Optical Follow-Up of High-Energy Neutrinos

Dissertation
zur Erlangung des akademischen Grades

– doctor rerum naturalium (Dr. rer. nat.) –

im Fach Physik

eingereicht an der
Mathematisch-Naturwissenschaftlichen Fakultät der
Humboldt-Universität zu Berlin
von

M.Sc. Simeon Reusch

Präsidentin der Humboldt-Universität zu Berlin:
Prof. Dr. Julia von Blumenthal

Dekanin der Mathematisch-Naturwissenschaftlichen Fakultät:
Prof. Dr. Caren Tischendorf

Gutachterinnen:
Prof. Dr. Katie Auchettl
Prof. Dr. Anna Franckowiak
Prof. Dr. Andreja Gomboc

Verteidigt am 22. Juli 2024

This thesis was typeset with the help of `KOMAScript` and LaTeX using the `kaobook` class

The code created to typeset this thesis and generate some of its figures can be accessed at github.com/simeonreusch/thesis

Contents

List of Figures

List of Tables

Abstract

This thesis is concerned with the origin of high-energy neutrinos detected by the IceCube Observatory at the South Pole since a decade. It summarizes the young field of neutrino astronomy and details the systematic high-energy neutrino follow-up program we have been conducting with the Zwicky Transient Facility (ZTF), an optical survey telescope, for the past four years.

One major finding is the establishment of the astrophysical transient *AT2019fdr* as counterpart to a high-energy neutrino. This transient is studied in detail, aided by the collection and reduction of data across the electromagnetic spectrum. Based on these observations and modeling of the light curve I conclude that this event constitutes a Tidal Disruption Event (TDE), albeit an unusual one. The chance coincidence of such an association is 0.034 % when including another previous association. Furthermore, I discuss the infrared dust echo from this transient in the context of two further possible associations of candidate TDEs with high-energy neutrinos, which also display such a dust echo.

This study is appended by the creation of the ZTF nuclear sample, the first systematic sample of transient events found near the cores of their host galaxies within the ZTF survey. One goal of this sample is to enlarge the number of TDEs found so far. This is achieved by the development of a machine-learning based photometric typing algorithm. That classifier is trained on a survey of bright ZTF transients, including augmentation of that sample to account for the fainter nature of the nuclear sample. When applying the classifier to the nuclear sample, an additional 27 new candidate TDEs are found. Furthermore, the identification of candidate TDEs via their infrared dust-echo signal is also successful, resulting in 16 previously unpublished TDE candidates.

Diese Dissertation untersucht den Ursprung der hochenergetischen Neutrinos, welche das IceCube-Observatorium am Südpol seit einer Dekade detektiert. Sie stellt das junge Feld der Neutrino-Astronomie vor und beschreibt das systematische Follow-Up-Programm für hochenergetische Neutrinos, das wir seit vier Jahren mit der Zwicky Transient Facility (ZTF) durchführen, einem optischen Teleskop.

Ein wesentliches Resultat ist die Identifikation des astrophysikalischen Objekts *AT2019fdr* als mögliche Quelle eines hochenergetischen Neurinos. Dieses Objekt wird im Detail untersucht; so werden Daten quer durch das elektromagnetische Spektrum zusammengetragen und analysiert. Basierend auf diesen Beobachtungen und einer Modellierung der Lichtkurve komme ich zu dem Schluss, dass *AT2019fdr* ein sogenanntes Tidal Disruption Event (TDE) darstellt, wenn auch ein ungewöhnliches. Die Wahrscheinlichkeit, dass eine solche Assoziation nur Zufall ist, liegt bei 0.034 %, wenn man eine weitere TDE-Neutrino-Assoziation mit einberechnet. Weiterhin diskutiere ich das Infrarot-Staubecho von diesem Objekt im Zusammenhang mit zwei weiteren Assoziationen von möglichen TDEs mit hochenergetischen Neutrinos, die ebenfalls ein solches Staubecho aufweisen.

Diese Studie wird begleitet von der Erstellung des ZTF nuclear sample, der ersten systematischen Sammlung innerhalb des ZTF-Datensatzes von solchen astrophysikalischen Ereignissen, die sich nahe dem Nukleus ihrer Wirtsgalaxie ereignen. Eines der Ziele dieser Untersuchung war es, die Zahl der TDEs zu vergrößern. Ich bewerkstellige dies durch die Entwicklung eines Algorithmus mit Verfahren des maschinellen Lernens zur photometrischen Typisierung astrophysikalischer Ereignisse. Dieser Klassifikator wird mit einem Datensatz nahegelegener astrophysikalischer Ereignisse trainiert, der zusätzlich künstlich verrauschter und lichtärmer gemacht wird, um dem nuclear sample mehr zu entsprechen. In Anwendung dieses Klassifikators auf das nuclear sample finde ich 27 neue TDE-Kandidaten. Die Identifikation von TDE-Kandidaten mittels ihres Infrarot-Staubechos ist ebenfalls erfolgreich und resultiert in 16 bisher nicht publizierten TDE-Kandidaten.

Introduction

For millennia, astronomy exclusively belonged to the domain of optics. Until the 17th century, the unaided eye, sensitive to wavelengths from 380 to 800 nm, was the sole 'instrument' [1] in observing the sky. Astronomy started with the practice of observing and predicting the motions of the Sun to forecast seasons in ancient Babylonia and Egypt[1] [2]. Over the centuries, observing the sky evolved into a mature science, leading to the Copernican Revolution, in which the Sun replaced the Earth as the center of the solar system. This pre-instrumental period also already saw the observation of galactic transient events, like the supernova (SN) *SN185* observed by Chinese astronomers almost two millennia ago, or the widely observed *SN1054*.

Astronomy was greatly accelerated by the invention of the telescope by Hans Lippershey, which allowed for the observation of much fainter objects [3]. Two important developments in the 19th century further advanced the field: Spectroscopy, allowing the identification of chemical elements in extraterrestrial objects, and photography, which vastly increased the depth of observations by gathering much more light as is possible when using a telescope with an eyepiece.

Radio and infrared astronomy extended the observational window to higher wavelengths in the 1930s and 1950s. To peek into the high-frequency domain with UV, X-ray and gamma-ray instruments was technologically even more challenging. It required satellites to escape the absorption of high-energy (HE) photons by the Earth's atmosphere, or—for the most energetic photons—necessitated the development of ground-based Imaging Atmospheric Cherenkov Telescopes (IACTs) to observe very-high-energy gamma rays.

After the invention of the telescope, it took another 360 years for the photon to lose its supremacy as the only messenger around. The advent of neutrino astronomy in the 1960s and lastly, gravitational wave astronomy in the 2010s initiated the era of multimessenger astronomy.

Besides those three messengers, the early 20th century saw the identification of a fourth potential messenger: Cosmic rays. These predominantly consist of highly energetic protons, electrons, positrons and helium nuclei.

The sky disk of Nebra. With an age of roughly 4000 years, this bronze disk is the oldest known depiction of the night sky. The group of seven stars has been interpreted as the Pleiades. Image source: Wikimedia Commons.

[1]: Wall (2018), *A History of Optical Telescopes in Astronomy*

[2]: Linton (2004), *From Eudoxus to Einstein: A History of Mathematical Astronomy*

[3]: Beckman (2021), *Multimessenger Astronomy*

---

[1] In Egypt, the seasons were not particularly prominent. Therefore, the practice of tracking the Sun relative to Sirius—the brightest star visible—to predict the Nile flooding had a direct impact on people's livelihood.

Cosmic rays come with a large drawback, though: Due to their charge, the particles constituting cosmic rays are deflected by magnetic fields on their way through interstellar and intergalactic space, which makes pinpointing their origin a very hard task [3].

High-energy astrophysical neutrinos could come to the rescue, as they are thought to stem from the very same processes that create cosmic rays. Their existence has been firmly established a decade ago, but their sources remain—at least partially—a mystery.

This thesis is concerned with those very sources. Chapter 1 contains a discussion of neutrino astronomy, including the subjects of neutrino physics, cosmic rays, production mechanisms and potential high-energy neutrino sources. After this, the two major instruments relevant for this thesis are detailed: IceCube (Chapter 2) and the Zwicky Transient Facility (Chapter 3). This is followed by an outline of the high-energy neutrino follow-up program and a discussion of the sources detected so far (Chapter 4). This is succeeded by a detailed discussion of a candidate Tidal Disruption Event that was found coincident with a high-energy neutrino: *AT2019fdr* (Chapter 5). Finally, the ZTF Nuclear Sample created by the author to further study nuclear transients is presented and analyzed in Chapter 6.

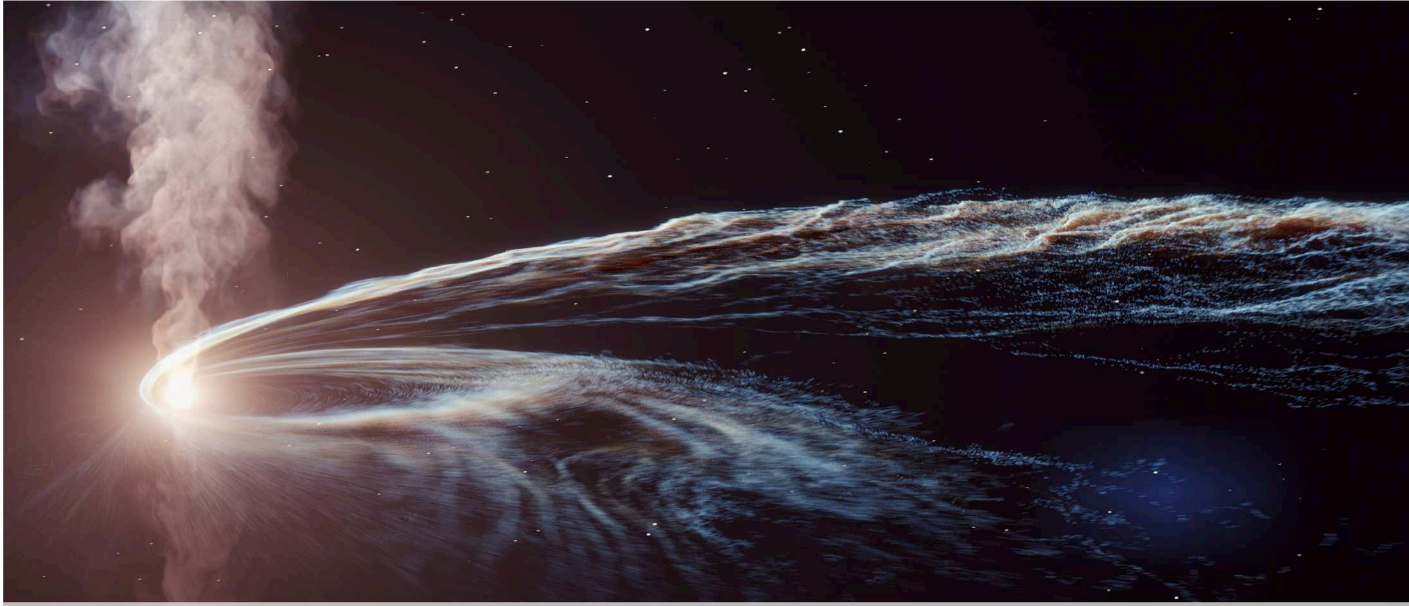

1 Neutrino Astronomy

Artistic impression of a tidal disruption event. Material from the shredded star starts accreting around a black hole, and a jet is launched. Image Credit: DESY and Science Communication Lab.

Contrary to photons—especially optical ones—neutrinos are extremely hard to detect. This explains why neutrino astronomy is such a recent development. As this thesis is concerned with the identification of the astrophysical sources of high-energy neutrinos, this chapter will first detail what neutrinos are, and then explain how high-energy neutrino astronomy has advanced. This will be followed by a discussion of different types of proposed source classes of high-energy neutrinos. Lastly, established counterparts to neutrinos will be discussed, as well as upper limits on most of the proposed source classes.

1.1 Neutrinos

Neutrinos were first predicted by Pauli in the 1930s [4], given firm theoretical footing by Fermi in 1934 [5], and experimentally detected by Reines and Cowan in 1955 [6].

[4]: Giunti et al. (2007), *Fundamentals of Neutrino Physics and Astrophysics*

[5]: Fermi (1934), *Versuch einer Theorie der β-Strahlen. I*

[6]: Reines et al. (1956), *The Neutrino*

1.1.1 The Neutrino Hypothesis

Wolfgang Pauli reluctantly proposed the neutrino to explain the continuous β-spectrum of decaying uranium. When nuclear elements decay, they emit different particles. Different types of emission were dubbed α-, β- and γ radiation [7, 8]; the first being helium-4 nuclei, the second electrons or protons, and lastly energetic photons. Contrary to α- and γ-radiation, the spectrum of β-radiation turned out to be continuous, while one would expect discrete energy levels for the emitted particles.

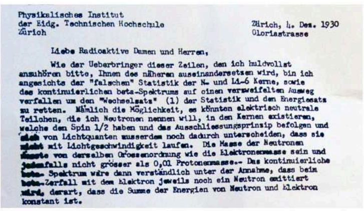

Physikalisches Institut
der Eidg. Technischen Hochschule
Zürich

Zürich, 4. Dez. 1930
Gloriastrasse

Liebe Radioaktive Damen und Herren,

Wie der Ueberbringer dieser Zeilen, den ich huldvollst
anzuhören bitte, Ihnen des näheren auseinandersetzen wird, bin ich
angesichts der "falschen" Statistik der N- und Li-6 Kerne, sowie
des kontinuierlichen beta-Spektrums auf einen verzweifelten Ausweg
verfallen um den "Wechselsatz" (1) der Statistik und den Energiesatz
zu retten. Nämlich die Möglichkeit, es könnten elektrisch neutrale
Teilchen, die ich Neutronen nennen will, in den Kernen existieren,
welche den Spin 1/2 haben und das Ausschliessungsprinzip befolgen und
sich von Lichtquanten ausserdem noch dadurch unterscheiden, dass sie
nicht mit Lichtgeschwindigkeit laufen. Die Masse der Neutronen
müsste von derselben Grossenordnung wie die Elektronenmasse sein und
jedenfalls nicht grösser als 0,01 Protonenmasse.- Das kontinuierliche
beta- Spektrum wäre dann verständlich unter der Annahme, dass beim
beta-Zerfall mit dem Elektron jeweils noch ein Neutron emittiert
wird, derart, dass die Summe der Energien von Neutron und Elektron
konstant ist.

Figure 1.1: Pauli's open letter from December 1930, proposing the existence of the neutrino (he called it 'neutron' at the time) to the community. Image credit: Pauli Letter Collection, CERN.

In β^--decay, the nucleus of an atom emits an electron when undergoing the nuclear transition $(A, Z) \rightarrow (A, Z+1) + e^-$. Here, A is the mass number (the number of nucleons in the atom), and Z is the atomic number (the number of protons in the nucleus)—a neutron-proton conversion occurs, as is known nowadays. In this scenario, the decay is a two body problem, involving the nucleus and an electron. When converting a neutron into a proton, the resulting binding energy of the atom is lower, and this

[7]: Rutherford (1899), *Uranium radiation and the electrical conduction produced by it*

[8]: Villard (1900), *Sur le rayonnement du radium*

energy needs to be accounted for. Also, the positive charge gained by the nucleus during the transition must be equalized somewhere.

Therefore, the emitted electron carries the quantized energy lost to the atom due to its transition to the final state. As was already known at that time, this tight distribution of energy is fixed by the difference in binding energy, so one would expect a narrow energy spectrum for the emitted electron, similar to α- and γ-radiation. But the observed electron spectrum was continuous, hence the puzzlement of the community. Bohr proposed loosening the requirement of energy conservation—a suggestion Pauli was strongly opposed to[1] [9].

Pauli instead took refuge in what he considered a desperate solution, and suggested that the β-decay was in fact a three body problem, with a neutral particle carrying a variable fraction of the released energy [10]. This would explain the continuous electron spectrum. He also assumed that the hypothetical second particle interacted extremely rarely, which explains why it had not yet been observed.

[9]: Jensen (2000), *Controversy and Consensus: Nuclear Beta Decay 1911–1934*

[10]: Bilenky (2012), *Neutrino. History of a unique particle*

The next advancement in the understanding of the neutrino was due to Enrico Fermi in 1934. By then, the picture of the atomic nucleus had been complemented by the neutron, discovered by James Chadwick two years earlier. Fermi developed the first theory of β-decay, analogous to the description of the emission of photons from excited nuclei [5]. He assumed that the electron-neutrino pair is produced when a neutron within the nucleus transitions into a proton: $n \longrightarrow p + e^- + \bar{\nu}$. The predictions from his theory were in fair agreement with observations, provided the neutrino mass was quite small [5].

1.1.2 Neutrino Detections

It took another 22 years until the neutrino was discovered experimentally. In the meantime, the nuclear bomb had been conceived, proposed, built, and dropped twice. The Savannah River Plant nuclear reactors were constructed as part of the nuclear arms race effort to develop more fission bombs, producing plutonium and tritium. On a more pacifistic note, Reines and Cowan used the hypothesized flux of neutrinos from the reactors to first experimentally detect a neutrino [6] (with a confirmation in [11]).

[11]: Cowan et al. (1956), *Detection of the Free Neutrino: A Confirmation*

The experiment was based on the inverse beta decay, in which the predicted neutrino (in fact, an electron antineutrino) reacts with a proton, converting it into a neutron and releasing a positron ($\bar{\nu_e} + p \longrightarrow n + e^+$). The released positron then annihilates with a free electron within the target material (for this they used H_2O + $CdCl_2$). In this process, two 511 keV photons are released, emitted in opposite directions due to conservation of momentum. These photons are then detected by a liquid scintillator surrounding the target. The neutron diffuses through the target medium for a while until it is finally captured by the cadmium, emitting a delayed photon signal (see Fig. 1.2).

The extremely low interaction rate, or, in other words, the small interaction cross-section of the neutrino, required large amounts of target material and liquid scintillator to detect this reaction [4]. Ultimately, thanks to the

[1] He brought forth a though experiment of a closed box, within which β-decay happens. The weight of this box would change over time, a result he deemed paradoxical.

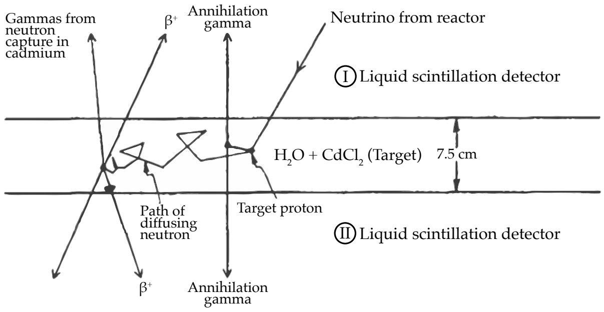

Figure 1.2: Experimental setup for the discovery of the neutrino. Adapted from [6].

unique event signature described above, Reins and Cowan succeeded, reporting a neutrino detection rate of $2.88^{+0.22}_{-0.22}$ counts h^{-1} [6].

1.1.3 Discovery of the Muon Neutrino

Another important step was the discovery of the muon neutrino: At the Brookhaven National Lab in the US, muons and kaons from an accelerator were used to prove the existence of this neutrino flavor. Prior to the experiment, it had already been pointed out by Lee and Yang that the experimental failure to detect the decay $\mu \longrightarrow e\gamma$ strongly hinted at the existence of a muon neutrino. This decay is only allowed if there is no difference between ν_e and ν_μ [12].

[12]: Ekspong (1993), *Nobel Lectures in Physics 1981 – 1990*

The goal of the Brookhaven experiment [13] was to directly detect the muon neutrino by producing π^+ through bombarding a Beryllium target with 15 GeV protons. π^+ primarily decay to $\mu^+ + \nu_\mu$, as the channel $\pi^+ \longrightarrow e^+ + \nu_e$ is strongly suppressed, see [10]. In a subsequent decay channel, almost all μ^+ decay, only allowing the muon neutrinos to pass. These were directed towards a neutrino detector, consisting of aluminum plates located in a spark chamber. The neutrinos would interact with the aluminum nuclei and produce charged leptons. If there was no difference between ν_μ and ν_e, one would expect detecting e^- and μ^- in equal numbers. However, the experiment observed a significant number of muons (29) and detected only a few electrons (6). The electrons could be attributed to background noise, therefore the experiment proved the existence of the muon neutrino [13].

[13]: Danby et al. (1962), *Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos*

1.1.4 The Solar Neutrino Problem

The neutrino was had been detected. This, combined with the prediction that the process of nuclear fission within the Sun should release neutrinos, triggered the hunt for solar neutrinos. The Sun creates neutrinos both via the pp chain and the CNO cycle, in both processes converting four protons and two electrons into a ^{4}He nucleus and two electron neutrinos: $4p + 2e^- \longrightarrow {}^4\text{He} + 2\nu_e + 26.7\,\text{MeV}$. The excess energy is then released via photons or neutrinos carrying kinetic energy [4].

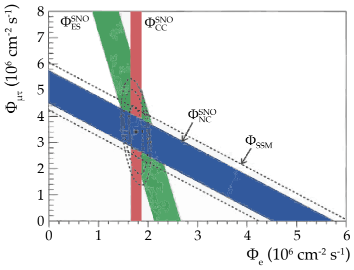

Figure 1.3: The solar neutrino flux as measured by SNO. The x-axis shows the ν_e flux, while the y-axis shows the flux of solar ν_μ and ν_τ. The intersection point shows the best-fit flux values for ν_e and $\nu_{\mu,\tau}$, with a resulting flavor ratio of ~ 1/3 for all types. Adapted from [14].

Solar electron neutrinos were first detected in the 1960s in the Homestake experiment by Davis. He operated a tank located underground in the Homestake mine in South Dakota in the US, filled with 380 m^3 of tetrachloroethene. This allowed antineutrino capture via the inverse beta decay $\nu_e + {}^{37}\mathrm{Cl} \longrightarrow {}^{37}\mathrm{Ar} + e^-$, transmuting chlorine to argon. First results after 150 days of data taking resulted in an upper neutrino flux limit much lower than the theoretical expectations [15]. The experiment would run continuously from 1970 to 1994, detecting a neutrino flux that never exceeded 50 % of the expected flux. The fact that the measured flux was consistently lower than the predicted one, despite various checks and improvements, became known as the solar neutrino problem [16].

[15]: Davis et al. (1968), *Search for Neutrinos from the Sun*

[16]: Bahcall et al. (1976), *Solar Neutrinos: A Scientific Puzzle*

A solution to the solar neutrino problem was first proposed by Gribov and Pontecorvo in 1969 [17]: The missing electron neutrinos from the Sun could have oscillated into neutrinos of a different flavor while traveling to Earth. This proposed neutrino oscillation was observed by the Sudbury Neutrino Observatory (SNO) in 2002; while evidence for the oscillation of atmospheric neutrinos had already been brought forth four years earlier by the Super-Kamiokande detector [18]. The SNO detector was sensitive to two types of interactions: One channel in which only ν_e could participate, and one open to neutrinos of all flavors. It was determined that the flux of solar ν_e was $\frac{1}{3}$ of the total neutrino flux (consisting of ν_e, ν_μ and ν_τ[2]), see Fig. 1.3. This result was in full agreement with theoretical predictions for the solar neutrino flux, and directly showed that neutrinos do oscillate [14].

[17]: Gribov et al. (1969), *Neutrino Astronomy and Lepton Charge*

[18]: Fukuda et al. (1998), *Evidence for Oscillation of Atmospheric Neutrinos*

[2] The existence of the tau neutrino had been confirmed one year earlier by the Direct Observation of the Nu Tau (DONUT) experiment [19].

[14]: Ahmad et al. (2002), *Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory*

These oscillations can be described by the 3 × 3 unitary transformation matrix known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, that works analogous to the matrix describing the quark flavor mixing (Cabbibo-Kobayashi-Maskawa, or short CKM, matrix). The values of the matrix elements must be determined experimentally, as there is no underlying theory predicting them.

1.1.5 Neutrino Mass

By observing solar neutrino oscillations, the SNO experiment also established the fact that neutrinos have mass—contradicting the Standard Model model of particle physics. Only massive particles allow oscillation between flavor eigenstates and mass eigenstates [20]. It therefore does not make sense to ask for the mass of e.g. an electron neutrino, as this flavor eigenstate is a superposition of different mass states. One can speak of the effective electron neutrino mass, though; this property is defined as

[20]: Fisher et al. (1999), *Neutrino Mass and Oscillation*

$$m_\nu = \sqrt{\Sigma_i |U_{ei}|^2 m_i^2}, \tag{1.1}$$

where U is the PMNS matrix describing the neutrino state mixing, i.e. a collection of nine experimentally determined values. If one could take individual mass measurements of the electron neutrino, one would measure the mass m_i with probability $|U_{ei}|^2$.

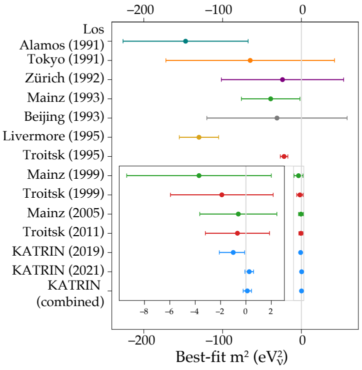

Figure 1.4: The history of upper limits on the neutrino mass. From [21].

Neutrino oscillations are only sensitive to the squared mass differences between the three mass states, as there is no zero point in Eq. 1.1. One way to measure the average proper neutrino mass is the close inspection of the beta-decay spectrum (see Section 1.1.1). The most recent in a long line of experiments to determine neutrino masses (see Fig. 1.4) is the Karlsruhe Tritium Neutrino Experiment (KATRIN). It uses the beta decay of tritium. If the neutrino is massless, the energy spectrum of the emitted electron will extend up to the total energy released in the decay (18.6 eV). On the other hand, if the neutrino mass is non-zero, the electron spectrum will always be short of that amount of energy. Note that in principle one could obtain all three neutrino masses from such an experiment if it weren't for current technological limitations. The latest upper limit from KATRIN on the effective electron antineutrino mass is $m_\nu < 0.8\,\mathrm{eV}$ at the 90 % confidence level [21].

[21]: M. Aker et al. (2022), *Direct neutrino-mass measurement with sub-electronvolt sensitivity*

1.1.6 Current Understanding

The Standard Model of particle physics knows three flavors of neutrinos, which are superposition of three mass states [22]. Their mass is small, but exists. Neutrinos have neither electromagnetic, nor color charge, but possess a weak hypercharge. This weak hypercharge lets them partake in weak interactions, which are mediated by exchanging W or Z bosons. Fig. 1.5 details some possible neutrino interaction channels, all constituting deep inelastic scattering with the quarks contained in nucleons. This is the dominant interaction mode for high-energy neutrinos with energies at or above several GeV in a medium.

[22]: Barger (2012), *The Pysics of Neutrinos*

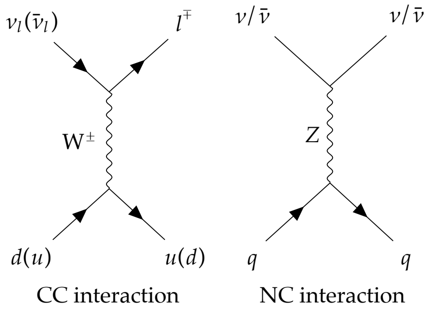

Figure 1.5: Different neutrino interaction channels for deep inelastic scattering. Left: interaction of neutrinos (antineutrinos) with quarks by exchange of a $W^\pm$ boson. Right: interaction of neutrinos of all types by exchange of a Z boson. The interaction on the left is called Charged Current (CC) interaction, while the one on the left is dubbed Neutral Current (NC) interaction.

Interactions that involve electromagnetically charged W^- bosons are called Charged Current (CC) interactions, while those involving neutral Z bosons are called Neutral Current (NC) interactions. For example, a neutron beta-decaying into a proton is a charged current interaction, as charge needs to be moved.

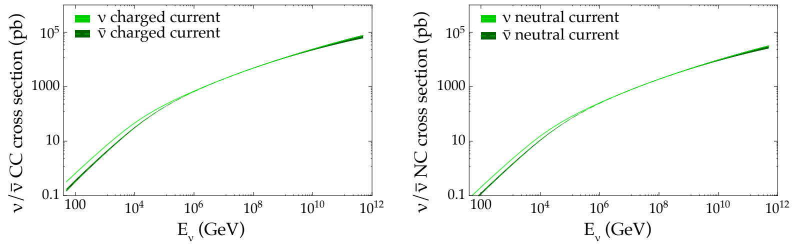

Figure 1.6: High-energy neutrino cross-section for interactions with matter equally composed of neutrons and protons. The unit on the y-axis are picobarn ($10^{-40}\,\mathrm{m}^2$). The figure on the left shows charged-current interactions, while the one on the right displays neutral-current interactions, both for neutrinos and antineutrinos. Adapted from [23].

The cross-section for neutrinos with matter significantly increases with energy, as can be seen in Fig. 1.6. This is helpful for the detection of high-energy neutrinos, as this effect at least partially counterbalances the probable decrease in numbers of neutrinos produced in cosmic sources towards higher energies. On the other hand, neutrino detectors like IceCube are increasingly insensitive to neutrinos of the highest energies that sometimes have to traverse the Earth's core and get absorbed along the way.

For a better understanding of neutrino astronomy, it is beneficial to broadly divide the subject into two main categories, as the processes and instruments involved are fairly different: low-energy and high-energy neutrino astronomy. As high-energy neutrino astronomy is intertwined with cosmic-ray astrophysics, a discussion of cosmic rays will be interspersed.

1.2 Low-Energy Neutrino Astronomy

Low-energy neutrinos, i.e. those carrying energies up to $\mathcal{O}$(MeV), are produced in stellar fusion processes and in supernovae. The primary stellar fusion process we can study is of course the Sun.

1.2.1 Stellar Neutrinos

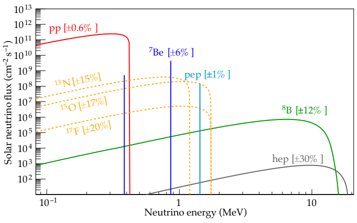

Figure 1.7: Predicted solar neutrino flux. From [24].

As briefly mentioned above, the solar neutrinos stem from two processes: The pp chain and the CNO cycle. Within the pp chain, neutrinos get produced by the proton-neutron conversion preceding the fusion of $p + n$ to ^{2}H. This conversion either happens via the β^{+} decay $p \longrightarrow n + e^{+} + \nu_{e}$, resulting in a continuous neutrino spectrum (red curve in Fig. 1.7), or via electron capture $e^{-} + p \longrightarrow n + \nu_{e}$, which results in a line spectrum (leftmost blue line).

As one can see in Fig. 1.7, the predicted solar neutrino spectrum is composed of different channels. Towards MeV energies, the solar neutrino flux is dominated by the decay of $^{8}\mathrm{B} \longrightarrow {}^{8}\mathrm{Be} + e^{+} + \nu_{e}$ (green curve) happening in one of the subsequent evolutions of the pp chain. This particular process does not contribute much to the radiative energy output of the Sun, but it was the responsible mechanism for the solar neutrino flux, of which $\frac{2}{3}$ were missing, constituting the aforementioned solar neutrino puzzle.

Fig. 1.7 also shows that the predicted flux cuts off around ~10 MeV, so only low-energy neutrinos are expected from the Sun and similar stars.

1.2.2 Supernovae

Only with the first extragalactic neutrino source the era of neutrino astronomy truly began. This source was *SN1987a*, an extremely close-by supernova in the Large Magellanic Cloud (LMC) that occurred in February 1987 [25].

[25]: Kunkel et al. (1987), *Supernova 1987A in the Large Magellanic Cloud*

[26]: Spurio (2018), *Probes of Multimessenger Astrophysics*

Since Kepler's supernova almost 500 years ago—a Type Ia supernova (see below) that occurred in 1604 in the Milky Way—*SN1987a* was the first supernova observable by the naked eye alone, sporting an impressive brightness of 3 – 4 mag due to its proximity [26]. *SN1987a* was the core-collapse of *Sanduleak-69202*, a blue supergiant, resulting in a Type IIP supernova[3] [28]. After the optical detection of the source the neutrino detectors Kamiokande-II, Irvine-Michigan-Brookhaven (IMB) and the Baksan Neutrino Observatory (BNO) measured a burst of neutrinos predating the light from the supernova by 3 hours [29, 30], see also Fig. 1.8.

[3] Though the rise instead of a plateau phase within its light curve gave it its own class, SN1987A-likes [27].

[28]: Utrobin et al. (2021), *Supernova 1987A: 3D Mixing and Light Curves for Explosion Models Based on Binary-merger Progenitors*

[29]: Hirata et al. (1987), *Observation of a neutrino burst from the supernova SN1987A*

[30]: Bionta et al. (1987), *Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud*

In total, 25 neutrinos were detected by the three instruments (had IceCube already been operational in 1987, it should have detected ~13,000 neutrinos [31]). These rates were in agreement with theoretical models for core-collapse supernovae, assuming that 99 % of the supernova energy is released in the form of neutrinos of all flavors.

[31]: Halzen et al. (2017), *Neutrinos from Core-Collapse Supernovae and Their Detection*

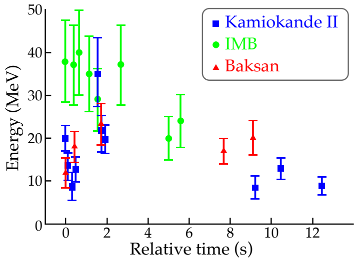

Figure 1.8: The neutrinos from *SN1987a*, as measured by Kamiokande-II, IMB and BNO (Baksan). Figure adapted from [32].

Historically, supernovae have been classified according to the presence of lines in their spectra. SNe of Type I lack hydrogen lines, in contrast to Type II SNe which do show hydrogen lines. Non-hydrogen SNe are further subdivided into SNe Ia that display both silicon and helium lines, SNe Ib that lack the silicon lines, and SNe Ic that contain neither silicon nor helium lines in their spectra.

SNe Ia

SNe Ia constitute a class sui generis. The underlying physical scenario for their explosion is often assumed to be a white dwarf accreting matter from a companion star, the so-called Single Degenerate (SD) scenario. The accretion eventually causes the star's mass to exceed the Chandrasekhar limit of $\sim 1.4\ M_{\odot}$. At this point the gravitational force overcomes the electron degeneracy pressure, and the star begins to burn carbon due to the increased pressure. Soon after that, a runaway thermonuclear reaction starts, and the white dwarf blows up violently [33].

[33]: Iben et al. (1984), *Supernovae of type I as end products of the evolution of binaries with components of moderate initial mass*

A second possible scenario that has recently gained traction is the Double Degenerate (DD) scenario. In this model, the SN Ia is the result of two gravitationally bound white dwarfs (a binary system) shedding gravitational energy in the form of gravitational waves and thereby slowly closing in on each other. Ultimately, they merge and cause the supernova explosion [33].

In either scenario (see [34] for a review), the luminosity of the resulting explosion can be standardized using the Philips relation [35]. This makes SNe Ia standard candles, excellent instruments to measure cosmological distances. As they are mainly powered by the decay $^{56}\text{Ni} \rightarrow {}^{56}\text{Co} \rightarrow {}^{56}\text{Fe}$, they are expected to produce low-energy neutrinos, but not in large quantities when compared to core-collapse SNe (see next Section 3) [36].

[34]: Maoz et al. (2014), *Observational Clues to the Progenitors of Type Ia Supernovae*

[35]: Phillips (1993), *The absolute magnitudes of Type IA supernovae*

[36]: Alsabti et al. (2017), *Supernovae and Supernova Remnants: The Big Picture in Low Resolution*

Core-Collapse Supernovae

In contrast to SNe Ia, all other types of SNe—no matter if their spectra show hydrogen lines or not—are Core-Collapse Supernovae (CCSNe), resulting from the gravitational collapse of massive stars.

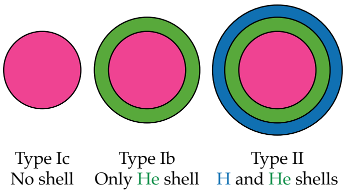

Figure 1.9: CCSN shells. The presence or absence of helium and hydrogen shells explains the differences in the respective spectra of CCSNe Type Ib, Ic and II. Because Ib and Ic Type SNe have lost parts of their outer shells, they are also referred to as Stripped-Envelope CCSNe. SNe IIb have almost lost their H shell, allowing them to quickly transform into a Type I SN.

The spectral differences between SNe of Type Ib, Ic and II can be explained by the presence or absence of the two outermost shells: hydrogen and helium (see Fig. 1.9). If both shells have already been blown away to a large distance from the star, the SN is of Type Ic, and neither H nor He lines are present in the spectrum. If only the helium shell remains, the type is Ib, and if both shells are still present, the collapse will result in a Type II supernova.

Because Ib and Ic supernovae have partly or entirely shed their outermost shells, they are referred to as Stripped-Envelope (SE) SNe. The outer layers are lost either due to strong solar winds, or due to the proximity of a companion star. SNe of Type IIb are also part of this class, as they initially display prominent hydrogen lines that fade over time until they disappear altogether. This behavior is most likely caused by the progenitor star having already lost all but a thin shell of its hydrogen [37].

[37]: Sravan et al. (2020), *Progenitors of Type IIb Supernovae. II. Observable Properties*

To complete the picture, some other subclasses of Type II SNe are motivated either by the presence of narrow lines in their spectra (SNe IIn) or by a peculiar light curve evolution: SNe IIP show a plateau phase, while SNe IIL display a linear decay.

Stellar Collapse

Stars generate energy by first burning hydrogen. When a massive star has exhausted its hydrogen, fusion subsides and the gravitational pull compresses it further, until temperatures suitable for helium burning are reached. This process repeats for carbon, neon, oxygen and silicon. The silicon burning produces nickel, which ultimately decays to iron.

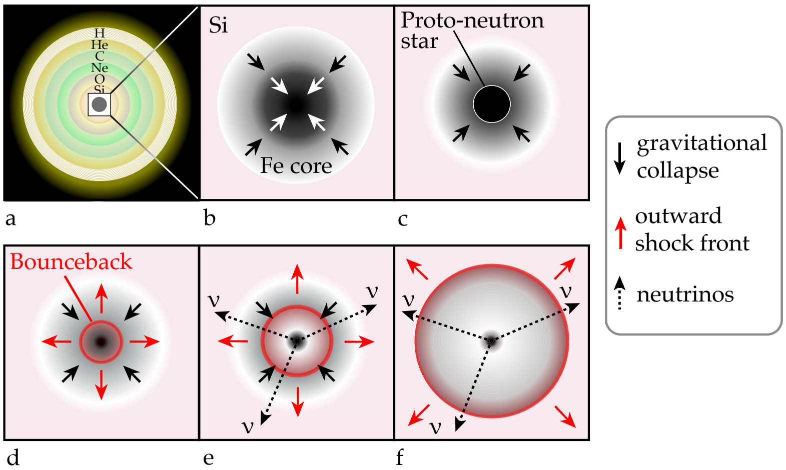

Figure 1.10: Schematic of a core-collapse supernova. Fusion stops in the core (b), causing it to collapse until neutron degeneracy prevents further collapse. A proto-neutron star forms in the center (c). The infalling material bounces back from the degeneracy border and forms a shock front rapidly moving outwards (d). The stalling shock is reinvigorated by the intense flux of neutrinos from the proto-neutron star (e) and expands further outwards (f). Adapted from `https://commons.wikimedia.org/wiki/File:Core_collapse_scenario.png` (which is most likely adapted from [38]).

A CCSN happens when a star with a mass between 8 and 140 $M_\odot$ has reached its final fusion stage—silicon burning—which ultimately produces iron. Iron has the highest binding energy of all elements. Therefore, fusing iron to heavier elements would release no energy, and the fusion chain stops. With the fusion in the core subsided, the radiation pressure that so far prevented the star from collapsing in on itself suddenly decreases, and the infalling material compresses the core of the star within a second. As soon as the core becomes incompressible due to neutron degeneracy pressure described by the Pauli exclusion principle, the infalling material bounces back from the incompressible core and blows out into space. The remaining neutron star at the center either survives, or the infalling material eventually causes it to overcome neutrino degeneracy pressure and collapse into a black hole [36].

The more massive the star, the higher the radiation pressure from fusion. This radiation pressure is what drives the shedding of outer layers, so it is assumed that the most-stripped events—SNe Ic—result from a collapse of the most massive stars, followed by SNe Ib and lastly 'regular' Type II supernovae, which are powered by stars having largely retained their outer shells.

In the centers of these core-collapse supernovae, neutron stars with masses of $1.5 - 2\ M_\odot$ are created. These need to release large amounts of gravitational binding energy during the short moment of their creation. Because the environment at the center of the CCSN is dense and optically thick, photons are unable to escape: They immediately produce electron-positron pairs, which are in turn quickly absorbed by the surrounding matter.

Instead, 99 % of the binding energy is carried away by low-energy neutrinos (see panels (e) and (f) in Fig. 1.10). The energy radiated by the nascent neutron star via neutrinos typically amounts to $\sim 1.5 \times 10^{53}$ erg, distributed equally over all three neutrino families, with a typical energy carried per neutrino of $\mathcal{O}(10\,\text{MeV})$ [39].

[39]: Lunardini (2017), *Diffuse Neutrino Flux from Supernovae*

The exact role neutrinos play in core collapses is still insufficiently understood, as one has to rely on data from *SN1987a* only, supplemented by numerical simulations. Eventually, a CCSN within our galaxy might come to the rescue [40].

[40]: Janka (2017), *Neutrino-Driven Explosions*

1.3 Cosmic Rays

Now, there might also be more energetic neutrinos from space. But as these are thought to be produced during the same processes as cosmic rays, high-energy neutrinos and cosmic rays are intricately linked. Therefore, one first needs a basic understanding of cosmic rays to understand high-energy neutrino production.

1.3.1 Discovery

Cosmic rays were discovered in 1912 by Victor Hess. At the time, it had already been discovered that electroscopes—devices measuring electric charge and voltage—sometimes spontaneously discharged. This was attributed to particles ionizing the air in the vessel, making it conductive and therefore allowing it to discharge. This was corroborated by the fact that reduced air pressure would slow the speed of the discharge. Another hint was that when shielding the isolated vessel, the ionization would decrease. The source of these ionizing particles was unknown, though.

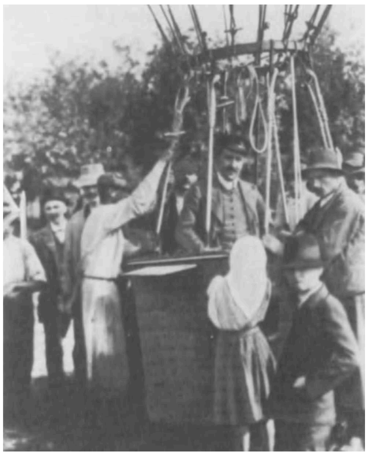

Figure 1.11: Hess in his balloon after landing in Brandenburg, Germany in 1912, having just discovered cosmic rays. From [41].

Hess measured the flux of ionizing radiation with three electrometers during several hot-air balloon journeys. In total, he carried out ten ascents in the years from 1911 to 1913. The first ascent over 5 km in August 1912 marked the discovery of cosmic rays, when he registered a 16-fold increase in ionization measured by the electrometers at that altitude. Hess consequently concluded that the ionizing radiation must be of extra-terrestrial origin [41].

[41]: Steinmaurer (1985), *Erinnerungen on V. F. Hess, Den Entdecker der Kosmischen Strahlung, und an Die ersten Jahre des Betriebes des Hafelekar-Labors*

1.3.2 Composition and Energy Spectrum

More than a century later, we know that this ionizing radiation consists of massive particles, spanning a broad range of energies. They are mostly protons, but all stable charged particles and nuclei with lifetimes of $\mathcal{O}(10^6\,\text{yr})$ can be found; most prominently protons, electrons, helium, but also e.g. carbon, oxygen and iron [42].

[42]: R L Workman et al. (2022), *Review of Particle Physics*

As can be seen in Fig. 1.12, the cosmic ray spectrum spans a wide range of energies, from a few MeV up to ZeV, with the high end of the energy spectrum constituting Ultra High-Energy Cosmic Rays (UHECR).

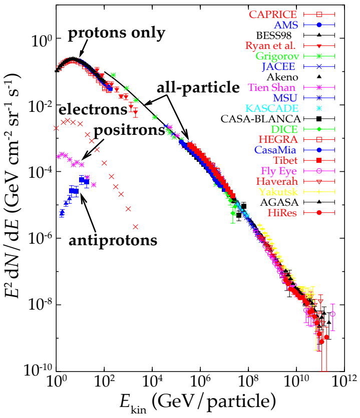

Figure 1.12: Cosmic ray spectrum, as seen by a range of experiments. Adapted from [43].

The UHECR part of the energy spectrum is shown in Fig. 1.13. It can be described by a series of power laws with different spectral indices, i.e. $E^{-\gamma}$, with γ being the spectral index. The UHECR spectrum shows several interesting features, dubbed the Knee, the anatomically challenging Second Knee, and the Ankle. These are the points where the spectral index changes, and which are possibly associated with a transition to a new source class, e.g. from galactic to extragalactic sources around the ankle (see [44]). This is still an active field of research [42].

[44]: Taylor (2016), *Cosmic rays beyond the knees*

If high-energy neutrinos are produced in the processes responsible for the creation of cosmic rays (see below), the neutrinos stemming from cosmic rays with energies higher than the ankle should be mainly produced by extragalactic sources [45].

[45]: Halzen (2006), *Lectures on High-Energy Neutrino Astronomy*

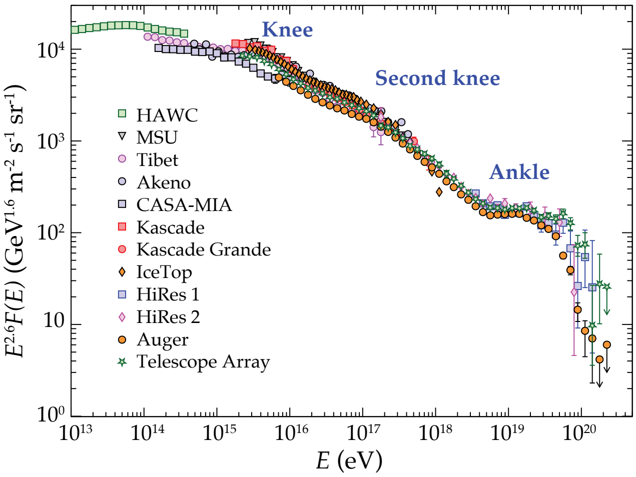

Figure 1.13: Ultra high-energy cosmic ray spectrum. It shows energy per particle/nucleus vs. the flux (multiplied by $E^{2.6}$ to highlight the changes in spectral index). The data stems from various air-shower experiments. Adapted from [42].

Identifying the sources of UHECR remains a challenge. As the cosmic-ray particles are charged, they are deflected by magnetic fields encountered on their way to Earth, obscuring their origin. What is known though is that their extremely high energies necessitate extreme production environments.

In the 1980s, Hillas proposed the requirement that the efficient acceleration of a charged particle needs an accelerator that is able to confine the particle during acceleration: This means that the source needs to be larger than the particle's Larmor radius [46], i.e. the radius of the circle a charged particle travels on given an external magnetic field. This introduces a bound on the possible sources for UHECR cosmic ray sources:

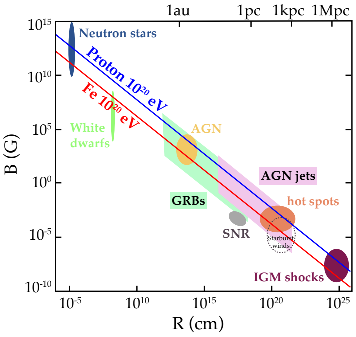

Figure 1.14: Possible sources for 10^{20} eV cosmic rays as a function of source radius R and the magnetic field strength B of the source. Adapted from [46], the original 'Hillas plot' can be found in [47].

$$E \le 10^{20} Z \left(\frac{B}{10\,\mu\text{G}} \right) \left(\frac{R}{10\,\text{kpc}} \right) \text{eV}, \tag{1.2}$$

where Z is the cosmic ray particle charge number, B is the magnetic field strength of the source, and R is the characteristic source dimension [46]. Sources that match this criterion for the production of UHECR protons and iron nuclei with an energy of 10^{20} eV are shown in Fig. 1.14. As one can see, a variety of extreme environments are possible sources of UHECR.

[46]: Rieger (2022), *Active Galactic Nuclei as Potential Sources of Ultra-High Energy Cosmic Rays*

As some of these also hot candidates for the production of high-energy neutrinos, it is instructive to look at the mechanisms involved.

1.3.3 Diffusive Shock Acceleration

There are several acceleration mechanisms for UHECR. The most important of these processes, Diffusive Shock Acceleration (DSA), was introduced in the late 1970s. DSA is thought to occur in the presence of shocks that are present e.g. in core-collapse supernovae (see Section 1.2.2). Such shocks are formed when matter, for example in the form of a plasma, is moving with supersonic speed through a surrounding medium.

To understand DSA, one has to differentiate between an upstream region in front of the moving shock front, and a downstream region behind the shock front.[4] Consider a shock front moving with a velocity $\vec{v}_s$ through a medium as seen in the observer's rest frame (red line in Fig. 1.15). If we switch to the rest frame of the shock front, it encounters the upstream medium as flowing towards it with speed $\vec{u}_u = -\vec{v}_s$. Downstream, behind the shock, the velocity of material u_d will be lower: Mass conservation requires that $\rho_u u_u = \rho_d u_d$. In astrophysical shocks containing a fully ionized plasma, the compression ratio $R = \rho_d/\rho_u$ can be as high as 4 [48]. Because $R = u_u/u_d$, the velocity of material downstream (behind the shock front) u_d will then only be $1/4u_u$.

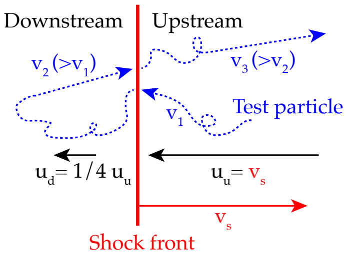

Figure 1.15: Sketch illustrating diffusive shock acceleration. A shock front is moving with velocity v_s with respect to an upstream medium. A test particle crosses the shock front twice, each time gaining energy. The length of the arrows are proportional to the velocity.

[4] The majority of this section is owed to [48].

[48]: Longair (2011), *High Energy Astrophysics*

Now consider some test particles upstream, ahead of the shock front rushing towards them. In the rest frame of the upstream gas their velocity distribution will be isotropic. Assume that the diffusion on both sides of the shock front is mediated via collisionless processes. This means that momentum and energy transfers between particles are mediated elastically by magnetic irregularities in the plasma, and Coulomb scattering can be neglected.

