are mean squared error loss and linear activation for regression tasks and softmax activation and cross entropy, eq. 5.11, for classification tasks. The softmax activation function for a vector of output values $\mathbf{z} = (z_1, \ldots, z_K) \in \mathbb{R}^K$ is defined as

$$p_i(\mathbf{z}) = \frac{e^{z_i}}{\sum_{j=1}^{K} e^{z_j}} \text{ for } i = 1, \ldots, K. \tag{5.9}$$

As a result the sum of the output values p_i is normalized to one, facilitating the interpretation of a vector of probabilities or *p-scores*, $\mathbf{p} = (p_0, p_1, \ldots, p_K)$ for each input sample. A visualization of the softmax function is shown in Figure 5.4. In accordance with the probabilistic interpretation, one usually uses a *cross entropy* loss function for classification tasks. For a simplified derivation, assume that the estimated probability for the class j of outcome i is p_{ij}, while the known probability distribution of the classes is given by q_{ij}. In this case the multinomial likelihood is given by

$$\mathcal{L} = \prod_{i=0}^{N} \prod_{j} p_{ij}^{q_j}. \tag{5.10}$$

For more convenient evaluation the equation can be simplified to get the cross entropy $H(\mathbf{P}, \mathbf{Q})$ with

$$H(\mathbf{P}, \mathbf{Q}) = -\frac{1}{N} \log \mathcal{L} = -\frac{1}{N} \sum_{i=0} \sum_{j} q_{ij} \log p_{ij} \tag{5.11}$$

$$= -\frac{1}{N} \sum_{i=0}^{N} \sum_{j} q_{ij} \log p_{ij}. \tag{5.12}$$

$\mathbf{P}$ and $\mathbf{Q}$ are vectors of the single samples, i.e. $\mathbf{P} = (\mathbf{p}_0, \ldots, \mathbf{p}_N)$ and $\mathbf{Q} = (\mathbf{q}_0, \ldots, \mathbf{q}_N)$; N is the number of samples per iteration or *batch size*. In machine learning classification we know the true class labels and hence every $\mathbf{q}_i$ in eq. (5.11) is just a vector of the form $\mathbf{q}_i = (0, \ldots, 1, \ldots, 0) \in \mathbb{R}^K$ with a unique $q_{ij} = 1$ representing the known true class label.

5.1.4 Convolutional Neural Networks

In many cases, it is not desirable that every node in one layer is connected to every node in the following layer. Specifically, for the case of pattern recognition tasks, a feature is usually spatially bounded, e.g., a dog in an image or a cascade in the IceCube detector. Therefore the network's complexity can be drastically reduced by searching for local features using a set of filters that move along the input grid. Let us assume an input grid H of the size $l \times m \times f \in \mathbb{N}^3$, where f is called the feature dimension, e.g., the colors RGB in an image or charge and time features for an IceCube event. Mathematically, moving a filter K of size $j \times k$ along this grid is called a *2D convolution*

$$G(l,m) = (H * K)(l,m) = \sum_j \sum_k \sum_f H(l-j, m-k, f) \cdot K(j,k,f). \quad (5.13)$$

In a simplified view those filters search for structures, e.g. lines or spheres, and always go along the *entire* feature dimension. Using n filters, the output grid G is therefore of the size $l \times m \times n$ [11]. Repeating those computations in subsequent layers, the initial features are combined to identify more and more complex structures. Note that the same filters are used one the entire input grid, intrinsically assuming translational invariance. A network that is based on convolutional layers is called a *convolutional neural network* (CNN).

As the precise location of a feature usually does not matter, the network can be further simplified through *pooling* operations. The idea is to shrink the image size by mapping several values in a bounded region to their maximum (or average) value. Similarly to the convolution, this operation is performed subsequently along the entire grid. In contrast to the convolutional filters, the pooling operation is thereby applied to each feature dimension separately. To get a significant reduction in grid size, the pooling operation is usually applied with a stride $s > 1$, i.e., the bounded region is moved s grid points after each pooling operation. The original grid is thereby shrunk by a factor s in each dimension. A special version of the pooling operation is the *gloabal average pooling*, where the pooling filter has the same size as the input grid. As a result, the input grid is reduced to the size $1 \times 1 \times n$, where n is the number of features of the input grid. Global average pooling layers are frequently used to summarize the extracted features before the output node. The $1 \times 1 \times n$ grid is then flattened to an input of size n that is fully connected to the output nodes. Note finally that the discussion of this section can be easily extended to higher dimensional convolutions.

---

[11] Assuming a zero padding at the edges.

5.2 Classifying Event Topologies in IceCube

Different neutrino flavors and interaction channels leave different traces in the IceCube detector. While those differences in the topology are challenging for event selection and reconstruction, they allow for a broader range of physics searches that facilitate each class's respective advantages. Cascades and starting tracks, for example, are characterized by their interaction vertex being inside the detector volume. This enables a clear distinction from atmospheric muons and a better reconstruction of the initial neutrino's energy. On the contrary, through-going tracks have a good angular resolution and large statistics, making them the preferred choice for neutrino point source searches.

The majority of events that pass the $\sim 2.8\,\mathrm{kHz}$ hardware trigger of IceCube have an atmospheric muon origin and are not useful for neutrino physics searches. Therefore, a major goal of the low-level processing is the identification of a sub-sample of potentially interesting neutrino events based on their topology and direction. In the current IceCube pipeline the event stream is therefore split into a *Cascade* and a *Muon* level 3 stream[12]. Despite the importance of the event topology in the selection process, there is no algorithm to classify event topologies directly based on the recorded pulses, yet. From processing level 2 to level 3, topological selections, therefore, rely on quality parameters of simple reconstructions and topological variables as the event's charge-weighted center of gravity. This is both laborious and counterintuitive as reconstructions are applied to event topologies they were not developed for. With a deep learning classifier, one can, in contrast, invert the procedure, i.e., *the data processing can follow a natural logic by running first the topological classification and then the optimal reconstruction.*

The classifier presented here is aimed to be *generic* in a sense that it covers simultaneously *all* the major event topologies of IceCube, as shown in Figure 5.5. In the following the definition of the event classes is discussed in more detail: [13] [14]

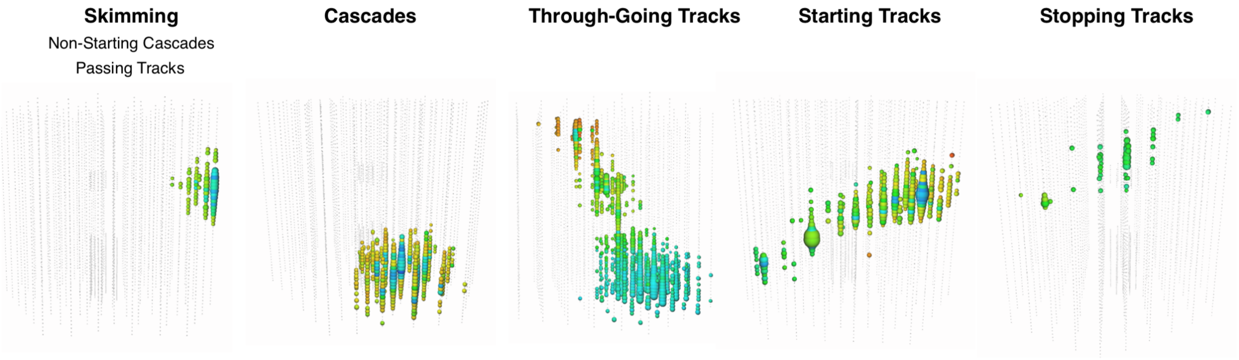

Figure 5.5: Event views of the five IceCube topologies predicted by the classifier. Colors range from red (early) to blue (late) and the size of the bubbles represents the amount of collected charge.

---

[12] Check section 3.5 for a review of the IceCube data pipeline.

[13] The corresponding software can be found under https://github.com/tglauch/DeepIceLearning/blob/master/lib/reco_quantities.py.

[14] see also section 3.3 for a review of the topologies.

- **Skimming Events**: This class contains all the events with no energy deposition in the detector. Major constituents are cascades with their primary vertex outside of the detector (uncontained cascades) and *passing* tracks, i.e., tracks that never enter the detector. Depending on the track direction, the two classes are not clearly distinguishable, e.g., a track passing a corner of the detector and leaving a single large energy deposition can appear as an uncontained cascade. In general skimming events are excluded from most IceCube analyses, as they are difficult to reconstruct and susceptible to atmospheric muon background.

- **(Starting) Cascades** are produced by neutral current interactions of neutrinos of any flavor, as well as charge current interactions of electron neutrinos with primary vertex inside the detector volume [15]. The hadronic and electromagnetic cascades produced in these interactions are usually indistinguishable. It has been suggested that a distinction is possible based on a neutrino echo,[16] which is, however, not visible using the standard IceCube SMT triggers. At high energies, an additional, but minor, contribution to this event class comes from Glashow resonance events, where the W-Boson decays either hadronically or into a $e\,\nu_e$ pair. Cascades are usually well contained in the detector, therefore providing a good measure for the primary neutrino energy. Besides, cascades have a larger probability of being of astrophysical origin than the muon-induced tracks, as the background of atmospheric ν_e is suppressed at neutrino energies above 100 GeV (see section 2.5). Therefore they are specifically important for the precise measurement of the astrophysical neutrino spectrum[17] and its flavor ratio.[18]

- **Through-Going Tracks**: This event class is unique to muons, which are either produced through ν_μ charged current (CC) interactions or in cosmic ray air showers in the atmosphere. Also, a much smaller number is produced via the subsequent decay of tau leptons from charged current ν_τ interactions. The unique feature of this topology is the long lever arm of the track that enters and leaves the detector. Through-going tracks constitute the basis of all major IceCube analyses searching for astrophysical sources,[19] as well as the high-statistic search for diffuse astrophysical muon-neutrinos.[20] Finally, through-going muons and muon bundles from the Southern Hemisphere can be used to probe the flux of muons from atmospheric air showers.[21]

- **Starting Tracks** are similar to through-going tracks but are exclusively produced by charged current ν_μ interactions with the interaction vertex inside the detector volume. This allows for clear separation from atmospheric muons, hence enabling selections of high-purity samples of astrophysical neutrinos over the entire sky. Starting tracks are one of the major channels for the IceCube realtime alert stream[22] of which an improved, DNN-based version

[15] Double bang topologies from CC tau neutrino interactions can be indistinguishable from cascades when the track length of the tau lepton is below the detector resolution of $\sim 5\,\mathrm{m}$. For consistency, they are, however, not defined as cascades here.

[16] Steuer, "Cascade type identification in IceCube and an application in a search for new physics".

[17] M. Aartsen et al., "Characteristics of the diffuse astrophysical electron and tau neutrino flux with six years of IceCube high energy cascade data".

[18] M. G. Aartsen et al., "Flavor Ratio of Astrophysical Neutrinos above 35 TeV in IceCube".

[19] M. G. Aartsen et al., "Search for steady point-like sources in the astrophysical muon neutrino flux with 8 years of IceCube data".

[20] M. G. Aartsen et al., "Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data".

[21] M. G. Aartsen et al., "Characterization of the Atmospheric Muon Flux in IceCube".

[22] Blaufuss et al., "The Next Generation of IceCube Realtime Neutrino Alerts".

is presented in chapter 6. Furthermore, starting tracks are used to measure the inelasticity of the interaction as both the starting cascade and the outgoing track are visible.[23]

- **Stopping Tracks** are produced in the same processes as through-going tracks but require the muon track to end inside the detector volume. As the track length increases with the muon energy, this event class is dominated by events with small energy on detector entry. Therefore, those events can be used to study the detector response in the minimum ionizing regime.[24]

Finally, double-bangs, another important event class, are not considered by the classifier as they are in most cases indistinguishable from cascade events and have low simulation statistics. A possible solution is to run an additional binary network with optimized input features that tries to identify double-bang events from a previously selected cascade stream. Feasibility studies for that are still on-going.

5.3 *From Waveforms to DNN Input*

On each of its 5160 DOMs, IceCube is measuring the photon arrival times using two different digitizers - ATWD and FADC. While the first one has a recording time of only 427 ns to capture light produced tens of meters away from the DOM, the latter one records for 6.4 µs to capture distant and scattered light.[25] The respective sampling rates are 3.3 ns (ATWD) and 25 ns (FADC). An example of a recorded waveform on a single DOM is shown in Figure 5.8 (top panel). Instead of using the waveform directly, IceCube reconstructions are based on the unfolded pulses, i.e., a series of charge and time tuples for each DOM [26]. In terms of the DNN, it is tempting to use a binned version of these pulses as input quantities. To estimate the input data size for this approach, we consider the simplified example of a muon traversing the detector with a track length of 1500 m [27]. Assuming the muon is traveling with vacuum speed of light c_0 this requires a time of

$$t_\mu = \frac{1500\,\mathrm{m}}{c_0} \approx 5000\,\mathrm{ns}. \tag{5.14}$$

Since there is an additional time delay until the detection of the photons, eq. (5.14) is only a lower-bound for the required time window of a typical event. This implies that, even when binning with the sampling rate of the FADC, each DOM would have more than $t_\mu/(25ns) = 400$ input bins. In total this sums up to more than $N = 400 \times 5160 = 2'064'000$ input data points or ~ 16.5 Megabyte per event, which is computationally unfeasible.

To reduce the amount of input data, one can use a set of suitable input features instead of a fixed time binning. The input can thereby be drastically compressed while keeping most of the (relevant) information. Examples for such features are the time of the first hit or the total collected charge, which are also used in other - likelihood-based - reconstructions (section 3.8). Such an approach has already

[23] M. G. Aartsen et al., "Measurements using the inelasticity distribution of multi-TeV neutrino interactions in IceCube".

[24] Tosi and Wendt, "Calibrating photon detection efficiency in IceCube".

[25] M. G. Aartsen, Ackermann, et al., "The IceCube Neutrino Observatory: instrumentation and online systems".

[26] see section 3.5 for a detailed review

[27] Approximated distance from one corner of the detector to the other

been presented for the reconstructions of muons and cascades[28].[29] In this thesis, we show that similar techniques can be used for the classification of events.

The basic structure of the DNN input consists of three spatial dimensions and a feature dimension, combining time and charge information. As discussed in chapter 3, the IceCube detector is arranged on a fairly regular, hexagonal grid of 86 strings, of which 7 are assigned to the deep core sub-detector optimized for low energy studies. As the DNN classifier aims for events above 100 GeV, it is convenient to neglect the DeepCore strings as they break with the detector's overall symmetry. A top view of the remaining In-Ice Array is depicted in Figure 5.6. In contrast to other approaches as graph-neural networks, CNNs are optimized for pattern recognition in images and hence require a fixed input grid of *pixels*. The required transformation is illustrated in Figure 5.6 and Figure 5.7. It consists of two steps: 1.) the detector is aligned with the x-axis, 2.) rows are alternately shifted to the left and right to align all the strings along the column direction. This procedure results in a 10×10 grid as shown in Figure 5.7. The transformation guarantees that string-to-string relations are preserved over the detector. Finally, every string of the IceCube Array has 60 DOMs in a regular distance of 15 m. Consequently, the spatial part of the input has a total size of $10 \times 10 \times 60$.

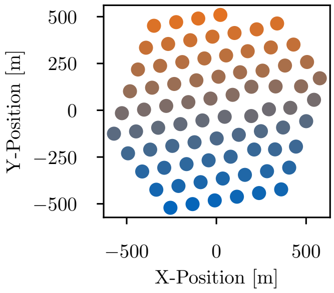

Figure 5.6: Top view of the hexagonal IceCube In-Ice Array. Colors represent the string numbers in increasing order.

The choice of the charge/time input features is discussed based on Figure 5.8. In addition to the waveform (top) and the unfolded pulses (middle), the bottom panel shows the accumulated charge against time. Evidently, the main fraction of the charge is collected early in the recording window, followed by a tail of scattered photons, as well as late and after pulses. In general, the most accurate timing information is carried by the first, unscattered photons. It is therefore useful to choose features as the deposited charge after *10 ns*, *50 ns* and *100 ns* from the first hit. A huge rise in collected charge on these time-scales is good indicator for a large energy loss in a DOM's proximity [30]. To capture most of the information on how quickly the charge piles up, we calculate *charge quantiles*, i.e., the time after which a certain fraction of the total charge has been collected. In contrast to the charge after fixed time intervals, charge quantiles are adaptive, as seen in Figure 5.8. Specifically, they get denser when more charge is collected. Additional timing features are the *time of the first pulse*, which contains precise distance information, as well as the *standard deviation* and the *spread* of the photon arrival times. Finally, the *total charge* is used as a measure for the importance of a DOM. The total charge also resolves the ambiguity between the zero timestamps for hit and non-hit DOMs [31]. Note that the absolute times inside the original trigger window are given with respect to an arbitrary reference t_0. For the DNN input this is resolved by calculating all times relative to the charge-weighted median pulse time of the event as a well-defined reference point. In summary, the CNN has 15 input features,

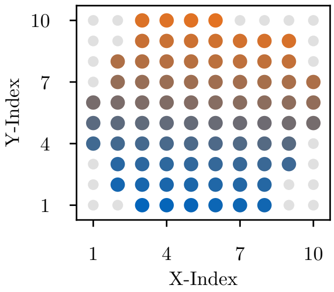

Figure 5.7: Top view of the DNN input grid in pixel coordinates X and Y. Some pixels stay empty after the transformation and are therefore padded with zeros, indicated as grey points in the plot. The other colors follow the string numbers in increasing order analogous to Figure 5.6.

[28] Huennefeld, "Deep Learning in Physics exemplified by the Reconstruction of Muon-Neutrino Events in IceCube".

[29] Huennefeld, "Reconstruction Techniques in IceCube using Convolutional and Generative Neural Networks".

[30] For a refractive index of $n \approx 1.3$ Cherenkov photons travel around 2.3 m in 10 ns

[31] All features are set to zero for non-hit DOMs, which produces ambiguous time values if the total charge is not known.

- Total charge & charge after 10ns, 50ns, 100ns [p.e.]
- Time of the first hit, time spread, standard deviation of the time [ns]
- Time after which 1%, 3%, 5%, 11%, 15%, 20%, 50%, 80% of the charge have been collected [ns]

resulting in a 4-dimensional input of the size $10 \times 10 \times 60 \times 15$. Neural networks perform best with input data that are roughly distributed between -1 and 1. Hence, the input is transformed by dividing charge values by a factor of hundred and time features by a factor of 10000 (event windows typically have a length of several thousand ns).

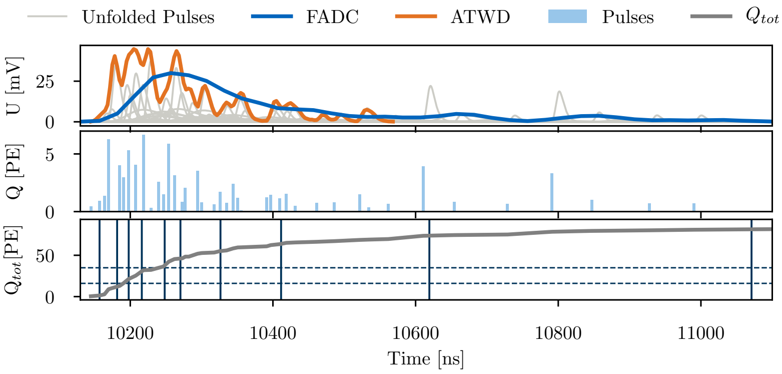

Figure 5.8: A sketch of the transformations from recorded waveforms (top), over pulses (middle) to the DNN input (bottom). The waveform information is unfolded into pulses using pulse templates (thin, grey lines in the top panel). Subsequently, the pulses can be used to calculate the measured charge versus time, indicated as the solid grey line in the bottom panel. Vertical solid lines indicate the times of successive 10% charge quantiles. Clearly, the quantiles are denser in areas where more charge is deposited. Dashed horizontal lines indicate the charge collected after 50 ns and 100 ns from the first pulse, respectively. The time range is clipped to show the most important part of the event.

5.4 *Training Dataset*

The quality of the training data is important and often limiting factor for the performance of DNNs. Fortunately, for the case of IceCube, large-scale Monte Carlo simulations are available, containing all the required information for high-statistics supervised learning. Two design choices need to be made: The composition of the dataset and the weighting of the events during the training. In general, there are three possibilities: 1.) an unbiased training dataset with event distributions following the IceCube best-fit flux and no weighting, 2.) an arbitrary biased training dataset with a best-fit weighting factor in the loss function, 3.) a biased dataset with event distributions being equal among the event topologies and no weighting factor.

As noted previously, the DNN classifier presented in this chapter is supposed to be generically applicable, implying that predictions are only based on the specific event topology. This corresponds to a training dataset with equal event distributions among the classes (3.). The dataset is further biased in the sense that it is constructed in a way that sufficiently covers all parts of the observable space, including high energies that rarely appear in experimental data but are of large importance for many analyses. For a specific analysis with a given flux model H, the probability of a certain event being

Dataset	Properties
11499 (00000-00999)	*Ice Model*: SPICELea CLSim *Flux Model*: CORSIKA-in-ice 5-component Spectrum: $E^{-2.6}$ (600 GeV - 100 TeV) *DOM efficiency*: 1.0
11362 (00000-01999)	*Ice Model*: SPICELea CLSim *Flux Model*: CORSIKA-in-ice 5-component Spectrum: $E^{-2.0}$ (100 TeV - 10 EeV) *DOM efficiency*: 1.0
Nancy's Simulation ν_μ: (00000-02999) ν_e: (00000-00999)	*Ice Model*: Spice 3.2.1 *Energy Range*: 100 GeV - 5 TeV (E^{-2} Spectrum) *Hole Ice*: Flasher-fit-model (p1=0.3, p2=0) *DOM efficiency*: 0.99
Nancy's Simulation ν_μ: (00000-02999) ν_e: (00000-01999) ν_τ: (00000-00999)	*Ice Model*: Spice 3.2.1 *Energy Range*: 5 TeV - 10 PeV ($E^{-1.5}$ Spectrum) *Hole Ice*: Flasher-fit-model (p1=0.3, p2=0) *DOM efficiency*: 0.99
Nancy's Simulation ν_μ: (00000-00999)) ν_e: (00000-00999) ν_τ: (00000-00999)	*Ice Model*: Spice 3.2.1 *Energy Range*: 1 PeV - 100 PeV (E^{-1} Spectrum) *Hole Ice*: Flasher-fit-model (p1=0.3, p2=0) *DOM efficiency*: 0.99

Table 5.1: Summary of the input Monte Carlo datasets used for the training of the DNN classifier. While the neutrino datasets ("Nancy's Simulation") use the latest ice model Spice 3.2.1, the atmospheric showers simulated with CORSIKA use a slightly older ice model. When using the DNN for any analysis these simulation sets should be excluded.

of true type X given the prediction score (p-score) p_X can then be approximated using Monte Carlo-based conditional probabilities

$$P(X|p_X, H) = \frac{P(X, p_X|H)}{P(p_X|H)}. \tag{5.15}$$

To reduce processing time during training, the input data are preprocessed. This includes both generating the 4-dimensional input grid as described in section 5.3 and calculating the true event label from the *I3MCTree*, a Monte Carlo variable describing the entire stack of particles produced during event simulation. In Table 5.1 the input Monte Carlos and their simulation properties are listed. The preprocessed data are finally stored in HDF5 files.

The resulting dataset covers events from the entire sky in an energy range between 100 GeV and 100 PeV. In addition, a 10% natural coincidence rate of atmospheric muons is kept, in order increase stability against this type of background contamination. The true label for an event with coincidence is calculated as if the coincidence was not there, e.g., a cascade with a coincident background muon would be classified as a cascade. Muon bundles are included by classifying them similarly to single muons. Specifically, a muon bundle with *any* muon reaching the other end of the detector is labeled as throughgoing, while a bundle with *all* the muons stopping in the detector is labeled as stopping.

In the last step, events are randomly selected to construct a training dataset that is equally distributed between the classes in zenith, azimuth, and the number of hit DOMs. In contrast to the deposited

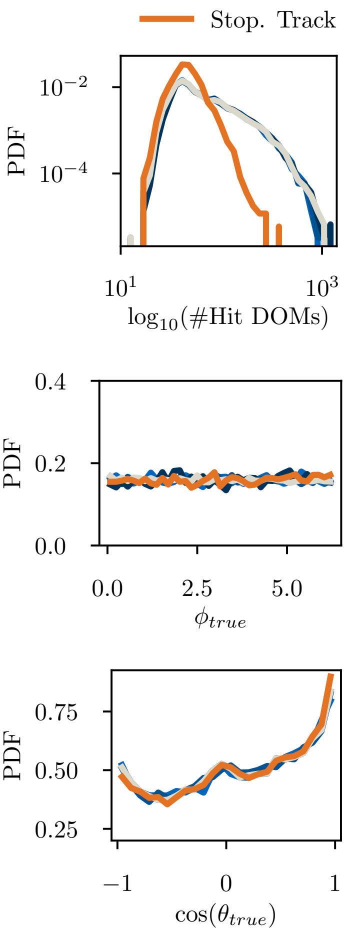

Figure 5.9: The distribution of event quantities in the training dataset. From top to bottom the number of hits, the azimuth and the zenith are shown. As discussed in the text the distributions are equal between the classes.

energy, which is not well defined for skimming events, the number of hit DOMs is a well-defined quantity for all event topologies and therefore preferred in this application. Figure 5.9 shows the resulting event distributions. An exception is visible for stopping tracks that follow a different distribution in the number of hit DOMs. The reason is that stopping tracks are by construction a lower energetic version of through-going tracks and therefore can not be perfectly aligned with the other classes. As a compromise, the distribution has been flattened as far as possible without losing event statistics. The final dataset contains 14’716’253 events in total ($\sim$2.9 million events per class) of which 90% are used for training and 10% for validation.

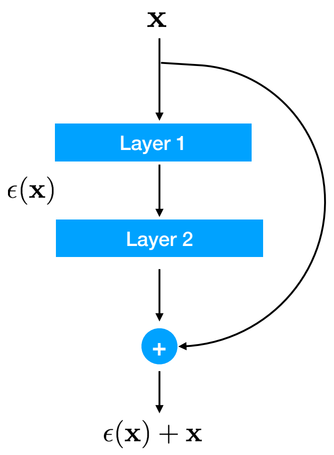

Figure 5.10: A schematic view of a residual unit. Starting from an input $\mathbf{x}$, a set of layers learns a mapping $\mathbf{x} \to \epsilon(\mathbf{x})$ and produces the output $\mathbf{x}' = \mathbf{x} + \epsilon(\mathbf{x})$.

5.5 CNN Architecture and Training Procedure

After the breakthrough of convolutional neural networks (CNNs) in image classification,[32] several authors have tried to further improve these networks by increasing their depth[33].[34] However, with increasing depth the gradient in the backpropagation becomes very small, a problem known as *vanishing gradient problem*.[35] Although this has been largely solved by adding normalization layers, it turned out that going deeper causes further problems. While the training accuracy starts saturating when adding more and more layers to a shallow network, it rapidly starts degrading when passing some threshold.[36] It has been shown that this effect is not caused by overfitting but rather by optimization problems. In fact, a deeper network should never perform worse than a shallow one, as the additional layers could always learn to perform a trivial identity mapping if no further improvement is possible. Based on this property a *residual structure* has been suggested.[37] The idea is to have shortcut connections and use each (set) of layers only to learn additional, residual features. This is schematically illustrated in Figure 5.10. It has been shown that residual architectures can go much deeper and thereby drastically improve performance while also being straightforward to optimize.

A general challenge for pattern recognition tasks is that the same structures can appear in different sizes. In the case of an IceCube event, this can be, for example, the difference between a small, low energy, and a large, high energy cascade. Therefore, it is not a priori clear which filter size one should use for the convolutions in each level of the network. A convenient solution for this has been presented by Google in a series of paper, introducing the *inception layers*.[38] Inception layers combine filters of varying sizes in one network level, thereby leaving it to the training procedure to optimize the kernel size that suits the problem best. While this increases flexibility in the kernel size, it adds free parameters to the model, increasing the risk of overfitting. This problem can be counteracted by another technique that reduces the number of free parameters in convolutional layers. The idea is to split up a $n \times n \times n$ filter into three filters of size $1 \times 1 \times n$, $1 \times n \times 1$ and $n \times 1 \times 1$. This reduces the amount for free parameters from n^3 to $3n$, equivalent to a reduction

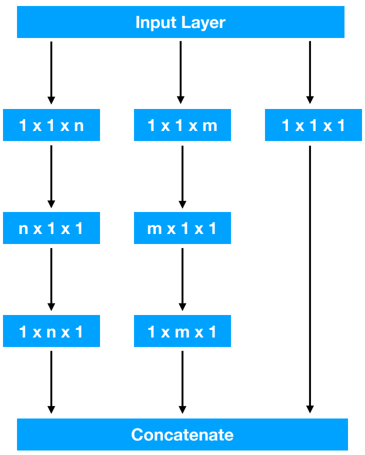

Figure 5.11: A schematic view of an inception unit. The first layer is the input for three branches with different filter sizes. In the last step, the branches are concatenated along the feature axis.

[32] Krizhevsky, Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks”.

[33] Ioffe and Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift”.

[34] Simonyan and Zisserman, “Very deep convolutional networks for large-scale image recognition”.

[35] Glorot and Bengio, “Understanding the difficulty of training deep feedforward neural networks”.

[36] He and Sun, “Convolutional neural networks at constrained time cost”.

[37] He, X. Zhang, et al., “Deep residual learning for image recognition”.

[38] Szegedy et al., “Inception-v4, inception-resnet and the impact of residual connections on learning”.

by $100 - (300/n^2)$ percent for $n \geq 2$. For the maximum filter size of the classifier presented in this chapter, $n = 10$, this approach reduces the number of free parameters by 97%. A sketch of the resulting inception architecture is shown in Figure 5.11.

Finally, both concepts can be combined to Inception-ResNet units that keep the general idea of inceptions layers while adding a residual connection. The DNN classifier presented in this work is a combination of pure inception units at the beginning to extract the most relevant features, followed by a large number of Inception-ResNet units to increase the depth. In between, average pooling layers are used to shrink the size of the feature maps (the detector) and thereby reduce the complexity in the subsequent layers. A sketch of the network is shown in Figure 5.12 and the corresponding Keras[39] implementation in appendix B.1. In contrast to the early developments of CNNs — where the convolutions were usually followed by fully-connected layers — later studies have shown that classifiers are more stable when using *global average pooling* before the output.[40] Therefore the DNN classifier uses 4096 $1 \times 1 \times 1$ convolutions to efficiently combine the features in each part of the detector before connecting everything with a global average pooling layer to the five output nodes corresponding to the five event topologies. In contrast to the rest of the network, which uses *ReLu* activations, the final layer uses a *softmax* activation to produce 5 outputs that are interpretable as probability scores (p-score) for each event class. Overall, the final network has a total of 1,968,581 parameters.

Training of the DNN classifier is conducted on 4 *Nvidia GeForce GTX 1080* in parallel, facilitating a mini batch size of 148 events. As optimizer *Adam* is used with standard settings.[41] Compared to standard gradient descent algorithms, Adam automatically controls the learning rate by calculating a moving average of the gradient's first and second moment. The overall training of the DNN takes around ~ 30 days. In this time, the network has processed a little more than 2 full epochs, or 188924 mini-batches. The training accuracy and loss as function of evaluation steps (20/epoch) are shown in Figure 5.13.

The final network can be used to make predictions on unseen data. While the training of the DNN is slow, the application to data is fast, specifically when compared to likelihood-based algorithms. Figure 5.14 shows the expected pre-processing and prediction times on a representative IceCube level 2 dataset. It can be seen that most of the total processing time for an event is required for the transformation from the pulsemaps to the DNN input (on average $\sim 100\,\text{ms}$). The prediction itself, on the other hand, only requires on average around $\sim 10\,\text{ms}$ on a GPU. To make perfect usage of the resources, it is therefore favorable to run the processing in parallel on several CPUs while making the predictions of event batches on a single GPU.

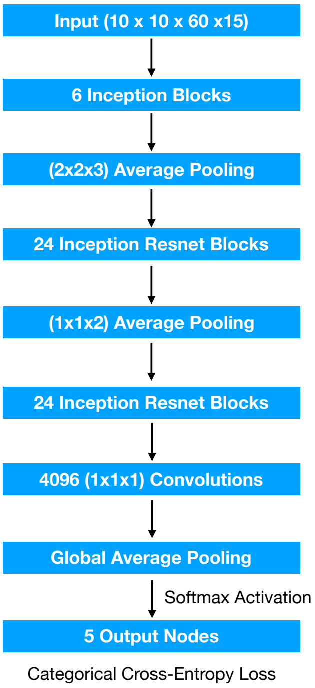

Figure 5.12: Schematic view of the DNN classifier's architecture. A pseudo-code implementation of the model definition can be found in appendix B.1.

[39] Chollet et al., *Keras*.

[40] Lin, Chen, and Yan, "Network In Network".

[41] Kingma and Ba, "Adam: A Method for Stochastic Optimization".

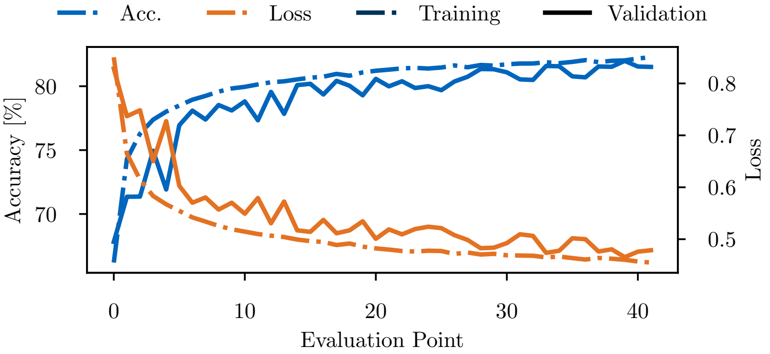

Figure 5.13: The learning curve of the DNN classifier as a function of the evaluation point. One evaluation point thereby equals to 1/20th of an epoch. In addition to the cross-entropy loss, also the accuracy, i.e., the fraction of correct predictions is shown on the training (dashed-dotted) and validation dataset (solid).

5.6 Performance Figures

In this section, the classifier's performance is evaluated based on a representative IceCube level 2 NuGen Monte Carlo dataset [42].[43] As level 2 is the first level of offline processing, it provides the largest variety of events and, therefore, a good benchmark scenario. Furthermore, the development of a new, DNN-based, level 3 processing is one of the major applications and long-term goals of the DNN classifier. In this regard, the following sections can be understood as a feasibility study for further development in this direction. In the following we assume a HKKM conventional atmospheric flux model[44] and the best-fit astrophysical powerlaw flux from IceCube's cascade analysis[45] ($\gamma = 2.53 \pm 0.07$ and $\phi_{astro} = 1.66^{+0.25}_{-0.27}$ at $E_0 = 100\,\text{TeV}$) with a 1:1:1 flavor ratio.

Prediction results are frequently presented in the form of *prediction matrices* as shown in Figure 5.15. Every entry in the confusion matrix shows the fraction of events with a given truth to be classified with a specific label under the normalization condition given in the respective title. Hereby the *classification* of an event is defined as the class with the maximum p-score p_{max}. From Figure 5.15 it can be seen, for example, that 81% of true through-going tracks are correctly identified while 9% are predicted as stopping track. Of the predicted through-going tracks, on the other hand, 75% are correctly classified, corresponding to an overall rate of 1.3 mHz. Note that Figure 5.15 only shows the neutrino flux and neglects the atmospheric muon contribution, as is usually the case when only considering the Northern Hemisphere, where events have to pass through the Earth before reaching the detector. A prediction matrix including the atmospheric muon tracks is shown in Figure C.2 and Figure C.3 in the appendix.

The prediction matrix in Figure 5.15 shows major entries on the diagonal, implying that event classes are generally correctly predicted. The diagonal elements of the middle plot are referred to as the **accuracy**, i.e., measuring the fraction of correct predictions among the total predictions for a specific (predicted) class. Off-diagonal elements,

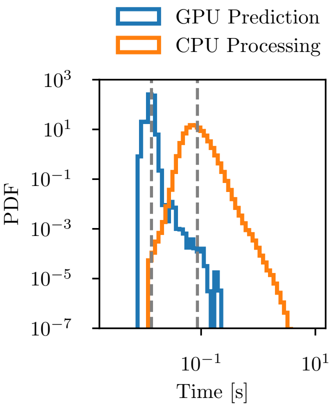

Figure 5.14: The CPU processing time and the GPU prediction time for a representative IceCube level 2 dataset. Dashed vertical lines show the average values.

[42] NuGen is an IceCube specific implementation of the ANIS neutrino generator.

[43] Gazizov and Kowalski, "ANIS: High energy neutrino generator for neutrino telescopes".

[44] Honda et al., "Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data".

[45] M. Aartsen et al., "Characteristics of the diffuse astrophysical electron and tau neutrino flux with six years of IceCube high energy cascade data".

in contrast, are mispredictions. Specifically, we use the term **confusion** for true events that are falsely predicted into another class (left plot). Note that, while the accuracy depends on the overall abundance between the classes, the confusion does not, e.g., 22% of all starting tracks get predicted as cascades, which makes 48% of the entire cascade predictions. Several aspects influence the accuracy and confusion, e.g., the event's energy and direction, the p-score threshold, the geometry of the detector, and particle physics effects like the inelasticity, i.e., the fractional energy transfer to the cascade. While the efficient combination of the DNN output with other observables helps to increase the accuracy further and generate high-purity samples, some confusions are unavoidable and expected, as shown in the following sections.

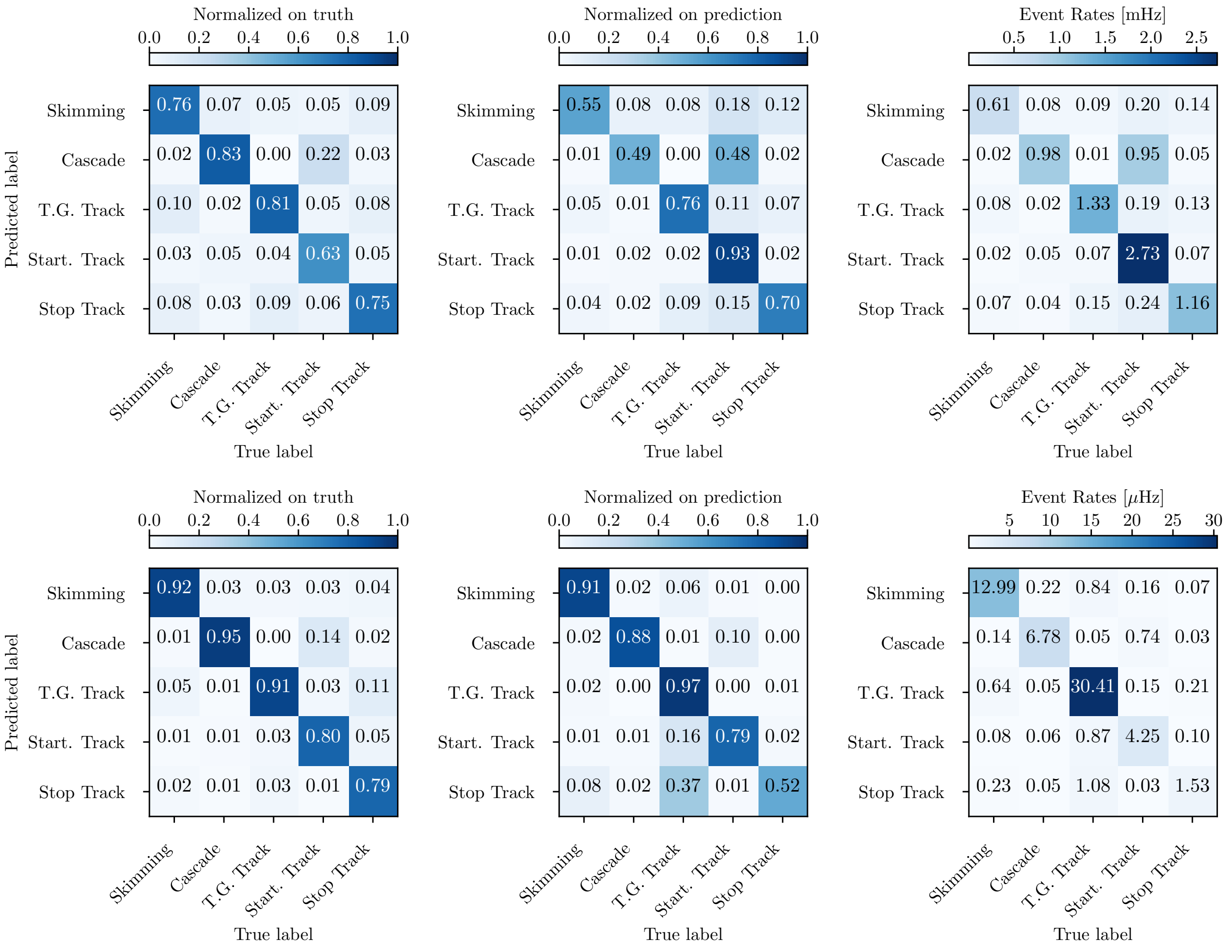

Figure 5.15: Prediction matrices for a representative level 2 neutrino Monte Carlo dataset and neutrino energies above 100 GeV (top) and 10 TeV (bottom). From left to right the classification performance is shown normalized on the truth, the prediction, and in total event numbers per year, respectively.

5.6.1 P-Score & Energy

As discussed in section 5.1, the p-score is a measure for the probability of a given event to belong to the predicted event class. However, due to the training on a biased dataset, there is no trivial correspondence between p-score and accuracy when applying the DNN to an experimental event sample. Yet, it is possible to perform a Monte Carlo-based calibration of the p-score for a given event selection and given flux assumptions. This is illustrated in Figure 5.16, showing the level 2 accuracy as a function of the maximum p-score, p_{max}, for two energy thresholds. While the lines do not align between the classes, it is evident that *the accuracy increases with larger* p_{max}, which is a minimal requirement for the classifier's consistency. Note also the difference between the top and the bottom panel of Figure 5.16. While starting tracks have higher accuracies when considering the full energy range, the prediction of through-going tracks is more accurate for neutrino energies above a few TeV. This can be understood as an effect of the shift in overall event rates, i.e., the larger contribution of through-going tracks at higher energies.

With increasing event energy there is also more information on the event topology. Hence we do not only expect the accuracy to increase with p_{max}, but also *the average* p_{max} *to increase with the energy*. This effect is shown in Figure 5.18. While the average p_{max} at a neutrino energy of 1 TeV is only $\sim 99\%$ it goes up to 99.999% at 1 PeV. This result has important implications for the usage of the DNN classifier in event selections. In particular, one can place more stringent purity cuts for high-energy events before losing in the efficiency of the selection, i.e. the fraction of selected events given the total number of events from a specific class.

All these effects are exemplarily summarized in the receiver operating characteristic (ROC) curve for starting tracks in Figure 5.17. It shows the true positive, i.e., the fraction of correctly identified starting tracks given all starting tracks, against the false positive, i.e., the fraction of non-starting tracks that get predicted as starting tracks. Evidently, the same false positive rate leads to a higher true positive rate when the event's deposited energy is larger. Furthermore, tighter $p_{starting}$ cuts are required for low energy events to reach the same level of background suppression while also removing more of the signal. Overall, Figure 5.17 shows the good separation power between starting tracks and other classes, motivating further studies for a dedicated starting track level 3 branch in the IceCube data pipeline as discussed in section 5.8.

A comprehensive summary of the energy-dependent performance of all event classes is given by Figure 5.19. For each topology, the accuracy and the confusion are shown as a function of the energy and two different p_{max} thresholds. Evidently, the confusion decreases for all classes with increasing energy or by applying a larger p_{max} threshold. In contrast to the confusion, the accuracy does not have such a clear tendency. For most of the classes, it is rather flat or even decreases above a few TeV. The explanation for this apparent contradiction is twofold: 1.) although the overall confusion of through-going and starting tracks is quite small, the overall dominance in event numbers strongly contaminates the other classes, 2.) the detector resolution and geometry, as well as physical effects are limiting the distinguishability of some classes at specific energy scales. The

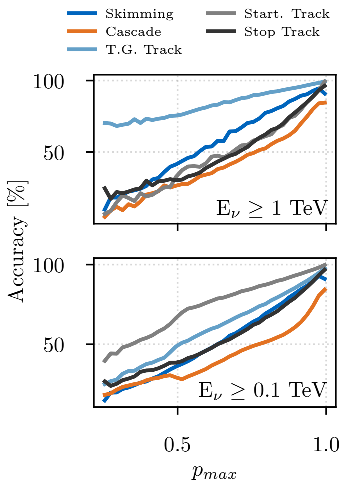

Figure 5.16: Accuracy of the different event topologies as a function of the maximum p-score. Top and bottom panel show two different energy thresholds, respectively.

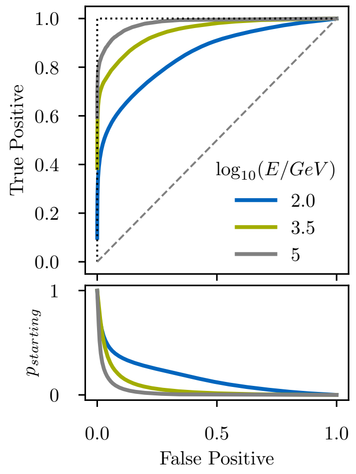

Figure 5.17: Receiver operating characteristic (ROC) curve for starting tracks at three different deposited energies. The top panel shows the true positive against the false positive, calculated from a sliding threshold on the starting track score (bottom panel). The black dotted line indicates a perfect classifier, while the grey dashed line represents a classifier without discriminating power.

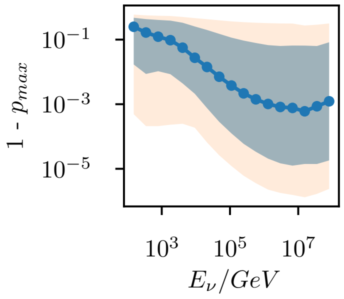

Figure 5.18: Maximum p-score as a function of the initial neutrino energy. The bands show the 68% and 90% quantiles, respectively.

latter is especially evident when comparing Figure 5.19 with Figure 5.15, where it can be seen that, independent of the energy, starting tracks are confused with cascades, through-going events with starting and stopping events, and basically all event classes with skimming events. In the next section, we will further analyze these confusions and reconcile them with physical expectations.

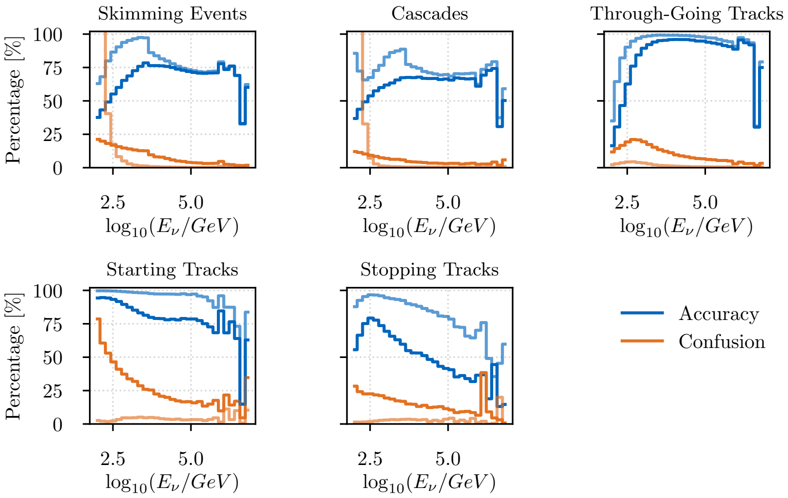

Figure 5.19: Accuracy and confusion (blue and orange) against the initial neutrino energy for $p_{max} \geq 0.20$ (all events) and $p_{max} \geq 0.95$, shown in dark and light colors, respectively. The name of the classes is given in the title. Note that accuracy refers to the fraction of events in a given class being correctly identified, while confusion refers to events with the respective true class to be predicted into any other class. This also implies that the lines do not add up to 100%.

5.6.2 *Effects of Detector Geometry & Particle Physics*

In a perfect world with an infinitely large detector and perfect resolution, one could train a neural network to arbitrarily high accuracy. In the real world, however, IceCube is limited in size, photons are scattered and absorbed, and DOMs have a non-negligible spacing. Altogether, this implies that some confusion is unavoidable. Based on Figure 5.19 and Figure 5.15 the major confusions are discussed in the following

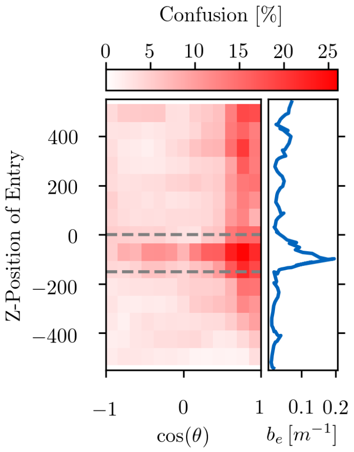

Figure 5.20: Left panel: The fraction of true through-going tracks that get falsely classified as starting tracks as a function of the zenith angle and entry point in the detector. Right panel: The scattering coefficient as a function of the depth. The dust layer is framed by the dashed grey lines.

- **Through-Going Tracks vs. Starting Tracks**: From the difference in the prediction matrices above 100 GeV and 10 TeV (Figure 5.15), it is evident that mis-classification between these classes mainly occurs at low energies. This is not surprising as low energy muons can sneak through the outer detector layers without leaving a visible trace and appear to be starting. Furthermore, the primary cascade, which usually indicates a starting event, is barely visible. A more detailed study of the direction and position of the mis-classified events in Figure 5.20 reveals additional effects that are related to the ice properties and the event position and direction in the detector. In fact, the ice in the detector is not uniformly clear. While it is very clear on the top and bottom, the dust layer in a depth between ∼ 1950 m and ∼ 2150 m significantly increases the scattering and absorption of photons.[46] This explains the band of high confusion in the middle of the detector visible in the middle of Figure 5.20. While confusion is generally enhanced in the dust

[46] M. G. Aartsen et al., "Measurement of South Pole ice transparency with the IceCube LED calibration system".

layer, it is more pronounced for down-going events ($\cos(\theta) > 0$). This is coherent with the IceCube DOMs facing downwards, therefore being blind to the direct light from events passing the DOM from the top. As a result, confusion is enhanced for down-going events independent of the entry position along the z-axis.

- **Cascades vs. Starting Tracks**: Starting tracks differ from cascades by an outgoing muon track. The energy of this muon is, however, not fixed but connected to the energy of the initial hadronic cascade E_{had} through the inelasticity, which is defined as

$$y = \frac{E_{had}}{E_\nu} = \frac{E_\nu - E_\mu}{E_\nu} = 1 - \frac{E_\mu}{E_\nu}. \tag{5.16}$$

Consequently, we find $y = 0$ if all the energy is transferred to the muon and $y = 1$ if all the energy is transferred to the hadronic cascade. In the latter case, the starting track gets physically indistinguishable from a cascade. Besides, for values very close to 1, the muon energy can be insufficient for the muon track length to exceed the extension of the cascade and/or create a visible amount of hits. Figure 5.21 shows the confusion of starting tracks with cascades against the energy and the inelasticity. In alignment with the previous discussion, the confusion increases strongly with growing inelasticity (more cascade-like) and decreases with the overall deposited energy (more event information).

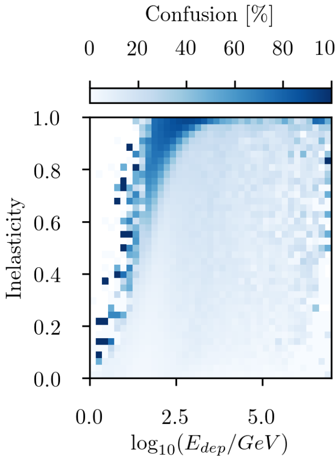

Figure 5.21: The confusion of true starting tracks with cascades as a function of deposited energy and inelasticity.

- **Cascades vs. Skimming Events**: Cascade-like topologies can appear with primary vertex inside and outside the detector. While the former are grouped into (starting) cascades, the latter are included in the skimming class. Figure 5.22 shows the vertex position of the true cascade events that are falsely classified as skimming. As expected, the confusion increases strongly the closer the vertex is to the detector hull. On the contrary, it can be seen that the confusion quickly drops to zero once the vertex is only a few meters inside the detector.

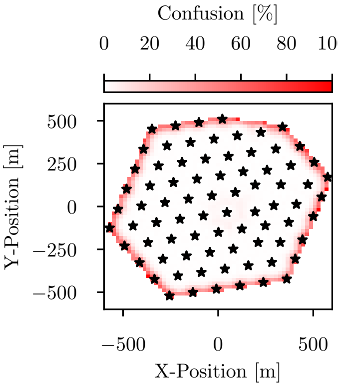

Figure 5.22: The confusion of cascades with skimming events against the interaction vertex. Only events entering from the side ($|z| < 450$ m) are considered for the plot. The black stars mark the position of the IceCube In-Ice Array.

- **Stopping vs. All Other Classes**: The maximum distance a muon can travel through the IceCube detector is ~ 1500 m. According to the energy loss equation for muons, eq. (3.11), and the resulting Figure 3.4, stopping muons therefore have a typical energy $\lesssim 300$ GeV on detector entry. During propagation, the muon keeps losing energy through ionization and thereby only produces very few hits along its trajectory. As a result, stopping tracks with a large track length are barely distinguishable from through-going tracks, as seen in Figure 5.23. For the maximal track length of around 1500 m the confusion goes up to around 50%, i.e., stopping and through-going tracks are predicted with the same probability. In the other extreme, stopping tracks with very short track length in the detector can appear as skimming. This effect strongly increases for track length below ~ 250 m, corresponding to only two strings in horizontal distance.

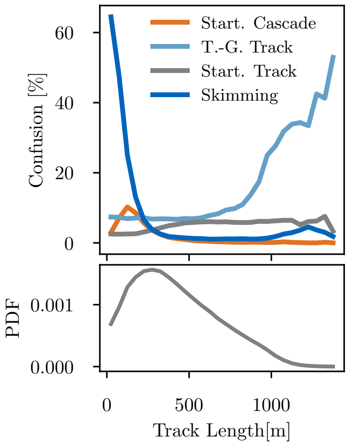

Figure 5.23: The confusion of true stopping tracks with the other topologies as a function of the muon track length in the detector. The bottom panel additionally shows the expected distribution of track length.

In summary, this section has shown that the origin of the classifier's confusions complies with physical expectations. This is especially important as the input to output mapping of deep neural networks is often referred to as a 'black-box', meaning that it is difficult to understand *why* a certain prediction was made. Having a coherent picture of the confusion proves, however, that the network has learned useful features independently of whether they are intuitively accessible for humans or not. Finally, understanding the classifier's confusions helps in the development of event selections, as it gives a guideline on how cuts should be placed for maximal efficiency. Assume, for example, the case of a starting track selection. Given the enhanced confusion in the dust-layer and the down-going region, a starting track selection would profit from p-score cuts that depend on the the events' direction and entry point. Such an approach will be further discussed for the new starting track selection in chapter 6.

5.7 Data/Monte Carlo Agreement and Systematic Checks

For the training of the DNN classifier a baseline Monte Carlo has been used which incorporates the best knowledge about the photon propagation in the ice and the detector response. Nevertheless, the Monte Carlo is only an approximation of the real experimental data. To further validate the performance of the DNN classifier on the experimental data, its predictions on a given data sample can be compared to the ones on a corresponding Monte Carlo dataset [47].

Due to the large variety of event topologies observed by IceCube and the uncertainty on the underlying fluxes, the agreement between data and Monte Carlo is, however, not optimal at processing level 2. Hence this early stage of the data pipeline is not suited to evaluate the stability of the DNN classifier. In contrast, the data to Monte Carlo agreement is usually required to be good on the final analysis level. Therefore the DNN classifier is exemplarily tested on the data sample of up-going (declination $\delta \geq -5°$) tracks used for the measurement of the diffuse astrophysical muon-neutrino spectrum[48] [49] which has a good data/MC agreement [50]. The results of this study are shown in Figure 5.24. It can be seen that, independently of the energy threshold, the p_{max} distribution is very well aligned between data and Monte Carlo. While this study is thereby a good indicator that the DNN predictions are trustworthy in the entire IceCube energy range, the check, in any case, needs to be redone for any new selection, as exemplarily shown for the case of the improved high-energy starting event (HESE) selection in chapter 6. In addition to the baseline data/MC agreement, chapter 6 also discusses the stability of the DNN against modifications of the pulse arrival times, systematic variations of the ice, and the detector efficiency.

[47] which should be independent of the training dataset.

[48] Stettner, "Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data".

[49] diffuse muon-neutrino sample

[50] See also Figure 5.27 and section 7.3 for more details on the selection

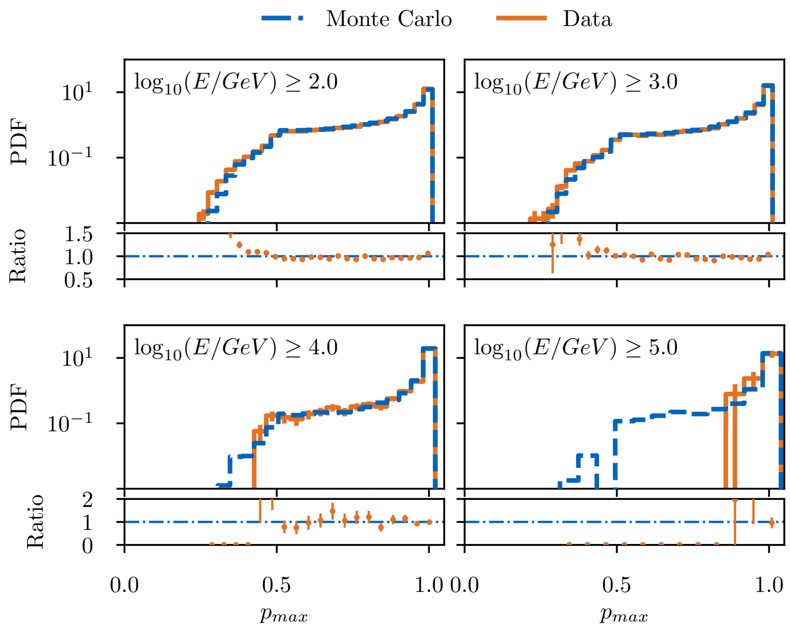

Figure 5.24: Data/Monte Carlo agreement of the DNN classifier on the diffuse up-going muon-neutrino sample for one year of data. From top-left to bottom-right the threshold in reconstructed energy increases from 100 GeV to 100 TeV. For each plot, the top panel shows the distribution of p_{max} and the bottom panel the ratio between data and Monte Carlo. Above 100 TeV event statistics are low and nearly all the events end up in the highest p-score bin.

5.8 Possible Applications

Event type classification is an important task in any stage of the IceCube data processing pipeline. Despite the importance, IceCube currently does not have an algorithm that provides a direct mapping from the pulses to the event topology in place. Therefore, most event selections are based on cuts combining a broad range of generic event variables (e.g., the center of charge) and reconstruction quantities (e.g., fit quality). In contrast to that, the DNN classifier provides a new framework for event type classification that combines performances and speed and can be used early in the processing chain. In Figure 5.25 a schematic view of the IceCube data pipeline, from online processing at the South Pole to final level selections, is shown. Marked in yellow are all the processing steps where the DNN classifier could be useful, although in many cases further studies are needed to quantify this. The following list is hence only an incomplete collection of ideas for a future, DNN-based processing pipeline

- The goal of the IceCube **filters** is the reduction of the data rate to the capacity of the satellite bandwidth while keeping the relevant events for physics analysis. Hence only events with clear topologies and good reconstruction quality should be selected. Two of IceCube's major filters are the cascade and the muon filter, with rates of around $\sim 40\,\mathrm{Hz}$. Without any topological classifier in place, there is, however, a large overlap between those streams. With the help of the DNN classifier, in contrast, topological information can be combined with the respective reconstructions. Subsequently, filter decisions are then based on a combination of

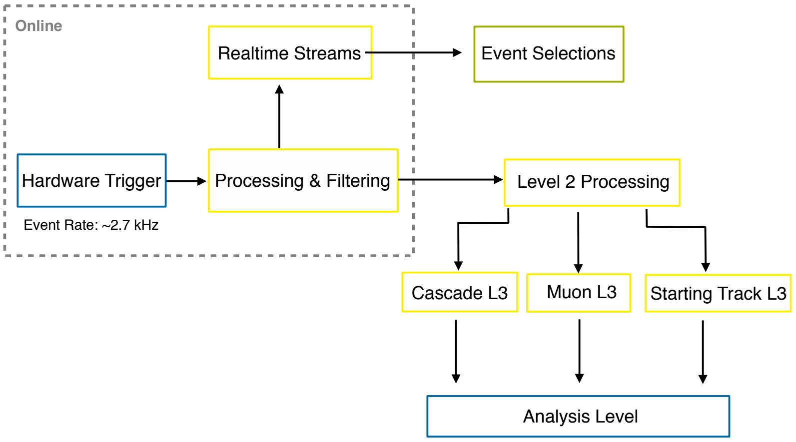

Figure 5.25: A sketch of the IceCube data processing pipeline. Processing steps that could profit from the DNN classifier are shown in yellow. The real-time event selection is already partially DNN-based and therefore marked in green. All steps included in the grey box are done *online*, i.e., using the computing facilities at the South Pole.

classification and reconstruction to efficiently select events and reduce satellite bandwidth [51]. Interestingly, the DNN p-score also provides an estimate of the reconstruction quality of the event, as shown in Figure 5.26 for through-going muons. The larger the through-going p-score, the smaller the opening angle between reconstruction and true muon direction. This is consistent with the picture of a high p-score indicating a clear event topology.

- IceCube's **realtime stream** is a prime example for a selection that requires fast and reliable reconstruction algorithms. With intended delays on the scale of seconds between detection at Pole and sending out a notification, it is unfeasible to run high-level likelihood reconstructions before deciding if an event is interesting or not. Using a combination of DNN algorithms, it is possible to reach a high performance without requiring several minutes of computation. As an example, a newly developed realtime selection of highly-energetic starting tracks (HESE) is presented in chapter 6. Compared to the previous version, it improves both efficiency and accuracy by solving problems that appear when trying to classify an event's topology on a short time scale, using (potentially) unsuitable reconstruction algorithms. Technically, the DNN realtime selection can be run either online or offline. However, there are currently no GPUs available at the South Pole. Hence, it is currently applied offline at the IceCube data center in the North. In addition to the new starting track selection, also a newly developed HESE realtime cascade selection is using the DNN classifier to identify interesting events [52].

- After the online filtering the data is transmitted to the IceCube data center for further processing. The major goal of the following **level 2/level 3** processing is the removal of events that are not useful for further analyses. Furthermore, the stream is split up into a cascade and muon level 3 stream, which is the starting point for all high-level analyses. The splitting is important as different

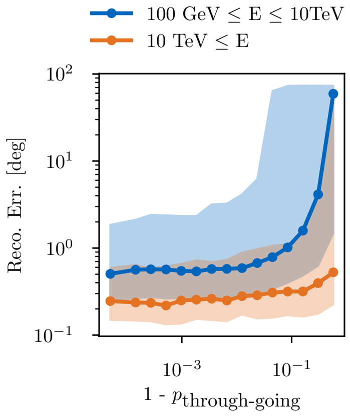

Figure 5.26: The opening angle between true muon direction and the *SplineMPE* reconstruction for two different deposited energy ranges versus the through-going p-score. Solid lines show the median values and the shaded bands the 68% quantiles. It can be seen that the resolution improves with the p-score, implying that it is a proxy variable for the reconstructability of a track.

[51] Which is even more important for future neutrino observatories as IceCube Gen-2 that will increase event rates significantly

[52] https://gcn.gsfc.nasa.gov/amon_icecube_cascade_events.html, accessed at 6th October 2020

event topologies can be more suitable for a specific analysis than others and each topology has its own reconstruction algorithms. In alignment with the five output branches of the DNN classifier, one could extend the current level 3 with an additional branch of starting tracks, and potentially stopping tracks and skimming events. Such a level 3 structure creates more flexibility in the subsequent combination of event streams to construct optimized final level samples. The performance discussions in section 5.6 is a first step into the direction of a deep learning-based level3 selection.

- To optimize the fraction of signal events, **final level selections** apply straight cuts or (boosted) decisions trees (BDTs) to remove remaining background events from level 3. As a performance check, the DNN classifier has been applied to the previously mentioned sample of up-going tracks.[53] The selection uses two BDTs to identify well-reconstructable up-going tracks and to remove the cascade background, respectively. The DNN prediction in Figure 5.27 shows a good agreement of the DNN classifier with the BDT-based selection. While more than 99% of the true tracks are also predicted as track, also 24% of the very small cascade component is identified. Albeit this is not a proof for the DNN to outperform the BDT itself, it indicates that it could be useful to add the five DNN p-scores as input to event selection BDTs. In this way, connections among the p-scores and synergies with other variables as direction and energy can be exploited optimally.

Figure 5.27: Confusion matrix for the final level up-going muon-neutrino sample. The respective normalization conditions are given in the titles of the plots. From the right plot — showing the absolute rates — it is evident that $\gtrsim 99\%$ of true tracks are also predicted as tracks. The left plots further shows that 24% of remaining cascade contamination is correctly identified.

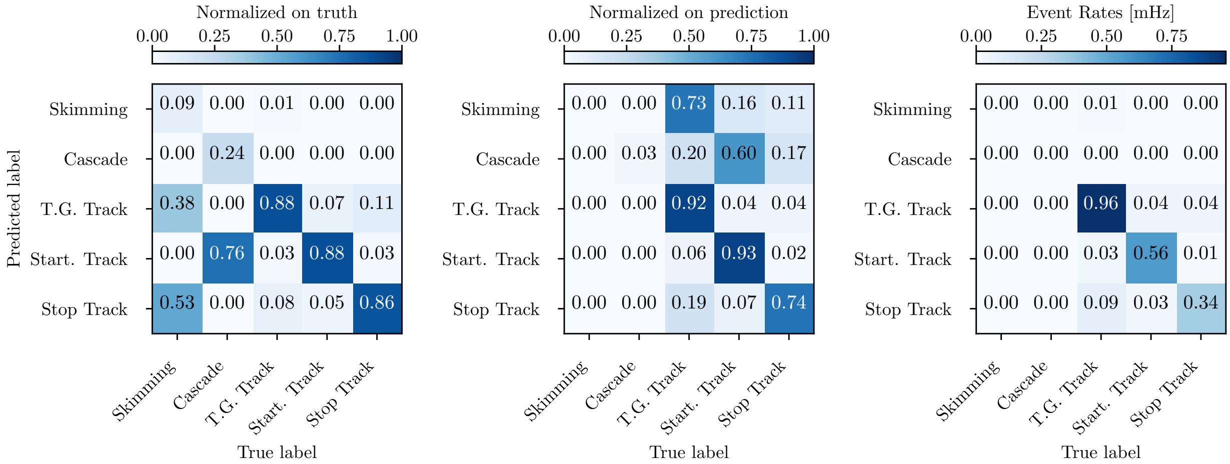

5.8.1 IceTray Module

In order to conveniently add the DNN Classifier to any IceCube data processing pipeline, an IceTray Module is available on the IceCube Open Source Github [54]. It is implemented in Python 3 and backward compatible with Python 2.X. After downloading the package, the module can be added to an existing IceTray via the following lines of python code:

---

[53] Stettner, "Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data".

[54] https://github.com/IceCubeOpenSource/i3deepice

```
from i3deepice.i3module import DeepLearningModule
tray.AddModule(DeepLearningModule,
               'dnn_classification',
               batch_size=128,
               cpu_cores=4,
               gpu_cores=1,
               model='classification',
               pulsemap='SplitInIcePulses',
               save_as='TUM_classification')
```

The batch size, as well as CPU and GPU cores are set here to default values, which should be varied depending on the respective hardware availability. In general it is advantageous for an optimal exploitation of resources to use several CPUs for the pre-processing and only one GPU for the prediction (see Figure 5.14).

Optionally, saturation windows, bright DOMs, calibration errata & saturated DOMs can be removed by giving the name of the respective keys in the I3Files. To apply the same DNN with different settings, all the previously mentioned options, as well as the pulsemap can also be a list. The following snippet applies the DNN twice, including and excluding pulses in saturation windows, respectively.

```
from i3deepice.i3module import DeepLearningModule
tray.AddModule(DeepLearningModule,
               'dnn_classification',
                batch_size=128,
                cpu_cores=4,
                gpu_cores=1,
                model='classification',
                saturation_windows = ['None',
                                      'SaturationWindows'],
                pulsemap=['SplitInIcePulses',
                          'SplitInIcePulses'],
                save_as=['TUM_classification_0',
                         'TUM_classification_1'])
```

Finally, note that it is also possible to use the tool on a CPU-only machine. In contrast to GPU machines this is, however, increasing the prediction times significantly, up to $\mathcal{O}(s)$ for a single event.

6 Realtime Selection of Highly-Energetic Starting Tracks

With the successful follow-up campaign of IceCube-170922A and the subsequent discovery of TXS 0506+056 as the first compelling source of astrophysical neutrinos,[1] it has become increasingly evident that realtime multi-frequency follow-ups can contribute crucial information to the identification of the sources of astrophysical neutrinos. Since telescope times are limited, it is important to have high-quality triggers that select IceCube events with good pointing and a high probability of being of astrophysical origin. IceCube currently has three different alert streams: High Energy Starting Events (HESE), Extremely-High Energy Neutrinos (EHE), and a selection of events originally developed for gamma-ray follow-ups (GFU).[2] In contrast to the latter two, which are optimized for through-going tracks, the HESE stream requires the primary event vertex to be inside the detector volume. While the event statistics of through-going tracks is larger, HESE events have the advantage of good all-sky coverage and a higher average signalness, due to the efficient reduction of atmospheric background fluxes. A sketch of the IceCube realtime pipeline is shown in Figure 6.1.

This chapter presents a major update to the current realtime HESE v2 starting track selection — HESE v3. In contrast to the old selection, the new version is entirely based on deep learning algorithms. At the core of the new selection, the DNN classifier presented in chapter 5 is used to identify promising starting tracks. [3] The new selection is a prime example for the advantages of deep learning-based classifications compared to standard selection approaches. Specifically, the output of the DNN classifier can be used to make clean cuts on the event topology, whereas the previous selections rely on goodness-of-fit parameters of reconstruction algorithms (as *Spline MPE*) that are not optimized for starting tracks in the first place. In short, HESE v3 has the advantage of first running the classification, followed by the corresponding reconstruction.

[1] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A”.

[2] Blaufuss et al., “The Next Generation of IceCube Realtime Neutrino Alerts”.

[3] Complementary to that also a new HESE cascade realtime stream has been implemented using the same classifier; https://gcn.gsfc.nasa.gov/amon_icecube_cascade_events.html.

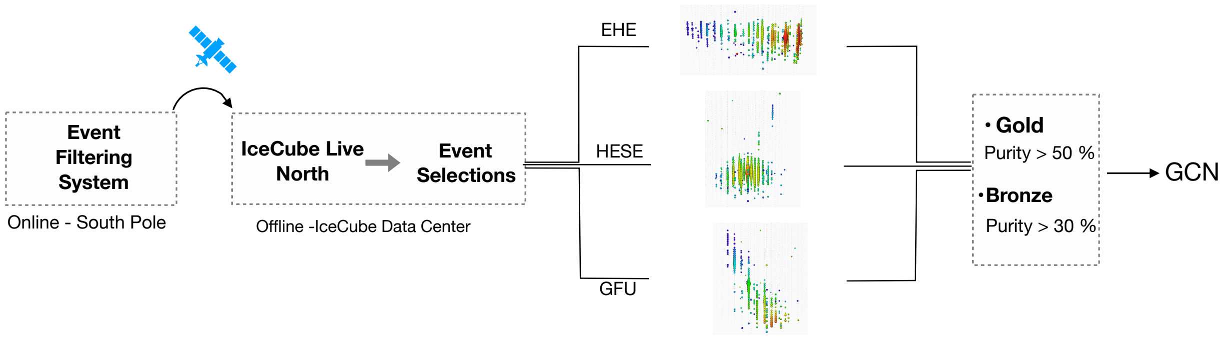

Figure 6.1: Sketch of IceCube's realtime pipeline. Events that pass any of the online filters are send to the IceCube data center via an Iridium satellite for further processing. If an event further passes any of the realtime selections, it is reconstructed and the *signalness* is determined. Finally, a summary of the event information is send to he Gamma-ray Coordinates Network (GCN). IceCube internally, this also triggers improved reconstructions algorithms for a refined direction and error contour that are finally published in a GCN Circular.

6.1 *Selecting HESE Tracks for the Realtime Stream*

On the lowest processing level, IceCube filters provide the first prediction of an event's topology. As those filers run online at the South Pole, they only apply simple and thereby computationally cheap algorithms. The HESE filter, for example, selects all events with a total of more than 6000 photoelectrons (p.e), of which less than 3 p.e of the first 250 p.e should be recorded in the veto region as shown in Figure 6.2. Being a robust pre-selection of potentially interesting astrophysical events, it is used as the starting point for the HESE realtime selections.

For the studies in this chapter, a dedicated HESE Monte Carlo is used, containing both neutrinos and atmospheric background muons. For the neutrinos we adopt the HESE 7.5 year best-fit flux with $\phi^{astro}_{\nu+\bar{\nu}} = 6.37 \cdot 10^{-18} \cdot (E_\nu/100\,\mathrm{TeV})^{-2.87}\,\mathrm{GeV}^{-1}\,\mathrm{s}^{-1}\,\mathrm{cm}^{-2}\,\mathrm{sr}^{-1}$ for the astrophysical component,[4] and a conventional atmospheric model base on the Honda parameterization.[5] For the steeply-falling atmospheric muon contribution, a dedicated *MuonGun* simulation is used to reach reasonable event statistics after the HESE filter. The input spectrum is modeled using a CORSIKA simulation[6] based on the Hillas-Gaisser H4a cosmic-ray model[7] and a SIBYLL 2.1 hadronic interaction model.[8] While this gives the correct shape of the spectrum, the normalization is highly uncertain. To estimate the normalization of the muon background contribution HESE uses a muon tagging strategy. Specifically, the passing fraction through the veto is estimated by muons that leave a trace in the veto's outer layer, but are not visible in the inner veto layer. This strategy suggests a scaling factor of 2.1, well consistent with the DNN predictions in Figure 6.11. The resulting event stream is composed mainly of cascades [9] (8.92 / year) and starting tracks (3.58 / year), and a smaller contribution of atmospheric background muons ($\sim$ 1.86 / year). Note that the rejection of incoming muons also reduces the number of atmospheric neutrinos as they are inherently accompanied by an atmospheric shower. This is called the atmospheric *self-veto*[10].[11]

In general, if an event passing any IceCube filter is potentially interesting for realtime follow-ups, additional reconstruction algorithms are run. This is specifically the case for everything that passes

[4] Schneider, "Characterization of the Astrophysical Diffuse Neutrino Flux with IceCube High-Energy Starting Events".

[5] Honda et al., "Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data".

[6] Heck et al., "CORSIKA: A Monte Carlo code to simulate extensive air showers".

[7] T. K. Gaisser, Stanev, and Tilav, "Cosmic Ray Energy Spectrum from Measurements of Air Showers".

[8] Ahn et al., "Cosmic ray interaction event generator SIBYLL 2.1".

[9] Double bangs are also counted as cascades here

[10] Schönert et al., "Vetoing atmospheric neutrinos in a high energy neutrino telescope".

[11] T. K. Gaisser, Jero, et al., "Generalized self-veto probability for atmospheric neutrinos".

the *OnlineL2* and the *HESE* filter. As this is still part of the online processing, the reconstructions are still required to be computationally inexpensive. In the absence of dedicated event classification algorithms, only the output of these reconstructions is available for realtime selections. For the current version of the HESE realtime stream (HESE v2), three cuts (based on five cut variables) are applied to all the events that pass the HESE filter

- *Direct track length of OnlineL2 Spline MPE* $> 200\,\mathrm{m}$
- $\Delta \log \mathcal{L}$ = *HESE_CascadeVertexFit.logL - HESE_SPE2Fit.logL* > 0
- $Q_{tot} > f(\theta_{OnlineL2\,Spline\,MPE})$ for some function f

While the direct track length is supposed to select events that have good angular reconstruction, the cut on the likelihood ratio is aimed to separate starting tracks from cascades. The final, zenith dependent, charge cut removes atmospheric neutrinos, as their rate strongly depends on the inclination angle. The HESE v2 stream comes in two purities, *gold* for cuts that give a 50% purity and *bronze* for cuts resulting in a 30% purity. The respective event rates are shown in Table 6.1.

Cutting on the track length is intuitive to select well-reconstructable tracks, but it introduces some severe problems by preferably selecting atmospheric muons. The problem originates from the underlying hypothesis of *SplineMPE*, which assumes an infinite muon track. Therefore, starting tracks (with large inelasticity, i.e., high energy transfer to the hadronic casades) are not covered by the hypothesis. In fact, *SplineMPE* is only based on hit probabilities and does not include no-hit DOMs. This allows the reconstructed direction to point far away from any recorded hits - and thereby the true direction - without being penalized in the likelihood. Examples of such experimental events are shown in Figure 6.3 (two event views on the right).

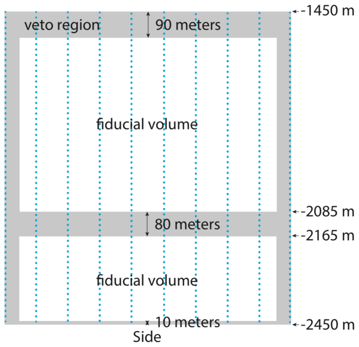

Figure 6.2: Side view of the IceCube detector with the HESE veto region indicated as grey shaded area. As the penetration of atmospheric muons is largest for vertical down-going events a large cap is needed on top of the detector. Adapted from M. Aartsen et al., "Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector".

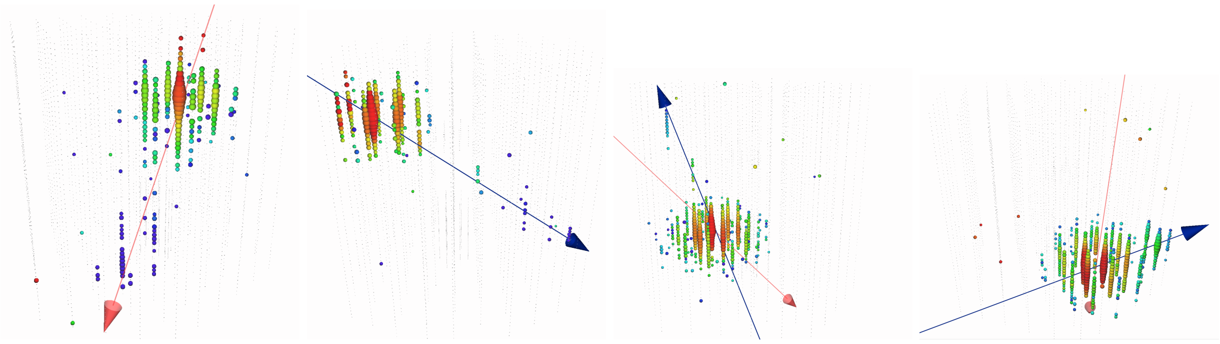

Figure 6.3: Event views indicating problems with the HESE v2 realtime selection: a) The two left plots show events (IceCube-120501A and IceCube-100912A) which get predicted as through-going tracks by the DNN classifier. The predictions are supported by visual inspection of the events, which show early hits (red) on the outer string layer. This is inconsistent with the starting track assumption as the photons from the primary cascade should be boosted in the event direction and hits in the backward direction should only appear later from scattered photons, b) The two plots on the right show clear starting tracks (IceCube-140216A and IceCube-200107A) where the *SplineMPE* reconstruction is visually far off the truth. The RNN reco on the contrary is aligned with the expectation.

We conclude, that *SplineMPE* wrongly reconstructs starting tracks systematically, resulting in the opening angle distribution between truth and reconstruction for HESE starting tracks as shown in Figure 6.6. The HESE v2 selection tries to compensate for the bulk of wrongly reconstructed events by selecting only events with a sufficiently large *Spline MPE* track length. As noted before, *Spline MPE* is, however, only reliable for through-going tracks. Consequently, the processing of archival HESE v2 data with the DNN classifier shows indications that the number of background muons in the HESE v2 data is higher than expected from Monte Carlo. Table 6.2 gives a complete list of archival events that pass the HESE v2 selection criteria with the respective predictions of the DNN Classifier. In the 8 years of data tested 4 (2), events that pass the bronze (gold) selection are predicted as through-going, i.e., are of atmospheric muon origin. This is clearly in tension with the predicted rates from Monte Carlo, which are 0.4/year (bronze) and 0.09/year (gold). A visualization of a subset of these events is shown in Figure 6.3, supporting the prediction of the DNN classifier. One reason for the discrepancy might be a weaker veto region between string 21 and 30. This spot that was originally detected in the medium energy starting cascade (MESC) point source search,[12] is clearly visible when plotting the detector entry points of all events that pass the MESE selection [13] and get predicted as through-going tracks, Figure 6.4. While the total event rates strongly decrease with increasing charge threshold, the leaking spot (shown in black) remains visible. Up to now, there is no clear explanation for the origin of this effect.

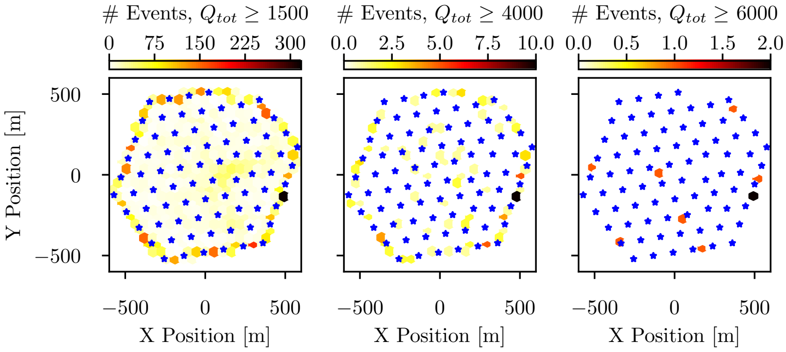

Figure 6.4: The *SplineMPE* detector entry points for all events in the 7.5 years of MESE sample (2010-2017) that are classified as through-going tracks by the DNN. From left to right the Q_{tot} threshold is increased. Blue stars show the position of IceCube strings. The black spot between string 21 and 30 on the right of the grid indicates a leaking spot of atmospheric muons.

In summary the HESE v2 selection is facing two major issues which are both connected to the starting track selection via *SplineMPE*

- Low efficiency due to the removal of events that are difficult to reconstruct for *Spline MPE*

- Higher than expected rate of background muons

Altogether, this motivates a revised version of the HESE starting track realtime stream incorporating the advantages of the DNN classifier. In fact, although not yet finally implemented, a first DNN-based GCN notice was send out for IceCube-200107A,[14] a high energy starting track that was found in coincidence with an HSP blazar, 3HSP J095507.9+355101[15],[16] see chapter 8.

[12] M. Aartsen et al., "Search for Sources of Astrophysical Neutrinos Using Seven Years of IceCube Cascade Events".

[13] HESE equivalent at lower energy with a veto that is adaptive to the reconstructed event vertex. As a result the background reduction is efficient down to ~ 1 TeV

[14] Stein, "IceCube-200107A: IceCube observation of a high-energy neutrino candidate event".

[15] Giommi, Padovani, Oikonomou, et al., "3HSP J095507.9+355101: a flaring extreme blazar coincident in space and time with IceCube-200107A".

[16] Paliya et al., "Multi-Frequency Observations of the Candidate Neutrino Emitting Blazar BZB J0955+3551".

6.2 Deep Learning-Based HESE Realtime Selection

The HESE realtime starting track selection aims to provide a high-purity sample of astrophysical events with starting track topology. Compared to cascades that constitute the major fraction of the HESE events, starting tracks have a better angular resolution and are therefore more interesting for follow-up observations.

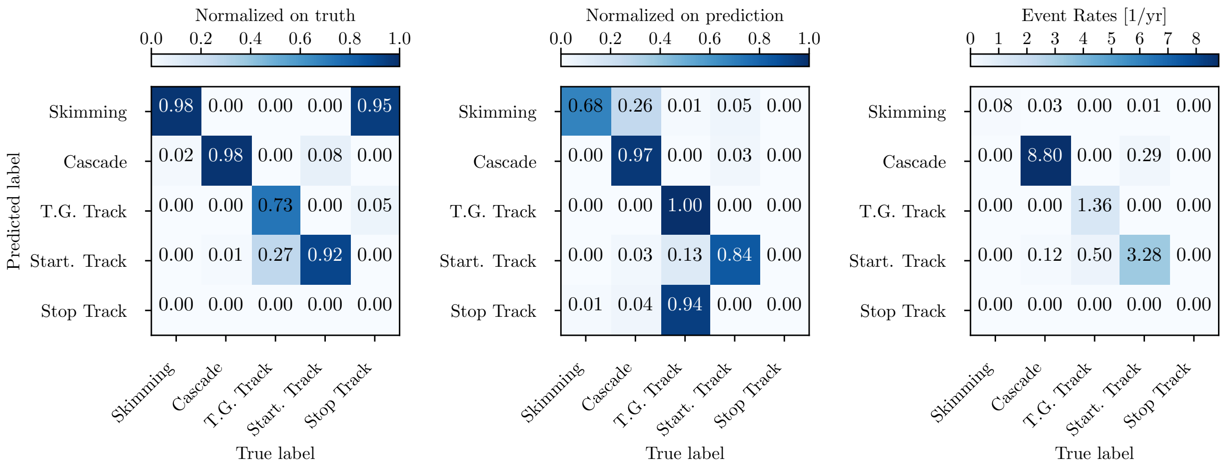

Figure 6.5: DNN confusion matrix for all events that pass the HESE filter. From left to right the confusion, accuracy, and total rates of the predictions are shown. Note that, while stopping tracks are never correctly classified, their overall rate is negligible.

Due to the large energy deposition of HESE events, the DNN classifier accurately predicts the events topologies, even without further cuts on the p-score, as shown in the confusion matrix in Figure 6.5. Specifically, the cascade contribution is well identified and through-going track predictions are accurate up to more than 99%. A smaller fraction of through-going track is, however, predicted as starting tracks. This is expected, as through-going track passing the HESE selection are by construction of the veto very similar to starting tracks. Going beyond the simple classification scheme, the p-scores of the DNN classifier can be used to construct a simple, straight cut-based selection that has a high efficiency for starting tracks and simultaneously reduces the muon background to a negligible number. Starting from all the events that pass the HESE 2015 online filter, three cuts are applied

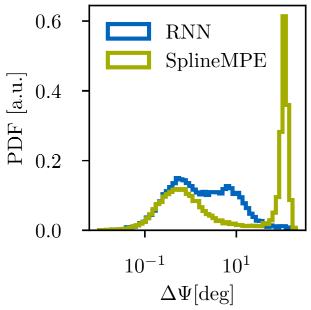

Figure 6.6: The reconstruction error of true starting tracks in the HESE event stream for the RNN Reco and *Spline MPE*.

1. The event is classified as starting track, i.e $p_{max} = p_{starting}$
2. The through-going p-score, $p_{through-going}$, is smaller than 0.01
3. $Q_{tot} > f(\theta_{RNN})$ with function f as illustrated in Figure 6.8.

Figure 6.7 shows how the event rates of the different topologies change with the cuts.

While step 1 efficiently removes nearly all the cascades from the stream, the hard cut on the through-going p-score in step 2 resolves the confusion between starting and through-going tracks. It thereby removes the remaining fraction of atmospheric muons. At this stage, the efficiency for astrophysical starting tracks is close to 100%, and

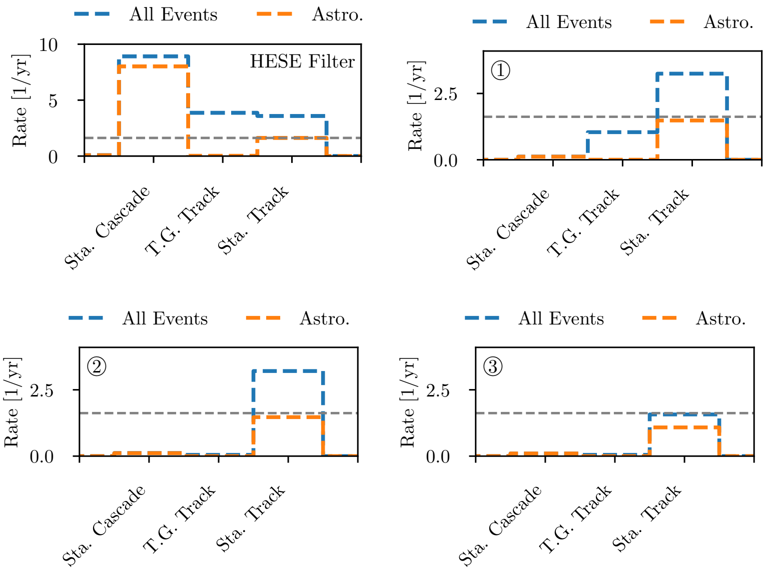

Figure 6.7: Yearly event rates per class after each of the cuts as explained in the text. From top-left to the bottom-right, cuts are placed on the HESE Filter, the starting track prediction (1), the through-going score (2), and the zenith and charge (3). The horizontal dashed line indicates the rate of astrophysical starting tracks after the initial HESE filter.

all other topologies are basically removed from the stream. There is, however, still large contamination of starting tracks produced by conventional atmospheric neutrinos. Since this contamination is highly zenith dependent and especially extreme around the horizon, a zenith dependent cut is applied to guarantee a purity of at least 50% in each zenith direction, see Figure 6.8. Instead of using *SplineMPE*, which isn't stable for starting tracks, a newly developed recurrent-neural network (RNN) reconstruction is used to estimate the direction [17]. A comparison of the two reconstructions is shown in Figure 6.6. For the RNN, the median reconstruction error on HESE starting tracks is around ~ 0.7 degrees, with smooth tails to high values. This is in clear contrast to *Spline MPE* that shows a second peak of reconstruction errors at around ~ 180 degrees.

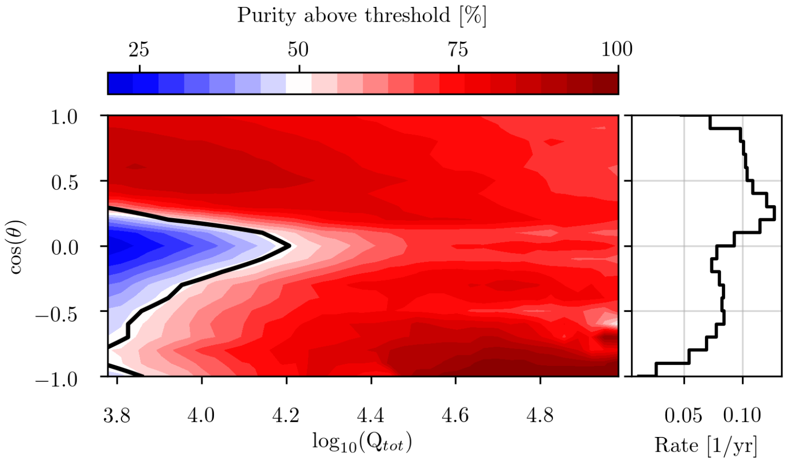

Figure 6.8: The cumulative purity above a given Q_{tot} threshold as a function of the reconstructed zenith angle. The reason for the decreasing purity at very high Q_{tot} is the confusion between starting tracks and cascades that increases with growing cascade energy. The black line shows the resulting cut boundary $f(\theta_{RNN})$ as explained in the text. In the right panel the zenith dependent yearly event rate after the selection cuts is shown.

[17] https://indico.physik.uni-muenchen.de/event/22/contributions/209/attachments/105/153/ErUM_Sept_2019_ECAP_malyshev.pdf, visited 12/2020.

A comparison of the event rates between HESE v2 and v3 selection is shown in Table 6.1. Evidently, signal event rates are increasing significantly, by a factor of ~ 2.5 from HESE v2 gold to HESE v3, reaching an overall astrophysical efficiency [18] of 67%. The major reduction in astrophysical efficiency originates from the charge cut, which reduces conventional atmospheric background and removes low energy astrophysical events. The overall high efficiency is also visible in the effective area, Figure 6.9. In fact, the HESE v3 effective area is fairly similar to the full HESE ν_μ charged-current effective area, i.e., the effective area for *all* HESE starting tracks. Deviations between the two are caused by cutting out atmospheric events (low energy) and confusing cascades with starting tracks (high energy).

[18] Fraction of astrophysical starting events that survive the selection cuts

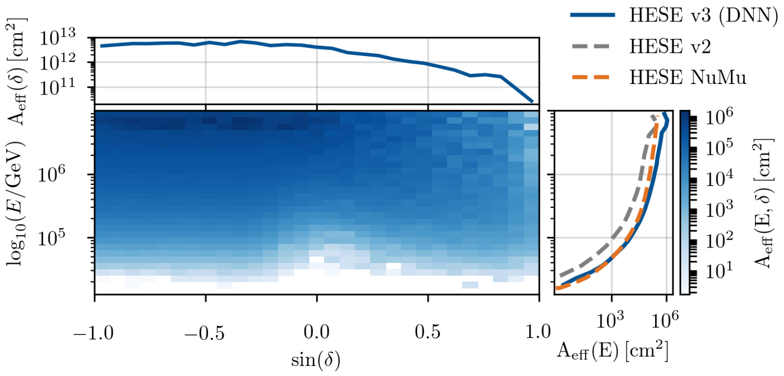

Figure 6.9: Differential effective area in energy and declination. The projections on the x- and y-axis show the marginal effective areas. For the energy-dependent effective area, a comparison of the new HESE v3 selection to the previous v2 selection and the overall HESE ν_μ CC effective area is shown.

In addition to the purity of the stream, an event-by-event signalness is defined as

$$S = \frac{N_{sig}(Q_{tot}, \theta)}{N_{tot}(Q_{tot}, \theta)} \tag{6.1}$$

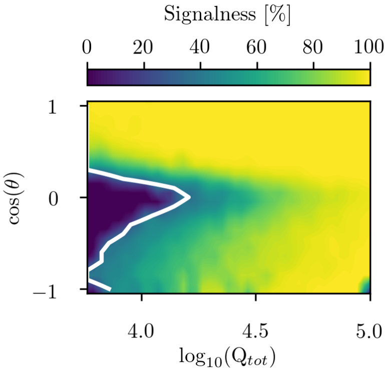

Figure 6.10: The signalness as a function of deposited charge and zenith angle for the HESE v3 selection. The white line represents the cut boundary of the selection. The calculation includes the atmospheric self-veto probability.

where N_{sig} is the expected signal rate and N_{tot} the total signal and background rate as a function of energy and zenith. In contrast to the v2 signalness, the v3 calculation also includes the atmospheric self-veto, boosting the signalness in the down-going region significantly, see Table 6.2. A plot of the signalness as a function of deposited charge and zenith angle is shown in Figure 6.10.

Applying the new and old classification to archival HESE data from 2010 to 2017 (7.5 years) reveals major differences between the old and the new selection, Table 6.2. As previously noted, there is some indication that the HESE v2 selection underestimates the rate of background muons. Besides those events that are likely muon background, all but one of the HESE v2 gold events also appear in the new v3 selection [19]. Furthermore, 8 events, some with signalness $> 90\%$, are added to the selection. Overall, 12 events pass the new HESE v3 selection in 8 years of data. This corresponds to a yearly rate of 1.6 ± 0.46, which is well consistent with the number estimated in Table 6.1.

[19] The missing event is just below the charge threshold of the new selection.

Table 6.1: Comparison of the yearly event rates between the HESE Filter, the old HESE v2 selection and the new HESE v3 selection. Uncertainties on the neutrino components can be calculate from systematics variations of the ice and the DOM efficiency, which significantly impact the total collected charged and thereby the rate of the initial HESE filter. The rates of HESE v2 are calculated for the same spectral assumptions as for HESE v3 and including the self-veto. Note that, as discussed in the text, the rate of atmospheric muons seems underestimated for the HESE v2 case. Hence the purities given here are only an upper limit, the truth can be significantly lower. In the last to rows the purity of the stream is given including and excluding astrophysical cascades, that technically don't belong to the stream, but are still part of the astrophysical signal.

Event Type	HESE Filter	HESE v2 B/G	HESE v3
Astro. Starting Track	1.60	0.66 / 0.38	1.07 ± 0.15
Astro. Cascade	8.05	0.01 / 0.01	0.10 ± 0.01
All Conv. Atmo. ν	2.91	0.71 / 0.07	0.50 ± 0.15
Atmo. Muons	~ 1.86	0.40 / 0.09	0.02
Purity (w/o cascades)	–	< 37 % / < 69%	64%
Purity (w cascades)	–	< 38% / < 71%	69%

6.3 Data/Monte Carlo Agreement and Consistency Checks

For the training of the DNN classifier and the development of the HESE v3 event selection, a baseline Monte Carlo set has been used. It incorporates the best knowledge of the detector calibration, including ice models and detection efficiencies. Nevertheless, there are systematic uncertainties on several of these parameters such as the scattering and absorption length of photons in the ice, or the photon detection efficiency. To ensure that the predictions of the DNN classifier, and thereby the event selection, performs consistently between data and Monte Carlo, the p-score distributions can be compared assuming a best-fit flux model as described in Figure 6.7. In Figure 6.11 a data to Monte Carlo comparison is shown for the starting and through-going p-score on 7.5 years of experimental data and varying systematics. Evidently, the distributions are consistent within statistical uncertainties. To further address the impact of Monte Carlo modeling uncertainties, the influence of some generic pulse modifications are tested on an event-to-event basis.

Figure 6.11: Data to Monte Carlo comparison for the DNN predictions on the HESE sample and 7.5 years of data. Different colors show the predictions for different systematics and black error bars show the experimental data.

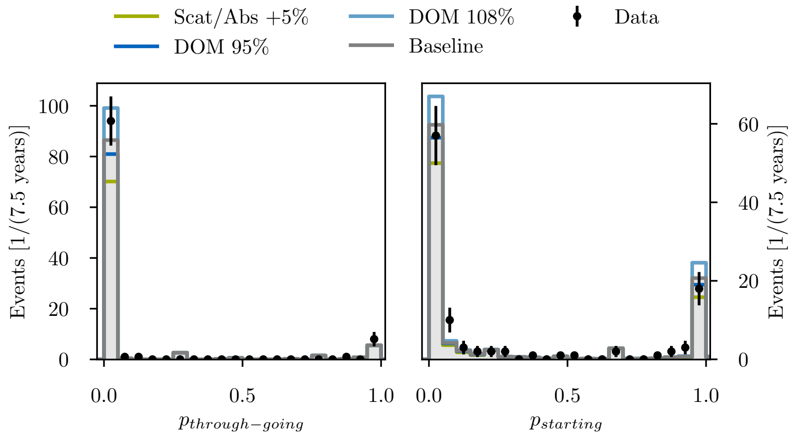

6.3.1 Modulations of the Photon Arrival Times

Uncertainties on the ice model and the Monte Carlo simulation can influence the prediction of the DNN classifier when applied to experimental data. For example, a shift in the scattering or absorption length of photons results in different arrival times. At the time of writing this thesis, no simulations containing the same events with varying ice systematics are available. Hence, we pursue a different strategy and evaluate the effect of *generic* pulse modifications on an event-to-event level. In the left and right plots of Figure 6.12 the shift of the relevant p-score distributions for HESE v3 is shown for smeared pulse arrival times and an increase in white noise, respectively. While the first one imitates the effect of photon propagation uncertainties in the ice, the latter is important to exclude an effect of the noise rate on the classification between through-going and starting events that could be sensitive to few hits in the outer veto layers. However, it appears that the prediction scores on HESE events are too extreme to be significantly impacted by those pulse modification. In fact, even a relative shift of 50% in the prediction, as seen in Figure 6.12, has a negligible effect on the overall event rate of the HESE v3 selection. Consider for example the case of starting tracks, which have a median $p_{starting}$ of 0.99997, i.e $1 - p_{starting} = 0.00003$. Even a relative shift of 100% on this number would only lead to a shift of $p_{starting}$ from 0.99997 to 0.99994. Similar arguments hold for the other event topologies. Consequently, propagating the pulse-shifted samples through the HESE v3 selection results in an overall difference in the event rates of less than 1% and thereby much smaller than the effect of systematics on the total collected charge as visible in Table 6.1.

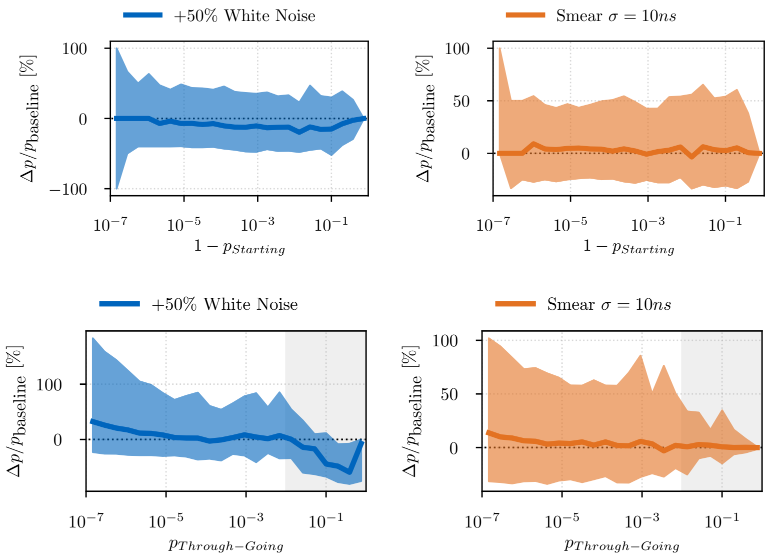

Figure 6.12: The influence of pulse modifications on the predictions of the DNN classifier on a full HESE Monte Carlo sample. The top panel shows the relative deviations for $1 - p_{Starting}$, i.e., events with high starting track probability for additional white noise and a pulse smearing of $\sigma = 10\,\mathrm{ns}$. In the bottom panels, the relative deviations are shown for $p_{Through-Going}$ which is used as a score to reject atmospheric muons. The solid lines show the median value, and the colored bands the 90% central quantile. The grey shaded area in the bottom plots indicates the HESE v3 cut threshold on $p_{Through-Going}$.

6.3.2 Missing DOMs

It occasionally happens that single DOMs of the detector drop during data runs. This usually means that they are not responsive at the beginning of the run and can, therefore, not take any reliable data. Although dropped DOMs are emulated during training, it is still important to ensure the predictions' stability against these effects. In Figure 6.13 the changes in the p-score distributions are shown for the cases of randomly dropping 75 DOMs and dropping the single DOM with the highest charge. Similar to the previous case of pulse modifications, the effect is small enough to guarantee the stability of the predictions for the HESE v3 selection.

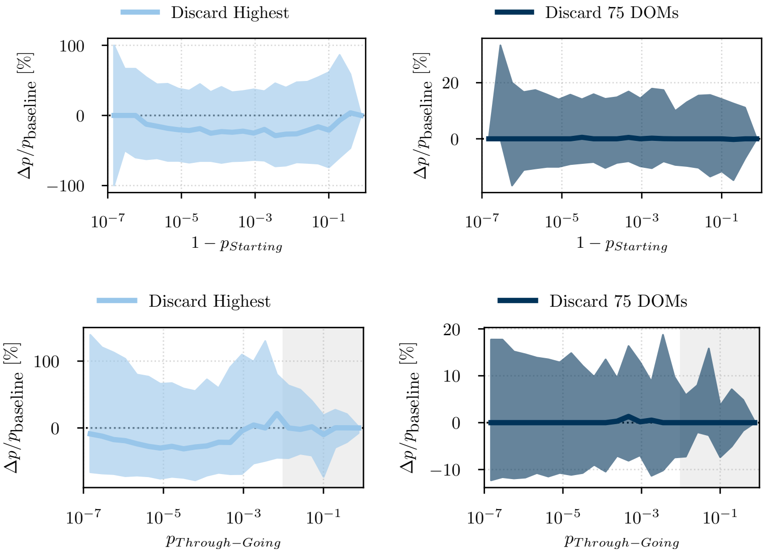

Figure 6.13: The influence of dropped DOMs on the predictions of the DNN classifier on a full HESE Monte Carlo sample. The top panels shows the relative deviations for $1 - p_{Starting}$, i.e., events with high starting track probability for dropping the highest DOM and randomly 75 DOMs, respectively. In the bottom panels the relative deviations are shown for $p_{Through-Going}$ which is used as a score to reject atmospheric muons. The solid lines shows the median value and the colored bands the 90% central quantile. The grey shaded area in the bottom plots indicates the HESE v3 cut threshold on $p_{Through-Going}$.

In summary, there is no indication that either detector or ice systematics have a relevant impact on the predictions. Most importantly, it is not expected that the mismodeling of pulses leads to a larger rate of background muons in the sample. The major impact on the uncertainty of the event rates comes, in fact, from the effect of systematics on the collected charge Q_{tot}, as well as the uncertainty on the atmospheric models.

6.4 Outlook

Although the new, DNN-based selection of HESE significantly improves the performance and stability of the respective realtime stream, it doesn't yet exploit the full realtime potential of starting tracks in general. As can be seen from Figure 6.10 there is still significant room for lowering the Q_{tot} threshold, especially in the down-going region while keeping a high purity. An even more fundamental change of the stream would be to replace Q_{tot} itself with a different energy estimator. There are good arguments to do so. Q_{tot} is a measure of the deposited energy and thereby systematically higher for event topologies that deposit most of their energy in the detector volume. This means that a Q_{tot}-based selection prefers events with dominant cascade contributions. Of particular interest for point source searches are, however, the opposite event types, i.e., starting tracks where only a small amount of energy is transferred to the cascade so that the track is well reconstructable. Those events have a systematically lower Q_{tot}. The problem could be overcome by using a reconstruction algorithm that estimates the initial neutrino energy itself. Again, the usage of deep learning methods can be crucial here to reach acceptable computing times. Such a new event selection could be developed starting from the previously mentioned MESE stream or the newly developed ESTES selection.[20]

Table 6.2: Table of events passing the old HESE v2 (B=bronze, G=gold) or the new HESE v3 selection. In contrast to the v2 signalness, the v3 calculation includes the atmospheric self-veto probability.