When a test particle with initial velocity v_1 (blue arrow in Fig. 1.15) eventually crosses the shock front to the downstream region, it will see a plasma rushing towards it with speed $w = |u_u - u_d| = (1 - 0.25)\,u_u$, which equals 3/4 of the shock front velocity v_s.

In the downstream region, they will again be scattered collisionlessly, and receive a small net energy gain of $\left\langle \frac{\Delta E}{E} \right\rangle = \frac{2}{3}\frac{w}{c}$ (see below). This energy gain results in a higher velocity v_2. After some time, the particle might cross the shock front again, this time into the upstream region. From the now isotropized particle's frame, the upstream region will again rush in with 0.75 v_s, and the process repeats, resulting in another gain in energy, and a higher v_3 [50], with an average round-trip energy gain of $\left\langle \frac{\Delta E}{E} \right\rangle = \frac{4}{3}\frac{w}{c}$. See Fig. 1.16 for a Monte Carlo simulation of a test particle crossing a shock front multiple times, accumulating a higher and higher velocity over time (blue line).

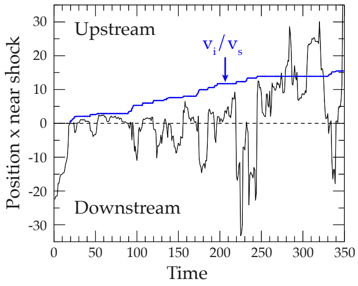

Figure 1.16: Monte Carlo simulation of a test particle near the shock front. The particle position wildly varies (solid line), but its velocity v_i (blue line) increases each time it crosses the shock front at $x = 0$. Adapted from [49]

[50]: Urošević et al. (2019), *Particle acceleration in interstellar shocks*

We still need to motivate the net energy gain per crossing of $\left\langle \frac{\Delta E}{E} \right\rangle = \frac{2}{3}\frac{w}{c}$. This can be explained as follows: First, one needs to switch to the test particle frame, i.e. perform a Lorentz transformation. The energy of the test particle will be:

$$E' = \gamma_w(E + p_x w), \tag{1.3}$$

where x is the coordinate perpendicular to the shock. The shock is non-relativistic ($w << c$ and $\gamma_w \approx 1$), but the test particle is, so $E \approx pc$ and $p_x \approx E/c\cos(\theta)$. From this it follows that

$$\Delta E = E' - E = E + \frac{E}{c} w \cos\theta - E = pw\cos\theta \tag{1.4}$$

and

$$\frac{\Delta E}{E} = \frac{p}{w}\cos\theta = \frac{w}{c}\cos\theta. \tag{1.5}$$

The probability of the test particle angle θ to lie between θ and $\theta + \mathrm{d}\theta$ is proportional to $\sin\theta\mathrm{d}\theta$, and the test particle approach rate is proportional to $v_x = \frac{p_x}{m} = \frac{E/c\cos\theta}{E/c^2} = c\cos\theta$. The probability of the test particle crossing the shock front is therefore proportional to $\sin\theta\cos\theta\mathrm{d}\theta$. The integral from 0 to $\frac{\pi}{2}$ is 0.5, so we need to multiply by 2 to normalize the cumulative probability to 1:

$$p(\theta) = 2\sin\theta\cos\theta\mathrm{d}\theta \tag{1.6}$$

By integrating $p(\theta)$ and using Eq. 1.5, we can calculate the average energy gain of each shock crossing:

$$\left\langle \frac{\Delta E}{E} \right\rangle = \int_0^{\pi/2} \frac{\Delta E}{E} p(\theta) = \frac{w}{c} \int_0^{\pi/2} 2\sin\theta\cos^2\theta\mathrm{d}\theta = \frac{2}{3}\frac{w}{c} \tag{1.7}$$

After another crossing, the test particle will gain another average energy increase of $\frac{2}{3}\frac{w}{c}$, resulting in an average energy gain of $\frac{4}{3}\frac{w}{c}$ per round trip [48].

During each crossing, a particle with initial energy E_0 will gain an energy fraction β, resulting in energy $E = \beta E_0$. If P denotes the probability that a particle stays within the accelerating region, the number of particles after k crossings can be written as $N = N_0 P^k$, and they will have energies $E = E_0\beta^k$.

If one solves both equations for k and sets them equal, one obtains

$$\frac{\ln(N/N_0)}{\ln(E/E_0)} = \frac{\ln P}{\ln\beta}. \tag{1.8}$$

By rearranging we get

$$\frac{N}{N_0} = \left(\frac{E}{E_0}\right)^{\ln P/\ln\beta} \tag{1.9}$$

To obtain the differential energy spectrum, we differentiate with respect to E and get

$$N(E)\,\mathrm{d}E = C \times E^{(\ln P/\ln\beta)-1}\,\mathrm{d}E \tag{1.10}$$

The result is a power law with index $\frac{\ln P}{\ln \beta}$, which is indeed what was already stipulated by experimental data.

We already know that $\left\langle \frac{\Delta E}{E} \right\rangle = \frac{4}{3}\frac{w}{c}$ per round trip. With this, $\beta = \frac{E}{E_0} = 1 + \frac{4}{3}\frac{w}{c}$. What is missing is an expression for P. To obtain this, we need to estimate the rate at which particles drop out of the system or are 'swept away'. As argued by [48] (and originally presented by [51]), the average number of particles crossing the shock is $\frac{1}{4}nc$, with n being the particle number density. This is true in both directions up- and downstream. The difference though is that downstream particles are 'advected' away from the shock further downstream and out of the system, as they are isotropic in that frame. The dropout rate is $nw = \frac{1}{4}nv_s$. From this it follows that the fraction of particles lost (per unit time) is $\frac{1}{4}nv_s/\frac{1}{4}nc = v_s/c$. The survival probability P is thus $P = (1 - v_s/c)$.

[51]: Bell (1978), *The Acceleration of Cosmic Rays in Shock Fronts*

With that, we can compute

$$\ln P = \ln\left(1 - \frac{v_s}{c}\right) \simeq -\frac{v_s}{c} \tag{1.11}$$

and

$$\ln \beta = \ln\left(1 + \frac{4}{3}\frac{w}{c}\right) \simeq \frac{4}{3}\frac{w}{c} = \frac{v_s}{c}. \tag{1.12}$$

This gives us $\frac{\ln P}{\ln \beta} \simeq -1$. Inserting this into Eq. 1.10 yields

$$N(E)\,\mathrm{d}E \propto E^{-2}\,\mathrm{d}E. \tag{1.13}$$

Therefore, the resulting differential energy spectrum has a spectral index of 2, a value that somewhat differs from the experimentally found spectral index of 2.7 up to the knee. This can be explained by the fact there are several assumptions at play here: The shock is assumed to be non-relativistic, situated in an ideal gas and with a constant escape probability. Inefficiencies in the shock, and the inclusion of shock-amplified magnetic fields can both harden the spectral index [26].

1.3.4 Interaction and Neutrino Production

As discussed above, the observed cosmic ray spectrum is readily explained by astrophysical shocks, as these account for the power-law structure of the energy distribution.

The particles accelerated within shocks might now interact with either hadrons and photons in their vicinity or on their way to Earth. For simplicity, this section focuses on the interaction of cosmic-ray protons, but it applies to heavier nuclei as well.

pp**−interactions** When cosmic-ray protons encounter gas, either at the source location or while crossing the universe, they can interact with this gas, here assumed to consist of protons only for simplicity. The resulting pp-interactions produce lots of unstable hadrons, which decay. Those decays are dominated by pion production. Ignoring all secondary hadrons X (which can also decay via pions), possible decay channels of these pions are:

$$\pi^{\pm} + X \longrightarrow \mu^{\pm} + \nu_{\mu}(\bar{\nu}_{\mu}) \longrightarrow e^{\pm} + \nu_{e}(\bar{\nu}_{e}) + \bar{\nu}_{\mu}(\nu_{\mu}), \quad (1.14)$$

$$\pi^{0} + X \longrightarrow 2\gamma. \quad (1.15)$$

As can be seen, $\pi^{\pm}$ decay into muons and a muon (anti-) neutrino. Those muons will decay into electrons (positrons), electron (anti-) neutrinos and muon (anti-) neutrinos. The decay of the neutral π^0 does not generate neutrinos. For pion decays, an average neutrino flavor content of ν_e:ν_μ:ν_τ = 1 : 2 : 0 is expected. Neutrino oscillations on the way to Earth wash out this initial difference, with an expected final flavor content of 1 : 1 : 1 [42].

The energy of the pion decay products is expected to trace the initial proton energy spectrum. If the proton interaction rate is fairly independent of their energy, they will then lose a fraction of their energy. This energy will then be converted into the production of pions, and subsequently muons. With this, the resulting neutrino spectrum is expected to have a similar shape as the initial cosmic ray spectrum [52].

[52]: Fiorillo et al. (2023), *Bump-hunting in the diffuse flux of high-energy cosmic neutrinos*

$p\gamma$**−interactions** Cosmic-ray protons can also interact with radiation fields. The cross-section for these interactions is large near the Δ^+ resonance. The dominant Δ^+ decay modes also contain pions:

$$\Delta^{+} \longrightarrow n + \pi^{+} \quad (1.16)$$

$$\Delta^{+} \longrightarrow p + \pi^{0} \quad (1.17)$$

The pions will again decay according to the channels already discussed for pp-interactions. The secondary neutrons might decay via beta decay, and produce electron antineutrinos. Here, the shape of the neutrino spectrum not only depends on the proton spectrum, but also the photon spectrum. Therefore, it is not expected to trace the shape of the initial cosmic-ray spectrum [53].

[53]: Fiorillo et al. (2021), *Unified thermal model for photohadronic neutrino production in astrophysical sources*

1.4 High-Energy Neutrino Astronomy

As shown above, high-energy cosmic ray particles can produce high-energy neutrinos in pp- or $p\gamma$-interactions. These high-energy neutrinos have energies in the MeV range and above. But what are possible source classes for these? There will be some overlap with the low-energy neutrino events, as some of those—like e.g. supernovae—are also probable sources of high-energy neutrinos.

1.4.1 Interacting Supernovae

As already detailed in Section 1.2.2, core-collapse supernovae are thought to produce copious amounts of MeV-neutrinos. Interestingly, certain types of CCSNe are thought not only to produce such comparably low-energy neutrinos, but also GeV to TeV neutrinos. Optical observations have stipulated that supernova progenitors often violently eject material into the circumstellar region in the years or centuries prior to their explosion [54]. When the supernova happens, ejecta are driven into this circumstellar material (CSM).

[54]: Ofek et al. (2014), *Precursors prior to Type IIn Supernova Explosions are common: Precursor Rates, Properaties and Correlations*

SNe of Type IIn are a subclass first stipulated by Schlegel [55]: The Balmer emission features of early IIn spectra, most prominently the H_α region, show narrow lines on top of a broad component [56]. The broad component is caused by the supernova itself, but the narrow lines result from interaction of the SN blast-wave with the CSM. This process photo-ionizes the CSM, which results in narrow hydrogen lines.

[55]: Schlegel (1990), *A new subclass of type II supernovae?*

[56]: Taddia et al. (2013), *Carnegie Supernova Project: Observations of Type IIn supernovae*

Such a system might be promising for the production of high-energy neutrinos, as highlighted in [57]. When the SN ejecta, expanding with velocities of 3000–10,000 $\mathrm{km\,s^{-1}}$, hit the CSM, a shock wave starts to propagate through the CSM. Particle acceleration within the shock is predicted to happen at the moment of shock breakout. This is the time when the optical depth of the shock drops below $\approx c/v$, with v being the shock velocity. As soon as the radiation can escape, the shock is no longer radiation-mediated, and protons can be accelerated to high energies via diffusive shock acceleration (see Section 1.3.3) [58]. This emission seizes after the shock decelerates or reaches the outer edge of the CSM.

[57]: Kurahashi et al. (2022), *High-Energy Extragalactic Neutrino Astrophysics*

[58]: Waxman et al. (2017), *Shock Breakout Theory*

These accelerated protons are expected to generate high-energy neutrinos, either by pp- or $p\gamma$-interactions (see Section 1.3.4). The resulting neutrinos' energies might range from GeV to hundreds of TeV, reaching their maximum flux shortly after shock breakout [59, 60].

[59]: Petropoulou et al. (2017), *Point-source and diffuse high-energy neutrino emission from Type IIn supernovae*

[60]: Murase (2018), *New prospects for detecting high-energy neutrinos from nearby supernovae*

The signature of such an event is a high-energy neutrino detected close to the optical peak of a supernova that shows signs of CSM interaction, i.e. the hydrogen lines in the spectrum have a narrow component.

1.4.2 Gamma-ray Bursts

Gamma-ray bursts (GRBs) were discovered by accident by the *Vela* satellites designed to monitor nuclear detonations as part of a partial test ban treaty between the US and the USSR in the 1960s. 16 bursts of high-energy photon events recorded by those satellites were published in 1973. This kicked off a flurry of research, as a terrestrial or solar origin of these gamma-ray bursts could be excluded based on light travel time differences allowing a rough triangulation of the signals [61].

[61]: Klebesadel et al. (1973), *Observations of Gamma-Ray Bursts of Cosmic Origin*

The gamma-ray light curves of these very bright and brief events last from a few milliseconds to tens, sometimes hundreds of seconds [62]. As more and more GRBs were registered, a bimodal population was recognized: Short GRBs (sGRBs), with an average duration of ∼0.3 s, and long GRBs (lGRBs) that typically last 10–20 s.

[62]: Vedrenne (2009), *Gamma-ray Bursts*

Short GRBs The most prevalent interpretation of these two populations is that short GRBs stem from the merger of two compact objects, e.g. neutron stars. This was hypothesized for a long time, and found spectacular confirmation with the association of the short gamma-ray burst *GRB170817A* detected 1.7 s after a gravitational wave signal was registered from the same region (*GW170817*), most likely stemming from the merger of two neutron stars [63].

[63]: Abbott et al. (2017), *GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral*

Long GRBs Meanwhile, long GRBs are hypothesized to be the signatures of jets launched by CCSNe [64]. The existence of such jets has been proposed already in the 1970s; see e.g. [65]. If these jets manage to punch through the stellar envelope and the surrounding CSM, one can detect them as long GRBs. Virtually all CCSNe that have been associated with a prior long GRB were Broad-Line Ic (Ic-BL) [66], i.e. collapses of highly stripped, massive stars with large ejecta velocities.

[64]: Galama et al. (1998), *An unusual supernova in the error box of the γ-ray burst of 25 April 1998*

[65]: Leblanc et al. (1970), *A Numerical Example of the Collapse of a Rotating Magnetized Star*

[66]: Pian et al. (2017), *Hydrogen-Poor Core-Collapse Supernovae*

Both short and long GRBs have been proposed as potential sources of high-energy neutrinos. For short GRBs, neutrinos are only attributable to a source if a GRB afterglow is detected, i.e. longer-lasting X-ray, optical and radio emission [67]. The event signature as accessible by IceCube and ZTF (see next two chapters) would be a prompt high-energy neutrino coincident with the gamma-ray signal, accompanied by optical emission quickly fading over the subsequent days.

[67]: Waxman et al. (1997), *High Energy Neutrinos from Cosmological Gamma-Ray Burst Fireballs*

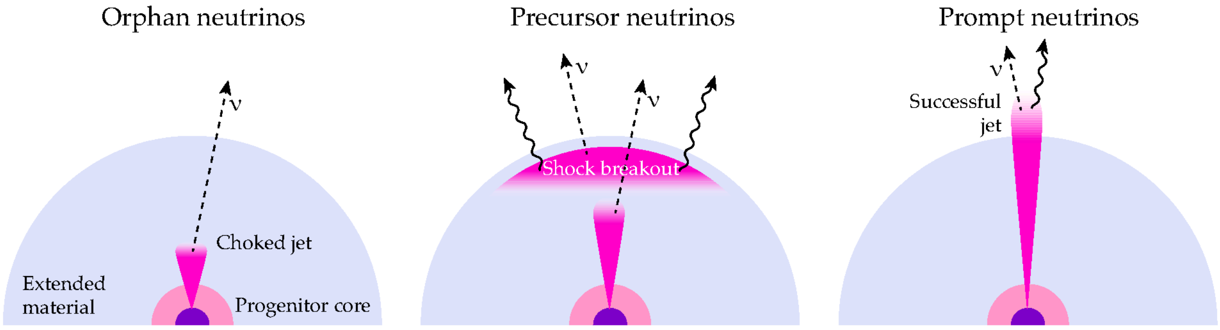

Figure 1.17: Concordance scenario for neutrinos from failed, semi-failed and successful jets launched by CCSNe. Left: The jet fails to punch through the environment, resulting in orphan neutrinos from the jet without any photons, as no GRB is produced. Middle: Here, the jet fails too, but triggers a shock breakout, resulting in a second wave of neutrinos, accompanied by a low-luminosity gamma-ray signal (low-luminosity GRB). Right: The jet launches successfully, generating prompt neutrinos accompanied by gamma-rays (long GRB). Adapted from [68].

Long GRBs on the other hand are also predicted to produce high-energy neutrinos [69]. One needs to distinguish three different scenarios, as illustrated in Fig. 1.17:

[69]: Mészáros et al. (2001), *TeV Neutrinos from Successful and Choked Gamma-Ray Bursts*

Choked jet (left panel) If the jet launched by the stellar collapse fails to penetrate the stellar and CSM environment and chokes, no light would is visible, only orphan high-energy neutrinos from the choked jet emerge, without an electromagnetic counterpart.

Choked jet with shock breakout (middle) If the choked jet manages to trigger a shock breakout (see Section 1.4.1), both the neutrinos from the choked jet (precursor neutrinos) and delayed high-energy neutrinos accompanied by a low-luminosity GRB and afterglow emission from the shock breakout are visible.

Successful jet (right) If the jet successfully launches, a long GRB accompanied by high-energy neutrinos is expected within minutes after the stellar collapse.

In all cases, further evolution will lead to an emerging supernova—most likely of Type Ic-BL—which will be detectable at the location of the neutrino in optical bands in the following weeks [68].

[68]: Senno et al. (2016), *Choked jets and low-luminosity gamma-ray bursts as hidden neutrino sources*

1.4.3 Active Galactic Nuclei

Another extreme environment are active galactic nuclei (AGN). At the center of most galaxies lies a supermassive black hole (SMBH). If the region around the SMBH is active, the system is considered an AGN. The key features of AGN explain what 'active' means: AGN are highly luminous, with bolometric luminosities up to 10^{48} erg s^{-1}. They vary on rapid timescales, which allows to infer that the emitting regions must be small (to not violate causality). Also, their luminosity function (the number of AGN per luminosity bin) evolves strongly with redshift. Lastly, AGN emit throughout the electromagnetic spectrum [70].

[70]: Padovani et al. (2017), *Active galactic nuclei: what's in a name?*

The history of AGN research began in the 1940s, when Grote Reber detected strong radio sources on the sky, among them Cygnus A, now known to be an AGN. Around the same time, Carl Seyfert discovered a type of galaxy with luminous nuclei [71]. In 1963 Maarten Schmidt discovered a point-like object (3C 273) with a redshift of 0.158, using the P200 on Mount Palomar (see Chapter 3). The large redshift ruled out that the source was a regular star. Rather, it was located in the nucleus of a galaxy with the same redshift. The remarkably large luminosity he inferred from that was about 100 times brighter than comparable galaxies [72].

[71]: Seyfert (1943), *Nuclear Emission in Spiral Nebulae.*

[72]: Schmidt (1963), *3C 273: A Star-Like Object with Large Red-Shift*

Today we assume that most galaxies contain a SMBH at their center. If the accretion rate, i.e. the rate at which the SMBH accumulates mass, is high, the galaxy appears active. We also assume that the myriad of different types of active galaxies (blazars, Seyfert I and II galaxies, radio-loud galaxies, etc.) are all appearances of intrinsically similar objects: They are merely a function of the viewing angle at which we look at the AGN system, plus the presence or absence of a jet launched by the system (and the accretion rate of the SMBH).

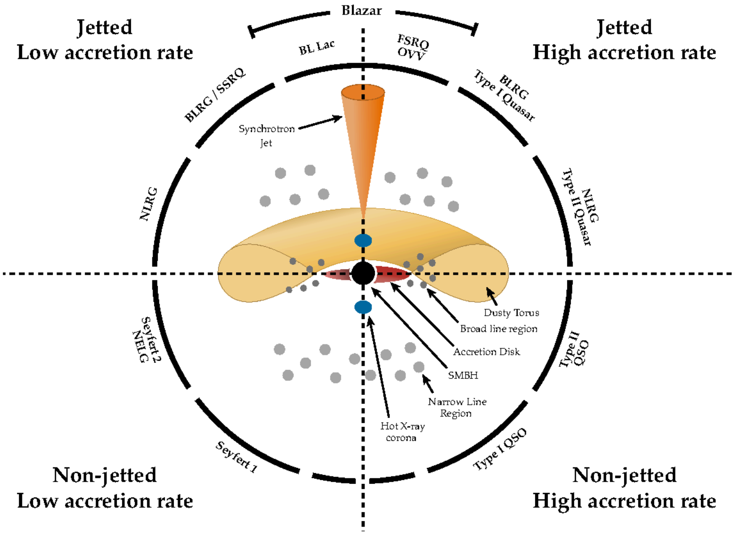

Figure 1.18: AGN unification scheme. Depending on the viewing angle towards the AGN system, different source classes emerge. Following the argumentation of [73], the upper (lower) segment label radio loud (radio quiet) was changed to jetted (non-jetted). Adapted from [74].

As can be seen in Fig. 1.18, these systems are thought to contain a SMBH at the center. The black hole is fed by a surrounding accretion disk of hot material, which in turn is enshrouded in a dusty torus.

Often, energetic jets of relativistic material are launched from the system, perpendicular to the accretion disk. When one looks directly into the jet, the system appears as blazar (top of 1.18). If one moves away from the jet direction, the AGN will in general appear as Type I AGN, displaying broad emission lines stemming from the so-called Broad-Line Region (BLR), lying at a distance of 0.01 − 1 pc to the SMBH. The BLR consists of gas clouds, and the Doppler-induced line-broadening suggests high velocities, resulting from the fact that the material is deep within the gravitational well of the SMBH.

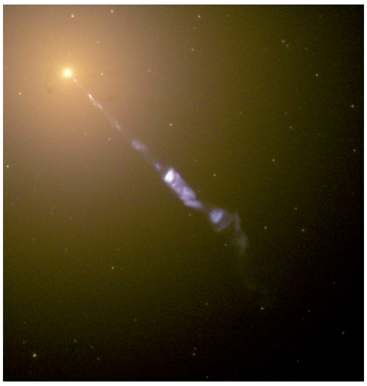

Figure 1.19: Hubble Space Telescope composite image of the jet launched by the AGN in *M87*, 17 Mpc away. The large size of the SMBH in *M87* made it the first target of the Event Horizon Telescope, resulting in the first direct image of a black hole ever taken [75]. Image credit: NASA / Hubble Heritage Team.

Moving yet further away from the jet, one gradually enters the regime of the Type II AGN: The BLR close to the core gets hidden behind the dust torus, and only the Narrow-Line Region (NLR) is still visible, as it is more distant from the core ($10^2 - 10^4$ pc). Here too, gas heated by the core radiates, but the gas density is lower compared to the BLR, and from the line-broadening one can infer much lower velocities of $\sim 500\,\mathrm{km\,s^{-1}}$ [76].

[76]: Beckmann et al. (2012), *Active Galactic Nuclei*

If a jet is present, the AGN strongly emits in radio wavelengths. This is caused by the synchrotron emission of highly relativistic electrons. These AGN were labeled ‘radio loud’, in contrast to ‘radio quiet’ AGN without a jet.

AGN have been proposed as sources of high-energy neutrinos as early as the 1970s [77]. The most promising site of high-energy neutrino production is the relativistic jet. In some cases such jets can be visually detected, like the prominent jet visible in *M87* (see Fig 1.19). Jets are thought to be highly collimated, extremely energetic plasma structures launched from the accretion disk. Only a minority of 10 − 20 % of observed AGN show evidence of such jets [78]. The intrinsic fraction of jetted AGN is considerably lower, with an estimate of 0.1 − 0.2 %.

[77]: Eichler (1979), *High-energy neutrino astronomy - A probe of galactic nuclei*

[78]: Padovani (2010), *The Microjansky and Nanojansky Radio Sky: Source Population and Multiwavelength Properties*

When the electrons and protons contained in the jet plasma are clumped together in relativistically moving ‘blobs’ and are spinning in magnetic fields, they emit synchrotron radiation. This is the usual explanation for the first hump in a typical blazar spectral energy distribution (SED) [79].

[79]: Combes (2021), *Active Galactic Nuclei*

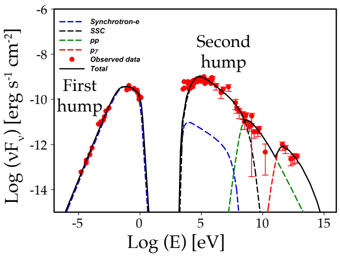

Figure 1.20: Typical AGN spectrum, showing two humps. Adapted from [80].

The explanation for the second hump is still debated. One scenario is that the synchrotron photons themselves form a target field for their ‘parent’ electrons and protons. The electrons and protons can interact with these photons via inverse Compton Scattering (CS), transferring energy to the photons. This process is called Synchrotron Self-Compton (SSC) [26], giving rise to the second hump. This was dubbed the Leptonic Model, and due to the absence of accelerated hadrons like protons, it does not predict the presence of high-energy neutrinos.

An alternative model to explain the second hump is the Hadronic Model. Here, in addition to leptons, the jet also contains a sizable fraction of hadrons [81]. So not only the leptonic processes take place (synchrotron radiation, inverse Compton scattering), but also the protons generate synchrotron radiation. Additionally, the hadrons can interact with target photons in the vicinity, producing charged and neutral pions. These decay as detailed in Section 1.3.4, producing high-energy neutrinos in the process. Therefore, the detection of high-energy neutrinos from blazars is considered a smoking gun signature for hadronic acceleration processes [82].

[81]: Reimer (2012), *On the Physics of Hadronic Blazar Emission Models*

[82]: Gao et al. (2018), *Modelling the coincident observation of a high-energy neutrino and a bright blazar flare*

As the AGN's optical brightness is driven by the synchrotron emission, it only traces the photon target field density for potential hadronic interactions, but not the proton luminosity itself.

1.4.4 Tidal Disruption Events

The last class of potential high-energy neutrino sources discussed here are Tidal Disruption Events (TDEs). A TDE marks the end of a very unlucky star. When a star happens to approach the SMBH at the galactic center too closely, the resulting gravitational forces might overcome the star's self-gravity, tidally disrupting and destroying the star in the process [83]. Roughly half the mass of the star is accreted around the black hole, and the bright electromagnetic flare created by this can shine for months. Nevertheless, the exact mechanisms at play causing this emission are still widely discussed [84].

[83]: Rees (1988), *Tidal disruption of stars by black holes of* $10^6 - 10^8$ *solar masses in nearby galaxies*

[84]: Gezari (2021), *Tidal Disruption Events*

TDEs have already been predicted in the 1970s [85], and the first few of them were discovered in the 1990s and early 2000s in X-ray wavelengths. Optically detected TDEs arrived in the last decade [86], with steadily rising numbers in the last years, mainly driven by optical all-sky survey telescopes like the Zwicky Transient Facility (ZTF, see Chapter 3). Fig. 1.21 shows the number of detected TDEs since 1995. As of 2022, roughly 100 TDEs have been discovered in total, with more than 30 detected during the 2.6 years of ZTF Phase I alone [87].

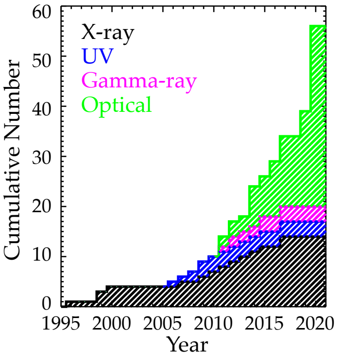

Figure 1.21: Cumulative number of TDE detections, with the color encoding the discovery wavelength. The relative increase in the detection rate is driven by ZTF. Adopted from [84].

[85]: Hills (1975), *Possible power source of Seyfert galaxies and QSOs*

[86]: Velzen et al. (2011), *Optical Discovery of Probable Stellar Tidal Disruption Flares*

[87]: Hammerstein et al. (2022), *The Final Season Reimagined: 30 Tidal Disruption Events from the ZTF-I Survey*

TDEs normally emit a continuous spectrum that can be well-approximated by a blackbody (see e.g. [88] for a spectrum of *PS1-10jh*, the prototypical optical/UV TDE). Peculiarly, there seems to be a somewhat bimodal distribution of blackbody temperatures: One optical/UV population, and a second population peaking in the X-ray, better described by blackbodies of higher temperature [84].

[88]: Gezari et al. (2015), *PS1-10jh Continues to Follow the Fallback Accretion Rate of a Tidally Disrupted Star*

As the accretion disk hypothesized to form from the infalling material after the disruption will most likely be very hot, it is unclear what the cause of the optical/UV emission is. Furthermore, the blackbody radii inferred from the optical/UV emission exceed the size of the freshly formed accretion disk by $1 - 2$ orders of magnitude.

To date, multiple models have been brought forward to motivate the optical/UV emission, such as semi-relativistic outflows or winds, or diffusive shock acceleration stemming from the tidal stream intersecting itself [84].

If the material from the disrupted star is circularized rapidly, one interesting approach to unify the two populations is a viewing angle dependent model, akin to the unified AGN model presented in Section 1.4.3.

In this model, which is shown in Fig. 1.22, TDEs appear X-ray bright when one looks into the direction of a shocked jet perpendicular to the accretion disk. Optical/UV TDEs are systems viewed more edge-on, where X-rays are obscured, and emission stems mainly from X-rays reprocessed in the outer disk or in outflows [89]. Intermediate viewing angles will produce a mixture of both signals.

[89]: Hayasaki (2021), *Neutrinos from tidal disruption events*

About 1 % of TDEs are expected to launch relativistic jets, like e.g. the recently discovered *AT2022cmc* [90]. Such jets, denoted as (1) in Fig. 1.22,

[90]: Andreoni et al. (2022), *A very luminous jet from the disruption of a star by a massive black hole*

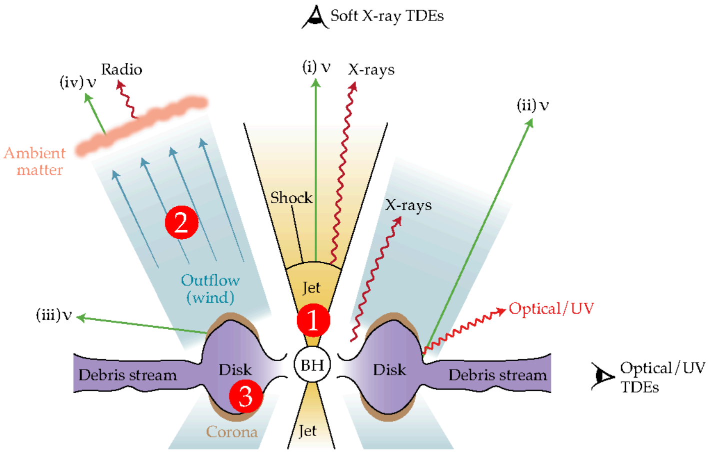

Figure 1.22: TDE unification model. The peak wavelength visible is a function of the viewing angle. Seen from above, the X-rays are dominating, while viewed edge-on, the TDE appears in the optical/UV. Possible sites of high-energy neutrino production are shown as numbered red circles. Adopted from [89].

as well as the possible winds/outflows (2) or a potentially present disk corona (3), have all been proposed as production sites of high-energy neutrinos, shown as green lines (see e.g. [91] for the non-jet production sites, and [92] for a jet scenario).

[91]: Murase et al. (2020), *High-energy Neutrino and Gamma-Ray Emission from Tidal Disruption Events*

[92]: Liu et al. (2020), *Neutrino emission from an off-axis jet driven by the tidal disruption event AT2019dsg*

The timescales involved in the neutrino production are poorly constrained, as the systems are not very well understood yet. It is to be expected though that neutrino production would not predate the TDE and optical emission from the event.

1.5 Established Counterparts and Limits

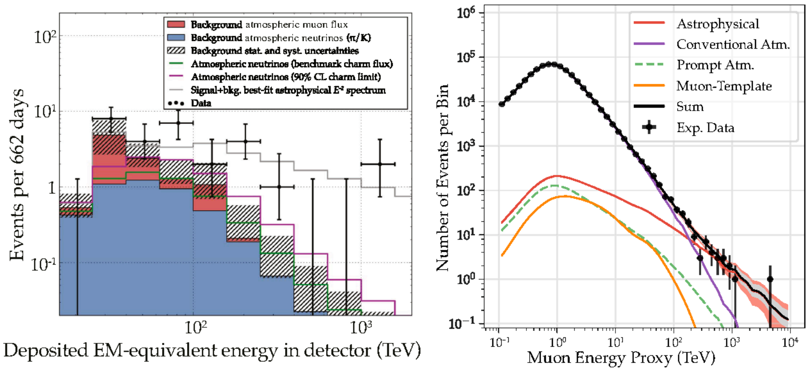

Figure 1.23: Left: Energy spectrum of the first astrophysical neutrinos, as detected by IceCube in 2013. The black points are the neutrinos measured, binned in energy deposited in the detector (a lower bound on the neutrino energy). The expected background rate of atmospheric muons (neutrinos) is shown in red (blue), and the gray line shows the best-fit astrophysical spectrum (E^{-2} plus background). Right: Latest measurement of astrophysical neutrinos from a selection of track-like events (see Section 2.3), including 9 years of IceCube data. Also here, the binned measurements are shown in black. Adapted from [93, 94].

In 2013, IceCube, a cubic-kilometer scale Cherenkov detector located within the Antarctic ice at the South Pole (see next Chapter 2 for details) first detected two likely astrophysical PeV neutrinos at a significance of $2.8\,\sigma$ [95]. Later that year, IceCube published more data, detecting high-energy astrophysical neutrinos at a confidence level of $4\,\sigma$ [93].

[95]: Aartsen et al. (2013), *First Observation of PeV-Energy Neutrinos with IceCube*

[93]: Aartsen et al. (2013), *Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector*

Fig. 1.23 shows the first energy spectrum of high-energy neutrinos published by IceCube (left), as well as a more recent measurement from 9 years of IceCube data (right). At energies of roughly 100 TeV, the atmospheric background begins to recede, revealing a flux of astrophysical neutrinos that follows a power-law spectrum with a spectral index $\gamma \approx 2.4$ [94].

[94]: Abbasi et al. (2022), *Improved Characterization of the Astrophysical Muon–neutrino Flux with 9.5 Years of IceCube Data*

The flux is largely distributed isotropically over the sky, which stipulates an extra-galactic origin for the majority of high-energy neutrinos. The majority of this flux is still unaccounted for. Nevertheless, there are some prominent source candidates found since 2013. These, as well as upper limits on other hypothesized source populations, will be discussed next.

1.5.1 AGN Counterparts and Limits

Blazar *TXS 0506+056*

In 2017, IceCube identified the blazar *TXS 0506+056*, at the time flaring in gamma rays, as the probable source of high-energy neutrino *IC170922A* [96, 97]. The neutrino had an estimated energy of 290 TeV, and a best-fit sky location separating it 0.1° from the blazar, see left plot in Fig. 1.24.

[96]: Kopper et al. (2017), *IceCube-170922A - IceCube observation of a high-energy neutrino candidate event.*
[97]: Aartsen et al. (2018), *Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A*

The blazar had a redshift of $z = 0.337$ [98], and a search for additional excess neutrinos from its sky location in 9.5 years of IceCube data found another episode of a $3.5\,\sigma$ excess of ~13.5 neutrinos between September 2014 and March 2015, originating from the direction of *TXS 0506+056*. However, during this period no significant gamma-ray flare was detected from *TXS 0506+056* [99].

[98]: Paiano et al. (2018), *The Redshift of the BL Lac Object TXS 0506+056*

[99]: Aartsen et al. (2018), *Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert*

TXS 0506+056 (in its gamma-ray flaring state) has been shown to be able to produce a flux of high-energy neutrinos compatible with the IceCube detection [82]. An exemplary SED is shown in the right plot of Fig. 1.24. Here, a 'classical' hadronic scenario that would attribute the second hump entirely to decaying pions, was discarded to not overshoot the X-ray flux. Instead, the second hump was interpreted as a combination of leptonic and hadronic processes, with the former being the major contributor and the latter as strong as the X-ray bounds allowed.

[82]: Gao et al. (2018), *Modelling the coincident observation of a high-energy neutrino and a bright blazar flare*

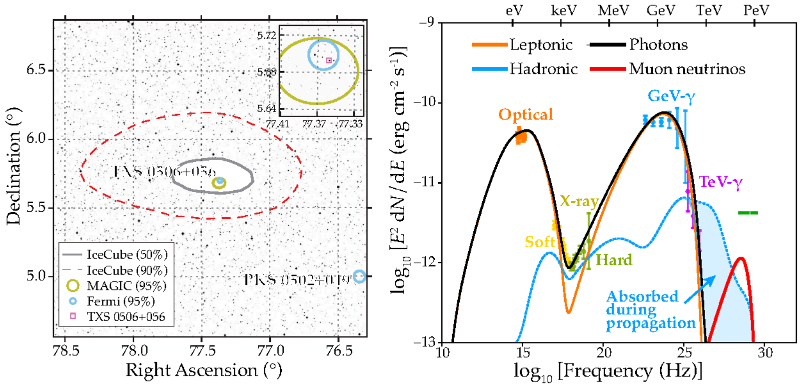

Figure 1.24: Left: Localization of the flaring blazar *TXS 0506+056*, found in coincidence with high-energy neutrino *IC170922A*. The 50 % (90 %) localization contours of the neutrino location are shown in grey (dashed red). Right: Hybrid lepto-hadronic emission model to explain the SED, including the detected high-energy neutrino. The leptonic component is shown as orange line, while the hadronic component is shown in blue. The resulting combined photon spectrum is displayed in black, and the muon neutrino spectrum in red. Adopted from [82, 97].

This lepto-hadronic model predicted a 14 % probability of detecting a neutrino from the source without violating the bounds imposed by observations [82]. If one takes into account the large Eddington Bias [100] expected for a single neutrino detection, this is still plausible. A consequence of this bias is that the neutrino flux of a single detection source could be systematically overestimated if that source is part of a large population emitting just below the IceCube detection limit, with occasional statistical overfluctuations.

[100]: Strotjohann et al. (2019), *Eddington bias for cosmic neutrino sources*

However, producing the 2014–2015 neutrino flare without violating the gamma-ray limits proved to be a much harder challenge. No lepto-hadronic model[5] able to produce more than 2–5 neutrino events during the time period could be found [101] without violating the observational constraints. It remains a challenge to reconcile both periods: The neutrino flux from the 2014–2015 period without a simultaneous gamma-ray flare on the one hand, and the single neutrino from 2017 accompanied by a gamma-ray flare on the other hand.

[5] To keep consistency with modeling the gamma-ray flaring state.

[101]: Rodrigues et al. (2019), *Lepto-hadronic Blazar Models Applied to the 2014–2015 Flare of TXS 0506+056*

Type II AGN *NGC 1068*

The second AGN-neutrino association was *NGC 1068*. This active galaxy is located relatively nearby, with a distance of ~14 Mpc [102]. It was classified as Type II AGN, i.e. an AGN without signs of a jet and viewed relatively edge-on, with dust obscuring the SMBH and the BLR (see Section 1.4.3) [103]. It has an exceptionally high rate of star formation [104], and hosts outflows [105]. Due to these features, it had already been proposed as a production site of high-energy neutrinos in the 1970s [106].

[102]: R. Abbasi et al. (2022), *Evidence for neutrino emission from the nearby active galaxy NGC 1068*

[103]: Rosas et al. (2022), *Thermal imaging of dust hiding the black hole in NGC 1068*

[104]: Eichmann et al. (2016), *The Radio-Gamma Correlation in Starburst Galaxies*

[105]: Cecil et al. (1990), *Imaging spectrophotometry of ionized gas in NGC 1068. I - Kinematics of the narrow-line region*

[106]: Silberberg et al. (1979), *Neutrinos as a Probe for the Nature of and Processes in Active Galactic Nuclei*

In 2022, an archival study comprising IceCube data from 2011 to 2020 found an excess of 79^{+22}_{-20} neutrinos with TeV energies from the location of *NGC 1068*. This results in a neutrino luminosity between 1.5 and 15 TeV of $L_\nu = 2.9^{+1.1}_{-1.1} \times 10^{42}\,\mathrm{erg\,s^{-1}}$, exceeding the equivalent gamma-ray luminosity between 100 MeV and 100 GeV by a factor of 18 [102].

[102]: R. Abbasi et al. (2022), *Evidence for neutrino emission from the nearby active galaxy NGC 1068*

NGC 1068 and *TXS 0506+056* each contribute ~1 % of the overall diffuse flux of astrophysical neutrinos measured by IceCube in their respective energy range. So far, it remains unclear if the diffuse flux is mainly composed of bright and nearby sources akin to *NGC 1068*, a large population of faint sources with high redshifts ($z \geq 1$), or a mixture of both. Given the big difference in distance between *NGC 1068* and *TXS 0506+056*—the latter is about 100 times further away—and the differences in their respective spectra, it seems plausible to assume at least two distinct AGN source populations [102].

Gamma-ray Blazar Limits

However, there are several constraints on the contribution from AGN. A stacking analysis from 2019 investigated 1301 gamma-ray blazars from 3FHL [107], the third catalog of hard gamma-ray sources issued by the *Fermi* Large Area Telescope (LAT) [108]. After stacking and searching for spatial correlations between through-going high-energy muon tracks from the northern hemisphere, no excess of neutrinos was detected [109]. Assuming a spectral source index of $\gamma = 2$, an upper limit on the contribution of blazars to the high-energy neutrino flux between 119 TeV and 4.9 PeV was found: Not more than 17 % of the diffuse flux can be attributed to these sources [109].

[107]: Ajello et al. (2017), *3FHL: The Third Catalog of Hard Fermi-LAT Sources*

[108]: Atwood et al. (2009), *The Large Area Telescope on the Fermi Gamma-Ray Space Telescope*

[109]: Huber (2019), *Searches for steady neutrino emission from 3FHL blazars using eight years of IceCube data from the Northern hemisphere*

[109]: Huber (2019), *Searches for steady neutrino emission from 3FHL blazars using eight years of IceCube data from the Northern hemisphere*

MeV Blazar Limits

The contribution of MeV blazars was also tested. 137 blazars from the *Fermi* Low Energy Catalog (1FLE) [110], detected below 100 MeV, were stacked. The result was compatible with a non-detection. The upper limit derived from this—assuming a spectral index tracing the diffuse neutrino

[110]: Principe et al. (2018), *The first catalog of Fermi-LAT sources below 100 MeV*

spectral index $\gamma = 2.37$ and evaluated in an energy range between 30 and 100 MeV—was ~1 % of the IceCube $\nu_\mu + \bar{\nu_\mu}$ flux [111].

[111]: Abbasi et al. (2022), *Search for Astrophysical Neutrinos from 1FLE Blazars with IceCube*

Jetted AGN Limits

The correlation between high-energy alert neutrinos above 200 TeV with jetted (radio bright) AGN was also tested. The 3388 jetted AGN were selected from the Radio Fundamental Catalog[6] by requiring X-band (8−12 GHz) flux densities above 0.15 Jy. These were correlated with 56 published IceCube events with directional uncertainties below 10 deg^2. With these selections, a significant correlation was found, with a p-value of 0.2 % [112].

[6] `https://astrogeo.org/rfc`

[112]: Plavin et al. (2020), *Observational Evidence for the Origin of High-energy Neutrinos in Parsec-scale Nuclei of Radio-bright Active Galaxies*

However, a more recent study correlating this set of jetted AGN with the IceCube diffuse flux data comprising 10 years of muon tracks found no correlation [113]. The unbinned maximum-likelihood-ratio method the study employed gave no significant correlation between the neutrinos and the jetted AGN, resulting in an upper limit on the overall contribution to the TeV–PeV neutrino flux of 30 %.

[113]: Abbasi et al. (2023), *Search for correlations of high-energy neutrinos detected in IceCube with radio-bright AGN and gamma-ray emission from blazars*

Furthermore, a recent study correlating blazars in the southern hemisphere (as the earth's core gets opaque for neutrinos of the highest energies) from the Roma-BZCat Data Release 5 blazar catalog (5BZCat [114]) with the 7-year all-sky map from IceCube found a correlation between blazars and high-energy neutrinos with a significance of 4.6 σ. However, a follow-up study by IceCube was able to replicate the result, but when employing a 10-year neutrino all-sky map, the correlation vanished [115]. This rather hints at a statistical fluctuation responsible for the 7-year map correlation.

[114]: Massaro et al. (2016), *VizieR Online Data Catalog: The Roma BZCAT - 5th edition (Massaro+, 2015)*

[115]: Bellenghi et al. (2023), *Correlating high-energy IceCube neutrinos with 5BZCAT blazars and RFC sources*

Non-jetted AGN Contribution

The picture changes significantly when looking at non-jetted AGN: A recent study found a 2.6 σ excess when correlating an infrared-selected subpopulation of non-jetted AGN with the astrophysical neutrino flux [116]. The accretion disk luminosity of the AGN was estimated with their soft X-ray flux, and the source neutrino flux was weighted by that. The AGN population was selected based on the AGN radio emission and their infrared colors, and a 2.6 σ excess (post trial) was obtained for this infrared-selected sample of 32,249 AGN.

[116]: Abbasi et al. (2022), *Search for neutrino emission from cores of active galactic nuclei*

If this excess is interpreted as constituting a physical signal, it contributes 10^{+5}_{-4} % of the diffuse flux at 100 TeV, see Fig. 1.25. Correcting for completeness, 27−100 % of the observed 100 TeV neutrinos could stem from particle acceleration within the accretion disks or coronae of AGN [116].

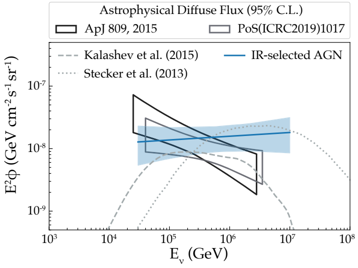

Figure 1.25: Contribution of non-jetted AGN to the diffuse IceCube neutrino flux. The best-fit power law muon neutrino flux is shown in blue, corrected for completeness. Adapted from [116].

1.5.2 GRB Limits

GRB limits can be drawn from a search for a correlation between 807 GRBs detected during a period of three years with IceCube high-energy neutrinos during that time. The search was constrained to prompt emission (see Section 1.4.2) from GRBs, excluding precursor events or GRB afterglows. This was achieved by looking for cascade events in the detector (see Section 2.3), which are created by neutrinos of all flavors. The neutrinos were required to be detected within the photon emission time as reported by the gamma-ray satellite for each GRB, the so-called prompt window. These prompt windows usually range from a few tens of seconds to a few minutes. Short windows are favorable for coincident searches, as the tight time-constraint strongly suppresses background events. Six events were found to be time-correlated with neutrinos, but also consistent with background. If they are interpreted as genuine signal events, GRBs contribute less than 1 % to the diffuse neutrino flux [117].

[117]: Aartsen et al. (2016), *An All-Sky Search for Three Flavors of Neutrinos from Gamma-Ray Bursts with the IceCube Neutrino Observatory*

An additional search extended the time window considered in order to search for a correlation between a subset of GRBs contained in `GRBweb`[7] and over 7 years of IceCube data. Also with this expanded time window of 1000 s sensitive to parts of the GRB afterglow, no time correlation could be found, with an upper limit on the contribution to the diffuse flux of 24 % [118].

[7] https://user-web.icecube.wisc.edu/~grbweb_public/

[118]: Abbasi et al. (2022), *Searches for Neutrinos from Gamma-Ray Bursts Using the IceCube Neutrino Observatory*

1.5.3 Supernova Limits

In 2023, a study [119] was published correlating 1040 core-collapse SNe with 7 years worth of IceCube neutrino events. The SN data was obtained from the Weizmann Interactive Supernova Data Repository `WiseREP`[8] [120] and the Open Supernova Catalog [121].

[119]: Abbasi et al. (2023), *Constraining High-energy Neutrino Emission from Supernovae with IceCube*

[8] https://wiserep.org/

[120]: Yaron et al. (2012), *WISeREP – An Interactive Supernova Data Repository*

[121]: Guillochon et al. (2017), *An Open Catalog for Supernova Data*

The study looked at correlations between individual supernovae, as well as the stacked full sample. Both methods yielded results that were compatible with the null hypothesis. In the study the SN neutrino energy spectrum was assumed to have a spectral index $\gamma = 2.5$. Within the neutrino energy range of 1 to 100 TeV, the different types of SNe cannot contribute more than the following: SN IIP: 59.9 %, SN IIn: 33.9 % and stripped-envelope SNe (Ibc and IIb, see Section 1.2.2): 14.6 %, assuming a choked-jet emission model [119].

The first two classes were tested with a set of emission time windows, ranging between 100 and 1000 days after the first detection, while the stripped-envelope SNe were tested for choked-jet emission, in which the neutrinos would predate the optical emission (see Section 1.4.2). In this case, the time window started 20 days prior to, and ended with the first optical detection. Neither SNe IIn nor choked jet neutrino emission in stripped-envelope SNe can dominate the diffuse IceCube neutrino flux, while SNe IIP could still be dominant [119].

1.5.4 TDE Limits

Lastly, there is a limit from stacking TDEs. A search comprising 13 non-jetted and 3 jetted TDEs, correlating with 9.5 years of IceCube muon neutrino data, found no significant excess. Under the assumption of constant TDE neutrino luminosity, upper limits for their contribution to the diffuse were derived. These constrain jetted TDEs to less than ~1 %, and non-jetted TDEs to less than 26 % of the diffuse flux detected by IceCube [122].

[122]: Stein (2019), *Search for Neutrinos from Populations of Optical Transients*

1.6 Conclusion

The neutrino has been proposed, discovered and studied intensively. The detection of solar neutrinos, supernova neutrinos and finally the detection of a flux of high-energy neutrinos has firmly established the field of neutrino astronomy.

Nevertheless, the origin of the majority of the cosmic high-energy neutrino flux remains unclear. As can be seen in Fig. 1.26, no dominant source class has yet been established, though there are strong hints that non-jetted AGN contribute significantly to the diffuse neutrino flux. It is very much possible that the diffuse flux is composed of multiple source classes, each subdominant. The most stringent limits existing are those on prompt GRB emission and jetted TDEs, given their low rate density.