Run ID	Event ID	MJD	v2 B/G	v3	Sig. v2	Sig. v3	Start. Score	T.G. Score
115994	29874216	55351	no / no	yes	-	79%	0.979	<0.001
116528	52433389	55451	yes / no	no	45%	-	0.013	0.979
118435	58198553	55756	yes / yes	yes	77%	82%	>0.99	0.002
119214	8606380	55925	no / no	yes	-	98%	0.528	<0.001
120045	22615214	56048	yes / yes	no	55%	-	0.008	0.991
122152	30701331	56390	yes / yes	no	52%	-	>0.99	<0.001
122604	17469985	56470	yes / yes	no	66%	-	0.832	0.167
123867	11659459	56671	yes / no	no	33%	-	>0.99	<0.001
123986	77999595	56679	yes / yes	yes	82%	99%	>0.99	<0.001
124244	7608824	56704	no / no	yes	-	47%	0.943	<0.001
125071	31397276	56859	yes / yes	no	56%	-	0.002	0.997
127853	67093193	57505	yes / no	yes	45%	49%	>0.99	<0.001
128290	6888376	57600	yes / no	no	44%	-	>0.99	<0.001
128340	58537957	57614	yes / yes	yes	60%	97%	>0.99	<0.001
128582	62891970	57667	no / no	yes	-	66%	>0.99	<0.001
128672	38561326	57695	yes / no	no	31%	-	0.022	0.978
129112	14865593	57782	no / no	yes	-	93%	>0.99	<0.001
129510	63482845	57888	no / no	yes	-	97%	0.691	<0.001
130111	70719303	58036	no / no	yes	-	99%	>0.99	<0.001
130126	56068624	58041	yes / yes	no	55%	-	0.665	0.333
130171	34032434	58054	no / no	yes	-	98%	0.889	<0.001

[20] Silva and Mancina, "Measurement of the Diffuse Muon Neutrino Flux using Starting Track Events in IceCube".

7 An Improved Method for Neutrino Point Source Searches

Ever since the beginning of the AMANDA detector in 1996, the identification of neutrino point sources has been among the major scientific goals of neutrino telescopes. While early analyses used binned Poisson statistics to search for excesses of neutrino counts,[1] following analyses used an unbinned likelihood formalism.[2] This increased the overall performance of the analysis and enabled the estimation of source spectra. The unbinned likelihood approach has now been the standard method for nearly a decade of point source analyses in IceCube[3,4].[5] Although the maximum likelihood approach itself guarantees optimal statistical power, the current implementation is non-optimal in its description of the underlying probability density functions (pdfs), thereby introducing biases in the fit parameters. To resolve these problems, this chapter presents a new, rigorous statistical approach to construct the likelihood function and to infer the pdfs directly from Monte Carlo simulations using kernel density estimation. Besides, further improvements are made through new event reconstruction methods. Most importantly, we present a deep learning-based energy reconstruction that increases the energy resolution by up to $\sim$ 50%, while being unbiased below 1 TeV, where previous approaches had no discrimination power. Overall, the new analysis reduces the variance in the fit parameters, removes biases, and improves the 5σ discovery potential by 20% to 30 % for a $\gamma = 2.0$ spectrum compared to the previous publication.[6] At the end of the chapter, we apply the new point source analysis on a sample of 9-years of muon tracks from the Northern Hemisphere.

7.1 Derivation of the Point Source Likelihood

7.1.1 Maximum Likelihood Estimators

The concept of maximum likelihood estimators (MLE) forms the basis for the analysis method discussed in this chapter. Therefore, before looking into the specific construction of the point source likelihood, a short review of MLEs is presented. The discussion is mainly inspired by the books of Glen Cowen[7] and Casella & Berger.[8]

---

[1] J. Ahrens et al., “Search for Extraterrestrial Point Sources of Neutrinos with AMANDA-II”.

[2] Braun et al., “Methods for point source analysis in high energy neutrino telescopes”.

[3] M. G. Aartsen, K. Abraham, et al., “All-sky Search for Time-integrated Neutrino Emission from Astrophysical Sources with 7 yr of IceCube Data”.

[4] M. G. Aartsen et al., “Search for steady point-like sources in the astrophysical muon neutrino flux with 8 years of IceCube data”.

[5] M. G. Aartsen et al., “Time-integrated Neutrino Source Searches with 10 years of IceCube Data”.

[6] Ibid.

[7] Cowan, *Statistical Data Analysis*.

[8] Casella and Berger, *Statistical Inference*.

For the beginning, consider a situation with a set of finite measurements (observations) which we want to use to infer a set of model parameters. Mathematically speaking this corresponds to a set of independent random variables $x = (x_1, \ldots, x_n)$, with each of them distributed according to the same probability density function *pdf* $f(x|\theta)$. The vector hereby θ describes a set of free model parameters $\theta = (\theta_1, \ldots, \theta_m)$. Intuitively, the optimal model parameters θ are those that maximise the likelihood to see the observed data. With f being a *pdf*, the probability of observing a certain x_i is

$$p(x_i) = f(x_i|\theta) \cdot dx, \tag{7.1}$$

and since the measurements x_i are independent, the overall probability is the product of the single probabilities

$$\mathcal{L}(\theta|x) = \prod_{i=1}^{n} f(x_i|\theta). \tag{7.2}$$

The infinitesimals, dx, have been dropped as they pose arbitrary factors that are independent of the model parameters. The function $\mathcal{L}$ is called the *likelihood*. In order to estimate the optimal values for the free parameters θ_i, the maximum likelihood value is determined through differentiation, i.e.,

$$\frac{\partial \mathcal{L}}{\partial \theta_i} = 0, \;\; i = 1, \ldots, m. \tag{7.3}$$

The resulting best-fit parameters are denoted as $\hat{\theta}_i$. MLEs have two important properties, *consistency* and *asymptotic normality*[9]

- **Consistency**: An estimator is consistent if it converges against the true unknown parameter θ_t for the sample size going to infinity

  $$\hat{\theta} \to \theta_\mathrm{t} \text{ as } n \to \infty. \tag{7.4}$$

- **Asymptotic Normality**: In the asymptotic case of the sample size approaching infinity, $n \to \infty$, the estimated $\hat{\theta}$ is normal distributed around the truth θ_t with some variance σ_t

  $$\sqrt{n}(\hat{\theta} - \theta_\mathrm{t}) \to N(0, \sigma_t^2), \text{ in distribution.} \tag{7.5}$$

As noted before and further discussed in section 7.5, the parameter estimation of the current IceCube point source analysis is biased in spectral shape and signal strength. In fact, the bias doesn't reduce for a larger sample size, which is a clear hint on the mismodeling of the likelihood function and a major motivation for the development of the improved method presented in this chapter.

Finally, the variance σ_t^2 of a MLE can be related to the Fisher-Information $I(\theta_t)$ as $\sigma_t^2 = I(\theta_t)^{-1}$. Empirically, the Fischer-Information provides a measure of how fast f changes when slightly varying θ in a repeated experiment with same truth θ_t, i.e.

$$I(\theta_t) = E_{\theta_t} \left(\frac{\partial}{\partial \theta} \log f(x|\theta)|_{\theta=\theta_t} \right)^2, \tag{7.6}$$

with expectation value E.

---

[9] Casella and Berger, *Statistical Inference*, p. 470 ff.

7.1.2 *Hypothesis Test*

A frequent problem in statistics is the discrimination between two hypotheses, the background hypothesis H_0 and the signal hypothesis H_1, for a given set of observations x. It was originally proven by Jerzy Neyman and Egon Pearson, that for simple hypotheses, i.e., hypothesis with fixed parameters θ_0 (background) and θ_1(signal), the likelihood-ratio provides the most powerful statistical test.[10]

Lemma 7.1.1 (Neyman-Pearson Lemma) *Suppose one wants to do a statistical test with two simple hypothesis H_0 and H_1 and corresponding model parameters $\theta = \theta_0$ and $\theta = \theta_1$. The set of observation is denoted as x with corresponding probability density function $f(x|\theta)$. In this case the ratio of the likelihood functions $\mathcal{L}(\theta|x) = f(x|\theta)$*

$$\Lambda(x) = \frac{\mathcal{L}(\theta_0|x)}{\mathcal{L}(\theta_1|x)}$$

provides the most powerful test at significance level α, where

$$\alpha = P(\Lambda(x) \leq \eta | H_0)$$

for any threshold η, which rejects H_0 in favor of H_1.

In the cases of the neutrino point source search in this chapter, the signal and background hypothesis are defined as follows

- **Signal Hypothesis** H_1 ($\theta = \theta_s$): A neutrino-emitting point source with spectrum $\phi(E) = \phi_0 \times E^{-\gamma}$ is located at a sky position $\vec{\omega} = \vec{r}_{src} = (\alpha_{src}, \delta_{src})^T$. Its emission is on top of the background of (conventional) atmospheric- and diffuse astrophysical neutrinos. The overall set of model parameters is then denoted as $\theta_s = (\alpha_{src}, \delta_{src}, \phi_0, \gamma)$, with α_{src} and δ_{src} describing the right ascension and declination in the J2000 coordinate system, respectively.
- **Background Hypothesis** H_0 ($\theta = \theta_b$): The flux coming from the direction $\vec{\omega}$ is composed solely of (conventional) atmospheric- and diffuse astrophysical neutrinos.

The limitation to single power laws is technically not necessary, but consistent with the prediction of neutrino spectra from theoretical modeling of emission regions with photo-meson or proton-proton production, see section 2.2.

While the background hypothesis is fixed by choosing an atmospheric and diffuse astrophysical flux model - and thereby doesn't have any free parameters - the signal hypothesis has at least two free parameters - the flux normalization ϕ_0 and the spectral index γ. Strictly speaking, the hypothesis test thereby doesn't fulfill the requirement of a simple test in Lemma 7.1.1. Nevertheless, the description of the likelihood ratio can be trivially extended for these problems by

$$\Lambda(x) = \frac{\sup_{\theta \in \theta_b} \mathcal{L}(\theta)}{\sup_{\theta \in \theta_s} \mathcal{L}(\theta)}. \tag{7.7}$$

[10] Neyman and Pearson, "On the Problem of the Most Efficient Tests of Statistical Hypotheses".

The supremum is thereby taken over all the allowed model parameters of the background and signal hypothesis, θ_b and θ_s, respectively. Λ is usually referred to as test statistic value, and its distribution is called the *test statistic distribution*. Note that for the specific case of the point source analysis, the background hypothesis is nested into the signal hypothesis, i.e., they converge against each other for $\phi_0 \rightarrow 0$. This has further important implications as it suggests a simple approximation for the test statistic distribution using *Wilks' Theorem*[11]

Theorem 7.1.2 (Wilks' Theorem) *For testing an hypothesis $H_0 : \theta = \theta_0$ versus $H_1 = \theta \neq \theta_0$, suppose $x_1, \ldots, x_n$ are independent and identically distributed random variables with $f(x|\theta)$, $\hat{\theta}$ is the MLE of θ, and $f(x|\theta)$ is sufficiently regular. Then under H_0, as $n \rightarrow \infty$*

$$\mathcal{TS} = -2 \times \log \Lambda(x) \rightarrow \chi_1^2 \textit{ in distribution}$$

While the above theorem is given for a hypothesis with only one degree of freedom, it can be extended to N-degrees of freedom, showing that in this case the test statistic values are distributed as χ_N^2 for the sample size $n \rightarrow \infty$.[12]

7.1.3 Deriving the Point Source Likelihood Function

In order to derive a likelihood function for the point source search, first of all the observable space needs to be defined. On the most fundamental level, IceCube DOMs measure a time series of pulses, i.e., tuples of charge and time, $p_i = (q_i, t_i)$, produced by Cherenkov photons. Naively, one could hence define the observable space as $p = p_1, \ldots, p_n$, representing all the observed pulses of an event. The process from the production of a neutrino to the observation of the pulses can then be described by a hierarchical model: production → interaction → propagation of secondaries in the detector → detection of pulses. The corresponding likelihood function can be written as

$$\mathcal{L}(\theta_s|p) = f(p|\theta_s) = \int f(p|\mu) f(\mu|\nu) f(\nu|\theta_s) \,\mathrm{d}\mu \,\mathrm{d}\nu. \tag{7.8}$$

Evidently, a series of marginalization terms appear, that highlight the complexity of the problem. The different terms in eq. (7.8) are described in the following:

- $f(\nu|\theta_s)$ is the production probability of a certain neutrino ν with energy E_ν and direction $\vec{r}_\nu$ under the source hypothesis, i.e.

  $$f(\nu|\theta_s) = \hat{\Phi}(E_\nu; \Phi_0, \gamma) \cdot \delta(\vec{r}_\nu - \vec{r}_{src}) \tag{7.9}$$

  with normalized flux *pdf* $\hat{\Phi}$ and delta-distribution δ.

- $f(\mu|\nu)$ describes the cross-section and kinematics of the muon production, i.e., the probability of seeing a muon μ with energy E_μ and direction $\vec{r}_\mu$ given the original neutrino. This term thereby includes the energy transfer from neutrino to the muon, the resulting kinematic angle between the two, and the energy loss before entering the detector.

---

[11] Casella and Berger, *Statistical Inference*, p. 489 & p. 516.

[12] Ibid., p. 490.

- $f(p|\mu)$ is the probability of observing the specific set of pulses given the muon in the previous step. This term is effectively what is used in IceCube reconstruction, i.e., a mapping between pulses and muon quantities. As discussed in section 3.8 the reconstruction procedure is quite complex and computationally intensive.

While eq. (7.8) consistently treats all the steps from the neutrino production process to the observation, it is computationally unfeasible as already a single reconstruction can take up to several minutes. A more efficient approach is therefore to separate reconstruction and hypothesis testing. In fact, by using reconstructed quantities, the observable space can be reduced to a few quantities which carry the majority of the event information.

For the case of a muon track-based point source analysis, there are two natural choices of observables: the muon energy E_μ and the muon direction $\vec{r}_\mu$. In addition, it is advantageous to add another observable, σ, that describes the reconstruction quality of the event. Note that this does not necessarily need to be an angular error estimator, but should correlate with one. The resulting likelihood function reads

$$\mathcal{L}(\theta|x = \{E_\mu, \vec{r}_\mu, \sigma\}) = f(x|\theta). \tag{7.10}$$

In the point source hypothesis as described above there are two data generating processes: a neutrino emitting point source, and background neutrinos (and muons) produced in the atmosphere or coming from the diffuse astrophysical background. Unfortunately, there is no *a priori* way to distinguish signal and background events, adding a latent variable ξ, with $\xi = 1$ for signal and $\xi = 0$ for background events, to the hierarchical model. The likelihood in eq. (7.10) can then be written as marginalisation over ξ via

$$f(x|\theta) = f(x|\xi = 1, \theta) f(\xi = 1|\theta) + f(x|\xi = 0, \theta) f(\xi = 0|\theta). \tag{7.11}$$

The second term of each addend describes the *a prior* probability to observe a signal or background event, respectively. This can be expressed over the normalized expectation values of the number of signal and background events, μ_s and μ_b, respectively. Hence,

$$f(\xi|\theta) = \begin{cases} \mu_b \cdot (\mu_s + \mu_b)^{-1}, & \text{for } \xi = 0 \\ \mu_s \cdot (\mu_s + \mu_b)^{-1}, & \text{for } \xi = 1 \end{cases} \tag{7.12}$$

with $\mu_s, \mu_b \geq 0$. Thereby eq. (7.11) can be rewritten as

$$f(x|\theta) = f(x|\theta_s) \frac{\mu_s}{\mu_s + \mu_b} + f(x|\theta_b) \frac{\mu_b}{\mu_s + \mu_b}. \tag{7.13}$$

This is the form of the point source likelihood for a single observed event. Extending it to a sample of events, one needs to consider that the size of the observed sample N is not fixed, but needs to be treated as an additional random variable that follows a Poisson distribution, i.e.,

$$N \sim \text{Poisson}(\mu_s + \mu_b). \tag{7.14}$$

The final likelihood for the entire set of observables becomes

$$f(\{x\}, N|\theta) = f(N|\theta) \cdot f(\{x\}|\theta) \tag{7.15}$$

where $\{x\} = \{x_1, \ldots, x_N\}$ is the set of events and their observables. Using eq. (7.14), eq. (7.13) and the fact that the x_i are independent random variables, the final expression for the point source likelihood becomes

$$\begin{aligned} f\left(\{x\},\, N \,|\, \theta\right) = {} & \frac{\left(\mu_s + \mu_b\right)^N}{N!} \exp\left(-\mu_s - \mu_b\right) \\ & \times \prod_{i=1}^{N} \left\{ f_s\left(x_i \,|\, \theta_s\right) \frac{\mu_s}{\mu_s + \mu_b} + f_b\left(x_i \,|\, \theta_b\right) \frac{\mu_b}{\mu_s + \mu_b} \right\}. \end{aligned} \tag{7.16}$$

Up to this point the likelihood is completely independent of the specific choice of observables.

As previously noted, a natural choice of observables for a muon track-based neutrino point source analysis are the reconstructed muon energy E_μ, and the reconstructed muon direction $\vec{r}_\mu$ with corresponding reconstruction quality estimator σ. Following eq. (7.16), the likelihood then requires the estimation of two probability density functions, $f_s\left(E_\mu, \vec{r}_\mu, \sigma \,|\, \theta_s\right)$ and $f_b\left(E_\mu, \vec{r}_\mu, \sigma \,|\, \theta_b\right)$, for signal and background, respectively. Note that the directional reconstruction, $\vec{r}_\mu$, is defined on the 2-dimensional manifold of a sphere. Hence it can be useful to work directly on the natural scale of spherical coordinates, $(\sin\delta, \alpha)$ [13], in order to full fill the normalization conditions [14].

As the background hypothesis does not involve any free parameters, the derivation of the corresponding likelihood function is straight forward. Due to the location of IceCube at the geographic South Pole, the *pdf* is further invariant against rotations in right ascension for the case of sufficiently large integration times (> 1 day). The background *pdf* therefore takes the form of

$$f_b\left(E_\mu, \vec{r}_\mu, \sigma \,|\, \theta_b\right) = \frac{1}{2\pi} f_b\left(E_\mu,\, \sin\delta_\mu,\, \sigma \,|\, \theta_b\right) \tag{7.17}$$

with muon declination δ_μ and the background hypothesis θ_b as defined in section 7.1.2. The signal case, in contrast, needs some additional considerations for two reasons: 1.) the right ascension of the event, α_μ, can not be neglected and 2.) the hypothesis contains additional free parameters, i.e., $\theta_s = (\alpha_{src},\, \delta_{src},\, \phi_0,\, \gamma)$. Assume now that we want to test a fixed but arbitrary sky position with $\vec{r}_{src} = (\alpha_{src},\, \delta_{src})^T$. Since the hypothesis test should not depend on the specific choice of the coordinates one can simply use the coordinate system for which the source is placed on one of the poles, i.e., $\tilde{\vec{r}}_{src} = (0,0)^T = \mathbf{R} \cdot \vec{r}_{src}$ with some rotation matrix $\mathbf{R}$. In this coordinate system the signal *pdf* becomes

$$f_s\left(E_\mu,\, \tilde{\vec{r}}_\mu,\, \sigma \,|\, \tilde{\theta}_s = (0, 0, \gamma)\right) = f_s\left(E_\mu, (\tilde{\delta}_\mu,\, \tilde{\alpha}_\mu),\, \sigma \,|\, \tilde{\theta}_s\right) \tag{7.18}$$

with rotated reconstructed right ascension and declination, $\tilde{\alpha}_\mu$ and $\tilde{\delta}_\mu$, respectively. Note that there is no more dependence on ϕ_0 as it has been absorbed in eq. (7.12), describing the relative strength

[13] declination δ, and right ascension α

[14] $$\int_0^{2\pi} \int_{-1}^{1} f(\sin\delta, \alpha)\, \mathrm{d}\sin\delta\, \mathrm{d}\alpha = 1$$

between signal and background flux. Assuming the reconstruction uncertainty σ to be circular, i.e., independent of the azimuth angle, eq. (7.18) gets independent of $\tilde{\alpha}_\mu$ and hence only depends on the (rotated) declination $\tilde{\delta}_\mu$, which is equivalent to the angular distance between source and reconstruction $\tilde{\psi}$. Applying this substitution and marginalizing over $\tilde{\alpha}_\mu$ the signal *pdf* simplifies to

$$f_s\left(E_\mu,\, (\cos\tilde{\delta}_\mu,\, \tilde{\alpha}_\mu), \sigma \,|\, \tilde{\theta}_s\right) = \frac{1}{2\pi\,\sin\tilde{\psi}}\, f_s\left(E_\mu,\, \tilde{\psi},\, \sigma \,|\, \tilde{\theta}_s\right) \tag{7.19}$$

with $f_s\left(E_\mu, \tilde{\psi}, \sigma \,|\, \tilde{\theta}_s\right)$ being normalized in a flat Euclidean space. One can easily confirm the normalization factors by integrating over the sphere

$$\begin{aligned}
I &= \int_0^{2\pi}\int_0^{\pi} \frac{1}{2\pi\,\sin\psi}\, f_s\left(E_\mu,\, \tilde{\psi},\, \sigma \,|\, \tilde{\theta}_s\right)\, \mathrm{d}\cos\psi\, \mathrm{d}\tilde{\alpha}_\mu & (7.20)\\
&= \int_0^{\pi} \frac{1}{\sin\psi}\, f_s\left(E_\mu,\, \tilde{\psi},\, \sigma \,|\, \tilde{\theta}_s\right) \sin\psi\, \mathrm{d}\psi & (7.21)\\
&= \int_0^{\pi} f_s\left(E_\mu,\, \tilde{\psi},\, \sigma \,|\, \tilde{\theta}_s\right)\, \mathrm{d}\psi = 1. & (7.22)
\end{aligned}$$

Finally, the coordinate system can be transformed back by using the inverse rotation matrix $\mathbf{R}^{-1}$ with $det(\mathbf{R}^{-1}) = 1$ and $\mathbf{R}\cdot\tilde{\psi} = \psi$ as angular distances are invariant under rotations. As a result equation (7.19) becomes

$$f_s\left(E_\mu,\, \vec{r}_\mu,\, \sigma \,|\, \tilde{\theta}_s\right) = \frac{1}{2\pi\,\sin\psi}\, f_s\left(E_\mu,\, \psi,\, \sigma \,|\, \theta_s\right). \tag{7.23}$$

For comparison with the previous implementation of the point source analysis[15] [16] one can further split this equation up by using the law of total probability

$$\begin{aligned}
f_s\left(E_\mu,\, \vec{r}_\mu,\, \sigma \,|\, \theta_s\right) &= \frac{1}{2\pi\,\sin\psi}\, f_s\left(\psi \,|\, \sigma,\, E_\mu,\, \theta_s\right) \cdot f_s\left(E_\mu,\, \sigma \,|\, \theta_s\right) & (7.24)\\
&= \mathcal{S}\left(E_\mu,\, \psi,\, \sigma\,|\theta_s\right) \cdot \mathcal{E}\left(E_\mu,\, \sigma\,|\theta_s\right). & (7.25)
\end{aligned}$$

Comparing to the previous method,[17] one finds the general structure to be pretty similar, i.e. one term describing the spatial and another the energy likelihood. There are, however, important differences. Firstly, the spatial term is not modelled as a spectral index independent Gaussian

$$\mathcal{S}\left(E_\mu,\, \psi,\, \sigma\,|\theta_s\right) = \mathcal{S}\left(\psi,\, \sigma\,\right) = \frac{1}{2\pi\sigma^2}\exp^{-\frac{\psi^2}{2\sigma^2}}, \tag{7.26}$$

as this approximation is only valid in the case of high energy events and hard spectral indices, see section 7.4. In addition, the energy term of the previous method is approximated as

$$\mathcal{E}\left(E_\mu,\, \sigma\,|\theta_s\right) = \mathcal{E}\left(E_\mu\,|\theta_s\right) \tag{7.27}$$

[15] Braun et al., "Methods for point source analysis in high energy neutrino telescopes".

[16] And numerical reasons, see section 7.4.

[17] Braun et al., "Methods for point source analysis in high energy neutrino telescopes".

and thereby neglecting the dependence on the reconstruction quality parameter σ. While it is desirable to keep the σ dependence to have a complete description of the problem, available Monte Carlo statistic are not sufficient to get a stable inference of this high-dimensional *pdf*. The σ parameter is therefore also dropped in this analysis from the energy term with only minor influence on the analysis performance as shown in section 7.5. To keep the parameter space consistent, the background term in eq. (7.17) also simplifies to

$$f_b\left(E_\mu, \vec{r}_\mu \,|\, \theta_b\right) = \frac{1}{2\pi} f_b\left(E_\mu, \sin\delta_\mu \,|\, \theta_b\right). \tag{7.28}$$

Note, however, that the spatial term still *conditionally* depends on σ. We will see in chapter 7.5 that it is, in fact, mainly the mismodeling of the spatial term that leads to the biases in the old analysis approach. The numerical inference of the likelihood terms using Monte Carlo-based kernel density estimation is discussed in section 7.4.

Putting everything together, the test statistic, eq. (7.7), for the point source analysis is given by

$$\begin{aligned}
\mathcal{TS} &= 2 \times \log(\Lambda) \\
&= 2 \times \log\left(C \times \frac{\prod_{i=1}^{N}\left\{f_s\left(x_i \,|\, \theta_s\right)\frac{\mu_s}{\mu_{tot}} + f_b\left(x_i \,|\, \theta_b\right)\frac{\mu_b}{\mu_{tot}}\right\}}{\prod_{i=1}^{N} f_b\left(x_i \,|\, \theta_b\right)}\right) \\
&= 2 \times \log\left(C \times \frac{\prod_{i=1}^{N}\left\{\frac{f_b(x_i \,|\, \theta_b)}{\mu_{tot}}\left[\mu_s\left(\frac{f_s(x_i \,|\, \theta_s)}{f_b(x_i \,|\, \theta_b)} - 1\right) + \mu_{tot}\right]\right\}}{\prod_{i=1}^{N} f_b\left(x_i \,|\, \theta_b\right)}\right) \\
&= 2 \times \log\left(C \times \prod_{i=1}^{N}\left\{\frac{\mu_s}{\mu_{tot}}\left(\frac{f_s\left(x_i \,|\, \theta_s\right)}{f_b\left(x_i \,|\, \theta_b\right)} - 1\right) + 1\right\}\right)
\end{aligned} \tag{7.29}$$

with $u_{tot} = \mu_s + \mu_b$ and a factor C, which is the ratio of the two leading Poisson terms in eq. (7.16). Specifically,

$$C = \left(\frac{\mu_{tot}}{N}\right)^N \times \exp\left(N - \mu_{tot}\right), \tag{7.30}$$

with the total number of events N, and $\mu_b = N$ for the background likelihood. This holds always true as there are no free parameters in the background case and therefore by definition $\mu_s = 0$. Hence the background likelihood will always be maximal for $\mu_b = N$. Note that the event sample contains $\mathcal{O}(10^5)$ events and hence eq. (7.30) gets rapidly small when $\mu_s + \mu_b$ deviates from the number of total events N. For this reason previous implementations of the point source analysis fixed $\mu_s + \mu_b$ to N by default. It can be useful to rewrite eq. (7.29) by applying the log function to get

$$\mathcal{TS} = 2 \times \log(C) + 2 \times \sum_{i=1}^{N} \log\left\{\frac{\mu_s}{\mu_{tot}}\left(\frac{f_s\left(x_i \,|\, \theta_s\right)}{f_b\left(x_i \,|\, \theta_b\right)} - 1\right) + 1\right\} \tag{7.31}$$

Finally, the maximum TS value is given as

$$\mathcal{TS}_{max} = \sup_{\theta_s \in \Theta_s} \mathcal{TS}(\theta_s). \tag{7.32}$$

Note that maximizing the test statistics is only correct because the background hypothesis does not have any free parameters. In any other case one would need to take the supremum separately for the signal *and* background likelihood, as stated in eq. (7.7).

7.2 Deep Learning-Based Muon Energy Reconstruction

In chapter 5, a deep neural network (DNN) for the classification of IceCube events has been presented. The neural network is thereby a mapping from a set of DOMs and pulses to an output value — event types in this case. Similar to the classification task, DNNs can be trained to predict other quantities as the muon energy or direction.[18] In this section, we present a slightly revised version of the DNN classifier, which provides an improved estimation of the muon energy on entry for up-going tracks. A comprehensive discussion of the general network structure and pre-processing can be found in section 5.3 and section 5.5.

7.2.1 Training & Architecture of the Network

The DNN energy estimator uses a Inception-ResNet architecture of which a pseudo-code implementation can be found in the appendix B.2 [19]. As a major difference to the DNN classifier, the average pooling operation used in the final part of the classifier is replaced by a small dense network. This allows for a different treatement of energy contributions from different parts of the inhomogeneous detector. Finally, a linear activation function is used to connect the dense layers to the output node.

For the training of the network, a dataset of ~ 6 million up-going muon events [20] in a reconstructed zenith range between 85°and 180°is used, see Table 7.1 for a detailed summary of the Monte Carlo datasets. The mix of two ice models, the newer *Spice 3.2.1* and the older *Spice Lea*, thereby adds more stability against ice systematics to the prediction. For the training, uncleaned *InIceDSTPulses* are used and the dataset is split into training (80%) and validation (20%) set. The estimated common logarithm of the muon energy of entry is evaluated against its truth using a mean-square error loss. In comparison to the deposited energy, the energy on entry has the advantage of being independent of the actual track length in the detector. Similarly to the training of the DNN classifier in section 5.5, the Adam optimizer[21] is used with default settings. In total, the DNN is trained for 23 epochs (5 evaluation steps per epoch), and the network weights with the best validation loss are selected for the final application. The evolution of the training and validation loss is shown in Figure 7.1.

As convolutional neural networks make heavy use of pooling operations to reduce complexity, they wash out the precise location of a feature, manifesting, for example, in a bad *directional* reconstructions of CNNs compared to standard likelihood approaches. For

[18] Huennefeld, “Deep Learning in Physics exemplified by the Reconstruction of Muon-Neutrino Events in IceCube”.

[19] As with the DNN classifier, an IceTray module for the DNN energy reconstruction is available through the *i3deepice* project https://github.com/IceCubeOpenSource/i3deepice.

[20] Only events without atmospheric coincidence are used during training, optimizing it for the usage with split pulsemaps later.

[21] Kingma and Ba, “Adam: A Method for Stochastic Optimization”.

Table 7.1: The set of Monte Carlo training datasets with respective simulation properties. Although coincident background has been simulated for Nancy's datasetrs these events are discarded for the training, as the DNN is only expected to be used on splitted pulsemaps, i.e. single, isolated events.

Dataset	Properties
11029 (00000-01999)	*Ice Model*: SPICELea CLSim *Energy Range*: 100 GeV - 10 PeV (E^{-2} Spectrum) *DOM efficiency*: 1.0
11069 (00000-00999)	*Ice Model*: SPICELea CLSim *Energy Range*: 100 GeV - 10 PeV (E^{-1} Spectrum) *DOM efficiency*: 1.0
11070 (00000-00999)	*Ice Model*: SPICELea CLSim *Energy Range*: 100 GeV - 10 PeV (E^{-1} Spectrum) *DOM efficiency*: 1.0
Nancy's Simulation (00000-04999)	*Ice Model*: Spice 3.2.1 *Energy Range*: 100 GeV - 5 TeV (E^{-2} Spectrum) *Hole Ice*: Dima's flasher-fit-model (p1=0.3, p2=0) *DOM efficiency*: 0.99
Nancy's Simulation (00000-11999)	*Ice Model*: Spice 3.2.1 *Energy Range*: 5 TeV - 10 PeV ($E^{-1.5}$ Spectrum) *Hole Ice*: Dima's flasher-fit-model (p1=0.3, p2=0) *DOM efficiency*: 0.99
Nancy's Simulation (00000-04999)	*Ice Model*: Spice 3.2.1 *Energy Range*: 1 PeV - 100 PeV (E^{-1} Spectrum) *Hole Ice*: Dima's flasher-fit-model (p1=0.3, p2=0) *DOM efficiency*: 0.99

energy reconstructions, however, the approximated position of an energy loss is sufficient as the spatial resolution is only relevant on the scales of detector inhomogeneities.

Overall, the DNN combines a larger amount of event information than truncated energy reconstruction and therefore significantly improves the energy resolution as shown in Figure 7.4 and Figure 7.5. As for the case of the DNN classifier, applying the DNN to experimental data is thereby extremely fast with predictions times on the order of 100 ms per event.

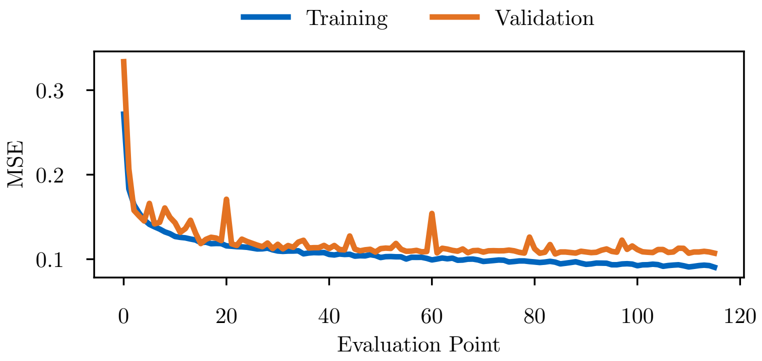

Figure 7.1: Evolution of the mean-squared error loss of the DNN energy estimator for training and validation dataset.

7.2.2 *Performance*

Compared to truncated energy, the DNN energy estimator improves the energy resolution by around 50% in a large part of the energy range, while also resolving the energy degeneracies at muon energies below ∼ 1 TeV, see Figure 7.4 and Figure 7.5. In the low energy regime, muon neutrino losses are dominated by ionization losses compared to stochastic losses that dominated at higher energies. This leads to very continuous (stopping) track signatures with only a few hits in the detector, see Figure 7.2. While likelihood estimators as truncated energy are stabilized against the stochasticity of high-energy events, the truncation of the small information content at low energies leads to instabilities in the energy estimation. As a result, there is no discrimination power for events below ∼ 1 TeV. The DNN, in contrast, can use the full set of event information while also being independent of a specific track hypothesis. Hence it can further gain from stopping tracks as the stopping length is a good proxy for the muon energy on entry, see Figure 3.4.

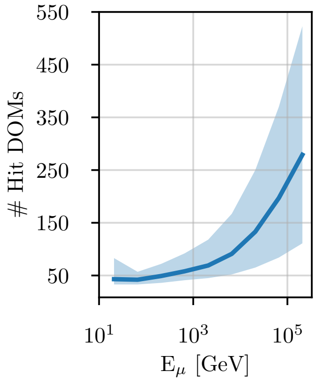

Figure 7.2: Number of hit DOMs against the muon energy, averaged over all events for a representative final level muon neutrino sample. The contours show the 90% central quantile.

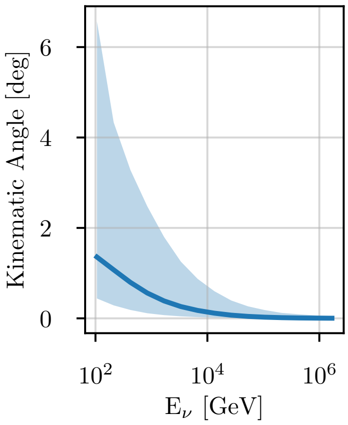

Figure 7.3: Kinematic angle between neutrino and muon for ν_μ CC interactions as a function of the true neutrino energy. The shaded area shows the 90% central quantile.

Although those events with energies below ∼ 1 TeV usually do not carry a strong astrophysical weight, they are an important fraction of the overall number of events as seen in Figure 7.8. Furthermore the energy resolution at these energies also matters for the overall angular resolution of the event. While the reconstruction error σ gives a measure of the quality of the reconstruction between muon direction and reconstructed direction, it does account for the energy-dependent kinematic angle between neutrino and muon, which approximates to

$$\Delta\Psi = \frac{0.7^\circ}{(E_\nu/\mathrm{TeV})^{0.7}} \tag{7.33}$$

and thereby becomes relevant for below a few TeV, see Figure 7.3. The old point source analysis accounts for this additional uncertainty by applying a *pull-correction* that shifts the median opening angle between neutrino and reconstruction to its Monte Carl expectation value for a E^{-2} spectrum. In the formalism of the new point source analysis (section 7.1.3), in contrast, the kinematic angle is naturally folded in during the numerical inference of the pdfs as explained in section 7.4.

The improvement in energy-dependent angular resolution is shown in Figure 7.6. The degenerated energy prediction of truncated energy manifests in a plateau of the median angular resolution that deviates strongly from the expectation assuming the true muon energy. This is the result of the many overlapping true energies that contribute to the same reconstructed energy bin. On the contrary, for the DNN energy estimation, the influence of the kinematic angle is continuously decreasing for increasing reconstructed energies, being well consistent with the true muon energy curve. The result is an *effectively* better angular resolution that impacts all source spectra with a significant contribution of events below a few TeV in energy and

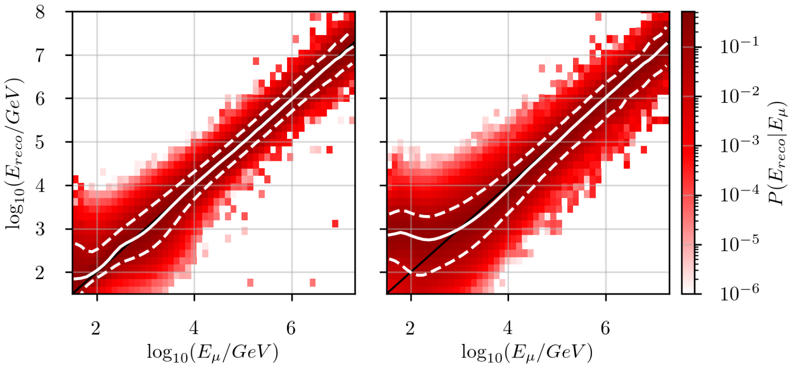

Figure 7.4: Reconstructed vs. true energy for the DNN (left) and truncated energy (right). The white solid lines show the median prediction for a given true energy on entry, E_μ, and the dashed lines the 90% central quantile. The black solid line gives the expectation for an unbiased estimation.

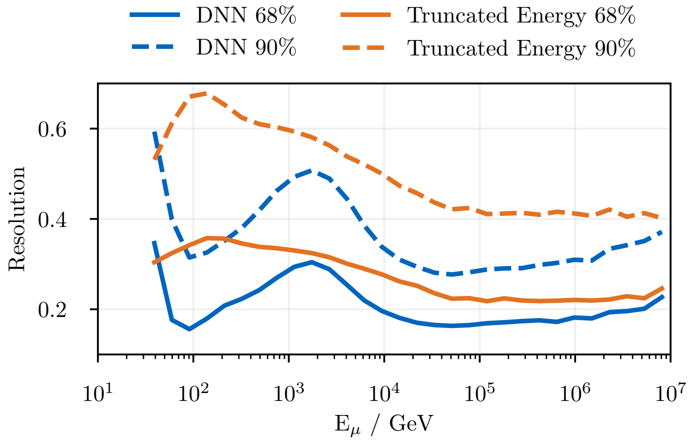

Figure 7.5: Energy resolution of the DNN reconstruction and truncated energy. The resolution is defined as the width of the 68% (90%) quantile of the distribution $f(\log_{10}(E_{reco}) \mid E_\mu)$. The DNN bump at $\sim$ 1 TeV is an effect of the transition between ionization and stochastic losses. At these energies a muon with the same energy on entry can have very different signatures depending on whether or not it has a large stochastic loss. With increasing energy — and thereby increasing number of stochastic losses — the effect averages out.

thereby is of specific importance for soft spectra ($\gamma \gtrsim 3$), as the previously published case of the 2.9σ neutrino excess around the Seyfert II starburst galaxy NGC 1068.[22]

Finally, the improvement of the energy reconstruction is not only affecting the astrophysical signal itself but also enables a better description of the conventional atmospheric neutrino flux and the discrimination between the two. Figure 7.8 shows the energy distribution of experimental data and Monte Carlo for the diffuse up-going track sample used in this point source analysis. It can be seen that, in fact, $\sim 79\%$ of the events have an energy ≤ 1 TeV and are therefore strongly impacted by the new energy reconstruction.

Further discussions of the impact of the new energy reconstruction on the point source analysis's parameter estimation and especially on the fit of the spectral index can be found in section 7.5.

7.2.3 *Influence of Simulation- & Detector Uncertainties*

While Monte Carlo simulations try to model the real world as good as possible, they can never provide complete accuracy. As the DNN is trained on these Monte Carlo events, it is therefore important to ensure that it is insensitive against small deviations from the simulation that could lead to inaccuracies when applying the DNN to experimental data. To test the stability of the DNN against systematic modeling errors, a set of pulse modifications are applied to a test set of Monte Carlo events. Specifically, the impact of pulse shifts, additional white noise, and broken and saturated DOMs are studied. The relative deviation in the prediction on an event-to-event basis as a function of the predicted baseline energy is shown in Figure 7.7. In general, it is visible that all the systematic effects influence the energy estimation only on a few percent level. Additional white noise is relevant for very low energetic events, while the removal of 50 random DOMs [23] has the largest effect in the transition region between ionization and stochastic losses. Saturation windows are time periods in which a DOM collects so many photons that it deviates from a linear response, i.e., saturates. Consequently, removing saturation windows only affects events at the high energy end of the sample. Finally, global Gaussian smearing of the pulse arrival times by 10 ns does not affect the result much, showing the stability of the DNN against mismodeling in the photon arrival time.

Overall, we conclude that the DNN is stable against those generic simulation and detector uncertainties. The specific effect of ice model variations is further discussed in the next section.

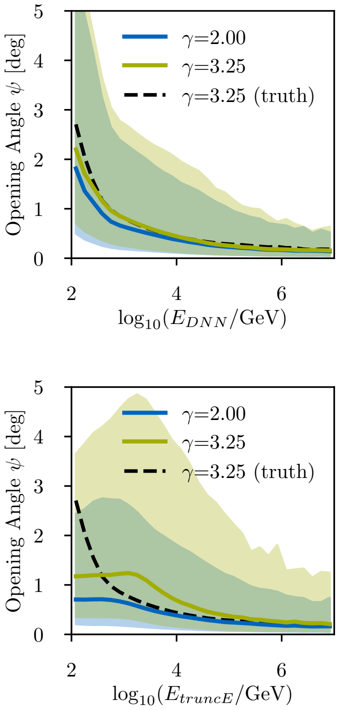

Figure 7.6: Energy-dependent opening angle between neutrino and SplineMPE reconstruction for the DNN energy estimator (top) and truncated energy (bottom) for spectral indices of $\gamma = 3.25$ and $\gamma = 2.0$. The median and central 90% quantile are shown as solid line and shaded area, respectively. Additionally, the expectation assuming the true muon energy on entry is shown as black, dashed line.

[22] M. G. Aartsen et al., "Time-integrated Neutrino Source Searches with 10 years of IceCube Data".

[23] Note that on top of that around 30 DOMs are permanently dead, which has also been accounted for during the training of the DNN.

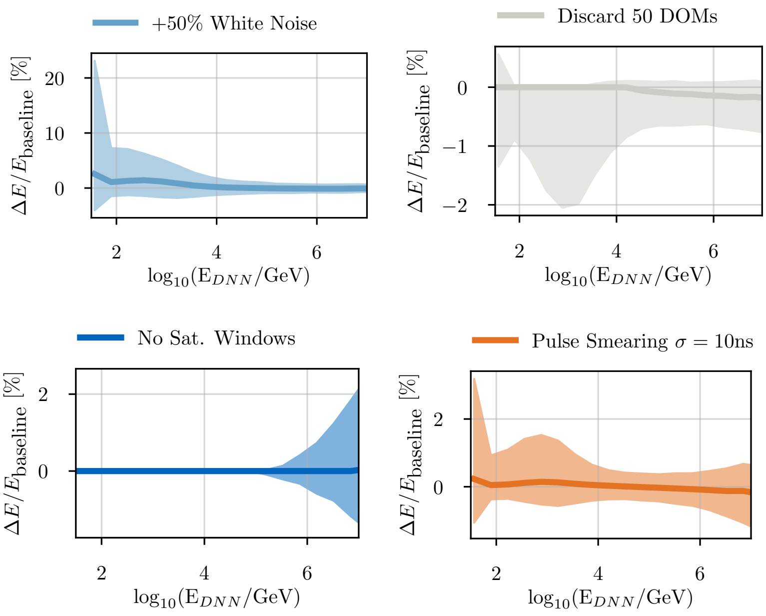

Figure 7.7: Stability of the energy DNN against pulse modifications. The shaded region shows the 90% central quantile of relative deviations from the baseline prediction.

7.3 *Muon Neutrino Dataset and Systematic Uncertainties*

The point source analysis methods presented in this thesis strongly relies on the usage of Monte Carlo simulations for the inference of the likelihood function (section 7.4), as well as the generation of pseudo-experiments for evaluating the analysis's performance and generating the background test statistic distributions (section 7.7). It is therefore required that the event observables are distributed consistently between data and Monte Carlo.

An event sample that is specifically designed to fulfill this criterion is the up-going[24] muon-neutrino sample, which has been originally developed for the measurement of the diffuse astrophysical muon-neutrino flux from the Northern Hemisphere[25].[26] Starting from Muon level 3, two boosted decisions trees (BDTs) are applied to select only well-reconstructed tracks from the Northern Hemisphere ($\delta > -5°$). While the first BDT rejects mis-reconstructed atmospheric muons by requiring quality cuts on the reconstructions, the second BDT removes the remaining fraction of cascades through an (anti-) cascade score. We basically adopt this processing chain but replace the existing cut on events above 100 GeV in truncated energy by a cut on the DNN energy estimation. Furthermore, all runs with less than 85 strings or more than 100 dead DOMs are removed to guarantee stable predictions of all reconstructions. As a result, the total livetime of the sample is lowered by around 0.3%. The remaining sample has an event rate of $\sim$ 2.53 mHz [27] and a very large track purity of $> 99.7\%$ (see also Figure 5.27 for the distribution of event topologies in the sample). In our analysis, we use experimental data from nine years of full-detector configuration (IC86) with the respective livetimes and rates shown in Table 7.2[28]. Compared to the previous point source analyses,[29] a new Monte Carlo simulation

[24] declination $\delta \geq -5°$

[25] Rädel, "Measurement of High-Energy Muon Neutrinos with the IceCube Neutrino Observatory".

[26] M. G. Aartsen et al., "Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data".

[27]

- conventional atmo.: $\sim$ 2.51 mHz,
- prompt atmo.: $\sim$ 0.01 mHz,
- astrophysical: $\sim$ 0.02 mHz

[28] The data have been re-processed to Pass2 which includes charge corrections and a consistent filtering among all detector seasons.

[29] M. G. Aartsen et al., "Search for steady point-like sources in the astrophysical muon neutrino flux with 8 years of IceCube data".

Table 7.2: Summary of the experimental datasets including MJD start and end time, the total livetime of the good runs, and the number of total events.

Season	MJD Start	MJD Stop	Events	Livetime [days]
IC86 2011	55694.4	56062.4	70185	338.09
IC86 2012	56062.4	56414.4	68446	325.43
IC86 2013	56400.5	56783.2	73513	352.28
IC86 2014	56757.4	57160.0	74892	360.88
IC86 2015	57136.1	57528.9	76138	364.60
IC86 2016	57528.9	57891.2	74332	355.78
IC86 2017	57891.2	58309.1	85726	409.91
IC86 2018	58288.8	58682.0	76804	368.09
IC86 2019	58682.0	58998.8	65257	311.04

Table 7.3: Properties of the Monte Carlo datasets that are combined for the point source analysis to have sufficient statistics throughout the entire energy range. The number of events is given for the final level sample.

Dataset	Properties
21002	*Ice Model*: Spice 3.2.1 *Energy Range*: 100 GeV - 100 PeV (E^{-2} Spectrum) #Events: 9’989’091
21124	*Ice Model*: Spice 3.2.1 *Energy Range*: 10 TeV - 500 PeV (E^{-1} Spectrum) #Num Events: 81’208
21217	*Ice Model*: Spice 3.2.1 *Energy Range*: 100 GeV - 100 PeV ($E^{-1.5}$ Spectrum) #Num Events: 2’233’814
21220	*Ice Model*: Spice 3.2.1 *Energy Range*: 100 TeV - 100 PeV (E^{-1} Spectrum) #Num Events: 95’418

set with the latest ice models and 4 times larger statistics has been processed (12,399,531 events in total). A summary of the dataset properties is given in Table 7.3. In accordance with the latest spectral measurements on the up-going muon-neutrino sample,[30] we adopt a conventional atmospheric flux model which is based on the cascade-equation solver MCEq [31] with the GST cosmic ray model[32] and the SIBYLL 2.3c hadronic interaction model.[33] To match the experimental event rates, the conventional spectrum is further scaled with a normalization factor of ~ 1.178. The best-fit astrophysical power law flux, $\phi_{\nu_\mu+\bar{\nu}_\mu} = \phi_0 \times (E/100\,\mathrm{TeV})^{-\gamma}$, for this model is given by $\phi_0 = 1.44 \times 10^{-18}\,\mathrm{GeV}^{-1}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}\,\mathrm{sr}^{-1}$ and $\gamma = 2.28$.[34]

The new point source analysis uses three observables: the DNN energy as presented in section 7.2, the SplineMPE direction as discussed in section 3.8.1 and a new BDT estimator for the angular uncertainty. The BDT uses several high-level input variables as different energy & angular reconstructions, the paraboloid angular error estimator (section 3.8.2), and common variables as the center of gravity of the event to estimate the median opening angle, $\psi_{\mu,\,median}$, between the Spline MPE and the reconstructed muon. To have an angular error estimator that describes the *distribution* of the opening angle ψ_μ, we are, however, not interested in the median opening angle but rather the parameter σ_R in a Rayleigh distribution [35]

$$R(\psi_\mu, \sigma) = \frac{\psi_\mu}{\sigma_R^2} e^{-\psi_\mu^2/(2\sigma_R^2)} \tag{7.34}$$

---

[30] Stettner, “Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data”.

[31] https://github.com/afedynitch/MCEq

[32] Dembinski et al., “Data-driven model of the cosmic-ray flux and mass composition from 10 GeV to 10^{11} GeV”.

[33] Riehn et al., “The hadronic interaction model SIBYLL 2.3c and Feynman scaling”.

[34] Stettner, “Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data”.

[35] 1D analogon of the 2D Gaussian

as an approximation for the spatial term in the point source likelihood. Conveniently, the estimated median opening angle can be converted into the Rayleigh σ_R via $\sigma_R = \psi_{\mu, median}/1.177$. Most importantly, the BDT provides zenith-independent predictions across the entire declination range of the analysis (see Figure 7.12 and Figure 7.13), which is a requirement of the spatial part in the signal likelihood, eq. 7.24.

The experimental data and Monte Carlo distributions of energy, zenith, and angular uncertainty are shown in Figure 7.8, Figure 7.9 and Figure 7.10, respectively. Evidently, the baseline Monte Carlo matches the data very well for all cases. Additionally, systematic datasets with 5% increased scattering and absorption length and two hole ice [36] variations are shown. It can be seen that there is only one variation, the absorption length, that has a larger impact as it systematically shifts the number of measured photons and thereby the energy estimation to lower energies. With the current knowledge of the atmospheric models, it is hard to conclude if the mismatch is consistent with atmospheric flux uncertainties or if an absorption length of +5% can be excluded from this data.

For the last point source observable, the opening angle between reconstruction and original neutrino, ψ, it is not possible to make a dedicated data/MC study as currently no point source for calibration is available. Up to today, the best study for the pointing verification is the measurement of the cosmic-ray shadow of the Moon.[37] Besides, we can use Monte Carlo datasets to test the influence of systematics on the reconstruction error, as shown in Figure 7.11. It can be seen that the track reconstruction quality is largely unaffected by systematics ($< 10\%$), which is consistent with the large lever arm of these events. Overall we conclude, that all observables used in this analysis give consistent baseline predictions between data and Monte Carlo and can therefore safely be used in the analysis.

While the data to Monte Carlo agreement indicates the reliability of our reconstructions on experimental data, the remaining mismatches can still pose a problem for the analysis as a small mismodeling of the large background flux can have a significant impact on the (small) signal. As the last step, we, therefore, apply a correction spline to remove the remaining residual. This results in a perfect background modeling for all further applications such as the inference of the point source likelihood and the generation of pseudo-experiments from Monte Carlo.

7.4 Kernel Density Estimation

As the different terms of the point source likelihood can not be described through simple analytical expression, numerical methods as kernel density estimation (KDE) have to be used. KDEs are thereby a convenient, non-parametric, way to infer a probability density function for a given data sample. In order to use the same data sample for different flux models, IceCube's Monte Carlo simulations assigns

[36] the refrozen ice from the drilling holes

[37] M. Aartsen et al., "Observation of the cosmic-ray shadow of the Moon with IceCube".

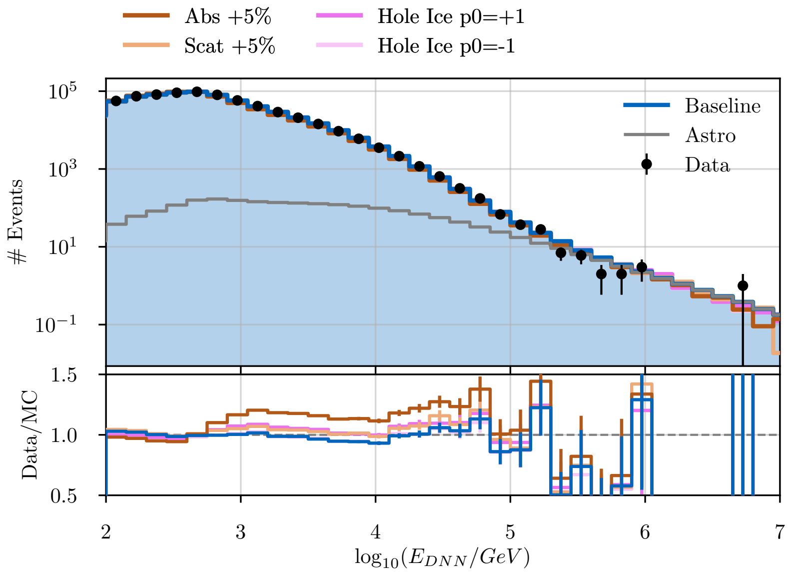

Figure 7.8: Event distribution and Data/MC agreement for the DNN energy. Black points show the experimental data and colored lines the prediction from the baseline and various systematic Monte Carlos, respectively. In addition, the bottom panel shows the data to Monte Carlo agreement for each case.

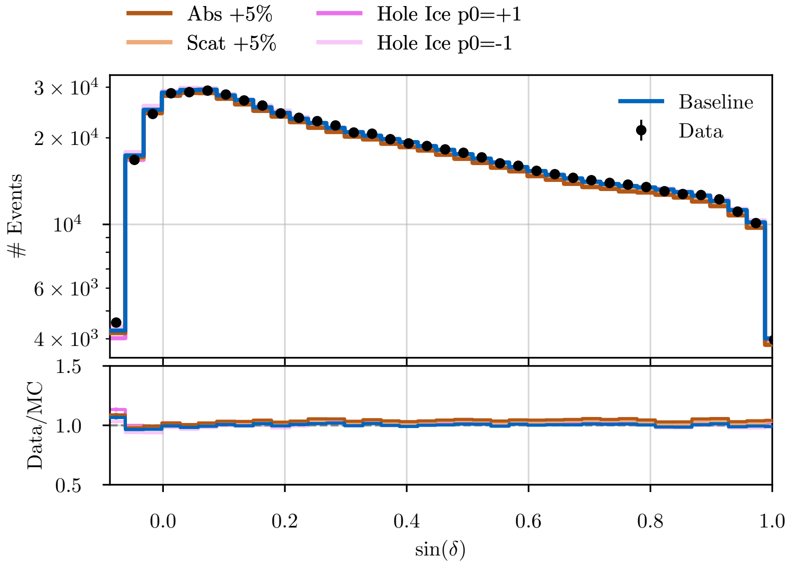

Figure 7.9: Event distribution and Data/MC agreement for the Spline MPE zenith angle. Black points show the experimental data and colored lines the prediction from the baseline and various systematic Monte Carlos, respectively. In addition, the bottom panel shows the data to Monte Carlo agreement for each case.

Figure 7.10: Event distribution and Data/MC agreement for the BDT angular error. Black points show the experimental data and colored lines the prediction from the baseline and various systematic Monte Carlos, respectively. In addition, the bottom panel shows the data to Monte Carlo agreement for each case.

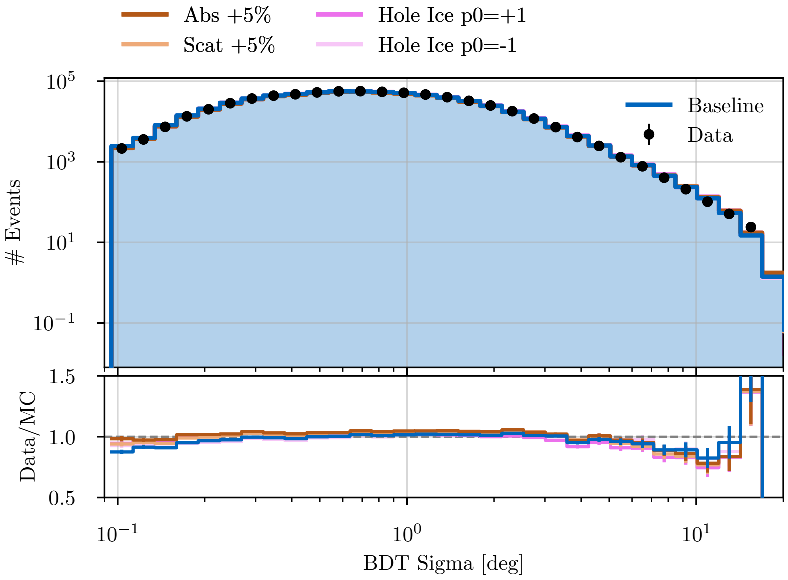

Figure 7.11: Opening angle between Spline MPE and true neutrino direction for different ice models as a function of the neutrino energy and assuming a best-fit spectrum. Solid lines, dashed lines, and dashed-dotted lines show the 50%, 90% and 10% quantiles, respectively.

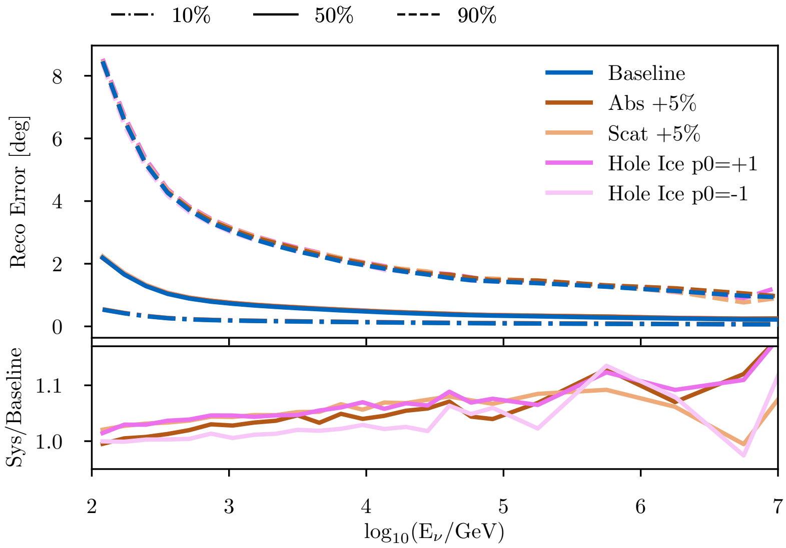

each event with a quantity called *one-weight* that can be used to reweight the event to *any* flux model. The one-weight is defined as

$$w_o = \frac{P_{int}}{E^{-\gamma}} \cdot A \cdot \Omega \int\limits_{E_{min}}^{E_{max}} E^{-\gamma} \mathrm{d}E \tag{7.35}$$

with the surface of the detection volume A, the solid angle of the event generation Ω, and the spectral index of the simulation spectrum γ. In addition, the zenith, flavor and energy-dependent propagation through the Earth, as well as the probability of interaction are summarized in the total interaction probability P_{int}. Using the one-weight one can calculate the event weights for any given flux $\phi(E, \theta)$ via

$$w = w_o \times \phi(E_i, \theta_i) / N_{sim} \tag{7.36}$$

where the factor N_{sim} accounts for the total number of simulated events in the dataset. Event rates can then be calculated by summing up the event weights. Hence, a single Monte Carlo dataset can be used to deduce all the *pdfs* needed in the point source likelihood, i.e., eq. (7.17) and eq. (7.24).

7.4.1 KDE properties

Let $(x_1, \ldots, x_n)$ be an independent and identically distributed sample of observables drawn from some unknown *pdf*. The kernel density estimator (KDE) is then defined as[38]

$$\widehat{f}_h(x) = \frac{1}{n} \frac{1}{\sum_{i=1}^{n} w_i} \sum_{i=1}^{n} w_i \times K_h(x - x_i) = \frac{1}{nh} \frac{1}{\sum_{i=1}^{n} w_i} \sum_{i=1}^{n} w_i \times K\Big(\frac{x - x_i}{h}\Big) \tag{7.37}$$

with the sample (event) weights w_i and a kernel function K_h centered on each x_i. The parameter h is a smoothing parameter which accounts for limited statistics in the sample. KDEs usually do not depend much on the choice of the kernel and hence a n-D Gaussian is used in the following. As the optimal bandwidth is not a-priori clear, a grid search is performed. The fit quality of each grid point (bandwidth) is thereby calculated through a likelihood-based k-fold cross validation. For this procedure — that is widely used in machine learning and data science — the original data sample is split into k parts. In subsequent k iterations the KDE is then trained on $k-1$ parts of the dataset (training dataset) and evaluated on the remaining part (test dataset). This is done uniquely for all the k dataset parts. Finally, the median of the evaluation metric can be used as a measure for the fit-quality of the KDE given a specific bandwidth h. As evaluation metric we use the (weighted) likelihood of the events in the test dataset given the KDE model build on the training dataset

$$\mathcal{L}_{cv} = \prod_i \widehat{f}_h(x_i)^{w_i}. \tag{7.38}$$

[38] Poluektov, "Kernel density estimation of a multidimensional efficiency profile".

7.4.2 Deducing the Point Source Likelihood

The point source likelihood contains a signal and background term, eq. (7.17) and eq. (7.24), respectively. While the background term does not have any free parameters, but is unambiguously defined by the background mdoel, the signal term consists of two conditional pdfs that depend on the spectral index γ. In general, a conditional pdf $p(A \mid B)$ can be calculated as the ratio of two pdfs $p(A, B)$ and $p(B)$ via

$$p(A \mid B) = \frac{p(A, B)}{p(B)}. \tag{7.39}$$

Consequently, the spatial term of the point source likelihood can be calculated as the ratio of

$$f_s\left(\psi \mid E_\mu, \sigma, \gamma\right) = \frac{f_s(\psi, E_\mu, \sigma \mid \gamma)}{f_s(E_\mu, \sigma \mid \gamma)}. \tag{7.40}$$

and the energy term as

$$f_s\left(E_\mu \mid \delta_{src}, \gamma\right) = \frac{f_s\left(E_\mu, \delta_{src} \mid \gamma\right)}{f_s\left(\delta_{src} \mid \gamma\right)}. \tag{7.41}$$

In summary, the point source likelihood can therefore be constructed from 4 KDEs *per spectral index* plus 1 additional KDE for the background. As discussed in chapter 7.1.3, the full energy term has an additional dependency on the reconstruction quality parameter σ. It turns out, however, that the generation of the required 3D KDEs is pretty unstable with the limited Monte Carlo statistics available. In addition, the σ distribution has harsh energy dependent boundaries, driving the optimal bandwidth to much smaller values than desirable for a large part of the phase space. We therefore drop σ from the likelihood with only a marginal effect on the analysis performance as discussed in section 7.5.

For the spatial part, on the other hand, there is a workaround that allows for stable production of 3D KDEs. From the discussion in section 7.1.3 it is known that the opening angle between neutrino and reconstructed direction, ψ, can be approximated by a Rayleigh distribution [39]. Using this approximation, a *relative kernel density estimation*[40] can be used. This specific type of KDEs use a slowly varying kernel to solely model deviations from an analytical approximation, i.e.,

$$f_s(x = (\psi, E_\mu, \sigma) \mid \gamma) = \frac{\sum K\left(\frac{x - x_i}{h}\right)}{(f_{approx} \otimes K)(x)} \times f_{approx}(\psi, E_\mu, \sigma \mid \gamma). \tag{7.42}$$

In our case the approximated function f_{approx} is the product of the 2D energy and sigma pdf and the Rayleigh approximation

$$f_{approx}(\psi, E_\mu, \sigma \mid \gamma) = f_s(E_\mu, \sigma \mid \gamma) \times R(\psi, \sigma_{tot}) \tag{7.43}$$

$$= f_s(E_\mu, \sigma \mid \gamma) \times \frac{\psi}{\sigma_{tot}^2} e^{-\psi^2/(2\sigma_{tot}^2)}. \tag{7.44}$$

---

[39] The 1D projection of the absolute value of the 2D circularized Gaussian in eq. (7.26)

[40] Poluektov, "Kernel density estimation of a multidimensional efficiency profile".

σ_{tot} is constructed to account for both angular uncertainties, the reconstruction uncertainty and the average energy-dependent kinematic angle, i.e.

$$\sigma_{tot} = \sqrt{\sigma_{reco}^2 + \sigma_{kin}^2(E_\mu)}. \tag{7.45}$$

7.4.3 KDE generation pipeline

The KDE generation pipeline for the point source analysis is based on a software package called *Meerkat*.[41] Since the energy, the opening angle ψ and the angular error estimator σ extend over several orders of magnitude, it is advantageous to generate the *pdfs* on a $\log_{10}$-scale, i.e., shifting the observables in the point source likelihood from $E_\mu \rightarrow \log_{10} E_\mu$, $\psi \rightarrow \log_{10}\psi$ and $\sigma \rightarrow \log_{10}\sigma$. As the E_μ and σ transformations affect the signal and background likelihood in the same way, no additional transformation terms appear. In contrast, ψ is an observables that is unique to the signal likelihood and by construction — eq. (7.24) — its distribution expected to be on the ψ scale. The required transformation term is

$$\frac{\mathrm{d}P}{\mathrm{d}\psi} = \frac{\mathrm{d}P}{\mathrm{d}\log_{10}(\psi)}\frac{\mathrm{d}\log_{10}(\psi)}{\mathrm{d}\psi} = \frac{\mathrm{d}P}{\mathrm{d}\log_{10}(\psi)}\frac{1}{\log(10)\cdot\psi}, \tag{7.46}$$

showing that a factor of $1/(\log(10)\cdot\psi)$ needs to be multiplied to the spatial term when generating KDEs on the $\log_{10}$ scale. As the transformation in eq. (7.46) works in both directions, it can also be used to convert the Rayleigh in eq. (7.44) on a $\log_{10}$-scale as required for the usage of the relative kernel density estimation technique in $\log_{10}$ space. Hence,

$$\frac{\mathrm{d}R}{\mathrm{d}\log_{10}(\psi)} = \log(10)\cdot\psi\frac{\mathrm{d}R}{\mathrm{d}\psi} = \log(10)\frac{\psi^2}{\sigma_{tot}^2}e^{-\psi^2/(2\sigma_{tot}^2)} \tag{7.47}$$

The KDE generation pipeline is then based on the following steps

1. A bandwidth optimization for the background term and the higher-dimension *pdfs*, $f_s(\psi, E_\mu, \sigma\,|\,\gamma)$ and $f_s\left(E_\mu, \delta_{src}\,|\,\gamma\right)$, in the spatial and energy term, respectively. In order to save computational resources, the cross-validation is done on a sparse spectral index grid, i.e. $\gamma \in [0.5, 1.0, 1.5, \ldots, 4.5]$. For each γ and each *pdf* the optimal bandwidth is determined from a grid scan with 5-fold cross validation at each grid point. An a priori range of bandwidth is chosen from the approximated scale on which variations are expected to happen, e.g. the energy resolution.

2. The results of the first step are collected and for each observable in both signal terms a γ-dependent 1D interpolator is constructed to independently estimate the optimal bandwidth between the grid points of step 1.

3. The final set of KDEs is generated on a fine grid between $\gamma = 0.5$ and $\gamma = 4.5$ in steps of $\Delta\gamma = 0.05$. The respective bandwidth are taken from the interpolator in step 2. For the lower dimensional denominator in the spatial and energy term, $f_s(E_\mu, \sigma\,|\,\gamma)$

[41] Ibid.

and $f_s(\delta_{src} \mid \gamma)$, the same bandwidth as in the corresponding higher dimensional term are used. This choice ensures the same behaviour towards the boundaries for numerator and denominator in eq. (7.40) and eq. (7.41).

4. Monte Carlo statistics are limited and consequently there are only few events with small opening angle ψ as the phase space close to the true direction gets infinitesimally small. It is, however, desirable that the KDEs are also stable for events that fall very close to the source, as this can not only happen through signal events, but also through background events that happen to be reconstructed close to a source position. A simple and stable solution is to fall back to the analytical Rayleigh approximation in this case. Conveniently, the Rayleigh can be approximated for small ψ as

$$R(\psi) = \frac{\psi}{\sigma^2} e^{-\frac{\psi^2}{2\sigma^2}} = \frac{\psi}{\sigma^2}(1 - \frac{\psi^2}{2\sigma^2} + \mathcal{O}(\psi^4)) = m\psi + \mathcal{O}(\psi^3) \quad (7.48)$$

with some slope m. Given a critical value of $\sigma_{crit} = \alpha \cdot \sigma$ the KDE can hence be continuously extrapolated to arbitrary small ψ values via

$$f_s(\psi \mid E_\mu, \sigma, \gamma) = \frac{f(\psi = \sigma_{crit} \mid E_\mu, \sigma, \gamma)}{\sigma_{crit}} \psi, \ \psi \leq \sigma_{crit} \quad (7.49)$$

As KDEs itself are slow in their evaluation, they are evaluated once on a very fine grid and finally stored as penalized B-splines using the *photospline* package.[42] As photosplines are optimized for high-dimensional spline representations the spectral index dimensions with its 80 individuals *pdfs* can be concatenated. Hence, all likelihood terms can be stored in 3 photosplines, spatial term, energy term, and background term.

A major difference of the new KDE method compared to the previous point source method is the proper modeling of the spatial pdfs in the signal likelihood — including their dependence on the spectral index. A collection of plots comparing the Monte Carlo expectation with the Rayleigh approximation of the old method and the new KDEs are shown in Figure 7.12 and Figure 7.13, as well as in the appendix D.3. Most importantly, it can be seen that, specifically at lower energies, the Monte Carlo expectation strongly deviates from the Rayleigh assumption but is correctly modeled by the KDEs. This has two major effects: 1.) The correct description of the tails allows the KDE method to better recover events that are further away from a source, 2.) Background events that fall close to a source are less pronounced because the KDE is less peaked for small opening angles. Altogether this can explain the fit biases of the old analysis approach, as discussed in section 7.5.

Finally, note that in many applications it is sufficient to evaluate the likelihood function on a subset of events that are in a given circle around the position of interest where the spatial term of the signal likelihood has a non-negligible effect [43]. The resulting simplifications and a description of the software tools for the analysis are discussed in the appendix, section D.1.

[42] Whitehorn, Santen, and Lafebre, "Penalized Splines for Smooth Representation of High-dimensional Monte Carlo Datasets".

[43] In this analysis we use a circle of 15° radius.

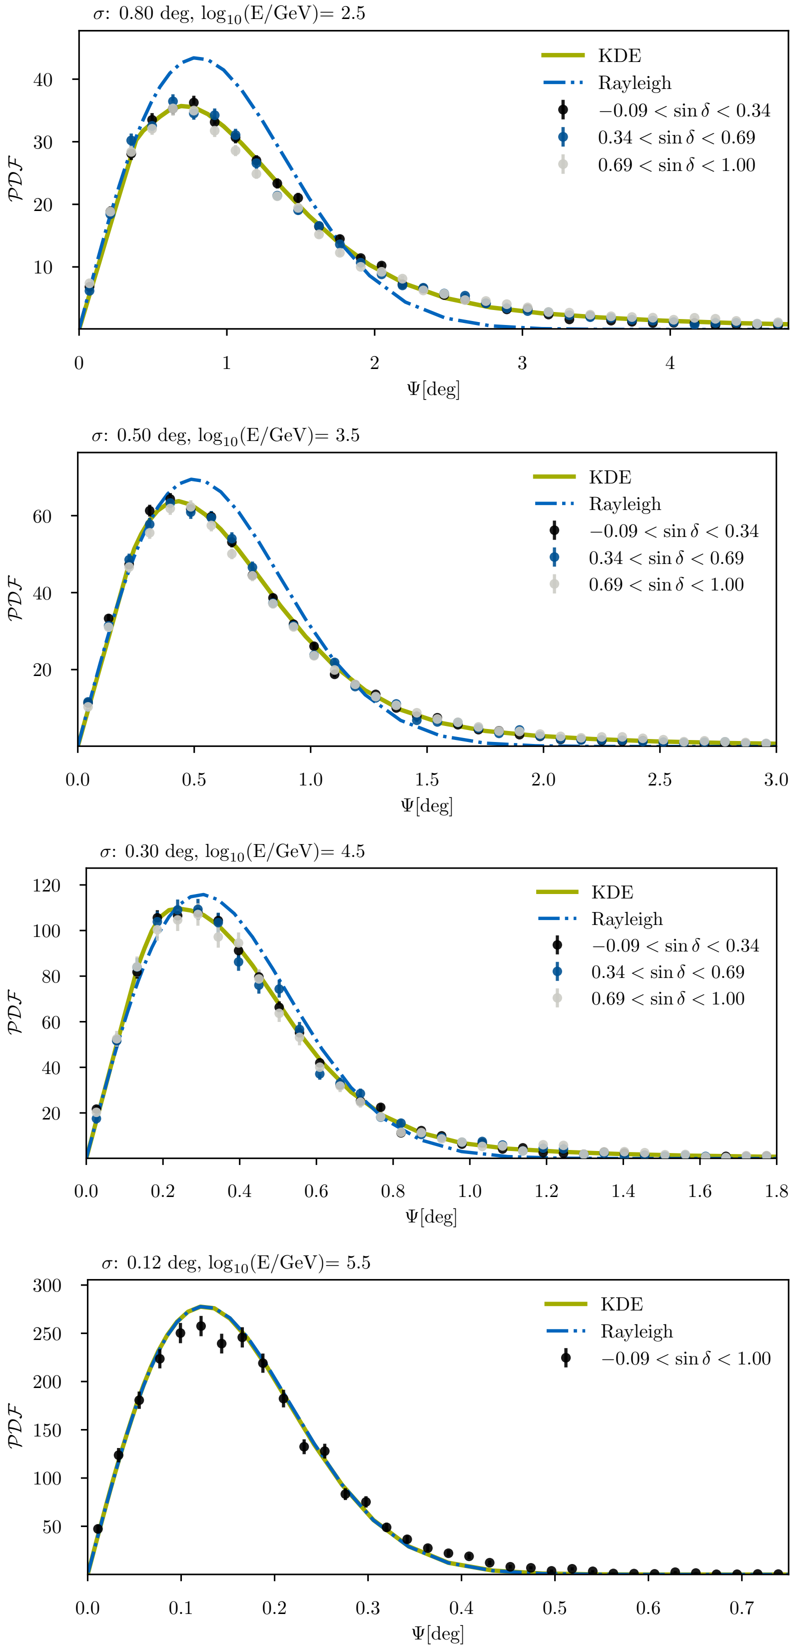

Figure 7.12: Comparison of the KDE spatial term and the Rayleigh approximation for a spectral index of $\gamma = 2.0$. From top to bottom, the plots are shown for increasing DNN energy and a common BDT angular uncertainty. The Monte Carlo truth is given for different zenith bands to highlight its invariance.

Figure 7.13: Comparison of the KDE spatial term and the Rayleigh approximation for a spectral index of $\gamma = 3.25$. From top to bottom, the plots are shown for increasing DNN energy and a common BDT angular uncertainty. The Monte Carlo truth is given for different zenith bands to highlight its invariance.

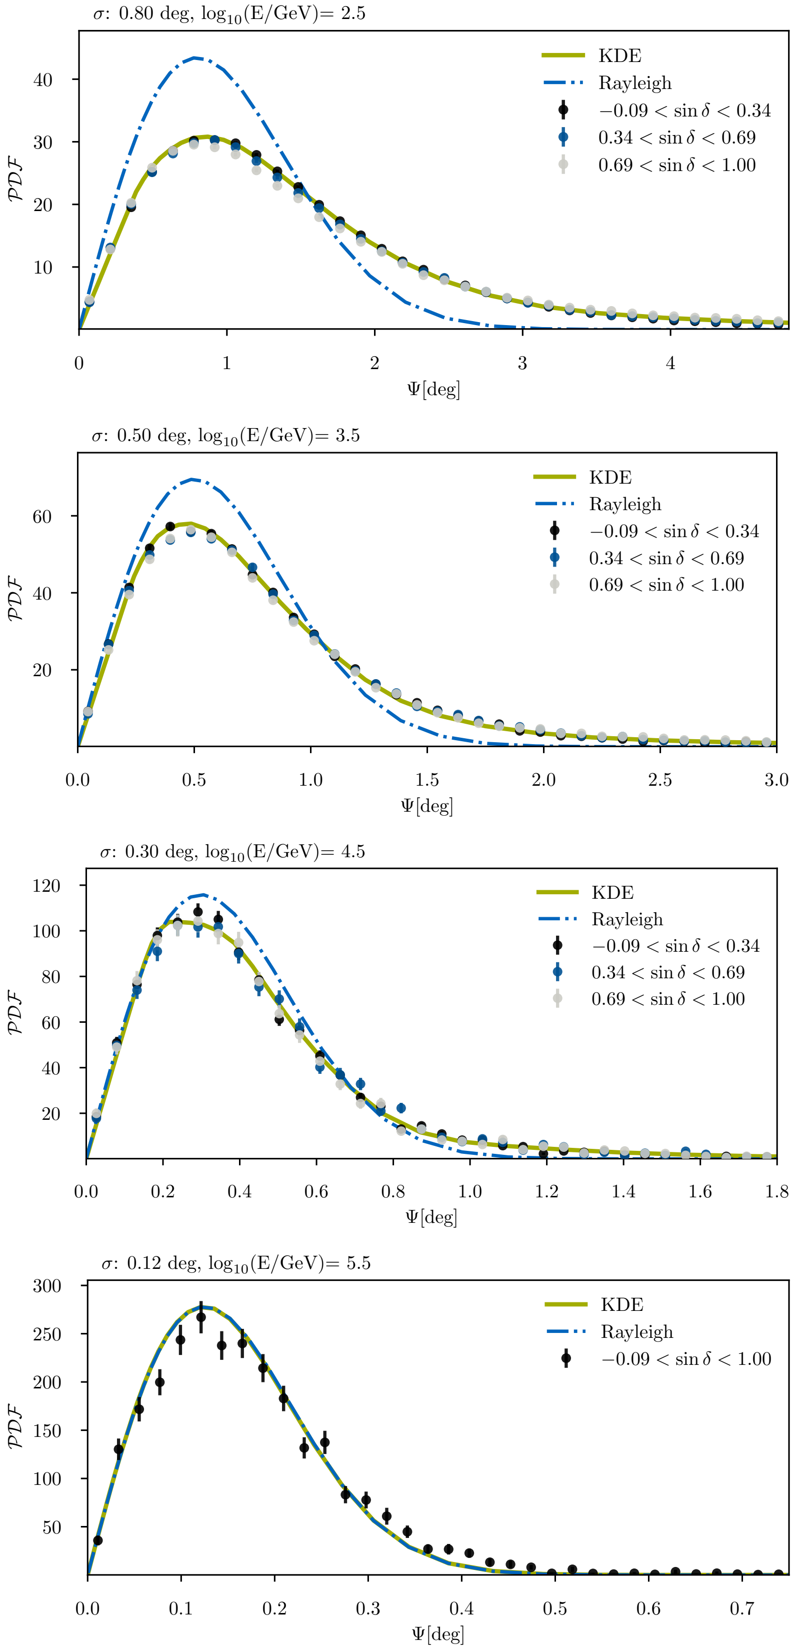

7.5 *Biases in the Fit Parameters*

The consistency of maximum likelihood estimators [44] ensures that the estimated parameters converge against the truth for a large sample size if the likelihood function correctly models the underlying data. In fact, the consistency is an important property for the reliable determination of physical parameters of a detected neutrino source and provides a good cross-check for the reliability of the analysis method. Since the new, KDE-based, point source analysis is based on pdfs directly inferred from Monte Carlo, we expect it to recover the true source parameters when sufficient signal is present. To verify these properties, pseudo-experiments are generated from Monte Carlo, emulating 9 years of recorded data. For each trial, events are injected by randomly drawing from a Monte Carlo sample according to their background weights and adding signal events for a given signal spectrum and position.

It is important for all the following discussions to keep in mind that the point source analysis is a strongly background-dominated analysis. To illustrate this, Figure 7.14 shows a simulated trial with 9 years of background and a soft spectrum ($\phi(E) \propto E^{-3}$) point source with 55 injected events. Most importantly, the large background stabilizes the analysis against mismodeling of the likelihood, which is why approximations in the old point source analysis do not cause larger problems. Signal events that get falsely assigned a small spatial weight can always be interpreted as a background event without being in large tension with the signal hypothesis at a specific sky position. In other words: the large background allows for some mis-identification, some exchange, between signal and background events that comes with a resulting fit bias. In the following, we will show based on a few examples that the new analysis approach, in contrast, produces unbiased results — for spectral index and source strength. Plots showing additional evidence can be found in the appendix, section D.4 and section D.5.

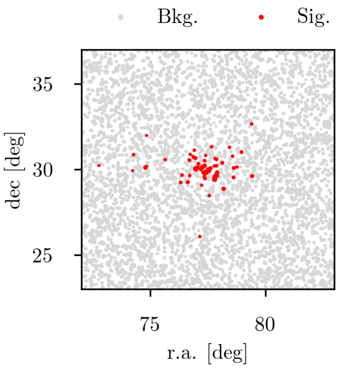

Figure 7.14: Counts map of an exemplary point source trial emulating 9 years of data. Grey dots indicate background events and red dots events from an injected source with $\phi(E) \propto E^{-3}$ and 55 events. The size of the dots scales with the event energy.

7.5.1 *Spectral Index Bias*

The spectral index is an important quantity to study the emission properties of a neutrino source. Figure 7.15 and Figure 7.17 show the recovery of the spectral index of the old and new analysis for injected sources with $\gamma = 2.0$ and $\gamma = 3.25$, as well as two source strength, respectively. The soft spectrum case is specifically motivated from the recent publication of a soft-spectrum 2.9σ excess at the position of NGC 1068.[45]

[44] section 7.1.1

[45] M. G. Aartsen et al., "Time-integrated Neutrino Source Searches with 10 years of IceCube Data".

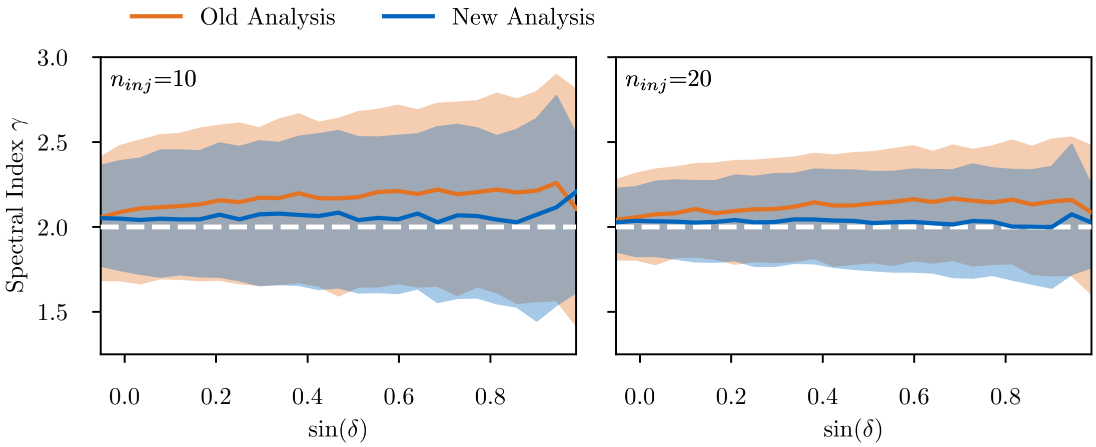

Figure 7.15: Recovery of the spectral index for the old and new method against the declination. The injected spectral index, $\gamma = 2.0$, is shown as white dashed line. The median and central 68% quantiles of the outcomes are shown as solid line and shaded band, respectively. In the left plot 10 events are injected and in the right plot 20 events.

For the harder spectral index, $\gamma = 2.0$, the old point source method shows a general trend to fit softer-than-injected source spectra. This behavior can be understood by the overly peaked spatial likelihood term (see Figure 7.12 and Figure 7.13) that assigns large spatial signal probabilities for (low-energy) background events that happen to fall close to the source. The effect is sufficiently strong not to be compensated by the high-energy signal events that are more accurately modeled. Furthermore, with increasing declination, more and more high-energy events get absorbed by the Earth, effectively softening the astrophysical spectrum observed at the detector and thereby enhancing the bias's effect. Comparing the two source strength in Figure 7.15 it becomes evident that increasing the signal does not significantly impact the median prediction, but only the variance — for both the old and the new method.

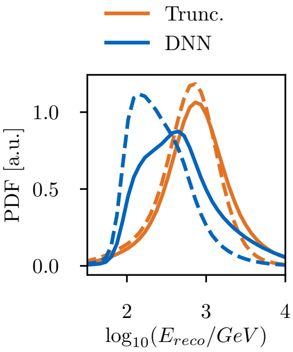

Figure 7.16: Distribution of reconstructed energies for the DNN and truncated energy and two spectral indices, $\gamma = 3.25$ (solid line) and $\gamma = 4.0$ (dashed line), respectively.

In the case of a soft spectrum source, $\gamma = 3.25$, the spectral index is recovered in the median for both methods. However, while the 90% central interval is rather small for the new analysis, it extends to the fit boundary at $\gamma = 4.0$ for the old analysis at 60 injected events (median significance $\sim 3\sigma$) and only drops below that for 90 injected events. Similar to the hard spectrum case, the old analysis performs worse for higher declinations. The KDE method is instead stable across the declination band. While for the estimation of hard spectral indices the improvement of the new method is mainly caused by the improved description of the likelihood function, the improved energy estimation gets particularly relevant for softer spectra. This is evident by looking at Figure 7.16, where the distribution of the reconstructed energy is shown for $\gamma = 3.0$ and $\gamma = 4.0$ and the two energy estimators. While the distributions are similar for truncated energy due to the low-energy degeneration, they differ significantly for the DNN energy estimator.

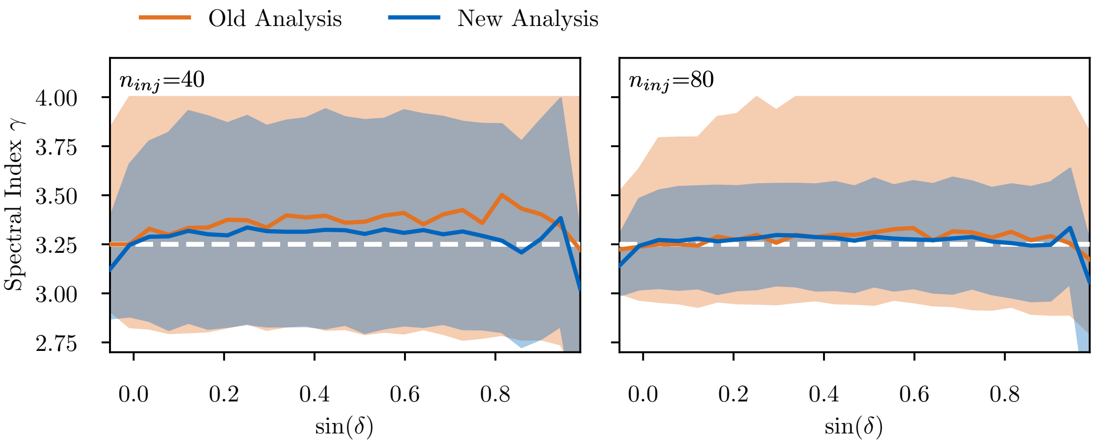

Figure 7.17: Recovery of the spectral index for the old and new method against the declination. The injected spectral index, $\gamma = 3.25$, is shown as white dashed line. The median and 68% quantiles of the outcomes for both methods are shown as solid line and shaded band, respectively. In the left plot 40 events are injected and in the right plot 80 events.

Although only two examples have been shown, the results are representative for a general trend, i.e., the new analysis provides (nearly) unbiased spectral index fits up to very high declinations ($\delta \sim 81$ degrees, $\sin(\delta) = 0.9876$) while having a variance that is lower than the previous analysis. The remaining biases result from neglecting the angular uncertainty dependency in the energy and background term of the likelihood, eq. (7.27). At very high declinations, fitting problems are caused by rapid changes in event rates and reconstruction quality as many events go only along a single string of the detector. This problem can be potentially cured in the future by removing single-string events from the event selection.

7.5.2 *Source Strength Bias*

Apart from the spectral index, the number of signal events, i.e., the flux normalization is the second important property of a neutrino source. In Figure 7.18 and Figure 7.19 the fitted source strength of the old and new analysis are shown in comparison for an injected spectrum of $\gamma = 2.0$ and $\gamma = 3.25$ at two different declinations, respectively. In both figures, the left panels show the fitted signal strength against the number of injected events, while the right plot shows the corresponding evolution of the spectral index. As expected, the new, KDE-based method provides unbiased fits in all cases. It thereby solves major problems of the old analysis as discussed in the following.

In the case of a $\gamma = 2.0$ spectrum, the old method is clearly overfitting the number of signal events, a trend that increases with the number of injected signal events. While this appears to be contradictory to the underestimation of the tails in the spatial term, it can be understood through the bias towards softer spectral indices in Figure 7.15, as tilting the spectrum towards softer indices will necessarily manifest in a larger number of low-energy events and thereby

Figure 7.18: Fitted versus injected signal for a $\gamma = 2.0$ source and two different declinations. Left and right plot show the recovery of the number of signal events and the spectral index as a function of the number of injected events, respectively. The unbiased expectations are shown as white dashed lines. The results for the new and old analysis are shown in different colors. Solid lines and shaded bands show the median and central 68% quantile.

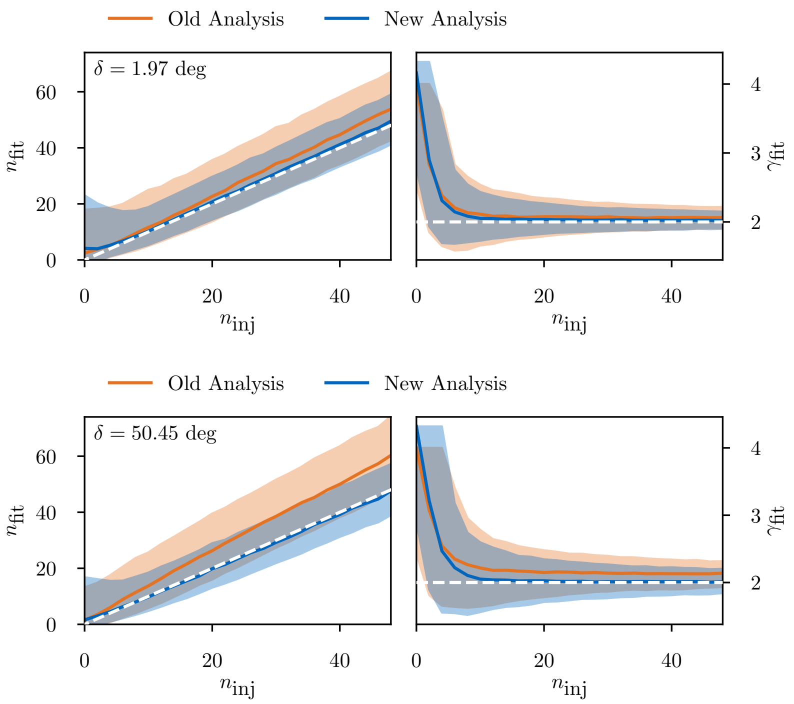

Figure 7.19: Fitted versus injected signal for a $\gamma = 3.25$ source spectrum and two different declinations. Left and right plot show the recovery of the number of signal events and the spectral index as a function of the number of injected events, respectively. The unbiased expectations are shown as white dashed lines. The results for the new and old analysis are shown in different colors. Solid lines and shaded bands show the median and central 68% quantile.

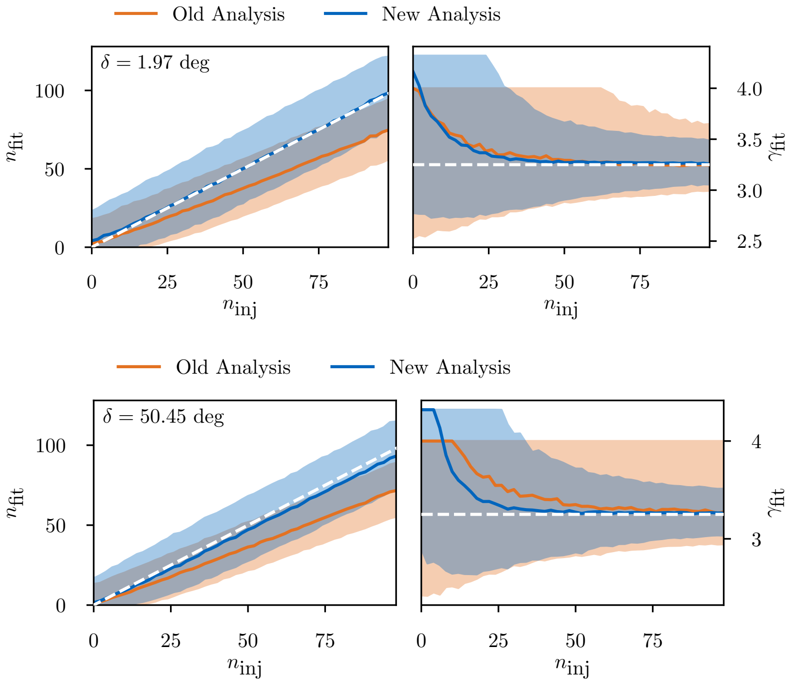

a higher n_{fit}. Hence, in consistency with the declination behavior of the spectral index, also the source strength bias increases with higher declination.

For soft spectral indices, as the exemplary $\gamma = 3.25$ case, the old analysis continuously underestimates the true number of injected signal events. Since we had seen previously that for these injection spectra the spectral index fit is unbiased, this is inevitably an effect of the mismodeling of the spatial term in the likelihood. Specifically, events that are reconstructed further away from the source position can not be correctly recovered. As for the case of the harder spectrum before, the bias increases with injected source strength. In fact, going beyond 100 injected signal events, the truth is not even contained inside the central 90% range of outcomes.

7.6 Impact of Systematic Uncertainties on the Analysis

In section 7.3 the impact of ice uncertainties on the three point source observables — DNN energy, Spline MPE direction, and BDT angular uncertainty — have been discussed. In this section, we want to discuss how the uncertainties impact the parameter estimation of the full analysis. Therefore we repeat the evaluation of fit biases, as shown for the baseline set in the previous sections, on the different Monte Carlo datasets with systematic ice model variations. As the data sample of the point source analysis is dominated by atmospheric neutrinos, it is important to have consistent background modelling to prevent the analysis from artificially picking up unmodeled background components. Specifically, it is required that the background model complies with the strong data constraints. Similarly, to the baseline model, all systematic variations are therefore corrected to match the experimental data perfectly. This procedure is thereby equivalent to using a hybrid of baseline and systematics Monte Carlo for background and signal injection, respectively.

The resulting fit bias plots are shown in Figure 7.20 and Figure 7.21 for two exemplary declinations. Most importantly, it can be seen that the influence of ice systematics on the fit parameters is small - independent of the spectral index and the declination. This guarantees that the analysis results are stable even if the true ice properties deviate from their baseline values. The same also applies for a hypothetical prompt component that is found negligible for the analysis.

Figure 7.20: Fitted versus injected signal for a $\gamma = 2.0$ source spectrum and two different declinations. The left and right plot show the recovery of the number of signal events and the spectral index as a function of the number of injected events, respectively. The unbiased expectation is shown as a white dashed line. The results for different systematic effects are shown in different colors. Solid lines and shaded bands show the median and central 68% quantile.

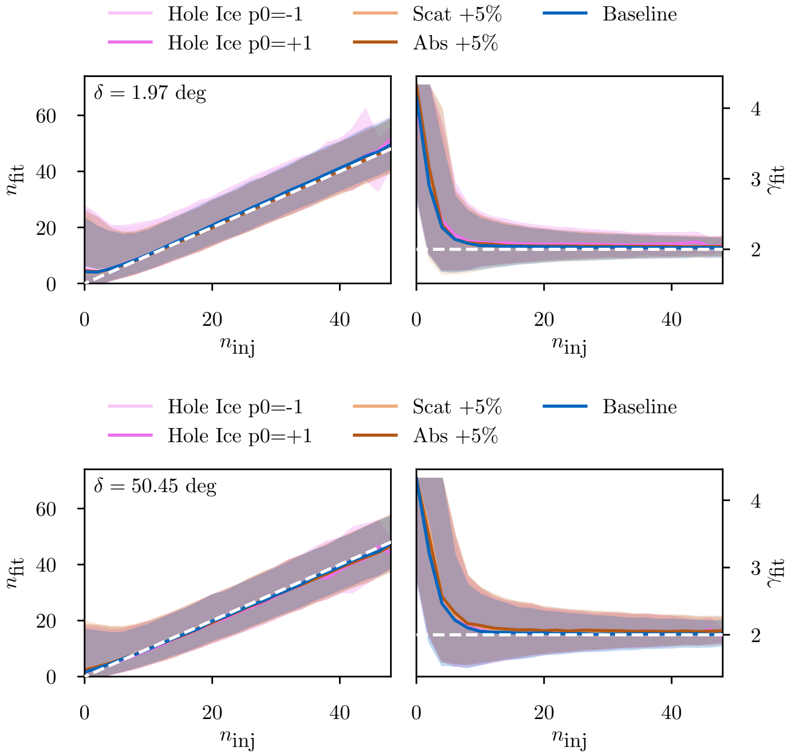

Figure 7.21: Fitted versus injected signal for a $\gamma = 3.25$ source spectrum and two different declinations. The left and right plot show the recovery of the number of signal events and the spectral index as a function of the number of injected events, respectively The unbiased expectation is shown as a white dashed line. The results for different systematics effects are shown in different colors. Solid lines and shaded bands show the median and central 68% quantile.

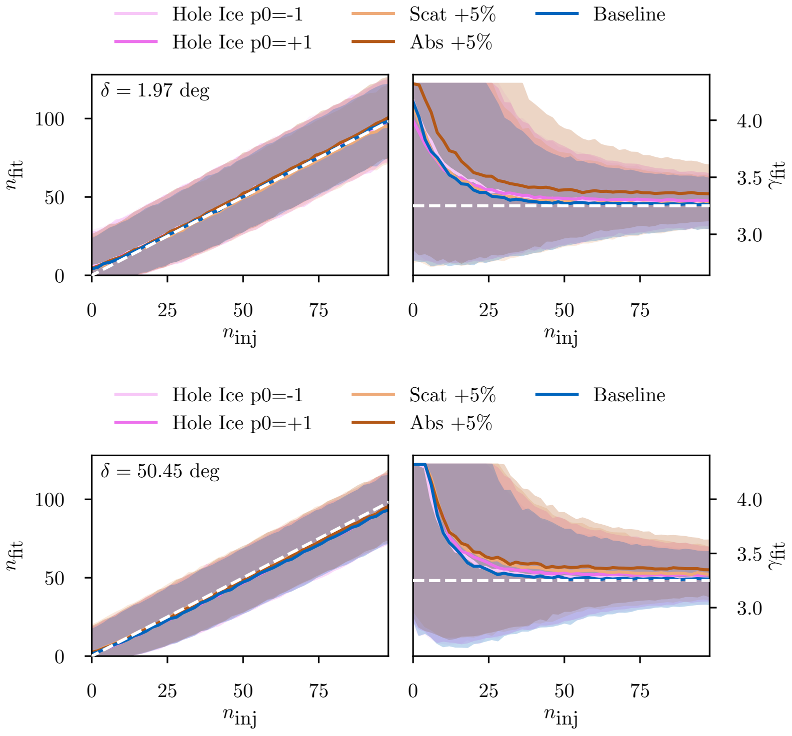

7.7 *Sensitivity and Discovery Potential*

Besides the recovery of the signal properties, the sensitivity and discovery potential are two important measures for the analysis performance. They are defined as follows

- **Sensitivity** is the source strength μ_s needed to exceed the median value of the background test-statistic distribution in 90% of the cases. Thereby it makes a statement to the median 90% upper limit assuming a non-observation of a signal.
- The **discovery potential** describes the power of the analysis to *detect* a signal. It is defined as the source strength needed to exceed a certain quantile (usually 3σ or 5σ) of the background test statistic distribution in 50% of the cases.

As a first step, both discovery potential and sensitivity, as well as the calculations of p-values in the final analysis, require the generation of background test statistics distributions as shown in Figure 7.22 for various declinations and 500,000 trials, respectively. It can be seen that the distributions are slightly steeper than the naive $0.5 \times \chi^2_2$ expectation from Wilks' theorem (theorem 7.1.2) [46]. In addition to the previous discussions, we will also consider source hypothesis where the spectral index is fixed, as most astrophysical scenarios expected that neutrino source spectra are harder than the conventional atmospheric neutrino spectrum. Shock acceleration, for example, predicts sources with spectral index $\gamma \approx 2$. Under these circumstances, the likelihood has only one degree of freedom. According to Wilks' theorem, we, therefore, expect that the background test statistic distribution follows a $0.5 \times \chi^2_1$ distribution. It turns out, however, that for hard spectra ($\gamma = 2.0$) the fraction of 0 trials increases to around 80% (70% for $\gamma = 2.5$) and the test statistic distribution is falling down more steeply, see Figure D.1 and Figure D.2 in the appendix.

The overall tendency to flatter test statistic distributions for all cases - fixed and free spectral index - can be explained by Wilks' theorem requiring the sample size of $n \to \infty$. Specifically, in the case of hard, fixed spectral indices the (lower energy) background events do not sufficiently sample the signal space to fulfill this requirement. While the deviations from the χ^2 expectation require additional simulations to derive accurate background test statistic distribution, their steepness also increases the analysis's background rejection power.

[46] Note that for the analysis a truncated gamma function is fitted as it has more flexibility to describe the data. This is specifically important to correctly calculate the p-value for the hottest spots deep in the tail of the test statistic distributions.

Figure 7.22: Background test statistic distribution for various declinations. The χ_2^2 expectation — as expected from Wilks' Theorem — is indicated as a black dashed line after normalizing for the number of zero trials N_0, which is around 50% to 60%.

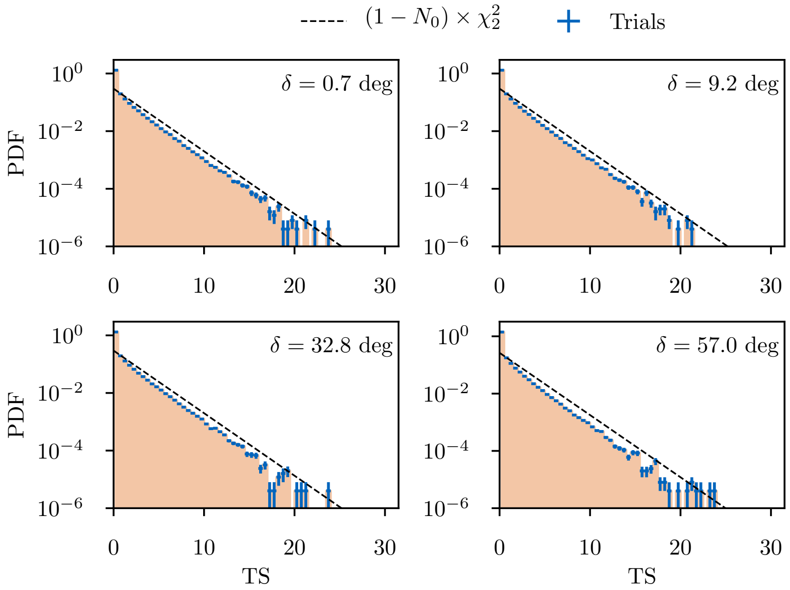

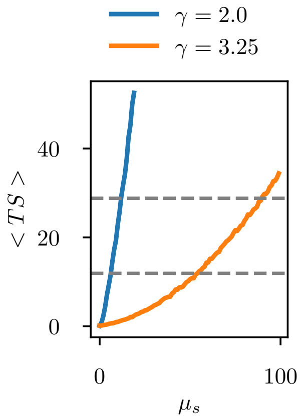

Figure 7.23: The median test statistic as a function of the signal strength at declination $\delta = 5°$ and 9 years of data. Different colors represent different spectral indices, and the dashed lines show the threshold for 5σ (top) and 3σ (bottom) significance.