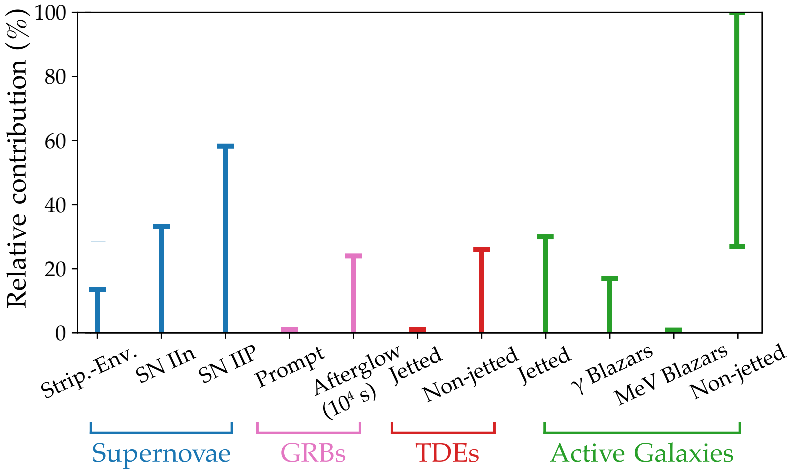

Figure 1.26: Potential source contributions to the diffuse IceCube high-energy neutrino flux. The values shown are mainly taken from [123], with some updates by the author. All numbers are discussed in the main text. Stripped-envelope SNe are denoted 'Strip.-Env.' and a choked-jet emission model is assumed for these. Note that the unclear connection between the electromagnetic and expected neutrino flux, as well as uncertainties regarding the exact shape of the neutrino spectrum add significant uncertainties to these estimates. Adapted from a figure by Foteini Oikonomou.

The absence of significant clustering of neutrinos, i.e. the paucity of point sources in the data, disfavors the hypothesis that rare and luminous objects are responsible for the majority of the flux [123]. It is entirely possible that the bulk flux stems from numerous faint objects, which could render establishing a dominant source class a challenging task.

[123]: Guépin et al. (2022), *High-energy neutrino transients and the future of multi-messenger astronomy*

The flux could be fairly equally shared by AGN, TDEs and other sources, including blazars, as [124] stipulate. The authors of that study estimated the contributions from source classes that have a known association (blazars: *TXS 0506+056*, AGN: *NGC 1068* and TDEs: *AT2019dsg*) versus all other classes. They included statistical detection uncertainties accounting for the varying neutrino luminosity within source classes, as well as the redshift evolution of the number density and uncertainties of the detection process. The individual contributions are shown in Fig. 1.27, with the flux shared fairly equally between AGN, TDEs and other events, plus a noteworthy subdominant contribution by blazars.

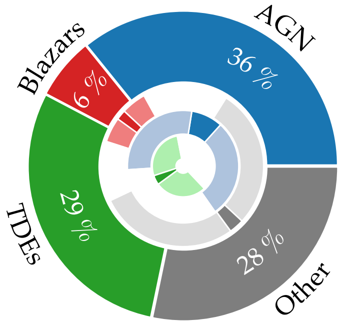

Figure 1.27: Pie chart of the contribution of known neutrino source classes as well as 'other', comprising all source classes without association (main circle). The inner charts show the minimum (dark) and maximum (light) contributions within the 90 % credible regions. Adapted from [124].

Programs trying to establish a connection between individual high-energy neutrinos and sources within their localization are one good instrument in the available toolbox to solve the origin question at least partially. Such programs have the advantage of being time sensitive and allowing for follow-up observations necessary to classify ambiguous transients, contrary to archival studies. Chapter 4 will present one such program, a dedicated optical follow-up to high-energy IceCube neutrinos.

[124]: Bartos et al. (2021), *The IceCube Pie Chart: Relative Source Contributions to the Cosmic Neutrino Flux*

But first one needs to now how these two instruments work. The next two chapters will introduce the neutrino observatory IceCube (see Chapter 2) used in this study, as well as the optical facility (ZTF, see Chapter 3).

2 The IceCube Detector

One of the two most relevant instruments for this thesis is the IceCube Neutrino Observatory, a neutrino detector located at the geographic South Pole. It is the successor to the Antarctic Muon And Neutrino Detector Array (AMANDA) at the same location [125, 126], and has been operational for about a decade.

The IceCube Laboratory, part of the IceCube Neutrino Observatory at the South Pole. Image credit: IceCube/NSF.

[125]: Andrés et al. (1999), *The AMANDA neutrino telescope*
[126]: Andrés et al. (2000), *The AMANDA neutrino telescope: principle of operation and first results*

The basic operational principle of IceCube (and already of AMANDA) is the detection of Cherenkov light within the Antarctic ice. When charged secondary particles created by neutrino interactions travel through the ice with an energy high enough, their speed can exceed the phase velocity of light in ice, and they start to emit Cherenkov radiation. The detector consists of 5160 individual Digital Optical Modules (DOMs), buried deep in the ice, dedicated to detect this Cherenkov radiation.

2.1 Cherenkov Radiation

Cherenkov radiation was first detected in 1934 by Soviet scientist Pavel Cherenkov [127]. It occurs when charged particles travel within a medium with a velocity exceeding the speed of light in that very medium. The refractive index in a medium is defined as $n = \frac{c_0}{c_m}$, where c_0 is the speed of light in vacuum and c_m is the phase velocity of light in that medium. Note that the phase velocity of light in a medium can exceed c_0, so $n < 1$ is possible.

[127]: Cherenkov (1934), *Visible emission of clean liquids by action of gamma radiation*

When charged particles cross an electrically neutral dielectric medium, atoms along the particle's path are briefly polarized. When they relax back to the ground state, the atoms emit electromagnetic radiation.

For non-relativistic particles, this radiation destructively interferes with itself, canceling out all signals (left panel of Fig. 2.1). But if the particle is traveling faster than the speed of light within the medium c_m, this destructive interference does not happen. Rather, a cone-shaped wavefront gets created (right panel of Fig. 2.1). This wavefront constitutes Cherenkov radiation.

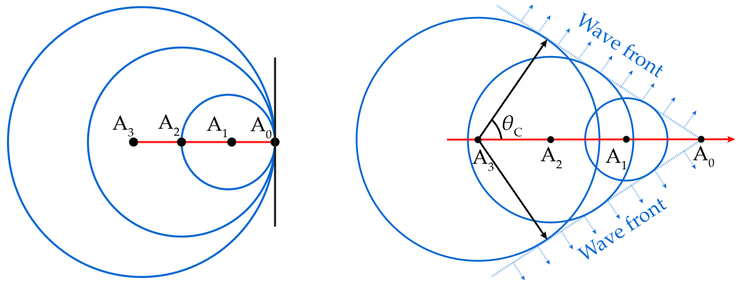

βn = 1
No Cherenkov radiation is emitted

βn > 1
Cherenkov radiation is emitted

Figure 2.1: The principle of Cherenkov radiation. In the left diagram, a particle travels through a medium along the red line and is not moving faster than light within that medium. All radiation from different points in space and time (A_3 to A_0, radiation shown as blue circles) is cancelled out by destructive interference (all circles are contained within the first one on the left). Right diagram: The particle does emit Cherenkov radiation, as the different circles form a mutual, cone-shaped wavefront, radiating away at the Cherenkov angle θ_C. Adapted from [128].

If the particle has speed $v = \beta c_0$, the angle θ between the particle trajectory and the direction of the Cherenkov radiation can be calculated as [128]

[128]: L'Annunziata (2020), *Handbook of Radioactivity Analysis. Volume 1: Radiation Physics and Detectors*

$$\cos\theta = \frac{\beta}{n}. \tag{2.1}$$

If the medium is ice, to first order the refractive index n is ≈ 1.31.[1] A secondary muon traveling through the ice at $0.999\,c_0$ will therefore emit Cherenkov light at an angle of $\theta = \cos^{-1}\left(\frac{0.999}{1.31}\right) \approx 40°$.

Cherenkov radiation has a smooth spectrum, with a relative intensity roughly proportional to the frequency. The refractive index of a medium also depends on the frequency, dropping below 1 in the X-ray [130]. From this it follows that Cherenkov radiation appears blue to the human eye (the high-frequency part dominates) and its intensity peaks in the Ultra Violet (UV), before it sharply drops off in the X-ray regime [129], see Fig. 2.2.

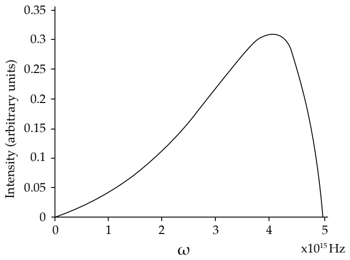

Figure 2.2: Cherenkov spectrum for a particle with $v = 0.8\,c_0$ in water. The intensity peaks around 4×10^{15} Hz, corresponding to a wavelength of 75 nm, lying at the high-frequency end of the UV spectrum. Adapted from [129].

[130]: Warren et al. (2008), *Optical constants of ice from the ultraviolet to the microwave: A revised compilation*

[129]: Filöp et al. (1992), *Cherenkov radiation spectrum*

2.2 Instrumentation

IceCube detects neutrinos by observing the optical and UV part of their secondary particle Cherenkov spectrum with 5160 DOMs in the ice. To understand how this is done, one first needs to look at the working principle of a Photomultiplier Tube (PMT), the basic instrument within the DOMs.

2.2.1 Photomultiplier Tubes

PMTs are devices used to detect very faint light signals by amplifying them. They consist of vacuum tubes and were successfully realized for the first time in the 1930s [131].

[131]: Iams et al. (1935), *The Secondary Emission Phototube*

As one can see in Fig. 2.3, there are three principal components: a cathode, several dynodes and an anode. When photons hit the cathode, they can release electrons via the photoelectric effect [133]. These photoelectrons are then accelerated (towards the right side in Fig. 2.3) by an electric field within the tube. This field is generated by applying a high voltage between the cathode and the anode.

[133]: Einstein (1905), *Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt*

---

[1] This assumption is rather crude. The n of Antarctic glacial ice depends e.g. on depth; a fact we will come back to later when discussing directional reconstruction of high-energy neutrinos.

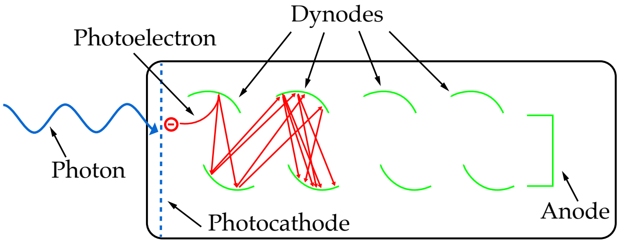

Figure 2.3: A photomultiplier tube. The incident photon creates a photoelectron on the left, which gets multiplied at each dynode along the way to the anode on the left, creating an avalanche of electrons. Adapted from [132].

To amplify the signal, a number of dynodes are placed between cathode and anode. These are additional electrodes with subsequently higher voltages. When the photoelectron hits the first dynode, a number of secondary electrons are generated, which are then accelerated towards the next dynode by the electric field. This process repeats for every dynode, generating an electron avalanche exponentially amplifying the original single photoelectron signal. The number of secondary electrons hitting the anode is proportional to the number of incident photons, resulting in a linear detector response (as long as the detector stays below its saturation limit) [134].

[134]: Wright (2017), *The Photomultiplier Handbook*

IceCube uses PMTs made by Hamamatsu Photonics (R7081–02) and which are sensitive to photons between 300 and 650 nm. They have a quantum efficiency (QE) at 390 nm of 25 %, are operated with a voltage of 1500 V and have a gain (electron multiplication factor) of 10^7. The photon-sensitive surface area is typically 530 cm^2 [135].

[135]: Abbasi et al. (2010), *Calibration and characterization of the IceCube photomultiplier tube*

2.2.2 The Digital Optical Module

The individual IceCube PMTs for detecting the Cherenkov radiation are enclosed in the DOMs. Each DOM consists of a pressure-resistant glass sphere, several controller boards and the PMT, facing downward (see Fig. 2.4). The glass sphere can withstand long-term pressure of 250 bar. The optical transmission of the spheres was measured to be 93 % at 400 nm, decreasing to 10 % at 315 nm.

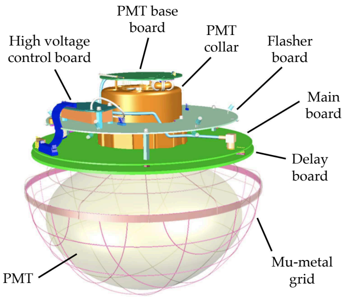

Figure 2.4: The IceCube DOM seen from the side. The detecting side of the PMT is facing downwards, with the main board and the PMT base board on top. From [136].

The circular main board hosts data acquisition and control systems, as well as units for communication and a power converter. Another board, the PMT base board, interfaces with the PMT, while additional boards delay the PMT signals and generate the high voltage current powering the PMT. Additionally, there is a so-called Flasher Board that controls Light-Emitting Diodes (LEDs). These are used to generate light flashes which can be received by neighboring DOMs for calibration purposes [136].

[136]: Aartsen et al. (2017), *The IceCube Neutrino Observatory: instrumentation and online systems*

Because of data storage restrictions, the DOMs only record and digitize the photon signal after several trigger criteria have been met. The digitized voltage over time detected by the PMTs is called a waveform, and combined with a timestamp it comprises the basic datum in IceCube, a hit.

To fully record the waveform after a hit, the signal needs to be stored in a buffer. This is realized with the delay board, which routes the analog PMT signal through a 10 m long, serpentine copper trace to delay it by 75 ns.

When a hit is detected, the DOM sends a tag signal to the neighbor and next-to-nearest neighbor DOMs. If no neighbors detect a hit, the isolated hit will only contain the timestamp, the amplitude and charge information extracted from the waveform. When at least two neighboring DOMs also detect a hit within 0.25 µs, a Local Coincidence (LC) is triggered. If the signal passes a threshold of 0.25 photoelectrons, the recorded waveforms are digitized and appended to the hits tagged as LC [136].

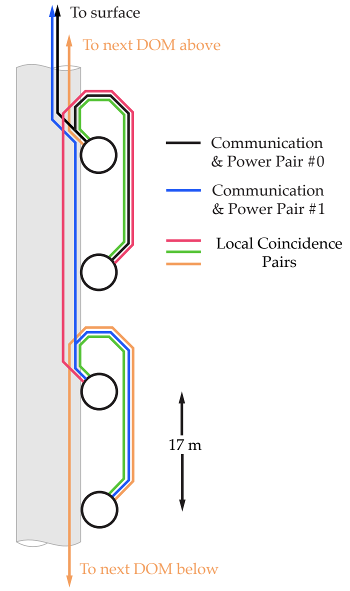

Figure 2.5: Connection scheme for four IceCube DOMs along one string. Pairs of DOMs share one twisted-pair cable. Also, each DOM is directly connected to its direct neighbors above and below. Adapted from [136].

The digitization of the PMT waveform is done with the Analog Transient Waveform Digitizer (ATWD), a custom-built Application Specific Integrated Circuit (ASIC). Usually lying dormant, the ATWDs start to capture the delayed waveform when the PMT discriminator initiates it. The captured waveforms are only digitized in case a hit (i.e. local coincidence) is registered [136].

The DOMs are connected to the IceCube Laboratory (ICL) with twisted-pair copper cables. The power for the DOM is also transmitted with this cable. Two DOMs share one twisted-pair cable, and each DOM is also directly connected to its two neighbors on the same string. Fig. 2.5 shows the connection layout.

The Flasher Board houses 12 LEDs operating at a wavelength of ~400 nm. These are used to verify the DOM timing response, to measure the DOM in-ice position, to determine the optical properties of the ice, and to verify the reconstruction algorithms [136].

2.2.3 Detector Layout

In total, approximately 5800 DOM units were built and tested, 300 failing tests and the rest being delivered to the South Pole. The vast majority of these were ultimately deployed (5160 in total). The final detector layout (since the last drilling campaign 2010/2011, see below) consists of 86 strings. The DOMs were deployed along those strings, like pearls on a necklace. Each string contains 60 DOMs, with an average horizontal spacing between strings of 125 m [136].

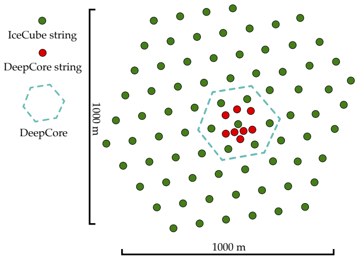

Figure 2.6: Top-down view of the IceCube detector, spanning 1 km^2 on the surface. From [137].

The instrumented part of the strings starts at 1450 m below surface, with one DOM every 17 m to a depth of 2450 m, just above the bedrock at a depth of 2820 m. In Fig. 2.7 the layout of the in-ice array can be seen. The strings follow a roughly hexagonal layout (see Fig. 2.6), with a side length of 1 km^2. The total instrumented volume of glacial ice is thus 1 km^3 [136]. Of the 5160 deployed DOMs, 92 are dead as of March 2023, a loss of 1.7 %[2].

One can see in Fig. 2.6 that there is a region in the center of the detector which is more densely instrumented: The strings are closer to each other, and also the spacing between DOMs on these strings is reduced from 17 m to 7–10 m. This part of the detector is DeepCore, designed to have a lower energy threshold of 10 GeV, a significant improvement over the 100 GeV for the rest of the detector. The DOMs within DeepCore are also modified for this goal, as they are equipped with PMTs that have a 35 % higher QE compared to the 'normal' DOMs [136].

[2] This is better than the predicted failure percentage, which was projected to be 2 % by 2023 [136].

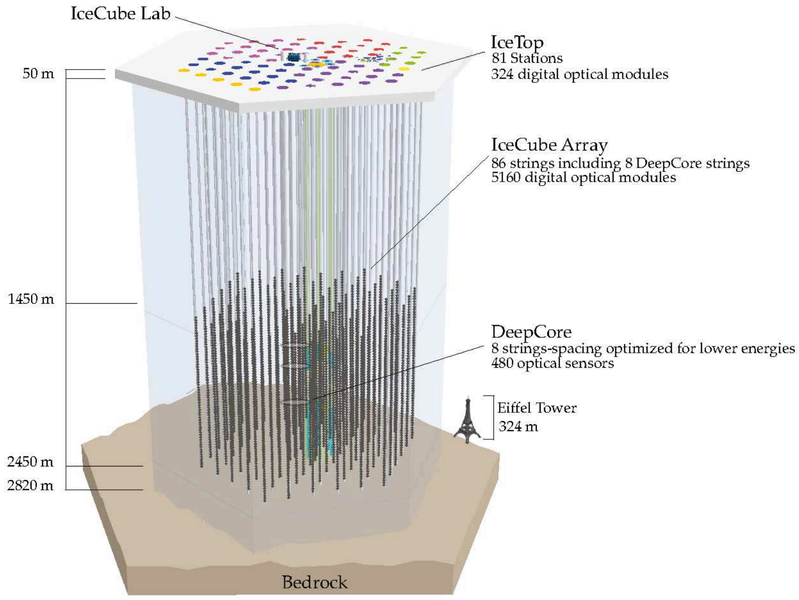

Figure 2.7: Side-on view of the IceCube detector, showing the instrumented array deep in the Antarctic glacial ice. In the center on top is the IceCube Laboratory, where data processing takes place. From [137].

2.2.4 Deployment

As one can imagine, embedding the DOMs within the ice was a highly nontrivial task. It required drilling 86 boreholes with a diameter[3] of roughly 60 cm and a length of 2500 m. This was achieved over several drilling campaigns with the Enhanced Hot Water Drill (EHWD) specifically built for this task. This drill had a total power of 5 MW and was able to drill with a maximum speed of 2.2 m min^{-1}. With these performance characteristics, one hole was drilled every 48 h on average, with drill operations round the clock [136]. It took seven drilling seasons to deploy the final IceCube86 setup, from the Antarctic summers 2004/2005 to 2010/2011. Fig. 2.8 shows the tower operations site directly above the borehole [138].

[138]: Benson et al. (2014), *IceCube Enhanced Hot Water Drill functional description*

The water for drilling the holes was heated to 88 °C with 35 water heaters working in parallel, each providing 125 kW power. The average amount of fuel used per drill hole was 27,000 L [138].

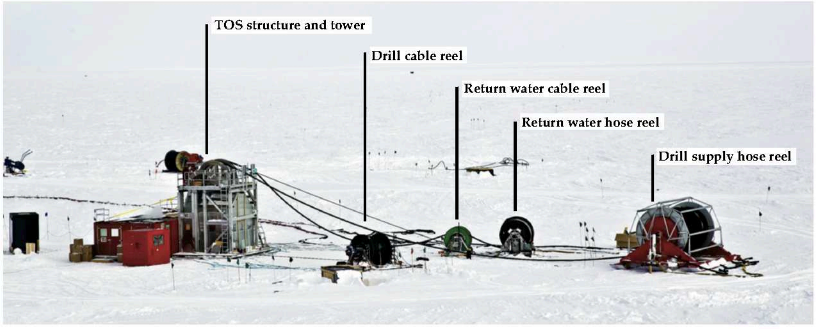

Figure 2.8: The hole drilling part of the IceCube Enhanced Hot Water Drill, excluding the supply for hot, pressurized water. One can see the tower operations site above the hole and the hoses providing hot water and returning cooled water from the borehole back to the generators in a closed loop. From [138].

2.2.5 The IceTop Surface Array

In addition to the in-ice detector, IceCube has a surface air shower detector for cosmic-ray physics (see Section 1.3), named IceTop. It was designed to study the cosmic-ray mass composition by correlating the energy it

[3] The hole diameter was larger than the DOM diameter (35 cm) to account for partial refreezing of the borehole.

measures on the surface with the energy deposited by muons in the ice as measured by the in-ice detector [139]. The energy sensitivity range of IceTop is 300 TeV to 1 EeV [140].

[139]: Abbasi et al. (2013), *IceTop: The surface component of IceCube*

[140]: Aartsen et al. (2013), *Measurement of the cosmic ray energy spectrum with IceTop-73*

The IceTop surface array consists of 2 × 81 ice-filled Cherenkov tanks. These are placed in pairs on the same hexagonal grid as the DOM strings for the in-ice array. Each tank is equipped with two standard IceCube DOMs (see Section 2.2.2) [139] which are operated at two different gain levels to increase the dynamic range.

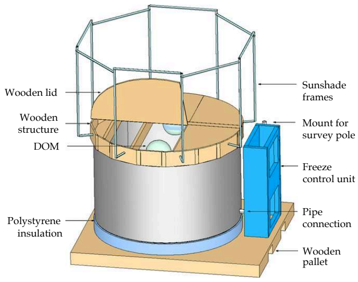

Figure 2.9: IceTop surface Cherenkov detector tank. From [139].

Results from three years of IceTop data show good agreement between models describing the transition from galactic to extragalactic cosmic rays at energies above the knee. In these, the spectrum of lighter elements softens earlier towards higher energies compared to heavier elements. This is indeed reflected in the data [141].

Furthermore, IceTop can be used to increase the purity astrophysical neutrino samples by vetoing muon events that mostly stem from cosmic-ray induced air showers, as opposed to muons created by neutrino interactions [142] (more details on background rejection will follow in Section 2.4.1).

[141]: Aartsen et al. (2019), *Cosmic ray spectrum and composition from PeV to EeV using 3 years of data from IceTop and IceCube*

[142]: Amin (2021), *Implementation of IceTop data in the IceCube Realtime Alert System*

2.2.6 Data Acquisition

As noted above, only for locally coincident hits in multiple detectors the full waveforms are digitized by the DOMs. These are then sent to the IceCube Laboratory on the surface via the twisted-pair cable data link. In the laboratory all DOM data is ingested into the data acquisition system (DAQ). Hits throughout the detector are investigated by the system to establish common causality by temporal and sometimes spatial patterns. All hits for which common causality can be established form an ‘event’. The rate of these events varies seasonally with the atmospheric muon flux, with a median event rate of 2.7 kHz and a total data rate of 1 TB d^{-1} (roughly 100 Mbit s^{-1}) [136].

Satellite bandwidth is limited and costly. Therefore, further on-site software triggers reduce the data rate to 15 % of the initial rate. These events are then transmitted via satellite to the IceCube data center at the University of Wisconsin-Madison for further analysis. The full event stream is also written to redundant disks, which are moved twice per year to Madison [136].

2.2.7 Time Synchronization

Precise timing information is crucial to reconstruct an event (see Section 2.3). For this reason, all DOMs need to be synced to a common clock. This is achieved by syncing the whole system to a Symmetricom ET6000 GPS receiver. The synchronization of individual DOMs is performed once per second, while data transfer is paused during the process (the calibration sequence takes $\lesssim$1.3 ms) [143].

[143]: Abbasi et al. (2009), *The IceCube data acquisition system: Signal capture, digitization, and timestamping*

IceCube ensures temporal synchronicity with an algorithm called Reciprocal Active Pulsing (`RAPcal`): A bipolar pulse is initiated on the surface and sent to the DOM. The sender saves the local time when it sends the pulse and starts a timer. Upon reception down the string, the DOM also saves its current local time, saves the received pulse waveform, starts a timer, responds with a bipolar pulse of its own and stops the timer. Upon reception, the surface station stops its timer and requests the received pulse waveform and all timing information from the DOM.

With these six pieces of information—the two transmit timestamps, the two receive timestamps and both waveforms—a transformation from the GPS-synchronized surface to the local DOM time domain and vice versa can be calculated, with a precision of 1 − 2 ns [143].

2.3 Angular Reconstruction

The goal of IceCube reconstruction is twofold: Reconstructing the deposited neutrino energy, and reconstructing the neutrino arrival direction. With this in mind, one can sort the events seen by the detector broadly into two categories: Well-localized track events with poorly reconstructed energy, and cascade events with well-determined energy, but very imprecise origin.

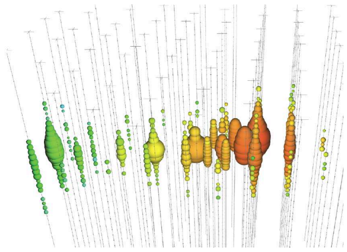

Figure 2.10: Track event: The long track allows for good angular reconstruction, with high uncertainty on the event energy. From `masterclass.icecube.wisc.edu`.

2.3.1 Event Types

Track Events (Fig. 2.10) are produced by secondary muons resulting from the charged-current interaction of muon neutrinos with the Antarctic glacial ice (see section 1.1.6). The secondary muons leave tracks in the ice with a length on the order of kilometers. Muons with energies above ∼300 GeV create tracks that exceed the detector length. In general, these allow for a good angular resolution, ranging from 1° for a 1 TeV muon to 0.3° for a 1 PeV muon [102]. The drawback is the large energy uncertainty of $0.25 \log E_\mu$ for a muon of energy E_μ [144], as parts of the high-energy muon tracks lie outside the instrumented volume [145].

Cascade Events (Fig. 2.11) on the other hand are initiated by the charged-current interactions of ν_e and ν_τ, as well as by neutral-current interactions from neutrinos of all flavors. They are usually relatively isotropic and contained within the detector, as typical track lengths are of $\mathcal{O}(10\,\mathrm{m})$. Their relative isotropicity only allows for poor angular resolution (10 − 15°), but comparably good energy reconstruction ($\frac{\delta E}{E} \approx 15\,\%$) [145].

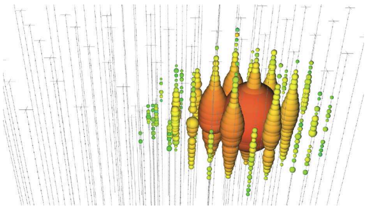

Figure 2.11: Cascade event: The energy is fully contained in the detector, as the event is relatively isotropic. The angular uncertainty is quite large though. From `masterclass.icecube.wisc.edu`.

[102]: R. Abbasi et al. (2022), *Evidence for neutrino emission from the nearby active galaxy NGC 1068*

[144]: Aartsen et al. (2014), *Energy reconstruction methods in the IceCube neutrino telescope*

[145]: Aartsen et al. (2017), *The IceCube realtime alert system*

As this thesis is concerned with the sources of high-energy astrophysical neutrinos and the angular resolution of track events allows for a better pointing accuracy compared to cascade events, the next section will focus on the angular reconstruction of the former.

2.3.2 Likelihood Approach

The main angular reconstruction algorithm for muon tracks used in IceCube is based on the work done for AMANDA. It employs a maximum-likelihood method [146]. This can be understood as follows: Given a set of unknown track parameters $\vec{a}$ and a set of experimentally determined

[146]: Ahrens et al. (2004), *Muon track reconstruction and data selection techniques in AMANDA*

values $\vec{x}$, what values of the unknown parameters $\vec{a}$ do maximize the probability of measuring the actually observed values $\vec{x}$?

This likelihood is denoted $\mathcal{L}(\vec{x}|\vec{a})$. If the components x_i of $\vec{x}$ are independent, it can be expressed as

$$\mathcal{L}(\vec{x}|\vec{a}) = \prod_i p(x_i|\vec{a}). \tag{2.2}$$

Here, $p(x_i|\vec{a})$ is the probability density function (PDF) of measuring x_i given a set of parameters $\vec{a}$. The reconstruction consists of obtaining the set of unknown track parameters $\vec{a}$ that maximizes the likelihood $\mathcal{L}$ (or, for technical reasons, minimizes $-\log\mathcal{L}$).

2.3.3 Parametrization

To simplify matters, we assume that we are dealing with a muon with maximum allowed speed ($\beta = 1$), traveling along a track of infinite length. Furthermore, we neglect stochastic photon losses, which are mainly caused by impurities within the Antarctic ice (see below). The set of parameters $\vec{a}$ needed to describe the physical situation in the detector is visualized in Fig. 2.12: $\vec{a} = (\vec{r}_0, E_0, t_0, \vec{p})$. It describes the trajectory of a muon located at position $\vec{r}_0$ with energy E_0 at time t_0, and traveling in the direction $\vec{p}$.

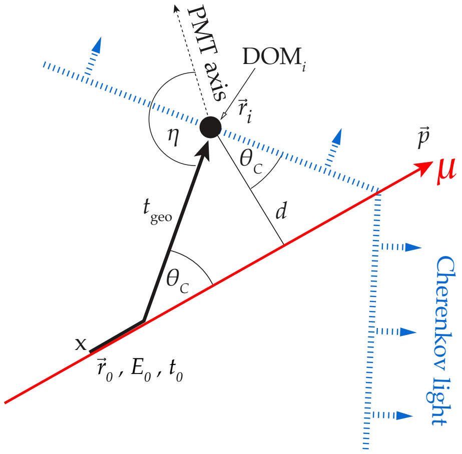

Figure 2.12: Parametrization for the angular muon reconstruction, describing the trajectory of a muon (red line) with energy E_0 that is in position $\vec{r}_0$ at time t_0, traveling towards direction $\vec{p}$. Cherenkov light from that muon is detected by a DOM at position $\vec{r}_i$ at distance d to the trajectory with a light travel time t_{geo}. Adapted from [146].

So we are dealing with one time parameter, one energy parameter and positional parameters in 5 dimensions, where the vertex position $\vec{r}_0$ can be expressed in (x,y,z) and the muon direction $\vec{p}$ is usually described by two angles, zenith and azimuth[4]. In the context of IceCube muon reconstruction, this parameter set $\vec{a}$ is known as the 'track hypothesis'.

Now, a DOM in the detector at position $\vec{r}_i$ at a distance d to the track can be hit by Cherenkov photons emitted by the muon traveling along the track with the Cherenkov angle θ_C, arriving at an angle η with respect to the PMT axis. Without scattering, the photon reaches the DOM$_i$ at the geometrical time t_{geo}, which can be expressed as

$$t_{geo} = t_0 + \frac{|\vec{p}\cdot(\vec{r}_i - \vec{r}_0)| + d \cdot \tan\theta_C}{c_0}, \tag{2.3}$$

[4] The zenith angle is the angle to a line vertical to the Earth, centered on the detector surface and pointing 'southwards' (away from the Earth). An azimuth angle of 90° describes a particle traveling from North (prime meridian pointing towards Greenwich) to South, while a particle with an azimuth angle of 0° travels along the 90th Meridian from East to West.

with c_0 the vacuum speed of light. As this does not take scattering into account, it is useful to define a relative time of arrival, the time residual $t_{\mathrm{res}} = t_{\mathrm{hit}} - t_{\mathrm{geo}}$. This is the additional time introduced by scattering as opposed to a Cherenkov photon traveling directly from the muon to the DOM.

To first order, the scattering does not loose light, but delays the photons. It is dominated by Mie scattering on impurities located within the ice. These impurities are thought to mainly comprise mineral dust, salt, acid droplets and soot by volcanic activity, all deposited by snow fall during the last 100,000 years [147].

[147]: Abbasi et al. (2022), *In-situ estimation of ice crystal properties at the South Pole using LED calibration data from the IceCube Neutrino Observatory*

As the position of the DOM is known, t_{res} is the most significant observable for each DOM. We therefore simplify the x_i in the likelihood's PDF to $t_{\mathrm{res},i}$ and express $\vec{a}$ as a function of the individual DOM parameters: $\vec{a} = (d_i, \eta_i, \ldots)$, where η_i is the angle to the DOM PMT axis (see Fig. 2.12).

2.3.4 First (Single) Photoelectron Fit

Matters can be simplified further. While the muon is traveling and emitting Cherenkov light, multiple photons can hit each DOM. One approximation is to only regard the first photon hitting an individual DOM, as it has usually been scattered less than the average photon. If the reconstruction is using this simplification, it is called Single Photoelectron (SPE) fit. The likelihood function for this is[5]

$$\mathscr{L}_{\mathrm{1st}}(\vec{x}|\vec{a}) = \prod_{i}^{\text{1st hits}} p(t_{\mathrm{res},i}|\vec{a} = d_i, \eta_i, \ldots), \tag{2.4}$$

where the probability density function $p(t_{\mathrm{res},i}|\vec{a})$ is obtained from simulations modeling the photon propagation through the Antarctic ice (this is necessary because the photon scatter needs to be accounted for). The simulation results are either stored in look-up tables or approximated by analytical functions [146].

2.3.5 Multi Photoelectron Fit

A complication is that the first photon in a DOM detecting multiple photons tends to hit the DOM earlier than a photon detected by a DOM only registering this very photon. This is because more photon hits mean that the DOM is closer to the event, and therefore receives a higher signal—which means that the event is detected earlier on.

This complication leads to the Multi Photoelectron (MPE) fit. Here, the single-photon part of the likelihood is modified by the cumulative PDF (CDF), which is given time-integrating the photon arrival PDF from the t_{res} to infinity:

$$P(t_{\mathrm{res},i}|\vec{a}) = \int_{t_{\mathrm{res},i}}^{\infty} p(t|\vec{a})\, dt. \tag{2.5}$$

Using this, the MPE likelihood is given as

[5] As stated above, this already includes the reduction of x_i to $t_{\mathrm{res},i}$.

$$\mathscr{L}_{\mathrm{MPE}} = \prod_{i}^{\text{1st hits}} p(t_{\mathrm{res},i}|\vec{a}) \cdot N_i \cdot (1 - P(t_{\mathrm{res},i}|\vec{a}))^{N_i - 1}, \tag{2.6}$$

where N_i is the total number of photons recorded by DOM$_i$ [146]. This is almost what is used in the angular reconstruction of the alerts IceCube distributes (see Section 2.4.3). The only difference is the PDF in question, which is described in the following Section 2.3.6.

2.3.6 `SplineMPE` Reconstruction

So far, nothing has been said about the photon arrival time PDF $p(t_{\mathrm{res}}|\vec{a})$. The most straightforward approach is using a Gaussian distribution, which models the muon moving through the ice as plane wave of constant velocity. This can be enriched by assuming a more physically correct minimally ionizing muon track. When doing so, the PDF becomes a Gamma distribution of the form

$$p(t_{\mathrm{res}}) = \frac{\beta^{\alpha} t_{\mathrm{res}}^{\alpha-1} e^{-\beta t_{\mathrm{res}}}}{\Gamma(\alpha)}, \tag{2.7}$$

where $\alpha = \frac{d_{\mathrm{eff}}}{\lambda}$, $\beta = \frac{1}{\tau} + \frac{c_m}{\lambda_\alpha}$. Here λ is the scattering length, λ_α the absorption length, c_m is the speed of light in a transparent medium and Γ is the Gamma function. d_{eff} is a modified version of the distance to the DOM d that takes into account that the PMTs face downwards, and light from a track above a DOM needs to scatter around the DOM first (thus introducing an additional angle describing this delay). The scattering and absorption lengths λ and λ_a, as well as the unspecified parameter τ have been determined by Monte Carlo simulations [148].

[148]: Abbasi et al. (2021), *A muon-track reconstruction exploiting stochastic losses for large-scale Cherenkov detectors*

`SplineMPE` uses a more sophisticated PDF, as the Gamma PDF above assumes an optically homogenous medium. This is not the case for Antarctic ice. For this reason, a fitted ice model derived from measurements with the Flasher Board (see Section 2.2.2) is used.

The basic idea is to create a large lookup table of simulated minimally ionizing muon tracks of infinite length with many different positions and orientations, i.e. high-dimensional histograms. A comprehensive lookup table of these simulations would be too large (hundreds of GB) and numerically problematic due to empty bins or interpolation artifacts [149]. To mitigate this, the histograms are normalized and interpolated with multidimensional basis splines. These splines then represent the photon arrival PDF dependent on the track vertex and orientation in the detector. They are defined by knots of fixed positions, so the PDF becomes

[149]: Whitehorn et al. (2013), *Penalized splines for smooth representation of high-dimensional Monte Carlo datasets*

$$p(t_{\mathrm{res}}) = \sum_{i=1}^{T-k-1} w_i B_{i,k}(t_{\mathrm{res}}, \vec{a}, \kappa). \tag{2.8}$$

Here T is the total number of knot positions, B is the i-th basis spline of order k, $\vec{a}$ again describes the track, and κ denotes the parameters of the ice model [149].

2.3.7 `Millipede` Reconstruction

High-energy neutrino events that pass the realtime alert selection criteria (see Section 2.4) are subjected to a more sophisticated algorithm, which uses the information from all the detected photons. This algorithm is dubbed `Millipede`. Due to its high demand of computational resources it is only run for selected events at the IceCube data center in Madison.

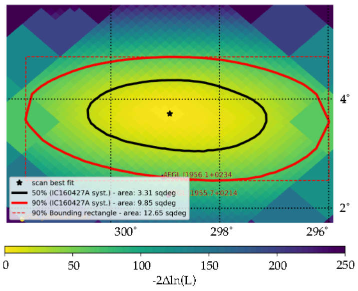

Figure 2.13: `Millipede` reconstruction of *IC221124A*.

This reconstruction consists of a maximum likelihood scan covering the whole sky. To allow scanning, the sky is pixelated into grids of increasing resolution following the Hierarchical Equal Area isoLatitude Pixelation (`HEALPix`) scheme [150]. Each pixel of a specific resolution covers the same area on the sky. For each scanned pixel, the muon direction is fixed to originate from that sky location.

[150]: Gorski et al. (2005), *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere*

For this fixed direction, the likelihood of the deposited energy resulting from the best fit and the neutrino interaction vertex in the detector are then computed. Pixels near the likelihood maximum are then scanned again with a finer `HEALPix` resolution. This procedure ultimately results in a likelihood map of the sky with increasing granularity towards the global maximum [151].

[151]: Abbasi et al. (2023), *IceCat-1: the IceCube Event Catalog of Alert Tracks*

To compute the 50 % and 90 % confidence level uncertainty contours, Monte Carlo resimulations from the high-energy neutrino event *IC160427A* are used. For this event, Pan-STARRS found a possible counterpart, *PS16cgx*[6] [152].

In the resimulations, the systematic parameters of the Antarctic ice used to model photon propagation were varied. Additionally, the simulated directions and energies were varied, with cuts to ensure that the light deposition of the resimulations resembled the original event (±2° in direction, ±20 % in deposited charge).

Each resimulated event was fit with `Millipede`. The distribution of differences between the best-fit likelihood and the ground truth of the simulated event were then employed to convert the change in log-likelihood over the map into a confidence level.

Due to the systematic uncertainties involved, Wilk's theorem does not apply here. To account for this fact, the contours derived for individual events (which make use of the theorem) are scaled up with correction values obtained from the resimulations of *IC160427A*. Then, all pixels that satisfy $\log \mathscr{L}_{\mathrm{min}} - \log \mathscr{L}_{\mathrm{pixel}} = -11.3\ (-32.1)$ form the 50 % (90 %) error contours. Currently, these correction values are 22.2 and 64.2 for the 50 % and 90 % uncertainty contours [153]. In Fig. 2.13 these contours are displayed for an example event. The black line shows the 50 % uncertainty region, while the red line shows the 90 % area.

[153]: Lagunas Gualda et al. (2021), *Studies of systematic uncertainty effects on IceCube's real-time angular uncertainty*

However, resimulations of newer high-energy neutrino events have shown that the method of scaling up the errors with correction values obtained from resimulating *IC160427A* does in some cases not faithfully capture the errors of those newer resimulations: They are sometimes under- and sometimes overestimated, depending on the topography of the event. For details on this, see [153].

[6] Spectroscopic follow-up revealed that this event was either an SN Ic, which would be compatible with neutrino production, or—more likely—an SN Ia, which would exclude it as a neutrino source [152].

[152]: Kankare et al. (2019), *Search for transient optical counterparts to high-energy IceCube neutrinos with Pan-STARRS1*

2.4 The Realtime Alert Program

Since 2016, IceCube hosts a realtime alert program, providing the astrophysical community with low-latency and high-quality astrophysical neutrino alerts [145]. This program saw a major revision in 2019, when two new alert streams, named 'Gold' and 'Bronze', were created [154] ('Silver' has been reserved for a Cascade stream that never saw the light of day). As these were designed based on cuts reducing the IceCube background, one needs to understand the background contamination first.

[154]: Blaufuss et al. (2019), *The Next Generation of IceCube Real-time Neutrino Alerts*

2.4.1 Background

When searching for the sources of astrophysical neutrinos, there are two major sources of background events in the IceCube detector.

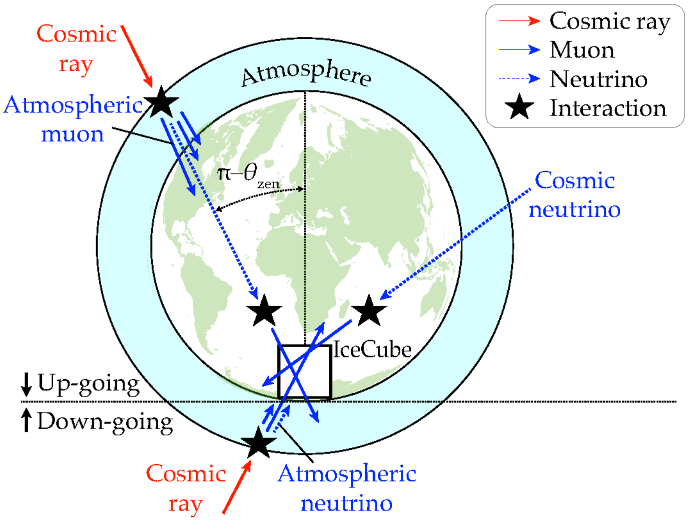

Figure 2.14: Background events in the detector. Cosmic rays (red arrows) hit the atmosphere around the globe and produce muons (solid blue arrows), as well as neutrinos (dotted blue arrows). When constraining to *up-going* events from the northern hemisphere, the detector is shielded from atmospheric muons, but not atmospheric neutrinos, as these can traverse the Earth. Adapted from [137].

Both types of background events stem from secondary cosmic-ray particles created within the Earth's atmosphere, showering down on the surface.

Atmospheric Muons are created by cosmic rays hitting the atmosphere. One can efficiently filter this background by restricting the analysis to *up-going* muon tracks. These are tracks that come from the bottom of the detector, i.e. the northern hemisphere. Atmospheric muons stemming from cosmic-ray events in the northern hemisphere are filtered out by the Earth's core (see Fig. 2.14), as the mean free path of muons within the Earth is much smaller than the distance they have to cross ([155], see Fig. 2.15). Note that due to light scattering within the ice, some down-going tracks can be misclassified as up-going [137].

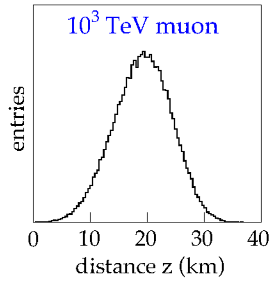

Figure 2.15: Free path length for 1 PeV muons in ice. The mean free path in ice is slightly longer than in rock. Adapted from [155].

[155]: Chirkin et al. (2004), *Propagating leptons through matter with Muon Monte Carlo (MMC)*

[137]: Ahlers et al. (2018), *Probing particle physics with IceCube*

One complication here stems from the fact that the Earth starts to become opaque for neutrinos of higher energies. Studies interested in PeV neutrinos therefore must deal with the fact that these get the more suppressed the longer their path through the Earth is. For example, a 1 PeV neutrino with a zenith angle of 140° is absorbed with 90 % probability [156].

[156]: Aartsen et al. (2017), *Measurement of the multi-TeV neutrino interaction cross-section with IceCube using Earth absorption*

Atmospheric Neutrinos also stem from cosmic-ray induced air showers. These cannot be suppressed by directional cuts, creating an irreducible background: When atmospheric neutrinos cross the Earth and interact with the matter close to the detector, they can produce muons indistinguishable from muons created by 'proper' cosmic neutrinos [137].

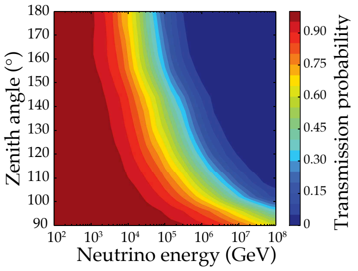

Figure 2.16: Neutrino transmission probability through the Earth. The longer the distance traveled (higher zenith angles) and the higher the neutrino energy, the more likely is absorption. Adapted from [156].

2.4.2 Event Selection

Gold and Bronze alerts are drawn from three different selection schemes, originally designed to cater to different science goals. These are High Energy Starting Events (HESE) [157], Extremely High Energy (EHE) events [158] and Gamma-Ray Follow Up (GFU) events [159].

HESE are events that start within the detector. This is guaranteed by using the outer regions of the detector as a veto region, in which (almost) no Cherenkov light must be detected [93]. The sizes of those regions can be seen in Fig. 2.17. The majority of HESE events are cascade-type events and therefore not well suited for observational follow-up due to their poor angular reconstruction. Because of this drawback, additional cuts are applied: At least 6000 photoelectrons are required, the reconstruction must favor a track interpretation of the event, and the reconstructed track length must be at least 200 m [151]. Note that the selection criteria used in the Gold and Bronze HESE selection are slightly different from those in the original paper cited above.

EHE aims at neutrino energies of 0.5–10 PeV. To reject atmospheric background events, a two-dimensional cut depending on the reconstructed zenith angle and the log of detected photoelectrons is applied. Additionally, a χ^2-based goodness-of-fit cut is applied to select track-like events with good reconstructions [151].

GFU events are selected based on a boosted decision tree trained to identify through-going (as opposed to 'starting') track events with astrophysical origin. Energy cuts are applied: Northern hemisphere events are selected based on their reconstructed muon energies, while events from the southern hemisphere are selected based on the total photoelectron charge deposited in the detector.

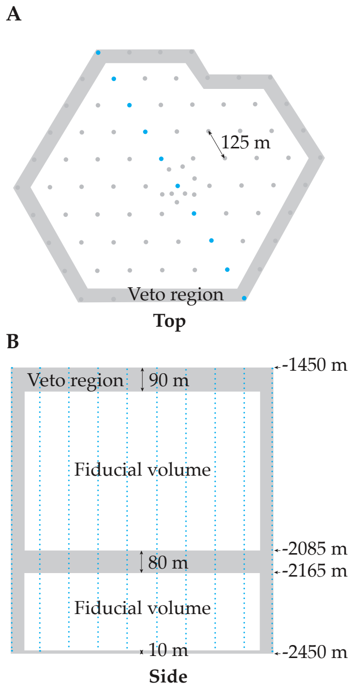

Figure 2.17: High-energy starting events veto regions. The strings marked in blue in the top-down view at the top (A) show the location of the side view, displayed at the bottom (B). Adapted from [93].

The cuts from all three event pools are gauged to achieve two different average values of signalness, which is a proxy for the probability that the event is of astrophysical origin [151]. It is defined as

Definition 2.4.1 $$Signalness(E, \theta_{zen}) = \frac{N_{signal}(E,\theta_{zen})}{N_{signal}(E,\theta_{zen})+N_{background}(E,\theta_{zen})}$$

Here E is the reconstructed neutrino energy, and $N_{\text{signal}}(E, \theta_{\text{zen}})$ and $N_{\text{background}}(E, \theta_{\text{zen}})$ are the number of signal and background events at zenith angle θ_{zen} above energy E as determined by simulations [151].

The cuts on the individual event selections are tuned to ensure that *Gold* alerts on average have a signalness of 50 %, and *Bronze* alerts have an average signalness of 30 %.

[157]: Abbasi et al. (2021), *IceCube high-energy starting event sample: Description and flux characterization with 7.5 years of data*

[158]: Aartsen et al. (2016), *Constraints on Ultrahigh-Energy Cosmic-Ray Sources from a Search for Neutrinos above 10 PeV with IceCube*

[159]: Aartsen et al. (2016), *Very high-energy gamma-ray follow-up program using neutrino triggers from IceCube*

2.4.3 Alert Distribution

All events that pass a first stage of filtering are sent to the IceCube data center (see Section 2.2.6) via Iridium satellite to minimize latency. There, their signalness is computed. After this, they are globally distributed with the General Coordinates Network[7] (GCN) in the form of GCN Notices.

Each notice contains the discovery time, a unique event number, the reconstructed direction in right ascension (RA) and declination (Dec) of the candidate neutrino computed by `SplineMPE` (see Section 2.3.6), a statistical error for the direction, the reconstructed neutrino energy, the signalness and the false alarm rate [154].

This information is later appended by `Millipede`, a computationally more demanding, but more sophisticated reconstruction (see Section 2.3.7). The results of these reconstructions (i.e. the angular uncertainty) are typically distributed a few hours after the initial notice in the form of a GCN circular, and an updated GCN notice [154].

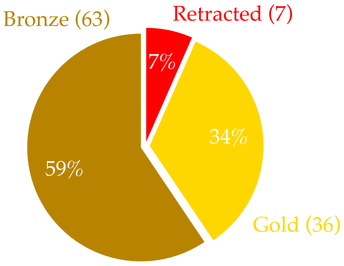

Figure 2.18: High-energy neutrino alerts issued by IceCube since start of the new alert stream in June 2019, as of March 2023.