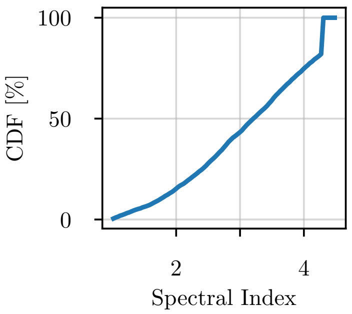

Figure 7.24: Cumulative spectral index distribution of background fluctuations with more than 4σ significance.

To calculate the sensitivity and the discovery potential, signal trials have to be generated for increasing source strength μ_s. In Figure 7.23 it is depicted how the 50% quantile of the signal test statistic distribution increases with μ_s. In a brute force approach, one could simply generate sufficient trials for various μ_s and calculate discovery potential and sensitivity through interpolation. This is, however, computationally ineffective as many unnecessary trials are generated. Instead, we use a more sophisticated binary search that generates a larger number of trials close to the respective test statistic threshold, increasing accuracy while generating fewer trials overall.

In Figure 7.26 the resulting sensitivity and discovery potential are shown for injected source spectra of $\gamma = 2.0$ and $\gamma = 3.0$, as well as three signal hypothesis with the different assumption on the spectral index: free, fixed to $\gamma = 2.0$, and fixed to $\gamma = 2.5$, respectively. Most prominently, the new analysis improves the discovery potential by $20\% - 30\%$ for an injected signal with $\gamma = 2.0$ (free spectral index case). Improvements are tendentially larger towards the poles. For the softer spectral index, the average improvement is smaller with only $\sim 10\%$. On the contrary, the sensitivity does not improve much compared to the old analysis — independent of the injected spectral index. This results from the sensitivity being dominated by statistical fluctuations of a few events that push the TS away from 0.

Figure 7.26 also shows the results for the fixed spectral index cases. While $\gamma = 3.0$ sources are better identified with the free spectral index analysis, the fixed spectral index of $\gamma = 2.0$ improves the discovery potential by another $\sim 10\%$ when searching for $\gamma = 2.0$ sources. Visually, fixing the spectral index to $\gamma = 2.0$ penalizes the background test-statistic values as background fluctuations are usually fitted to softer values, see Figure 7.24. On the contrary, signal test statistic values are not much affected by the spectral index constraint as the fit anyway recovers in the median the injected source spectrum.

The energy dependence of the discovery potential can finally be studied by using the same construction as above, but limiting the flux to a specific energy range. This results in the *differential discovery potential* as shown in Figure 7.25 for two bins in each energy decade. Most notably, the absorption of neutrinos in the Earth strongly reduces the discovery potential at higher declinations and energies.

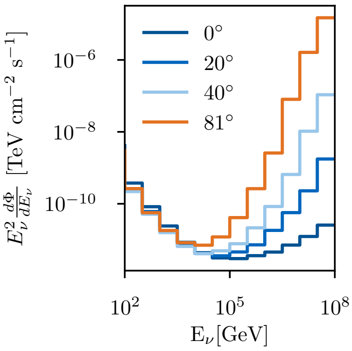

Figure 7.25: Differential 5σ discovery potential for different declinations, assuming an E^{-2} spectrum in each energy bin.

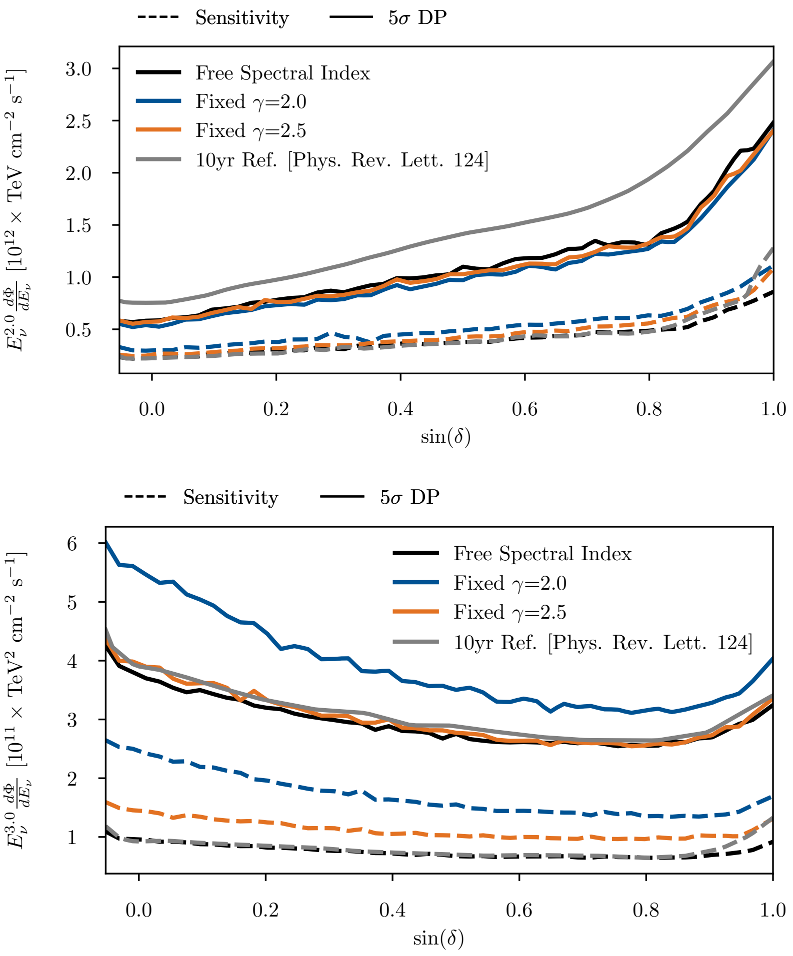

Figure 7.26: Sensitivity and discovery potential for an injected spectral indices of $\gamma = 2.0$ (top) and $\gamma = 3.0$ (bottom). The results for this 9 year IC86 analysis are also compared to the previously published 10 year point source result (including IC40, IC59 and IC79 seasons) with the old method (grey) as adopted from M. G. Aartsen et al., "Time-integrated Neutrino Source Searches with 10 years of IceCube Data".

7.8 Experimental Results

Note: The results presented in this section are preliminary. While the general picture will not change, numbers can still vary on the order of numerical precision in future publications.

Finally, in this section, the new point source likelihood is used to analyze the 9 years sample of up-going muon tracks[47] for a neutrino point source signal. Two types of tests are performed:

[47] as described in section 7.3

1. A **sky scan** that covers the Northern Hemisphere between declination $-3^\circ \leq \delta \leq 81^\circ$ and evaluates the point source likelihood — eq. (7.24) with eq. (7.27) — on a grid with a pixel area of 0.052 square degree [48]. The resulting test statistic values are converted to p-values using background test statistic distributions generated at 230 declinations. For the 20 hottest spots from the initial scan, a fine scan is performed in a region of 0.75 degrees around the hot spot and a HEALPix nside of 2048 (0.0008 sqdeg per pixel). In addition to the case of the free spectral index, the same procedure is repeated for the two fixed spectral indices, $\gamma = 2.0$ and $\gamma = 2.5$. The hottest spot is reported as the most significant spot among all maps.

2. A **source catalog** search including 110 known gamma-ray sources. The choice of 110 sources guarantees that a 5σ pre-trial detection will not drop below a 4σ discovery after accounting for trials. For the construction of the catalog, the integrated gamma-ray flux above 1 GeV of all 4FGL sources in the declination range $-3^\circ \leq \delta \leq 81^\circ$ is weighted with the IceCube sensitivity at the respective source declination. Subsequently, the 5% highest weighted BL Lacs and FSRQs are then added to the source list. On top, all 6 4FGL starburst galaxies with suitable declination are added as they have been predicted to produce neutrinos via proton-proton interaction.[49] Finally, one Galactic source, MGRO J1908+06, is added as its expected neutrino emission assuming $\phi_\nu \propto \phi_\gamma$ is compatible with the differential sensitivity in Figure 7.25. We report on the hottest spot of the catalog.

The results of the sky scans are shown in Figure 7.31. When interpreting the sky scan's local p-values, it is important to keep in mind that each new position in the scan corresponds to a new source hypothesis. Although not all of these positions are statistically independent, scanning the entire sky, therefore, produces a large number of independent trials. Consequently, the interpretation of the hottest local p-value requires a correction for the probability of observing a hot spot at least as significant as the experimental one somewhere in a sky scan, assuming the background hypothesis to be true. The required distribution of smallest p-values from 1000 background pseudo-experiments and subsequent sky scans is shown in Figure 7.27 for the case of the free spectral index and in the appendix, Figure D.11 and Figure D.12, for the fixed spectral index cases.

In order to describe the distribution of p-values, one can use order statistic, which states that the kth smallest p-value from an independent sample of N trials follows a beta distribution,[50] i.e.,

$$p_{pre} \sim Beta(k, N+1-k). \tag{7.50}$$

In our case we search for the hottest spot in a skymap and hence $k = 1$. N can be interpreted as an *effective* number of independent trials. Fitting eq. (7.50) to the respective p-value distributions results in N = 202,299 (free spectral index), N=333,683 ($\gamma = 2.0$) and

[48] corresponding to a HEALPix grid with $nside = 256$, see Gorski et al., "HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere"

[49] Murase, Ahlers, and Lacki, "Testing the Hadronuclear Origin of PeV Neutrinos Observed with IceCube".

[50] Gentle, *Computational Statistics*, p.63.

N=409,883 ($\gamma = 2.5$) independent trials. As expected the number of independent trials is fairly large, being on the order of one independent spot per 0.1-0.2 sqdeg — approximately the size of the area covered by a median directional uncertainty of ~ 0.2 deg. Noticeably, the number of independent trials is smaller for the free spectral index case. This is consistent with the results of the sky scan in Figure 7.31 which clearly shows that neighbouring pixels are much more washed out and therefore more correlated in this case. Finally, Figure 7.27 also shows the resulting cumulative distributions that is needed to convert the pre-trial into a post-trial p-value. Analytically it can be described as[51]

$$p_{post} = 1 - (1 - p_{pre})^N \tag{7.51}$$

which for $p_{pre} \ll 1$ can be approximated as

$$p_{post} \approx N \cdot p_{pre}. \tag{7.52}$$

Hence, a pre-trial p-value of 7σ is need for a 5σ discovery.

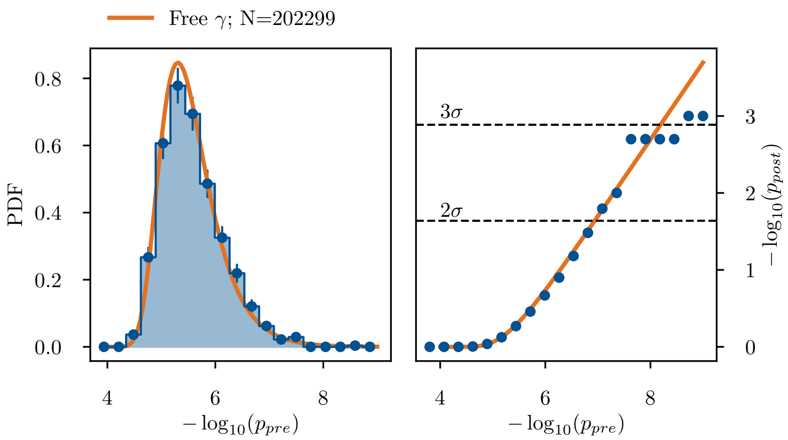

Figure 7.27: The trial correction function for the case of a free spectral index hypothesis. The left plot shows the distribution of p-values from 1000 pseudo-experiments assuming the background hypothesis to be true. The fitted beta distribution is shown as orange line. The corresponding cumulative function that can be used to convert the pre-trial into a post-trial p-value is shown in the right plot.

After correcting the smallest pre-trial p-value from each sky scan with the respective trial correction function, the analysis's hottest spot is given as the most significant spot among all three maps. It is found in the free-spectrum spectrum map at right ascension 40.69° and declination 0.09° (J2000) with a local pre-trial p-value of $p_{loc} = 8.64 \times 10^{-8}$ (5.2 σ). In addition to the previously discussed trial correction for the sky scan, also a trial correction factor for testing three different spectral index hypotheses needs to be considered. Using simulations, we find that this correction factor is 1.6, which is smaller than 3 as the maps are highly correlated. Combining all trial correction factors, the post-trial p-value is $\sim 1.9\sigma$ and therefore by itself not significant when the entire sky is scanned without any prior assumption. A fine scan around the best-fit position of the hottest spot is shown in Figure 7.28.

In addition to the neutrino hotspot, Figure 7.28 also shows the presences of a 4FGL source spatially consistent with the neutrino emission. This source — a Seyfert type-II starburst galaxy NGC 1068

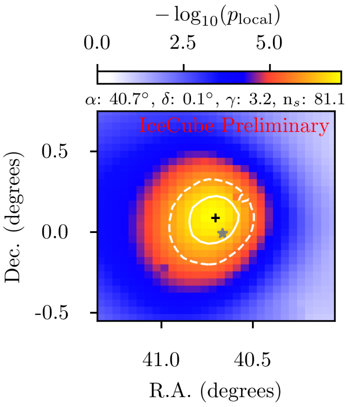

Figure 7.28: A fine scan around the location of the sky scan's hottest spot. The black cross indicates the hottest spot position with respective coordinates given above the plot. 4FGL sources are indicated as grey stars. The solid and dashed white contours give the 68% and 95% confidence regions of the hot spot localization as calculated from Wilks' Theorem with two degrees of freedom.

---

[51] Casella and Berger, *Statistical Inference*, p.229.

— is also part of our source catalog and, therefore, unsurprisingly its most significant candidate. At the multi-frequency position of NGC 1068 (r.a.: $-0.007°$ dec: $40.667°$) we find a local p-value of $p_{loc} = 1.64 \times 10^{-7}$ (5.1 σ), with best-fit values of spectral index $\gamma = 3.19$ and number of signal events $n_{fit} = 78.78$. In fact, this hotspot has already been seen in a previous publication but was shifted by around $0.35°$ further south in declination.[52] With the new analysis, the localization of the neutrino hotspot is well consistent with the position of NGC 1068, which manifest in a jump in the post-trial p-value from $2.9\,\sigma$ to $4.2\,\sigma$ after accounting for the 110 trials from the source catalog. **Hence, NGC 1068 is the first candidate of astrophysical neutrinos with post-trial significance above 4σ.**

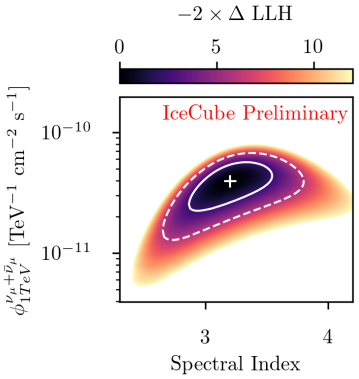

Figure 7.29: Profile likelihood scan for the flux parameters of NGC 1068. The cross shows the best-fit values, solid and dashed lines represent 68% and 95% confidence levels derived from Wilks' Theorem, respectively. The contours include only statistical uncertainties.

To better characterize the spectrum of the source, the likelihood function is scanned around the minimum in source normalization and spectral index resulting in the profile likelihood shown in Figure 7.29. Using Wilks' Theorem, 68% and 95% confidence intervals are drawn assuming the likelihood difference to follow a χ^2_2 distribution, which has been validated using Monte Carlo simulations.

Finally, the profile likelihood enables us to assemble the source's full SED, Figure 7.30, including all the data from the multi-frequency analysis tool as described in section 8.2. The uncertainty contour of the neutrino emission is thereby derived by calculating the energy-dependent minimum and maximum flux of all fluxes that lay on the 68% contour of Figure 7.29. While the soft neutrino spectrum of NGC 1068 is against the usual expectation of hard neutrino source spectra, there are already models that predict this kind of emission for neutrino production close to the supermassive black hole in the center of the galaxy[53].[54] These models can also explain the discrepancy between neutrino and gamma-ray flux through absorption of gamma-rays in the X-ray coronae of the source. Nevertheless, further theoretical studies of the source will be necessarily to unambiguously identify the underlying processes of the neutrino production.

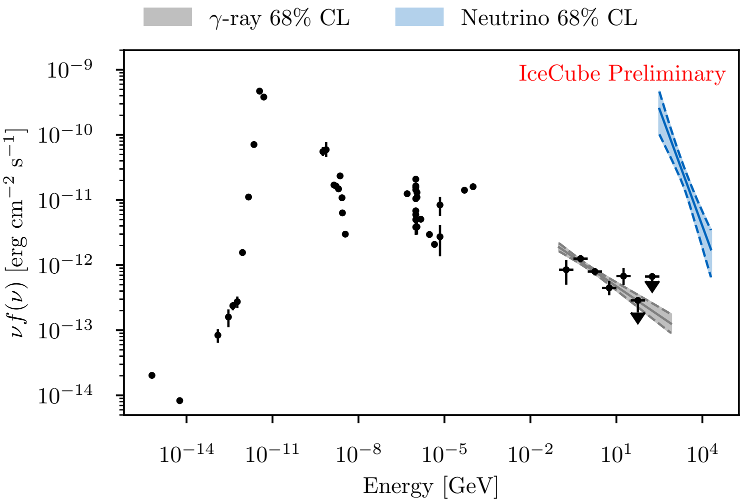

Figure 7.30: Spectral energy distribution of NGC 1068. Black points show archival multi-frequency measurements. The *Fermi*-LAT gamma-ray spectrum (grey with black SED points) is based on an analysis integrating over the same 9 years as the neutrino data. The $\nu_\mu + \bar{\nu}_\mu$ neutrino spectrum is shown as the blue band in the energy range between 300 GeV and 20 TeV.

[52] M. G. Aartsen et al., "Time-integrated Neutrino Source Searches with 10 years of IceCube Data".

[53] Inoue, Khangulyan, and Doi, "On the Origin of High-energy Neutrinos from NGC 1068: The Role of Nonthermal Coronal Activity".

[54] Murase, Kimura, and Meszaros, "Hidden Cores of Active Galactic Nuclei as the Origin of Medium-Energy Neutrinos: Critical Tests with the MeV Gamma-Ray Connection".

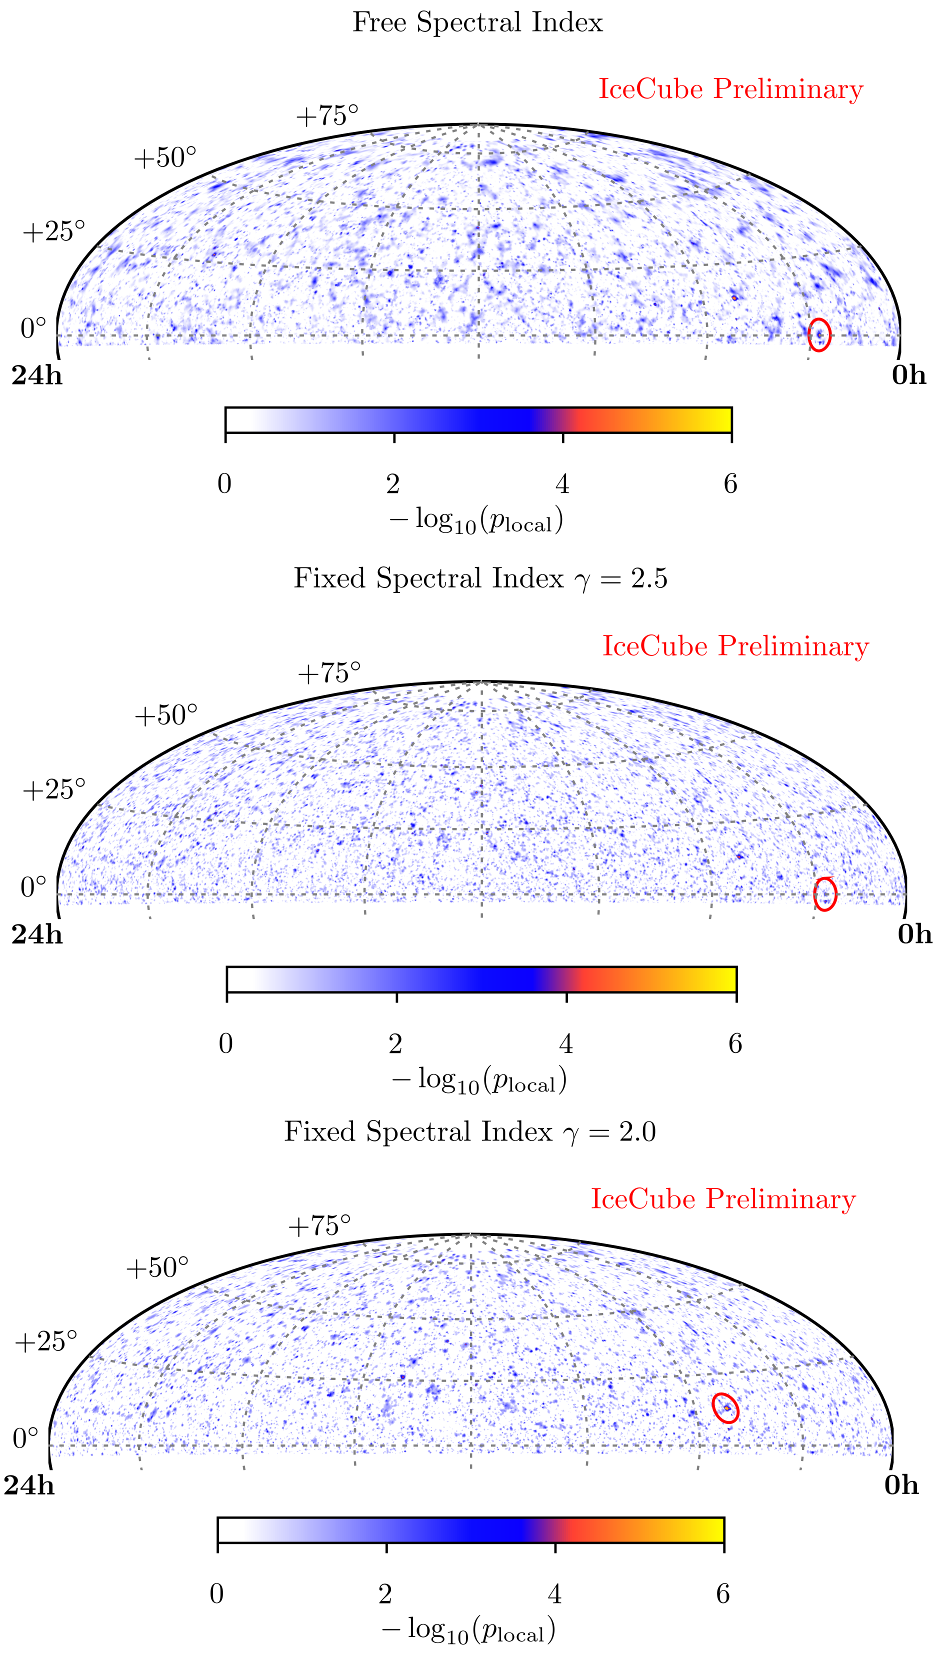

Figure 7.31: Skymaps for the different signal hypotheses. The maps are shown in Hammer-Aitoff projection with the color scale representing the local p-value. The red circles indicate the hottest spot positions.

8 Multi-Messenger Searches for Neutrino-Emitting Blazars

In this thesis, we have presented the first compelling sources of astrophysical neutrinos — the masqueraded BL Lac TXS 0506+056 in chapter 4 and the starburst galaxy NGC 1068 in section 7.8. Overall, those two sources can, however, only explain a small fraction of the entire astrophysical neutrino flux. Population studies of hot spots in the neutrino sky actually indicate contributions from more than 1000 neutrino sources[1].[2] Additional search strategies are therefore required to identify a larger population of neutrino sources. In general, astrophysical neutrinos are expected to be produced along with a broad range of multi-frequency electromagnetic emission (see section 2). Hence, requiring a spatial connection of neutrinos with other messengers, for example, gamma rays, can significantly reduce trials compared to model-independent neutrino-only searches.

Blazars, as a subclass of active galactic nuclei with a jet pointing towards the observer, have often been proposed as good candidates for the production of astrophysical neutrinos (see section 2.3.3). Despite their acceleration power, which is manifested in the bright synchrotron peak of their spectral energy distribution, they are also known to produce a significant fraction of the extragalactic 100 MeV to 100 GeV gamma-ray emission (see Figure 2.6). In fact, TXS 0506+056, the first compelling neutrino source candidate (see chapter 4), is a blazar. Dedicated IceCube Collaboation studies have shown that the contribution from the 3FHL [3] blazar catalog to the diffuse IceCube neutrino flux is limited to 14.4 % ($\gamma = 2.0$) and 22.3% ($\phi(E) \propto E^{-1.0} \cdot \exp(E/E_0)$; $E_0 \geq 1\,\mathrm{PeV}$).[4] In contrast, the slightly older, but lower energy threshold catalog, 2LAC, allows for a contribution of up to 66.9% ($\phi(E) \propto E^{-1.0} \cdot \exp(E/E_0)$; $E_0 \geq 1\,\mathrm{PeV}$).[5] While these results indicate that blazars are not the only neutrino emitters, they can still be responsible for a large majority of the flux — specifically at energies above 100 TeV.

This chapter presents a new statistical analysis, searching for gamma-ray blazars in the error regions of the IceCube Observatory's high-energy neutrinos. It is thereby the natural extension of the dissection analysis presented in chapter 4. For future applications, the analysis pipeline has been automatized and optimized to perform a semi-realtime dissection of all the publicly available data in the region around neutrino realtime alerts. It thereby poses a powerful tool to

[1] Glauch and Turcati, "Search for weak neutrino point sources using angular auto-correlation analyses in IceCube".

[2] M. G. Aartsen, K. Abraham, et al., "All-sky Search for Time-integrated Neutrino Emission from Astrophysical Sources with 7 yr of IceCube Data".

[3] Blazars detected in *Fermi*-LAT above a photon energy of 10 GeV

[4] Huber, "Searches for steady neutrino emission from 3FHL blazars using eight years of IceCube data from the Northern hemisphere".

[5] Huber, "Multi-Messenger correlation study of Fermi-LAT blazars and high-energy neutrinos observed in IceCube".

identify source targets for multi-frequency follow-ups in the proximity of IceCube neutrinos. In addition to the analysis of the available multi-frequency data, the pipeline also generates gamma-ray test statistic maps to identify previously unknown sources. Finally, the multi-wavelength SED is assembled and analyzed using machine learning methods for all interesting candidates.

Making extensive use of multi-frequency data, the work presented in this chapter is complementary to the analysis presented in chapter 7. It also focuses only on the few publicly available events with the highest energies. While this reduces the sample size significantly, a large fraction of the (hard spectrum) astrophysical signal is preserved. In fact, if we assume that neutrinos are produced with hard spectra (e.g., $\phi(E) \propto E^{-1}$) — as expected from photo-meson production — IceCube would only observe an astrophysical signal of scattered high-energy events. This can be easily seen by looking at the IceCube event rates for a hard spectrum with $\phi(E) \propto E^{-1}$. In this case, the rate of neutrinos reaching the Earth is the same in every log-decade, i.e.,

$$\frac{\mathrm{d}N}{\mathrm{d}\log_{10}(E)} = E\frac{\mathrm{d}N}{\mathrm{d}E} = E \cdot \phi(E) \propto E \cdot E^{-1} = const. \tag{8.1}$$

Furthermore, the effective area of IceCube increases strongly in energy and hence hard spectrum sources manifest in single high energy neutrinos without accompanying low energy events. This is particularly true for the region around the horizon where the $\nu_\mu + \bar{\nu}_\mu$ effective area is around $200\,\mathrm{m}^2$ at $1\,\mathrm{TeV}$ while going up to around $4000\,\mathrm{m}^2$ at $100\,\mathrm{TeV}$ and $9000\,\mathrm{m}^2$ at $1\,\mathrm{PeV}$.[6] Hence, cross-correlation studies between multi-wavelength emitters and high-energy neutrinos provide an important, complementary channel to identify the common sources of the astrophysical messengers. This chapter summarizes and extends the results of a series of papers that are establishing new methods for multi-wavelength searches of astrophysical neutrino sources[7],[8].[9]

8.1 *Catalogs of Blazars*

Over the years, objects identified as blazars in various wavelength bands have been collected in catalogs. Three major catalogs that are frequently used throughout this chapter are briefly summarized in the following.

- **BZCAT:**[10] The latest, 5th-edition of the Roma-BZCAT is based on the multi-frequency detection of blazars with and without gamma rays. It contains a total of 3561 sources and is thereby the largest catalog of blazars. In addition to the source position and its redshift, also radio, microwave, optical, X-ray, and gamma-ray fluxes are given wherever available. Sources are further grouped into four classes: 5BZB (BL Lacs), 5BZQ (FSRQ), BSG (Blazars

[6] M. G. Aartsen et al., "Search for steady point-like sources in the astrophysical muon neutrino flux with 8 years of IceCube data".

[7] Padovani, Giommi, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source".

[8] Giommi, Glauch, et al., "Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection".

[9] Giommi, Padovani, Oikonomou, et al., "3HSP J095507.9+355101: a flaring extreme blazar coincident in space and time with IceCube-200107A".

[10] Massaro et al., "The 5th edition of the Roma-BZCAT. A short presentation".

with strong galaxy), and BZU (Blazars of uncertain type). The completeness of the BZCAT is unknown as the underlying multi-frequency surveys are not uniform.

- **4LAC**:[11] The fourth catalog of active galactic nuclei (*Fermi*-4LAC) contains the gamma-ray sources of the 4FGL catalog[12] that have an AGN counterpart, in total 2863 objects. 98% of those counterparts are blazars. As it is based on the 4FGL catalog, it has a uniform exposure and, therefore, the same completeness everywhere (outside the Galactic plane). In addition to the counterpart name, it also provides an estimate of the synchrotron peak.
- **3HSP**:[13] This catalogue contains high synchrotron peaked blazars, i.e., blazars with synchrotron peak ν^S_{peak}larger than 10^{15} Hz. The catalog contains 2013 objects in total. In addition to the radio, X-ray, and gamma-ray flux (where available), the redshift is given in 88% of the cases.

8.2 Blazar Counterparts of High-Energy Neutrinos

Historically, IceCube has published several lists of high-energy track-like events from different event selections including highly-energetic starting tracks (HESE)[14,15],[16] up-going muon tracks from the Northern Hemisphere[17] and the realtime stream[18] [19] [20]. With the publication of the discovery of TXS 0506+056, an additional list of archival alerts was provided.[21] While this sums up to a sample of ~ 100 events, no in-depth study of counterpart candidates had been presented so far. By combining all of those streams we have compiled a list of 70 unique events with angular resolution ≤ 3 deg and being off the Galactic plane ($|b| > 10$ deg), Table 8.3. For events appearing in two streams, the most recent reconstruction and errors are used. Note that errors are, however, not treated consistently across the sample. Overall, four different methods are being used by IceCube

- For a few very special events as *IceCube-141106A*, a ~ 6 PeV through-going muon neutrino,[22] and *IceCube-170922A*, the event in coincidence with TXS 0506+056,[23] a full **event re-simulation** was performed. Given the reconstructed quantities as direction, the entry point in the detector, and energy, similar events are generated from Monte Carlo with varying ice systematics. Subsequently, each of the events is reconstructed, and the $\Delta \log \mathcal{L}$ between reconstruction and truth is recorded. From this distribution, the 90% contours can be calculated as the region where the $\Delta \log \mathcal{L}$ of the experimental event is smaller than the 90% quantile of the distribution. Note that the choice of systematic ice variations is an ad-hoc approximation and not deduced from a clear procedure, leaving some uncertainty on the coverage of the contour.
- For all the events from the realtime stream, the likelihood space is scanned around the minimum. In contrast to the re-simulation, however, no dedicated calculation of the event-dependent $\Delta \log \mathcal{L}$

---

[11] Ajello et al., "The Fourth Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope".

[12] Abdollahi et al., "Fermi Large Area Telescope Fourth Source Catalog".

[13] Y.-L. Chang, Arsioli, et al., "The 3HSP catalogue of extreme and high-synchrotron peaked blazars".

[14] M. Aartsen et al., "Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector".

[15] M. Aartsen et al., "Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data".

[16] M. Aartsen et al., "The IceCube Neutrino Observatory - Contributions to ICRC 2017 Part II: Properties of the Atmospheric and Astrophysical Neutrino Flux".

[17] M. G. Aartsen et al., "Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data".

[18] M. Aartsen et al., "The IceCube Realtime Alert System".

[19] https://gcn.gsfc.nasa.gov/amon_icecube_gold_bronze_events.html

[20] https://gcn.gsfc.nasa.gov/gcn/amon_ehe_events.html

[21] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

[22] M. G. Aartsen et al., "Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data".

[23] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

distribution is performed. Instead, the $\Delta \log \mathcal{L}$ distribution of a previously re-simulated HESE event is used as a baseline. While this is presumably better than a simple $\mathcal{X}^2$ assumption, it is not clear how well this suits the individual event, as the $\Delta \log \mathcal{L}$ distribution *could* depend on parameters as the deposited energy, the energy loss pattern, the direction, and the event's position of entry in the detector. While there are some hints on how the mapping looks like,[24] no dedicated large-scale study has been performed yet. Similarly to this approach, the error for the high-energy tracks in the diffuse muon-neutrino search[25] is calculated by using a $\Delta \log \mathcal{L}$ scaling that only includes *statistical* uncertainties.

- For the remaining fraction of HESE events, a 50% uncertainty is calculated from the point-spread function as an ensemble-averaged quantity.

In summary, there is a non-negligible uncertainty on the coverage of IceCube's error estimations [26]. For this reason, the following cross-correlation study of IceCube high-energy tracks with blazar candidates will not only search in the published 90% error region (50% for a small number of HESE events) but also increase the contour by factors of 1.1, 1.3, 1.5 to not miss any signal due to overly optimistic error contours. Additionally, by definition, 10% of sources are expected outside the 90% error contour. While in the background-dominated case an enlargement of the contours would mainly increase the number of false associations, in the presence of small catalogs mainly signal is added, leading to an improvement in sensitivity. The error regions are denoted with $\Omega_{90\times i}$ where $i \in \{1.0, 1.1, 1.3, 1.5\}$.

To search for counterpart blazars, we have dissected all the 70 regions of the IceCube high-energy tracks as follows:

- Search for multi-frequency blazar candidates in 1.0, 1.1., 1.3, 1.5 times the elliptical error region using the *VOU Blazar* tool.[27] In the first phase, the tool queries $\sim$ 30 catalogs to collect known multi-frequency blazars or new blazar candidates based on a co-spatial emission of radio and X-ray.
- For the subset of candidates spatially consistent with a 4FGL source, the overall SED from radio to VHE gamma-rays is requested from $\sim$ 70 catalogs and visually inspected in detail. If the SED is consistent, the source is added to Table 8.4.
- Finally, the synchrotron peak, ν^S_{peak}, is estimated using a polynomial fit to the SED after removing contributions from the host galaxy (usually in IR), the blue bump of the accretion (optical+UV), and inverse Compton emission (X-Ray). The result is cross-checked with the machine learning SED classifier in chapter 8.4.1. Note that this leads to relevant deviations to the predictions in the 4LAC catalog,[28] where in some cases the blue bump is falsely interpreted as the ν^S_{peak}. Wherever possible, the redshift of the sources is also added, preferably from spectroscopic measurements; otherwise, from a fit to the blue bump of the host galaxy.

---

[24] Rädel, "Measurement of High-Energy Muon Neutrinos with the IceCube Neutrino Observatory".

[25] M. G. Aartsen et al., "Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data".

[26] not on the best-fit direction itself, which is - at least for tracks - fairly stable against systematics, due to their long lever arm.

[27] Y.-L. Chang, Brandt, and Giommi, "The Open Universe VOU-Blazars tool".

[28] Abdollahi et al., "Fermi Large Area Telescope Fourth Source Catalog".

The resulting list of blazars associated with high-energy neutrinos is shown in Table 8.4. While the list itself is a major step by providing for the first time a list of possible counterpart objects to IceCube's high-energy neutrinos, it also allows searching for statistical excesses in the classes of LBLs or IBLs/HBLs. The classification into these classes is purely based on the accretion power measured via the ν^S_{peak}, i.e., objects are classified as LBL if $\nu^S_{\text{peak}} < 10^{14}$ Hz, IBL for 10^{14} Hz $\leq \nu^S_{\text{peak}} < 10^{15}$ Hz and HBL 10^{15} Hz $\leq \nu^S_{\text{peak}}$. Due to the optical peak of the host galaxy at around 10^{14} Hz, the sparse multi-wavelength data, and the time variability of the ν^S_{peak}it is often difficult to separate between IBL and HBL. Hence we group them into one single class.

8.3 Counterpart Statistics and Discussion

In the previous section, a method for the identification of gamma-ray visible multi-wavelength blazars has been presented. In contrast to a simple catalog search, additional blazar counterparts have been identified, and the synchrotron peaks have been determined more precisely.

To compare the number of observed objects in the class of LBLs and IBLs/HBLs with the background expectation, we first need to calculate the background source density for each of these classes following the same procedure. Assuming an isotropic distribution of extragalactic sources, we redo the same dissection in 27 square degrees control regions centered $\pm 6^\circ$ in declination from the observed neutrinos. For consistency, we still avoid regions where the center lays inside the Galactic plane ($|b| < 10$ deg) or overlaps with any other neutrino error region. The resulting control area covers a total of 2,573 square degrees. In total, 103 gamma-ray LBL blazars and 103 gamma-ray IBL/HBL blazars are found in the control area, leading to the same expected average density of one object every 27.4 square degrees for both classes. Two tables summarizing the results of the dissections of the neutrino error regions and the control sample are shown in Table 8.1 and 8.2.

Area searched (1)	γ-ray IBL/HBL found in neutrino error region (2)	Expectation from control sample (3)	Likelihood-test p-value (4)
Ω_{90}	20	11.9	7.4×10^{-3}
$\Omega_{90\times1.1}$	24	14.4	1.4×10^{-2}
$\Omega_{90\times1.3}$	35	20.1	1.9×10^{-4}
$\Omega_{90\times1.5}$	47	26.8	2.0×10^{-4}

Table 8.1: Summary of γ-ray HBL/IBL blazars within the 70 IceCube tracks and comparison with the expectations due to random coincidences as estimated from the control sample. Adapted from Giommi, Glauch, et al., "Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection"

While the numbers of observed sources for the case of LBL is pretty consistent with the background expectation there is a clear over fluctuation of IBL/HBL sources. In order to calculate the statistical significance we have developed a method that properly takes

Table 8.2: Summary of γ-ray LBL blazars within the 70 IceCube tracks and comparison with the expectations due to random coincidences as estimated from the control sample. Adapted from Giommi, Glauch, et al., "Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection"

Area searched (1)	γ-ray LBL found in neutrino error region (2)	Expectation from control sample (3)	Likelihood-test p-value (4)
Ω_{90}	9	11.9	0.43
$\Omega_{90\times1.1}$	15	14.4	0.44
$\Omega_{90\times1.3}$	17	20.1	0.48
$\Omega_{90\times1.5}$	24	26.8	0.33

into account that the same error regions can include several counterpart candidates while the neutrino itself can only be associated to one object. The general idea is to count the number of regions with $0, 1, 2, \ldots$ counterparts and add them to a vector $\theta = (n_0, n_1, n_2, n_3, n_{\geq 4})$. Given the low background expectation, regions with more than 4 sources are extremely rare and are therefore combined to a single bin in order to simplify computations, while not losing significant information. Based on this, a likelihood ratio test is designed via

$$\lambda = \frac{\mathcal{L}_{H0}}{\mathcal{L}_{H1}} = \frac{p(n_1, n_2, n_3, n_{\geq 4} \,|\, \mathrm{H0})}{\max_{N_{srcs}} p(n_1, n_2, n_3, n_{\geq 4} \,|\, \mathrm{H1})} \tag{8.2}$$

with test-statistic $\mathcal{TS} = -2 \times \log \lambda$. The probability density function p is independent of n_0 as the number of regions with zero counterparts is unambiguously determined given the number of regions with 1,2,3, $\geq$ 4 counterparts. The signal and background hypothesis are defined as follows

- **Background Hypothesis (H0)**: The sources found in the neutrino error regions are random coincidences. Consistently the expectation value for the number of sources in a region is the product of the average source density as estimated from the control region, and the size of the error region.
- **Signal Hypothesis (H1)**: In addition to the random occurrence of background sources, a total of N_{srcs} signal sources are observed in the error regions.

Evidently, H_0 is a special case of H_1 with $N_{srcs} = 0$ and hence the hypothesis is nested. In order to calculate the pdfs in equation (8.2) Monte Carlo simulations are run for every relevant number of signal sources $N_{srcs} \in \{0, 1, \ldots, 40\}$ and scaling of the error region ($\Omega_{90\times1.0}$, $\Omega_{90\times1.1}, \Omega_{90\times1.3}, \Omega_{90\times1.5}$) using the 70 events in Table 8.3 as a representative sample. The simulation procedure for each trial is based on the following steps:

1. For each of the event the background expectation μ_{bg} is calculated as the product of the average source density (as estimated from the control region) and the size of the error region.
2. Background sources are drawn for each neutrino event based on a Poisson distribution with expectation value μ_{bg}

3. N_{srcs} unique signal sources are randomly distributed over the events.

To get a smooth distribution of the 4-dimensional pdfs, around 10 million trials are generated for each N_{srcs} [29]. Using the calculated pdfs, a background test statistic distribution can be generated. As expected for the case of a nested hypothesis with a single free parameter, Wilks' theorem (theorem 7.1.2) is applicable. Hence, the test statistic distribution follows a $\mathcal{X}_1^2$ distribution with one degree of freedom.

Finally, the experimental test statistic values are calculated and compared to the background test statistic distribution. The resulting p-values are shown in the last column of Table 8.1 and Table 8.2 for the various cases. The strongest signal is found for the class of IBLs/HBLs in 1.3 times the 90% error region, with a best-fit of $\hat{N}_{srcs} = 15 \pm 3.6$ and a pre-trial p-value of 0.019%, equivalent to 3.56 σ in a one-sided test. The corresponding profile likelihood is shown in Figure 8.1 with the 1,2 and 3 sigma regions calculated on basis of a $\mathcal{X}_1^2$ distribution which has been validated using signal simulations[30].

In the last step, the results are corrected by the *effective* number of independent trials. As the regions, and thereby the p-values of the different scalings are highly-correlated, each class's trial factor should be smaller than 4 as expected for the uncorrelated case. Again, the actual number can be calculated through Monte Carlo simulation, where for each trial the *overall* p-value is the *minimum* of all the p-values of the 4 different error regions ($\Omega_{90\times1.0}$, $\Omega_{90\times1.1}$, $\Omega_{90\times1.3}$, $\Omega_{90\times1.5}$). Asking how many times we find a p-value as low as 0.019% from pure background considerations in *any* of the regions, we obtain a trial correction factor of 1.64. To further account for the two independent blazar scenarios, this is multiplied by two, resulting in an overall trial correction factor of 3.28 and a post-trial p-value for the IBL/HBL excess of 6.2×10^{-4} (3.23 σ). In contrast, for the case of LBL objects, no excess has been found, allowing to set an upper limit of 3.48 sources at 90% C.L.

In order to evaluate if the statistical excess of the IBL/HBL objects is consistent with a signal, the expected distribution of counterparts $\theta = (n_0, n_1, n_2, n_3, n_{\geq 4})$ is shown in Figure 8.3 for the case of $\Omega_{90\times1.3}$. While the grey band indicates the pure background case, the blue band adds the best-fit number of 15 signal sources. In both cases, the dotted line represents the median value and the shaded area the standard deviation. In addition, grey and blue points show the experimental results for the HBL/IBL and the LBL class, respectively. While the latter one is, as expected, well compatible with the background expectation, the HBL/IBL distribution aligns well with the signal case without showing any systematic deviations.

In summary, the analysis of the 70 archival IceCube tracks provides convincing arguments that HBL/IBL gamma-ray blazars produce a significant fraction of the TeV to PeV astrophysical neutrino flux. Using the best-fit number of 15 sources, a lower limit on the fraction can be set to $\sim 21\%$ of the astrophysical high-energy events.

[29] This is another reason to use the overflow bin for regions with ≥ 4 sources. Going in even higher dimensions would make the generation of smooth pdfs quite challenging.

[30] See Cowan et al., "Asymptotic formulae for likelihood-based tests of new physics" p. 130 for the statistical theory

Figure 8.1: Profile likelihood for the identified LBL counterparts. The solid blue line shows the $-2 \times \Delta LLH$ values for different number of signal sources. The red point shows the best-fit value. Using a $\mathcal{X}_1^2$ distribution, upper limits can be calculated as indicated through the grey, dashed lines.

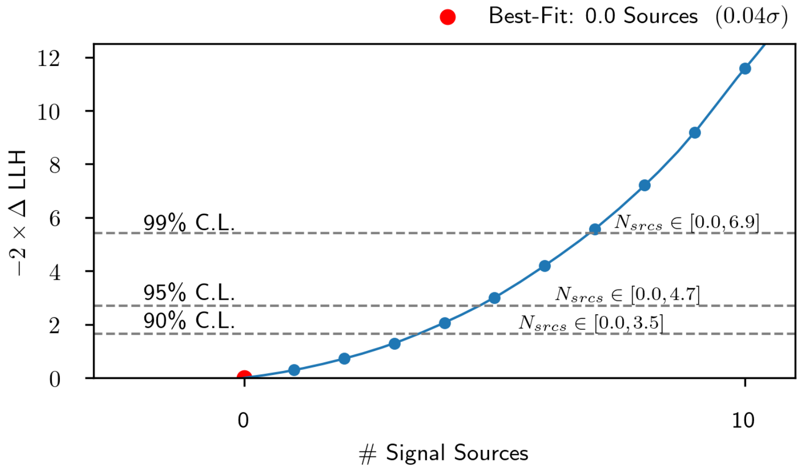

Figure 8.2: Profile likelihood for the identified IBL/HBL counterparts. The solid blue line shows the $-2 \times \Delta LLH$ values for different number of signal sources. The red point shows the best-fit value. Using a $\mathcal{X}_1^2$ distribution, central limits can be calculated as indicated through the grey, dashed lines.

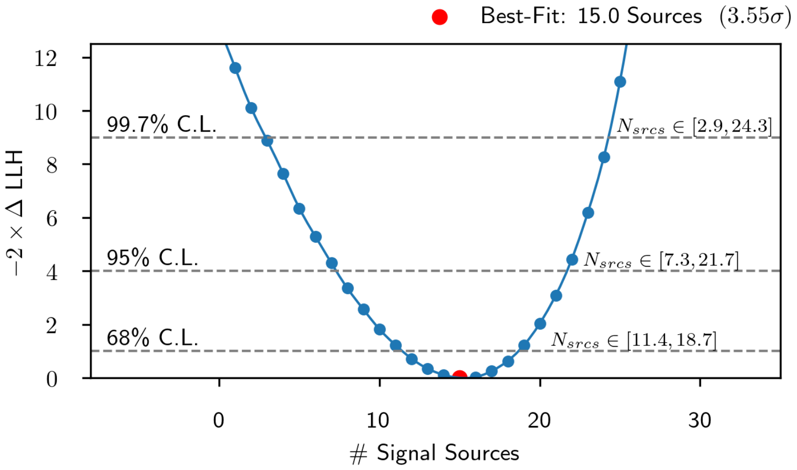

Figure 8.3: The expected distribution of the number of counterparts for pure background and background plus 15 signal events in grey and blue, respectively. The dashed lines show the mean and the shaded area the standard deviation from Monte Carlo simulation. The experimental data are shown as dots. In the bottom panel the deviation of the data from the background expectation is shown in Gaussian sigma.

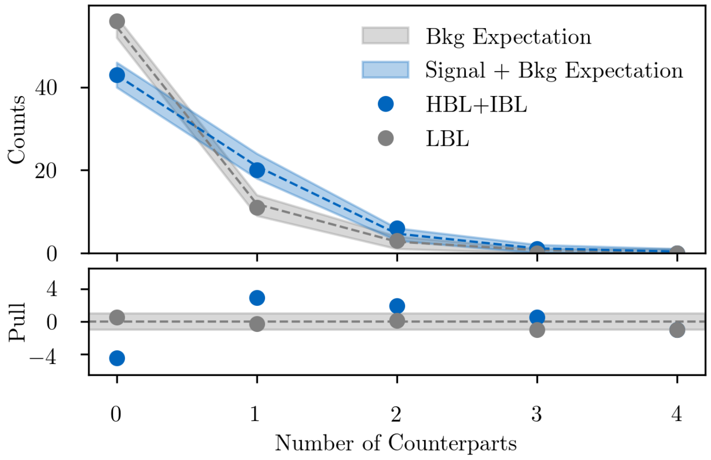

The true value could be higher as not all the 70 events are expected to be of astrophysical origin. The calculated limit is consistent with the IceCube limit for the 2LAC catalog which constrains hard-spectrum ($\phi(E) \propto E^{-1.0} \cdot \exp(E/E_0)$; $E_0 \geq 1\,\mathrm{PeV}$) gamma-ray blazars to contribute not more than 66.9% to the (high-energy) flux.[31]

Interestingly, Figure 8.3 also allows to extract an interesting subsample of the HBL/IBL class for further studies. While from background expectation only 10 events with a single HBL/IBL counterpart are expected, 20 are found. Therefore, we can identify a subsample of 20 sources with a single, clear counterpart of which around 10 (50%) are expected to be signal. The characterization and identification of common features in these objects are major goals for further work on this topic.

Finally, after submitting the corresponding paper[32] in November 2019, additional neutrino realtime alerts have been sent out by IceCube. For one of these events — IceCube-200107A — a plausible counterpart in the form of a rare HBL blazar has been found as discussed in section 8.5. This observation further strengthens the case for IBL/HBL blazars as emitters of highly-energetic neutrinos.

8.4 An Automatic Multi-Wavelength Follow-Up Pipeline for IceCube Realtime Alerts

The previous section has shown the potential of large-scale multi-frequency dissections to identify plausible blazar counterparts to IceCube neutrinos. As blazars can be highly-variable it is, however, even more favorable to run the entire dissection in near-realtime to request follow-up observations by other observatories. This is particularly important as only a small minority of telescopes have full sky coverage. Other telescopes, such as the X-ray satellite Swift (see section 2.6.3), have a small field of view and therefore need a trigger to observe sources of interest.

The automatic multi-frequency follow-up pipeline presented in this section [33] combines techniques for multi-frequency blazar detection with a fast and reliable, time-dependent analysis of *Fermi*-LAT gamma-ray data. The goal is to collect all the relevant data about a region of interest (ROI) and the blazar candidates on a time-scale similar to the one required for sophisticated IceCube reconstructions ($\sim$ 1 hour). Once triggered by a neutrino alert, the pipeline goes through the following analysis steps:

1. Multi-frequency search for blazar counterparts using the VOU Blazar tool[34] and download of additional optical data from the OVRO database[35] if available.
2. Calculation of the blazars ν_{peak}^{S} using a machine learning estimator, see section 8.4.1.
3. Automatic download of all *Fermi*-LAT data of the ROI from the NASA webpage [36]

---

[31] Huber, "Multi-Messenger correlation study of Fermi-LAT blazars and high-energy neutrinos observed in IceCube".

[32] Giommi, Glauch, et al., "Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection".

[33] https://github.com/tglauch/Realtime_Dissection.

[34] Y.-L. Chang, Brandt, and Giommi, "The Open Universe VOU-Blazars tool".

[35] https://www.astro.caltech.edu/ovroblazars/.

[36] https://fermi.gsfc.nasa.gov/cgi-bin/ssc/LAT/LATDataQuery.cgi

4. A time-dependent gamma-ray analysis of the *Fermi*-LAT data, including light curves of all (potential) gamma-ray blazars, test-static maps and association probabilities for the highest-energy photons (see section 8.4.2). To get quick results, computations are optimized for CPU parallelization and high-performance clusters.

5. Generation of a summary website and PDF document with gamma-ray test statistic maps of the region, hybrid-SEDs, and multi-frequency light curves.

Since some of the computations are more time consuming than others, the summary website and the PDF document are continuously updated whenever calculations finish. For a working example of how the pipeline can be used, see the discussion about the detection of 3HSP J095507.9+355101 in chapter 8.5.

The multi-frequency search for blazar counterparts (step 1 in the list above) has already been discussed in section 4.3. Hence the following sections focus on the discussion of the ν^S_{peak}classifier and the gamma-ray pipeline.

8.4.1 Deep Learning-Based Evaluation of Blazar SEDs

The calculation of the synchrotron peak $\nu^S_{\text{peak}}$is of major importance for the classification of a blazar. As discussed in section 2.3.3 the first bump of the typical blazar SED is a superposition of emission features from the galaxy and the accretion around the black hole with the synchrotron emission. It is therefore not (always) straightforward to fit a polynomial model to determine the ν^S_{peak}. Alternatively, one can use a deep neural network (DNN) that predicts the ν^S_{peak}based on a binned version of the SED [37]. In contrast to the polynomial fit, the DNN provides guaranteed convergence while considering all the features in the spectrum. On the other hand, however, it needs to be trained on a sufficiently large set of SEDs with well-determined ν^S_{peak}. This makes the case specifically complicated for blazars, as there are only a few thousand well-identified and classified blazars in the sky. In order to get a reasonably sized training dataset, we combine the 3HSP catalog of high ν^S_{peak}blazars[38] with the 4LAC catalog.[39] While the ν^S_{peak}estimation of the 4LAC catalog occasionally suffers from confusion with the emission of the host galaxy, it provides a good accuracy on average.