As of March 2023, a total of 106 events have been distributed in this format, with 7 later being retracted. Since the start of the Gold and Bronze alert streams in June 2019, this amounts to 2.2 non-retracted alerts per month, with 0.8 Gold alerts and 1.4 Bronze alerts—this is quite close to the 2.5 alerts per month predicted in [154]. For a full list of all high-energy neutrino alerts issued by IceCube, not only those since introduction of the Gold and Bronze format, see [151].

If the neutrino is most likely astrophysical and its origin reasonably well pinpointed, one can scan the sky localization with a telescope and look for potential sources of the neutrino. But where to obtain optical images from? The telescope used for this, the Zwicky Transient Facility, will be described in the next Chapter 3.

[7] https://gcn.nasa.gov/

[154]: Blaufuss et al. (2019), *The Next Generation of IceCube Real-time Neutrino Alerts*

3 The Zwicky Transient Facility

The P48 dome at Mount Palomar, housing the Zwicky Transient Facility. Image credit: Caltech.

The second instrument relevant for this thesis is the Zwicky Transient Facility (ZTF). It is named after the notorious Swiss-American astronomer Fritz Zwicky[1].

ZTF is a wide-field optical survey telescope. Its wide field means that its use case is to periodically scan the full sky accessible to it, in contrast to taking pictures of specific objects. As its angular resolution is poor, ZTF is mainly used to register changes in brightness and therefore to detect transient events (short: transients).[2]

ZTF is located at Mount Palomar in California, United States, at 1700 m above sea level, roughly 130 km southeast of Los Angeles. Its optical system, the 1.2 m (48 inch) Samuel Oschin telescope, follows a Schmidt design (see Section 3.1) and was inaugurated in 1948 [162]. At first light and for years to come, it was the largest Schmidt telescope in the world. Originally, the telescope exposed photographic plates, covering a field of view (FoV) of 44 deg^2. As such photographic plates have obvious drawbacks, and because technological progress made it possible, the Near-Earth Asteroid Tracking (NEAT) program [163] replaced them with a charge-coupled device (CCD) camera in the early 2000s.

[163]: Pravdo et al. (1999), *The Near-Earth Asteroid Tracking (NEAT) Program: An Automated System for Telescope Control, Wide-Field Imaging, and Object Detection*

In 2003, the NEAT camera was replaced by the Quasar Equatorial Survey Team (QUEST) [164] camera. The immediate predecessor of ZTF, the Palomar Transient Factory (PTF) [165], began operation in 2009. Equipped with a 96 Megapixel camera, it already had many of the characteristics of ZTF: It was a fully automated survey, searching for optical transients with a CCD camera.

See `https://sites.astro.caltech.edu/palomar/about/telescopes/oschin.html` for a historical overview.

[164]: Baltay et al. (2007), *The QUEST Large Area CCD Camera*

[165]: Law et al. (2009), *The Palomar Transient Factory: System Overview, Performance, and First Results*

ZTF follows the concept of PTF, but is equipped with a larger and more sensitive camera. With 47 deg^2, the 600 Megapixel camera fully covers the large FoV of the P48. The main design metric for ZTF was volumetric survey speed. This is the volume within which an object of given absolute magnitude can be detected in one exposure, divided by the total time for the exposure (observation plus overhead) [166]. The system saw first light in 2017, and started scientific operations in 2018 (the first survey data was taken on 2018–03–20). As of August 2023, ZTF is still operational.

[166]: Bellm (2016), *Volumetric Survey Speed: A Figure of Merit for Transient Surveys*

[1] Zwicky first employed the Virial theorem to infer the existence of dark matter [160]. Furthermore, together with Walter Baade, he posited the existence of supernovae and the creation of neutron stars in such events [161].

[2] Transient astrophysical events can be defined as the time-dependent signal (electromagnetic or in the form of gravitational waves) that is usually caused by the partial or total destruction of an astrophysical object; supernovae are a prime example.

Figure 3.1: View of Mt. Palomar with the three telescopes highlighted in the text. The P48 hosts ZTF and discovers transients, while the P60 with SEDM allows for quick classification with its robotic spectroscopic capabilities. The P200 allows for photometry of faint objects and high-resolution spectroscopy. Image credit: Caltech, annotations by the author.

There are two other telescopes located on Mount Palomar: The 1.5 m (60 inch) P60 telescope houses the SED Machine (SEDM) [167], a fully robotic, low-resolution spectrograph used for automatic classification of transients. The largest facility on the mountain is the 5.1 m (200 inch) Hale Telescope, commonly referred to as 'P200'. It is used for optical and infrared photometry as well as mid- and high-resolution spectroscopy of fainter sources. Together, these telescopes form a natural hierarchy: ZTF is the discovery engine for optical transients. Promising sources are then classified with SEDM. If a source warrants it, deeper photometry and better resolved spectroscopy can then be obtained with the P200. All three telescopes are shown in Fig. 3.1.

[167]: Blagorodnova et al. (2018), *The SED Machine: A Robotic Spectrograph for Fast Transient Classification*

3.1 Telescope Design

A Schmidt telescope like ZTF is by design dedicated to taking images, contrary to earlier designs allowing to observe through an eyepiece [168]. For this reason, it is also referred to as a Schmidt camera.

[168]: Schmidt (1938), *Ein lichtstarkes komafreies Spiegelsystem*

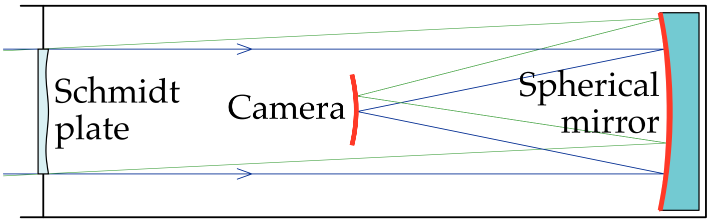

Figure 3.2: Schmidt telescope schematic. Light enters from the left, passes the Schmidt plate (an aspherical correcting lens), gets reflected by a spherical mirror at the end onto a curved photographic plate or camera halfway down the tube. Figure adapted from `https://commons.wikimedia.org/wiki/File:Schmidt-Teleskop.svg`.

The design goal of the Schmidt telescope was a wide FoV. This made it the ideal instrument for sky surveys, where a large FoV maximizes the on-sky area that can be monitored. A Schmidt telescope combines a spherical mirror at the end of the telescope tube with an aspherical correcting lens (Schmidt plate) at the tube's entrance. The use of a spherical mirror combined with the correcting lens gets rid of comatic aberration and astigmatism, which increases the image quality.

Until the end of the 20th century, around ten Schmidt telescopes have been built, most of them to conduct sky surveys [169]. At least two of them were space telescopes: ESA's astrometry mission *HIPPARCOS* [170] (1989–1993) and the NASA exoplanet mission *Kepler* [171] (2009–2018). In both cases, the mission entailed monitoring of large areas of the sky; prime territory for Schmidt telescopes.

[169]: Cannon (1995), *Schmidt Telescopes: Their Past, Present and Future*

[170]: ESA (1997), *The HIPPARCOS and TYCHO catalogues. Astrometric and photometric star catalogues derived from the ESA HIPPARCOS Space Astrometry Mission*

[171]: Koch et al. (2010), *Kepler Mission Design, Realized Photometric Performance, and Early Science*

3.2 Camera

The ZTF camera is a CCD design, consisting of 16 individual CCDs by commercial manufacturer e2v (now Teledyne, the model is Science CCD231-C6), each having 6144 × 6160 pixels, resulting in a total camera resolution of ~600 Megapixel [173]. As one can see in Fig. 3.7, the array of 16 CCDs is slightly bent. This is necessitated by the Schmidt design, where the camera needs to be spherical, matching the spherical mirror. As individual CCDs are flat, each of the 16 sensors is installed slightly tilted, tracing the overall curvature. To get rid of residual deviations from the global curvature, a field flattener lens is mounted in front of each sensor [174].

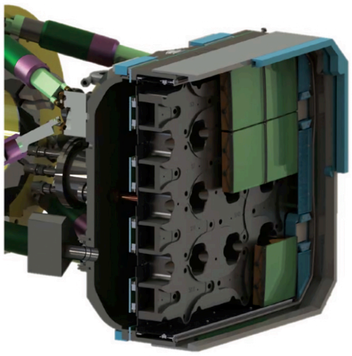

Figure 3.3: The ZTF camera in detail. From [172].

[173]: Dekany et al. (2016), *The Zwicky Transient Facility Camera*

[174]: Bellm et al. (2019), *The Zwicky Transient Facility: System Overview, Performance, and First Results*

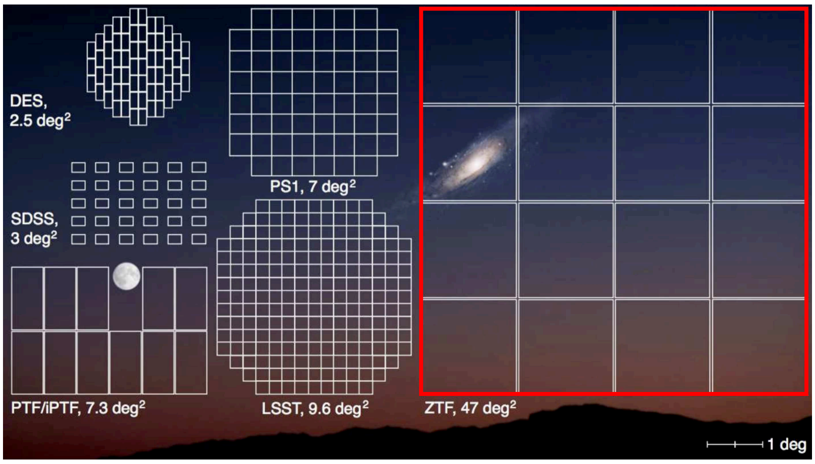

Figure 3.4: ZTF field of view (highlighted in red) in comparison to other sky survey telescopes, including the future Vera C. Rubin observatory (denoted as 'LSST') [175]. Note also the 6.5-fold increase with respect to ZTF's predecessor, PTF/iPTF. From [176], highlighting by the author.

3.2.1 CCDs

CCDs are silicon-based light sensors. They consist of arrays of coupled metal-oxide semiconductor (MOS) capacitors. Each one of those is able to store the charge created by incident photons; spatial resolution is achieved by having one capacitor per pixel of the sensor array. The array is then exposed to light for a specified amount of time. During the exposure, incident photons create a charge proportional to the amount of light hitting each capacitor via the photo-electric effect. This charge is accumulated in each capacitor until the exposure is finished. To read out the CCD, the charges need to be moved to neighboring capacitors. When the MOS capacitors are tightly placed, one can move the charges from one capacitor to the next by changing the voltages on the capacitor's gates.

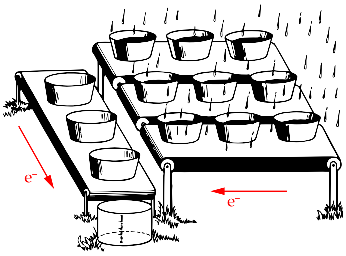

Figure 3.5: CCD operational principle, explained with buckets measuring precipitation. From [177].

In Fig. 3.5 the principle of a CCD is explained with little buckets collecting rain water. Each bucket symbolizes one capacitor or one pixel of the sensor array respectively. After the rain has stopped (the exposure is finished), each bucket naturally contains an amount of water proportional to the amount of water that rained down over it. Now the amount of water in each bucket (the charge deposited by incident light in each capacitor) needs to be measured.

To do this, the buckets in each row are moved one position to the left with horizontal conveyor belts. Each bucket at the left end of the horizontal conveyor belt is then emptied into the bucket on the single vertical conveyor belt. The buckets of this vertical belt are then one by one drained into the measuring bucket on the bottom left. After all buckets on the vertical belt are emptied, the process starts anew, until all buckets are empty. As one can see, the time this process takes is quadratic with respect to the amount of buckets (or pixels) [177]. To speed up the process, one can subdivide the sensor area into smaller sections, which are read out in parallel.

[177]: Janesick et al. (1987), *Scientific Charge-Coupled Devices*

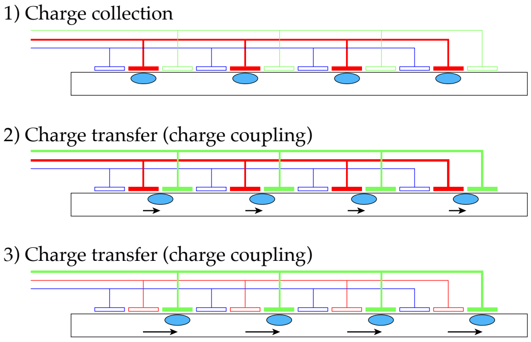

Figure 3.6: Charge transport in a CCD. High voltage electrodes are shown filled, while low voltage electrodes are empty. By stepping through different combinations of electrodes, the electric potential changes, and the charge shown as blue ellipses moves through the semiconductor.

Each CCD pixel actually consists of three electrodes, which can be seen in Fig. 3.6. The transport of charge within the CCD is realized by subsequently applying a higher voltage to different electrodes, moving the charge step-wise.

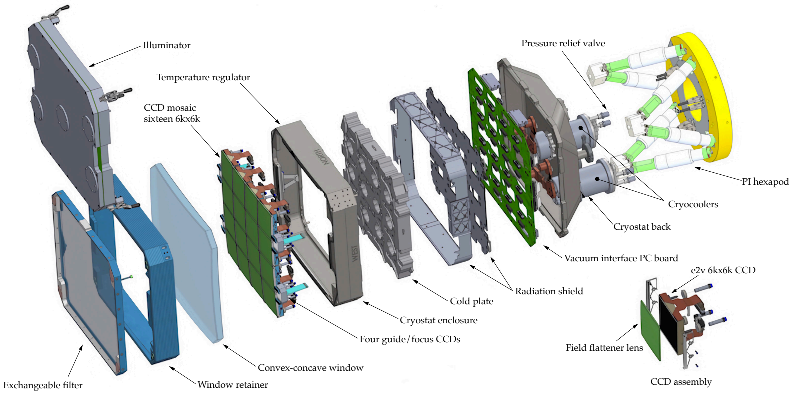

Figure 3.7: The ZTF camera. One can see the CCDs sandwiched between the filter on the front and the cryostat on the back. From [174].

The typical single image exposure time for the ZTF camera is 30 s, while the readout time (transferring and measuring the charge in each capacitor) is 8.2 s. Readout and digitization is done in parallel for four quadrants, each containing 4 CCDs. The four readout devices for the quadrants are Archon CCD controllers by Semiconductor Technology Associates (STA), operated at 1 MHz. Each of these operates 16 simultaneous readout channels, four for each CCD. In total, the camera is simultaneously read out in 64 independent regions to speed up the process.

Four additional smaller CCDs (2k x 2k pixels) are used as guidance, tip, tilt and focus sensors, with one additional Archon controller to read them out [172]. To reduce the thermal noise of the camera, it was placed on a cold plate, which is cooled by a cryocooler to 160 K [173].

[172]: Dekany et al. (2020), *The Zwicky Transient Facility: Observing System*

3.3 Optical System

ZTF is operated with three different color filters, and uses a fixed grid on the sky; both of which will be detailed in this Section.

3.3.1 Filters and Shutter

Because the CCD is only sensitive to the total amount of collected photoelectrons, the wavelength information of the incident photons is lost. To circumvent this, different optical bandpass filters must be placed in front of the camera. Color information can then be obtained by imaging the object in question with a least two filters. ZTF employs three different filters: A *g*-band filter with a median wavelength of 472 nm (corresponding to blue light), an *r*-band filter (median wavelength: 634 nm, red light) and an *i*-band filter (789 nm, near-infrared). These filters can be exchanged with a robotic arm, securely stowing the replaced filter and magnetically attaching the new one [172]; a process taking 110 s [174].

The main decision goal for the filter selection was to maximize the signal-to-noise ratio by avoiding major sky emission lines at Mt. Palomar while avoiding excessive costs. ZTF does not exactly match the filters of potential calibrators, e.g. the Sloan Digital Sky Survey (SDSS) [178], PanSTARRS (PS1) [179] or *Gaia* [180].

[178]: York et al. (2000), *The Sloan Digital Sky Survey: Technical Summary*

[179]: Kaiser et al. (2002), *Pan-STARRS: A Large Synoptic Survey Telescope Array*

[180]: Prusti et al. (2016), *The Gaia mission*

The transmission of each filter and the CCD quantum efficiency curve can be seen in Fig. 3.8. As one can see, the quantum efficiency starts to decrease within the *r*-band towards higher wavelengths, rendering the *i*-band the least sensitive of the three filters. The median sensitivity of 5 σ for a 30s exposure reflects that fact: It is 20.8 (21.1) mag in the *g*-band, 20.6 (20.9) mag in the *r*-band and 19.9 (20.2) mag in the *i*-band; with values for optimal visual conditions during new moon given in brackets. The resulting median image quality is 2.1 arcsec (*g*-band), 2.0 arcsec (*r*-band) and 2.1 arcsec (*i*-band) full width at half maximum (FWHM) of the point spread function (PSF, see Section 3.4.6 below) [174].

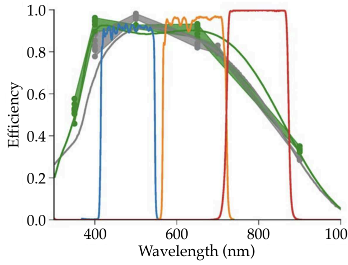

Figure 3.8: ZTF filter transmission for the three different bands (*g*-band: blue, *r*-band: orange, *i*-band: red). The green and gray datapoints show the CCD quantum efficiency measurements (single and double-layer reflective coating). From [174].

To decrease light obstruction, the ZTF shutter was newly developed and is mounted in front of the aperture, outside the telescope tube. It was developed by Deutsches Elektronen-Synchrotron (DESY) in cooperation with industry partner Bonn-Shutter and allows opening and closing in 290 ms [172].

3.3.2 ZTF Grid

One exposure during regular operations—in contrast to e.g. deep target of opportunity (ToO) images—lasts 30 s. There is an additional ~15 s overhead for readout and slewing the telescope. Also, some additional time is needed to exchange the filters. Therefore, a typical night lasting 8.67 h [181] results in roughly 700 exposures. In total, these amount to a sky area of over 32,500 deg^2, allowing to cover the full visible sky at Mount Palomar 15° above the horizon at least once.

[181]: Masci et al. (2019), *The Zwicky Transient Facility: Data Processing, Products, and Archive*

To enable robotic control, ZTF operates on a fixed primary on-sky grid of so-called 'fields'. Each field corresponds to a fixed sky location with an area of 47 deg^2; the telescope exclusively points to those fields. With such a system, some parts of the sky will always fall into the chip gaps, i.e. the parts of the FoV that lie between the 16 CCDs. To mitigate that, there exists a secondary grid that is diagonally offset from the primary grid, with 0.29° average overlap in right ascension, and 0.26° in declination [182]. Fig. 3.9 shows the primary grid and the number of visits per field in the g-band in a typical week.

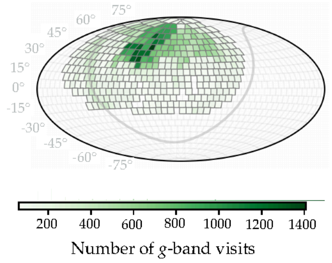

Figure 3.9: Number of ZTF g-band field visits during the first week of May 2020. The primary grid fully tiles the sky accessible at Mount Palomar. As one can see, there are fields which are visited more often; these are part of special surveys (see Section 3.5).

[182]: Bellm et al. (2019), *The Zwicky Transient Facility: Surveys and Scheduler*

3.4 Calibration and Image Processing

The ZTF image processing can be divided into two parts: The science exposures are taken at Mount Palomar, while calibration, extraction of transients and archival storage happens at the Infrared Processing and Analysis Center (IPAC)[3] situated on the campus of the California Institute of Technology (Caltech). IPAC ultimately hosts the ZTF images as part of IRSA[4], the Infrared Science Archive.

[3] `https://www.ipac.caltech.edu`

[4] `https://irsa.ipac.caltech.edu`

The full calibration and image processing pipeline is shown in Fig. 3.10. First come the different calibration steps: overscan and bias correction, flat fielding as well as astrometric and photometric calibration. This is followed by creating the difference images and finally extracting the sources. I will briefly explain the different steps in the next sections.

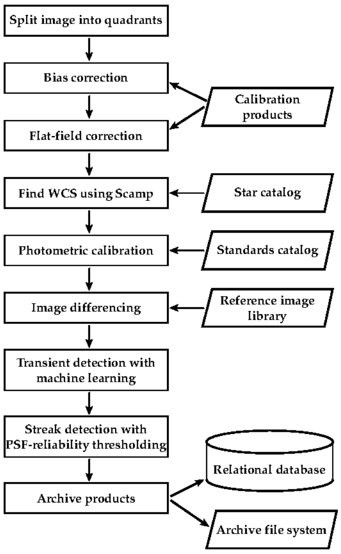

Figure 3.10: Flowchart of the ZTF calibration, starting with the raw image on top and ending with the final science products on the bottom. Adapted from [176].

3.4.1 On-site Processing and Datalink

Each image taken on site (calibration and science exposures alike) is stored as a FITS file and is subsequently compressed losslessly. The compressed images on average use 5 bit per pixel, so the full image is roughly 380 MB large. These images are immediately sent to IPAC with the High Performance Wireless Research & Educational Network (HPWREN)[5], a microwave-based data network, linking Palomar Observatory with the IPAC post-processing site. Each transfer typically takes 20 s, keeping up with the pace of ZTF observations [172].

[5] `https://hpwren.ucsd.edu/`

3.4.2 Overscan Correction

As the temperature of the CCDs changes over time, each image is subject to a time-dependent global offset induced by thermal noise. To correct for this, an overscan region is used. In the case of the ZTF CCDs, this region does not correspond to physical pixels, but is created during CCD readout. If one reads out more clock cycles (each cycle corresponds to one pixel readout and a global charge shift, see Section 3.2.1) than pixels are available, the charge that has gathered during science readout can be accessed.

For ZTF, this overscan is performed with 24 additional cycles, corresponding to 24 overscan pixels for each sensor row. After this, the median of these 24 pixels is taken and the full overscan column is fitted with a quadratic function, resulting in one column of overscan values. For each and every image taken by the camera, this column of overscan values is computed and subtracted from each row of the image [183].

3.4.3 Bias Correction

The CCD pixels also have different intrinsic noise levels: In the absence of external photoelectrons they still generate a signal, which is different from pixel to pixel—i.e. their zero point differs. Luckily, this variation has a higher degree of time stability than the thermal noise [184]. To correct it, at the beginning of each night at least 10 so-called 'bias images' are taken and overscan corrected. These images are zero-second exposures which are then stacked. After this, the truncated mean of each pixel is calculated, constituting the final bias image for the night and filter. This bias image is subtracted from each science exposure using the respective filter taken during the night [183].

[184]: Howell (2006), *Handbook of CCD Astronomy*

3.4.4 Flat Fielding

The pixels in the CCDs do not only differ in zero point, they also have slightly different gains and quantum efficiencies. This means that their response to light is not uniform. To account for this, the calibration needs an unstructured, uniformly bright light source against which the individual pixel response can be measured [184]. In the case of ZTF, this is achieved by using a flat field illuminator consisting of a round screen illuminated by eight identical boards with LEDs, each board housing $4 \times 15 = 60$ LEDs of 15 different colors, covering the full wavelength region of ZTF [172].

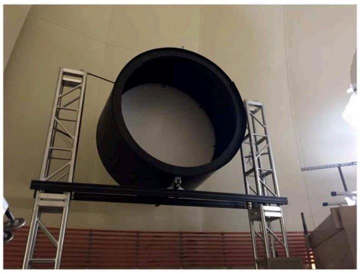

Figure 3.11: The ZTF flat field illuminator. From [172].

Each afternoon, before science operations begin, at least 20 images per filter are taken of the flat field illuminator. These are then overscan corrected, the bias image is subtracted, and the pixel values are normalized to a truncated global mean of 1 over the image to allow for later division. After this individual treatment, all flat field images per filter are stacked to a truncated mean per pixel and an outlier rejection algorithm is applied to isolate additional noisy pixels. All science images taken during the night are divided by this flat field image [183].

3.4.5 Astrometric Calibration

Astrometric calibration is the mapping of image pixel coordinates to an on-sky coordinate system. For ZTF, this is performed with stars contained in the *Gaia* Data Release 1 (*Gaia* DR1 [185]). The *Gaia* satellite

[185]: Brown et al. (2016), *Gaia Data Release 1. Summary of the astrometric, photometric, and survey properties*

was launched in 2013 and is still operational ten years later. Its main goal is a high-precision astrometric measurement of ~ 1 billion stars, therefore providing an ideal reference for the ZTF calibration. For this task, *Gaia* sources were selected from the catalog that are neither too faint, nor run risk of saturating the detector ($12 \leq G \leq 18$ mag). The astrometric solution is derived using the `SCAMP` [186] package. The stars are extracted from the image with `SExtractor` [187] and matched to the *Gaia* stars. The pointing, rotation and polynomial distortion needed to match the stars constitute the astrometric solution of the image [181].

[186]: Bertin (2006), *Automatic Astrometric and Photometric Calibration with SCAMP*
[187]: Bertin et al. (1996), *SExtractor: Software for source extraction*
[181]: Masci et al. (2019), *The Zwicky Transient Facility: Data Processing, Products, and Archive*

3.4.6 PSF Photometry

After astrometric calibration, one needs to extract sources contained in the image in order to perform photometric calibration (i.e. one first needs to identify stars to compare their brightness to reference measurements). To do so, point spread function (PSF) photometry is employed. Each telescope has a finite aperture and therefore suffers from diffraction. Additionally, the Earth's atmosphere is constantly and turbulently changing, with different atmospheric layers having different refractive indices, depending on their temperature. This smears out the light from the original point source, an effect known as 'seeing' (see Fig. 3.12).

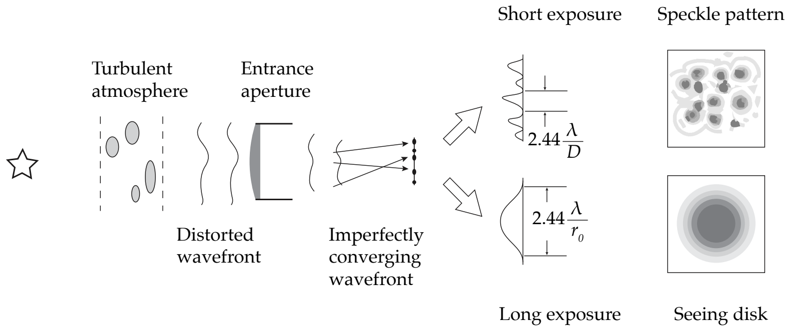

Figure 3.12: Seeing: Atmospheric distortions cause time-dependent distortions in the wavefronts reaching earth. These distortions average out during longer exposures and form a Gaussian seeing disk on the sensor (bottom right image). Adapted from [188].

The PSF describes the way an ideal point source is smeared out spatially after being subjected to atmospheric and optical effects in the telescope, and seeing is the FWHM of the PSF. The algorithm used to reconstruct the PSF of a ZTF exposure is `DAOphot` [189]. It fits an analytic Gaussian profile to all point sources within an image, and subtracts this analytic model. After subtraction, the model is updated with the residuals remaining. This process is repeated until the residuals are reasonably small. The background estimate needed for this procedure is extracted from the most common brightness value of a histogram of evenly distributed pixels within the image [189].

[189]: Stetson (1987), *DAOPHOT – A computer program for crowded-field stellar photometry*

After the PSF has been reconstructed, one can measure the brightness of all objects contained in the image. To do so, the global PSF is scaled to the source with a least-squares fit, and then its flux is extracted by integrating over the source's PSF profile [189].

Of course, PS1 needs to be calibrated in itself. That is done in two steps: A self-consistent relative calibration is created with `Ubercal` [190]. This relative calibration is in turn anchored to precisely measured Calspec standard stars. For details, see [191].

3.4.7 Photometric Calibration and Magnitudes

Not only the source positions in the images need to be calibrated (astrometric calibration), but their brightness as well (photometric calibration).

This means that one needs to convert the flux values derived by integrating over the PSF to physical flux.

To simplify matters, ZTF is photometrically calibrated against a reference survey, namely PS1. To perform the ZTF photometric calibration, a catalog of suitable calibrator stars from the PS1 survey has been curated. These stars were required to fulfill some basic quality criteria: They should be stable over multiple PS1 survey epochs in all PS1 filters excluding the y-band (g, r, i, z) and should be fairly bright, but not so bright that they saturate the ZTF sensors. Furthermore, they were required to be isolated to avoid blending with neighboring objects and need to have a high probability of being in fact stars, as opposed to galaxies [183].

The brightness of both ZTF and PS1 are measured in magnitudes. Magnitudes are somewhat counterintuitive, as a higher value corresponds to a fainter source.[6] The following definition holds:

Definition 3.4.1 *A star one magnitude brighter than another star is* $\sqrt[5]{100} \approx 2.512$ *times brighter*

If one uses the flux density (power per unit area) on Earth as a measure of brightness, it follows from Def. 3.4.1 that the difference between two objects with magnitudes m_1 and m_2 and respective flux densities f_1 and f_2 is:

$$m_1 - m_2 = -2.5 \log_{10} \frac{f_1}{f_2}. \tag{3.1}$$

Now we have a relative definition of a source's magnitude, but we need an absolute one. In other words, one needs to know what constitutes a magnitude of 0 (the zero point of the magnitude scale). ZTF uses AB magnitudes [194], which are defined via the spectral flux density f_ν with units of $\mathrm{W\,m^{-1}\,Hz^{-1}}$. In this system, the magnitude is a logarithm of the spectral flux density:

[194]: Oke et al. (1983), *Secondary standard stars for absolute spectrophotometry*

Definition 3.4.2 $m_{AB}(\nu) = -2.5 \log_{10} (f_\nu / 3631\,\mathrm{Jy})$

As one can see, a source with constant spectral flux density $f_{\nu,0} = 3631 \times 10^{-23}\,\mathrm{W\,cm^{-1}\,Hz^{-1}} = 3631\,\mathrm{Jy}$ has a magnitude of 0.[7] Telescopes like ZTF and PS1 use bandpass filters (see Section 3.3.1), so the spectral flux density needs to be integrated over the filter wavelengths. Therefore, the magnitude definition changes to [195]:

[195]: Tonry et al. (2012), *The Pan-STARRS1 Photometric System*

Definition 3.4.3 $m = -2.5 \log_{10} \frac{\int f_\nu (h\nu)^{-1} A(\nu)\, d\nu}{\int 3631\,\mathrm{Jy} (h\nu)^{-1} A(\nu)\, d\nu}$

Here, h is Planck's constant and $A(\nu)$ is the capture cross-section (i.e. the chance of an incoming photon to produce an electron in the detector[8]). ZTF does not use a precisely modeled response function $A(\nu)$, but relies on PS1. To first order, ZTF magnitudes are tied to the PS1 system via their zero point (ZP), which is an arbitrary offset that needs to be calculated for each image:

$$m_{\text{cal}} = m_{\text{instr}} + \text{ZP}. \tag{3.2}$$

[6] Their use can be traced back over 2000 years to Greek/Roman astronomers Hipparchus and Ptolemy [192]. They classified the brightest stars to be of 'first order' or 'first magnitude', with subsequently dimmer stars assigned lower magnitudes until the dimmest stars visible to the naked eye were of 'sixth magnitude'. The system stuck and was put on firm footing by Norman Pogson in 1856. He defined a star being 5 magnitudes brighter than another one to be 100 times brighter [193].

[7] The value of $f_\nu, 0$ is not entirely random: The traditional zero point was the star Vega. Vega's magnitude in the AB system as defined above, integrated over the V-band, is 0.03, close to the traditional 0. The AB system has the advantage of not relying on a physical source.

[8] More general, $A(\nu)$ is $A(\nu, \theta, t)$, also depending on the angle of the incoming photon and therefore the atmospheric column along the line of sight, as well as time.

To obtain m_{cal}, all extracted sources are spatially matched to the PS1 calibrator catalog. After creating a one-to-one relation between ZTF stars and PS calibrator stars, one can to first order calculate m_{cal}.

But there is a potential complication: The ZTF and PS1 filters are fairly similar, but do not exactly match. Eq. 3.2 assumes a constant spectrum which is not the case for most celestial objects. Therefore, an object will have slightly different brightness values when measured through the ZTF and the corresponding PS1 filter. To account for this, a linear, filter-dependent color correction c_f needs to be applied:

$$m_{\mathrm{cal}} = m_{\mathrm{instr}} + \mathrm{ZP}_f + c_f \times \mathrm{PS1}_{\mathrm{clr}}, \tag{3.3}$$

where $\mathrm{PS1}_{\mathrm{clr}}$ is filter dependent ($g_{\mathrm{PS1}} - r_{\mathrm{PS1}}$ for the ZTF g and r-band, as well as $r_{\mathrm{PS1}} - i_{\mathrm{PS1}}$ for the i-band).

In the photometric calibration step, ZP_f and c_f are chosen to globally minimize $\Delta m_f = \mathrm{ZP}_f + c_f \times \mathrm{PS1}_{\mathrm{clr}}$ for all calibrator stars in the respective image [183].

3.4.8 Image Subtraction

Because ZTF is a survey telescope deeply rooted in time-domain astronomy, many of the science goals concern observing sources that are new or changing. To do so, one needs to detect changes in the nightly observations with regard to reference images. From all new science images, these reference images are subtracted (see Fig. 3.13). All remaining detections constitute a temporal evolution with respect to the epoch of reference image creation. All reference images in ZTF are stacked images of 15–40 individual high-quality images, most of them created early after the start of telescope operations. Quality criteria for the individual images comprise good seeing, low errors on the astrometric and photometric calibrations, as well as background levels falling into filter-specific ranges [181].

[196]: Bertin (2010), *SWarp: Resampling and Co-adding FITS Images Together*

[197]: Zackay et al. (2016), *Proper Image Subtraction – Optimal Transient Detection, Photometry, and Hypothesis Testing*

`SWarp` [196] is used to interpolate and resample the reference image onto the science image, while subsequent subtraction and PSF photometry (see Section 3.4.6) is performed with the `ZOGY` algorithm [197].

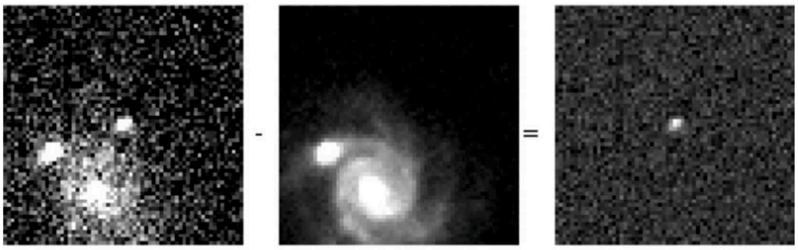

Figure 3.13: ZTF image subtraction. The new science image is on the left; in the middle is the reference image, which is subsequently subtracted from the science image. This results in a difference image, seen on the right. From [198].

3.4.9 Alert Packages

The last step in the imaging pipeline is alerting a group of upstream alert brokers. The information on new transients or updates on existing ones therefore need to be packaged into a convenient format. All positive detections after image subtraction with a Signal-to-Noise Ratio (SNR) greater than 5 are subjected to `RealBogus`. This is a machine learning (ML) algorithm trained to discriminate between 'real' events (most likely astrophysical) and 'bogus' events (e.g. subtraction artifacts, see Fig. 3.14) [198]. This algorithm generates an `rbscore`, ranking the probability that the detection is real.

[198]: Mahabal et al. (2019), *Machine Learning for the Zwicky Transient Facility*

Another ML algorithm has been employed to assign all PS1 sources an `sgscore`, separating between stars and galaxies based on their morphology and flux in all PS1 bands [199]. Both `rbscore` and `sgscore`, as well as cutouts of the science, reference and difference images are packed together and shipped as an alert package. These alerts also include the distances to the three nearest PS1 sources, the closest *Gaia* source, up to 30 days of previous detections if these exist, and some quality metrics like the limiting magnitude. The file format for distribution is Apache `Avro`[9], and the method of distribution is an Apache `Kafka`[10] stream [200].

[199]: Tachibana et al. (2018), *A Morphological Classification Model to Identify Unresolved PanSTARRS1 Sources: Application in the ZTF Real-time Pipeline*

[9] `https://avro.apache.org`

[10] `https://kafka.apache.org`

[200]: Patterson et al. (2018), *The Zwicky Transient Facility Alert Distribution System*

3.5 Surveys and Cadence

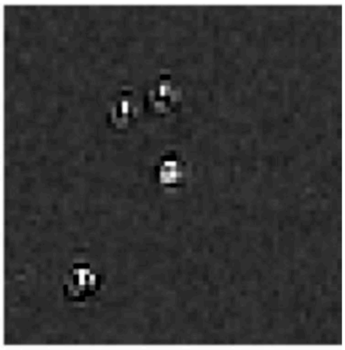

Figure 3.14: ZTF subtraction artifact, resulting in a bogus transient. From [198].

ZTF is supporting three main survey programs: First, there is the NSF Mid-Scale Innovations Program (MSIP) survey, which is allocated 40 % of telescope time. The MSIP survey is subdivided into the Northern Sky Survey, using a large fraction (85 %) of MSIP time. This survey covers the entire 23,675 deg^2 of the northern sky 7° above the galactic plane. As long as a field (see Section 3.3.2) is accessible, it is observed once in the g- and once in the r-band every 3 nights, with both images separated by at least 30 minutes to reject transients and moving (i.e. solar system) objects. The rest of the northern sky (galactic latitude $|b| \leq 7°$, with a footprint of 2800 deg^2) is visited twice per night, with the same observational parameters as the Northern Sky Survey [182].

Another 40 % are allocated to the partners of the ZTF collaboration. During the first year of ZTF operation, this allocation was dedicated to the Extragalactic High Cadence Survey, the i-band Survey, to ToO observations like the Neutrino Follow-up Program (see Chapter 4), to the Twilight Survey, the High-Cadence Plane Survey and the Asteroid Rotation Period Survey [182].

The last 20 % are private Caltech time, with surveys selected each semester by the Caltech Time Allocation Committee [182]. A prominent example is the observational campaign to follow up gravitational waves (GWs) detected by the LIGO and Virgo interferometers [201, 202] during their 3rd observational run (O3) [203].

[201]: Aasi et al. (2015), *Advanced LIGO*

[202]: Acernese et al. (2014), *Advanced Virgo: A second-generation interferometric gravitational wave detector*

[203]: Kasliwal et al. (2020), *Kilonova Luminosity Function Constraints Based on Zwicky Transient Facility Searches for 13 Neutron Star Merger Triggers during O3*

Left: The IceCube detector (Image credit: IceCube). Center: Mount Palomar (image credit: Caltech). Right: The `AMPEL` framework.

4 The ZTF Neutrino Follow-Up Program

After introducing the IceCube detector (see Chapter 2) and the Zwicky Transient Facility (see Chapter 3), we now have all the ingredients to introduce the ZTF high-energy neutrino follow-up pipeline.

IceCube sends out ~2.2 astrophysical high-energy neutrino alerts on average per month (see Section 2.4.3), mostly originating in the northern sky. Due to its large FoV, its location in the northern hemisphere and its completely robotic operation, ZTF is the ideal follow-up instrument for these alerts. It allows to cover the typically reported IceCube localization regions with only one pointing of the telescope.

The follow-up procedure can be outlined as follows: An IceCube alert is received. If the alert meets our alert quality criteria, we perform an observability check. If the sky region is accessible to ZTF, we observe. After observations, we filter candidates, followed by visual inspection. Lastly, we acquire forced photometry and—if needed—trigger additional follow-up.

4.1 Source Classes

Given the limit observational time at our disposal, we are only sensitive to a subset of possible high-energy neutrino source classes. For a complete overview, see Section 1.4. These are:

Optical AGN flares As we are decidedly an optical and time-domain program, we restrict ourselves to AGN (including blazar) undergoing significant optical flaring activity around the time of neutrino detection.

CCSNe We are sensitive to the subtypes of CCSNe that have been proposed as potential high-energy neutrino sources, i.e. CCSNe with signs of CSM interaction, and CCSNe that launch jets. This class is heavily contaminated by SNe Ia. Because of this, and because not all types of CCSNe are predicted to emit neutrinos, the exact type of CCSN needs to be determined with spectroscopic follow-up (see Section 4.7).

Short and long GRBs We are both sensitive to jets launched by CCSNe (long GRBs) and binary mergers (sGRB), given they are optically detectable. Both the emerging CCSN and the rapidly fading red optical emission by sGRBs are good signatures to test, though the sGRB afterglow might evolve too quickly for us to catch it.

TDEs Lastly, we are also sensitive to TDEs. Due to the significant model uncertainties, we do not restrict ourselves in terms of how old the TDE must be at the time of neutrino detection, provided it is still active.

All other possible sources are not captured by this follow-up program, because they are either evolving too quickly (sGRBs might fall under this category), too faint, or the neutrino emission might not be correlated to optical emission.

4.2 Alert Cuts

Two factors motivate imposing additional cuts on the alert stream received via GCN notices: We only have limited telescope time, and too large uncertainty areas would result in the significance of potential counterparts dropping too much. As outlined in Section 2.4.2, there are two alert streams: Gold alerts with an average purity of 50 % (36 % of all non-retracted alerts so far belonged to that category), and Bronze alerts with an average purity of 30 % (64 % of all alerts).

In general, we only follow up if the reported bounding rectangle comprising the 90 % uncertainty region (see Section 2.3.7) covers less than 30 deg^2. All Gold alerts making this cut qualify for follow-up. There is a more stringent cut on the uncertainty areas of Bronze alerts, requiring their bounding rectangle to be $< 10\,\mathrm{deg}^2$ in size. To avoid contamination by foreground stars, we implemented a cut on galactic latitude of $|b| > 10°$.

4.3 Observation Planning with `planobs`

If an alert makes these cuts, one needs to check if observations with ZTF are feasible, and create an observation plan. As we are interested in the time evolution of potential source candidates, usually observations within a 10-day window are triggered. To obtain deep images during the first night, we first trigger 300 s exposures in the g- and r-band. These are followed by shallower 30 s observations in the g-band during nights 2, 3, 5 and 7, and finally shallow observations in both g- and r-band during night 9.

To reduce the potential for errors, the feasibility checks and creation of an observation schedule are done with the `planobs` [204] tool, developed and maintained by the author to automate triggering follow-up observations as much as possible.

[204]: Reusch et al. (2023), *planobs: Release 0.6.4*

`planobs` was written in Python, and has been deployed on a virtual private server. The backend runs on `Flask`[1] behind an `nginx`[2] reverse proxy exposed to the internet. The `nginx` endpoint serves a Slack[3] bot

[1] `https://flask.palletsprojects.com`

[2] `https://nginx.com`

[3] `https://slack.com`

integrated into the DESY multimessenger group chat for ease of use. It can also be run locally.

The tool was designed to be used with a simple command-line style interface in the working group Slack chat. For example, issuing

```
Plan IC221223A -multiday
```

in the `plan_observations` Slack channel will create an observation plan for IceCube high-energy neutrino *IC220501A*.

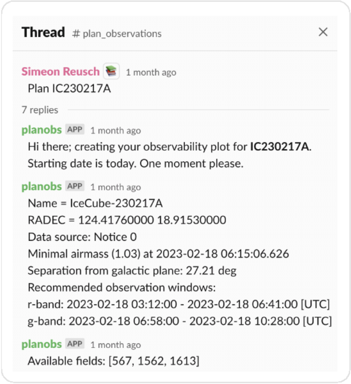

Thread # plan_observations

Simeon Reusch 1 month ago
Plan IC230217A

7 replies

planobs APP 1 month ago
Hi there; creating your observability plot for **IC230217A**.
Starting date is today. One moment please.

planobs APP 1 month ago
Name = IceCube-230217A
RADEC = 124.41760000 18.91530000
Data source: Notice 0
Minimal airmass (1.03) at 2023-02-18 06:15:06.626
Separation from galactic plane: 27.21 deg
Recommended observation windows:
r-band: 2023-02-18 03:12:00 - 2023-02-18 06:41:00 [UTC]
g-band: 2023-02-18 06:58:00 - 2023-02-18 10:28:00 [UTC]

planobs APP 1 month ago
Available fields: [567, 1562, 1613]

Figure 4.1: Sample interaction with `planobs` in Slack, checking the observability of *IC230217A*.

To obtain the positional and error information on the neutrino in question from the respective GCN Circular, `planobs` is searching for the GCN on a server hosted by NASA[4]. As circulars are written by humans, it fuzzily parses them to extract the relevant information. After this, observability for ZTF at Mt. Palomar is calculated, defaulting to the current time (optionally, a desired observation time can be requested).

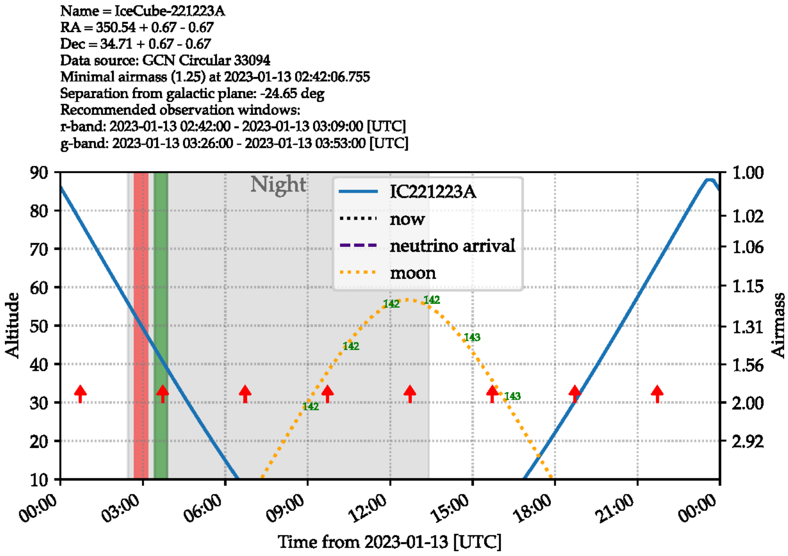

Name = IceCube-221223A
RA = 350.54 + 0.67 - 0.67
Dec = 34.71 + 0.67 - 0.67
Data source: GCN Circular 33094
Minimal airmass (1.25) at 2023-01-13 02:42:06.755
Separation from galactic plane: -24.65 deg
Recommended observation windows:
r-band: 2023-01-13 02:42:00 - 2023-01-13 03:09:00 [UTC]
g-band: 2023-01-13 03:26:00 - 2023-01-13 03:53:00 [UTC]

Figure 4.2: Observation plot created by `planobs` for the follow-up of IceCube neutrino *IC221223A*. The altitude and airmass of the alert region at Mt. Palomar are shown in blue. The red and green shaded regions are proposed observation windows in the *r*- and *r*-band. The red arrows show the airmass limit of 2.0, and the moon is displayed as yellow dotted curve. The gray shaded region marks nighttime at the telescope site. All information automatically extracted from the GCN Circular are shown on top.

An exemplary observability plot is shown in Fig. 4.2. The blue curve shows the altitude of the neutrino sky region above Mt. Palomar, and the red and green shaded regions mark the two proposed observation windows in the *r*- and the *g*-band.

Additionally, for each field of the primary and secondary grid (see Section 3.3.2) that overlaps with the uncertainty region, `planobs` checks if reference images are available (see Section 3.4.8), calculates the resulting coverage and selects the field with the highest coverage. Fig. 4.3 shows such an overlay plot for *IC221223A* and ZTF field 693.

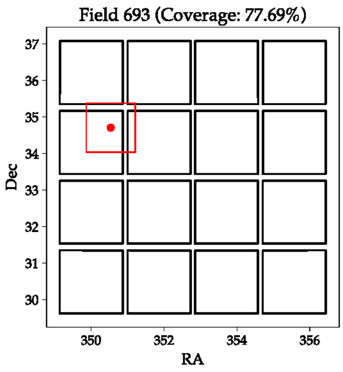

Figure 4.3: The bounding rectangle of the 90 % uncertainty area of *IC221223A* overlaid onto the ZTF grid. The coverage does not equal 100 % because chip gaps are taken into account.

If the plan looks good, i.e. if the object is observable above an airmass of 2.0 and the moon is not too close, one needs to invoke a command like

```
Plan IC221223A -trigger
```

to submit the observation request via a dedicated API to the ZTF telescope scheduler. There is additional functionality to ensure that the trigger has been added to the telescope queue. If all goes well and weather permits, the observations are carried out, and one can proceed to do candidate vetting.

---

[4] https://heasarc.gsfc.nasa.gov/wsgi-scripts/tach/gcn_v2/tach.wsgi

4.4 The `AMPEL` Broker

The next step in the pipeline is the selection of good candidates. ZTF typically serves on the order of 200,000 alerts per night. As only a fraction of these is relevant for the neutrino follow-up, software is needed to cut down the number of alerts. This is precisely what `AMPEL` [205] is doing, a streaming data analysis framework developed at Humboldt-University Berlin and DESY Zeuthen with contribution of the author.

[205]: Nordin et al. (2019), *Transient processing and analysis using AMPEL: Alert management, photometry, and evaluation of light curves*

The main design goals of `AMPEL` comprise scalability, modularity and provenance tracking. It was built with the data rate of the future Vera C. Rubin observatory [175] in mind, and was subsequently selected as one of seven brokers[5] for Rubin observatory. The full software stack is written in Python.

[175]: Ivezić et al. (2019), *LSST: From Science Drivers to Reference Design and Anticipated Data Products*

[5] See `https://www.lsst.org/scientists/alert-brokers`.

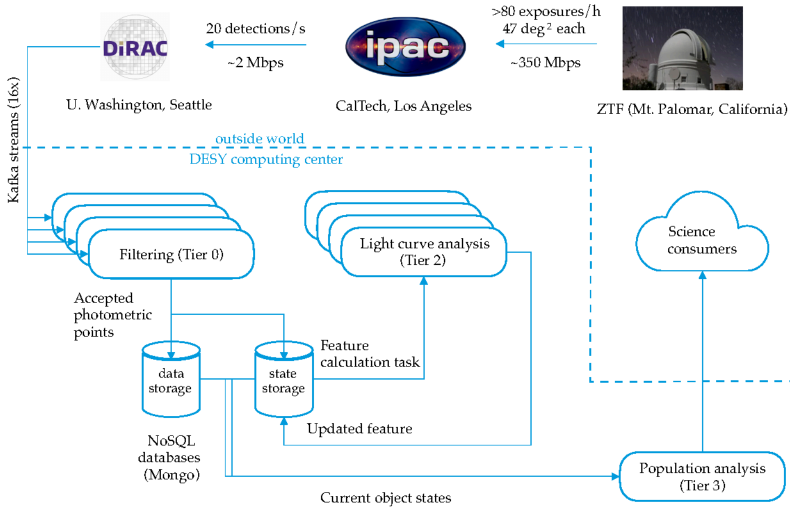

Figure 4.4: Overview of the `AMPEL` data processing. Alerts from DiRAC are ingested into `AMPEL`, where they are processed, combined and analyzed; any subsequent results can be disseminated to science consumers. From [205]

Fig. 4.4 shows the design and information flow of `AMPEL`. On the top, data from Mt. Palomar is transmitted to IPAC (see Section 3.4.1), where detections are extracted from the difference images. These are then sent to the Institute for Data Intensive Research in Astrophysics & Cosmology (DiRAC) at the University of Washington, where they are distributed via parallel `Kafka` streams. This is the live data stream the ZTF instance of `AMPEL` listens to.