The procedure of the data preparation is exemplarily shown in Figure 8.4. First of all, a set of frequency bins need to be chosen. While a finer binning is usually favored, data coverage is sparse across the SED. To avoid many missing data points, the bins are chosen such that there is data available for most SEDs. As seen in Figure 8.4 this results in one bin in radio, microwave, and many bins in optical, X-ray, and gamma-ray, respectively. For each bin, the 68% central flux range is calculated as input, and all values are concatenated to an array of the length $n = 26$ (13 bins with 2 values each).

[37] See section 5.1 for a short review of DNNs

[38] Y.-L. Chang, Arsioli, et al., "The 3HSP catalogue of extreme and high-synchrotron peaked blazars".

[39] Ajello et al., "The Fourth Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope".

As neural networks perform better when their input values are distributed on a scale $\mathcal{O}(1)$ around zero, the flux values are transformed to the input I_n via

$$I_n = \frac{\log_{10}(\phi_n)}{< \log_{10}(\phi) >}, \tag{8.3}$$

where $< \phi >$ is the median flux value of the entire SED. Finally, bins without data are padded with -5 and 5, representing maximal uncertainty.

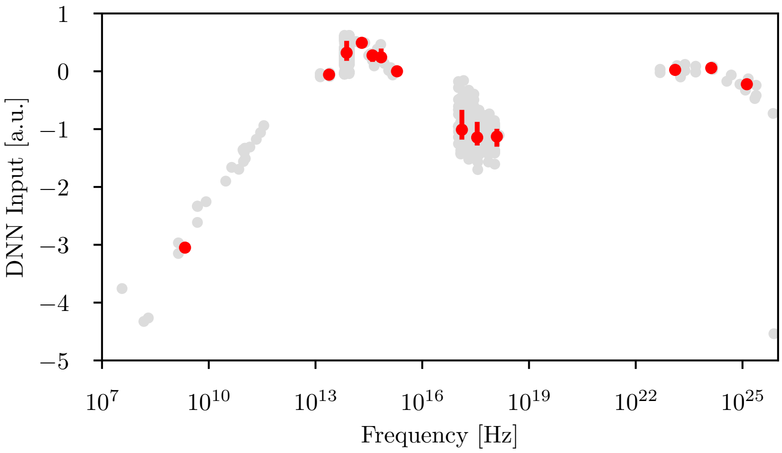

Figure 8.4: Input scheme of the DNN ν^S_{peak}estimator. The grey points exemplarily show the SED data for the blazar TXS 0506+056 based on all available multi-frequency data. The upper and lower points of the red bars represent the central 68% quantiles used as input.

In the end, we also want to re-apply the trained network to the input SED data to refine and consistently determine their ν^S_{peak}. Hence, a special cross-validation strategy is adopted for the training. In the first step, the dataset is divided into 10 parts. Subsequently, for each 9/10th of the dataset, a separated network with the same hyperparameters is trained and applied to the remaining dataset (test dataset). This procedure ensures that the prediction on the test dataset is not affected by overfitting. For the training, we choose a split of training to validation dataset of 80% to 20 % and subsequently minimize the loss function using the Adam optimizer.[40] Early stopping is used after 100 epochs of no improvement in the validation loss. To increase the prediction's stability for yet unclassified objects — and thereby potentially sparse SEDs — up to 4 SED points are randomly dropped for each SED during training. For the loss function, a standard Gaussian likelihood is used. Hence, two parameters are predicted for each SED, a mean value $\hat{y}$ and an uncertainty σ. The corresponding loss function is given by

$$\log \mathcal{L} = -0.5 \sum \left(\frac{y_{0,i} - \hat{y}_i}{\sigma_i} \right)^2 - \sum \log(\sigma_i) - 0.5 \cdot \log(2\pi) \tag{8.4}$$

with target value y_0. The architecture of the DNN is based on a simple fully connected network, followed by some residual layers, leading to a total of around 1000 free parameters.

[40] Kingma and Ba, "Adam: A Method for Stochastic Optimization".

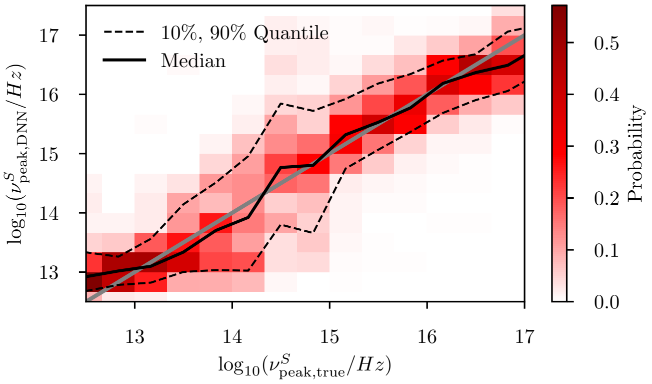

Figure 8.5: Performance of the DNN ν^S_{peak}estimator. The predicted $\nu^S_{\text{peak}}$is shown against the true value from the training catalogs. The probability is normalized along each column of true ν^S_{peak}. The expectation for an unbiased estimator is shown as a solid grey line. The prediction median, as well as the central 90% quantile are indicated as black solid and dashed lines.

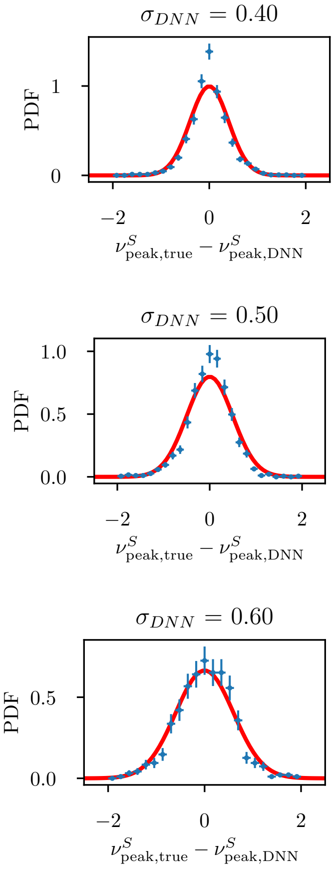

Figure 8.6: Distribution of the ν^S_{peak}prediction errors for different estimated uncertainties σ. The distributions are shown for $\sigma \in 0.4, 0.5, 0.6$ and are reasonably consistent with the Gaussian likelihood.

After the training, the set of DNNs is re-applied to the dataset with the result shown in Figure 8.5. Evidently, the prediction is largely unbiased with respect to the ground truth. Deviations mainly appear at very low or very high ν^S_{peak}, where training statistic is limited, and the uncertainty on the ground truth is high. While this can be improved with better training data in the future, it does not significantly impact the simple LBL/IBL/HBL scheme. Furthermore, Figure 8.6 shows that the predictions are consistently scattered around the truth with their respective uncertainty σ. The uncertainty σ is thereby the superposition of three effects: the uncertainty on the ν^S_{peak}in the training data, the prediction error from the DNN itself, and the time variability of the source emission. While the first two can be reduced by improving the methodology, the third one is unavoidable when calculating an average ν^S_{peak}. In practice, the DNN ν^S_{peak}estimation has been used for a survey of blazars in Swift-XRT fields (see Giommi, Y. L. Chang, et al., "The Open Universe survey of Swift-XRT GRB fields: a complete sample of HBL blazars").

8.4.2 *Gamma-Ray Analysis*

The pipeline's gamma-ray analysis uses all *Fermi*-LAT Pass 8 data [41] acquired between August 4, 2008 (the beginning of the *Fermi*-LAT mission) and the neutrino trigger time. Data preparation follows the standard procedures suggested by the *Fermi*-LAT team. Specifically, only the events with a high probability of being photons (evclass = 128, evtype = 3 [FRONT+BACK]) in the energy range of 100 MeV - 300 GeV from a region of interest (ROI) defined as a circle of radius 12 degrees centered at the multi-wavelength position of the respective source are analyzed. Possible contamination from the Earth limb is removed by cutting out all the events with zenith angle > 90 degrees and only including time intervals with stable data acquisition (DATA QUAL>0 && LAT CONFIG==1). Consistently with the event selection, standard Galactic (gll_iem_v06_V11) and isotropic

[41] with the P8R2 SOURCE V6 instrument response functions

(iso_P8R2_SOURCE_V6_v06) models are used to describe the diffuse background emissions.

Residual test-statistic maps

In general, test statistic maps can identify emission on a coarse grid in a region of interest. For each point on the grid, the test statistic is defined as

$$\mathcal{TS} = 2 \times [\log \mathcal{L}(H_{source}) - \log \mathcal{L}(H_{background})] \tag{8.5}$$

for a given source and background hypothesis. In the case of the residual test statistic maps, the signal hypothesis is given by a point source with spectrum $\phi(E) = \phi_0 \times (E/E_0)^{-\gamma}$ at a certain position in the sky. The background is the superposition of the diffuse gamma-ray background (isotropic and Galactic) and all the known sources in the region. For the latter, the respective flux normalizations are added to the background model if the source is located inside a circle with the radius of the 95% quantile of the *Fermi*-LAT point spread function. The spectral indices for all background sources are fixed to the values in the 4FGL catalog. The rest of the calculations, including exposure maps and data preparation, follows the description in section 8.4.2 and the Fermi Cicerone [42].

Light curves

There are two ways to analyze the time behavior of a gamma-ray source: fixed and adaptive binning light curves.[43] While the first tends to dilute important emission features, the latter is computationally slow and only applicable for a limited number of objects. As the multi-wavelength pipeline is supposed to be fast and applied to many sources simultaneously, a hybrid approach between fixed and adaptive binning is used. Specifically, the length of the time windows is fixed for a single source but varies *between* different sources depending on the time-integrated gamma-ray emission.

Conveniently, the a measurement of the time-integrated emission is already provided by the 4FGL catalog.[44] To deduce a reasonable binning for the light curve, we use the asymptotic behavior of counting experiments. Specifically, for a counting experiments with χ^2_1 background test statistic distribution in the asymptotic limit, the median test statistic quantile of a signal behaves as[45]

$$\mathcal{TS} = -2 \times \left[(s+b)\log\left(1+\frac{s}{b}\right) - s\right] \tag{8.6}$$

where s and b are the number of signal and background photons, respectively. For all the sources considered in this analysis we can safely assume the signal dominated case with $s >> b >> 1$ and $s + b \to \infty$. In this limit, the previous equation simplifies to

$$\mathcal{TS} \to 2 \times s\left[\log\left(\frac{s}{b}\right)\right] \tag{8.7}$$

Assuming further a somewhat steady emission, the signal and background can be approximated to $s = s_0 \cdot t$ and $b = b_0 \cdot t$, with emission rate s_0 and background rate b_0, respectively. Thereby equation (8.7)

---

[42] https://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation/Cicerone/

[43] Lott et al., "An adaptive-binning method for generating constant-uncertainty/constant-significance light curves with Fermi-LAT data".

[44] Abdollahi et al., "Fermi Large Area Telescope Fourth Source Catalog".

[45] Cowan et al., "Asymptotic formulae for likelihood-based tests of new physics".

becomes

$$\mathcal{TS} \to 2 \times s_0 t \left[\log\left(\frac{s_0}{b_0}\right)\right] \propto t \tag{8.8}$$

which scales linearly in time. Using eq. (8.8) we can therefore calculate the integration time t needed to reach a certain significance Σ via

$$t(\Sigma) = \frac{TS(\Sigma)}{\mathcal{TS}_{4FGL}} \cdot 2920\,[\mathrm{days}]. \tag{8.9}$$

In this equation, the 2920 days correspond to the integration time of the 4FGL catalog and $TS(\Sigma)$ to the test-statistic value needed for a certain significance Σ under the assumption of a χ^2_1 distribution, i.e., $TS(\Sigma) = \Sigma^2$. The test statistic information in the 4FGL catalog is given in 7 energy bands. In order to calculate the $\mathcal{TS}_{4FGL}$ for a given energy threshold E_{min} one can sum up the independent energy bands with the lower energy bound being equal or higher energy than E_{min} of the analysis, i.e.,

$$\mathcal{TS}_{4FGL} = \sum_{E_i > E_{min}} \mathcal{TS}^{E_i}_{4FGL} \leq \mathcal{TS}^{true}_{4FGL}. \tag{8.10}$$

Note that this approximation is conservative, as the test statistic value in each energy band ,$\mathcal{TS}^{E_i}_{4FGL}$, is given for a standard power law assumption $\phi = \phi_0 \times (E/E_0)^{-2}$. Hence the test statistic value calculated from equation (8.10) is optimal for a *true* source spectrum with $\gamma = 2$, but underestimated otherwise. Optimally one would aim for a 5σ detection in every time window, in order to characterize the spectral variations of the source. The required integration times is given by

$$t(5\sigma) = \frac{25}{\mathcal{TS}_{4FGL}} \cdot 2920\,[\mathrm{days}]. \tag{8.11}$$

As stated before, this is a conservative approximation assuming a steady-source.

In reality, sources can be strongly variable, in which case the key assumption of the calculation does not hold anymore. There are, however, two reasons why the procedure still gives good results in a large majority of cases: a) a jump from 5σ to 3σ (which is around a factor 3 smaller in required integration time) still gives meaningful spectral fits, i.e., there is certain stability against lower-than-average flux windows; b) the methods ensure that especially time windows with particular large photon flux — which are particularly interesting — are well described. In summary, the method provides a quick, efficient, and reliable way to pick an ad-hoc binning for light curves that preservers most of the important information while being computationally inexpensive.

8.5 The Case of 3HSP J095507.9+355101 - a Flaring Blazar in the Error Region of IceCube-200107A

The detection of an X-ray flaring 3HSP source as a probable counterpart to the high-energy neutrino IceCube-200107A is a prime example of the interplay between the new DNN-based HESE v3 selection (chapter 6) and the automatic follow-up pipeline (section 8.4).

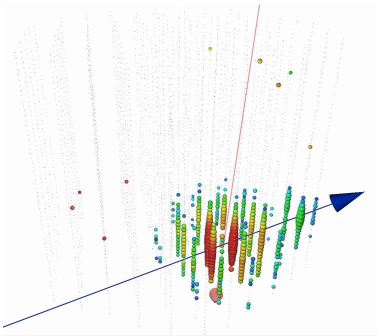

Figure 8.7: Event view of IceCube-200107A. The color indicates the time, from early (red) to late (blue). The size of the bubbles represents the total collected charge. In addition, two reconstruction are shown: Spline MPE in red, and the RNN reco in blue.

On the 7^{th} of January 2020, a highly-energetic starting event passed the IceCube HESE filter (see Figure 8.7 for an event view). While the *SplineMPE* reconstruction falsely reconstructs a vertically downgoing direction and therefore rejects the event, it was clearly identified as a starting track by the DNN classifier. Consequently, a GCN was sent out[46] to alert the astronomical community for follow-up observations. Based on the new HESE v3 realtime selection, the event's signalness, i.e., the probability of being of astrophysical origin, is around 52 %. With the limited public information in the GCN it is, however, difficult to make precise statements om the event's energy. Using the $\nu_\mu + \bar{\nu}_\mu$ public HESE effective area as a good approximation for the HESE v3 stream (Figure 6.9), we can nevertheless approximate the expected neutrino energy given the reconstructed direction and an assumed astrophysical neutrino spectrum. The calculation requires the normalized cumulative density function of neutrino energies E_ν which is given by

$$\mathcal{C}(E_\nu) = \frac{\int\limits_{-\infty}^{E_\nu} A_{eff}(E) \times E^{-\gamma}\,\mathrm{d}E}{\int\limits_{-\infty}^{\infty} A_{eff}(E) \times E^{-\gamma}\,\mathrm{d}E}. \tag{8.12}$$

The resulting neutrino energy distributions for the case of IceCube-200107A are shown for three different spectral indices in Figure 8.8. Assuming an E^{-2} ($E^{-1}/E^{-2.7}$) spectrum we can read off energy ranges between 65 TeV - 2.5 PeV (87 TeV - 7.2 PeV / 39 TeV - 981 TeV), respectively.

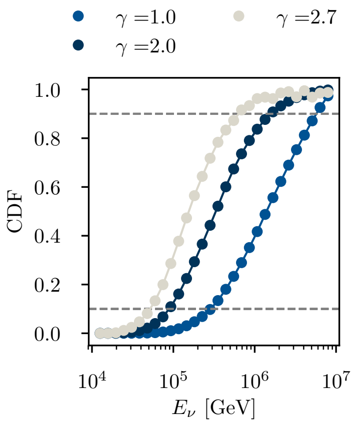

Figure 8.8: Estimation of the neutrino energy assuming different flux models. The cumulative density function (CDF) is calculated according to eq. (8.12). The 10% and 90% quantile are given by the crossing points of the CDFs with the grey dashed lines.

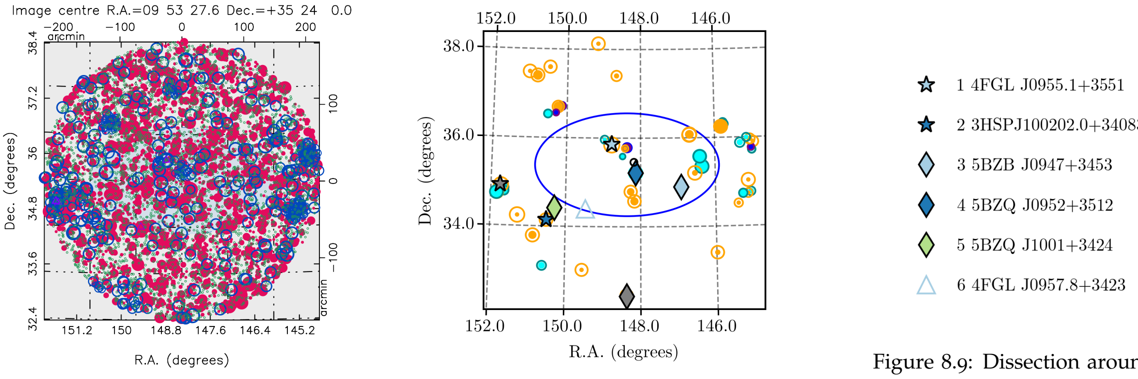

Figure 8.9: Dissection around IceCube-200107A using the VOU Blazar tool. The left plot shows the radio (red) and X-ray emission (blue). After cross-matching the emissions a few blazar-like sources remain (right). Dark blue circles represent LBL, cyan symbols IBL, and orange symbols HBL candidates, respectively. Known blazars are marked by diamonds or stars if they are part of the BZCAT or 2WHSP catalog.

Based on the GCN, the automatic realtime pipeline (section 8.4) was triggered for the region around IceCube-200107A. The VOU-Blazars output with all radio and X-ray sources, as well as the blazar

---

[46] https://gcn.gsfc.nasa.gov/gcn3/26655.gcn3

candidates, is shown in Figure 8.9. It can be seen that inside the 90% confidence region many cataloged and non-cataloged blazars candidates are identified.

According to usual photo-meson scenarios, neutrino production in blazars is assumed to happen together with the production of gamma rays. Hence the residual gamma-ray maps in Figure 8.10 are used to identify plausible candidates. The map is residual in a sense that known gamma-ray sources, as *4FGL J0955.1+3551*, are subtracted to identify yet undiscovered weak emitters. For neither the short term TS-map or the long term one, any significant excess is found. Hence the only gamma-ray detected candidate in the region remains *4FGL J0955.1+3551*.

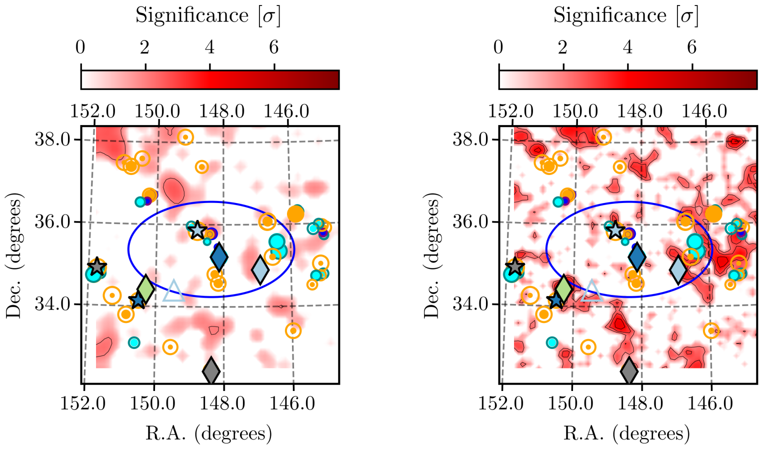

Figure 8.10: *Fermi*-LAT gamma-ray significance maps in the region around IceCube-200107A after subtracting known 4FGL sources. The respective integration times are 250 days before the alert (left), and over the full *Fermi*-LAT mission (right). The multi-frequency blazar candidates identified in Figure 8.9 are shown with their respective symbols.

The full, time-dependent multi-frequency SED of the source is presented in Figure 8.11. In the optical range between 1.3 eV - 3.1 eV the accretion of the galaxy is clearly visible, while the non-thermal synchroton peak goes up to X-ray energies. In fact, the source is included in the 3HSP catalog (*3HSP J095507.9+355101*)[47] with a ν^S_{peak} larger than 5×10^{17} Hz. Overall, the 3HSP catalogue includes only 80 *Fermi*-LAT detected sources with a higher ν^S_{peak}. The corresponding source density of those objects is therefore about one every 425 square degrees. Consequently, the probability that an extreme source like that is included in the 7.3 square degrees error region of IceCube-200107A is around 1.7%. Overall, the automatic follow-up pipeline has thereby laid out a detailed scientific justification to ask for further follow-up observations of the one primary good candidate in the error contour, 3HSP J095507.9+355101. A target of opportunity (ToO) request for the X-ray observation of the source was therefore submitted to Swift and granted shortly after [48]. The results from these optical and X-ray observations are then fed back into the pipeline and combined to the SED in Figure 8.11. Compared to the archival data (grey) the source appears in a high X-ray and slightly increased gamma-ray state. Quantitatively, a dedicated analysis of the Swift data show that the 2 – 10 keV X-ray flux is around

[47] Y.-L. Chang, Arsioli, et al., "The 3HSP catalogue of extreme and high-synchrotron peaked blazars".

[48] https://www.swift.psu.edu/toop/summary.php

$5 \times 10^{-12}\,\mathrm{erg\,cm^{-2}\,s^{-1}}$, which is a factor of 2.5 larger than the average value observed during a previous observation in 2012/2013. Fitting a log-parabola model we find a spectral index of $\gamma = 1.8 \pm 0.06$ without curvature and therefore conclude that the synchrotron peak ν^S_{peak} must be beyond 2×10^{18} Hz. The gamma-ray spectral indices in Figure 8.11 are $\gamma = 1.73 \pm 0.31$ for the 250 days around the neutrino arrival time and $\gamma = 1.88 \pm 0.15$ for the full mission. Together with the redshift estimation of 0.557[49] the gamma-ray luminosity between 100 MeV and 100 GeV is then given by $3.96 \times 10^{45}\,\mathrm{erg\,s^{-1}}$ and $1.47 \times 10^{45}\,\mathrm{erg\,s^{-1}}$, respectively. In a similar way the redshift can also be used to calculate the neutrino luminosity required to expect on average one IceCube neutrino on a given timescale using the same effective area as in eq. (8.12) and assuming an E^{-2} flux between 65 TeV and 2.5 PeV. The results are $5 \times 10^{47}\,\mathrm{erg\,s^{-1}}$ and $3 \times 10^{46}\,\mathrm{erg\,s^{-1}}$, respectively.[50]

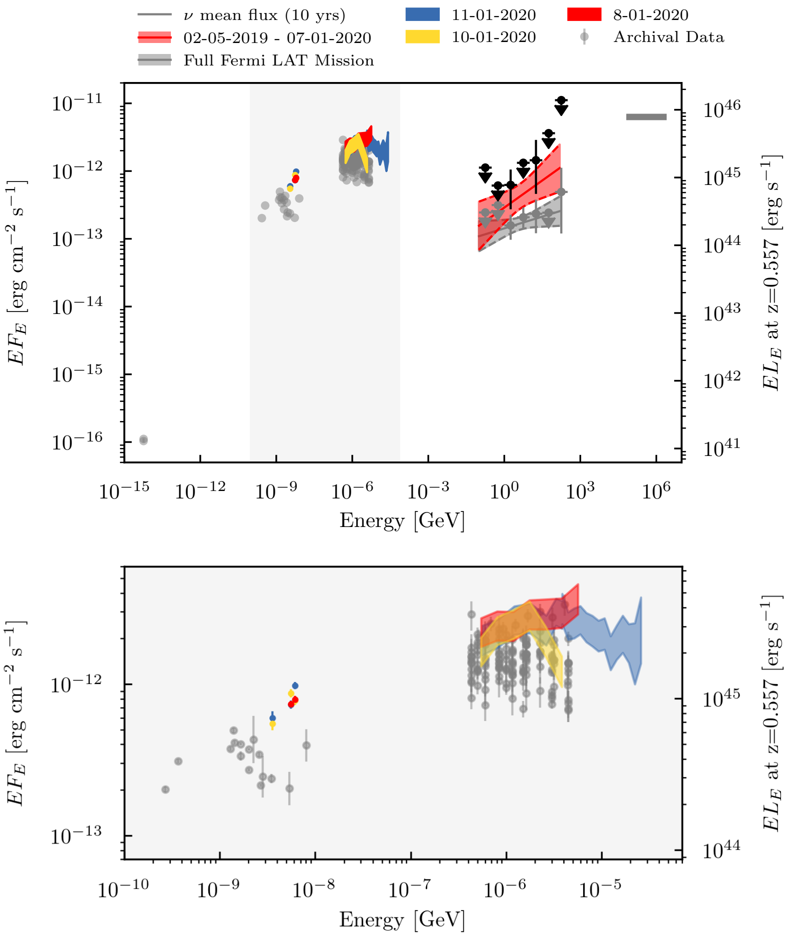

Figure 8.11: The hybrid SED of 3HSP J095507.9+355101. The gamma-ray emission is shown as grey and red bow ties for the full *Fermi*-LAT mission and the 250 days before the detection of IceCube-200107A. The neutrino flux upper limit, calculated from the observation of a single event over the entire *Fermi*-LAT emission period, is shown as grey solid line. In all other wavelengths grey points indicate archival data. Data from the three Swift follow-up observations are shown as colored points and bands. The bottom panel is a zoom in view of the optical and X-ray emission.

8.5.1 Theoretical Considerations

Neutrino production in Blazars is usually expected to happen through photo-hadronic interaction of accelerated protons with ambient photons. While the detailed modelling of the sources is challenging and

---

[49] Paiano, Falomo, Padovani, et al., "The redshift and the host galaxy of the neutrino candidate 4FGL J0955.1+3551 (3HSP J095507.9+355101)".

[50] Giommi, Padovani, Oikonomou, et al., "3HSP J095507.9+355101: a flaring extreme blazar coincident in space and time with IceCube-200107A".

depends on many parameters,[51] some generic arguments can be laid out.[52] First of all we note that 3HSP J095507.9+355101 is likely a blazar of type BL Lac. This is mainly based on two arguments a.) the non detection of strong emission lines in the optical spectrum[53] (see also Figure 2.14) and b.) the Eddington ratio of $L/L_{edd} \lesssim 0.02$, which is at the boarder between BL Lacs and FSRQs. The calculation of the luminosity is based on the estimated black hole mass $M_{BH} \sim 3 \times 10^8 M_\odot$ (typical for a BL Lac),[54] with resulting Eddington luminosity $L_{edd} \sim 4 \times 10^{46}\,\mathrm{erg\,s^{-1}}$ and the upper limit on the oxygen OII and OIII emission lines of $L_{\rm OII}$ and $L_{\rm OIII} < 2 \times 10^{40}\,\mathrm{erg\,s^{-1}}$. This upper limit can be converted into a limit on the bolometric luminosity $L_{bol} \lesssim 8 \times 10^{44}\,\mathrm{erg\,s^{-1}}$,[55] giving $L/L_{edd} \lesssim 0.02$.

In the following, the energetics of the gamma-ray and neutrino production are discussed assuming the source being a BL Lac, implying that the target photons are produced comoving in the jet. Following section 2.2.1 and specifically eq. (2.29) a fraction of 3/8th of the proton energy lost in photo-hadronic interactions goes to the neutrinos and hence neutrino and proton luminosity are related over

$$\epsilon_\nu L_\nu = \frac{3}{8} f_{p\gamma} \epsilon_p L_p \tag{8.13}$$

with respective optical depth $f_{p\gamma}$ and neutrino energy $\epsilon_\nu = 0.05\epsilon_p$. For the case of the target photons being produced comoving with the jet, the observed neutrino energy can be related to the *observed* energy of the target photons ϵ_t via

$$\epsilon_{\nu,\rm obs} \;\approx\; 7.5\ \mathrm{PeV} \left(\frac{2\ \mathrm{keV}}{\epsilon_{t,\rm obs}}\right) \left(\frac{\Gamma}{20}\right)^2 \frac{1}{(1+z)^2} \tag{8.14}$$

The remaining 5/8 of the energy goes into the production of gamma rays and the subsequent production of electron positron pairs. The resulting synchrotron cascade is shifting the gamma ray photons to lower energies. The luminosity of the neutrino and the gamma rays can hence be related over[56]

$$\epsilon_\nu L_{\epsilon_\nu} \approx \frac{6(1+Y_{\rm IC})}{5} \epsilon_\gamma L_{\epsilon_\gamma}|_{\epsilon^{p\pi}_{\rm syn}} \approx 8 \times 10^{44}\ \mathrm{erg\ s^{-1}} \left(\frac{\epsilon_\gamma L_{\epsilon_\gamma}|_{\epsilon^{p\pi}_{\rm syn}}}{7 \times 10^{44}}\right) \tag{8.15}$$

with the compton dominance $Y_{\rm IC} \in \mathcal{O}(1)$ and the typical energy of the observed synchrotron photons

$$\epsilon^{p\pi}_{\rm syn,obs} \approx 39.4\ \mathrm{GeV} \left(\frac{B}{0.3\ \mathrm{G}}\right) \left(\frac{\epsilon_{\nu,\rm obs}}{7.5\ \mathrm{PeV}}\right)^2 \left(\frac{20}{\Gamma}\right) \left(\frac{1}{1+z}\right) \tag{8.16}$$

From Figure 8.8 we see that for an assumed spectral index $\gamma = 1$ the most likely neutrino energy is around 1 PeV. This is consistent with eq. (8.14) assuming that the target photons are provided by the synchroton peak at $\nu_t \approx 10^{18}$ Hz ($\epsilon_t \approx 2\,$keV) and that the jet has a typical bulk Lorentz factor of $\Gamma \approx 20$.[57] Using these values in eq. (8.16) for a typical magnetic field strength of $B = 20\,$G the synchroton gamma-ray flux appears at $\epsilon^{p\pi}_{\rm syn,obs} \approx 30\,$GeV consistent with an

[51] Petropoulou, Oikonomou, et al., "Comprehensive Multimessenger Modeling of the Extreme Blazar 3HSP J095507.9+355101 and Predictions for IceCube".

[52] Giommi, Padovani, Oikonomou, et al., "3HSP J095507.9+355101: a flaring extreme blazar coincident in space and time with IceCube-200107A".

[53] Paiano, Falomo, Padovani, et al., "The redshift and the host galaxy of the neutrino candidate 4FGL J0955.1+3551 (3HSP J095507.9+355101)".

[54] Ibid.

[55] Punsly and S. Zhang, "Calibrating emission lines as quasar bolometers".

[56] Murase, Oikonomou, and Petropoulou, "Blazar Flares as an Origin of High-Energy Cosmic Neutrinos?"

[57] Saikia, Körding, and Falcke, "Lorentz factor distribution of blazars from the optical Fundamental plane of black hole activity".

observations in the *Fermi*-LAT band. Based on the gamma-ray luminosities in section 8.5 we then estimate the gamma-ray luminosity per decade via

$$\epsilon_\gamma L_\gamma \sim \frac{L_\gamma}{\log(100\,\mathrm{GeV}/100\,\mathrm{MeV})} \tag{8.17}$$

assuming a spectral index of $\gamma = 2$. As a result $\epsilon_\gamma L_\gamma$ is $5.73 \times 10^{44}\,\mathrm{erg\,s^{-1}}$ and $2.13 \times 10^{44}\,\mathrm{erg\,s^{-1}}$ for the 250 day period and the full mission, respectively. Plugging this into eq. (8.15), the expected neutrino luminosities $\epsilon_\nu L_{\epsilon_\nu}$ are hence on the order of $10^{44}\,\mathrm{erg\,s^{-1}}$, which is around two orders of magnitude smaller than the flux required to expect the observation of one neutrino. While this is puzzling, it can be understood as the result of an underlying population of sources[58] or the presence of multiple emission zones. Note finally, that as for the case of TXS 0506+056, 3HSP J095507.9+355101 is a clear outlier from the blazar sequence (Figure 8.12).

In summary, 3HSP J095507.9+355101 is a plausible counterpart to IceCube-200107A in the sense that it provides sufficient target photon material to produce neutrinos and gamma-rays in the expected energy range. The discrepancy between the neutrino and gamma-ray luminosities can be understood through the limited knowledge of the underlying source population producing astrophysical neutrinos.

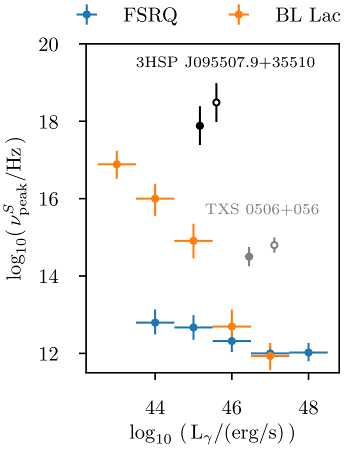

Figure 8.12: 3HSP J095507.9+355101 in the blazar sequence. Similar to TXS 0506+056, which is very much off the blazar sequence for FSRQs, 3HSP J095507.9+355101 is inconsistent with the blazar sequence for BL Lacs. Blazar sequence data taken from Ghisellini et al., "The Fermi blazar sequence".

[58] Strotjohann, Kowalski, and Franckowiak, "Eddington bias for cosmic neutrino sources".

Table 8.3: The sample of 70 IceCube tracks up to November 2019 off the Galactic plane. Columns 1 and 2 give the standard IceCube name and previous namings and the other columns give the MJD, right ascension and declination with 90% error (if available) and the Galactic latitude. Wherever no 90% error is given, a fixed *median* angular resolution of 1.3 degrees is used. The events are sorted by right ascension. Whenever a new data release has provided an updated reconstruction we add a † to the event name. Adapted from Giommi, Glauch, et al., "Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection".

IceCube Name	Other IceCube Name	MJD	RA J2000.0 (deg)	Dec J2000.0 (deg)	Galactic Latitude (deg)
IceCube-160331A	DIF35	57478.60	$15.60^{+0.45}_{-0.58}$	$15.60^{+0.53}_{-0.60}$	-47.19
IceCube-090813A	DIF1	55056.70	$29.51^{+0.40}_{-0.38}$	$1.23^{+0.18}_{-0.22}$	-57.42
IceCube-131014A	DIF23	56579.91	$32.94^{+0.63}_{-0.62}$	$10.20^{+0.34}_{-0.49}$	-47.90
IceCube-111216A	DIF16	55911.28	$36.65^{+1.85}_{-1.71}$	$19.10^{+2.21}_{-2.21}$	-38.34
IceCube-161103A	AHES4	57695.38	$40.83^{+1.10}_{-0.70}$	$12.56^{+1.10}_{-0.65}$	-41.92
IceCube-161210A	EHE3	57732.84	$46.58^{+1.10}_{-1.00}$	$14.98^{+0.45}_{-0.40}$	-36.67
IceCube-150831A		57265.22	$54.85^{+0.94}_{-0.98}$	$33.96^{+1.07}_{-1.19}$	-17.09
IceCube-141109A	HES61†	56970.21	$55.63^{+0.79}_{-1.53}$	$-16.50^{+0.81}_{-0.68}$	-49.11
IceCube-190504A		58607.77	$65.79^{+1.23}_{-1.23}$	$-37.44^{+1.23}_{-1.23}$	-44.68
IceCube-120922A		56192.55	$70.75^{+1.56}_{-1.63}$	$19.79^{+1.37}_{-0.68}$	-16.90
IceCube-151114A	DIF34	57340.90	$76.30^{+0.75}_{-0.74}$	$12.60^{+0.61}_{-0.58}$	-16.79
IceCube-170922A	EHE5	58018.87	$77.43^{+0.95}_{-0.65}$	$5.72^{+0.50}_{-0.30}$	-19.56
IceCube-150428A	HES71†	57140.47	$80.77^{+1.12}_{-1.23}$	$-20.75^{+0.45}_{-0.83}$	-28.33
IceCube-101028A	DIF9†	55497.30	$88.68^{+0.54}_{-0.55}$	$0.46^{+0.33}_{-0.27}$	-12.38
IceCube-170321A	EHE4	57833.31	$98.30^{+1.20}_{-1.20}$	$-15.02^{+1.20}_{-1.20}$	-10.75
IceCube-140721A	HES58	56859.76	102.10	-32.40	-14.73
IceCube-140611A	DIF27†	56819.20	$110.30^{+0.66}_{-0.45}$	$11.57^{+0.14}_{-0.24}$	11.79
IceCube-190503A		58606.72	$120.28^{+0.57}_{-0.77}$	$6.35^{+0.76}_{-0.70}$	18.37
IceCube-160806A	EHE2	57606.51	$122.81^{+0.50}_{-0.50}$	$-0.81^{+0.50}_{-0.50}$	17.29
IceCube-130907A		56542.79	$129.81^{+0.48}_{-0.28}$	$-10.36^{+0.36}_{-0.31}$	18.35
IceCube-150904A	DIF32	57269.80	$134.00^{+0.39}_{-0.58}$	$28.00^{+0.47}_{-0.47}$	38.35
IceCube-100623A	DIF4	55370.74	$141.25^{+0.46}_{-0.45}$	$47.80^{+0.56}_{-0.48}$	45.16
IceCube-180908A		58369.83	$144.58^{+1.55}_{-1.45}$	$-2.13^{+0.9}_{-1.2}$	35.09
IceCube-141209A	HES63†	57000.14	$160.05^{+0.84}_{-1.04}$	$6.57^{+0.64}_{-0.56}$	52.68
IceCube-171015A		58041.07	$162.86^{+2.60}_{-1.70}$	$-15.44^{+1.60}_{-2.00}$	38.43
IceCube-130408A	HES37	56390.19	$167.17^{+2.87}_{-1.90}$	$20.67^{+1.15}_{-0.89}$	65.69
IceCube-121026A	DIF20	56226.60	$169.61^{+1.16}_{-1.11}$	$28.04^{+0.67}_{-0.66}$	69.40
IceCube-140923A		56923.72	$169.72^{+0.91}_{-0.86}$	$-1.34^{+0.73}_{-0.66}$	53.85
IceCube-120523A		56070.57	$171.03^{+0.81}_{-0.90}$	$26.36^{+0.49}_{-0.30}$	70.51
IceCube-190819A		56070.57	$148.80^{+2.07}_{-3.24}$	$1.38^{+1.00}_{-0.75}$	70.51
IceCube-141126A	HES62	56987.77	187.90	13.30	75.41
IceCube-150926A		57291.90	$194.50^{+0.76}_{-1.21}$	$-4.34^{+0.70}_{-0.95}$	58.49
IceCube-151017A	DIF33	57312.70	$197.60^{+2.46}_{-2.09}$	$19.90^{+2.82}_{-2.21}$	81.57

IceCube-120515A	DIF17	56062.96	$198.74^{+1.44}_{-1.09}$	$31.96^{+0.81}_{-0.85}$	82.97
IceCube-160814A	AHES3	57614.91	$200.30^{+2.43}_{-3.03}$	$-32.40^{+1.39}_{-1.21}$	30.05
IceCube-121011A	DIF19†	56211.77	$205.22^{+0.59}_{-0.65}$	$-2.39^{+0.51}_{-0.57}$	58.17
IceCube-131202A	HES43†	56628.57	$206.63^{+2.04}_{-1.56}$	$-22.02^{+1.69}_{-1.04}$	39.08
IceCube-120123A	HES23	55949.57	208.70	-13.20	46.84
IceCube-140216A	HES47	56704.60	209.40	67.40	48.49
IceCube-160731A	EHE1, AHES2	57600.08	$214.50^{+0.75}_{-0.75}$	$-0.33^{+0.75}_{-0.75}$	55.55
IceCube-170506A	AHES6	57879.53	$221.80^{+3.00}_{-3.00}$	$-26.00^{+2.00}_{-2.00}$	30.01
IceCube-130817A	DIF22	56521.83	$224.89^{+0.87}_{-1.19}$	$-4.44^{+1.21}_{-0.94}$	45.80
IceCube-181014A		58405.50	$225.15^{+1.40}_{-2.85}$	$-34.80^{+1.15}_{-1.85}$	20.95
IceCube-190730A		58694.87	$225.79^{+1.28}_{-1.43}$	$10.47^{+1.14}_{-0.89}$	54.83
IceCube-110521A	DIF12	55702.77	$235.13^{+2.70}_{-1.76}$	$20.30^{+1.00}_{-1.43}$	50.88
IceCube-120301A		55987.81	$238.01^{+0.60}_{-0.59}$	$18.60^{+0.46}_{-0.39}$	47.76
IceCube-140420A	HES53†	56767.07	$238.98^{+1.81}_{-1.91}$	$-37.73^{+1.47}_{-1.31}$	12.06
IceCube-150911A	HES76†	57276.57	$240.20^{+1.29}_{-1.38}$	$-0.45^{+1.17}_{-1.23}$	36.83
IceCube-160427A	AHES1, HES82†	57505.25	$240.57^{+0.60}_{-0.60}$	$9.34^{+0.60}_{-0.60}$	41.68
IceCube-151122A		57348.53	$262.18^{+0.90}_{-1.21}$	$-2.38^{+0.73}_{-0.43}$	17.16
IceCube-110930A		55834.45	$266.48^{+2.09}_{-1.55}$	$-4.41^{+0.59}_{-0.86}$	12.43
IceCube-100925A	DIF7	55464.90	$266.29^{+0.58}_{-0.62}$	$13.40^{+0.52}_{-0.45}$	20.64
IceCube-110610A	DIF13	55722.43	$272.22^{+1.23}_{-1.19}$	$35.55^{+0.69}_{-0.69}$	23.50
IceCube-131204A		56630.47	$289.16^{+1.08}_{-0.94}$	$-14.25^{+0.91}_{-0.81}$	-11.94
IceCube-131023A		56588.56	$301.82^{+1.10}_{-0.93}$	$11.49^{+1.19}_{-1.09}$	-11.10
IceCube-170312A	AHES5	57824.58	$305.15^{+0.50}_{-0.50}$	$-26.61^{+0.50}_{-0.50}$	-30.40
IceCube-100710A	DIF5	55387.54	$306.96^{+2.70}_{-2.28}$	$21.00^{+2.25}_{-1.56}$	-10.13
IceCube-190124A		58507.15	$307.40^{+0.80}_{-0.90}$	$-32.18^{+0.70}_{-0.70}$	-33.76
IceCube-110128A	DIF11†	55589.56	$307.53^{+0.82}_{-0.81}$	$1.19^{+0.35}_{-0.32}$	-21.22
IceCube-150714A	DIF30	57217.90	$325.50^{+1.77}_{-1.46}$	$26.10^{+1.68}_{-1.85}$	-19.93
IceCube-150812A	DIF31†	57246.76	$328.19^{+1.01}_{-1.03}$	$6.21^{+0.44}_{-0.49}$	-35.44
IceCube-120807A	DIF18	56146.21	$330.10^{+0.65}_{-0.82}$	$1.57^{+0.46}_{-0.42}$	-39.84
IceCube-101009A	DIF8†	55478.38	$331.09^{+0.56}_{-0.72}$	$11.10^{+0.48}_{-0.58}$	-34.30
IceCube-140114A	HES44	56671.88	336.71	0.04	-45.92
IceCube-190331A		58573.29	$337.68^{+0.23}_{-0.34}$	$-20.70^{+0.30}_{-0.48}$	-57.31
IceCube-171106A		58063.77	$340.00^{+0.7}_{-0.5}$	$7.40^{+0.35}_{-0.25}$	-43.05
IceCube-140108A		56665.31	$344.53^{+0.67}_{-0.48}$	$1.57^{+0.35}_{-0.32}$	-50.41
IceCube-140203A		56691.79	$349.54^{+2.21}_{-1.97}$	$-13.71^{1.23}_{-1.38}$	-64.43
IceCube-160510A		57518.66	$352.34^{+1.63}_{-1.31}$	$2.09^{+0.99}_{-0.85}$	-54.72
IceCube-190104A		58487.36	$357.98^{+2.30}_{-2.10}$	$-26.65^{+2.20}_{-2.50}$	-76.73

Table 8.4: Table of IceCube tracks with blazar counterpart candidates within 1.5 times the 90% error ellipses ($\Omega_{90\times1.5}$). In addition to the source names we also give the 1.4 GHz radio flux, the ν^S_{peak} for the SED classification and the redshift. Bold event names indicate tracks with at least one plausible counterpart, while the letters a, b, c, d indicate whether the source was found in Ω_{90}, $\Omega_{90\times1.1}$, $\Omega_{90\times1.3}$, $\Omega_{90\times1.5}$ respectively. Adapted from Giommi, Glauch, et al., “Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection”

IceCube Name	Source Name	$S_{1.4\mathrm{GHz}}$ [mJy]	ν^S_{peak} [Hz]	Redshift
IceCube-160331A	3HSP J010326.0+152624 [a]	225	15.0	0.25
IceCube-090813A	5BZU J0158+0101 [b]	82	14.1	0.4537
IceCube-131014A				
IceCube-111216A	5BZQ J0225+1846 [a]	461	12.5	2.69
	3HSP J023248.5+201717 [a]	82	18.5	0.14
	VOU J022411+161500 [d]	13	14.5	0.3
IceCube-161103A	VOU J024445+132002 [a]	200	14.5	0.90
	3HSP J023927.2+132738 [d]	20	15.0	0.5
IceCube-161210A				
IceCube-150831A	3HSP J034424.9+343017 [c]	13	15.7	—
	5BZQ J0336+3218 [d]	2677	12.	1.26
IceCube-141109A	3HSP J033913.6-173600 [d]	171	15.6	0.07
IceCube-190504A	5BZB J0428-3756 [a]	753	12.8	1.11
	4LAC J0420.3-3745 [a]	60	<13.5	0.3
IceCube-120922A				
IceCube-151114A	5BZB J0502+1338 [d]	545	13.2	–
IceCube-170922A	5BZB J0509+0541 [a]	536	14.5.	0.34
IceCube-150428A	VOU J052526-201054 [c]	231	14.5	0.12
IceCube-101028A				
IceCube-170321A	3HSP J062753.3-151957 [c]	43	17.3	0.29
IceCube-140721A	3HSP J064933.5-313920 [a]	8	17.0	>0.56
	5BZQ J0648-3044 [c]	898	12.5	1.15
IceCube-140611A				
IceCube-190503A				
IceCube-160806A				
IceCube-130907A				
IceCube-150904A	3HSP J085410.1+275421 [a]	15	16.1	0.49
IceCube-100623A				
IceCube-180908A				
IceCube-141209A	VOU J104031+061721 [a]	35	14.5	–
	5BZB J1043+0653 [b]	8	14.5	0.43
IceCube-171015A	VOU J105603-180929 [d]	12	14.1	–
IceCube-130408A	3HSP J111706.2+201407 [a]	103	16.5	0.14
	5BZQ J1059+2057 [b]	121	13.0	0.39
	3HSP J112405.3+204553 [d]	9	15.3	0.54

	3HSP J112503.6+214300 [d]	8	15.8	0.36
IceCube-121026A				
IceCube-140923A				
IceCube-120523A	5BZQ J1125+2610 [c]	921	12.5	2.34
IceCube-190819A	3HSP J094620.2+010451 [a]	15	$>$ 18	0.58
	3HSP J100326.6+020455 [c]	6	15.8	0.48
	5BZQ J0948+0022 [d]	70	12.8	0.59
IceCube-141126A	M87 [a]	138488	–	0.004
	3HSP J123123.9+142124 [a]	54	16.0	0.26
IceCube-150926A	3HSP J125848.0-044745 [a]	4	17.0	0.59
IceCube-151017A	5BZB J1314+2348 [d]	184	$\geq$14	0.15?
	3HSP J130008.5+175537 [d]	16	14.5	0.55
	5BZQ J1321+2216 [d]	314	12.0	0.943
	3HSP J125821.5+212351 [d]	26	16.7	0.6265
IceCube-120515A	5BZU J1310+3220 [b]	1687	12.5	0.997
	5BZQ J1310+3233 [b]	374	12.0	1.64
	5BZB J1322+3216 [c]	906	14.5	–
IceCube-160814A	5BZQ J1316-3338 [b]	1277	12.5	1.21
IceCube-121011A	5BZQ J1340-0137 [d]	175	13.0	1.62
IceCube-131202A	5BZQ J1342-2051 [a]	399	12.0	1.58
IceCube-120123A	VOU J135921-115043 [d]	48	14.0	0.27
IceCube-140216A	3HSP J140449.6+655431 [c]	15	16.0	0.36
	5BZQ J1344+6606 [d]	639	12.3	1.35
IceCube-160731A				
IceCube-170506A	3HSP J144656.8-265658 [a]	41	17.6	0.32
	VOU J143934-252458 [a]	35	14.0	0.18
	3HSP J143959.4-234140 [c]	101	16.2	0.25
IceCube-130817A				
IceCube-181014A	5BZB J1505-3432 [a]	138	12.5	–
	5BZQ J1457-3539 [a]	675	13.5	1.42
	VOU J150720-370902 [d]	74	14.5	–
IceCube-190730A	5BZQ J1504+1029 [a]	1775	12.8	1.84
IceCube-110521A	3HSP J155424.1+201125 [c]	80	17.3	0.27
	3HSP J153311.2+185429 [d]	23	17.0	0.30
	3HSP J152835.7+200420 [d]	5	16.2	0.52
IceCube-120301A				
IceCube-140420A				
IceCube-150911A	5BZQ J1557-0001 [a]	1107	12.2	1.77
IceCube-160427A				
IceCube-151122A				
IceCube-110930A	5BZQ J1743-0350 [b]	1411	12.5	1.06
IceCube-100925A				
IceCube-110610A	VOU J180812+350104 [a]	94	14.5	0.4
	3HSP J180849.7+352042 [a]	31	15.	0.142
IceCube-131204A	VOU J191651-151902 [b]	166	13.0	–
IceCube Name	Source Name	$S_{1.4\mathrm{GHz}}$ [mJy]	ν^{S}_{peak} [Hz]	Redshift

IceCube-131023A				
IceCube-170312A				
IceCube-100710A	3HSP J203031.7+223439 [a]	5	16.2	–
	3HSP J203057.1+193612 [a]	57	15.8	0.27
IceCube-190124A				
IceCube-110128A				
IceCube-150714A	3HSP J213314.3+252859 [c]	40	15.2	0.29
	VOU J213253+261144 [d]	211	12.0	0.8
IceCube-150812A				
IceCube-120807A				
IceCube-101009A				
IceCube-140114A	5BZB J2227+0037 [a]	102	14.5	–
	5BZQ J2226+0052 [a]	615	12.5	2.26
	3HSP J222329.5+010226 [b]	7	15.5	0.51
IceCube-190331A				
IceCube-171106A				
IceCube-140108A				
IceCube-140203A				
IceCube-160510A	VOU J232625+011147 [c]	204	14.0	0.53
IceCube-190104A	IC 5362 [a]	90	14.5	0.03
	VOU J235815-285341 [b]	169	14.0	–
	3HSP J235034.3-300604 [d]	39	15.7	0.23

9 Conclusion & Outlook

The quest for the origin of high energy cosmic radiation is both a scientific and a technological challenge. While astronomy and astrophysics provide the models and search strategies to achieve this goal, data science and statistics are mandatory to identify, with confidence, a signal in the large amount of data provided by modern telescopes such as the IceCube Neutrino Observatory. This has been the focus of this thesis's research. Consequently, we have tackled the quest from an interdisciplinary standpoint: Novel deep learning methods are used to improve the event selection and energy reconstruction, statistical methods are refined to better identify a signal of point-like neutrino sources above the atmospheric background, and multi-frequency techniques help to identify counterpart candidates.