The first of several execution layers ('tiers') is the Filtering stage (Tier 0). Here, different filters can be implemented, reducing the large number of alerts by different criteria. These comprise e.g. `RealBogus` and `sgscore` (see Section 3.4.8), color evolution, host galaxy properties and the (non-)existence of a detection history. All alerts surviving the filtering stage are then stored in a `MongoDB`[6] database collection. Additionally, processing states are stored in another collection.

[6] `https://mongodb.com`

A description of Tier 1 can be skipped here, as it serves mainly technical purposes (for details, see [205]). The next relevant stage is the Light curve analysis stage (Tier 2). Here, additional information on the transients are either obtained or generated. Possible steps are querying external catalogs for host galaxy or redshift information, fitting light curves with various models or photometrically classifying transients with machine learning methods [205].

The last level, Tier 3, executes population analyses. These are schedulable actions, triggered on request or at pre-defined times (ranging from yearly data dumps, through daily updates to nearly real-time execution). A typical use case is the automated ranking of different transients for a specific science goal. An example is the daily posting of new supernova candidates to a Slack channel, ranked by how promising they are for spectroscopic follow-up [205].

To simplify matters significantly and to allow reprocessing as well as full replayability, all ZTF alerts received via the `Kafka` streams since June 2018 are also stored in an archival alert database hosted at DESY. The database is based on `Postgresql`[7] and can be accessed via a web frontend and a REST API[8].

[7] https://postgresql.org/

[8] https://ampel.zeuthen.desy.de/api/ztf/archive/v3/docs

4.5 Candidate Filtering with `nuztf`

To streamline the neutrino follow-up process, `nuztf` [206] was created for filtering and inspecting candidate counterparts, with significant contribution by the author. It was written in Python, and relies heavily on `AMPEL` for filtering purposes (see the previous Section 4.4). The main use case is the follow-up of high energy neutrinos, but it can also handle skymaps from LIGO-Virgo-Kagra (LVK) gravitational wave (GW) alerts or GRB skymaps, and was used in the ZTF follow-up campaign during the third observational campaign (O3) of the Ligo and Virgo interferometers [203].

[206]: Stein et al. (2023), *nuztf: v2.6.5 Release*

When run, `nuztf` executes the following steps: (1) It obtains the IceCube alert information from the GCN, (2) constructs a `HEALPix` map from the 90 % uncertainty rectangle, (3) queries the `AMPEL` archive for all alerts within the uncertainty area `HEALPix` map[9] in a given time range, (4) applies the `AMPEL` `DecentFilter` [205] with custom parameters (see next paragraph), (5) crossmatches the surviving counterpart candidates to a list of catalogs, (6) pushes the candidates to `Fritz`[10] and finally (7) creates an overview PDF file with details and a light curve for each candidate.

[9] Because the archive database contains `HEALPix` indices with different resolutions, this kind of query is much faster than a cone search.

[10] https://fritz.science

4.5.1 `DecentFilter` Parameters

The GCN parsing is done akin to `planobs` (see Section 4.3). After extraction of the alert information and querying the Archive database with a `HEALPix` map, a first filter, `DecentFilter`, is run with the following parameters:

Time window The transient must have shown activity in a 14-day window after neutrino detection.

`RealBogus` The transient must have a `RealBogus` score of > 0.3, i.e. a high probability of being real.

Positive subtraction Sometimes, subtraction from reference images result in negative flux. This criterion ensures that there is excess flux.

Detections The candidate must have at least 2 detections, separated by at least 15 min. Note that this needs to be reflected in the observation planning. `planobs` (see Section 4.3) takes care of that.

`sgscore` The transient must have an `sgscore` < 0.8, which means a low probability of being a star.

Maximum distance to PS1 As noted in Section 3.4.9, the `sgscore` star-galaxy classifier is trained on data from PS1. To ensure correct association of the `sgscore` value which is based on the closest PS1 source, we require the source to be closer than 1°. The veto is also ignored if 3 PS1 sources are closer than 3° and their `sgscore` lies between 0.4 and 0.6, as this hints at possible PS1 source confusion.

***Gaia* star veto** Veto the source if the probability of it being a star is high, based on parameters from the *Gaia* survey.

The vast majority of alerts do not pass these filters, which significantly reduces the human effort needed in manual candidate vetting.

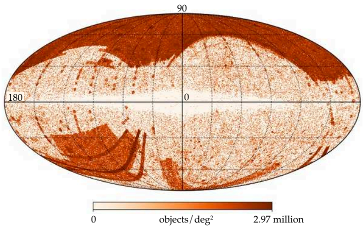

Figure 4.5: Distribution of the 8.9 million NED objects with spectroscopic redshift (as of November 2021). From `https://ned.ipac.caltech.edu/Documents/Holdings/graphics`.

4.5.2 Catalog Crossmatching

All surviving candidates are then spatially crossmatched to a set of catalogs. These are:

CRTS A catalog derived from the Catalina Realtime-Transient Survey (CRTS) [207], the Catalina Surveys Periodic Variable Star Catalog [208], is queried to check if the source is a variable star.

Gaia *Gaia* Data Release 2 (*Gaia* DR2) [209] is also used to crossmatch with known stars.

MILLIQUAS The Million Quasar Catalog (MILLIQUAS) [210] contains 1.4 million quasi-stellar object (QSO) candidates. Crossmatching is done to see if the source candidate is most likely associated to an AGN.

NED The NASA/IPAC Extragalactic Database (NED)[11] hosts information on 1.7 billion objects, compiled from various surveys and catalogs. Almost 9 million of those have a spectroscopic redshift.

SDSS Data Release 10 (DR10) [211] of SDSS is used to check if the object is flagged as potential star.

TNS The Transient Name Server (TNS)[12] is the repository of astrophysical transients and contains over 100,000 objects, both classified and unclassified. It is queried to check if the candidate source is a known and possibly classified transient.

WISE Color information from the Wide-Infrared Survey Explorer (*WISE*) Mission [212] can be used to identify likely AGN, as these predominantly occupy a small section of the WISE color-color diagram (see e.g. [213] for this).

None of those catalog matches serve as veto, but are rather appended to the final output. This can be motivated by the insecure classification of MILLIQUAS (QSOs) and SDSS (stars), which are partly derived from machine learning and therefore warrant human inspection. Also, some of the additional information can strengthen the case for a candidate source—for example, an existing classification on TNS as a certain type of supernova.

Fig. 4.6 shows a page from the final PDF file generated for neutrino *IC220624A*, displaying candidate source *ZTF21aauwmgr*. All information required for quickly deciding if the candidate warrants further scrutiny is displayed: Image cutouts of the science, reference and difference image are shown to allow identifying potential subtraction artifacts. The different

[207]: Drake et al. (2009), *First Results from the Catalina Real-Time Transient Survey*

[208]: Drake et al. (2014), *The Catalina Survey Periodic Variable Star Catalog*

[209]: A. G. A. Brown et al. (2018), *Gaia Data Release 2*

[210]: Flesch (2023), *The Million Quasars (Milliquas) Catalogue, v8*

[11] `https://ned.ipac.caltech.edu`

[211]: Ahn et al. (2014), *The Tenth Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the SDS-III Apache Point Observatory Galactic Evolution Experiment*

[12] `https://wis-tns.org/`

[212]: Wright et al. (2010), *The Wide-Field Infrared Survey Explorer (WISE): Mission Description and Initial On-Orbit Performance*

[213]: Hviding et al. (2022), *A New Infrared Criterion for Selecting Active Galactic Nuclei to Lower Luminosities*

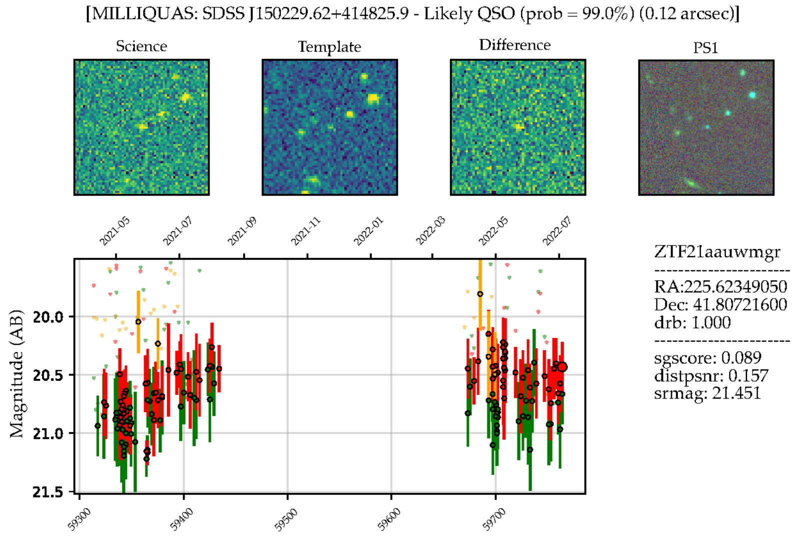

Figure 4.6: Sample output from `nuztf` showing the light curve of ZTF21aauwmgr, a transient selected as potential source in the follow-up of *IC220624A*. The cutouts at the top show the science image, the reference image (template), the resulting difference image and a PS1 cutout of the same region. The bottom shows the transient light curve, as well as the position, the `RealBogus` score (`drb`) and the `sgscore`. On the top, crossmatching information is shown. Here, the source is a source discovered by SDSS and flagged as likely QSO in MILLIQUAS.

ML scores (star-galaxy separation, `RealBogus`) are also displayed, as well as results from the catalog matching and of course the transient light curve.

In addition to the PDF file, the draft of a GCN Circular is created, pre-filled with the observation times, the coverage and all final candidates to allow for quick distribution of promising sources. Also, all candidates passing the final round of filtering are pushed to the collaborative time-domain astronomy portal `Fritz` for bookkeeping and further discussion.

The code can either be run locally, or through the interaction with a bot running on Slack. This bot is hosted by the author on a virtual private server. The endpoint exposed to Slack is run with `gunicorn`[13], a Python WSGI server, behind an `nginx` reverse proxy. The bot can parse the names of IceCube alert neutrinos and gravitational wave event names, extract the information from the GCN or GW map and upload the overview PDF and the GCN draft to Slack.

[13] `https://gunicorn.org`

4.6 Forced Photometry with `fpbot`

The last step in the pipeline is the acquisition of forced photometry. This is achieved with `fpbot` [214], a forced photometry pipeline written by the author for ZTF, built upon `ztflc`[14].

[214]: Reusch (2023), *fpbot: Release 1.1.2*

[14] `https://github.com/mickaelrigault/ztflc`

4.6.1 Forced Photometry Explained

The usual extraction of transient flux, as performed by the ZTF imaging pipeline (see section 3.4.8), cannot a priori know at which position to expect flux in the difference image. In the case of ZTF, there is a signal-to-noise threshold of 5. When the flux detected at a position in the image exceeds that threshold, an alert is generated.

If the position of the transient is known, as it has already generated some alerts, one can use this location to reprocess the difference images and extract flux at precisely this position. With this method, one can extract a signal that is fainter than the signal-to-noise threshold. In most cases, this means that the transient can be detected at a younger age than its earliest alert photometry detection, or that the tail of the light curve of a fading transient can be detected longer. This process is dubbed 'forced photometry', as the known position is 'forced' upon the flux extraction.

The possibility to extract flux prior to the first alert detection makes forced photometry especially useful to accept or reject some candidate neutrino sources: If e.g. a stripped-envelope supernova (these are compatible with high-energy neutrino production, see Section 3) shows an early forced photometry detection days prior to the neutrino arrival, the source can be rejected as candidate (see e.g. Section 4.9.1).

4.6.2 The `fpbot` Pipeline

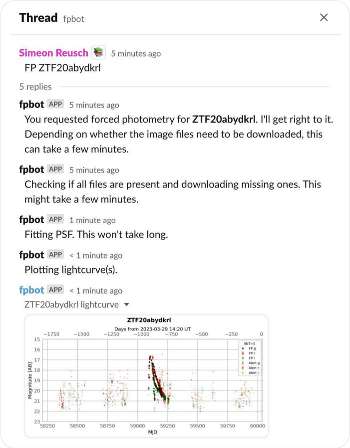

Simeon Reusch 5 minutes ago
FP ZTF20abydkrl

5 replies

fpbot 5 minutes ago
You requested forced photometry for **ZTF20abydkrl**. I'll get right to it. Depending on whether the image files need to be downloaded, this can take a few minutes.

fpbot 5 minutes ago
Checking if all files are present and downloading missing ones. This might take a few minutes.

fpbot 1 minute ago
Fitting PSF. This won't take long.

fpbot < 1 minute ago
Plotting lightcurve(s).

fpbot < 1 minute ago
ZTF20abydkrl lightcurve ▾

Figure 4.7: Sample interaction with the `fpbot` Slack bot, obtaining forced photometry for *ZTF20abydkrl*.

`fpbot` was written in Python by the author and provides a multithreaded pipeline to obtain forced photometry extracted from ZTF difference images. It employs a `MongoDB` database to store the download and processing status for transients to avoid multiple downloads and unnecessary refits. Fitting requests can either be issued locally with a command-line interface, or within a dedicated Slack channel to serve the broader community.

For each object that is processed, first the `AMPEL` archive is queried for alerts of this object. From these, the median sky location is computed independently for each band (g, r and i), as the stable astrometric solution deviated from band to band when testing the pipeline.

Difference image cutouts at the desired location as well as the PSF shape images are then downloaded from IPAC in parallel, maximizing throughput. After this stage, the `ztflc` package is used to measure the flux in the difference image cutout, given the median band location and the PSF image. This is done using the `iminuit` package [215], running the least-square fit algorithm `MIGRAD` [216].

[215]: Dembinski et al. (2023), *scikit-hep/iminuit*

[216]: James et al. (1975), *Minuit – A system for function minimization and analysis of the parameter errors and correlations*

The fit results are stored as CSV files and returned either via Slack channel or email. If another forced photometry request is issued for the same object again, only new epochs are downloaded and analyzed to speed up the process. An instance of the service, hosted at DESY, is also used by the cosmology and supernova ZTF working groups.

4.7 Spectroscopic Resources

To classify promising counterpart candidates, our group has successfully submitted proposals for spectroscopic resources during the run of the neutrino follow-up program. Several of these proposals were led by the author.

The instruments that were available to our group for ToO observations during the last four years comprised the following instruments, in ascending order with respect to mirror size:

SEDM on the Palomar P60 telescope (1.5 m), the Alhambra Faint Object Spectrograph and Camera (ALFOSC) on the Nordic Optical Telescope (NOT, 2.6 m) [217] on La Palma, the Device Optimized for the LOw RESolu-

[217]: Djupvik et al. (2010), *The Nordic Optical Telescope*

tion (Dolores) on the Telescopio Nazionale Galileo (TNG, 3.6 m) [218], also on La Palma, the Gemini Multi-Object Spectrograph (GMOS-N) [219] on the Gemini North Telescope (8.1 m) in Hawaii, Multi-Object Double Spectrograph (MODS) [220] on the Large Binocular Telescope (LBT, 2x8.4 m) in Arizona, the Low-Resolution Imaging Spectrometer (LRIS) [221] on the W. M. Keck Observatory (Keck, 10 m) in Hawaii and lastly the Optical System for Imaging and low-Intermediate-Resolution Integrated Spectroscopy (OSIRIS) [222] on the Gran Telescopio Canarias (GTC, 10.4 m) on La Palma.

The image reduction pipelines used were either `pyraf`[15], a Python wrapper to IRAF [223], or `pypeit` [224], a Python-based reduction package. The latter was used for NOT/ALFOSC, TNG/Dolores, Gemini/GMOS-N, LBT/MODS and GTC/OSIRIS, while `pyraf` was used for Keck/LRIS.

4.8 Follow-Up Performance

As of March 2023, we have followed up 34 of the 108 non-retracted high-energy neutrino alerts issued by IceCube (see Section 2.4) since the start of the program, which amounts to 31 % of all alerts. An overview is shown in Fig. 4.8. The majority of alerts not followed up were only visible during the day ('Proximity to the sun'), or did not meet our alert quality criteria (see Section 4.2). The rest were either located too close to the galactic plane ($|b| \leq 10°$) or the telescope was on hold due to bad weather or technical problems.

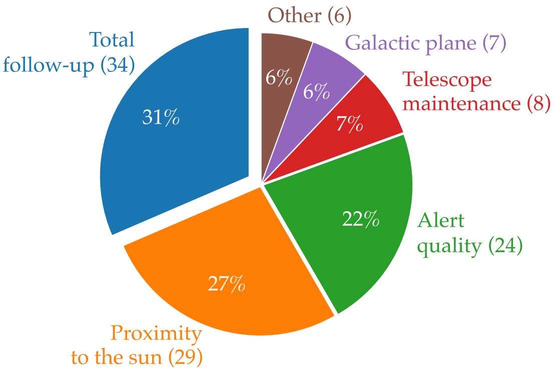

Figure 4.8: Breakdown of the follow-up performance as of March 2023. Out of 108 neutrino alerts, 34 were followed up. The 'Other' category comprises low altitude (below our airmass limit of 2.0 or entirely in the southern sky) and bad weather.

An overview of all follow-up campaigns is shown in Table A.1 in the appendix. In total, 205 candidate sources were selected by `nuztf`. Of these, 27 (13 %) were distributed via GCN Circulars. The classification of these candidates will be discussed in Section 4.9.1.

Additionally, the follow-up performance over time is displayed in Fig. 4.9, with the followed-up alerts shown in green. As one can see, the performance has slightly deteriorated in the last two years, which can be mostly attributed to telescope downtime due to maintenance or technical faults.

[218]: Mancini et al. (1997), *Italian National Galileo Telescope (TNG) system description and tracking performance in the workshop*

[219]: Hook et al. (2004), *The Gemini–North Multi-Object Spectrograph: Performance in Imaging, Long-Slit, and Multi-Object Spectroscopic Modes*

[220]: Pogge et al. (2010), *The multi-object double spectrographs for the Large Binocular Telescope*

[221]: Oke et al. (1995), *The Keck Low-Resolution Imaging Spectrometer*

[222]: Cepa et al. (2003), *OSIRIS tunable imager and spectrograph for the GTC. Instrument status*

[15] `https://github.com/iraf-community/pyraf`

[223]: Tody (1986), *The IRAF Data Reduction And Analysis System*

[224]: J. X. Prochaska et al. (2020), *PypeIt: Release 1.0.0*

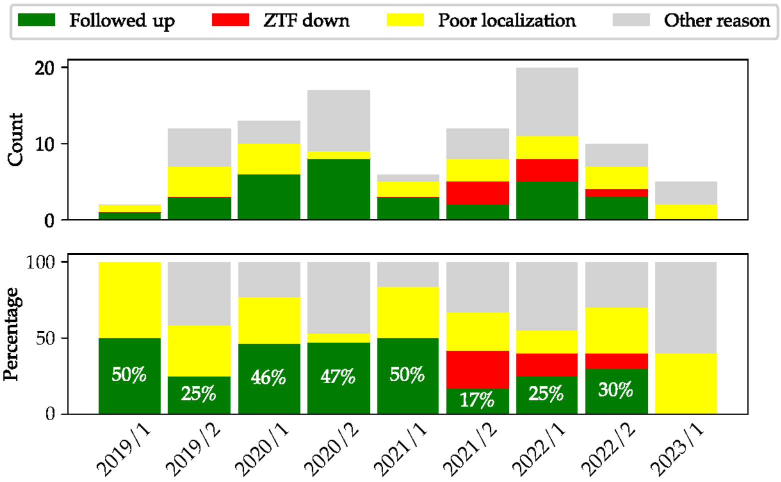

Figure 4.9: Follow-up performance over time, shown in half-year bins. The alerts we did follow up on are shown in green, while poorly localized alerts not meeting our quality cuts are shown in yellow, alerts not followed up due to telescope downtime are shown in red, and alerts rejected for other reasons are shown in gray.

4.9 Notable Sources

A full overview over the first 24 follow-up campaigns can be found in [225]. In the following, those sources will be highlighted to which I contributed either by requesting spectroscopic observations, reducing spectra or vetting candidates.

[225]: Stein et al. (2023), *Neutrino Follow-Up with the Zwicky Transient Facility: Results from the first 24 Campaigns*

4.9.1 *SN2020lls*: Type Ic Supernova

SN2020lls was identified in coincidence with high-energy neutrino *IC200530A* [226]. 87 % (22.05 deg^2) of the 90 % rectangular uncertainty area were serendipitously observed by ZTF roughly 10 min after the neutrino detection. These observations were later appended by deep 300 s observations Three candidate sources were identified, including *SN2020lls* [227, 228].

We took a spectrum of *SN2020lls* with NOT / ALFOSC on June 12, 2020 and used the Supernova Identification (`SNID`) code [229] to fit spectra of different types supernovae to the spectrum. The best fit spectrum was a Type Ic supernova at redshift $z = 0.041$, 14 days post peak. The spectrum and the template are shown in Fig. 4.10.

[226]: Stein (2020), *IceCube-200530A – IceCube observation of a high-energy neutrino candidate event*

[227]: Reusch et al. (2020), *IceCube-200530A: Candidate Counterparts from the Zwicky Transient Facility*

[228]: Reusch et al. (2020), *IceCube-200530A – SN2020lls likely unrelated*

[229]: Blondin et al. (2007), *Determining the Type, Redshift, and Age of a Supernova Spectrum*

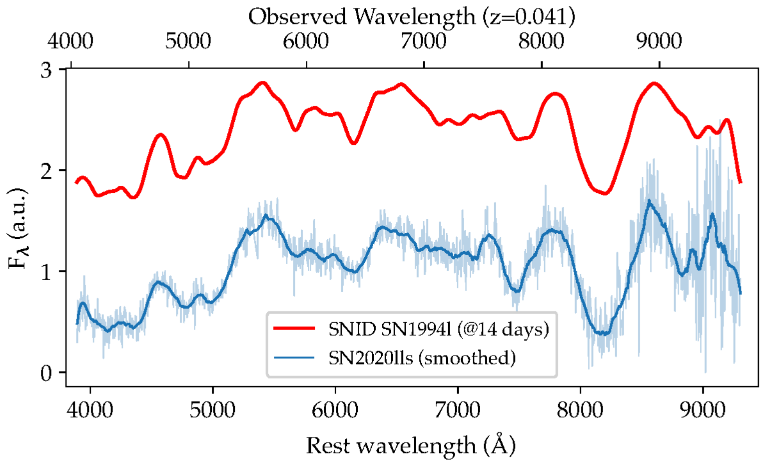

Figure 4.10: Spectrum of *SN2020lls* taken with NOT / ALFOSC. The smoothed spectrum is shown in blue on the bottom, while the template spectrum of the Ic type supernova *SN1994I* is shown on top in red. Figure by the author, see also [225].

SNe Ic are possible emitters of neutrino, but only shortly after explosion time (see Section 3). We therefore used the Modular Open Source Fitter for Transients (`MOSFiT`) [230] to fit the light curve. The light curve, including an early forced photometry detection in the *i*-band, can be seen in Fig. 4.11.

[230]: Guillochon et al. (2018), *MOSFiT: Modular Open Source Fitter for Transients*

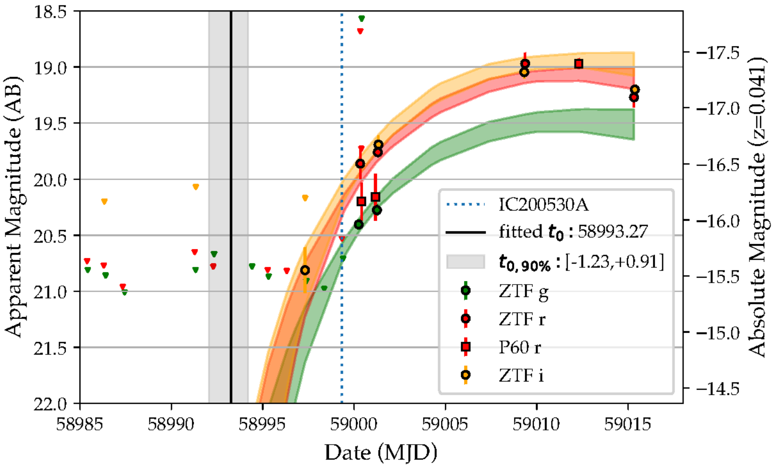

Figure 4.11: Light curve of *SN2020lls*. Detections are displayed as filled circles or squares, while upper limits are shown as triangles. The early forced photometry detection in the *i*-band constrained the explosion time estimated by `MOSFiT` fairly well to a time days prior (black line) to the detection of the neutrino (blue dotted line). This ruled out *SN2020lls* as candidate source. Figure by the author, see [225].

As can be seen, the fit estimates the explosion time (black line) at 6 days earlier then the neutrino arrival time (blue dotted line). This allowed us to rule out *SN2020lls* as source of *IC200530A*.

4.9.2 *SN2020lam*: Type IIP SN without CSM-interaction

A second source candidate—also coincident with *IC200530A*—was *SN2020lam* [231], displaying a supernova-like light curve. To obtain a classification, we took a spectrum with NOT/ALFOSC on June 6, 2020. Template fitting with `SNID` and `GELATO` [232] identified the source as a Type IIP supernova (see Fig. 4.12).

[231]: Reusch et al. (2020), *IceCube-200530A – SN2020lam classified as SN II without evidence of CSM interaction*

[232]: Harutyunyan et al. (2008), *ESC supernova spectroscopy of non-ESC targets*

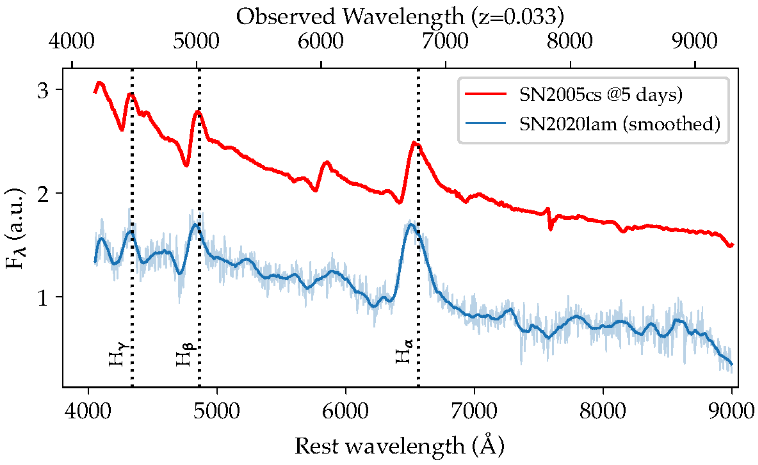

Figure 4.12: Spectrum of *SN2020lam* taken with NOT/ALFOSC. The smoothed spectrum is shown in blue on the bottom, while the template spectrum of the IIP type supernova *SN2005cs* is shown on top in red. Balmer lines are marked with black dotted lines. Figure by the author, see [225].

The light curve and the best-fit template show that the neutrino was detected close to the peak of the supernova, a Type IIP. The viable production mechanism in this case is CSM-interaction (see Section 1.4.1). As the spectrum showed no narrow lines, we inferred that there is no sign of CSM interaction, and consequently ruled out *SN2020lam* as a source candidate.

4.9.3 *AT2020ybb*: Unclassified

AT2020ybb was discovered when following up *IC201021A* [233]. Observations of the sky area started 44 h after the neutrino detection and covered 91 % of it (6.3 deg^2). The transient was a promising candidate, as the light curve of the source was compatible with a supernova around peak [234]. The other possibility was AGN activity, as the source was flagged as QSO with 98 % probability within MILLIQUAS.

[233]: Lagunas Gualda (2020), *IceCube-201021A: IceCube observation of a high-energy neutrino candidate event*

[234]: Stein et al. (2020), *IceCube-201021A: One candidate counterpart from the Zwicky Transient Facility*

We triggered GTC/OSIRIS and reduced the spectrum with `pyraf`. The spectrum is shown in Fig. 4.13. Unfortunately, the spectrum did not allow a robust classification of the source, neither confirming nor rejecting the neutrino source hypothesis. If the line observed at 7293 Å is interpreted as SII, the resulting redshift is $z = 0.087$.

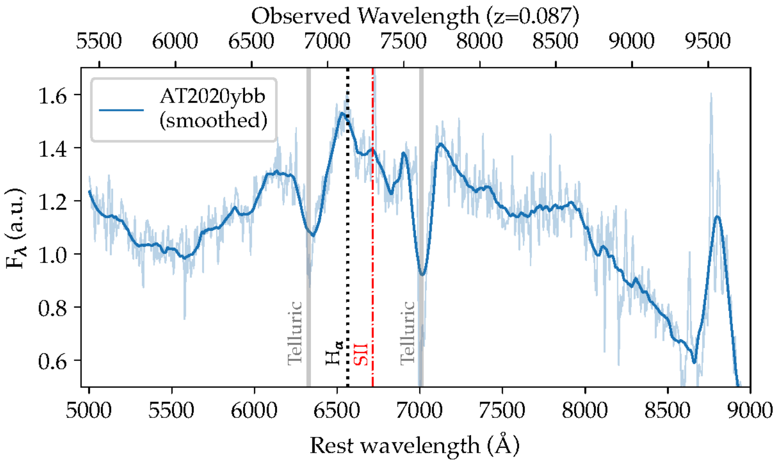

Figure 4.13: Spectrum of *AT2020ybb* taken with GTC/OSIRIS. The telluric absorption regions are shown in gray, potential hydrogen and SII lines are shown as black dotted and red dash-dotted lines. The resulting redshift is $z = 0.087$; the spectrum allowed no source classification.

4.9.4 *AT2021osi*: Regular AGN Activity

AT2021osi was detected when following up *IC210629A* [235]. It was first reported by `ALeRCe` [236], about two weeks prior to the neutrino's arrival.

[235]: Santander (2021), *IceCube-210629A – IceCube observation of a high-energy neutrino candidate track-like event*

[236]: Sánchez-Sáez et al. (2021), *Alert Classification for the ALeRCE Broker System: The Light Curve Classifier*

The proximity of the source to the nucleus of its host galaxy stipulated a TDE origin [237]. To classify the transient, we triggered two 750 s observations with Gemini/GMOS-N and reduced the spectrum with `pypeit`.

[237]: Necker et al. (2021), *IceCube-210629A: One Candidate Counterpart from the Zwicky Transient Facility*

The spectrum showed features typical for AGN, most importantly broad Balmer lines (H_α, H_β and H_γ). These allowed also to infer a redshift of $z = 0.194$ [238]. As there was no contemporaneous flaring activity in the optical light curve, we ruled out this candidate as regular AGN activity.

[238]: Reusch et al. (2021), *IceCube-210629A – Spectroscopic observation of AT2021osi*

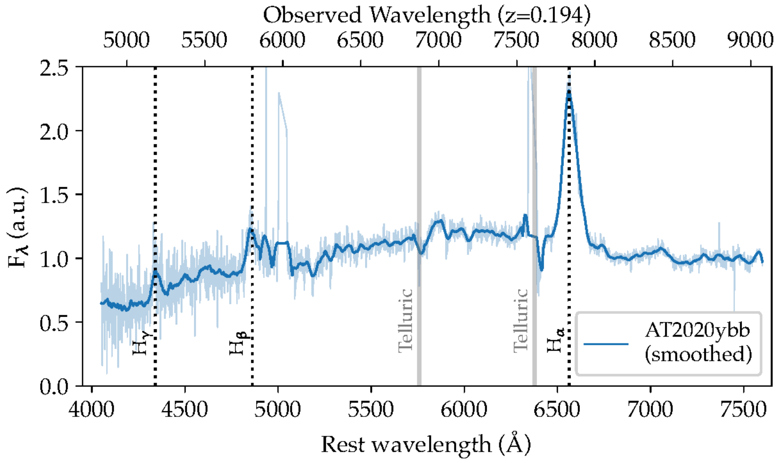

Figure 4.14: Spectrum of *AT2020osi*, taken with Keck/GMOS-N. The telluric absorption areas are shown in gray, and Balmer lines are shown as black dotted lines. The redshift is $z = 0.194$, and we classified the source as AGN activity.

4.9.5 *SN2022oyn*: Type Ia Supernova

When following up *IC220907A* [239], *SN2022oyn* emerged as notable source candidate, as its light curve was consistent with a supernova about 40 days post peak [240].

We obtained a spectrum with Keck/LRIS on September 20, 2022. Based on a fit with `SNID` (see Fig. 4.15), we classified the source as Type Ia supernova. We therefore ruled it out as neutrino source, as SNe Ia are not predicted to emit high-energy neutrinos (see Section 3).

[239]: Lincetto (2022), *IceCube-221223A – IceCube observation of a high-energy neutrino candidate track-like event*

[240]: Stein et al. (2020), *IceCube-201021A: One candidate counterpart from the Zwicky Transient Facility*

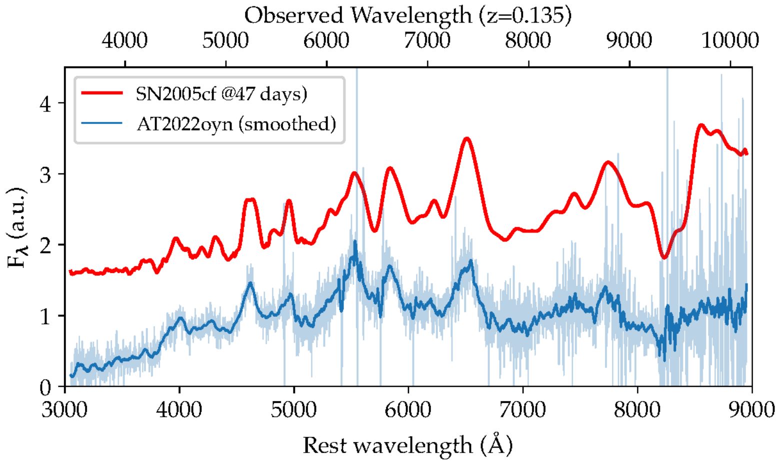

Figure 4.15: Spectrum of *SN2022oyn*, taken with Keck/LRIS. The best-fit `SNID` template (*SN2005cf*), a SN Ia, is shown for comparison in red on top. The redshift inferred from the template match is $z = 0.135$.

4.9.6 TDE *AT2019dsg*

I did not contribute to the analysis of the TDE *AT2019dsg*, but as it is important in the context of the work on *AT2019fdr* highlighted in the next Chapter (5), the most important findings will be presented in this section.

AT2019dsg was identified as a possible counterpart to *IC191001A*, an alert with a 90 % bounding rectangle of 25.5 deg^2, an estimated neutrino energy of 217 TeV and a signalness (see Section 2.4.2) of 0.59 [241, 242].

[241]: Stein (2019), *IceCube-191001A – IceCube observation of a high-energy neutrino candidate event*

[242]: Stein et al. (2019), *Candidate Counterparts to IceCube-191001A with ZTF*

AT2019dsg had already been classified as TDE and was 6 months old at the time of neutrino arrival, peaking in the ZTF *g*-band on May 3, 2019 (151 days prior to the neutrino). As a ZTF-detected TDE, it had already been observed by the Ultra-Violet/Optical Telescope (UVOT) [243] aboard NASA's *Neil Gehrels Swift Observatory* (*Swift*) [244]. These observations revealed bright UV emission, tracing the optical light curve of *AT2019dsg* (see Fig. 4.16).

[243]: Roming et al. (2005), *The Swift Ultra-Violet/Optical Telescope*

[244]: Gehrels et al. (2004), *The Swift Gamma-Ray Burst Mission*

The blackbody temperature, inferred from the optical/UV datapoints, was measured at $10^{4.6}$ K, which is a bit higher than typical for TDEs [245]. Together with the transient redshift of $z = 0.0512$, a peak luminosity of $10^{44.5}$ erg s^{-1} was calculated, which rendered *AT2019dsg* one of the most luminous TDEs discovered so far.

[245]: Stein et al. (2021), *A tidal disruption event coincident with a high-energy neutrino*

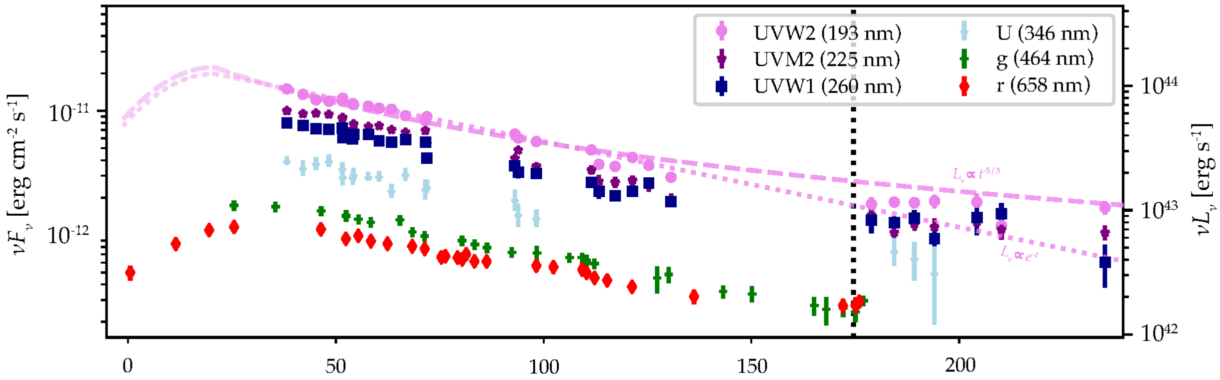

Figure 4.16: Optical (rotated red squares and green crosses) and UV (all other datapoints) light curve of *AT2019dsg*. The neutrino arrival time of *IC191001A* is marked with a black dotted line. From [245].

Furthermore, *AT2019dsg* was observed by the X-ray Telescope (XRT) [246] aboard *Swift*, starting about 40 days after discovery. Observations revealed a bright source that quickly faded to non-detection within the next 30 days [245]. Possible explanations for this are either cooling of the newly formed accretion disk, or some kind of obscuration along the line of sight. The X-ray detection of *AT2019dsg* is not uncommon: 9 of the 30 TDEs detected during Phase I of ZTF operations have been detected in X-rays [87].

[246]: Burrows et al. (2005), *The Swift X-Ray Telescope*

[87]: Hammerstein et al. (2022), *The Final Season Reimagined: 30 Tidal Disruption Events from the ZTF-I Survey*

More peculiar, *AT2019dsg* was also detected in radio wavelengths. Observations by the Large Array of the Arcminute MicroKelvin Imager (AMI) [247, 248], MeerKAT [249] and the Karl G. Jansky Very Large Array (VLA) [250] at several epochs following the neutrino detection showed a radio signal that evolved over the course of the next months.

[247]: Zwart et al. (2008), *The Arcminute Microkelvin Imager*

[248]: Hickish et al. (2018), *A digital correlator upgrade for the Arcminute MicroKelvin Imager*

[249]: Jonas (2018), *The MeerKAT Radio Telescope*

[250]: Thompson et al. (1980), *The Very Large Array*

In [245] the evolving radio signal was interpreted as the signature of a mildly relativistic outflow launched by the TDE. This interpretation was later disputed: [251] added later radio observations to the dataset and found the radio data more compatible with a non-relativistic outflow, excluding neutrino production. [252] performed very-long-baseline interferometry (VLBI) and also concluded on a non-relativistic outflow as the origin of the radio emission, while emphasizing the compatibility of that scenario with high-energy neutrino production. Nonetheless, the observations confirm long-lived non-thermal emission from the transient.

[251]: Cendes et al. (2021), *Radio Observations of an Ordinary Outflow from the Tidal Disruption Event AT2019dsg*

[252]: Mohan et al. (2022), *High-resolution VLBI Observations of and Modeling the Radio Emission from the Tidal Disruption Event AT2019dsg*

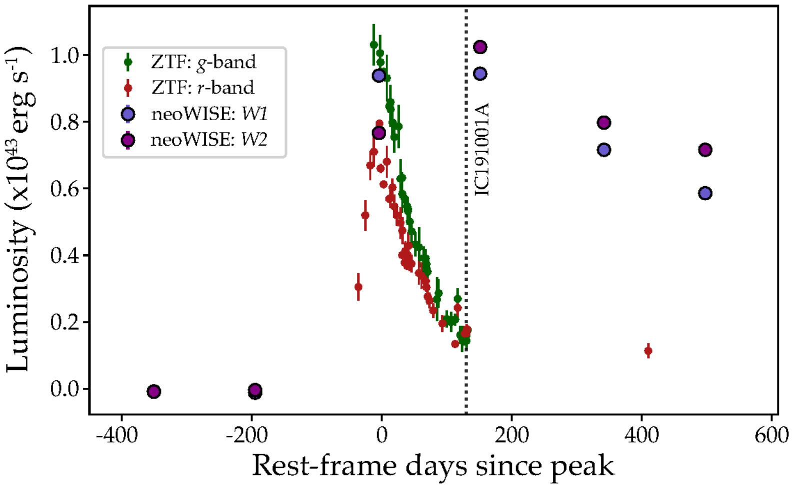

Figure 4.17: *AT2019dsg* light curve with the *WISE* infrared datapoints shown as bigger violet and magenta circles. The time of neutrino arrival is shown as black dotted line. From [253].

Lastly, a strong infrared signal was detected when evaluating NEOWISE [254] data that was not yet published at the time the source paper was finalized. The two *WISE* bands can be seen in Fig. 4.17. The infrared flare trails the optical evolution—this can be interpreted as a dust echo, where the optical to X-ray light is reprocessed by surrounding dust, re-emitting it at infrared wavelengths after a delay introduced by the light travel time. This will be discussed in detail in Section 5.4 in the context of *AT2019fdr*, the second neutrino associated TDE.

[254]: Mainzer et al. (2011), *Preliminary Results from NEOWISE: An Enhancement to the Wide-Field Infrared Survey Explorer for Solar System Science*

The chance coincidence to observe a high-energy neutrino in spatial and temporal coincidence with a radio-emitting TDE like *AT2019dsg* was computed to be 0.5 %, while the probability of finding one with a bolometric flux as high as *AT2019dsg* was 0.2 % [245]. This suggested that Tidal Disruption Events do in fact contribute to the flux of high-energies measured by IceCube.

Lightning rarely strikes twice. So the fact that we identified *AT2019fdr* as the second likely neutrino-associated TDE candidate half a year later further strengthened the case for a TDE-neutrino connection. The next Chapter 5 is dedicated to that event.

4.10 Classification Performance

The classification for the 205 candidates selected by `nuztf` can be seen in Fig. 4.18. The vast majority of transients has some kind of classification. Some of these classifications stem from spectroscopic follow-up performed by us. In some cases, promising candidates were already spectroscopically classified by the ZTF Bright Transient Survey (BTS) [255, 256], which aims to classify every transient detected by ZTF brighter than 19 mag. The rest of the classifications stem from catalog matches (.e.g MILLIQUAS or *Gaia* DR1, see Section 4.5.2).

[255]: Fremling et al. (2020), *The Zwicky Transient Facility Bright Transient Survey. I. Spectroscopic Classification and the Redshift Completeness of Local Galaxy Catalogs*

[256]: Perley et al. (2020), *The Zwicky Transient Facility Bright Transient Survey. II. A Public Statistical Sample for Exploring Supernova Demographics*

Roughly half of the transients were classified as AGN variability, i.e. stochastic AGN activity without a prominent optical flare visible around the time of neutrino detection. 12 % (24) of the candidates were classified as stars, and hence ruled out as potential neutrino emitters. We classified 9 % (18) of alerts as subtraction artifacts, as their difference images showed

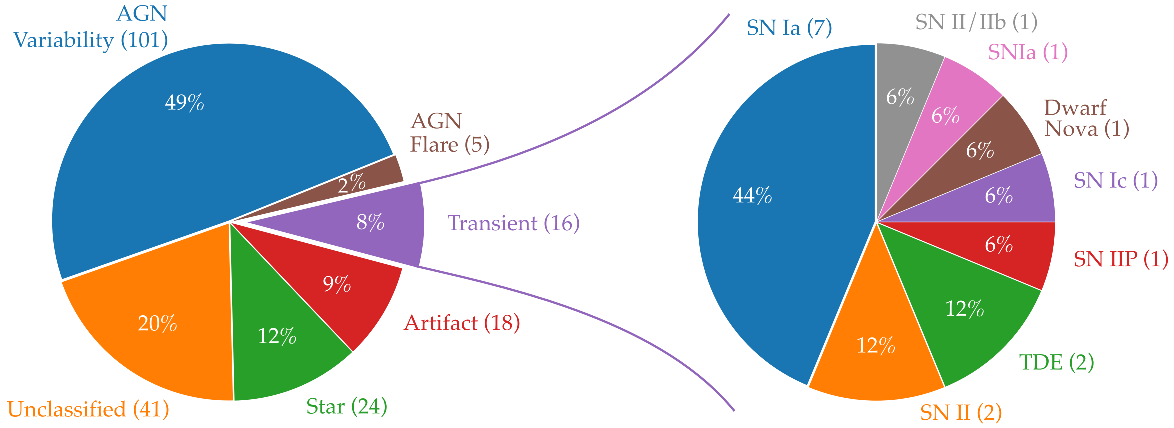

Figure 4.18: Overview of the classification performance of the neutrino follow-up program as of March 2023. The figure on the left shows all transients, while the figure on the right shows only the subclasses of the *Transient* category.

clear signs of erroneous subtractions (see Section 3.4.8). For 20 % (41) of the candidates we were not able to obtain classifications.

This leaves 5 candidates that were classified as AGN, but did show coincident optical flares, as well as 16 bona fide transients, totaling 21 classified candidates.

Most of the classified candidates were SNe Ia. This is not surprising, as the majority of transients detected by ZTF are SNe Ia, due to their intrinsic brightness: As of April 2023, 63 % of all classified transients in the BTS were SNe Ia. Among the other transients, only the two TDEs were identified as source candidates; these will be detailed subsequently. The remaining non-Ia supernovae could be ruled out as potential sources, as they either showed no sign of CSM interaction, or—in the case of the SN Ic—the explosion predated the neutrino detection significantly (see Section 4.9.1).

5 Candidate TDE *AT2019fdr*: A Possible Source?

Light from the candidate TDE *AT2019fdr* illuminating the surrounding dust in an artistic impression. Image credit: DESY and Science Communication Lab.

AT2019fdr emerged as a candidate when following up the high-energy IceCube neutrino *IC200530A*. The Gold alert (see Section 2.4.2) neutrino was detected on May 30, 2020 with a 90 % rectangular uncertainty area of 25.4 deg^2, a signalness of 0.59 and an estimated neutrino energy of 82.2 TeV [226].

Two more sources were initially published by us as candidate sources (*SN2020lls* and *SN2020lam*), but both could be ruled out with follow-up observations. For details on these events, see Section 4.9.1 and 4.9.2. This left *AT2019fdr* as the only remaining neutrino source candidate, a long-lived optical transient located at RA = 257.278575 and Dec = +26.855758, with a distance of 1.72° to the reported neutrino best-fit location.

At the moment of neutrino arrival, the transient was already 400 days old. It was originally discovered May 13, 2019 by `AMPEL` [257], with a first detection on May 3, 2019. The neutrino arrived about 10 months after the optical peak. The peak flux in the ZTF g-band was $1.3 \times 10^{-12}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$. The redshift of $z = 0.2666$ was inferred from the Balmer lines visible in a spectrum taken on July 3, 2019 with the Double Spectrograph (DBSP) on the P200 telescope. The peak g-band luminosity calculated with this redshift, $L_{\mathrm{peak}} = 2.9 \times 10^{44}\,\mathrm{erg\,s^{-1}}$, showed that *AT2019fdr* was an extraordinarily luminous event[1].

[257]: Nordin et al. (2019), *ZTF Transient Discovery Report for 2019-05-13*

The long-duration flare was initially tentatively classified as a superluminous supernova of type IIn (SLSN IIn). A peculiarity of this event is that it occurred close to the core of its host galaxy SDSSCGB 6856.2, a Narrow-Line Seyfert 1 (NLSy1) galaxy, a type of AGN (see Section 1.4.3). The mean angular separation to the host position as reported in the *Gaia* DR2 was $0.03^{+0.15}_{-0.15}$ arcsec. Due to this proximity, it was systematically followed up by a ZTF program dedicated to study flares in NLSy1 galaxies [258].

[258]: Frederick et al. (2021), *A Family Tree of Optical Transients from Narrow-line Seyfert 1 Galaxies*

In a paper describing the sample of such flares detected during ZTF Phase I, a TDE interpretation of *AT2019fdr* was stipulated [258]. The authors outright rejected a SLSN interpretation, and also argued against an AGN origin of the flare. The reasons for disfavoring the SLSN hypothesis were the long-lived U-band and UV emission, emission at the blue end of the

[1] Here and in the following sections a standard flat cosmology will be used, with $\Omega_\Lambda = 0.7$ and $H_0 = 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$.

Balmer line profiles, the longevity of the flare and its nuclear location close to the center of its host galaxy [258].

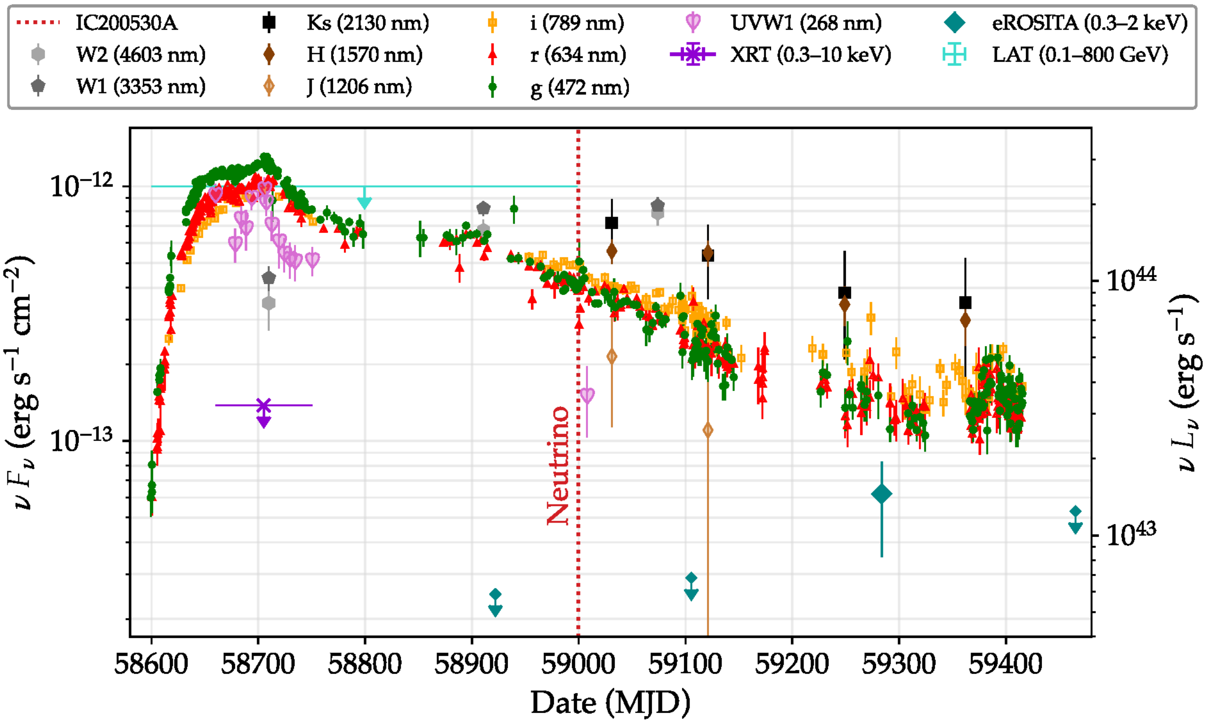

Figure 5.1: Light curve of *AT2019fdr*. The arrival time of *IC200530A* is marked with a red dotted line. It comprises host-subtracted measurements from *WISE*, P200/WIRC, ZTF and *Swift*/UVOT. Also shown are measurements without host subtraction from *SRG*/eROSITA, as well as upper limits from *Swift*/XRT and *Fermi*/LAT. The left y-axis shows νF_ν, where F_ν is the spectral flux density at frequency ν; the right y-axis shows νL_ν with L_ν being the luminosity at frequency ν. Figure by the author, from [259].

In the following, an overview will be given over the observations in different wavelengths, the reductions of the near infrared observations triggered by our group, the modeling of the spectral energy distribution (SED), and the interpretation of the delayed infrared signal as dust echo. This will be followed by a discussion of the TDE interpretation and the calculation of the chance of observing an event comparable to *AT2019fdr* in temporal and spatial coincidence with a high-energy neutrino.

Due to its multi-messenger nature and the large number of instruments involved, the study of *AT2019fdr* was a collaborative effort. Therefore, I have highlighted computing work, reductions or analyses that were done by me personally by explicitly writing 'I'. When I was only involved in the discussion and interpretation of results, I wrote 'we'.

5.1 Multi-wavelength Observations

As *AT2019fdr* was discovered close to the core of its host galaxy, it was followed up as part of regular ZTF TDE group activities from early on, including ToO observations by *Swift* UVOT and XRT.

After observing the neutrino, I triggered additional follow-up with *Swift*/UVOT and requested multi-epoch near-infrared observations with the Wide Field Infrared Camera (WIRC) [260] on the P200. Furthermore, our group triggered multi-epoch radio observations with VLA.

[260]: Wilson et al. (2003), *A Wide-Field Infrared Camera for the Palomar 200-inch Telescope*

Additionally, we analyzed data by the LAT [108] aboard the *Fermi Gamma-Ray Space Telescope* (*Fermi*), and investigated four epochs of observations by the Extended Roentgen Survey with an Imaging Telescope Array (eROSITA) telescope [261] aboard the *Spectrum Roentgen Gamma* (*SRG*) [262] space mission.

[261]: Predehl et al. (2021), *The eROSITA X-ray telescope on SRG*

[262]: Sunyaev et al. (2021), *SRG X-ray orbital observatory*

All detections and upper limits except for the radio data are shown in Fig. 5.1. As one can see, the flare was very long-lived, with a moderately fast rise (~100 d until optical peak), a short decay phase followed by a plateau, and lastly an exponential decay, apparently with a constant decay rate [259].

[259]: Reusch et al. (2022), *Candidate Tidal Disruption Event AT2019fdr Coincident with a High-Energy Neutrino*

5.1.1 Gamma-ray Limits

The details of the *Fermi*/LAT analysis were published in [263]. The one-year time window analyzed for *AT2019fdr* spanned from the first optical detection of the source on May 3, 2019, until the arrival of the neutrino on May 30, 2020.

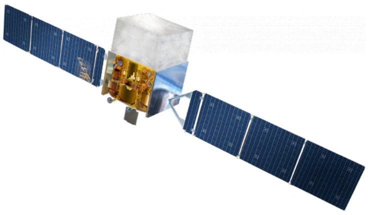

Figure 5.2: The *Fermi* satellite. Image credit: NASA.

No gamma-ray sources were detected significantly ($\geq 5\,\sigma$), neither new sources nor previously known sources from the 10 year *Fermi* LAT Fourth Source Catalog (4FGL-DR2) [264, 265]. The 95 % confidence level (CL) upper limit derived for this timeframe within the energy bin of 0.1–800 GeV was $1 \times 10^{-12}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$. This upper limit was calculated by testing a point-source hypothesis at the location of *AT2019fdr*, assuming a power-law spectrum of the form $\mathrm{d}N/\mathrm{d}E \propto E^{-\gamma}$, with a spectral index $\gamma = 2$ [263].

[263]: Velzen et al. (2021), *Establishing accretion flares from massive black holes as a major source of high-energy neutrinos*

[264]: Abdollahi et al. (2020), *Fermi Large Area Telescope Fourth Source Catalog*

[265]: Ballet et al. (2020), *Fermi Large Area Telescope Fourth Source Catalog Data Release 2*

5.1.2 X-ray Observations

The sky region of *AT2019fdr* was visited four times by *SRG* with a six-month cadence. The first visit was March 13–14, 2020. Only on the third visit one year later the source was detected, which can be counted as evidence for temporal evolution in X-ray wavelengths. All measurements and the upper limits and detections derived from these are displayed in Fig. 5.1 and listed in Table 5.1.

Figure 5.3: The *SRG* satellite. The eROSITA instrument aboard *SRG* was put into hibernation following the Russian invasion of Ukraine in February 2022. Image credit: Roskosmos.

The single detection from March 11, 2021 revealed a very soft thermal spectrum. As can be seen in Fig. 5.4, the best-fit blackbody temperature was 56^{+32}_{-26} eV, with the errors being 68 % for one parameter of interest [266]. In the source rest frame, this measurement corresponds to a temperature of 71^{+41}_{-33} eV. This renders the X-ray spectrum of *AT2019fdr* one of the softest spectra of the TDEs discovered by *SRG*/eROSITA so far.

[266]: Sazonov et al. (2021), *First tidal disruption events discovered by SRG/eROSITA: X-ray/optical properties and X-ray luminosity function at z<0.6*

Table 5.1: *SRG*/eROSITA upper limits and detection of *AT2019fdr* in the 0.3–2.0 keV-band. From [259].

MJD	Date	Upper limit ($\mathrm{erg\,s^{-1}\,cm^{-2}}$)	Energy flux ($\mathrm{erg\,s^{-1}\,cm^{-2}}$)
58922	2020–03–14	2.5×10^{-14}	—
59105	2020–09–13	2.9×10^{-14}	—
59284	2021–03–11	—	$6.2^{+2.7}_{-2.1} \times 10^{-14}$
59465	2021–09–08	5.3×10^{-14}	—

The equivalent neutral hydrogen column density had a best-fit value of NH = $1.47^{+2.80}_{-1.25} \times 10^{21}\,\mathrm{cm}^{-2}$. This value is a measure of the amount of neutral hydrogen atoms along the line of sight to the source. Within the errors it was consistent with the galactic value of $\mathrm{NH}_{\mathrm{Gal}} = 0.40 \times 10^{21}\,\mathrm{cm}^{-2}$. There was some degeneracy between the neutral hydrogen column density and the inferred blackbody temperature; this is usual for soft sources. However, with T_{bb} = 131 eV at the 95 % confidence level, the temperature upper bound was still fairly low [259].

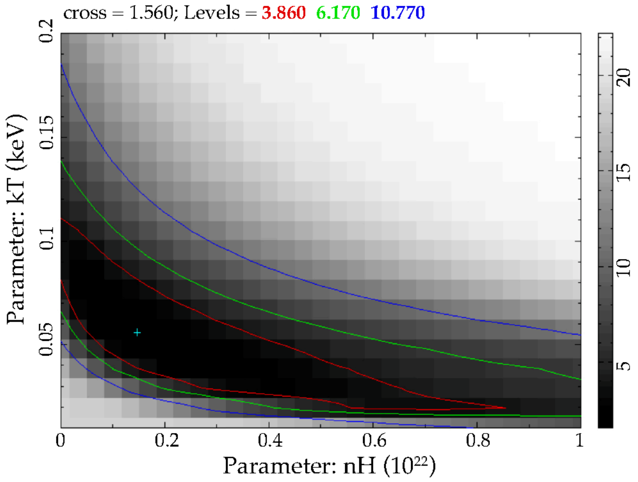

Figure 5.4: Temperature fit for the *SRG*/eROSITA detection. Figure by M. Gilfanov, slight modifications by the author.

As noted above, *Swift* XRT had already observed *AT2019fdr* 14 times prior to the neutrino arrival [258]. We requested a prompt ToO observation after the source emerged as a neutrino candidate, which was carried out on June 7, 2020 with a 2000 s exposure.

The data was reduced with the publicly available `Swift XRT data products generator` [267] in the energy range of 0.3–10 keV. As all 14 observations prior to the neutrino arrival were non-detections, those were binned (total exposure: 20,700 s) to compute a $3\,\sigma$ energy flux upper limit of $1.4 \times 10^{-13}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$. The upper flux limit for the post-neutrino observation was $4.7 \times 10^{-12}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$. To convert photon counts to energy flux, I employed the `HEASARC WebPIMMS` tool [268]. I set the blackbody temperature to the 56 eV measured by *SRG*/eROSITA, and used the best-fit neutral hydrogen column density from the same measurement [259].

[267]: UK *Swift* Science Data Centre (2020), *Swift XRT Data Products Generator*

[268]: Arida et al. (2020), *HEASARC WebPIMMS*

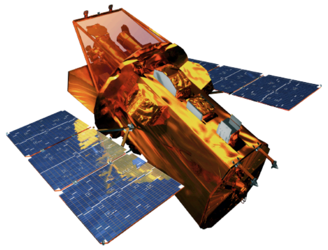

Figure 5.5: The *Swift* satellite. Image credit: NASA.

5.1.3 Optical/UV Observations

I processed the ZTF optical observations with `fpbot` (see Section 4.6). The science-ready *Swift* data was retrieved from the *Swift* archive[2]. All exposures of individual epochs were co-added, and filtered to boost the signal-to-noise ratio with `uvotsim`, contained in `HEAsoft`[3]. The brightness of the transient was measured using `uvotsource` of the same package. The aperture used was 3 arcsec, with a far bigger radius used for the background level extraction. We calibrated the photometry with files made available by September 2020 and employed the methods of [269] to convert the measured magnitudes to the AB system (see Section 3.4.7). I corrected the extracted and converted magnitudes by subtracting the synthetic host model described in Section 5.2.3 [259].

[2] `https://www.swift.ac.uk/swift_portal`

[3] `https://heasarc.gsfc.nasa.gov/docs/software/heasoft`

[269]: Breeveld et al. (2011), *An Updated Ultraviolet Calibration for the Swift/UVOT*

As the extraction of the infrared measurements was more involved due to the presence of host galaxy blending, this will be detailed in Section 5.2 below.

5.1.4 Radio Observations

After the neutrino detection, we applied for Director's Discretionary Time (DDT) on the VLA. Three individual measurements were taken on July 3, September 13 and November 7, 2020. The first epoch was carried out in the 2–4 GHz and 8–12 GHz bands. In the subsequent epochs, we added the 1–2 GHz and 4–8 GHz bands. Delay, bandpass and flux calibration were performed on the source *3C 286*, and the nearby source *J1716+2152* was used for the complex gain calibration. For this, the VLA calibration pipeline in the Common Astrometry Software Application (CASA) [270] was used, and imaging was performed with the CASA task `tclean`. Finally, the target flux density was measured by fitting a point source in the image plane [259].

[270]: Bean et al. (2022), *CASA, the Common Astronomy Software Applications for Radio Astronomy*

Figure 5.6: VLA in New Mexico. Image credit: NSF.

In all three epochs, we did measure flux from the source location. Furthermore, the flux density declined with higher frequencies in all epochs. While epoch 1 and 2 showed no evolution in the 4–8 GHz and 8–12 GHz bands, the final epoch was suggestive of reduced flux and subsequently a spectral steepening in these bands. Intensive testing revealed that the reduced flux was not intrinsic to the source, but rather stemmed from significant atmospheric phase changes between the calibrator scans. These reduced the measured flux densities during epoch 3 due to decorrelation. As the source flux density was too low for self-calibration, the effect could not be corrected. We concluded that no evidence for temporal evolution in the radio can be found during the 5 months of observation. All measurements are shown in Fig. 5.7 and listed in Table A.2.

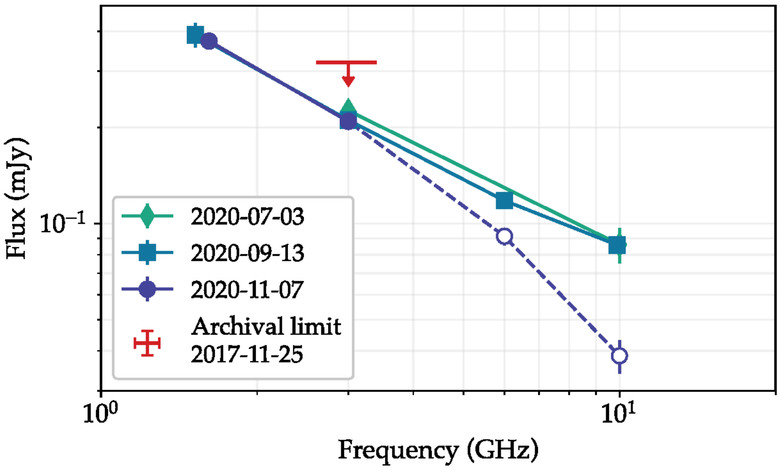

Figure 5.7: Radio observations of *AT2019fdr*. An archival upper limit is shown in red. The apparent spectral steepening in the third epoch (dashed blue lines) are a calibration artifact from decorrelation due to atmospheric phase changes, and are not intrinsic to the source. Figure by the author, from [259].

To compare to pre-flare host radio emission, we also obtained an archival upper limit from the Very Large Array Sky Survey (VLASS) [271]. We exclusively used VLASS, as it was the only survey with sufficient sensitivity and angular resolution to probe *AT2019fdr*. The quicklook continuum `fits` file for tile T17t23We was obtained from the archive[4]. The observation was carried out on November 25, 2017 in the 2–4 GHz band. At

[271]: Lacy et al. (2020), *The Karl G. Jansky Very Large Array Sky Survey (VLASS). Science Case and Survey Design*

[4]: `https://archive-new.nrao.edu/vlass/quicklook`

a significance level of 3 σ, no emission from the source was detected. The local RMS noise was $\sigma = 0.11$ mJy/beam, and the beam size was 2.46 arcsec × 2.28 arcsec at a position angle of −37°. This resulted in an upper limit of 0.32 mJy. A second VLASS observation was taken a few days prior to our second trigger epoch, but the resulting 3 σ upper limit was less constraining (0.4 mJy) [259].

5.2 Near-infrared Observations

After the detection of the neutrino, I triggered near-infrared (NIR) observations to extend the SED to longer wavelengths. We observed during four epochs with WIRC on the Palomar P200 in the *J*-, *H* and *Ks*-band, centered on 1.2 µm, 1.6 µm and 2.1 µm. The dates of observations were July 1 and September 27, 2020, as well as February 2 and May 28, 2021.

All WIRC measurements were reduced using a custom pipeline (see [272]), which performs flat fielding, background subtraction and fits an astrometric solution using *Gaia* DR2. Subsequently, the individual images for each filter and epoch were stacked to increase the signal-to-noise ratio. The photometric zero-point calibration of the stacked images was achieved within the pipeline by comparison to the Two Micron All Sky Survey (2MASS) [273].

[272]: De et al. (2020), *Palomar Gattini-IR: Survey Overview, Data Processing System, On-sky Performance and First Results*

[273]: Skrutskie et al. (2006), *The Two Micron All Sky Survey (2MASS)*

5.2.1 `GALFIT` flux extractions

As another galaxy lies close to *AT2019fdr*'s host galaxy (angular separation: 7 arcsec), I performed a manual flux extraction with `GALFIT` [274]. This two-dimensional fitting algorithm allows defining different galactic types and components, as well as point sources, and simultaneously fit the science images with these source types.

[274]: Peng et al. (2002), *Detailed Structural Decomposition of Galaxy Images*

To improve photometric accuracy, I first derived the point spread function (see Section 3.4.6) of individual stacked images with `photutils` [275], a package for the Python astronomy toolkit `astropy` [276]. All stars within the surrounding of *AT2019fdr* were selected for this. From this selection, I visually selected a subset of stars neither too dim nor too bright (see Fig. 5.8 for an example). Using these stars, I calculated the PSF of the image.

[275]: Bradley et al. (2020), *astropy/photutils: 1.0.0*

[276]: Price-Whelan et al. (2022), *The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package**

I retrieved the stars with `DAOStarFinder`, while I used `EPSFBuilder` for fitting the PSF. I verified the quality of the extracted PSF by inspecting the residuals of 4 nearby reference stars from SDSS (see Fig 5.9).

I then ingested the resulting PSF into `GALFIT`. Two Sérsic profiles describing the galaxy flux as a function of the distance to the galaxy core were then defined[5], centered on the host and the neighboring galaxy location. Additionally, I used a point source with variable position to account for the transient flux. These three sources were then fit together, resulting in flux values for the host galaxy, the neighboring galaxy and the transient.

[5] For details, see e.g. [277].

I fixed all parameters except for the point source flux after fitting one epoch (reference epoch). The point source flux was then fit in the remaining three epochs. Unfortunately, the choice of reference epoch had an impact on

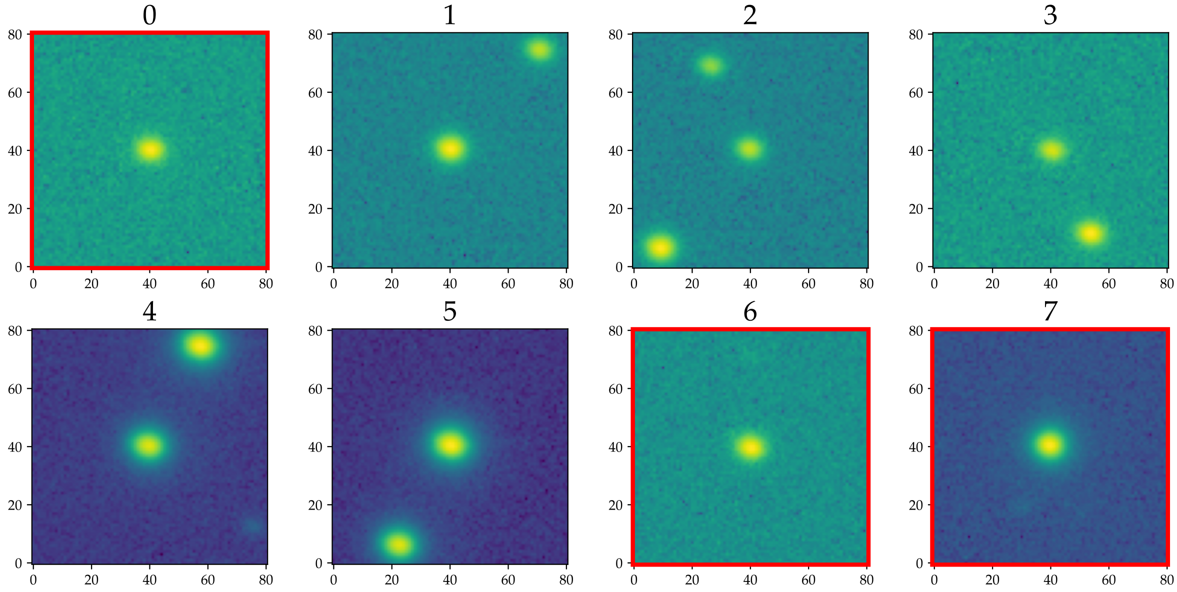

Figure 5.8: Star selection to fit the PSF of the stacked P200/WIRC *H*-band image from Februar 2, 2021. Each cutout comprises 80 × 80 pixels. As can be seen, some stars have neighboring stars. As this can hamper a clean extraction of the PSF, such stars were rejected. All stars included in the PSF extraction are marked with a red rectangle (star 0, 6 and 7). One can also see that during this epoch observing conditions were not ideal, resulting in horizontally elongated images of the stars.

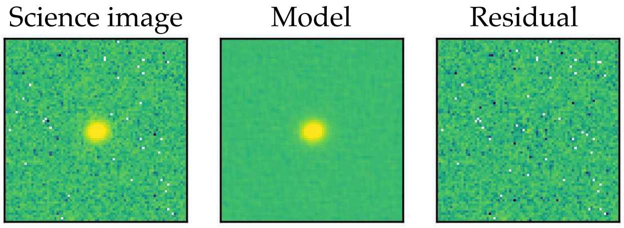

Figure 5.9: P200 *H*-band reference star. The `GALFIT` point source model constructed from the image PSF (center) was subtracted from the science image (left), with the residual shown on the right.

the point source flux values. This can be explained by differing observing conditions at different epochs. To account for this variation, both epoch 1 and epoch 4 were used as reference epochs. I used the difference in the resulting transient flux (when using epoch 1 as reference epoch on the one hand, and when using epoch 4 as such on the other hand) as systematic uncertainty on the extracted transient flux [259].

As there was no image without transient flux available, I added the point source flux to the host galaxy flux, and then subtracted the synthetic host model (see Section 5.2.3). The model-subtracted AB magnitudes derived by this procedure are shown in Table 5.2.

5.2.2 *WISE* Detections

In addition to the ground-based NIR observations with P200, images from the mid-infrared (MIR) *WISE* satellite were obtained in the *W1*- and *W2*-band (3.4 µm and 4.6 µm). *WISE* has cadence of 6 months, and the available images comprised 13 observational epochs prior to the detection of *AT2019fdr*, as well as 3 epochs during the transient's life.

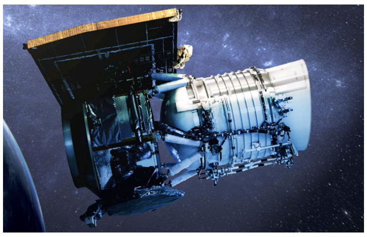

Figure 5.10: The *WISE* satellite. Image credit: NASA.

As the *WISE* MIR photometry—like the NIR photometry detailed above—suffered from blending with the nearby galaxy, forced PSF photometry was used on the co-added images (done with `ICORE` [278]) of each observational epoch. The package employed for this was `PythonPhot` [279], a Python adaptation of the `DaoPhot` [189] package.

[278]: Masci (2013), *ICORE: Image Co-addition with Optional Resolution Enhancement*

[279]: Jones et al. (2015), *PythonPhot: Simple DAOPHOT-type photometry in Python*

The root mean square (RMS) of the pre-flare datapoints was only 18 µJy, which was significantly smaller than the peak difference flux, measured at 0.9 mJy. A robust baseline can therefore be constructed from the pre-flare *WISE* detections.

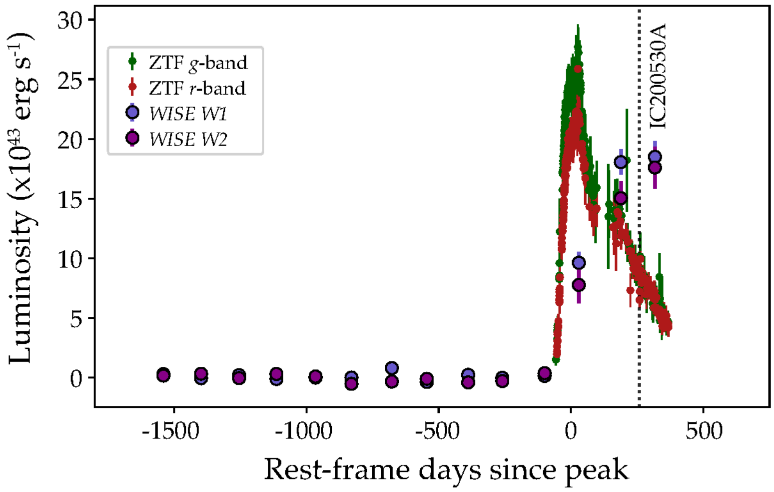

Figure 5.11: *NEOWISE* + ZTF forced photometry light curve at the location of *AT2019fdr*. The arrival time of neutrino *IC200530A* is marked with a dotted black line. Note the small scatter of the pre-flare datapoints, allowing the construction of a robust baseline. From [263]

5.2.3 Synthetic Host Model

As not all measurements from the various instruments had the luxury of readily available reference images for easy extraction of transient light (as is e.g. the case for ZTF), we constructed a synthetic model of the host galaxy SED, which can be seen in Fig. 5.12. The best-fit values were a galaxy mass of $\log(M/M_\odot) = 10.5$, a metallicity of $\log(Z/Z_\odot) = 0.1$, dust extinction of $E(B-V) = 1$, a star formation start time of $t = 1.6$ Gyr after the Big Bang, and a star formation rate e-folding time of $\tau = 0.8$ Gyr.

The fit procedure is described in [280]; we used archival measurements from different sources: To measure the UV flux, we used images from the Galaxy Evolution Explorer (*GALEX*) [281]. The `gPhoton` [282] package with an aperture of 4 arcsec was used for flux extraction.

[280]: Velzen et al. (2020), *Optical-Ultraviolet Tidal Disruption Events*

[281]: Martin et al. (2005), *The Galaxy Evolution Explorer: A Space Ultraviolet Survey Mission*

[282]: Million et al. (2016), *gPhoton: The GALEX Photon Data Archive*

To measure the optical flux, we employed the model magnitudes from SDSS. These are computed either from a de Vaucouleurs galaxy profile, or an exponential galaxy profile, with `mag_model` using the better of the two fits (see `https://www.sdss4.org/dr12/algorithms/magnitudes/mag_model` for details).

Lastly, we included the baseline *WISE* datapoints detailed above, as well as an archival infrared measurement from the UKIRT Infrared Deep Sky Survey (UKIDSS) [283].

[283]: Lawrence et al. (2007), *The UKIRT Infrared Deep Sky Survey (UKIDSS)*

MJD	**Date**	***J*-band** (1.2 µm)	***H*-band** (1.6 µm)	***Ks*-band** (2.1 µm)
59031	2020–07–01	19.06 ± 0.50	17.73 ± 0.12	17.13 ± 0.26
59121	2020–09–29	19.78 ± 0.98	17.76 ± 0.12	17.45 ± 0.35
59249	2021–02–04	—	18.26 ± 0.19	17.81 ± 0.48
59362	2021–05–28	—	18.42 ± 0.22	17.91 ± 0.53

Table 5.2: NIR AB magnitudes after subtracting the synthetic host model. Only systematic uncertainties are provided, as the photometric uncertainties should be negligible in comparison, at least in the *J*- and *H*-band. The flux of the third and fourth *J*-band epochs after host-model subtraction was negative, which was counted as non-detection. From [259]

After gathering archival flux measurements from these sources, the `prospector` toolkit [284] was used to sample synthetic galaxy models built by Flexible Stellar Population Synthesis (FSPS) [285]. This package is written in FORTRAN and was employed via the Python wrapper `python-fsps` [286].

[284]: Johnson et al. (2021), *Stellar Population Inference with Prospector*

[285]: Conroy et al. (2010), *The Propagation of Uncertainties in Stellar Population Synthesis Modeling. III. Model Calibration, Comparison and Evaluation*

[286]: Foreman-Mackey et al. (2014), *python-fsps: Python bindings to FSPS (v0.1.1)*

The archival values used in constructing the synthetic host model can be found in the Appendix in Table A.3.

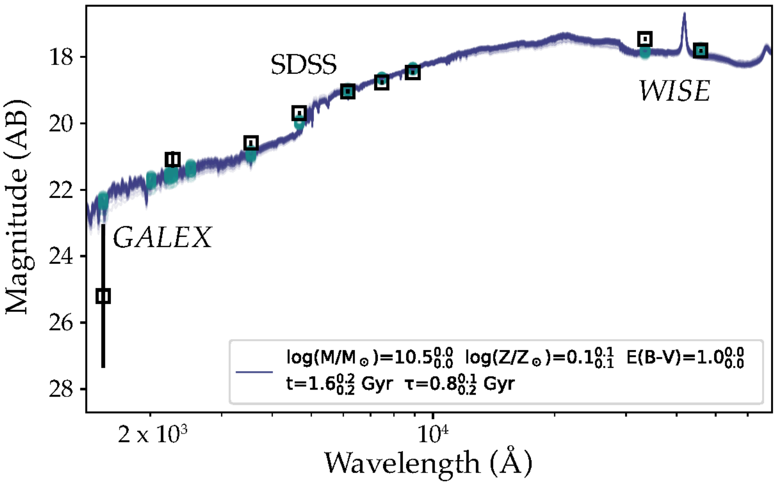

Figure 5.12: Synthetic host spectrum constructed with FSPS. Figure by S. van Velzen, annotations by the author.

I subsequently subtracted the host model from all measurements possibly containing host light within their respective bandpass using `SNCosmo` [287], a Python framework dedicated to supernova analysis.

[287]: Barbary et al. (2022), *SNCosmo*

5.3 SED Blackbody Modeling

After compiling and reducing the individual light curve measurements, a physical model motivating the detected light needed to be applied.

To explain the SED, I employed different models: A power law, a single blackbody, a broken power law[6] and lastly a double blackbody model. In the following, the derivation of this model and the energy output of the system described by it will be detailed, including a discussion of dust extinction.

5.3.1 Fit Epochs

For the fit I isolated three regions of interest. These comprised (1) the 20 days between August 5 and 25, 2019, as this windows covered the

---

[6] I.e. a power law with a break at a certain wavelength where the spectral index changes.

Table 5.3: Priors for the double blackbody minimization. A_V is the extinction parameter (only fitted in epoch 1 and fixed for epoch 2 and 3), T_1 and T_2 are the temperatures of the two blackbodies, while S_1 and S_2 are the 'Scale' parameters (which can be converted into radii.)

Parameter	Initial value	Bounds
A_V	0.45	$[0, 2]$
T_1	1730 K	[7000 K, 50,000 K]
T_2	1650 K	[500 K, 2300 K]
S_1	2.3×10^{23}	$[1 \times 10^{20}, 1 \times 10^{27}]$
S_2	1×10^{20}	$[1 \times 10^{18}, 1 \times 10^{21}]$

optical peak of the light curve, as well as the first *WISE* measurement. The second, longer epoch (2) ranged from June 6 to October 8, 2020 (124 days), shortly after the neutrino detection. It covered one *WISE* measurement epoch and two P200 NIR observations. The last epoch (3) covered January 1 to February 26, 2021 (51 days) and contained another P200 NIR observation.

As the light curve did not significantly change within those epochs, I used the mean flux values for each observed bands as input for the fitting procedure.

5.3.2 Fit Models

As stated above, I tried a variety of models of increasing complexity. The least complex model that captured the measurements reasonably well was a double blackbody model. These models were, with increasing number of free parameters:

Unbroken power law A single, unbroken power law with the power law index as free parameter.
Single blackbody This model is described by two parameters, the blackbody radius and its temperature.
Broken power law The two spectral indices and the wavelength of the break are the free parameters here.
Double blackbody Two unmodified blackbodies, resulting in four free parameters (2× temperature, 2× radius).

All fits were performed with the `lmfit` [288] Python package, which implements the Levenberg-Marquardt algorithm for solving non-linear least-squares problems.

[288]: Newville et al. (2021), *lmfit/lmfit-py 1.0.2*

To account for host extinction, I used the `extinction` [289] Python package, namely the Calzetti attenuation law [290] contained therein. I fixed the extinction R_V at 3.1, the value typical for the Milky Way [291] and left A_V free in epoch 1. The best-fit value of $A_V = 0.45^{+0.14}_{-0.14}$ was fixed in epoch 2 and 3. I decided on using epoch 1 to determine A_V as it provided the best optical and UV coverage of the three epochs. For the double blackbody model, this procedure ultimately resulted in 5 (4) free parameters in epoch 1 (2 and 3).

[289]: Barbary (2016), *Extinction V0.3.0*

[290]: Calzetti et al. (2000), *The Dust Content and Opacity of Actively Star-forming Galaxies*

[291]: Schultz et al. (1975), *Interstellar reddening and IR-excesses of O and B stars*

I imported the measurements, averaged the flux in each band per epoch, and passed them to the minimizer (the fit priors are shown in Table 5.3).

The model for the blackbodies was provided by `astropy`[7]. As this model works with a *scale* parameter[8] instead of a blackbody radius, this needed to be converted via

[7] `astropy.modeling.physical_models.BlackBody`

[8] This `astropy` terminology is slightly confusing, as the scale has units of spectral radiance ($\mathrm{erg\,s^{-1}\,cm^{-2}\,Hz^{-1}\,sr^{-1}}$).

$$\text{Scale} = d_L^2/(r^2\pi), \tag{5.1}$$

where r is the blackbody radius and d_L is the luminosity distance, calculated with the redshift $z = 0.2666$ and assuming a standard cosmology. The luminosity distance is a measure for how far away something 'looks'. If the luminosity L is the total amount of energy isotropically radiated by an object per unit time, it is related to the flux F measured at luminosity distance d_L via this relation, describing an expanding sphere of light:

$$d_L = \sqrt{L/4\pi F} \tag{5.2}$$

The luminosity distance is equal to the proper distance (as in the amount of space a photon had to traverse to reach us) if and only if the universe is geometrically flat and the universe is static, so it is neither shrinking nor expanding. As the universe in fact *is* expanding, the luminosity distance is similar to the proper distance only for small redshifts (i.e. if there was little time needed for the photon to reach us, and therefore little expansion). In general, the luminosity distance depends on the evolution of the universe. Here, I calculated the luminosity distance with the `astropy.cosmology` module.

5.3.3 Minimization

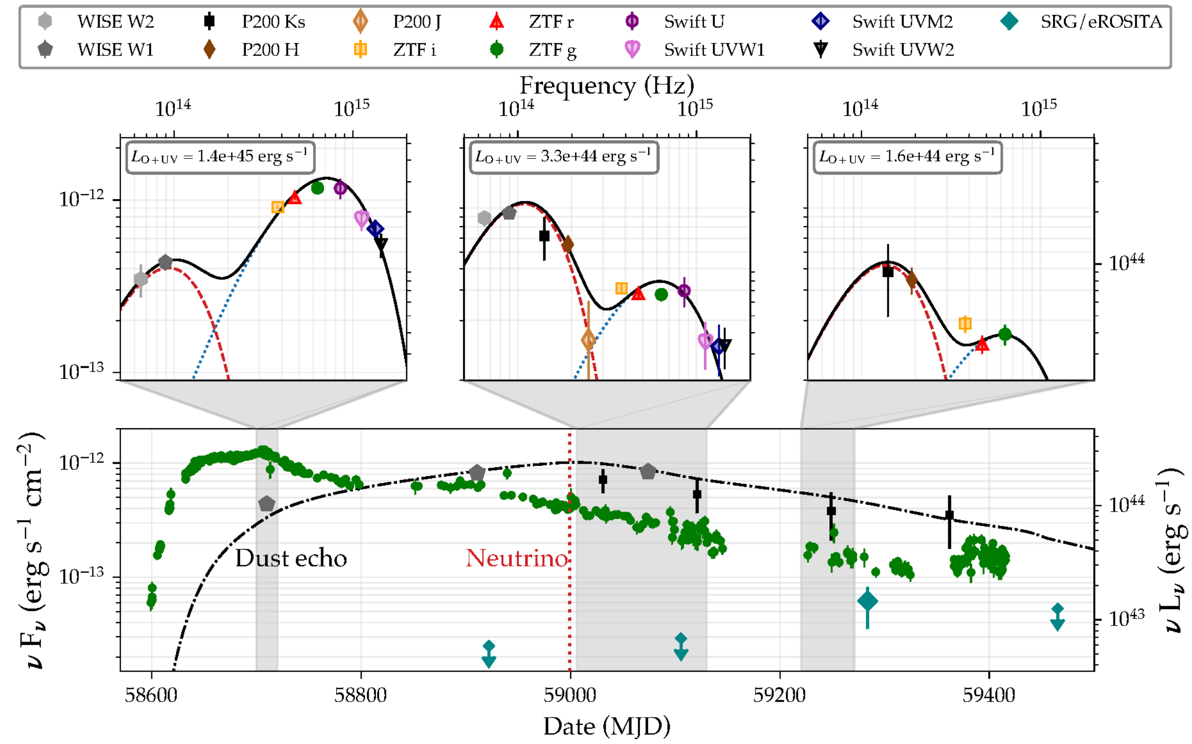

Figure 5.13: Light curve and SED of *AT2019fdr*. The three panels on the top show the double blackbody fits for the three epochs marked in light gray. On the bottom, the ZTF *g*- and *WISE W1*-band light curve is shown (green circles and gray pentagons), alongside the P200 *Ks* band data (black squares) as well as the three upper limits and one X-ray detection by *SRG*/eROSITA (rotated cyan squares), and the modeled dust echo emission (black dash-dotted line). The left y-axes show νF_ν (where F_ν is the spectral flux density at frequency ν), and the right y-axes display νL_ν (with L_ν being the luminosity at frequency ν). Figure by the author, from [259].

Table 5.4: Blackbody best-fit values for three epochs (1–3). *O+UV* denotes the 'blue' blackbody in the optical/UV and *IR* denotes the 'red' infrared blackbody. The luminosity is given dereddened and in the source frame. The uncertainties are at the 68% confidence level. Note that the O+UV temperature and radius (and therefore the luminosity) in the third epoch are not well constrained, as no late-time UV measurements were available. The same holds true for the IR blackbody in the first and the last epoch, as only two data points were available.

Epoch	**Band**	**Temp.** (K)	**Radius** (cm)	**Luminosity** ($\mathrm{erg\,s^{-1}}$)
1	O+UV	13526^{+569}_{-574}	$7.8^{+0.4}_{-0.4} \times 10^{15}$	$1.4^{+0.1}_{-0.1} \times 10^{45}$
	IR	1505^{+421}_{-313}	$2.2^{+1.6}_{-1.4} \times 10^{17}$	$1.7^{+2.2}_{-1.0} \times 10^{44}$
2	O+UV	11731^{+663}_{-683}	$4.9^{+0.4}_{-0.4} \times 10^{15}$	$3.3^{+0.3}_{-0.4} \times 10^{44}$
	IR	1762^{+121}_{-124}	$2.5^{+0.2}_{-0.2} \times 10^{17}$	$4.3^{+0.5}_{-0.7} \times 10^{44}$
3	O+UV	10230^{+2373}_{-1645}	$4.3^{+3.3}_{-1.0} \times 10^{15}$	$1.5^{+1.2}_{-0.4} \times 10^{44}$
	IR	2237^{+402}_{-462}	$1.0^{+0.6}_{-0.4} \times 10^{17}$	$1.9^{+1.4}_{-0.5} \times 10^{44}$

In each iteration of the minimization process, the spectrum of both blackbodies constructed as such were then reddened, and redshifted. After this, the flux of the two blackbody spectra was added up.

To extract the flux as observed in each band, I used `SNCosmo` to evaluate all bands given their transmission function. For this I used the integrated bandpass library of `SNCosmo`. Bandpasses that were not shipped with the package were added manually from the SVO Filter Profile Service [292]. This was the case for *Swift*, P200 and *WISE*.

[292]: Rodrigo et al. (2020), *The SVO Filter Profile Service*

The best fit values from all three epochs, as well as their uncertainties at the 68 % confidence level can be seen in Table 5.4. The best-fit temperatures resulted in a 'blue' blackbody peaking in optical/UV wavelengths, and a 'red' blackbody peaking in the infrared.

I derived an optical/UV peak luminosity of $L = 1.4^{+0.1}_{-0.1} \times 10^{45}\,\mathrm{erg\,s^{-1}}$. By integrating this component over time by scaling it with the shape of the ZTF g-band light curve, I obtained a total bolometric luminosity of $E_{\mathrm{bol}} = 3.4 \times 10^{52}$ erg. I did not add the infrared blackbody, as dust absorption was already accounted for by fitting for dust extinction. This pushed *AT2019fdr* to the class of the brightest transients ever detected. Its inferred bolometric energy was almost twice as high as that of ASASSN-15lh, which was one of most luminous transients ever reported [293] and which was suggested to be a TDE [294].

[293]: Dong et al. (2016), *ASASSN-15lh: A highly super-luminous supernova*

[294]: Leloudas et al. (2016), *The superluminous transient ASASSN-15lh as a tidal disruption event from a Kerr black hole*

5.3.4 Uncertainty Estimate

As the fit routine was not able to generate stable covariance matrices in all epochs, I estimated the uncertainties manually.

For each epoch and fit parameter, this was done by letting all parameters vary except for the parameter of interest. I then varied this parameter around the best-fit value, and the resulting χ^2 was evaluated. Following this, I approximated the χ^2 distribution with a polynomial, and calculated the 68 % confidence level—the two parameter values with $\chi^2 - \chi^2_{\mathrm{min}} = 1$ confine the 68 % confidence level.

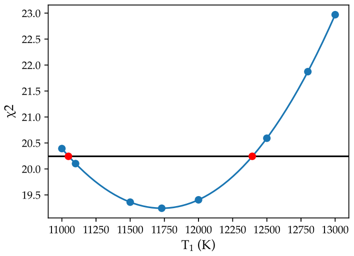

Figure 5.14: Uncertainty estimation for the double blackbody fits. Here, the uncertainty of the temperature T_1 of the hot (blue) blackbody during epoch 2 is estimated. The red datapoints mark the 68 % confidence level.

From the fit parameter uncertainties calculated as such, I obtained uncertainties on the radius and the luminosities via Gaussian error propagation.

As can be seen in Table 5.4, the third epoch saw a temperature rise in the infrared blackbody. As these values were in tension with the dust-echo model discussed in the next Section 5.4, I investigated this behavior.

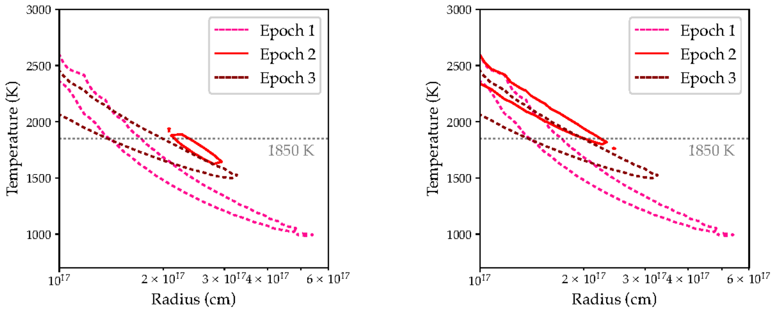

(a) Infrared blackbody, 90 % CL contours for all three epochs (optical+UV blackbody fit parameters remain free).

(b) Infrared blackbody, 90 % CL contours for all three epochs after removing the two *WISE* datapoints for epoch 2.

Figure 5.15: 90 % CL contours for the infrared blackbody temperature fits, once including the *WISE* data in all epochs (plot on the left), and once after removing them in epoch 2 (plot on the right). The contours of epoch 2 (solid red line) move to a region of higher temperatures comparable to epoch 3 when the *WISE* datapoints are removed. This shows that they are pushing the temperature down.

The most likely reason for the higher temperature in epoch 3 is that the infrared blackbody was less constrained in that epoch, as there were no *WISE* measurements available, and the P200 *J*-band extraction resulted in a non-detection. In Fig. 5.15a, one can see that all three epochs were compatible within their 90 % confidence level with a blackbody temperature of 1850 K.

Furthermore, all epochs showed the expected linear correlation between increasing blackbody temperature and decreasing radius of the blackbody. Fig. 5.15 shows how the 90 % contour of epoch 2 gets moved to a higher temperature and lower radius when removing the two *WISE* datapoints. This reinforces the case that the higher best-fit temperature in epoch 3 is merely an artifact of having less constraints for the fit and not some unexplained physical process.

5.4 Dust Echo Model

As one can see in the light curve of *AT2019fdr* (see Fig. 5.1), the emission in all infrared bands (*WISE* and P200) was delayed with respect to the optical signal. Furthermore, the infrared emission is well approximated by a blackbody, as was shown in the last section.

As the infrared blackbody with a delayed increase in brightness with respect to the optical light curve is well explained by such a model, I interpreted this infrared emission as a dust echo. In this model, the X-ray to optical light from the flare is reprocessed by dust. This dust is usually located in a torus around the supermassive black hole (SMBH) at the center of the host galaxy [295]. The high luminosity of the flare causes the dust in the vicinity of the black hole to evaporate, up to a safe distance, at which the incident flux heats the dust, but not above its sublimation temperature at around 1850 K.

[295]: Velzen et al. (2016), *Discovery of Transient Infrared Emission from Dust Heated by Stellar Tidal Disruption Flares*

Outside the sublimation radius, the grain temperature is coupled to the incident radiation, as it cools rapidly. The cooling time is $10^{-4} a_{0.1}^{-1} T_{1850}^{-5}$ s, which equals to 0.1 ms for a grain size of 0.1 µm at 1850 K [295].

According to [295], the infrared light curve produced this way can be approximated as follows:

$$L_{\mathrm{IR}} = \int d\tau\ \Psi(\tau)\ L(t-\tau), \tag{5.3}$$

where $L(t)$ is the light curve of an isotropic flare, $\Psi(\tau)$ is the response function and τ is the delay. In the case of a spherical dust shell, $\Psi(\tau)$ can be approximated by a square wave response function, ranging from $\tau = 0$ to $\tau = 2R_0c$ (where R_0 is the radius of the dust shell). One can therefore write

$$L_{\mathrm{IR}} = \int d\tau L(t-\tau)\ \Pi(\tau), \tag{5.4}$$

where Π is the rectangular function

$$\Pi(\tau) = \begin{cases} 0 & \text{if } |\tau| > \frac{1}{2} \\ \frac{1}{2} & \text{if } |\tau| = \frac{1}{2} \\ 1 & \text{if } |\tau| < \frac{1}{2}. \end{cases} \tag{5.5}$$

The delay induced by light travel time can thus be approximated by convolving the optical light curve with a rectangular function of width $2\times \Delta_{t_c}$, where t_c is the light travel time from the transient to the surrounding dust sphere.

This can be understood if one looks at the geometry of the system. After a delay of Δ_{t_c}, light from the flare reaches the inner edge of the sphere that survives the incident radiation, directly along our line of sight. Here it gets re-emitted in the infrared and reaches us with that initial delay of cR_0. Re-emitted light from the side of the system takes even longer. Lastly, re-emitted light from the back of the system reaches us with a delay of $3 \times cR_0$—first the flare light needs to travel one radius away from us, and then the re-emitted light needs to cross the whole system (twice R_0) towards us.

To calculate R_0, I used the ZTF g-band light curve, as it had the best sampling. I then convolved this light curve with the box function, where the width of the box function in days was left as free parameter. The initial delay with respect to the optical light curve was then half the size of the box function. The second free parameter was a dimensionless amplitude used to scale the optical light curve to match the amplitude of the *WISE* datapoints.

Following this procedure, I obtained a best-fit light travel time of 193 days, which translates to $R_0 = 0.16\,\mathrm{pc}$. The best fit parameters can be seen in Fig. 5.16 and result in the black dash-dotted curve in Fig. 5.13.

To crosscheck this result, I used the fact that according to [295], the radius of this dust shell can be approximated as

$$R = \left(\frac{L_{45}}{a_{0.1}^2 T_{1850}^{5.8}}\right)^{1/2} \mathrm{pc}, \tag{5.6}$$

where L_{45} is the flare's luminosity integrated over all wavelengths where the dust absorbs, normalized to $1 \times 10^{45}\,\mathrm{erg\,s^{-1}}$, $a_{0.1}$ is the dust grain size normalized to $0.1\,\mu\mathrm{m}$, and T_{1850} is the grain temperature normalized to $1850\,\mathrm{K}$.

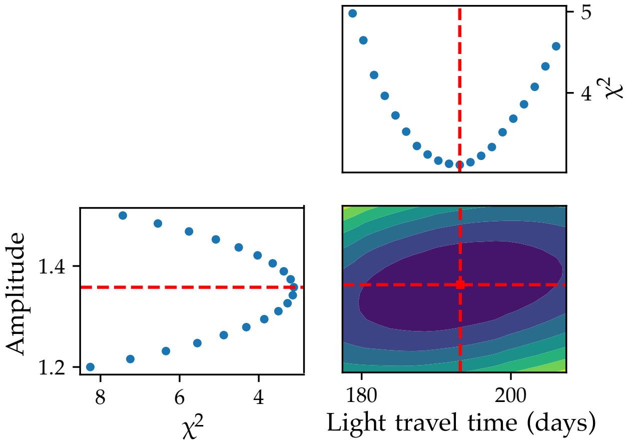

Figure 5.16: Corner plot of the dust model fit for *AT2019fdr*. The two free parameters (dimensionless amplitude and light travel time in days) are plotted together with the marginalized χ^2 values.

When inserting the infrared blackbody temperature from the second epoch (as it is the best constrained one) with a temperature of 1762 K, using the default grain size and the peak optical luminosity of 1.4×10^{45} erg s, this equation yielded a sublimation radius of 0.2 pc, which is in good agreement with the best-fit value of 0.16 pc.

To obtain the covering factor, i.e. the fraction of the shell around the transient that actually contains dust, I needed to derive the total bolometric energy of the dust echo. This was achieved by time-integrating the fitted dust echo light curve, scaled to the peak luminosity of 4.3×10^{44} erg s. This resulted in a dust echo bolometric energy of 1.1×10^{45} erg. The ratio of the dust echo bolometric energy to the optical/UV bolometric energy yielded a covering factor of 1/3, assuming a 100 % efficiency in the absorption-reemission process. This value was unusually high (TDEs in quiescent galaxies usually have covering factors ~1 %, see [295]). The high covering factor can be explained by the fact that—contrary to usual TDEs—the system is an active galactic core, with a pre-existing dust torus.

Lastly, I used this covering factor and the radius of the dust shell to determine the dust mass. Again assuming a grain size of 0.1 µm, spherical dust grains and a typical dust mass density of 2.5 g cm^{-1} [296] resulted in a mass of 0.017 $M_\odot$ contributing to the echo. When assuming a typical mass-to-gas ratio of 1:100 [297], this corresponds to 1.7 $M_\odot$ in gas.