While many questions remain still unresolved, the field of multimessenger astronomy has seen significant progress over the period during which this Ph.D. was completed. Most notably, the association of the IBL/HBL blazar TXS 0506+056[1,2,3] with a high-energy neutrino has provided the first glimpse beyond the curtain of the large background of atmospheric events. Although the physical processes of the neutrino emission from TXS 0506+056 are still disputed, we have seen in chapter 4 that, based on all the available multifrequency information, especially the timing and the energetics of the emission, the association is consistent. Therefore, this observation motivates an even deeper look into the available data to associate a larger fraction of the astrophysical signal to their sources.

At trigger level, the IceCube Observatory measures events at a rate of $\sim 2.5\,\mathrm{kHz}$, of which the vast majority are of atmospheric origin. Thus, sophisticated reconstruction and classification algorithms are used to extract potential signal events. Even though the IceCube Observatory has been operational for a decade, rapid developments in machine learning and data science since then still provide room for improvements. In this thesis, two new deep learning-based algorithms have been presented: a deep neural network event-type classifier[4] and a deep neural network energy reconstruction for upgoing muons. The DNN classifier thereby provides — for the first time — a direct mapping between the measured pulses in the detector and the event's topology. This approach's potential has been discussed exemplarily through the development of a new realtime stream of high-energy starting events. Compared to its predecessor, the stream's accuracy and efficiency improve significantly, with event

[1] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Al Samarai, et al., "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A".

[2] IceCube Collaboration, M. G. Aartsen, Ackermann, Adams, Aguilar, Ahlers, M. Ahrens, Samarai, et al., "Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert".

[3] Padovani, Giommi, et al., "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source".

[4] Kronmueller and Glauch, "Application of Deep Neural Networks to Event Type Classification in IceCube".

rates more than doubled. The DNN energy estimator, on the other hand, increases the energy resolution of muons by up to 50% compared to previous, likelihood-based methods while also resolving biases at energies below 1 TeV. In addition to these improvements, DNN-based reconstructions are also extremely quick. Once the mapping between pulses and output quantity has been *learned*, applying the algorithm to new data only takes around 100 ms, allowing large scale applications at an early stage of the processing pipeline. Although the applications presented in this thesis have shed some light on the potential of deep learning algorithms in neutrino telescopes, many developments are yet to come. In recent years, graph neural networks, recurrent neural networks, and adversarial neural networks, together with large scale Monte Carlo simulations, have opened new opportunities to exploit further the potential of deep learning[5] and to extend their application to geometries other than the hexagonal shape of IceCube.

In addition to the search of multi-frequency counterparts to the most promising astrophysical neutrino candidates, the IceCube Collaboration has, over the last nine years, collected a sample of around 650,000 neutrino-induced tracks from the Northern Hemisphere ($\delta > -3^\circ$) in an energy range between 100 GeV and several PeV. Chapter 7 has presented a new approach to search this sample for accumulations of neutrinos that are consistent with a point-like astrophysical source. Compared to previous analyses, the novel modeling of the likelihood function provides unbiased fits and allows for better source localization and measurements of the source spectrum. In addition, the DNN energy estimator improves the energy resolution by up to 50%. Analyzing a catalog containing 110 potential neutrino sources reveals the emission from a starburst galaxy — NGC 1068 — at the level of 4.2σ, with a soft spectral index of $\gamma \approx 3.2$. This source is therefore the most significant candidate of astrophysical neutrino emission to date. Given the source's soft spectrum, this observation will also require more work on theoretical astrophysical models for neutrino production.

Even with the associations of neutrino emission from TXS 0506+056 and NGC 1068, a large fraction of the astrophysical neutrino flux remains unresolved. In chapter 8, we have therefore expanded the multi-frequency dissection strategy for TXS 0506+056 to the error regions of 80 high-energy tracks of IceCube. Most importantly, the study has found an excess of HBL/IBL blazars at a 3.23σ level.[6] While the study does not make clear associations of neutrinos with a source, a sample of 20 regions with a unique HBL/IBL counterpart have been compiled, of which $\sim$ 50% are expected to be neutrino emitters. To take this further, we currently conduct large-scale multi-frequency analyses, including dedicated optical observations with the *Very Large Telescope* and additional X-ray observations with Swift. Combining this data with source models, we hope to identify an emission profile of a subclass of HBL/IBL blazars that is consistent with neutrino production. An example of such a study has

[5] https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF3_CompF0-019.pdf.

[6] Giommi, Glauch, et al., "Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection".

been discussed to the scope of another energetically plausible HBL neutrino source candidate, 3HSP J095507.9+355101, that has been detected through the automatic version of our multi-frequency dissection analysis.[7]

It has now been shown by many analyses that the flux of astrophysical neutrinos is not dominated by a few strong sources, but rather by populations of weak objects[8,9].[10] With the associations of TXS 0506+056 and NGC 1068, it further appears that it is not only one, but rather different types of sources which are contributing. As observed many times in astronomy, source populations can appear at different energies with different spectra. For example, there could be a population of starburst galaxies at lower energies while blazars significantly contribute to the high-TeV to PeV regime. To finally settle the question of the origin of neutrinos and cosmic rays, it will thus be indispensable to further increase event rates and astronomical data coverage to get a deeper view into our universe. Conveniently, experience from simulation and data analysis techniques developed in IceCube can be transferred to newly constructed telescopes. Future neutrino telescopes will therefore likely rely on deep learning methods to optimize their data pipelines and coordinate observations in real time with other multi-frequency telescopes around the globe. Four new neutrino telescopes are currently in development: IceCube Gen-2 at the South Pole[11],[12] KM3NET in the Mediterranean sea,[13] Baikal-GVD at Lake Baikal in Russia,[14] and P-ONE in the Cascadia Basin in Canada.[15] Altogether, they will significantly increase event rates and provide good coverage over the entire sky. This will further accelerate the development of the field, bringing us a few steps closer to fully understanding the origin of cosmic rays.

[7] Giommi, Padovani, Oikonomou, et al., "3HSP J095507.9+355101: a flaring extreme blazar coincident in space and time with IceCube-200107A".

[8] Glauch and Turcati, "Search for weak neutrino point sources using angular auto-correlation analyses in IceCube".

[9] M. G. Aartsen, K. Abraham, et al., "All-sky Search for Time-integrated Neutrino Emission from Astrophysical Sources with 7 yr of IceCube Data".

[10] Huber, "Multi-Messenger correlation study of Fermi-LAT blazars and high-energy neutrinos observed in IceCube".

[11] M. Aartsen et al., "Neutrino astronomy with the next generation IceCube Neutrino Observatory".

[12] M. Aartsen et al., "IceCube-Gen2: The Window to the Extreme Universe".

[13] Margiotta, "The KM3NeT deep-sea neutrino telescope".

[14] Avrorin et al., "Baikal-GVD: status and prospects".

[15] Agostini et al., "The Pacific Ocean Neutrino Experiment".

Acknowledgement / Danksagung

The work presented in this thesis is by no means just the product of a single person, but rather the result of continuous exchange, support, and inspiration by many. Therefore, I want to take the opportunity to thank several people that have particularly accompanied me over the years of my Ph.D.

First of all, I want to deeply thank Elisa Resconi and Paolo Giommi for giving me the opportunity to spend three exciting years as an affiliate of the TUM Institute of Advanced Studies. Together with Paolo Padovani you have taught me a lot about research, life, and astronomy. Our scientific work has always been meaningful and interesting to me.

Several people have significantly contributed to the outcome of this thesis. Thanks a lot to Anna Franckowiak for sharing endless days (and nights) working on the interpretation of IceCube-170922A, to Hans Niederhausen, Tomas Kontrimas, and Chiara Bellenghi for always keeping up the spirit during our joint work on the point source analysis, and to my students Maximilian Kronmüller and Johannes Kager for developing the basis of the deep neural network event classifier. Thanks a lot also to Christian Haack, Foteini Oikonomou, and Andrea Turcati for doing major proofreading of this thesis.

My deep gratitude also to my colleagues. I've always appreciated the time we have spend over several coffee breaks, matches of Mario Tennis, group meetings, and scientific, as well as private discussion. Thanks specifically to Matthias & Andrea, who have been there since my first day and have accompanied me to several conferences and summer schools. It has always been a pleasure!

While science gives life an inspiring context, it requires great people to enjoy it in all its facets. Therefore I'd like to give my deep thanks to my flatmates Annika and Kira for always making home a wonderful place with joyful discussions and great food. Thanks to Fiona for being a great IAS-Buddy and to Franzi for providing harmony in times of stressful writing. Finally, it has been my family that helped me all the way to where I am now. Your faith and support have always been a major source of energy for me.

My special gratitude goes to Stefan Coenders for being the person that introduced me to the wonderful IceCube collaboration several years ago.

A Abbreviations

Table A.1: This table provides an overview of abbreviations that are used commonly throughout this thesis.

Acronym	Meaning
XFGL	The Xth generation of the Fermi FGL catalog
AGN	Active galactic nucleus
BZCAT	Multi-frequency catalog of blazars
CNN	Convolutional neural network
DNN	Deep neural network
DOM	Digital optical module
FSRQ	Flat spectrum radio quasar
GPU	Graphics processing unit
HESE	High-energy starting events
HBL	Blazar of type BL Lac with high synchroton peak frequency
HSP	Catalog of high synchrotron peaked blazars
IBL	Blazar of type BL Lac with intermediate synchroton peak frequency
ICXX	IceCube detector season with XX strings
KDE	Kernel density estimation
MESE	Medium-energy starting events
MLE	Maximum likelihood estimator
LBL	Blazar of type BL Lac with low synchroton peak frequency
RNN	recurrent neural network

B Pseudo-Code Implementations of the DNN Models

B.1 DNN Classifier

```
inputs = OrderedDict()

inputs["Branch_IC_time"] = {"variables":
                            ['IC_charge','IC_time_first', 'IC_charge_10ns',
                             'IC_charge_50ns','IC_charge_100ns',
                             'IC_time_spread', 'IC_time_std',
                             'IC_time_weighted_median',
                             'IC_pulse_0_01_pct_charge_quantile',
                             'IC_pulse_0_03_pct_charge_quantile',
                             'IC_pulse_0_05_pct_charge_quantile',
                             'IC_pulse_0_11_pct_charge_quantile',
                             'IC_pulse_0_15_pct_charge_quantile',
                             'IC_pulse_0_2_pct_charge_quantile',
                             'IC_pulse_0_5_pct_charge_quantile',
                             'IC_pulse_0_8_pct_charge_quantile'],
                     "transformations":
                            [tr.IC_divide_100, tr.IC_divide_10000,
                             tr.IC_divide_100, tr.IC_divide_100,
                             tr.IC_divide_100, tr.IC_divide_10000,
                             tr.IC_divide_10000, tr.IC_divide_10000,
                             tr.IC_divide_10000,tr.IC_divide_10000,
                             tr.IC_divide_10000, tr.IC_divide_10000,
                             tr.IC_divide_10000, tr.IC_divide_10000,
                             tr.IC_divide_10000, tr.IC_divide_10000]}


# define outputs for each branch
outputs = OrderedDict()
outputs["Out1"] = {"variables": ["classification"],
                   "transformations": [tr.oneHotEncode_new]}
loss_weights = {'Target1': 1.}
loss_functions = ["categorical_crossentropy"]
metrics = ["acc"]

output_names = {0: 'Skimming', 1: 'Cascade', 2: 'Through_Going_Track',
                3:'Starting_Track', 4:'Stopping_Track'}


def inception_block4(input_tensor, n, t0=2, t1=4,
```

```
                     t2=5, n_pool=3, scale=0.1):


    tower_0 = conv3d_bn(input_tensor, n, (t0,1,1),
                        padding='same')
    tower_0 = conv3d_bn(tower_0, n, (1,t0,1),
                        padding='same')
    tower_0 = conv3d_bn(tower_0, n, (1,1,t0),
                        padding='same')

    tower_1 = conv3d_bn(input_tensor, n, (t1,1,1),
                        padding='same')
    tower_1 = conv3d_bn(tower_1, n, (1,t1,1),
                        padding='same')
    tower_1 = conv3d_bn(tower_1, n, (1,1,t1),
                        padding='same')

    tower_4 = conv3d_bn(input_tensor, n, (1,1,t2),
                        padding='same')

    tower_3 = MaxPooling3D(
                (n_pool, n_pool, n_pool),strides=(1,1,1),
                padding='same')(input_tensor)
    tower_3 = conv3d_bn(tower_3, n, (1,1,1),
                        padding='same')

    up = concatenate(
        [tower_0, tower_1, tower_3, tower_4],
        axis = channel_axis)
    return up


def inception_resnet(input_tensor, n, t1=2, t2=3,
                     n_pool=3, scale=0.1):


    tower_1 = conv3d_bn(input_tensor, n, (1,1,1),
                        padding='same')
    tower_1 = conv3d_bn(tower_1, n, (t1,1,1),
                        padding='same')
    tower_1 = conv3d_bn(tower_1, n, (1,t1,1),
                        padding='same')
    tower_1 = conv3d_bn(tower_1, n, (1,1,t1),
                        padding='same')

    tower_2 = conv3d_bn(input_tensor, n, (1,1,1),
                        padding='same')
    tower_2 = conv3d_bn(tower_2, n, (t2,1,1),
                        padding='same')
    tower_2 = conv3d_bn(tower_2, n, (1,t2,1),
                        padding='same')
    tower_2 = conv3d_bn(tower_2, n, (1,1,t2),
                        padding='same')

    tower_3 = MaxPooling3D(
                    (n_pool, n_pool, n_pool),strides=(1,1,1),
                    padding='same')(input_tensor)
```

```
    tower_3 = conv3d_bn(tower_3, n, (1,1,1), padding='same')

    up = concatenate(
        [tower_1, tower_2, tower_3], axis = channel_axis)

    x = Lambda(
                lambda inputs, scale: inputs[0] + inputs[1] * scale,
                output_shape=K.int_shape(input_tensor)[1:],
                arguments={'scale': scale},)([input_tensor, up])

    return x


#Model definition

def model(input_shapes, output_shapes):
    # The Input
    input_b1 = Input(
        shape=input_shapes['Branch_IC_time']['general'],
        name = "Input-Branch1")


    # Hidden Layers
    z1 = inception_block4(input_b1, 24, t0=3, t1=7, t2=10)
    z1 = inception_block4(z1, 24, t0=2, t1=3, t2=7)
    z1 = inception_block4(z1, 24, t0=2, t1=4, t2=8)
    z1 = inception_block4(z1, 24, t0=3, t1=5, t2=9)
    z1 = inception_block4(z1, 24, t0=3, t1=8, t2=9)
    z1 = AveragePooling3D(pool_size=(2, 2, 3))(z1)
    z1 = BatchNormalization()(z1)
    for i in range(8):
        z1 = inception_resnet(z1, 32, t2=3)
        z1 = inception_resnet(z1, 32, t2=4)
        z1 = inception_resnet(z1, 32, t2=5)
    z1 = AveragePooling3D(pool_size=(1, 1, 2))(z1)
    z1 = BatchNormalization()(z1)
    for i in range(8):
        z1 = inception_resnet(z1, 32, t2=3)
        z1 = inception_resnet(z1, 32, t2=4)
        z1 = inception_resnet(z1, 32, t2=5)
    z1 = Conv3D(4096, (1, 1, 1), activation='relu',
            padding="same", name='conv1x1x1')(z1)
    z1 = GlobalAveragePooling3D()(z1)
    output_b1 = Dense(output_shapes["Out1"]["general"][0],
                                    activation="softmax",
                                    name="Target1")(z1)
    # The Output
    model= Model(inputs=[input_b1],
                 outputs=[output_b1])
    return model
```

B.2 DNN Energy Estimator

```
inputs = OrderedDict()

inputs["Branch_IC_time"] = {"variables":
                              ['IC_charge','IC_time_first',
                               'IC_charge_10ns',
                               'IC_charge_50ns','IC_charge_100ns',
                               'IC_time_spread', 'IC_time_std',
                               'IC_time_weighted_median',
                               'IC_pulse_0_01_pct_charge_quantile',
                               'IC_pulse_0_03_pct_charge_quantile',
                               'IC_pulse_0_05_pct_charge_quantile',
                               'IC_pulse_0_11_pct_charge_quantile',
                               'IC_pulse_0_15_pct_charge_quantile',
                               'IC_pulse_0_2_pct_charge_quantile',
                               'IC_pulse_0_5_pct_charge_quantile',
                               'IC_pulse_0_8_pct_charge_quantile'],
                       "transformations":
                              [tr.IC_divide_100, tr.IC_divide_10000,
                               tr.IC_divide_100, tr.IC_divide_100,
                               tr.IC_divide_100, tr.IC_divide_10000,
                               tr.IC_divide_10000, tr.IC_divide_10000,
                               tr.IC_divide_10000,tr.IC_divide_10000,
                               tr.IC_divide_10000, tr.IC_divide_10000,
                               tr.IC_divide_10000, tr.IC_divide_10000,
                               tr.IC_divide_10000, tr.IC_divide_10000]}


# define outputs for each branch
outputs = OrderedDict()
outputs["Out1"] = {"variables": ["e_on_entry"],
                   "transformations": [tr.log10]}
loss_weights = {'Target1': 1.}
loss_functions = ["mse"]
metrics = ['mae']


def inception_block4(input_tensor, n, t0=2, t1=4,
                     t2=5, n_pool=3, scale=0.1):

    tower_0 = conv3d_bn(input_tensor, n, (t0,1,1),
                        padding='same')
    tower_0 = conv3d_bn(tower_0, n, (1,t0,1),
                        padding='same')
    tower_0 = conv3d_bn(tower_0, n, (1,1,t0),
                        padding='same')

    tower_1 = conv3d_bn(input_tensor, n, (t1,1,1),
                        padding='same')
    tower_1 = conv3d_bn(tower_1, n, (1,t1,1),
                        padding='same')
    tower_1 = conv3d_bn(tower_1, n, (1,1,t1),
                        padding='same')

    tower_4 = conv3d_bn(input_tensor, n, (1,1,t2),
```

```
                         padding='same')

     tower_3 = MaxPooling3D(
                  (n_pool, n_pool, n_pool),strides=(1,1,1),
                  padding='same')(input_tensor)
     tower_3 = conv3d_bn(tower_3, n, (1,1,1),
                         padding='same')

     up = concatenate(
         [tower_0, tower_1, tower_3, tower_4],
         axis = channel_axis)
     return up


def inception_resnet(input_tensor, n, t1=2, t2=3,
                     n_pool=3, scale=0.1):

     tower_1 = conv3d_bn(input_tensor, n, (1,1,1),
                         padding='same')
     tower_1 = conv3d_bn(tower_1, n, (t1,1,1),
                         padding='same')
     tower_1 = conv3d_bn(tower_1, n, (1,t1,1),
                         padding='same')
     tower_1 = conv3d_bn(tower_1, n, (1,1,t1),
                         padding='same')

     tower_2 = conv3d_bn(input_tensor, n, (1,1,1),
                         padding='same')
     tower_2 = conv3d_bn(tower_2, n, (t2,1,1),
                         padding='same')
     tower_2 = conv3d_bn(tower_2, n, (1,t2,1),
                         padding='same')
     tower_2 = conv3d_bn(tower_2, n, (1,1,t2),
                         padding='same')

     tower_3 = MaxPooling3D(
                     (n_pool, n_pool, n_pool),strides=(1,1,1),
                     padding='same')(input_tensor)
     tower_3 = conv3d_bn(tower_3, n, (1,1,1), padding='same')

     up = concatenate(
         [tower_1, tower_2, tower_3], axis = channel_axis)

     x = Lambda(
               lambda inputs, scale: inputs[0] + inputs[1] * scale,
               output_shape=K.int_shape(input_tensor)[1:],
               arguments={'scale': scale},)([input_tensor, up])

     return x


#Model definition

def model(input_shapes, output_shapes):

     # The Input
     input_b1 = Input(
```

```
    shape=input_shapes['Branch_IC_time']['general'],
    name = "Input-Branch1")


# Hidden Layers
z1 = inception_block4(input_b1, 18, t0=2, t1=5, t2=8)
z1 = inception_block4(z1, 18, t0=2, t1=3, t2=7)
z1 = inception_block4(z1, 18, t0=2, t1=4, t2=8)
z1 = inception_block4(z1, 18, t0=3, t1=5, t2=9)
z1 = inception_block4(z1, 18, t0=3, t1=8, t2=9)
z1 = AveragePooling3D(pool_size=(2, 2, 3))(z1)
z1 = BatchNormalization()(z1)
for i in range(6):
    z1 = inception_resnet(z1, 24, t2=3)
    z1 = inception_resnet(z1, 24, t2=4)
    z1 = inception_resnet(z1, 24, t2=5)
z1 = AveragePooling3D(pool_size=(1, 1, 2))(z1)
z1 = BatchNormalization()(z1)
for i in range(6):
    z1 = inception_resnet(z1, 24, t2=3)
    z1 = inception_resnet(z1, 24, t2=4)
    z1 = inception_resnet(z1, 24, t2=5)
z1 = Conv3D(64, (1, 1, 1), activation='relu',
        padding="same", name='conv1x1x1')(z1)
z1 = Conv3D(4, (1, 1, 1), activation='relu',
        padding="same", name='conv1x1x1_2')(z1)
z1 = AveragePooling3D(pool_size=(1, 1, 2))(z1)
z1 = Flatten(data_format=K.image_data_format())(z1)
z1 = Dense(120, name='dense1')(z1)
z1 = Dense(64, name='dense2')(z1)
z1 = Dense(16, name='dense3')(z1)
output_b1 = Dense(output_shapes["Out1"]["general"][0],
                                activation="linear",
                                name="Target1")(z1)
# The Output
model= Model(inputs=[input_b1],
             outputs=[output_b1])
return model
```

C Confusion Matrices of the DNN Classifier

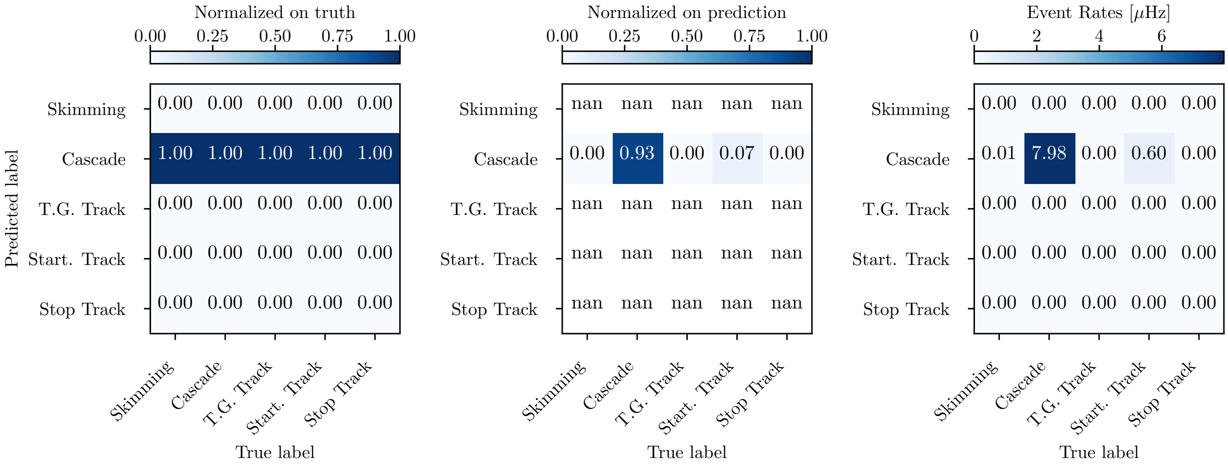

Figure C.1: Example of a simple cascade selection on level 2. Only events with true deposited energy larger than 5 TeV, cascade p-score larger than 0.5 and track p-score smaller 0.001 are chosen. The track p-score is defined as the maximum score of through-going and stopping prediction. As a result a fairly clean sample of cascades is selected without applying any further cuts on the direction. The plot includes a complete modelling of the atmospheric components and assumes the best-fit astrophysical flux of IceCube's latest cascade sample M. Aartsen et al., "Search for Sources of Astrophysical Neutrinos Using Seven Years of IceCube Cascade Events".

Normalized on truth

	Skimming	Cascade	T.G. Track	Start. Track	Stop Track
Skimming	0.82	0.04	0.02	0.03	0.01
Cascade	0.00	0.90	0.01	0.27	0.11
T.G. Track	0.17	0.01	0.92	0.04	0.08
Start. Track	0.00	0.03	0.02	0.64	0.05
Stop Track	0.01	0.01	0.03	0.03	0.76

Normalized on prediction

	Skimming	Cascade	T.G. Track	Start. Track	Stop Track
Skimming	0.85	0.00	0.14	0.00	0.01
Cascade	0.00	0.00	0.29	0.00	0.71
T.G. Track	0.02	0.00	0.97	0.00	0.01
Start. Track	0.00	0.00	0.75	0.00	0.25
Stop Track	0.01	0.00	0.22	0.00	0.78

Event Rates [Hz]

	Skimming	Cascade	T.G. Track	Start. Track	Stop Track
Skimming	1.51	0.00	0.25	0.00	0.01
Cascade	0.00	0.00	0.10	0.00	0.25
T.G. Track	0.31	0.00	13.36	0.00	0.17
Start. Track	0.00	0.00	0.32	0.00	0.11
Stop Track	0.02	0.00	0.48	0.00	1.72

Normalized on truth

	Skimming	Cascade	T.G. Track	Start. Track	Stop Track
Skimming	0.99	0.00	0.00	0.00	0.00
Cascade	0.00	0.97	0.00	0.02	0.01
T.G. Track	0.01	0.01	0.99	0.01	0.01
Start. Track	0.00	0.01	0.01	0.96	0.01
Stop Track	0.00	0.01	0.00	0.01	0.97

Normalized on prediction

	Skimming	Cascade	T.G. Track	Start. Track	Stop Track
Skimming	0.97	0.00	0.03	0.00	0.00
Cascade	0.00	0.02	0.15	0.00	0.83
T.G. Track	0.00	0.00	1.00	0.00	0.00
Start. Track	0.00	0.00	0.76	0.02	0.22
Stop Track	0.00	0.00	0.03	0.00	0.97

Event Rates [Hz]

	Skimming	Cascade	T.G. Track	Start. Track	Stop Track
Skimming	0.48	0.00	0.02	0.00	0.00
Cascade	0.00	0.00	0.00	0.00	0.01
T.G. Track	0.01	0.00	6.70	0.00	0.01
Start. Track	0.00	0.00	0.04	0.00	0.01
Stop Track	0.00	0.00	0.02	0.00	0.87

Normalized on truth

	Skimming	Cascade	T.G. Track	Start. Track	Stop Track
Skimming	1.00	0.00	0.00	0.00	0.00
Cascade	0.00	0.95	0.00	0.00	0.00
T.G. Track	0.00	0.01	0.99	0.00	0.00
Start. Track	0.00	0.01	0.00	1.00	0.01
Stop Track	0.00	0.02	0.00	0.00	0.99

Normalized on prediction

	Skimming	Cascade	T.G. Track	Start. Track	Stop Track
Skimming	0.99	0.00	0.01	0.00	0.00
Cascade	0.00	0.94	0.00	0.06	0.00
T.G. Track	0.00	0.00	1.00	0.00	0.00
Start. Track	0.00	0.00	0.75	0.06	0.19
Stop Track	0.00	0.00	0.01	0.00	0.99

Event Rates [Hz]

	Skimming	Cascade	T.G. Track	Start. Track	Stop Track
Skimming	0.18	0.00	0.00	0.00	0.00
Cascade	0.00	0.00	0.00	0.00	0.00
T.G. Track	0.00	0.00	3.07	0.00	0.00
Start. Track	0.00	0.00	0.01	0.00	0.00
Stop Track	0.00	0.00	0.00	0.00	0.46

Figure C.2: Confusion matrices above 100 GeV for a full IceCube level 2 Monte Carlo dataset including the dominant atmospheric muon background. From top to bottom a cut is place on the maximum p-score p_{max} of 0.5, 0.95 and 0.99, respectively.

Figure C.3: Confusion matrices above 10 TeV for a full IceCube level 2 Monte Carlo dataset including the dominant atmospheric muon background. From top to bottom a cut is place on the maximum p-score p_{max} of 0.5, 0.95 and 0.99, respectively.

D Additional Material for the Point Source Analysis

D.1 Efficient Minimization of the Likelihood Function

For a given source position the point source signal likelihood has two free parameters, the signal strength μ_s and the spectral index γ, which need to be maximized according to the maximum likelihood principle. Numerically, this requires the gradient of eq. (7.31), which is given by

$$\frac{\partial(\mathcal{TS})}{\partial\gamma} = \sum_{i=1}^{N} \frac{1}{(f_{s,i} - f_{b,i}) + \frac{\mu_{tot}}{\mu_s} f_{b,i}} \left[\frac{\partial \mathcal{S}_i}{\partial\gamma} \mathcal{E}_i + \mathcal{S}_i \frac{\partial \mathcal{E}_i}{\partial\gamma} \right] \tag{D.1}$$

for the spectral index and

$$\frac{\partial(\mathcal{TS})}{\partial\mu_s} = 2 \times \left(\frac{N}{\mu_{tot}} - 1 \right) + 2 \times \sum_{i=1}^{N} \frac{1}{\mu_s \left(\frac{f_{s,i}}{f_{b,i}} - 1 \right) + \mu_{tot}} \left(\frac{f_{s,i}}{f_{b,i}} - 1 \right) \tag{D.2}$$

for the signal strength. For readability, function arguments have been suppressed and $\mathcal{S}$ and $\mathcal{E}$ represent the spatial and energy term, respectively. It can be seen that while the μ_s gradient only depends on the function values of the pdfs in the likelihood, the γ gradient additionally requires their derivatives. Note that eq. (D.1) and eq. (D.2) are maximizing directly the test statistic. This is only valid as the background likelihood doesn't have any free parameters to be maximized and is mainly done for simplification. In a case where the background likelihood has additional free parameters, the signal and background likelihood need to be maximized separately.

As the minimization procedure requires several steps, and thereby many spline evaluations per event, performance considerations are of large importance. In a first step, computation time can be reduced by only evaluating events in a certain radius around the source position. For all the events outside this radius, f_b is much larger than f_s and therefore the addend in the test statistic definition, eq. (7.31), can be simplified to

$$\log \left\{ \frac{\mu_s}{\mu_{tot}} \left(\frac{f_s\left(x_i \,|\, \theta_s\right)}{f_b\left(x_i \,|\, \theta_b\right)} - 1 \right) + 1 \right\} \to \log \left\{ 1 - \frac{\mu_s}{\mu_{tot}} \right\} \approx -\frac{\mu_s}{\mu_{tot}}$$

which is independent of f_b and f_s. Events outside the radius thereby only contribute a constant factor of $-1/\mu_{tot}$ to the gradient in eq. (D.2) while completely vanishing in eq. (D.1). This simplification reduces the computing time for the generation of one pseudo-experiment, i.e., the injection of background and signal events and the minimization of the likelihoods, to $\sim 200\,\mathrm{ms}$ when using a box of 15 degrees radius around the source position, because only around 3% of the total number of events in the sky have to be evaluated. Further improvement can be reached by making the box a function of the reconstructed energy and the angular uncertainty estimation. In our case, for each E_{DNN} and σ_{BDT} bin the 99.8% quantile of the point-spread function is calculated. For stability reasons a lower threshold is set to 4 degrees, while an upper threshold of 15 degrees is used wherever the 99.8% quantile is larger or not sufficient Monte Carlo statistic is available to reliably calculate the quantile. Using only events that fall in the dynamic box further reduces the amount of events that are evaluated to ~ 2 % and thereby reducing the computing time to $\sim 150\,\mathrm{ms/trial}$.

All the calculations for the point source analysis are implemented in a new tool python-based tool called *SkyLLH* [1]. It is designed to perform likelihood analysis in a wide range of astrophysical applications, not being limited to neutrino point source searches.[2] It also provides an interface for reading and interpolating the previously generated pdfs and to perform a gradient-based minimization of the likelihood function.

[1] https://github.com/IceCubeOpenSource/skyllh

[2] Wolf, "SkyLLH – A generalized Python-based tool for log-likelihood analyses in multi-messenger astronomy".

D.2 Background Test Statistic Distributions

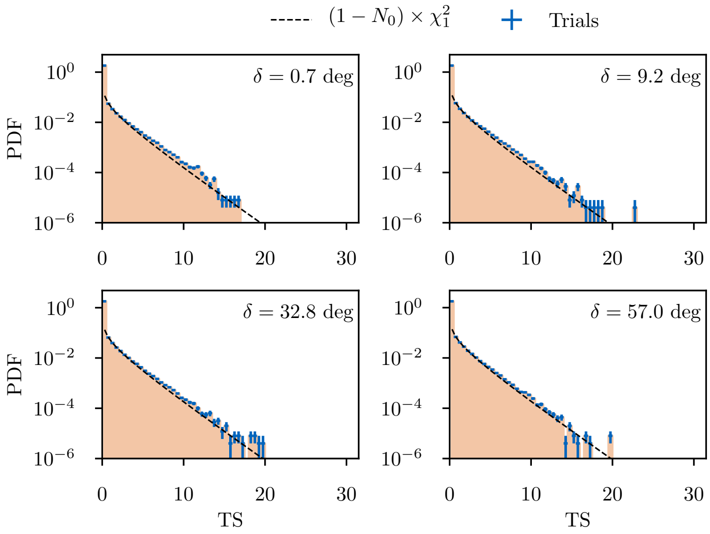

Figure D.1: Background test statistic distributions for various declinations assuming a fixed spectral index of $\gamma = 2.0$. A χ^2_1 is indicated as a black dashed line after normalizing for the number of zero trials N_0 which is around 80%.

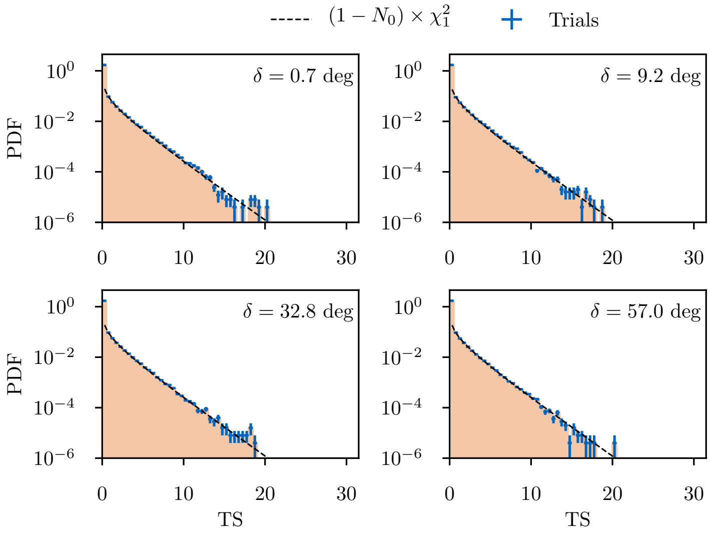

Figure D.2: Background test statistic distributions for various declinations assuming a fixed spectral index of $\gamma = 2.5$. A χ^2_1 is indicated as a black dashed line after normalizing for the number of zero trials N_0, which is around 70%.

D.3 Point Spread Functions

Figure D.3: Comparison of the KDE spatial term and the Rayleigh approximation for a **spectral index of $\gamma = 1.0$**. From top to bottom, the plots are shown for increasing DNN energy and a common BDT angular uncertainty. The Monte Carlo truth is given for different zenith bands to highlight its invariance.

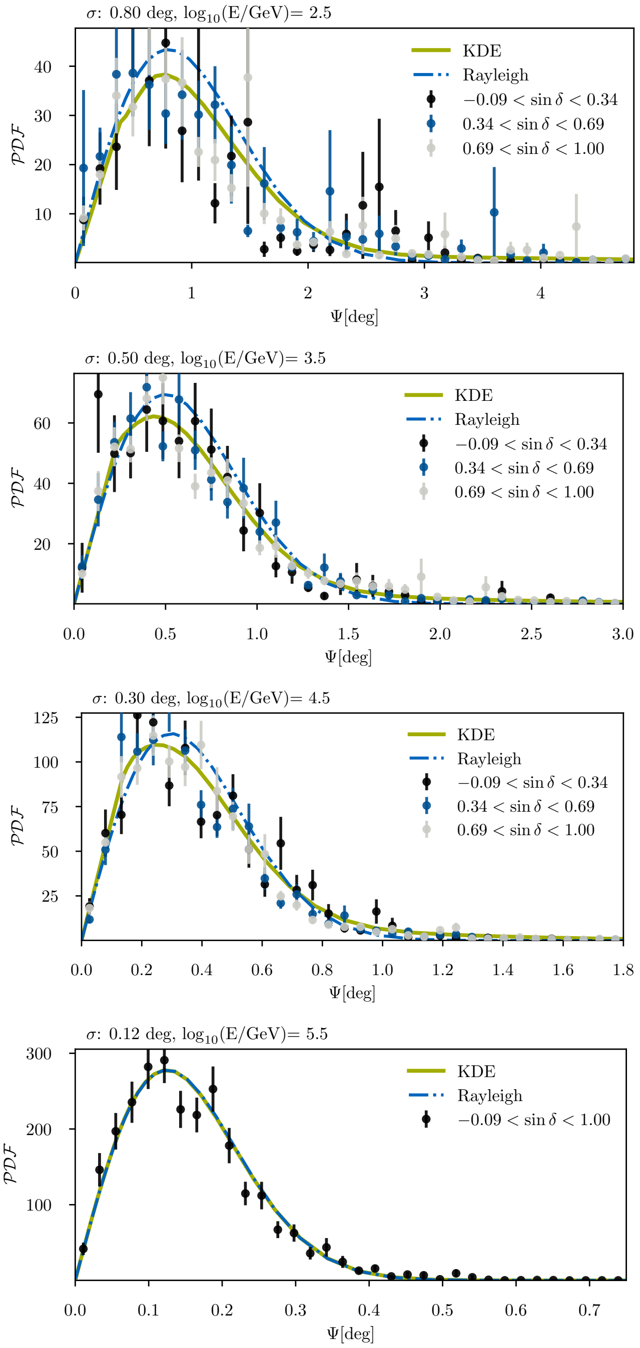

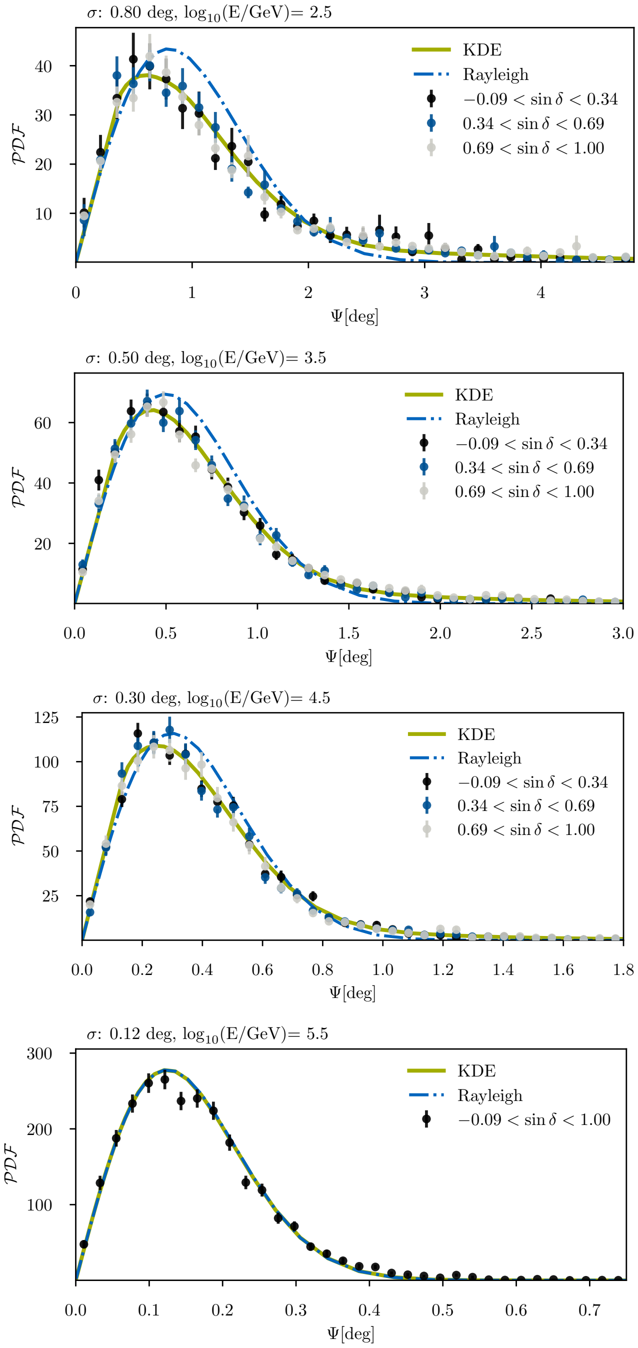

Figure D.4: Comparison of the KDE spatial term and the Rayleigh approximation for a **spectral index of $\gamma = 1.5$**. From top to bottom, the plots are shown for increasing DNN energy and a common BDT angular uncertainty. The Monte Carlo truth is given for different zenith bands to highlight its invariance.

Figure D.5: Comparison of the KDE spatial term and the Rayleigh approximation for a **spectral index of $\gamma = 2.5$**. From top to bottom, the plots are shown for increasing DNN energy and a common BDT angular uncertainty. The Monte Carlo truth is given for different zenith bands to highlight its invariance.

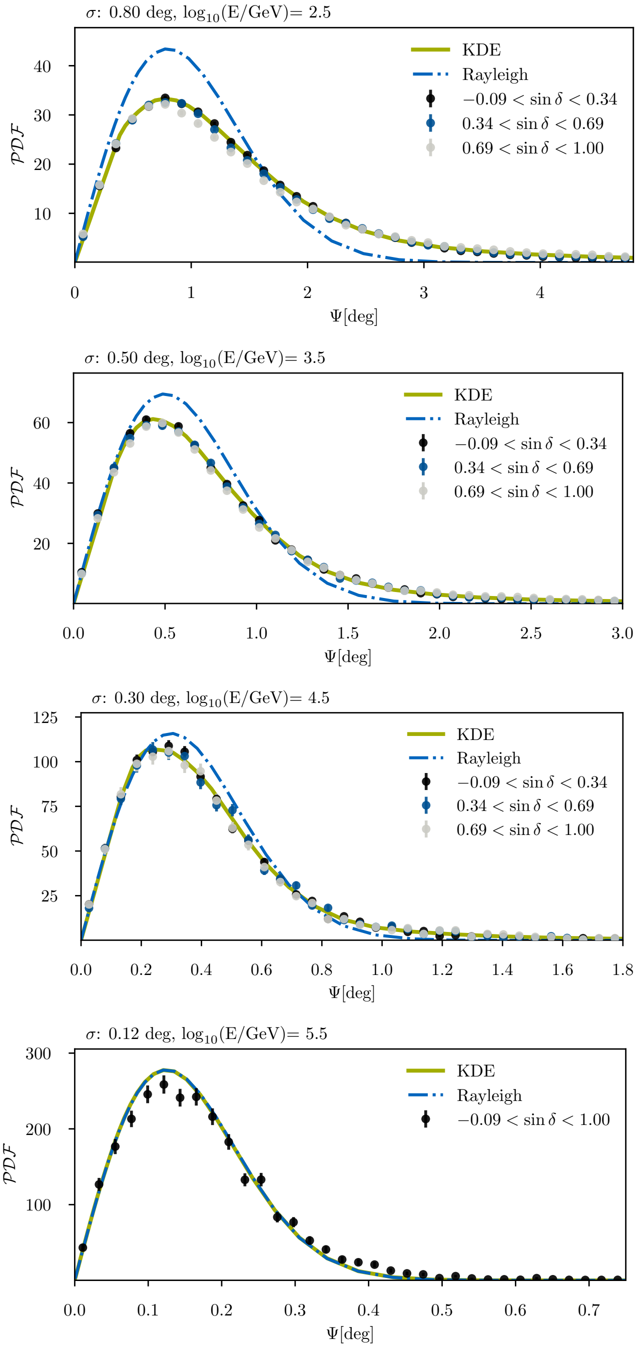

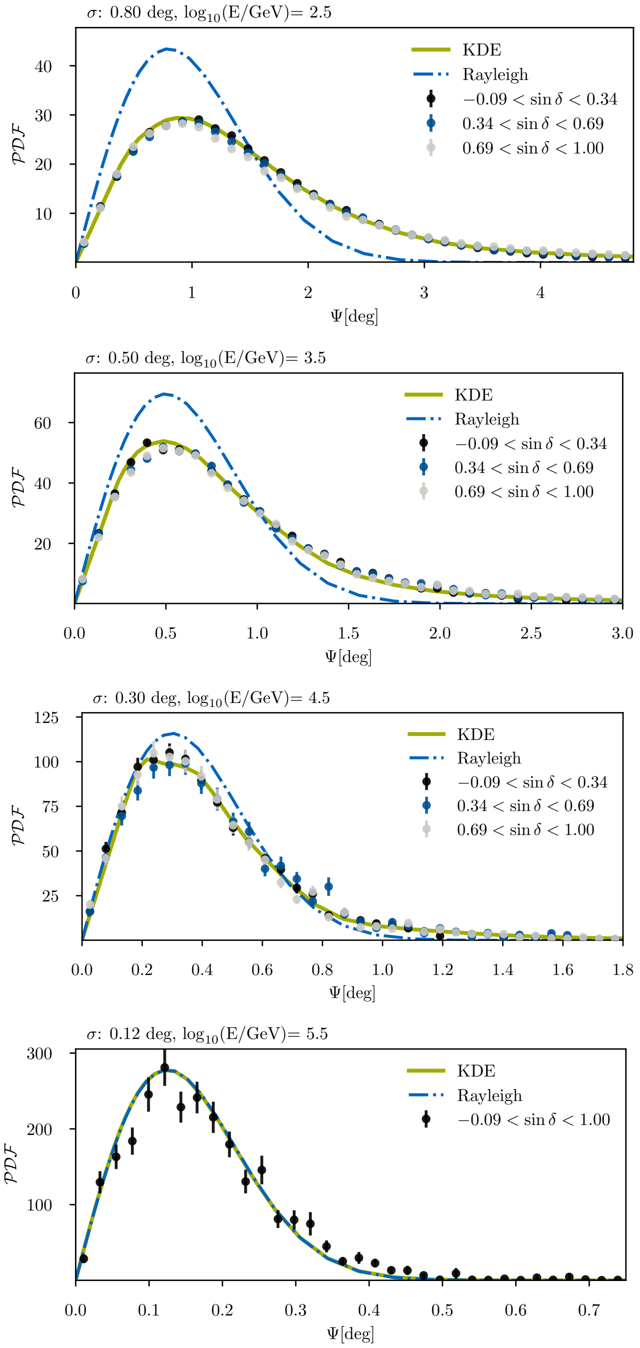

Figure D.6: Comparison of the KDE spatial term and the Rayleigh approximation for a **spectral index of $\gamma = 4.0$**. From top to bottom, the plots are shown for increasing DNN energy and a common BDT angular uncertainty. The Monte Carlo truth is given for different zenith bands to highlight its invariance.

D.4 Fit Biases at Various Declinations

Figure D.7: Fitted versus injected signal for a $\gamma = $ **2.0 source spectrum** and four different declinations. Left and right plot show the recovery of the number of signal events and the spectral index as a function of the number of injected events. The unbiased expectation is shown a white dashed lines. The results for the new and old analysis are shown in different colors as solid lines and shaded bands for the median and central 68% quantile, respectively.

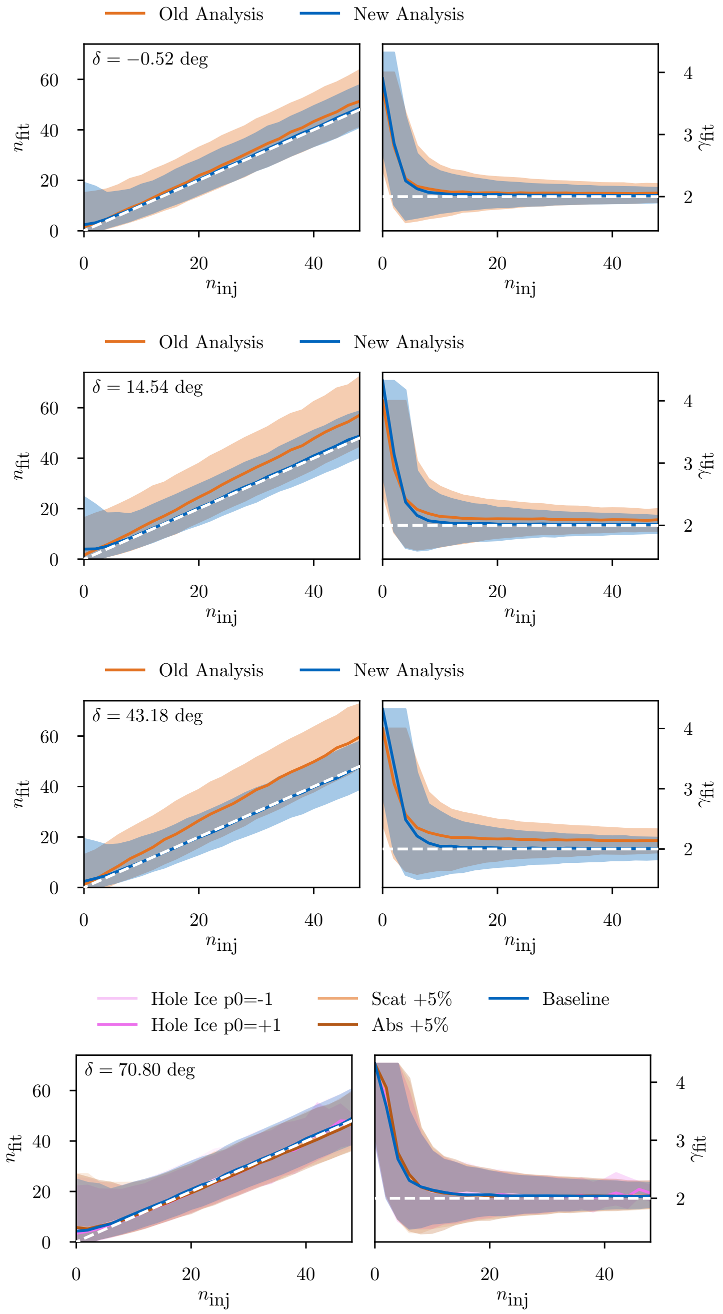

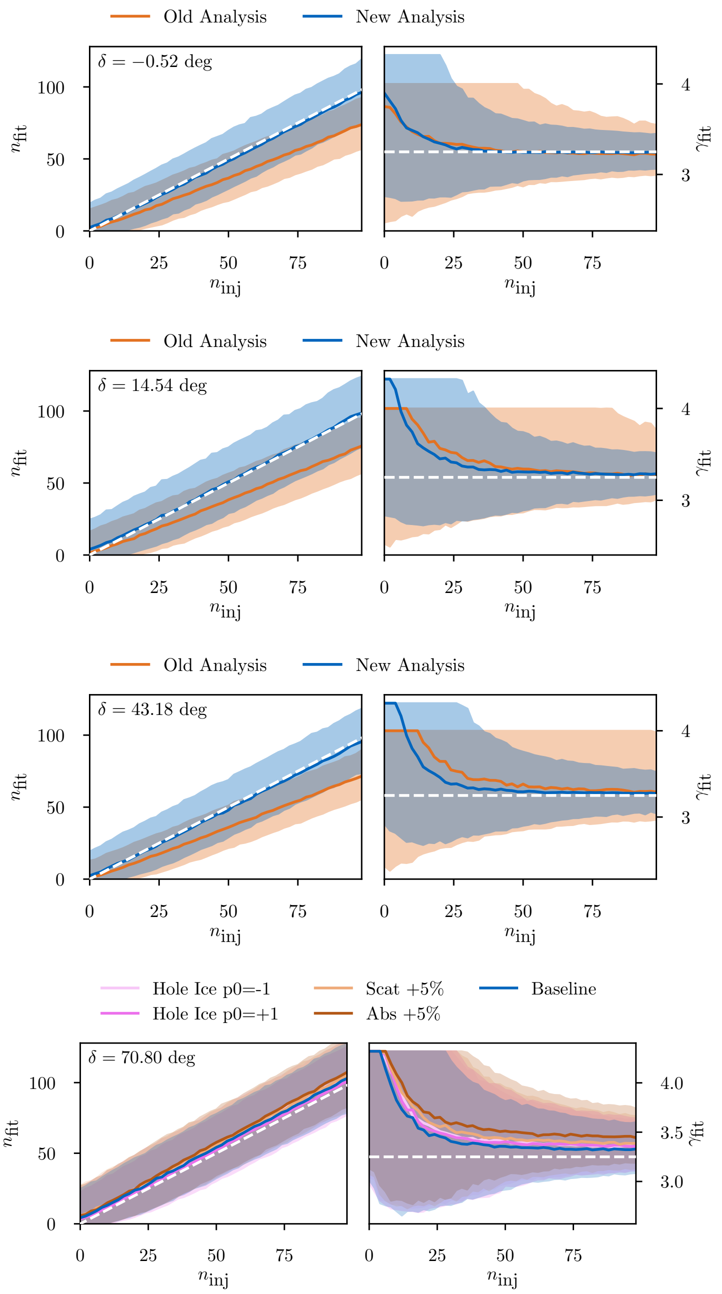

Figure D.8: Fitted versus injected signal for a $\gamma = 3.25$ **source spectrum** and four different declinations. Left and right plot show the recovery of the number of signal events and the spectral index as a function of the number of injected events. The unbiased expectation is shown as white dashed lines. The results for the new and old analysis are shown in different colors as solid lines and shaded bands for the median and central 68% quantile, respectively.