[296]: Mann et al. (2000), *Interstellar dust properties derived from mass density, mass distribution, and flux rates in the heliosphere*

[297]: Leroy et al. (2011), *The CO-to-H$_2$ Conversion Factor from Infrared Dust Emission Across the Local Group*

5.5 Classification

As stated at the beginning of this chapter, *AT2019fdr* was initially classified as a SLSN IIn, but a subsequent paper classified *AT2019fdr* as TDE [258]. In this section I will present their reasoning, as well as further evidence we brought forward in favor of this hypothesis, but also some data weakening this interpretation.

5.5.1 Original TDE Classification

Frederick et al. [258] outright rejected the initial SLSN IIn interpretation, but discussed that *AT2019fdr*—like all the 5 flares located in NLSy1 galaxies they studied—could either be an AGN flare or a TDE. They based their classifications on a set of criteria, detailed in Table 5.5. The green cells in the table show that the feature being present (✓) or absent (×) makes a TDE interpretation more likely, while red color means that it renders a flare origin more likely. For comparison, *AT2019dsg* is shown as an example of a bona fide TDE.

Table 5.5: Classification matrix of *AT2019fdr* and ambiguous transient *PS1–10adi* for comparison, as well as bona fide TDE *AT2019dsg*. ✓ means the property is present, while × marks an absence. Green means that the presence or absence favors a TDE interpretation, while red is evidence for an AGN flare interpretation. (✓) means the presence of soft X-rays. Adapted from [258], with additions by the author.

Name	$M_{\mathrm{BH}} <$ $10^8 M_\odot$	$\mathrm{H}_\beta < 2000$ $\mathrm{km\,s^{-1}}$	FeII	$[\mathrm{OIII}]/\mathrm{H}_\beta$ < 3	$\Delta g - r$ ~ 0 mag	UV	X-ray & Γ	*W1*-W2 > 0.7 mag	Re-brighten
AT2019fdr	✓	✓	✓	✓	×	✓	×	×	×
PS1–10adi	✓	✓	✓	✓	✓	✓	(✓)	×	×
AT2019dsg	✓	✓	×	✓	✓	✓	(✓)	×	×

In the case of *AT2019fdr*, the host galaxy black hole mass lies below the Hills mass ($\lesssim 10^8 M_\odot$), it was persistently detected in UV around the peak, but initially not in X-ray wavelengths, its *WISE W1-W2* color was below 0.7 mag, and it showed no signs of rebrightening at the time of publication.

On the other hand, its narrow H_β emission feature and its $[\mathrm{OIII}]/\mathrm{H}_\beta$ flux ratio, the strong FeII complex and the transient's lack of cooling rather stipulated a flare of the underlying NLSy1 host.

To conclude: At this point in time, a thin majority of features observed in *AT2019fdr* were 'TDE-like'. Also, it resembled *PS1–10adi*, a bright transient discovered by PS1 on August 1, 2010 [298], which has been classified as a TDE embedded in an AGN [299]. For these reasons, the transient was ultimately classified as TDE. Following the subclassification scheme proposed in [300], it was a H-only TDE, with the FeII complex most likely stemming from the NLSy1 host galaxy.

[298]: Kankare et al. (2017), *A population of highly energetic transient events in the centres of active galaxies*

[299]: Jiang et al. (2019), *Infrared Echo and Late-stage Rebrightening of Nuclear Transient PS-10adi: Exploring the Torus with Tidal Disruption Events in Active Galactic Nuclei*

[300]: Velzen et al. (2021), *Seventeen Tidal Disruption Events from the First Half of ZTF Survey Observations: Entering a New Era of Population Studies*

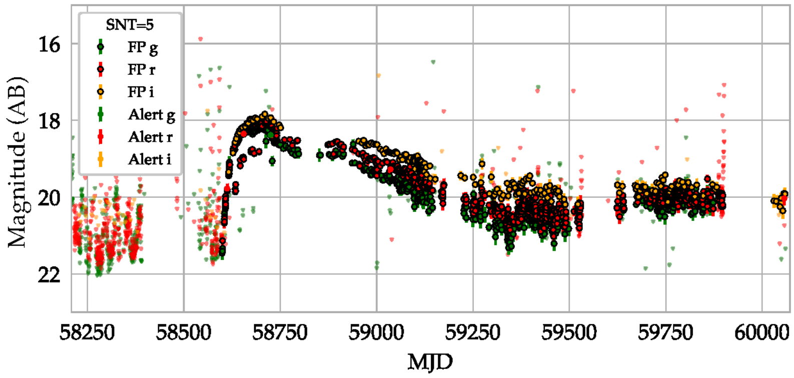

Figure 5.17: *AT2019fdr* late stage rebrightening. Starting around MJD = 59500, the flux in all ZTF bands increased, until it reached a plateau ~ 200 days later.

After the publication of our paper on *AT2019fdr*, the transient underwent significant rebrightening in the optical. About 2.2 years after the optical peak and a continuous decline ever since, the flux in all three ZTF bands started to increase again, until it hit a plateau roughly 200 days later. As of April 2023, *AT2019fdr* is still in this plateau. This is atypical behavior for TDEs and SNe, as was already pointed out in [258]. For now, it remains to be seen how long-lived that rebrightening is.

Interestingly, also *PS1–10adi* showed late-stage rebrightening [299]. For this transient, the authors discussed a variety of possible causes for that behavior. They concluded that late-stage accretion was unlikely, based on the different properties of the dust echoes of the initial flare and the rebrightening. They also discarded interaction between the debris and the SMBH torus, based on the fact that the complete conversion of the unbound debris' orbital energy was insufficient to explain the luminosity of the rebrightening. Instead, they favored a scenario where a mildly relativistic outflow from the TDE is responsible for the rebrightening [299].

It remains an open question if the long-lasting rebrightening of *AT2019fdr* can be explained by such an outflow. One could also hypothesize that a perturbation in the accretion disk, triggered by the TDE, is responsible for the increased optical flux; see e.g. [301].

[301]: Ricci et al. (2022), *Changing-look Active Galactic Nuclei*

5.5.2 Soft X-ray Signal and Dust Echo: More Evidence for a TDE

As already discussed in Section 5.1.2, an X-ray signal from *AT2019fdr* was detected by *SRG*/eROSITA on its third visit with an unusually soft thermal spectrum of 56^{+32}_{-26} eV. As AGN X-ray spectra are rarely soft [302], this significantly strengthens the TDE interpretation.

[302]: Saxton et al. (2020), *X-Ray Properties of TDEs*

NLSy1 galaxies in general exhibit softer spectra, but the temperature of *AT2019fdr* was atypically low also in this context: For example, it was lower than all NLSy1 temperatures in [303] and [304]. Additionally, SLSNe rarely emit X-ray radiation. Only the first SLSN ever detected, *SCP 06F6* [305], possibly showed an X-ray flux exceeding *AT2019fdr*'s luminosity [306].

The dust echo (see Section 5.4) we found, together with the energy budget and the bolometric evolution suggests that *AT2019fdr* belongs to an emerging class of strong TDE candidates located in AGN, as e.g. *PS1-10adi* (discussed above), *AT2017gbl* [307], or *Arp 299-B AT1* [308].

[303]: Leighly (1999), *A Comprehensive Spectral and Variability Study of Narrow-Line Seyfert 1 Galaxies Observed by ASCA. II. Spectral Analysis and Correlations*

[304]: Vaughan et al. (1999), *X-ray spectral complexity in narrow-line Seyfert 1 galaxies*

[305]: Barbary et al. (2008), *Discovery of an Unusual Optical Transient with the Hubble Space Telescope*

[306]: Gänsicke et al. (2009), *SCP 06F6: A Carbon-Rich Extragalactic Transient at Redshift* $z \simeq 0.14$

[307]: Kool et al. (2020), *AT 2017gbl: A dust obscured TDE candidate in a luminous infrared galaxy*

[308]: Mattila et al. (2018), *A dust-enshrouded tidal disruption event with a resolved radio jet in a galaxy merger*

5.6 Chance Coincidence

To compute the chance coincidence of observing a TDE comparable to *AT2019fdr* in spatial and temporal coincidence with a high-energy neutrino in addition to the first neutrino-TDE association, I created a sample of TDEs and candidate TDEs of similar or higher brightness.

The sample from which these sources were selected was the full sample of nuclear transients as selected by a filter implemented in `AMPEL` (see Section 4.4). This filter required at least 10 detections in both the g- and the r-band. Additionally, we required a weighted maximum distance of the flare to its host of ≤ 0.5 arcsec, and that most datapoints had positive flux after subtraction of the reference image. A detailed discussion of the nuclear filter can be found in Section 6.1.1.

I restricted the dataset to transients first detected after January 1, 2018 to ensure that all of them had high-quality reference images. I also required the transients to have peaked before July 2020. In total, 3172 flare candidates remained after these cuts.

I further required that the transients were not classified as variable stars, or bogus objects. Additionally, their rise (fade) e-folding time was required to exceed the uncertainty on this parameter. At this stage, 1628 candidates remained.

As only candidates with brightness comparable to or higher than *AT2019fdr* were of interest, I required their peak apparent magnitude to be ≤ *AT2019fdr*'s peak apparent magnitude; this left 157 transients. To select candidate TDE and accretion flares, I required the rise (fade) e-folding time to lie in a [15,80] ([30,500]) day interval, which left 25 transients.

Lastly, I excluded those transients from the sample of 25 transients which were spectroscopically classified as supernovae, or could be ruled out as short-timescale AGN variability by visual inspection. I also excluded candidates which displayed no consistent color or color evolution post peak, or which had a non-smooth light curve evolution post peak. This left a final sample of 12 transients.

I then calculated the effective source density ρ_{eff}—i.e. the density of sources per deg^2 of sky in the survey footprint at any given time—as follows:

$$\rho_{\text{eff}} = (N_{\text{flare}} \times \Delta t_{\text{flare}})/(A_{\text{ZTF}} \times \Delta t_{\text{search}}), \tag{5.7}$$

where N_{flare} is the total number of comparable sources (including *AT2019dsg* and *AT2019fdr*), Δt_{flare} is the typical duration of such events, A_{ZTF} is the sky area accessible to ZTF, excluding sources with a galactic latitude $|b| < 7°$, and Δt_{search} is the time range of the sources sample, i.e. 2.5 years. With the 12 events of the comparison sample, the effective source density was $\rho_{\text{eff}} = 1.71 \times 10^{-4}\,\text{deg}^{-2}$.

Now an expectation value for the number of neutrinos coincident with two sources (*AT2019dsg* and *AT2019fdr*) can be calculated as $\mu = \rho_{\text{eff}} \times A_{\text{IC}}$, where A_{IC} is the summed 90 % uncertainty area of the alert neutrinos we followed up. This was 154.33 deg^2 at that time, which resulted in $\mu = 0.026$. Employing a Poisson distribution, the probability of finding two sources by chance was $p(2) = 3.4 \times 10^{-4}$, while finding one source only had a p-value of $p(1) = 2.6 \times 10^{-2}$ [259].

To obtain the probability of finding two neutrinos coincident by chance, I calculated the probability of observing 0 and 1 coincidences, and subtracted these from 1—this yields the probability of observing two or more coincidences:

$$p(2) = 1 - \sum_{k=0}^{1} \frac{e^{-\mu}\mu^k}{k!}, \tag{5.8}$$

where the sum is the cumulative distribution function of the Poisson distribution. Using a right-tailed normal distribution, $p(2) = 3.4 \times 10^{-4}$ translates to 3.4 σ (1.94 σ for a single association).

5.7 Dust-Echo TDEs as Neutrino Sources?

I have so far detailed the data reduction for *AT2019fdr*, established that it likely was a TDE—albeit an unusual one—modeled its SED and the dust echo we discovered, calculated its energy output and showed that a chance coincidence of two high-energy neutrinos with sources comparable to *AT2019fdr* is quite improbable.

5.7.1 Adding a Third Event: *AT2019aalc*

During the work on *AT2019fdr*, another candidate emerged as a potential high-energy neutrino counterpart, also accompanied by a dust echo. This section is dedicated to that event, *AT2019aalc*.

As both high-energy neutrino associated TDEs had prominent dust echoes, this prompted a search for other optical flares accompanied by a strong dust echo and in coincidence with high-energy neutrinos [263]. That study found a third candidate counterpart that so far was missed as its sky location was inaccessible to ZTF at the time of neutrino detection due to its proximity to the sun.

5.7.2 Event Details

AT2019aalc is the possible counterpart to neutrino *IC191119A* [263]. It featured the largest dust echo luminosity of all three transients, while the dust echo strength as defined in [263] was a bit lower due to pre-flare variability.

In contrast to *AT2019dsg* and *AT2019fdr*, much less is known about *AT2019aalc*. Early scanning discarded it as probable AGN activity (which it still could be), and it was not studied further until it emerged as neutrino counterpart candidate many months later (the neutrino itself arrived about 5 months after optical peak). For this reason, no UV measurements or spectra were taken. This is why the classification as a TDE is a much more tentative one, but nevertheless compatible with the light curve.

AT2019aalc's peak g-band luminosity was comparable to *AT2019fdr* and it was also detected by *SRG*/eROSITA during one visit. Like *AT2019fdr* and *AT2019dsg*, it displayed a soft thermal spectrum with a blackbody temperature of 172^{+10}_{-10} eV; this counts as evidence favoring a TDE interpretation. Also, there was an archival radio detection by VLASS.

The chance coincidence of finding three such dust-echo events in spatial and temporal coincidence with a high-energy neutrino was found to be 3.7σ. The soft X-ray detections were not part of the initial selection criteria, but rather emerged ex post.

5.7.3 Comparing the Candidate Counterparts

Table 5.6 shows a comparison of the relevant measurements and inferred quantities of the three events. One commonality is that the neutrino was significantly delayed with respect to the optical peak, by about 5–10 months. All transients were detected in X-ray wavelengths with a low temperature, and all were accompanied by a prominent dust echo. There have been attempts to draw conclusions from these facts.

Table 5.6: Comparing the multi-messenger properties of the candidate TDEs coincident with high-energy neutrinos. The dust echo strength was defined as $\Delta F/F_{\mathrm{RMS}}$ (flux increase after optical/UV peak vs. the root mean square of the pre-flare flux); 'ν production' means that models exist showing that the source might be capable of producing a neutrino with the respective detected energy; for 'ν signalness' see Definition 2.4.1. '—' denotes insufficient modeling or missing data. Table by the author, adapted from [309].

Property	***AT2019dsg***	***AT2019fdr***	***AT2019aalc***
TDE	yes	strong candidate	candidate
Peak bol. luminosity	3.5×10^{44} erg s^{-1}	1.3×10^{45} erg s^{-1}	—
SMBH Mass	$10^{6} - 10^{6.7} M_{\odot}$	$10^{7.55} M_{\odot}$	$10^{7.2} M_{\odot}$
Radio	evolving	not evolving	archival det.
UV	very bright	bright	—
X-ray	early, soft spectrum	late, soft spectrum	soft spectrum
Dust echo strength	92.2	39.2	15.7
ν delay	$\sim$ 5 months	$\sim$ 10 months	$\sim$ 5 months
ν production	possible	possible	possible
ν energy	217 TeV	82 TeV	176 TeV
ν 90% uncertainty box	25.5 deg^2	25.2 deg^2	61.2 deg^2
ν signalness	0.59	0.59	0.45

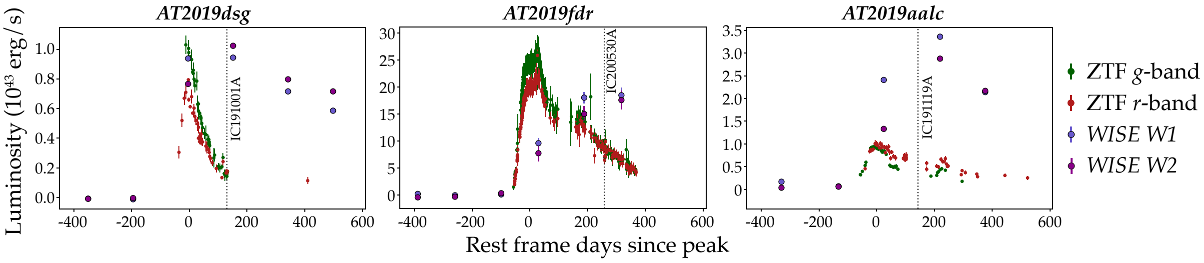

Figure 5.18: All three candidate counterpart TDEs with ZTF optical light curve and *WISE* detections, showing the relative strength of the dust echo. The neutrino arrival times are shown as black dotted vertical lines. Figure adapted from [263].

The observed delay might either be a statistical effect, or it could carry physical meaning. In [263] it was explained by the assumption that the debris first has to circularize before the first neutrinos can be produced. In [310] the observed time delay of the neutrino was combined with the fact that the dust echo was peaking around the neutrino detection time for all three transients, as one can see in Figure 5.18.

[310]: Winter et al. (2023), *Interpretation of the Observed Neutrino Emission from Three Tidal Disruption Events*

One of the models presented in [310] makes use of that fact, as the photons from the dust echo serve as targets for the protons accelerated within the source. Therefore, the neutrino delay arises naturally from the delay of the infrared emission. The energies involved in this model would make TDEs interesting candidates for the production of Ultra-High Energy Cosmic Rays (UHECR). The downside of this model is that it requires very high proton energies.

A companion model swaps the infrared target photons for X-ray photons, as all three sources showed signs of soft X-ray emission. That model can also explain the time delay, which arises from the confinement of moderate energy protons in this case. The observed neutrino energies are explained better here, but the observed neutrino delay is described less well in comparison to the IR model. Lastly, a third model uses optical/UV target photons (as was done in [263]). This yields the highest neutrino production efficiency, but does not explain the observed time delay of the neutrino.

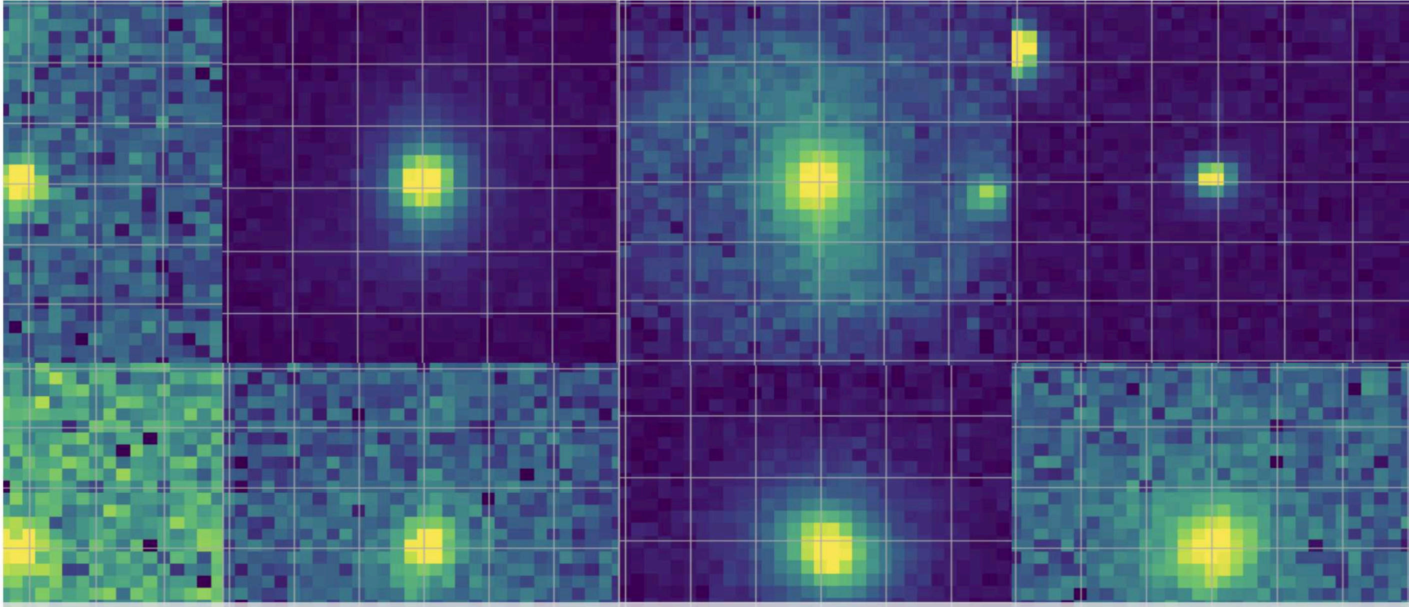

6 The Nuclear Sample

Cutouts of several transients in the nuclear sample, as seen in the ZTF g-band near optical peak.

Motivated by the three neutrinos coincident with accretion phenomena occurring in galactic nuclei, I created a systematic sample of nuclear transients in the ZTF footprint. No such sample existed yet, as AGN studies are usually interested in long-term variability of galaxies, while transient studies normally exclude phenomena too close to galactic cores as contaminants. Such a systematic sample of partially unclassified transient phenomena might provide additional insights into the physics of galactic nuclei.

Furthermore, with the advent of a deep high-cadence sky surveys in the form of the Legacy Survey of Space and Time (LSST, hosted by the Vera C. Rubin Observatory) [175] in the near future, photometric identification of transients will be crucial.

[175]: Ivezić et al. (2019), *LSST: From Science Drivers to Reference Design and Anticipated Data Products*

The rate of transients detected by LSST will by far exhaust the available spectroscopic resources, thus requiring informed decisions about when to rely on spectroscopy for classification and characterization. The vast majority of LSST-detected transients will either be photometrically classified, or not classified at all. Therefore, photometrically classifying the ZTF nuclear transients can serve as a precursor study to LSST-era astronomy.

This chapter is dedicated to transients similar to the three transients for which a neutrino association could be made: Nuclear transients, i.e. transients observed close to the cores of their galaxies. Because no clear picture on the composition of such a sample exists, we gathered one and made a first attempt at classifying it.

The routines used to create and classify the nuclear sample were written in Python, and are accessible online[1].

6.1 Sample Creation

The nuclear sample was created with `AMPEL` (see Section 4.4), using its capability to rerun analyses on archival data. To perform such a rerun, a modified version of the `AMPEL` nuclear filter was used[2].

---

[1] `https://github.com/simeonreusch/ztfnuclear`

[2] `https://github.com/AmpelAstro/Ampel-nuclear`

6.1.1 `AMPEL` Nuclear Filter

This filter is used primarily by the ZTF TDE working group to scan for transient activity that is compatible with an emerging TDE [300]. In the rerun, it was slightly modified, and used to re-evaluate each and every alert issued by ZTF.

The logic behind the filtering process can be explained as follows: The filter selects events with decent photometric quality, ensured by requirements on the number of detections and their brightness. The event at least once needs to be judged as being 'real' (opposed to bogus), and the host needs a high probability of being a galaxy, not a star. Furthermore, *Gaia* is consulted to veto against stars, and the events need to be nuclear, i.e. happen close to the core of their host galaxy.

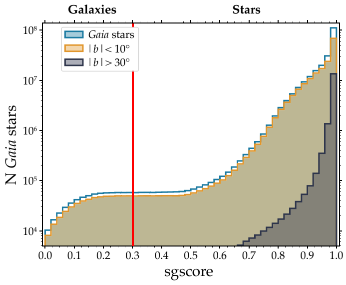

Figure 6.1: `sgscore` performance evaluated with known *Gaia* stars. At the chosen threshold of 0.3 (red line), the misidentification of stars as galaxies is negligible. Adapted from [199]

The criteria used in this study were as follows:

`sgscore` As detailed in Section 3.4.8, `sgscore` is a machine-learning based star-galaxy score for PS1 objects (low values: galaxy, high values: stars). The transient at least once needed to have an `sgscore` < 0.3 to pass the filter.

Number of detections At least 3 detections in both ZTF g- and r-band were required.

Proximity to galactic plane The object had to be separated by at least 5° from the galactic plane to avoid contamination by foreground stars.

PS1 photometry To avoid crowded areas, transients for which more than 100 objects in the vicinity had a counterpart in PS1 were removed.

Brightness At least one alert datapoint of the transient had to be brighter than 20 mag.

`rbscore` The real-bogus score separating erroneous detections (low values) from real ones (high values, see Section 3.4.9) must be larger than 0.3. Note that for more recent data, also `deep real bogus` is available, which promises considerably better results [311]. Unfortunately, alerts from 2018 or 2019 do not contain this information, and there is no direct translation from an `rbscore` threshold to a `drbscore` threshold or re-analysis of the older alerts available. For these reasons and to maximize consistency, I restricted myself to using only the older `rbscore`. The quite loose cut of 0.3 does not entail overly large contamination, as I also required the transient to have a PS1 counterpart. Figure 6.2 shows a comparison of the False Negative Rate (FNR) of `rbscore` vs. `drbscore`. At the chosen threshold of 0.3, the FNR for both algorithms lies at the percent level.

[311]: Duev et al. (2019), *Real-bogus classification for the Zwicky Transient Facility using deep learning*

Core distance For all objects that made it this far, their angular distance to the core was computed. To make this more robust, three different distance metrics were employed: The mean distance to the PS1 source in the reference images, the median distance to that, and lastly, a weighted distance. The latter was computed according to [312] and accounts for the fact that the RMS of the angular core distance scales linearly with magnitude. $\sigma_{\mathrm{dist}} = 0.24 + 0.04(m - 20)$, where m is the difference photometry magnitude. This stage accepted all transients for which at least one of the three angular distances lay below 0.5 arcsec.

[312]: Velzen et al. (2019), *The First Tidal Disruption Flare in ZTF: From Photometric Selection to Multi-wavelength Characterization*

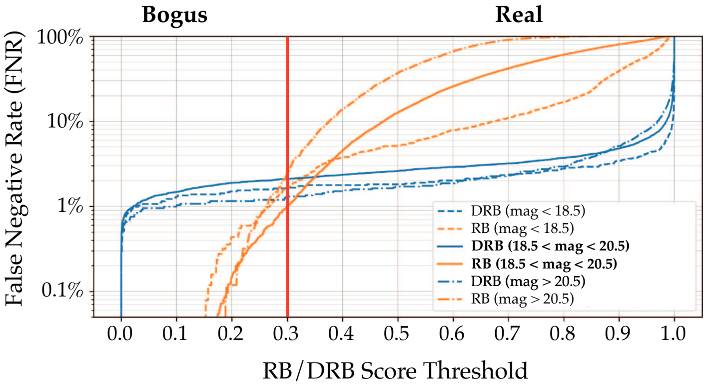

Figure 6.2: `rbscore` vs `drbscore` performance, evaluated in terms of false negative rate as a function of the threshold. Adapted from [311]. The chosen threshold of 0.3 is shown as red vertical line. At that value, `rbscore` has a False Negative Rate of ≈1 % in the relevant magnitude range of 18.5 < mag < 20.5.

This filter was applied to all ZTF alerts issued between 1 April, 2018 and 30 April, 2022, comprising over 4 years of data in total.

6.1.2 Rejection Statistics

From a total sample of ∼ 350 Million alerts issued by ZTF during these 4 years, ultimately 11,687 nuclear transients were selected. The survival rates during the filtering process are shown in Fig. 6.3, with the complete alert stream on top, different thematically grouped rejection stages in between and the final accepted alerts on the bottom. In total, 0.2 % of alerts issued by ZTF were accepted by the filter.

As one can see, the vast majority of alerts were rejected based on them being either too faint, likely being bogus, or likely being stars. The majority of alerts already got filtered out by the initial SQL query. The SQL-based filtering was also the most efficient one, so I ensured that the amount of filtering at that stage (on the archive server) was maximized.

This was followed by the live transfer of all surviving alerts to another machine and processing of the subsequent filtering stages. In total, the filtering process took roughly 5 weeks.

6.1.3 Forced photometry

To be sensitive to early and late-time light curve evolution, forced photometry was acquired for all 11,687 transients making the final cut. This was again done with `fpbot`, see Section 4.6 for details. The process proved somewhat cumbersome due to the enormous data volume (several hundreds of GB) that was required to be transferred from IPAC and stored and processed in batches due to computing center restrictions.

6.1.4 Infrared Data

Motivated by the strong dust echo detected for *AT2019dsg*, *AT2019fdr* and *AT2019aalc*, the optical forced photometry dataset was complemented with infrared light curves from the *WISE* mission. These were obtained to serve in a selection of interesting transients based on their infrared dust echo, see Section 6.9.

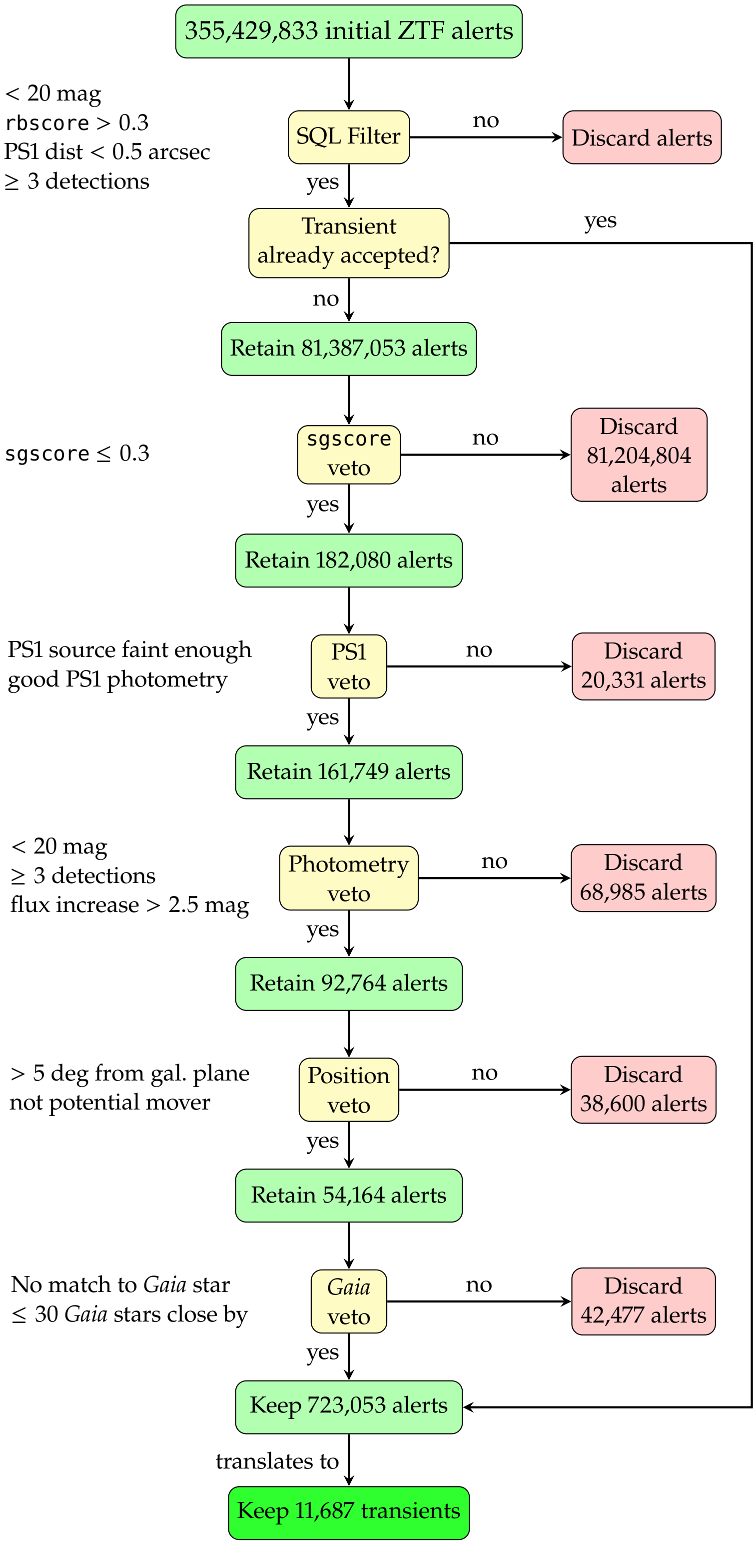

Figure 6.3: Flow chart showing the rejection statistics of the nuclear filter.

[313]: Necker et al. (2023), *timewise: v0.3.13 Release*

Data retrieval was performed with the `timewise` package [313] to download all datapoints available for each source location in the *W1*- and *W2*-band. The tool automatically crossmatched the ZTF transient position to the *WISE* light curve repository, downloaded the photometry for each source and binned the single source exposures of each epoch, spaced in half-year intervals.

Most sources did have infrared counterparts. To obtain a selection of likely dust-echo candidates, two `AMPEL` packages (`T2BayesianBlocks` and `T2DustEchoEval`) were used to analyze the infrared light curves with a Bayesian block algorithm in order to identify periods of flaring activity compared to a baseline; see also Section 6.3 for a discussion of the same algorithm utilized for ZTF optical data.

The significance of these flares was then calculated with a metric dubbed 'dust echo strength'. This was defined in [263] as $\Delta F_{\mathrm{IR}}/F_{\mathrm{RMS}}$, where ΔF_{IR} is the difference between the flaring period flux and the baseline flux, while F_{RMS} is the root mean square of the baseline flux.

6.1.5 Catalog Matching

To enrich the sample by information available on the transients, these were crossmatched to a variety of catalogs and services. The results of the crossmatches were stored locally in a `MongoDB` database for ease of retrieval. The crossmatches were either used to extract training features (e.g. *WISE* colors) or evaluate the quality of the trained models (e.g. GROWTH Marshal, Fritz or TNS). In detail, these crossmatches comprised:

Spectroscopic Redshifts To obtain spectroscopic redshifts, the `AMPEL` module `T2DigestRedshifts` was employed. This queries the following services: A local database of the spectroscopic redshifts contained in the NASA/IPAC Extragalactic Database (NED)[3], spectroscopic redshifts from SDSS, and finally spectroscopic redshifts from the Galaxy List for the Advanced Detector Era (GLADE) [314] v2.3.

Photometric Redshifts Additionally, `T2DigestRedshifts` also provided photometric redshifts from the Legacy Survey [315], the 2MASS Photometric Redshift Catalog [316], and photometric redshifts from the PS1 Source Types and Redshifts with Machine learning (PS1-STRM) catalog [317] as well as GLADE.

GROWTH Marshal/Fritz During ZTF Phase I, the GROWTH Marshal [318] served as a community hub to gather information on single transients. It provided a web interface to upload spectroscopy, discuss photometry and add classifications and redshift. This service has been replaced by Fritz [319] with the same design goal, but greater modularity and API support. Both services were queried for transient classifications and redshifts.

TNS The Transient Name Server (see Section 4.5.2) was also used to obtain classifications and redshifts of known transients.

AllWISE Additionally to the *WISE* light curves, archival data from the first part of the *WISE* mission was obtained. This dataset had the advantage of providing two additional bands reaching into the far infrared (*W3* and *W4*), both of which were deactivated after the nominal mission end of *WISE* due to lack of coolant. These allowed for the calculation of more colors, which were used in AGN rejection, see Section 6.4.3.

CRTS DR1 The Catalina Real-time Transient Survey Catalog contains cataclysmic variables which were crossmatched against to reduce contamination by foreground stars.

[3] `https://ned.ipac.caltech.edu`

[314]: Dálya et al. (2018), *GLADE: A galaxy catalogue for multimessenger searches in the advanced gravitational-wave detector era*

[315]: Zhou et al. (2020), *The clustering of DESI-like luminous red galaxies using photometric redshifts*

[316]: Bilicki et al. (2013), *Two Micron All Sky Survey Photometric Redshift Catalog: A Comprehensive Three-Dimensional Census of the Whole Sky*

[317]: Beck et al. (2020), *PS1-STRM: Neural network source classification and photometric redshift catalogue for PS1 3π DR1*

[318]: Kasliwal et al. (2019), *The GROWTH Marshal: A Dynamic Science Portal for Time-domain Astronomy*

[319]: Coughlin et al. (2023), *A Data Science Platform to Enable Time-domain Astronomy*

SDSS The Sloan Digital Sky Survey was also used to crossmatch against foreground stars.

6.2 Creating Features: Light Curve Fits

To photometrically classify the transients, the following strategy was employed: Fit all transients with a dedicated TDE light curve model, as well as a supernova Ia model. The features extracted from these fits, appended by additional ones, served as a base to classify the nuclear sample. Such a training task required a truth on which to train on. The truth base chosen here was the BTS sample, which was augmented by creating copies of existing objects, as well as adding new, underrepresented transients. All those steps will be explained below.

6.2.1 Fitting a TDE Model

To model a TDE-like source evolution, a variation of the parametrization presented in [300] was used.

TDE Parametrization

The basic idea is to fit the light curve evolution with a Gaussian rise and exponential decay of a blackbody, which is also allowed to linearly change its temperature for a number of days. The luminosity evolution is then given by

$$L(t,\lambda) = \frac{T_{\mathrm{peak}}^4}{T(t)^4} B_\lambda F(t), \tag{6.1}$$

where λ is the evaluated wavelength, T_{peak} is the blackbody temperature at peak, $T(t)$ is its temperature at time t and $B(t)_\lambda$ is the spectral radiance of the blackbody. $B(t)_\lambda$ is given by

$$B(t)_\lambda = \frac{2ch}{\lambda^5} \frac{1}{\exp(\frac{hc}{\lambda k_B T(t)}) - 1}, \tag{6.2}$$

and $F(t)$ is the Gaussian rise pre-peak and exponential decay post peak, given by

$$F(t) = \begin{cases} \exp[-(t - t_{\mathrm{peak}})^2/2\sigma^2] & \text{if } t \leq t_{\mathrm{peak}} \\ \exp[-(t - t_{\mathrm{peak}})/\tau] & \text{otherwise.} \end{cases} \tag{6.3}$$

Here, σ is the rise time of the transient, and τ is the decay time, both in days.

Lastly, $T(t)$ was allowed to change linearly during an interval after t_{peak} until t_{cutoff}, marking the end of the linear temperature evolution:

$$T(t) = \begin{cases} T_{\mathrm{peak}} & \text{if } t < t_{\mathrm{peak}} \\ T_{\mathrm{peak}} + t \cdot \Delta T & \text{if } t_{\mathrm{peak}} \leq t \leq t_{\mathrm{cutoff}} \\ T_{\mathrm{cutoff}} & \text{otherwise,} \end{cases} \tag{6.4}$$

with ΔT being the coefficient of temperature change. Also, there was a check that required at least 10 datapoints within the time interval 30 days prior to peak time and one year after that. If that check did not succeed, the fit was set to fail due to poor source sampling.

As a source luminosity can only be calculated when the redshift is known, and this was not the case for a majority of the nuclear transients, an arbitrary amplitude was used instead. This does not affect the results in any meaningful way, as TDEs can differ in brightness by at least two orders of magnitude (see e.g. [87]), and only the shape of the light curve is relevant for the goodness of fit[87].

This model has been realized as an instantiation of an `SNCosmo` source model, as this package had the advantage of built-in filter profiles which allowed the correct evaluation of the blackbody flux as seen through the ZTF bandpass filters. To make the fit procedure computationally feasible, the algorithm used was not a Markov-Chain Monte Carlo as in [300], but `iminuit/MIGRAD`.

To account for Milky Way dust extinction along the line of sight, the infrared dust maps by Schlegel, Finkbeiner & Davis were used [320], as provided by the `sfdmap`[4] Python package.

[320]: Schlegel et al. (1998), *Maps of Dust Infrared Emission for Use in Estimation of Reddening and Cosmic Microwave Background Radiation Foregrounds*

[4] `https://github.com/kbarbary/sfdmap`

TDE Fit Priors

To constrain the fits and mitigate runaway, fit priors were used, partly taken from [300]. These are shown in Table 6.1.

The prior on the time of peak, here dubbed t_0, was inferred from a peak-finding algorithm that identified the highest flux point in each band (given it had a signal-to-noise ratio of > 3) after smoothing the light curve with a rolling window, calculating the median flux within a window of 10 days.

Param.	Description	Prior	Bounds
t_{peak}	Peak time	t_0	$t_0 \pm 30$ days
$\log T_{\text{peak}}$	Peak temperature	4 K	[3.5, 5] K
$\log \sigma$	Gaussian rise time	1.6 days	[0,5] days
$\log \tau$	Exp. decay time	2.3 days	[0,5] days
ΔT	Temp. change / day	0	$T_{\text{peak}} \pm 15000$ K (total)
t_{cutoff}	End of temp. evol.	300	[100, 1200] days $+t_0$

Table 6.1: Priors for the TDE fit.

As the blackbody fits were not constrained by UV data, in some cases a temperature runaway occurred when allowing all fit parameters to vary freely within their bounds. This can be explained as follows: A decrease in brightness over time can either be achieved by exponential decay or by an increase in temperature. The latter gradually moves the blackbody spectrum outside the ZTF bands, which then translates to a decrease in brightness. As such excessive temperature changes are most likely unphysical, the daily temperature change ΔT was limited to an integrated maximum change of $T_{\text{peak}} \pm 1.5 \times 10^4$ K.

Another measure to mitigate the issue was to perform the fits as a two-stage process: First, the temperature evolution was neglected, i.e. ΔT was set to 0 and t_{cutoff} was removed as parameter, with the blackbody only having one temperature for the duration of the light curve, T_{peak}.

The best-fit values for t_{peak}, $\log T_{\mathrm{peak}}$, $\log \sigma$ and $\log \tau$ obtained in this first stage were then used as fit priors for stage 2. This procedure solved the temperature runaway in almost all cases.

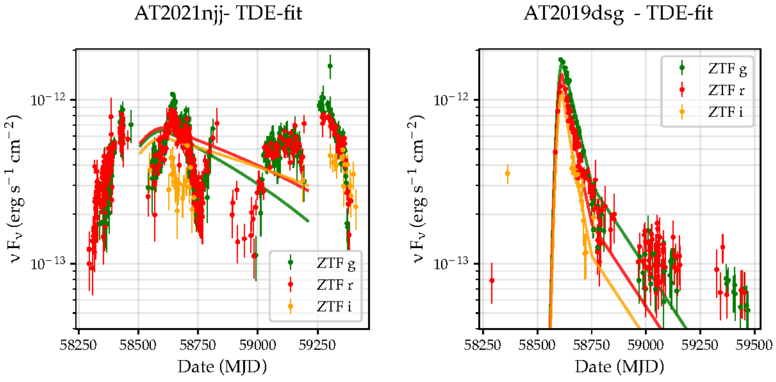

Figure 6.4: Two exemplary TDE light curve fits. Left: *AT2021njj*, displaying periodic AGN activity not well captured by the fit. Right: *AT2019dsg*, a TDE which is captured fairly well by the fit routine.

Fig. 6.4 shows two exemplary TDE fits; one for an object of unknown nature (*AT2021njj*) on the left and one for a confirmed TDE (*AT2019dsg*, see Section 4.9.6) on the right. As one can see, the fit captures the light curve evolution of the TDE fairly well, including the change in color due to the changing blackbody temperature. *AT2021njj* on the other hand—whatever it actually is (the object has no spectroscopic classification)—is clearly not a TDE. The fit cannot account for the periodic nature of the object, resulting in a reduced χ^2 of 26.3.

The successful fit for *AT2019dsg* meanwhile results in the following values: A reduced $\chi^2$of 4.1, a peak temperature of 8700 K which decays by 78 K per day for the following 138 days, an initial risetime of 21 days, and a characteristic light curve decay time of 216 days.

This object nevertheless highlights a restriction: The available optical to infrared data from ZTF does not constrain blackbodies well. When including additional UV data for *AT2019dsg* near light curve maximum, the inferred blackbody temperature is 4×10^4 K, significantly higher. This is not a failure of the fit procedure, but driven by the high UV flux pushing the blackbody to shorter wavelengths (the additional UV data can be seen in Fig. 4.16).

6.2.2 Fitting a SN Ia Model

As SNe Ia are a prominent contaminant for bona fide nuclear events (i.e. such events that can only occur in the centers of galaxies), it was crucial to identify them. To achieve this goal, all light curves were additionally subjected to a fit using the tried and tested Spectral Adaptive Lightcurve Template (SALT2) [321] SN Ia light curve model.

[321]: Guy et al. (200710.1088/1538-3873/aafbc2), *SALT2: Using Distant Supernovae to Improve the Use of Type Ia Supernovae as Distance Indicators*

SALT2 assumes that the light curve is indeed a supernova Ia and applies empirical corrections to the light curve color and stretch (i.e. the width of the light curve). This is achieved by matching templates generated by a set of well-sampled Ia light curves and spectra of varying distance. The resulting color and stretch correction parameters, as well as the peak brightness and the reduced χ^2 were saved and will later be used as features in training the classifier.

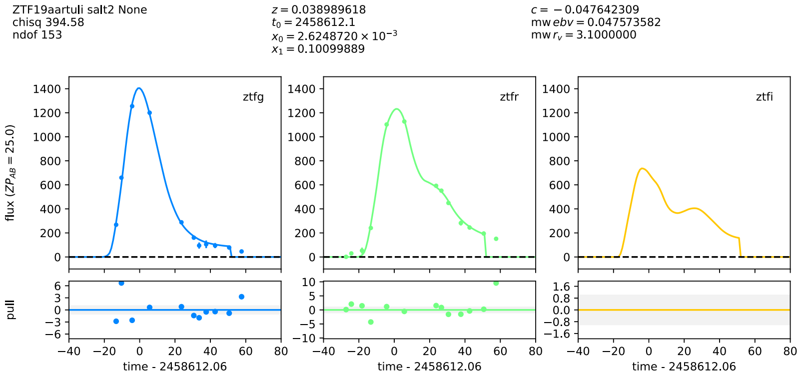

Figure 6.5: Exemplary SALT2 fit output. The three panels show three ZTF bands of *AT2019dzzo*, a spectroscopically confirmed SN Ia. Time 0 is the fitted peak of the assumed SN Ia; and as one can see the transient in question is fairly well approximated by SN Ia light curve, with a reduced $\chi^2 = 2.67$.

The fits were performed with `SNCosmo`, which itself fits the light curves using `iminuit/MIGRAD`.

6.3 Creating Features: Optical Flare Analysis

Another feature which will later be used in the classification of transients were simultaneous flares across the well-sampled ZTF g- and r-bands. The number and timing of optical flares within different bands can then be used to differentiate between different types of transients.

6.3.1 Bayesian Block Algorithm

To determine time periods of heightened activity, a Bayesian block algorithm was employed. This was a version of a package developed at DESY, modified by the author to analyze optical ZTF light curves instead of infrared *WISE* data.

The Bayesian block algorithm is explained in [322, 323]; the implementation used here is integrated in the `astropy` package. The light curves were first smoothed by calculating the median and standard deviation within a 10-day rolling window. After that, only data points lying within 3σ distance in magnitude space from the median were used. This got rid of flux outliers, which were quite frequent for the forced photometry light curves of the nuclear sample.

The Bayesian block algorithm uses a prior on the number of blocks N_{blocks}, which is computed as $P(N_{\text{blocks}}) = P_0 \gamma^{N_{\text{blocks}}}$, where P_0 is a normalization dependent on the number of datapoints N. For this analysis, the prior was computed with a slope γ of $\gamma = N^{-10/\ln 10}$, with N being the number of datapoints in the smoothed single-band light curve. The slope γ was determined empirically to yield robust results. It was a compromise between sensitivity to flux changes and detecting too many blocks, each too small to be meaningful.

[322]: Scargle (1998), *Studies in Astronomical Time Series Analysis. V. Bayesian Blocks, a New Method to Analyze Structure in Photon Counting Data*
[323]: Scargle et al. (2013), *Studies in Astronomical Time Series Analysis. VI. Bayesian Block Representations*

6.3.2 Block Coincidence

As they were much better sampled than the i-band, the g- and r-band were used to check for blocks temporally coincident in both bands.

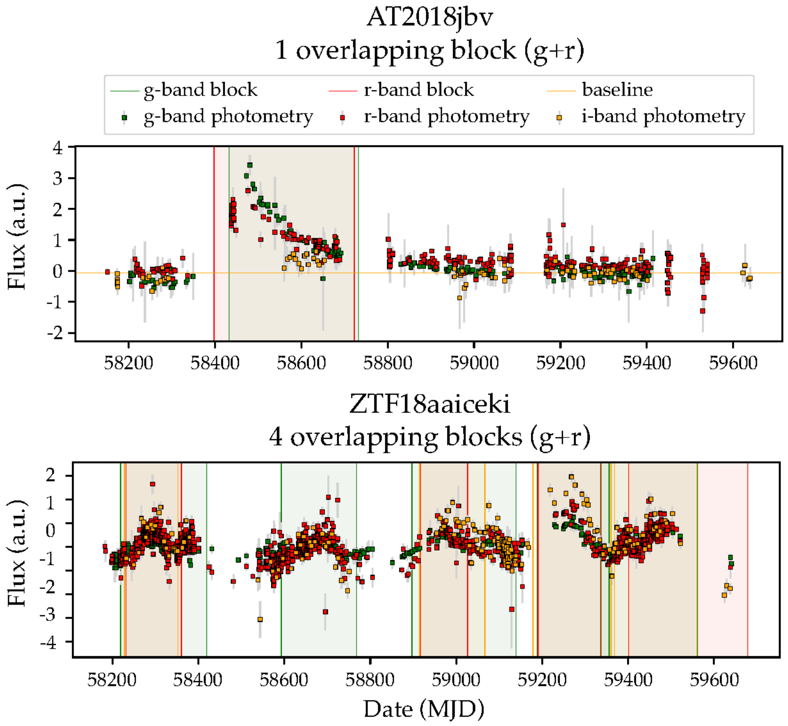

Figure 6.6: Top: Bayesian Blocks identified for *AT2018jbv*. There is one block both in g- and r-band, which overlaps. This is a strong indicator that the transient is a one-time flare, stemming from e.g. a supernova or a TDE (the transient is in fact a TDE), in contrast to stochastic AGN variability. Bottom: Bayesian Blocks for *ZTF18aaiceki*, showing AGN activity. This is correctly captured by a large number of overlapping regions.

For this, each region with increased flux compared to a baseline in the g-band was checked to see if there was a corresponding r-band block with increased flux overlapping in time.

One can see two examples of such overlapping blocks in Fig. 6.6, which shows the result of the Bayesian Block algorithm for *AT2018jbv* and *ZTF18aaiceki*. Both are transients from the nuclear sample, the first one being a TDE, the second an AGN displaying stochastic variability.

6.4 Training Sample

As one of the goals was the classification of the nuclear sample, a training sample that most closely resembled the target sample in brightness needed to be chosen.

6.4.1 The Bright Transient Survey

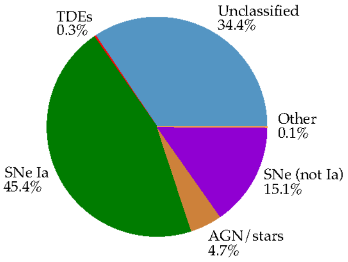

Figure 6.7: Composition of the Bright Transient Survey sample used in this study. The classified part of the sample is heavily biased towards SNe Ia, AGN are vastly under-sampled.