D.5 *Impact of Systematic Uncertainties on the Fit Biases*

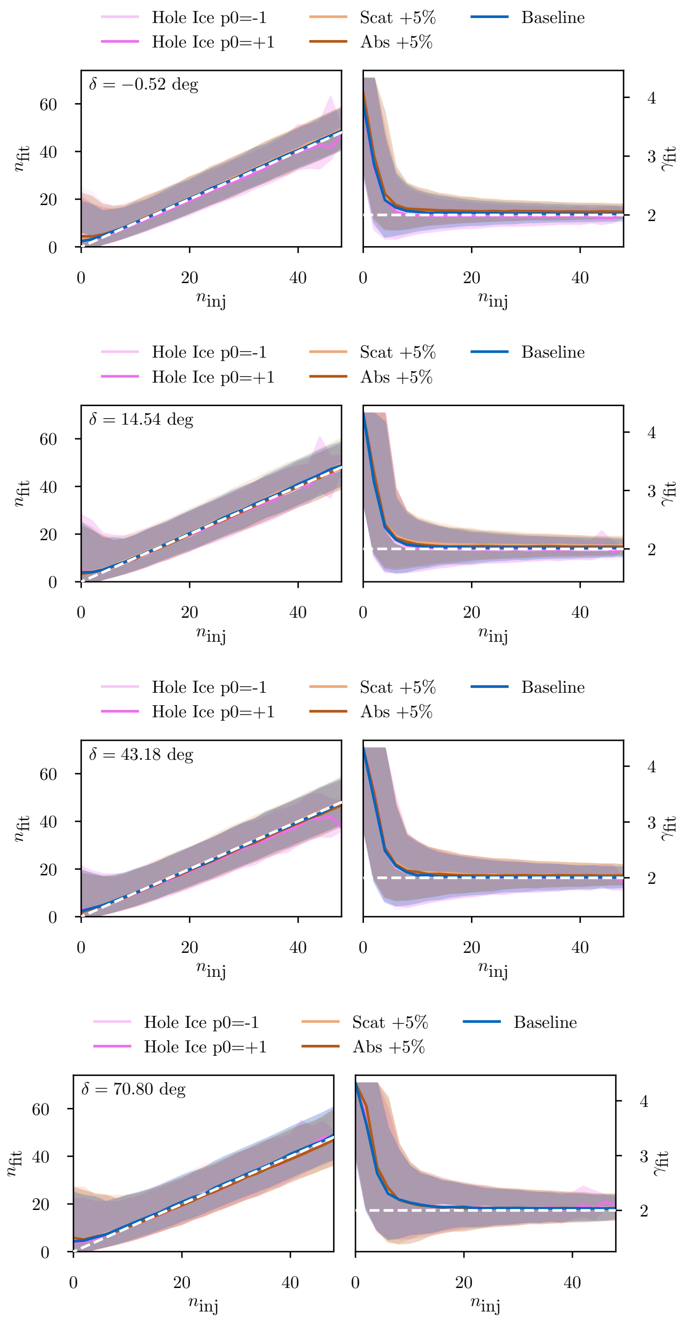

Figure D.9: Fitted versus injected signal for a $\gamma = 2.0$ **source spectrum** and four different declinations. Left and right plot show the recovery of the number of signal events and the spectral index as a function of the number of injected events. The unbiased expectation is shown as a white dashed line. The results for different systematic effects are shown in different colors, while solid lines and shaded bands represent the median and central 68% quantile, respectively.

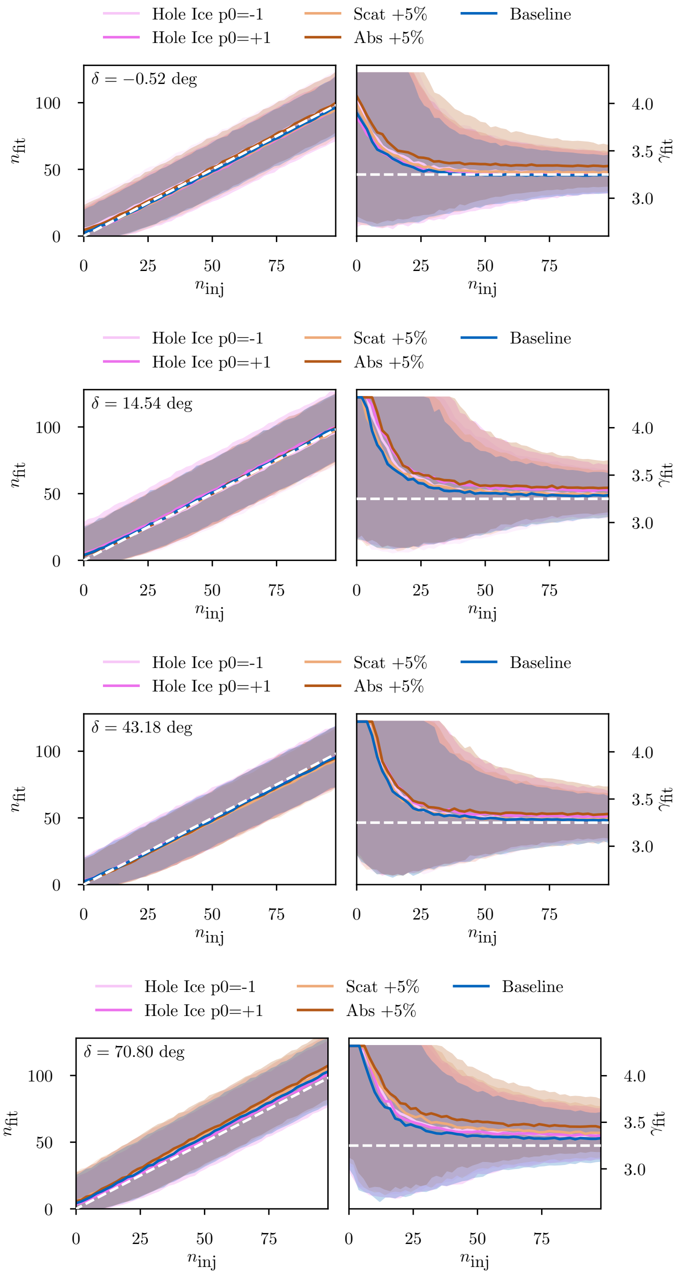

Figure D.10: Fitted versus injected signal for a $\gamma = 3.25$ **source spectrum** and four different declinations. Left and right plot show the recovery of the number of signal events and the spectral index as a function of the number of injected events. The unbiased expectation is shown as a white dashed line. The results for different systematic effects are shown in different colors, while solid lines and shaded bands represent the median and central 68% quantile, respectively.

D.6 Trial Correction Functions

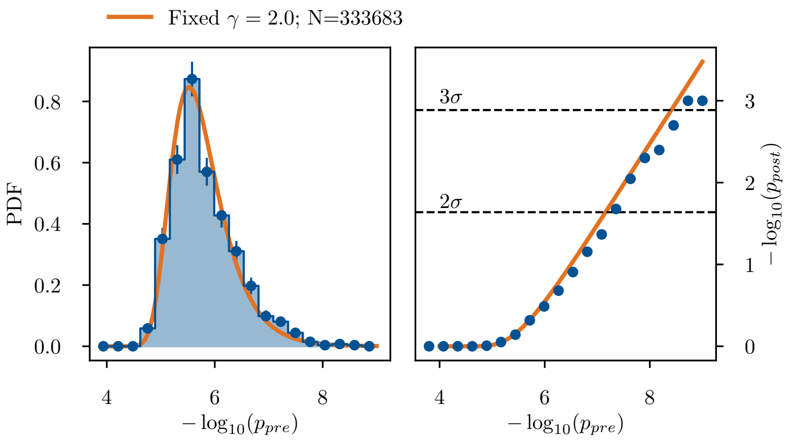

Figure D.11: The trial correction function for the case of a fixed spectral index of $\gamma = 2.0$. The left plot shows the distribution of p-values from 1000 pseudo-experiments assuming the background hypothesis to be true. The fitted beta distribution is shown as orange line. The corresponding cumulative function that can be used to convert the pre-trial into a post-trial p-value is shown in the right plot.

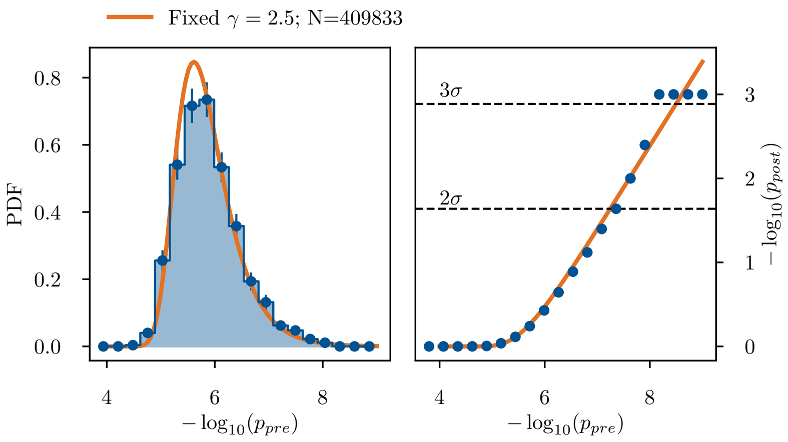

Figure D.12: The trial correction function for the case of a fixed spectral index of $\gamma = 2.5$. The left plot shows the distribution of p-values from 1000 pseudo-experiments assuming the background hypothesis to be true. The fitted beta distribution is shown as orange line. The corresponding cumulative function that can be used to convert the pre-trial into a post-trial p-value is shown in the right plot.

E Bibliography

Aab, A. et al. “Observation of a Large-scale Anisotropy in the Arrival Directions of Cosmic Rays above 8×10^{18} eV”. *Science* 357.6537 (2017) 1266–1270. arXiv:1709.07321 [astro-ph.HE] (cit. on p. 7).

Aartsen, M. G. et al. “Characterization of the Atmospheric Muon Flux in IceCube”. *Astropart. Phys.* 78 (2016) 1–27. arXiv:1506.07981 [astro-ph.HE] (cit. on p. 71).

– “Flavor Ratio of Astrophysical Neutrinos above 35 TeV in IceCube”. *Phys. Rev. Lett.* 114.17 (2015) 171102. arXiv:1502.03376 [astro-ph.HE] (cit. on p. 71).

– “Measurement of South Pole ice transparency with the IceCube LED calibration system”. *Nucl. Instrum. Meth.* A711 (2013) 73–89. arXiv:1301.5361 [astro-ph.IM] (cit. on pp. 44, 81).

– “Measurements using the inelasticity distribution of multi-TeV neutrino interactions in IceCube”. *Phys. Rev.* D99.3 (2019) 032004. arXiv:1808.07629 [hep-ex] (cit. on pp. 36, 72).

– “Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data”. *Astrophys. J.* 833.1 (2016) 3. arXiv:1607.08006 [astro-ph.HE] (cit. on pp. 71, 114, 141, 142).

– “Search for steady point-like sources in the astrophysical muon neutrino flux with 8 years of IceCube data”. *Eur. Phys. J.* C79.3 (2019) 234. arXiv:1811.07979 [hep-ph] (cit. on pp. 71, 101, 114, 140).

– “Time-integrated Neutrino Source Searches with 10 years of IceCube Data”. *Phys. Rev. Lett.* 124.5 (2020) 051103. arXiv:1910.08488 [astro-ph.HE] (cit. on pp. 101, 113, 125, 133, 136).

Aartsen, M. G., Abraham, K., et al. “All-sky Search for Time-integrated Neutrino Emission from Astrophysical Sources with 7 yr of IceCube Data”. *Astrophysical Journal* 835.2, 151 (Feb. 2017) 151. arXiv:1609.04981 [astro-ph.HE] (cit. on pp. 101, 139, 165).

Aartsen, M. G., Ackermann, M., et al. “The IceCube Neutrino Observatory: instrumentation and online systems”. *Journal of Instrumentation* 12.3 (Mar. 2017) P03012. arXiv:1612.05093 [astro-ph.IM] (cit. on p. 72).

Aartsen, M. et al. “Characteristics of the diffuse astrophysical electron and tau neutrino flux with six years of IceCube high energy cascade data” (Jan. 2020). arXiv:2001.09520 [astro-ph.HE] (cit. on pp. 11, 71, 78).

– “Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector”. *Science* 342 (2013) 1242856. arXiv:1311.5238 [astro-ph.HE] (cit. on pp. 5, 91, 141).

– “IceCube Search for High-Energy Neutrino Emission from TeV Pulsar Wind Nebulae” (Mar. 2020). arXiv:2003.12071 [astro-ph.HE] (cit. on p. 19).

Aartsen, M. et al. "IceCube-Gen2: The Window to the Extreme Universe" (Aug. 2020). arXiv:2008.04323 [astro-ph.HE] (cit. on pp. 44, 165).

– "Neutrino astronomy with the next generation IceCube Neutrino Observatory" (Nov. 2019). arXiv:1911.02561 [astro-ph.HE] (cit. on p. 165).

– "Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data". *Phys. Rev. Lett.* 113 (2014) 101101. arXiv:1405.5303 [astro-ph.HE] (cit. on p. 141).

– "Observation of the cosmic-ray shadow of the Moon with IceCube". *Phys. Rev. D* 89.10 (2014) 102004. arXiv:1305.6811 [astro-ph.HE] (cit. on p. 116).

– "Search for Sources of Astrophysical Neutrinos Using Seven Years of IceCube Cascade Events". *Astrophys. J.* 886 (2019) 12. arXiv:1907.06714 [astro-ph.HE] (cit. on pp. 12, 38, 92, C-177).

– "The IceCube Neutrino Observatory - Contributions to ICRC 2017 Part II: Properties of the Atmospheric and Astrophysical Neutrino Flux" (Oct. 2017). arXiv:1710.01191 [astro-ph.HE] (cit. on p. 141).

– "The IceCube Neutrino Observatory: Instrumentation and Online Systems". *JINST* 12.03 (2017) P03012. arXiv:1612.05093 [astro-ph.IM] (cit. on pp. 2, 33, 39–42).

– "The IceCube Realtime Alert System". *Astropart. Phys.* 92 (2017) 30–41. arXiv:1612.06028 [astro-ph.HE] (cit. on pp. 53, 141).

– "Very High-Energy Gamma-Ray Follow-Up Program Using Neutrino Triggers from IceCube". *JINST* 11.11 (2016) P11009. arXiv:1610.01814 [hep-ex] (cit. on p. 55).

Abbasi, R. et al. "An improved method for measuring muon energy using the truncated mean of dE/dx". *Nuclear Instruments and Methods in Physics Research A* 703 (Mar. 2013) 190–198. arXiv:1208.3430 [physics.data-an] (cit. on pp. 45, 46).

Abdalla, H. et al. "H.E.S.S. observations of RX J1713.7$-$3946 with improved angular and spectral resolution: Evidence for gamma-ray emission extending beyond the X-ray emitting shell". *Astron. Astrophys.* 612 (2018) A6. arXiv:1609.08671 [astro-ph.HE] (cit. on p. 18).

Abdollahi, S. et al. "*Fermi* Large Area Telescope Fourth Source Catalog". *Astrophys. J. Suppl.* 247.1 (2020) 33. arXiv:1902.10045 [astro-ph.HE] (cit. on p. 10).

Abdollahi, S. et al. "Fermi Large Area Telescope Fourth Source Catalog". *Astrophysical Journal Suppl. Ser.* 247.1, 33 (Mar. 2020) 33. arXiv:1902.10045 [astro-ph.HE] (cit. on pp. 56, 141, 142, 151).

Abraham, J. et al. "Observation of the suppression of the flux of cosmic rays above 4×10^{19}eV". *Phys. Rev. Lett.* 101 (2008) 061101. arXiv:0806.4302 [astro-ph] (cit. on p. 8).

Acero, F. et al. "Fermi Large Area Telescope Third Source Catalog". *Astrophys. J. Suppl.* 218.2 (2015) 23. arXiv:1501.02003 [astro-ph.HE] (cit. on p. 55).

Ackermann, M. et al. "Fermi-LAT Observations of the Diffuse γ-Ray Emission: Implications for Cosmic Rays and the Interstellar Medium". *Astrophysical Journal* 750.1, 3 (May 2012) 3. arXiv:1202.4039 [astro-ph.HE] (cit. on p. 10).

Ade, P. et al. "Planck 2013 results. XVI. Cosmological parameters". *Astron. Astrophys.* 571 (2014) A16. arXiv:1303.5076 [astro-ph.CO] (cit. on p. 26).

Ageron, M. et al. "ANTARES: The first undersea neutrino telescope". *Nuclear Instruments and Methods in Physics Research A* 656.1 (Nov. 2011) 11–38. arXiv:1104.1607 [astro-ph.IM] (cit. on p. 33).

Agostinelli, S. et al. "GEANT4–a simulation toolkit". *Nucl. Instrum. Meth. A* 506 (2003) 250–303 (cit. on p. 34).

Agostini, M. et al. "The Pacific Ocean Neutrino Experiment" (May 2020). arXiv:2005.09493 [astro-ph.HE] (cit. on pp. 33, 44, 165).

Ahmad, Q. R. et al. "Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory". *Phys. Rev. Lett.* 89 (1 June 2002) 011301 (cit. on p. 25).

Ahn, E.-J. et al. "Cosmic ray interaction event generator SIBYLL 2.1". *Phys. Rev. D* 80 (2009) 094003. arXiv:0906.4113 [hep-ph] (cit. on p. 90).

Ahrens, J. et al. "Muon track reconstruction and data selection techniques in AMANDA". *Nucl. Instrum. Meth. A* 524 (2004) 169–194. arXiv:astro-ph/0407044 (cit. on p. 47).

Ahrens, J. et al. "Search for Extraterrestrial Point Sources of Neutrinos with AMANDA-II". 92.7, 071102 (Feb. 2004) 071102. arXiv:astro-ph/0309585 [astro-ph] (cit. on p. 101).

Ajello, M. et al. "The Fourth Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope" (May 2019). arXiv:1905.10771 [astro-ph.HE] (cit. on pp. 141, 148).

– "The Origin of the Extragalactic Gamma-Ray Background and Implications for Dark-Matter Annihilation". *Astrophys. J. Lett.* 800.2 (2015) L27. arXiv:1501.05301 [astro-ph.HE] (cit. on pp. 10, 20).

Aker, M. et al. "Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN". *Phys. Rev. Lett.* 123.22 (2019) 221802. arXiv:1909.06048 [hep-ex] (cit. on p. 26).

Alves Batista, R. et al. "Open Questions in Cosmic-Ray Research at Ultrahigh Energies". *Front. Astron. Space Sci.* 6 (2019) 23. arXiv:1903.06714 [astro-ph.HE] (cit. on p. 7).

Apel, W. D. et al. "Kneelike Structure in the Spectrum of the Heavy Component of Cosmic Rays Observed with KASCADE-Grande". 107.17, 171104 (Oct. 2011) 171104. arXiv:1107.5885 [astro-ph.HE] (cit. on p. 7).

Atwood, W. B. et al. "The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission". *Astrophysical Journal* 697.2 (June 2009) 1071–1102. arXiv:0902.1089 [astro-ph.IM] (cit. on p. 30).

Avrorin, A. et al. "Baikal-GVD: status and prospects". *EPJ Web Conf.* 191 (2018). Ed. by Volkova, V. et al. 01006. arXiv:1808.10353 [astro-ph.IM] (cit. on pp. 33, 165).

Bahcall, J. N., Bahcall, N. A., and Shaviv, G. "Present Status of the Theoretical Predictions for the ^{37}Cl Solar-Neutrino Experiment". *Phys. Rev. Lett.* 20 (21 May 1968) 1209–1212 (cit. on p. 25).

Bechtol, K. et al. "Evidence against star-forming galaxies as the dominant source of IceCube neutrinos". *Astrophys. J.* 836.1 (2017) 47. arXiv:1511.00688 [astro-ph.HE] (cit. on pp. 19, 20).

Bell, A. R. "The acceleration of cosmic rays in shock fronts – I". *Monthly Notices of the Royal Astronomical Society* 182.2 (Feb. 1978) 147–156. ISSN: 0035-8711. eprint: https://academic.oup.com/mnras/article-pdf/182/2/147/3710138/mnras182-0147.pdf (cit. on p. 14).

Berezinsky, V. and Smirnov, A. Y. "Cosmic neutrinos of ultra-high energies and detection possibility". *Astrophysics and Space Science* 32.2 (1975) 461–482 (cit. on p. 16).

Blasi, P. "The Origin of Galactic Cosmic Rays". *Astron. Astrophys. Rev.* 21 (2013) 70. arXiv:1311.7346 [astro-ph.HE] (cit. on p. 6).

Blaufuss, E. et al. "The Next Generation of IceCube Realtime Neutrino Alerts". *PoS* ICRC2019 (2020) 1021. arXiv:1908.04884 [astro-ph.HE] (cit. on pp. 71, 89).

Bradascio, F. and Glüsenkamp, T. "Improving the muon track reconstruction of IceCube and IceCube-Gen2". *EPJ Web Conf.* 207 (2019). Ed. by Spiering, C. 05002. arXiv:1905.09612 [astro-ph.IM] (cit. on p. 47).

Braun, J. et al. "Methods for point source analysis in high energy neutrino telescopes". *Astroparticle Physics* 29.4 (May 2008) 299–305. arXiv:0801.1604 [astro-ph] (cit. on pp. 101, 107).

Burrows, D. N. et al. "The Swift X-ray Telescope". *Space Sci. Rev.* 120 (2005) 165. arXiv:astro-ph/0508071 (cit. on p. 31).

Casella, G. and Berger, R. *Statistical Inference*. Duxbury advanced series in statistics and decision sciences. Thomson Learning, 2002. ISBN: 9780534243128 (cit. on pp. 101, 102, 104, 135).

Cerruti, M. et al. "Leptohadronic single-zone models for the electromagnetic and neutrino emission of TXS 0506+056". *Mon. Not. Roy. Astron. Soc.* 483.1 (2019) L12–L16. arXiv:1807.04335 [astro-ph.HE] (cit. on p. 63).

Chang, Y.-L., Arsioli, B., et al. "The 3HSP catalogue of extreme and high-synchrotron peaked blazars". *Astron. Astrophys.* 632 (2019) A77. arXiv:1909.08279 [astro-ph.HE] (cit. on pp. 56, 141, 148, 154).

Chang, Y.-L., Brandt, C., and Giommi, P. "The Open Universe VOU-Blazars tool". *Astron. Comput.* 30 (2020) 100350. arXiv:1909.11455 [astro-ph.HE] (cit. on pp. 57, 142, 147).

Chirkin, D. and Rongen, M. "Light diffusion in birefringent polycrystals and the IceCube ice anisotropy". *PoS* ICRC2019 (2020) 854. arXiv:1908.07608 [astro-ph.HE] (cit. on p. 45).

Chollet, F. et al. *Keras*. https://keras.io. 2015 (cit. on p. 77).

Choma, N. et al. "Graph Neural Networks for IceCube Signal Classification". *arXiv e-prints*, arXiv:1809.06166 (Sept. 2018) arXiv:1809.06166. arXiv:1809.06166 [cs.LG] (cit. on p. 65).

Cooper-Sarkar, A., Mertsch, P., and Sarkar, S. "The high energy neutrino cross-section in the Standard Model and its uncertainty". *JHEP* 08 (2011) 042. arXiv:1106.3723 [hep-ph] (cit. on p. 36).

Cowan, G. *Statistical Data Analysis*. Oxford science publications. Clarendon Press, 1998. ISBN: 9780198501558 (cit. on p. 101).

Cowan, G. et al. "Asymptotic formulae for likelihood-based tests of new physics". *Eur. Phys. J. C* 71 (2011). [Erratum: Eur.Phys.J.C 73, 2501 (2013)] 1554. arXiv:1007.1727 [physics.data-an] (cit. on pp. 49, 145, 151).

Davis, R., Harmer, D. S., and Hoffman, K. C. "Search for Neutrinos from the Sun". *Phys. Rev. Lett.* 20 (21 May 1968) 1205–1209 (cit. on p. 25).

De Angelis, A., Galanti, G., and Roncadelli, M. "Transparency of the Universe to gamma rays". *Mon. Not. Roy. Astron. Soc.* 432 (2013) 3245–3249. arXiv:1302.6460 [astro-ph.HE] (cit. on p. 8).

Dembinski, H. P. et al. "Data-driven model of the cosmic-ray flux and mass composition from 10 GeV to 10^{11} GeV". *PoS* ICRC2017 (2018) 533. arXiv:1711.11432 [astro-ph.HE] (cit. on p. 115).

Desai, V. et al. "PAH Emission from Ultraluminous Infrared Galaxies". *Astrophys. J.* 669 (2007) 810–820. arXiv:0707.4190 [astro-ph] (cit. on p. 19).

Dolag, K. et al. "Mapping deflections of ultrahigh energy cosmic rays in constrained simulations of extragalactic magnetic fields". *JETP Lett.* 79 (2004) 583–587. arXiv:astro-ph/0310902 (cit. on p. 8).

Esteban, I. et al. "Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23}, δ_{CP}, and the mass ordering". *JHEP* 01 (2019) 106. arXiv:1811.05487 [hep-ph] (cit. on pp. 26, 27).

Fanaroff, B. L. and Riley, J. M. "The morphology of extragalactic radio sources of high and low luminosity". *Monthly Notices of the Royal Astronomical Society* 167 (May 1974) 31P–36P (cit. on p. 22).

Fedynitch, A. et al. "Calculation of conventional and prompt lepton fluxes at very high energy". *EPJ Web Conf.* 99 (2015). Ed. by Berge, D. et al. 08001. arXiv:1503.00544 [hep-ph] (cit. on pp. 28, 30).

Fermi, E. "On the Origin of the Cosmic Radiation". *Phys. Rev.* 75 (1949) 1169–1174 (cit. on p. 18).

Formaggio, J. and Zeller, G. "From eV to EeV: Neutrino Cross Sections Across Energy Scales". *Rev. Mod. Phys.* 84 (2012) 1307–1341. arXiv:1305.7513 [hep-ex] (cit. on pp. 35, 36).

Gaggero, D., Grasso, D., et al. "The gamma-ray and neutrino sky: A consistent picture of Fermi-LAT, Milagro, and IceCube results". *Astrophys. J. Lett.* 815.2 (2015) L25. arXiv:1504.00227 [astro-ph.HE] (cit. on p. 10).

Gaggero, D., Urbano, A., et al. "Gamma-ray sky points to radial gradients in cosmic-ray transport". *Phys. Rev. D* 91.8 (2015) 083012. arXiv:1411.7623 [astro-ph.HE] (cit. on pp. 9, 10).

Gaisser, T. "Neutrino astronomy: Physics goals, detector parameters". July 1997. arXiv:astro-ph/9707283 (cit. on p. 37).

Gaisser, T. K., Engel, R., and Resconi, E. *Cosmic Rays and Particle Physics*. 2nd ed. Cambridge University Press, 2016 (cit. on pp. 18, 26, 28, 29).

Gaisser, T. K., Jero, K., et al. "Generalized self-veto probability for atmospheric neutrinos". *Phys. Rev. D* 90.2 (2014) 023009. arXiv:1405.0525 [astro-ph.HE] (cit. on pp. 43, 90).

Gaisser, T. K., Stanev, T., and Tilav, S. "Cosmic Ray Energy Spectrum from Measurements of Air Showers". *Front. Phys. (Beijing)* 8 (2013) 748–758. arXiv:1303.3565 [astro-ph.HE] (cit. on pp. 7, 90).

Gazizov, A. and Kowalski, M. "ANIS: High energy neutrino generator for neutrino telescopes". *Computer Physics Communications* 172.3 (2005) 203–213. ISSN: 0010-4655 (cit. on p. 78).

Gentle, J. E. *Computational Statistics*. Springer New York, 2009 (cit. on p. 134).

Ghisellini, G. et al. "The Fermi blazar sequence". *Mon. Not. Roy. Astron. Soc.* 469.1 (2017) 255–266. arXiv:1702.02571 [astro-ph.HE] (cit. on pp. 24, 157).

Ginzburg, V. L. and Syrovatskii, S. I. *The Origin of Cosmic Rays*. 1964 (cit. on p. 6).

Giommi, P., Chang, Y. L., et al. "The Open Universe survey of Swift-XRT GRB fields: a complete sample of HBL blazars". *arXiv e-prints*, arXiv:2003.05153 (Mar. 2020) arXiv:2003.05153. arXiv:2003.05153 [astro-ph.HE] (cit. on p. 150).

Giommi, P., Glauch, T., et al. "Dissecting the regions around IceCube high-energy neutrinos: growing evidence for the blazar connection" (Jan. 2020). arXiv:2001.09355 [astro-ph.HE] (cit. on pp. 3, 140, 143, 144, 147, 158, 160, 164).

Giommi, P., Padovani, P., Oikonomou, F., et al. "3HSP J095507.9+355101: a flaring extreme blazar coincident in space and time with IceCube-200107A" (Mar. 2020). arXiv:2003.06405 [astro-ph.HE] (cit. on pp. 3, 56, 92, 140, 155, 156, 165).

Giommi, P., Padovani, P., and Polenta, G. "A simplified view of blazars: the gamma-ray case". *Mon. Not. Roy. Astron. Soc.* 431 (2013) 1914. arXiv:1302.4331 [astro-ph.HE] (cit. on pp. 24, 25).

Glauch, T., Padovani, P., et al. "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source". *EPJ Web Conf.* 207 (2019). Ed. by Spiering, C. 02003 (cit. on pp. 2, 60).

Glauch, T. and Turcati, A. "Search for weak neutrino point sources using angular auto-correlation analyses in IceCube". *PoS* ICRC2017 (2018) 1014 (cit. on pp. 139, 165).

Glorot, X. and Bengio, Y. "Understanding the difficulty of training deep feedforward neural networks". *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. 2010 p. 249 (cit. on p. 76).

Gorski, K. et al. "HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere". *Astrophys. J.* 622 (2005) 759–771. arXiv:astro-ph/0409513 (cit. on p. 134).

Gruppioni, C. et al. "The Herschel PEP/HerMES Luminosity Function. I: Probing the Evolution of PACS selected Galaxies to z~4". *Mon. Not. Roy. Astron. Soc.* 432 (2013) 23. arXiv:1302.5209 [astro-ph.CO] (cit. on p. 19).

Halzen, F. and Zas, E. "Neutrino Fluxes from Active Galaxies: A Model-Independent Estimate". *Astrophysical Journal* 488.2 (Oct. 1997) 669–674. arXiv:astro-ph/9702193 [astro-ph] (cit. on p. 20).

Harari, D., Mollerach, S., and Roulet, E. "On the ultrahigh energy cosmic ray horizon". *JCAP* 11 (2006) 012. arXiv:astro-ph/0609294 (cit. on p. 8).

Haverkorn, M. "Magnetic Fields in the Milky Way". *Magnetic Fields in Diffuse Media*. Ed. by Lazarian, A., de Gouveia Dal Pino, E. M., and Melioli, C. Vol. 407. Astrophysics and Space Science Library. 2015 p. 483 (cit. on p. 6).

Haykin, S. *Neural Networks: A Comprehensive Foundation*. 1st. USA: Prentice Hall PTR, 1994. ISBN: 0023527617 (cit. on p. 67).

He, K. and Sun, J. "Convolutional neural networks at constrained time cost". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015 p. 5353 (cit. on p. 76).

He, K., Zhang, X., et al. "Deep residual learning for image recognition". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016 p. 770 (cit. on pp. 65, 76).

Heck, D. et al. "CORSIKA: A Monte Carlo code to simulate extensive air showers" (Feb. 1998) (cit. on p. 90).

Hess, V. F. "Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten". *Phys. Z.* 13 (1912) 1084–1091 (cit. on pp. 1, 5).

Honda, M. et al. "Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data". *Phys. Rev.* D75 (2007) 043006. arXiv:astro-ph/0611418 [astro-ph] (cit. on pp. 78, 90).

Huber, M. "Multi-Messenger correlation study of Fermi-LAT blazars and high-energy neutrinos observed in IceCube". PhD thesis. Technical University of Munich, 2020 (cit. on pp. 139, 147, 165).

– "Searches for steady neutrino emission from 3FHL blazars using eight years of IceCube data from the Northern hemisphere". *PoS* ICRC2019 (2020) 916. arXiv:1908.08458 [astro-ph.HE] (cit. on p. 139).

Huennefeld, M. "Deep Learning in Physics exemplified by the Reconstruction of Muon-Neutrino Events in IceCube". *PoS* ICRC2017 (2018) 1057 (cit. on pp. 65, 73, 109).

– "Reconstruction Techniques in IceCube using Convolutional and Generative Neural Networks". *EPJ Web of Conferences* 207 (Jan. 2019) 05005 (cit. on pp. 48, 73).

IceCube Collaboration, Aartsen, M. G., Ackermann, M., Adams, J., Aguilar, J. A., Ahlers, M., Ahrens, M., Al Samarai, I., et al. "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A". *Science* 361.6398, eaat1378 (July 2018) eaat1378. arXiv:1807.08816 [astro-ph.HE] (cit. on pp. 2, 53–56, 60, 89, 141, 163).

IceCube Collaboration, Aartsen, M. G., Ackermann, M., Adams, J., Aguilar, J. A., Ahlers, M., Ahrens, M., Samarai, I. A., et al. "Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert". *Science* 361.6398 (July 2018) 147–151. arXiv:1807.08794 [astro-ph.HE] (cit. on pp. 2, 53, 57, 60, 163).

Inoue, Y., Khangulyan, D., and Doi, A. "On the Origin of High-energy Neutrinos from NGC 1068: The Role of Nonthermal Coronal Activity". *Astrophys. J. Lett.* 891.2 (2020) L33. arXiv:1909.02239 [astro-ph.HE] (cit. on p. 136).

Ioffe, S. and Szegedy, C. "Batch normalization: Accelerating deep network training by reducing internal covariate shift". *arXiv preprint arXiv:1502.03167* (2015) (cit. on p. 76).

Jackson, J. D. *Classical Electrodynamics, 3rd Edition*. 1998 (cit. on p. 33).

Keenan, M. et al. "The Relativistic Jet Dichotomy and the End of the Blazar Sequence" (July 2020). arXiv:2007.12661 [astro-ph.GA] (cit. on p. 24).

Keivani, A. et al. "A Multimessenger Picture of the Flaring Blazar TXS 0506+056: implications for High-Energy Neutrino Emission and Cosmic Ray Acceleration". *Astrophys. J.* 864.1 (2018) 84. arXiv:1807.04537 [astro-ph.HE] (cit. on p. 63).

Kelner, S. and Aharonian, F. "Energy spectra of gamma-rays, electrons and neutrinos produced at interactions of relativistic protons with low energy radiation". *Phys. Rev. D* 78 (2008). [Erratum: Phys.Rev.D 82, 099901 (2010)] 034013. arXiv:0803.0688 [astro-ph] (cit. on p. 13).

Kelner, S., Aharonian, F. A., and Bugayov, V. "Energy spectra of gamma-rays, electrons and neutrinos produced at proton-proton interactions in the very high energy regime". *Phys. Rev. D* 74 (2006). [Erratum: Phys.Rev.D 79, 039901 (2009)] 034018. arXiv:astro-ph/0606058 (cit. on p. 14).

Kingma, D. P. and Ba, J. "Adam: A Method for Stochastic Optimization". *arXiv e-prints*, arXiv:1412.6980 (Dec. 2014) arXiv:1412.6980. arXiv:1412.6980 [cs.LG] (cit. on pp. 77, 109, 149).

Kopper, C. "Performance Studies for theKM3NeT Neutrino Telescope". PhD thesis. Friedrich-Alexander-Universit Erlangen-Nuernberg, 2010 (cit. on p. 37).

Koyama, K. et al. "Evidence for shock acceleration of high-energy electrons in the supernova remnant SN1006". *Nature* 378.6554 (Nov. 1995) 255–258 (cit. on p. 18).

Krizhevsky, A., Sutskever, I., and Hinton, G. E. "Imagenet classification with deep convolutional neural networks". *Advances in neural information processing systems*. 2012 p. 1097 (cit. on p. 76).

Kronmueller, M. and Glauch, T. "Application of Deep Neural Networks to Event Type Classification in IceCube". *36th International Cosmic Ray Conference (ICRC2019)*. Vol. 36. International Cosmic Ray Conference. July 2019 p. 937. arXiv:1908.08763 [astro-ph.IM] (cit. on pp. 2, 66, 163).

Kronmüller, M. "Application of Deep Neural Networks on Event Type Classification in IceCube". MA thesis. Technical University of Munich, 2018 (cit. on p. 66).

LeCun, Y., Bengio, Y., and Hinton, G. "Deep learning". *nature* 521.7553 (2015) 436–444 (cit. on p. 65).

Lin, M., Chen, Q., and Yan, S. "Network In Network". *arXiv e-prints*, arXiv:1312.4400 (Dec. 2013) arXiv:1312.4400. arXiv:1312.4400 [cs.NE] (cit. on p. 77).

Litjens, G. et al. "A survey on deep learning in medical image analysis". *Medical image analysis* 42 (2017) 60–88 (cit. on p. 65).

Liu, R.-Y. et al. "Hadronuclear interpretation of a high-energy neutrino event coincident with a blazar flare". *Phys. Rev. D* 99.6 (2019) 063008. arXiv:1807.05113 [astro-ph.HE] (cit. on p. 63).

Loeb, A. and Waxman, E. "The Cumulative background of high energy neutrinos from starburst galaxies". *JCAP* 05 (2006) 003. arXiv:astro-ph/0601695 (cit. on p. 19).

Lott, B. et al. "An adaptive-binning method for generating constant-uncertainty/constant-significance light curves with Fermi-LAT data". *Astronomy & Astrophysics* 544, A6 (Aug. 2012) A6. arXiv:1201.4851 [astro-ph.HE] (cit. on pp. 60, 151).

Maki, Z., Nakagawa, M., and Sakata, S. "Remarks on the unified model of elementary particles". *Prog. Theor. Phys.* 28 (1962) 870–880 (cit. on p. 25).

Mannheim, K. "High-energy neutrinos from extragalactic jets". *Astroparticle Physics* 3.3 (May 1995) 295–302 (cit. on p. 20).

Margiotta, A. "The KM3NeT deep-sea neutrino telescope". *Nucl. Instrum. Meth. A* 766 (2014). Ed. by Sumiyoshi, T. et al. 83–87. arXiv:1408.1392 [astro-ph.IM] (cit. on pp. 33, 165).

Massaro, E. et al. "The 5th edition of the Roma-BZCAT. A short presentation". *Astrophys. Space Sci.* 357.1 (2015) 75. arXiv:1502.07755 [astro-ph.HE] (cit. on p. 140).

Morrison, P. "On gamma-ray astronomy". *Il Nuovo Cimento* 7.6 (Mar. 1958) 858–865 (cit. on p. 5).

Mücke, A. et al. "BL Lac objects in the synchrotron proton blazar model". *Astroparticle Physics* 18.6 (Mar. 2003) 593–613. arXiv:astro-ph/0206164 [astro-ph] (cit. on p. 20).

Murase, K., Ahlers, M., and Lacki, B. C. "Testing the Hadronuclear Origin of PeV Neutrinos Observed with IceCube". *Phys. Rev. D* 88.12 (2013) 121301. arXiv:1306.3417 [astro-ph.HE] (cit. on p. 134).

Murase, K., Guetta, D., and Ahlers, M. "Hidden Cosmic-Ray Accelerators as an Origin of TeV-PeV Cosmic Neutrinos". *Phys. Rev. Lett.* 116.7 (2016) 071101. arXiv:1509.00805 [astro-ph.HE] (cit. on p. 17).

Murase, K., Kimura, S. S., and Meszaros, P. "Hidden Cores of Active Galactic Nuclei as the Origin of Medium-Energy Neutrinos: Critical Tests with the MeV Gamma-Ray Connection". *Phys. Rev. Lett.* 125.1 (2020) 011101. arXiv:1904.04226 [astro-ph.HE] (cit. on p. 136).

Murase, K., Oikonomou, F., and Petropoulou, M. "Blazar Flares as an Origin of High-Energy Cosmic Neutrinos?" *Astrophys. J.* 865.2 (2018) 124. arXiv:1807.04748 [astro-ph.HE] (cit. on p. 156).

Neunhoffer, T. "Estimating the angular resolution of tracks in neutrino telescopes based on a likelihood analysis". *Astropart. Phys.* 25 (2006) 220–225. arXiv:astro-ph/0403367 (cit. on p. 49).

Neyman, J. and Pearson, E. S. "On the Problem of the Most Efficient Tests of Statistical Hypotheses". *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 231 (1933) 289–337. ISSN: 02643952 (cit. on p. 103).

Padovani, P. et al. "Active galactic nuclei: what's in a name?" *Astron. Astrophys. Rev.* 25.1 (2017) 2. arXiv:1707.07134 [astro-ph.GA] (cit. on pp. 20–22).

Padovani, P., Giommi, P., et al. "Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source". *Mon. Not. Roy. Astron. Soc.* 480.1 (2018) 192–203. arXiv:1807.04461 [astro-ph.HE] (cit. on pp. 2, 53, 58, 61, 140, 163).

Padovani, P., Oikonomou, F., et al. "TXS 0506+056, the first cosmic neutrino source, is not a BL Lac". *Mon. Not. Roy. Astron. Soc.* 484.1 (2019) L104–L108. arXiv:1901.06998 [astro-ph.HE] (cit. on pp. 24, 62).

Padovani, P. and Resconi, E. "Are both BL Lacs and pulsar wind nebulae the astrophysical counterparts of IceCube neutrino events?" *Monthly Notices of the Royal Astronomical Society* 443.1 (Sept. 2014) 474–484. arXiv:1406.0376 [astro-ph.HE] (cit. on p. 20).

Padovani, P. "Active Galactic Nuclei at All Wavelengths and from All Angles". *Frontiers in Astronomy and Space Sciences* 4 (2017) 35. ISSN: 2296-987X (cit. on pp. 22, 23).

Padovani, P. and Giommi, P. "The Connection between X-Ray– and Radio-selected BL Lacertae Objects". *Astrophysical Journal* 444 (May 1995) 567. arXiv:astro-ph/9412073 [astro-ph] (cit. on pp. 24, 58).

Padovani, P., Perlman, E. S., et al. "What types of jets does nature make: A New population of radio quasars". *Astrophys. J.* 588 (2003) 128–142. arXiv:astro-ph/0301227 (cit. on p. 58).

Paiano, S., Falomo, R., Padovani, P., et al. "The redshift and the host galaxy of the neutrino candidate 4FGL J0955.1+3551 (3HSP J095507.9+355101)". *Mon. Not. Roy. Astron. Soc.* 495.1 (2020) L108–L111. arXiv:2003.03634 [astro-ph.HE] (cit. on pp. 155, 156).

Paiano, S., Falomo, R., Treves, A., et al. "The redshift of the BL Lac object TXS 0506+056". *Astrophys. J. Lett.* 854.2 (2018) L32. arXiv:1802.01939 [astro-ph.GA] (cit. on pp. 58, 62).

Paliya, V. S. et al. "Multi-Frequency Observations of the Candidate Neutrino Emitting Blazar BZB J0955+3551" (Mar. 2020). arXiv:2003.06012 [astro-ph.HE] (cit. on p. 92).

Palladino, A. et al. "IceCube Neutrinos from Hadronically Powered Gamma-Ray Galaxies". *JCAP* 09 (2019) 004. arXiv:1812.04685 [astro-ph.HE] (cit. on p. 20).

Petropoulou, M., Dimitrakoudis, S., et al. "Photohadronic origin of γ-ray BL Lac emission: implications for IceCube neutrinos". *Mon. Not. Roy. Astron. Soc.* 448.3 (2015) 2412–2429. arXiv:1501.07115 [astro-ph.HE] (cit. on p. 61).

Petropoulou, M., Oikonomou, F., et al. "Comprehensive Multimessenger Modeling of the Extreme Blazar 3HSP J095507.9+355101 and Predictions for IceCube" (May 2020). arXiv:2005.07218 [astro-ph.HE] (cit. on pp. 23, 156).

Poluektov, A. "Kernel density estimation of a multidimensional efficiency profile". *Journal of Instrumentation* 10.2, P02011 (Feb. 2015) P02011. arXiv:1411.5528 [physics.data-an] (cit. on pp. 119–121).

Ptuskin, V., Zirakashvili, V., and Seo, E.-S. "Spectrum of Galactic Cosmic Rays Accelerated in Supernova Remnants". *Astrophysical Journal* 718.1 (July 2010) 31–36. arXiv:1006.0034 [astro-ph.CO] (cit. on p. 6).

Punsly, B. and Zhang, S. "Calibrating emission lines as quasar bolometers". *Monthly Notices of the Royal Astronomical Society* 412.1 (Mar. 2011) L123–L127. arXiv:1101.4833 [astro-ph.CO] (cit. on p. 156).

Rädel, L. "Measurement of High-Energy Muon Neutrinos with the IceCube Neutrino Observatory". PhD thesis. RWTH Aachen University, 2017 (cit. on pp. 52, 114, 142).

Radel, L. and Wiebusch, C. "Calculation of the Cherenkov light yield from electromagnetic cascades in ice with Geant4". *Astropart. Phys.* 44 (2013) 102–113. arXiv:1210.5140 [astro-ph.IM] (cit. on p. 34).

Reimer, A., Boettcher, M., and Buson, S. "Cascading Constraints from Neutrino Emitting Blazars: The case of TXS 0506+056" (Dec. 2018). arXiv:1812.05654 [astro-ph.HE] (cit. on p. 62).

Reynolds, S. P. "Supernova remnants at high energy." *Ann. Rev. Astron. Asrophys.* 46 (Sept. 2008) 89–126 (cit. on p. 18).

Richards, J. L. et al. "Blazars in the Fermi Era: The OVRO 40 m Telescope Monitoring Program". *Astrophysical Journal Suppl. Ser.* 194.2, 29 (June 2011) 29. arXiv:1011.3111 [astro-ph.CO] (cit. on p. 31).

Riehn, F. et al. "The hadronic interaction model SIBYLL 2.3c and Feynman scaling". *PoS* ICRC2017 (2018) 301. arXiv:1709.07227 [hep-ph] (cit. on p. 115).

Rodrigues, X. et al. "Leptohadronic Blazar Models Applied to the 2014–2015 Flare of TXS 0506+056". *Astrophys. J. Lett.* 874.2 (2019) L29. arXiv:1812.05939 [astro-ph.HE] (cit. on p. 63).

Rongen, M. "Calibration of the IceCube Neutrino Observatory". Other thesis. 2019. arXiv:1911.02016 [astro-ph.IM] (cit. on p. 45).

Rosenblatt, F. "The perceptron: a probabilistic model for information storage and organization in the brain." *Psychological review* 65.6 (1958) 386 (cit. on p. 66).

Saikia, P., Körding, E., and Falcke, H. "Lorentz factor distribution of blazars from the optical Fundamental plane of black hole activity". *Mon. Not. Roy. Astron. Soc.* 461.1 (2016) 297–303. arXiv:1606.06147 [astro-ph.HE] (cit. on p. 156).

Schneider, A. "Characterization of the Astrophysical Diffuse Neutrino Flux with IceCube High-Energy Starting Events". *PoS* ICRC2019 (2020) 1004. arXiv:1907.11266 [astro-ph.HE] (cit. on p. 90).

Schonert, S. et al. "Vetoing atmospheric neutrinos in a high energy neutrino telescope". *Phys. Rev. D* 79 (2009) 043009. arXiv:0812.4308 [astro-ph] (cit. on p. 43).

Schönert, S. et al. "Vetoing atmospheric neutrinos in a high energy neutrino telescope". *Phys. Rev. D* 79 (4 Feb. 2009) 043009 (cit. on p. 90).

Shannon, C. E. and McCarthy, J. *Automata Studies. (AM-34), Volume 34*. Princeton: Princeton University Press, 2016. ISBN: 978-1-4008-8261-8 (cit. on p. 66).

Silva, M. and Mancina, S. "Measurement of the Diffuse Muon Neutrino Flux using Starting Track Events in IceCube". *PoS* ICRC2019 (2020) 1010. arXiv:1908.06586 [astro-ph.HE] (cit. on p. 99).

Silver, D. et al. "Mastering the game of Go with deep neural networks and tree search". *Nature* 529 (2016) 484–503 (cit. on p. 65).

Simonyan, K. and Zisserman, A. "Very deep convolutional networks for large-scale image recognition". *arXiv preprint arXiv:1409.1556* (2014) (cit. on p. 76).

"South Pole glacial climate reconstruction from multi-borehole laser particulate stratigraphy". *J. Glaciol.* 59.218 (2013) 1117–1128 (cit. on p. 44).

Stachurska, J. "First Double Cascade Tau Neutrino Candidates in IceCube and a New Measurement of the Flavor Composition". *PoS* ICRC2019 (2020) 1015. arXiv:1908.05506 [astro-ph.HE] (cit. on p. 39).

Stein, R. "IceCube-200107A: IceCube observation of a high-energy neutrino candidate event". *GCN* 26655 (Jan. 2020) 1 (cit. on p. 92).

Stein, R. et al. "A high-energy neutrino coincident with a tidal disruption event" (May 2020). arXiv:2005.05340 [astro-ph.HE] (cit. on p. 56).

Stettner, J. "Measurement of the Diffuse Astrophysical Muon-Neutrino Spectrum with Ten Years of IceCube Data". *PoS* ICRC2019 (2020) 1017. arXiv:1908.09551 [astro-ph.HE] (cit. on pp. 11, 30, 83, 86, 115).

Steuer, A. K. "Cascade type identification in IceCube and an application in a search for new physics". PhD thesis. Mainz U., 2018 (cit. on p. 71).

Strotjohann, N. L., Kowalski, M., and Franckowiak, A. "Eddington bias for cosmic neutrino sources". *Astron. Astrophys.* 622 (2019) L9. arXiv:1809.06865 [astro-ph.HE] (cit. on p. 157).

Szegedy, C. et al. "Inception-v4, inception-resnet and the impact of residual connections on learning". *Thirty-first AAAI conference on artificial intelligence*. 2017 (cit. on p. 76).

Tamborra, I., Ando, S., and Murase, K. "Star-forming galaxies as the origin of diffuse high-energy backgrounds: Gamma-ray and neutrino connections, and implications for starburst history". *JCAP* 09 (2014) 043. arXiv:1404.1189 [astro-ph.HE] (cit. on p. 19).

Tanabashi, M. et al. "Review of Particle Physics". *Phys. Rev. D* 98 (3 Aug. 2018) 030001 (cit. on pp. 12, 28, 35, 37–39, 45).

Tosi, D. and Wendt, C. "Calibrating photon detection efficiency in IceCube". *PoS* TIPP2014 (2014) 157. arXiv:1502.03102 [physics.ins-det] (cit. on p. 72).

Urbach, F. "The Long-Wavelength Edge of Photographic Sensitivity and of the Electronic Absorption of Solids". *Phys. Rev.* 92 (5 Dec. 1953) 1324–1324 (cit. on p. 34).

Urry, C. M. and Padovani, P. "Unified schemes for radio-loud active galactic nuclei". *Publications of the Astronomical Society of the Pacific* 107.715 (1995) 803 (cit. on pp. 2, 21).

Vanden Berk, D. E. et al. "Composite Quasar Spectra from the Sloan Digital Sky Survey". *Astronomical Journal* 122.2 (Aug. 2001) 549–564. arXiv:astro-ph/0105231 [astro-ph] (cit. on p. 22).

Villante, F. and Vissani, F. "How precisely neutrino emission from supernova remnants can be constrained by gamma ray observations?" *Phys. Rev. D* 78 (2008) 103007. arXiv:0807.4151 [astro-ph] (cit. on p. 19).

Waxman, E. "Cosmological origin for cosmic rays above 10**19-eV". *Astrophys. J. Lett.* 452 (1995) L1–L4. arXiv:astro-ph/9508037 (cit. on p. 15).

Waxman, E. and Bahcall, J. N. "High-energy neutrinos from astrophysical sources: An Upper bound". *Phys. Rev. D* 59 (1999) 023002. arXiv:hep-ph/9807282 (cit. on pp. 14, 15, 17, 37).

Whitehorn, N., Santen, J. van, and Lafebre, S. "Penalized Splines for Smooth Representation of High-dimensional Monte Carlo Datasets". *Comput. Phys. Commun.* 184 (2013) 2214–2220. arXiv:1301.2184 [physics.data-an] (cit. on pp. 48, 122).

Wolf, M. "SkyLLH – A generalized Python-based tool for log-likelihood analyses in multi-messenger astronomy". *PoS* ICRC2019 (2020) 1035. arXiv:1908.05181 [astro-ph.IM] (cit. on p. D-182).