The natural starting point for the creation of a training sample was the so-called ZTF Bright Transient Survey (BTS)[5], see [255, 256] for details. The basic goal of the survey was the spectroscopic classification of all ZTF-detected transients brighter than 18.5 mag in either the g- or r-band. It even managed to push further, achieving a spectroscopic completeness of 75 % at 19 mag (93 % at the nominal cutoff of 18.5 mag) [256].

This was achieved by utilizing the fully robotic SED machine for quick classification, and employing other spectroscopic resources for ambiguous spectra. BTS is the largest spectroscopic supernova survey ever conducted, with over 8000 classified transients as of August 2023.

---

[5] `https://sites.astro.caltech.edu/ztf/bts/bts.php`

There were three potential issues that needed to addressed though when using the BTS as starting point for a training sample:

Brightness bias The BTS restriction to objects usually brighter than 18.5 mag means that the majority of the BTS sample is brighter than the nuclear sample with its magnitude cut of 20.

Class imbalance The BTS is heavily skewed towards supernovae. Also, because SNe Ia are brighter, the majority of the SNe contained in a flux-limited sample like the BTS are of type Ia [256], with 70 % of all classified BTS transients being SNe Ia. This results in a class imbalance of the training sample.

Anti-AGN bias To reduce contamination stemming from AGN variability, the BTS vetoes transient host galaxies that likely harbor an AGN. This is achieved by crossmatching candidate sources to known AGN, based on *WISE* color cuts, or by rejecting sources with previous variability most likely hinting at AGN activity. This poses a problem: As the nuclear sample will contain a large number of AGN, this also results in a class imbalance.

These issues needed to be solved somehow. Two procedures to do that will be described below: Augmentation by creating fainter copies addressed the brightness bias and partially the class imbalance, while rejecting likely AGN within the nuclear sample and expanding the training set with known AGN were both employed to deal with the anti-AGN bias.

6.4.2 Augmentation: Enhancing the Training Set

One promising strategy was to multiply the number of light curves available by simulating observations of the same object at higher redshifts.

The method of augmentation employed here was developed by A. Townsend and the author. The procedure was implemented as Python package [324] and worked as follows:

[324]: Reusch et al. (2023), *ztfparsnip: Release v0.3.0*

Draw new redshift

A transient light curve (from now on: parent light curve) was obtained, as was its redshift z_{parent} (all classified BTS transients have a spectroscopic redshift). A new redshift z_{child} was drawn for the child light curve (i.e. the noisified copy) from a cubic redshift distribution ranging from z_{parent} to $z_{\text{parent}} + 0.1$.

Scale and scatter the flux

After this, the parent flux was redrawn from a normal distribution centered around each parent flux value, scaled by its error. This was done to account for the fact that flux measurements are expected to scatter around their true value with their error, thereby simulating a 'new' measurement which is more noisy due to the increased distance. After this, the re-drawn flux measurements were rescaled with the new redshift z_{child}. The flux F scales according to $D_L = \frac{L}{4\pi F}$ i.e. with the square of the luminosity distance (see Section 5.3.2). Therefore, one can determine a flux scaling factor a according to

$$a = \frac{D_L(z_{\text{parent}})^2}{D_L(z_{\text{child}})^2}. \tag{6.5}$$

With this scaling factor, the child light curve flux is simply $f_{\text{child}} = a f_{\text{parent}}$.

The flux error got scaled accordingly, but with two additional empirical constant coefficients e_f and e_b determined by fitting the flux error as a function of flux. This procedure showed that the error F_{err} deviates from the expected $\sqrt{F}$ behavior, but can be well approximated by $F_{\text{err}} = \sqrt{F + e_f^2 \times F + e_b^2}$. The error as a function of flux, as well as the different approximations are shown in Fig. 6.8. The factors e_f and e_b were determined as $e_f = 0.54$ and $e_b = 5.37$.

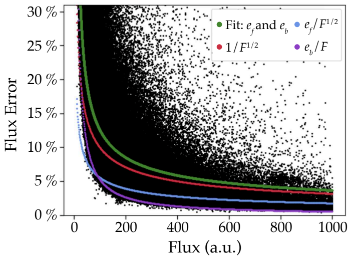

Figure 6.8: ZTF error distribution: F vs. F_{err} in percentage of flux. The red curve shows the expected error behavior ($F_{\text{err}} \propto \sqrt{F}$), while the green curve shows the improved version $F_{\text{err}} = \sqrt{F + e_f^2 \times F + e_b^2}$. Figure by A. Townsend with small modifications by the author.

Apply K-Correction

If an `SNCosmo` model for the transient type existed, a K-correction was applied. This procedure accounts for the fact that by arbitrarily redshifting the light curve of an object, the spectrum will also be redshifted. Due to this, the ZTF bandpasses will register different parts of the spectrum, which leads to a different observed flux (see [325] for an introduction).

[325]: Hogg et al. (2002), *The K Correction*

Scatter and sub-sample observations

As a last step, the newly created light curve was sampled with a default retention fraction of 0.9, i.e. 10 % of datapoints were dropped randomly. Additionally, the observation times were slightly randomized to avoid creating regularities a machine-learning algorithm could pick up. This was achieved by scattering the observation dates around their original values, with a σ of 0.03 days.

Reject faint light curves

Finally, for each generated light curve, at least 5 datapoints were required to lie above a detection threshold of 5σ. If that requirement was not met, the light curve was discarded. This was implemented to mimic the nuclear filter (see Section 6.1.1), which reacts to alerts issued by ZTF. Only datapoints with a signal-to-noise ratio of > 5 warrant an alert, so the noisified light curves needed at least some datapoints above that threshold. The minimal number of detections for the noisified light curves was a bit higher compared to the nuclear filter (5 vs. 3) to account for the fact that forced photometry is slightly more sensitive than the PSF photometry used in the alerts.

Sample dynamically

Depending on the number of light curves for a specific class, the number of desired child light curves per parent light curve was determined to yield an augmented training set with balanced classes. For example, only 3 children per SN Ia light curve were generated, but up to 14 children per light curve of all other types of supernovae. The only exception were variable stars: For these, the redshift ≈ 0 and could not be varied, so no children were generated.

The results of all these steps are exemplified in Fig. 6.9. It shows the light curve of an SN Ia, *SN2019qym*, detected with a redshift of $z = 0.02$. The original light curve can be seen, as well as a child light curve at higher redshift ($z = 0.09$). Note the increased errors (the original object was very bright, so only one datapoint in the g-band at 35 days after peak has visible error bars), the slightly different modified observation dates and the random removal of datapoints.

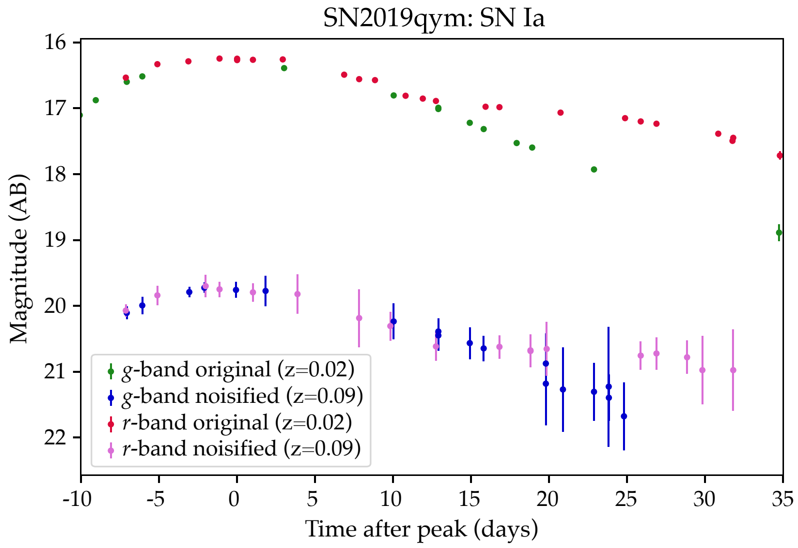

Figure 6.9: Original and 'child' light curve of *SN2019qym*, an SN Ia with a redshift of $z = 0.02$. The original two bands are shown in green and red, while the fainter, noisier copy at a redshift of $z = 0.09$ is displayed in blue and pink. Figure by A. Townsend, with slight modifications by the author.

To match the anti-AGN bias within the BTS sample, in principle two methods are viable: Increase the number of AGN in the training sample, or veto against AGN in the sample that needs classification. To maximize performance, a combination of both approaches was chosen here, and each will be detailed in the next paragraphs.

6.4.3 AGN Rejection

Only those nuclear transients that survived the AGN rejection were kept for classification. The AGN rejection was performed as a two-step process, based on *WISE*-colors and matches to the MILLIQUAS catalog.

WISE color selection

The cuts applied were taken from [213]. The authors of that study evaluated roughly half a million SDSS galaxies with matching *WISE* infrared data, resulting in a region within the *WISE* color–color diagram that robustly picks out AGN. The AGN 'box' is defined as follows:

$$1.734 < (W2 - W3) < 3.916 \tag{6.6a}$$

$$(W1 - W2) > 0.0771(W2 - W3) + 0.319 \tag{6.6b}$$

$$(W1 - W2) > 0.261(W2 - W3) - 0.260, \tag{6.6c}$$

with $W1$, $W2$ and $W3$ being the *WISE* Vega magnitudes in the respective bands. All objects lying within this box (dot-dot-dashed contours in Fig. 6.10) were rejected as likely AGN.

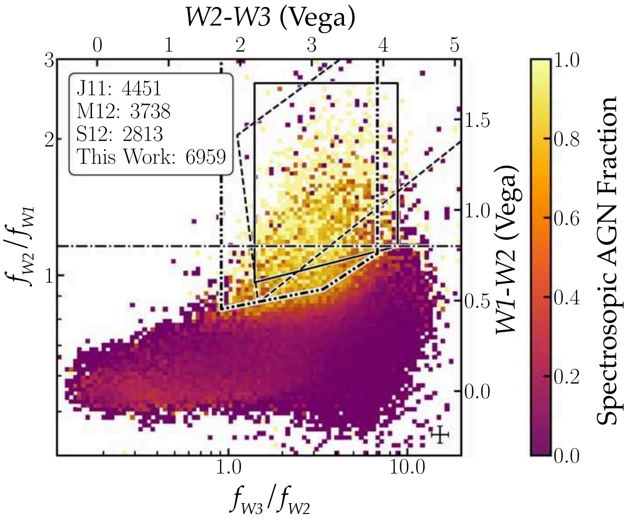

Figure 6.10: AGN Selection based on infrared *WISE* colors. It shows the $W2-W3$ vs. $W1-W2$ color; the fraction of spectroscopically classified AGN is encoded in datapoint color. The dot-dot-dashed box is the one used to flag likely AGN. Adapted from [213].

MILLIQUAS matching

Additionally to the *WISE* color selection, the MILLIQUAS [210] catalog[6] was also used to identify likely AGN. This catalog was also used in the regular high-energy neutrino follow-up (see Section 4.5.2). It contains over 900,000 Type I AGN, roughly 50,000 Type II AGN and again as many quasar candidates.

I was interested in maximizing the purity of the sample in terms of AGN, i.e. rejecting all likely AGN. Therefore, all objects that had a match in the MILLIQUAS catalog were rejected, regardless of their respective likelihood of the match being an AGN.

6.4.4 Adding non-BTS Sources

To counter the fact that almost no AGN were present in the BTS sample, forced photometry for an additional 1772 AGN light curves was acquired. This was done by crossmatching sources from a recent AGN study including redshifts [326] with ZTF alerts in a cone with a radius of 3 arcsec. All matched AGN with redshifts available for creating noisified copies were subsequently added to the training sample.

[326]: Mechbal et al. (2023), *Machine learning applications for the study of AGN physical properties using photometric observations*

As one of the desired features of the trained classifier was its ability to correctly identify TDEs, it was necessary to increase the number of TDEs in the training sample. For reasons not clear to me, the BTS underperformed significantly in terms of TDE detection rate. To maximize the number of TDEs in the training sample, all published ZTF-detected TDEs were included, resulting in a total number of 66 TDEs available for training the classifier.

[6] `https://quasars.org/milliquas.htm`

6.4.5 Creating a less biased Training Sample

It was not fully possible to control the number of child light curves: There were cases of low-flux objects of which all children failed the signal-to-noise cut. In total, the following number of objects were processed during the training sample augmentation:

SNe Ia The BTS sample contained 3230 SNe Ia, from which 8694 noisified children were created, i.e. 2.7 children per original light curve on average. The total number of light curves (original ones plus children) was 11924.

Core-collapse SNe 1075 CCSNe were available in the BTS sample. With an average of 10.6 children per light curve this amounted to 11,420 children and 12,495 light curves in total.

AGN 1893 AGN were originally present. 7177 children were created (3.8 per parent), resulting in 9079 light curves.

TDE Starting with 66 TDE, 13,097 children were creating with a large number of children per light curve (198). This amounted to 13,163 light curves in total.

Stars As already stated above, stars were not noisified due to their redshift being practically 0. Therefore, only the 525 parent light curves were used, rendering stars the most underrepresented class within the training sample.

Overall, the final training sample contained 47,186 light curves, generated from 6789 initial light curves. On average, $\sim$ 6 child light curves were created from each parent.

6.5 Training the Classifier

In recent years `XGBoost`[7] [327], an optimized gradient-boosted decision tree algorithm (this will be explained below), has performed well when classifying structured data; for this reason, it was chosen here.

[7] `https://xgboost.ai`

[327]: Chen et al. (2016), *XGBoost*

6.5.1 Boosted Decision Trees

In a regular decision tree, a set of iterative binary decisions (hence 'tree') are employed to decide on the classification of an object. Ensemble methods build upon this concept by either working in parallel by creating multiple decision trees with random subsets of features (random forest) or by working in sequence. The latter method qualifies boosted decision trees, meaning that decision trees are generated and evaluated in sequence. The residual squared errors from the first tree are fed into the second tree, and so on. The final classification then is the (weighted) sum of all individual tree's classifications. This method can be generalized to differentiable loss functions in general. If one uses the negative gradient of the squared error loss function instead of the squared error itself, the method is called gradient boosting. `XGBoost` is one variant of such a gradient boosting algorithm.

6.5.2 Hyperparameter Search

As is usually the case with such algorithms, there is a number of hyperparameters which determine how `XGBoost` behaves. `max_depth` for example avoids overfitting by restricting the size of individual decision trees, the `learning_rate` scales down the results by trees other than the initial tree, and `n_estimators` controls the number of trees.

These hyperparameters can be tuned by running a grid search in the parameter space, each time training with a small subset of the full data to speed up the process. The parameters yielding the best result are then used when training the full model with the complete training set. In this work, the best values for 9 hyperparameters were searched by randomly drawing from a grid of parameter combinations for 5000 times. This value was chosen as a compromise between computational feasibility and better sampling of the parameter search space. The grid employed here, as well as the chosen hyperparameters are shown in the appendix in Table A.4.

6.5.3 Training Procedure

After obtaining a classified training sample (BTS light curves, see Section 6.4.1), AGN and TDE light curves were added (Section 6.4.4) and the sample was augmented and enlarged by creating redshifted and noisified copies (Section 6.4.2). After this, the full training sample was fit with SALT2 (Section 6.2.2) and a TDE model (Section 6.2.1), as well as analyzed with a Bayesian block algorithm (Section 6.3) to extract features.

The set of features used for training was the following:

Peak magnitude This is the peak observed magnitude (in any filter). The value was scaled accordingly for the noisified and redshifted child light curves:

$$\begin{aligned} m_{\text{child}} &= m_{\text{parent}} - \mu_{\text{parent}} + \mu_{\text{child}} \text{ and} && (6.7a)\\ \mu &= 5(\log_{10} D_L(z) + 1), && (6.7b) \end{aligned}$$

where m is the parent or child magnitude, μ is the distance modulus, and D_L is the luminosity distance for redshift z.

SALT2 Fit Results This included the SALT2 fit parameters x_0, x_1 and c, as well as the reduced χ^2, the non-reduced χ^2 and the degrees of freedom.

TDE Fit Results The best-fit TDE Model parameters were also included. These were the rise- and decay times, the peak temperature, the temperature change and the amplitude. The temperature change cutoff time was not included, as the model never picked up on this parameter. The results used also included the χ^2 and the degrees of freedom.

Overlapping blocks The number of overlapping regions, extracted with the Bayesian block analysis, was also included to allow identification of non-repeating flares (in contrast to stochastic AGN behavior).

***WISE* colors** To identify likely AGN, the $W1 - W2$ and $W2 - W3$ colors of the parent host were added.

`sgscore` The `sgscore` of the parent light curve, based on machine learning with PS1 photometry (see Section 3.4.8), was included as a proxy for host galaxy photometry and morphology.

Core distance The median core distance was also included. To account for the fact that the core distance of the redshifted child light curves changed with respect to their parents, this value was recalculated for all children. This was done according to

$$\theta_{\text{child}} = \theta_{\text{parent}} \frac{D_L(z_{\text{parent}})}{D_L(z_{\text{child}})} \frac{(1 + z_{\text{child}})^2}{(1 + z_{\text{parent}})^2}, \tag{6.8}$$

where θ is the angular core distance of the child or parent light curve, D_L is the luminosity distance and z is the respective redshift.

6.6 Evaluating the Model

A large fraction of the nuclear sample was unclassified—this is why a classifier was needed in the first place. Therefore, evaluating the classifier's performance needed to happen with a part of the BTS training set. This was done in the usual way: A certain fraction (here: 30 %) of the training set was kept aside, never to be seen by the classifier beforehand. Such a sample is called a test set.

One pitfall needed to be avoided: The model might learn features of a noisified test light curve by having already seen its parent or one of its siblings in the training process, thereby cheating. To deal with this, the parent light curve and all of its children were always kept together: If e.g. one child light curve was part of the training set, so were its parent and all its siblings.

6.6.1 Performance

We are now ready to see how the model performs. This is done by first investigating the feature importance, and then by evaluating the confusion matrices.

Feature Importance

Feature importance can be used to decide on the usability of certain features, i.e. if they actually contribute to the final classification in a meaningful way. In this study, it was calculated by using a 'gain' metric. In this metric, one calculates the average improvement in loss when adding that feature to a tree during training, splitting one branch into two, which could result in more accurate predictions than the original branch.

As one can see in Fig. 6.11, the feature importance looks well-behaved. The feature importance was distributed fairly equally, as there was not one dominating feature. A dominant feature would raise concerns, as it could be indicative of the classifier finding a loophole regarding the actual task. Also, there were no features present that did not influence the classification at all.

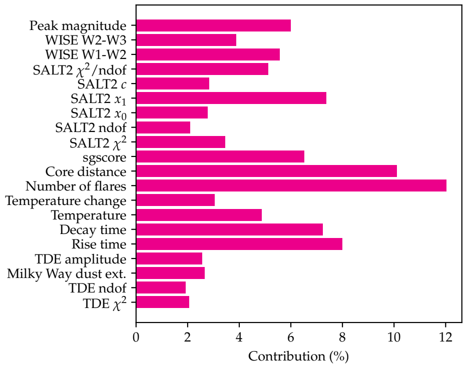

Figure 6.11: Feature importance for the classifier when trained with the augmented BTS sample. The relative contribution of each feature to the model is shown in percent, with all individual features adding up to 100 %.

The two most important features were the number of flares identified by the Bayesian block algorithm, as well as the core distance. This is not surprising, as the former is a good predictor for AGN behavior, and the latter might help in differentiating between AGN/TDEs and supernovae and stars.

Also important were the TDE fit rise- and decay time, as well as the SALT2 fit x_1 parameter, which encodes the stretch of the light curve. On equal footing were the peak magnitude and the `sgscore`. The *WISE* colors were also signficant, as the classifier probably captured the importance of *WISE* colors for AGN classification.

Interestingly, the temperature change did not seem to play a major role, and also the peak temperature was less important as one would have expected from previous studies (see e.g. [300]). This is indicative of the TDE model fit not capturing all physically present information. It has been shown that a $g - r$ color close to 0, as well as the absence of light curve color evolution are good predictors for TDEs. As the blackbody temperature and temperature change should in principle encode color and color evolution, it is somewhat surprising that these two features do not mirror the importance of color information.

Confusion Matrices

To see how the model performs with regard to the test set, one can either evaluate a version of the test set also containing child light curves created by the augmentation process, or evaluate only using light curves. There was no reason to favor one method over the other, so both evaluations were performed.

Test without Augmented Light Curves

The results from only allowing ‘real’ light curves (i.e. parent light curves) can be seen in Fig. 6.12.

As one can see, the absolute numbers displayed in brackets in Fig. 6.12 show that the test set—like the training set—is heavily biased towards SNe Ia, and only 15 TDEs are present.

The left matrix shows the prediction-normalized values: Each number is the percentage of the predicted classifications actually belonging to the predicted class. For example the purity[8] of the TDE selection is 90 %, meaning 9 out of 10 predicted TDEs are actually TDEs, while the

---

[8] Purity describes how many of the selected objects are in fact TDEs. This metric is also known as ‘precision’ in ML literature

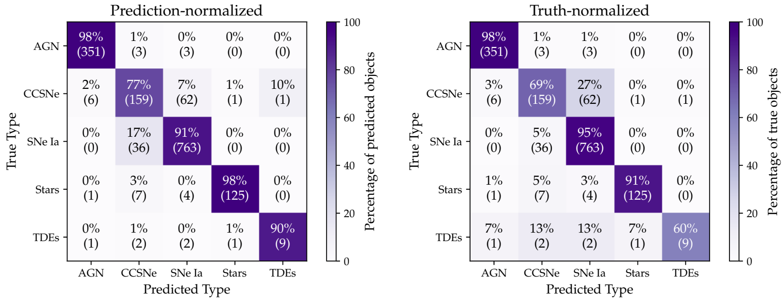

Figure 6.12: Confusion matrices of the test set predictions, excluding child light curves from the augmentation process. For each figure, the row corresponds to the true type, and the column to the predicted type of object. The matrix on the left shows prediction-normalized values, while the matrix on the right shows truth-normalized values. Absolute numbers are included below the percentages.

remaining 10 % are CCSNe in truth. In short, these 10 % constitute the false positive rate.

The truth-normalized matrix on the right shows the percentage with which objects that in truth belong to a class are correctly identified as members of that class. For example, the TDE completeness[9] is 60 %: 9 out of 15 TDEs were correctly identified as TDE, but 7 % were misclassified as AGN, 13 % were wrongly identified as SNe Ia, 13 % as CCSNe, and 7 % as stars. The sum of all misclassifications, in this case 40 %, corresponds to the false negative rate for TDE classification.

The performance looks decent, especially the false positive rate: Except for CCSNe, all predictions capture the truth in ≥90 % of all cases. Also, the TDE identification works fairly well: 60 % of all TDEs were identified by the classifier as such, and of the 10 objects predicted to be TDEs, only 1 was not a TDE. The validity of this result is of course limited by the small number of TDEs in the test sample. As the number of available ZTF light curves of confirmed TDEs is only in the double digits, there is no straightforward fix for this problem.

Test with Augmented Light Curves

The evaluation above was repeated, but this time including child light curves from the augmentation process. The results can be seen in Fig. 6.13.

The sample looks more balanced due to the inclusion of more noisified child light curves for underrepresented classes (except stars). As one would expect, the overall performance does increase when compared to the test set without child light curves: 97 % of all predicted TDEs were in fact TDEs (compared to 90 % without augmentation in the test set), and 79 % of all TDEs were identified as such (without augmentation: 60 %).

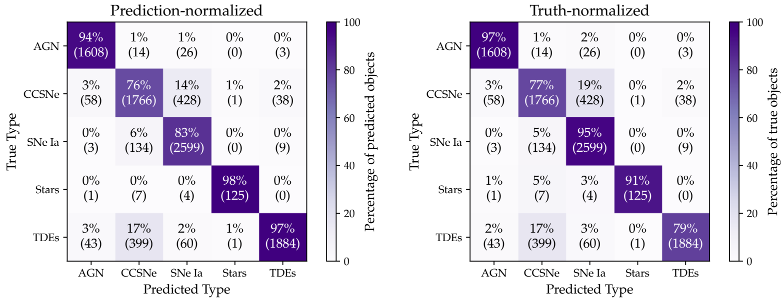

Figure 6.13: Confusion matrices of the test set predictions, including child light curves from the augmentation process. Left: prediction-normalized; right: truth-normalized.

---

[9] Completes details how many of the true class members are identified by a classifier; also known as 'recall' in the ML literature.

6.7 Finding Candidate TDEs

The traditional method of selecting TDE candidates consists of a sequence of two-dimensional cuts in the fit result parameter space (see e.g. [312]). One way of increasing the purity of the machine-learning selected TDEs is by defining such a set of 2D cuts by studying the BTS sample, and applying the very same cuts to the nuclear sample.

To do this, the classified part of the BTS sample enriched by additional TDEs was evaluated with exactly the same features as used in the training of the classifier model. The goal was to look for a set of cuts in the fit parameter space that would retain a high number of TDEs, while rejecting as many non-TDE as possible.

6.7.1 2D Cuts: Different Cut Stages

An overview over the sample after subsequent application of these two-dimensional cuts, as well as the TDE selection purity and completeness can be seen in Figure 6.15. These were the stages (the numbers in brackets correspond to the individual plots, with (1a) showing the sample without cuts):

AGN Veto (1b, 2a) Although the BTS is already biased against AGN, the BTS sample still contained objects which could be ruled out as likely AGN by crossmatching to MILLIQUAS (1b) and selecting based on *WISE* colors (2a), see Section 6.4.3.

Core distance (2b) As the BTS is not selecting nuclear transients per se, nuclearity needed to be enforced by requiring a maximum core distance of 0.4 arcsec.

`sgscore` cut (3a) Only transients with an `sgscore` < 0.3 were allowed at this cut stage.

SNIa diagonal cut (3b) The transient had to be to the right side of a diagonal cut in the rise decay time plane, described by $3.55 - 2.29\tau$ (with τ being the decay time, all values in $\log$ day).

Temperature cut (4a) The temperature was required to lie between $\log 3.9\,\mathrm{K}$ and $\log 4.4\,\mathrm{K}$. Furthermore, the daily temperature change ΔT was required to lie between $-100\,\mathrm{K\,day}^{-1} < \Delta T < 150\,\mathrm{K\,day}^{-1}$ (left).

Rise- and decay-time cut (4b) The rise and decay times were required to lie within $\log 0.8\,\mathrm{day} < \mathrm{risetime} < \log 2.05\,\mathrm{day}$ and $\log 1.1\,\mathrm{day} < \mathrm{decaytime} < \log 3.0\,\mathrm{day}$.

χ^2 cut (5a) Only those transients with a TDE fit reduced χ^2 outperforming their SALT2 fit reduced χ^2 were selected here to reject remaining likely SNe Ia. Also, the reduced TDE fit χ^2 was required to lie below 6.

Exactly one flare (5b) Finally, only transient with a single flare coincident in *g*- and *r*-band were allowed to reject likely stochastic AGN activity.

6.7.2 Visual Cuts: Evaluate the Cuts

Before all cuts, only requiring the TDE model fit to succeed, the sample consistsed of 4094 transients, of which 64 were TDEs. This corresponded to an initial purity of 1.6 % and a selection completeness of 100 % (all TDEs are of course retained before any cuts).

All Cuts Combined

The final selection consisted of 39 transients, of which 31 were TDE. This translates to a purity of 79.5 %, with a selection completeness of 48.4 %, i.e. roughly half of the TDE survive all cuts. The purity and completeness with subsequent cut stages are shown in Fig. 6.14.

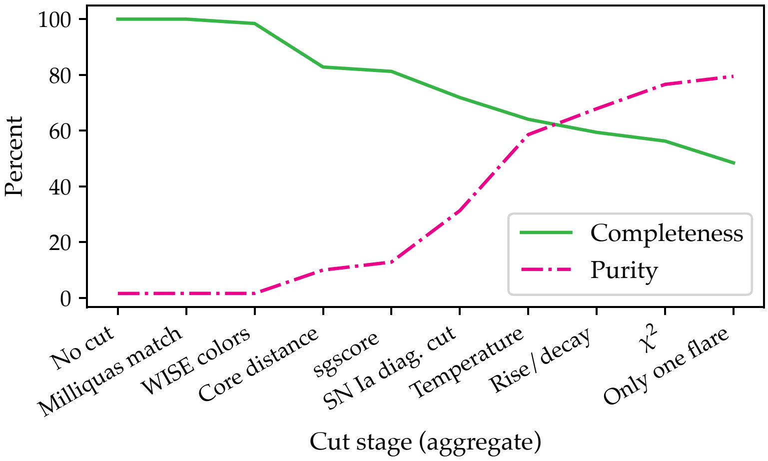

Figure 6.14: Completeness (green line) and purity (magenta, dash-dot) of the visual TDE selection. The cut stages are shown on the x-axis, with each new cut added on top of all previous cuts.

These numbers were somewhat promising, as objects lying within the final selection had a 4 in 5 chance of correctly being identified as TDE while only sacrificing half of the TDE population. Of course, there is always a trade-off between completeness and purity; i.e. a program designed to discover TDEs early on might need to accept a lower purity.

After establishing and verifying the 2D cuts, everything is ready to use them on the model predictions and investigate its performance further.

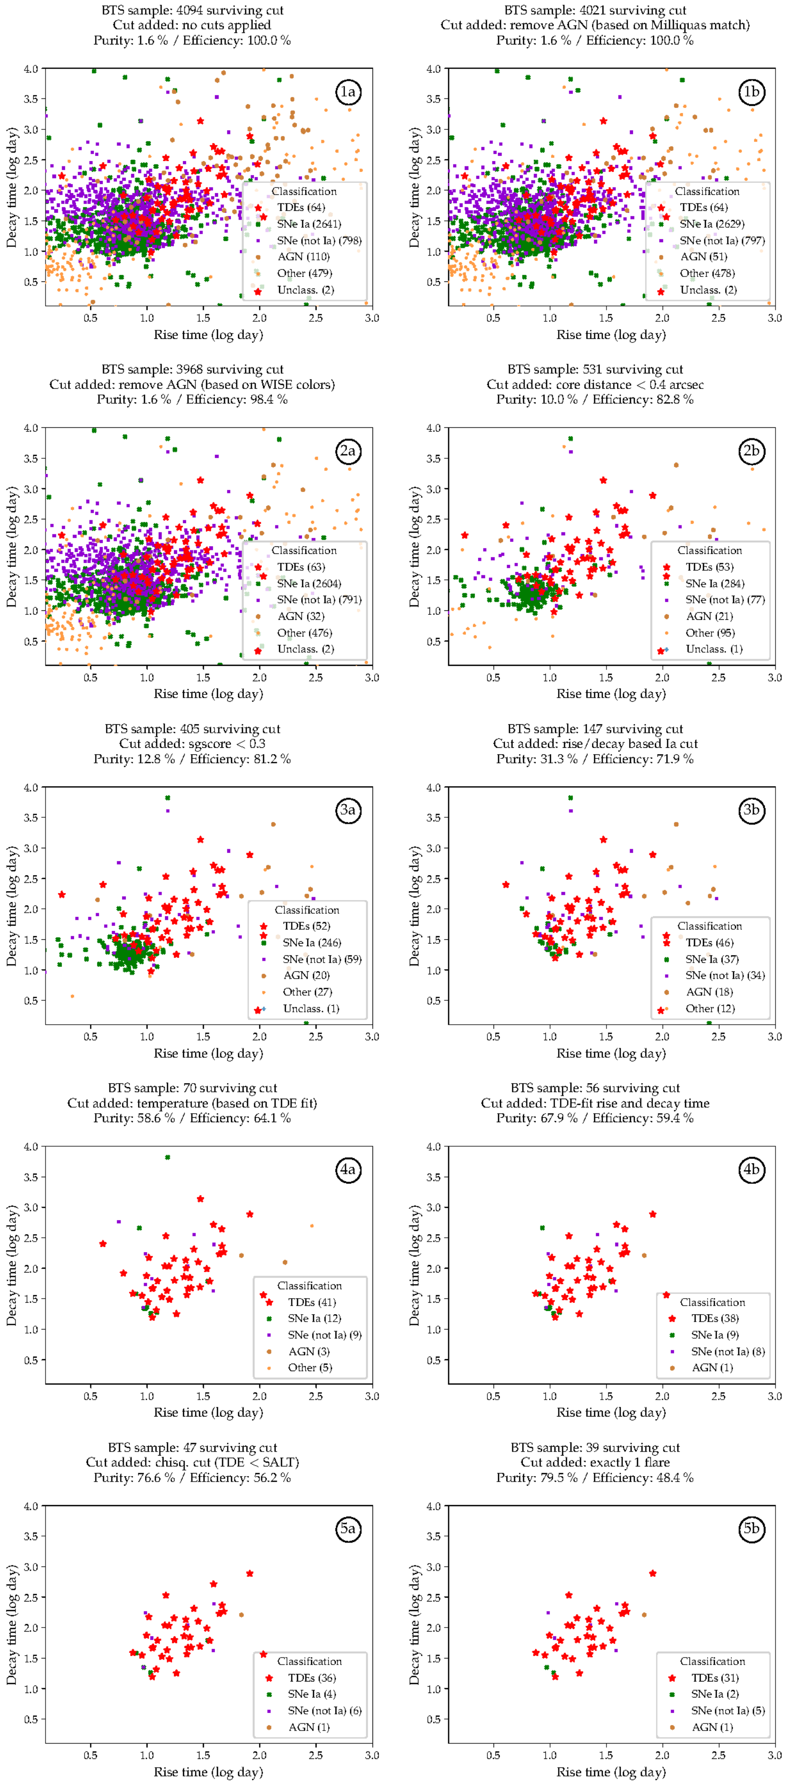

BTS sample: 4094 surviving cut
Cut added: no cuts applied
Purity: 1.6 % / Efficiency: 100.0 %

BTS sample: 4021 surviving cut
Cut added: remove AGN (based on Milliquas match)
Purity: 1.6 % / Efficiency: 100.0 %

BTS sample: 3968 surviving cut
Cut added: remove AGN (based on WISE colors)
Purity: 1.6 % / Efficiency: 98.4 %

BTS sample: 531 surviving cut
Cut added: core distance < 0.4 arcsec
Purity: 10.0 % / Efficiency: 82.8 %

BTS sample: 405 surviving cut
Cut added: sgscore < 0.3
Purity: 12.8 % / Efficiency: 81.2 %

BTS sample: 147 surviving cut
Cut added: rise/decay based Ia cut
Purity: 31.3 % / Efficiency: 71.9 %

BTS sample: 70 surviving cut
Cut added: temperature (based on TDE fit)
Purity: 58.6 % / Efficiency: 64.1 %

BTS sample: 56 surviving cut
Cut added: TDE-fit rise and decay time
Purity: 67.9 % / Efficiency: 59.4 %

BTS sample: 47 surviving cut
Cut added: chisq. cut (TDE < SALT)
Purity: 76.6 % / Efficiency: 56.2 %

BTS sample: 39 surviving cut
Cut added: exactly 1 flare
Purity: 79.5 % / Efficiency: 48.4 %

Figure 6.15: 2D-cut based selection of TDEs among the classified objects of the Bright Transient Survey (augmented with additional TDEs). These are displayed in the rise- and decay-time plane, as derived from the TDE fit (see Section 6.2.1). From top to bottom and left to right, more and more cuts are added.

All classes also used in the training of the classifier are shown. These were TDEs (red stars), SNe Ia (green diagonal crosses), CCSNe (purple squares), AGN (brown hexagons), other types (orange circles), and unclassified objects (blue crosses).

Top row: No cuts applied (1a) and the MILLIQUAS-based AGN veto (1b).

Second row: AGN cut based on *WISE* colors (2a) and cut based on the core distance (2b).

Third row: The `sgscore` is required to lie below 0.3 (3a) and SNe Ia are removed by a diagonal cut in the rise-/decay time plane (3b)

Fourth row: The temperature and temperature evolution were restricted (4a), as were the rise and decay times (4b).

Bottom row: Finally, it was required that the TDE model was a better fit than the SALT2 model by requiring the reduced χ^2 of the TDE fit to be smaller than the SALT2 fit's χ^2 (5a), and that the transient light curve had exactly 1 flare (5b).

6.8 A Photometric TDE sample

As we now obtained a trained classifier and a set of 2D cuts optimized for isolating TDE candidates, the next step was to classify the nuclear sample with the trained model, and subject it to the very same cuts.

6.8.1 Running the Classifier

The full classification of the nuclear sample is shown in Fig. 6.16. Here, the plots on top only show non-AGN hosts, while the bottom plots show likely AGN hosts. This differentiation was again achieved by crossmatching to the MILLIQUAS catalog and selecting based on *WISE* colors (see Section 6.4.3). The plots on the left show community classifications drawn from Fritz, the GROWTH Marshal and the TNS, while the right side plots display the `XGBoost` classifications.

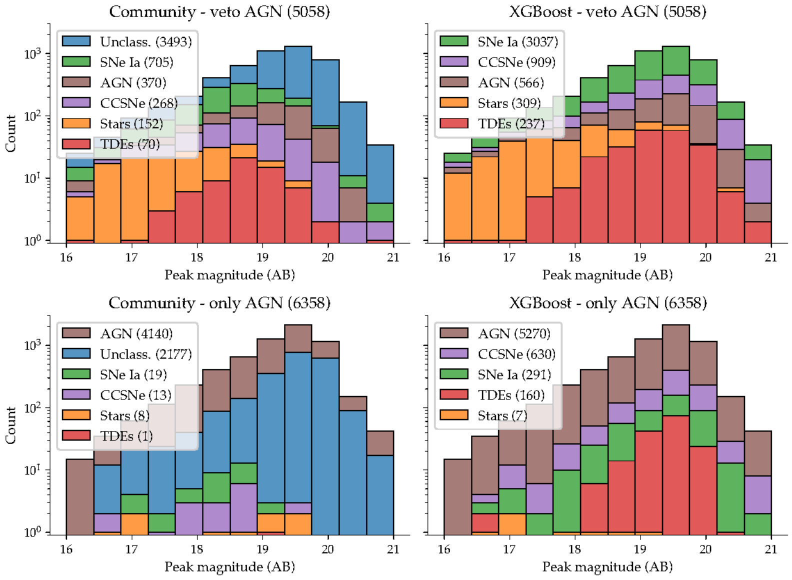

Figure 6.16: Classifications of the `XGBoost` model. This figure shows the nuclear sample, binned in magnitude steps of 0.5. There are no cuts applied. The classifications on the left side come from the community, i.e. Fritz, the GROWTH Marshal or TNS. The plots on the right side show the `XGBoost` classifications. The top row shows non-AGN hosts, while the bottom row shows likely AGN hosts.

Looking at the `XGBoost` classifications of the AGN part of the nuclear sample (bottom right plot), one can notice that—reassuringly—the majority (83 %) was indeed classified as AGN by the decision tree. As the non-AGN part of the nuclear sample is much more likely to contain bona fide transients, we will now focus on this selection (top right plot). At this stage, without further cuts on the sample, `XGBoost` classified 3037 objects as SNe Ia, 909 as CCSNe, 566 as AGN, 309 as stars, and 237 as TDEs.

To obtain a purer (but of course less complete) TDE sample, the 2D cuts established in Section 6.7 were applied. The classification of nuclear sample after application of these cuts can be seen in Fig. 6.17. The full list of all 79 transients from non-AGN hosts classified as TDE, including new and already known ones can be found in the appendix, see Table A.5. In total, 27 new TDE candidates could be added to the list of already known and spectroscopically classified TDEs. An overview of the candidates is shown in Table A.5 in the appendix.

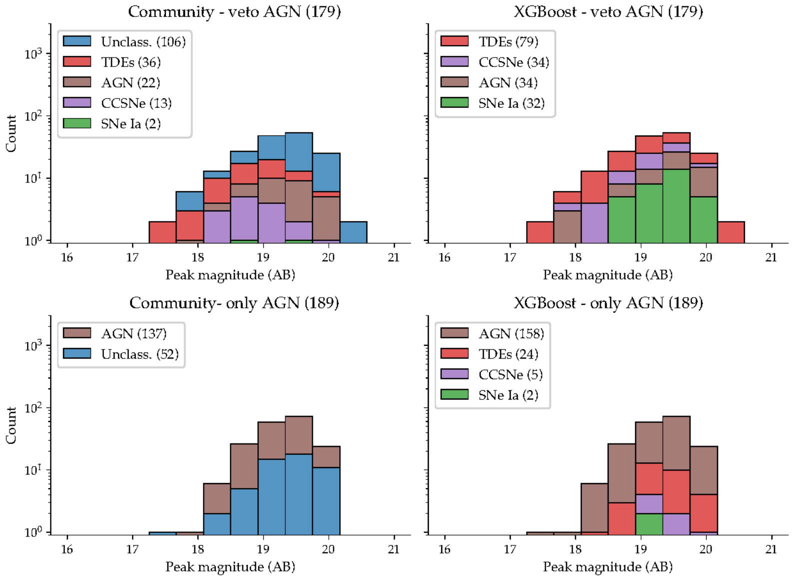

Figure 6.17: Classifications of the `XGBoost` model, this time with all the 2D cuts applied. Left side: community classifications, i.e. from Fritz, the GROWTH Marshal or TNS; right side: `XGBoost` classifications. Top row: non-AGN hosts; bottom row: likely AGN hosts.

The 24 objects classified as TDEs that were lying in AGN hosts (Fig. 6.17, bottom right plot) were also evaluated. Unfortunately, visual inspection of these events ruled out all but one (*AT2018ktr*) transient. Visually rejecting the transients was almost always possible based on light curve features alone. This highlights potentially lacking performance of the classifier. On the bright side, this selection shows that the AGN rejection seems to work, as significantly fewer events (24 vs. 79 for the non-AGN hosts) were initially classified as TDEs at that stage.

6.8.2 Inspection of Objects Classified as TDE

One pitfall in evaluating the classifier performance comes from the fact that 70 % of all TDEs were already seen by the classifier in the training stage—this was necessitated by the fact that there were not many to begin with. Therefore, a visual inspection of all objects classified as TDE that were not already known as TDEs during the training stage was performed.

This was done for the non-AGN part of the sample (see top right histogram in Fig. 6.16). Of the 237 objects classified as TDE, 29 were already seen by the classifier during the training stage and were therefore excluded, leaving 208 transients. These were then visually inspected, and rated according to the following scheme: 'Correct' meant that the transient looked like a plausible TDE, and no contrary evidence—like a classification based on a spectrum—existed. 'Possible' meant that either the light curve looked ambiguous, or the TDE fit result looked unsatisfying (for example, unphysically fast rise and fade timescales), but a clear misclassification could be excluded. Finally, 'wrong' was reserved for transients where the visual inspection rendered the TDE-classification unlikely, as e.g. stochastic AGN variability was present.

The exact rejection reasons for judging classifications as 'wrong' were: Misclassification based on an available spectrum, issues with the light curve quality, the existence of multiple recurring flares within the light

Spectrum says no TDE	Data quality	Multiple flares	1 flare, but variable baseline	No flare
16	4	24	26	23

Table 6.2: Rejection statistics for the visual inspection of the 93 transients analyzed and rated as 'wrong classification'.

curve, the existence of only one flare, but paired with a variable baseline suggestive of AGN variability, and finally not a clear flare, but only slightly elevated flux above a variable baseline. A breakdown of all 93 transients discarded due to these reasons is given in Table 6.2, while an example for each category in that Table is shown in Fig. 6.18.

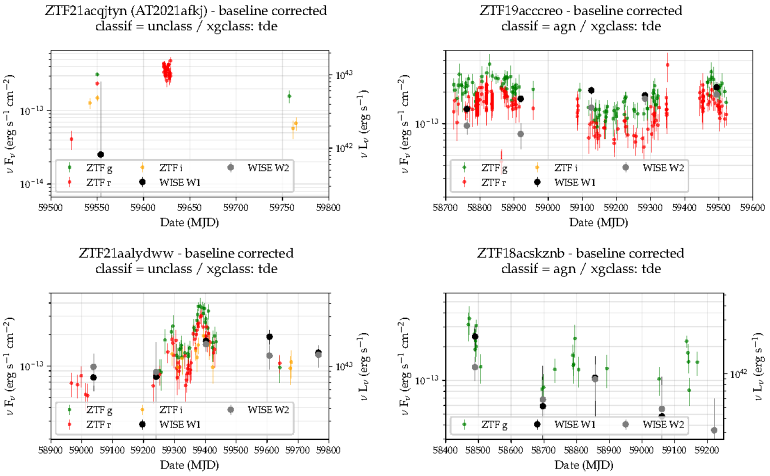

Figure 6.18: Examples for the different categories of visual AGN rejection. Top left: Data quality issues. Top right: There were multiple flares, which the Bayesian block algorithm failed to detect. Bottom left: Only one clear flare was detected, but there were signs of a variable baseline. Bottom right: No flare was discernible.

Of the 208 transients thus analyzed, 82 had a correct classification, for 33 a TDE nature was possible, and 93 were rated as wrong, i.e. misclassifications. When one generously rates both 'correct' and 'possible' as correct classification, a false positive rate of 44.7 % results. If one is more conservative and counts 'possible' as misclassification, a false positive rate of 60.5 % ensues.

In any case, these numbers mark a sharp drop compared to the 10 % false positive rate for the augmented BTS test sample (see Section 6.6.1). This gets alleviated a bit by the fact that TDEs which were present during the training phase were excluded—so the 'real' performance might as well be better. One can therefore safely assume that roughly half of the transients classified as TDE are in fact TDEs. All further calculations will be done with the assumption of a 50 % false positive rate.

On the bright side: Many of the clear misclassifications showed signs of AGN baseline variability. As can be seen in Table 6.2, 24 transients showed multiple flares, and 23 showed only slightly elevated flux above a variable baseline. It should be possible to develop a second classifier designed only to discern these features. If that were possible, the false positive rate could be reduced to roughly 22 %.

6.8.3 Conclusions for Photometric Classification

The performance of the classifier yielded somewhat mixed results. There was a number of problems identified that need to be addressed to improve the classification quality:

Data quality In some cases, inspection of transients classified as TDE showed a sparsely sampled light curve, e.g. *ZTF21aanubdr*. A human would recognize that the sparse sampling allowed only little conclusions to be drawn, while the decision tree had no adequate way of reacting to that fact. It would be beneficial to somehow infer from the data quality the range of conclusions that could possibly be drawn for a specific object.

Misclassifications There were examples where visual inspection allowed for easy and straightforward classification—especially in the case of stochastic AGN variability—but the classifier struggled to identify the AGN-nature of the transient. This highlights that there might still be room for improvement, i.e. further features that humans easily grasp and which could be exploited for automatic classification. The 'exactly 1 coincident flare' feature was apparently not sufficient for this task, and neither was a high reduced χ^2 of the TDE and SALT2 fits. It should be possible to add another classifier designed just to identify periodic features typical for such stochastic AGN behavior.

Effectiveness of Augmentation The TDE classification performance of the classifier when applied to the nuclear sample was investigated in Section 6.8.2. The results were worse than the confusion matrices of the test sample (see Section 6.6.1) suggested: At least 101 of 208 objects classified as TDE and not seen by the classifier during training are most likely something else, resulting in a much higher false positive rate (48.5 %) than the test sample suggested. This means that the augmentation by creating redshifted and noisified copies was not able to fully capture the transition from the BTS sample to the fainter nuclear sample.

6.9 Selection by Dust Echo

Another avenue explored for obtaining interesting and hitherto missed nuclear transients was a selection based on their infrared dust echo. For this, all transients with a `T2DustEchoEval` result (see Section 6.1.4) suggesting the existence of an infrared flare occurring simultaneous to or after the optical peak were visually inspected.

6.9.1 Inspection Tool

To aid in the visual inspection of light curves, I created a web-based frontend to interactively view the light curves of the nuclear sample, especially the ones showing an infrared flare qualifying for a dust echo. This tool, accessible under `https://ztfnuclear.simeonreusch.com`, was written in Python. It is based on the `Flask`[10] web framework, and operated behind an `nginx` web server also serving as reverse proxy, hosted on a virtual private server.

Fig. 6.19 shows the frontend. It was designed to display the transient light curve, results from the crossmatching, a redshift if one was found, as well as links to object pages on the GROWTH Marshal, Fritz and TNS. Lastly, it offered the functionality to leave comments (this is why a login

[10] `flask.palletsprojects.com`

system was created) and to rate the transient as ‘interesting’, ‘maybe interesting’, and ‘boring’.

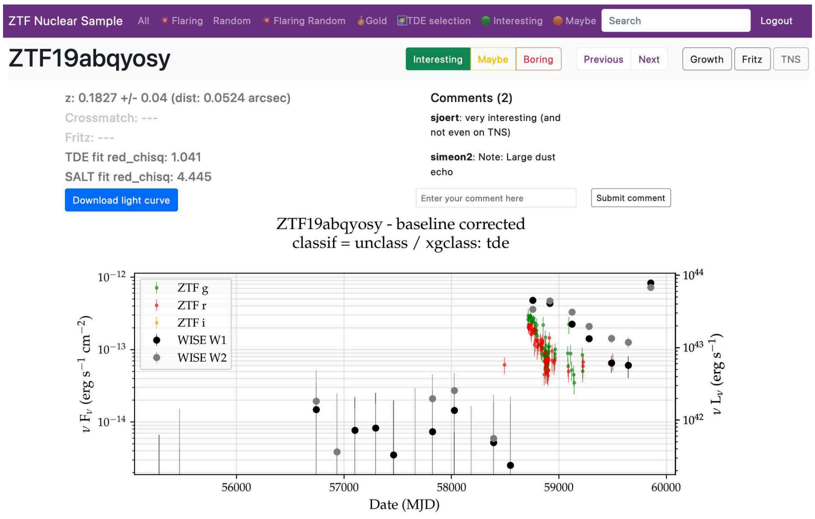

ZTF19abqyosy

z: 0.1827 +/- 0.04 (dist: 0.0524 arcsec)
Crossmatch: ---
Fritz: ---
TDE fit red_chisq: 1.041
SALT fit red_chisq: 4.445

Comments (2)

sjoert: very interesting (and not even on TNS)

simeon2: Note: Large dust echo

Figure 6.19: The frontend for inspecting and reviewing the nuclear sample. It shows an exemplary page for *ZTF19abqyosy*, including a light curve, a comment and a rating as ‘Interesting’. The fit results of the TDE and the SALT2 fit are displayed as well, plus potential redshift matches (none in this case).

6.9.2 Creating a Dust Echo Sample

To create a sample of dust echo transients that looked promising, multiple users reviewed all transients with a `T2DustEchoEval` score of 1, designed to select infrared flares compatible with stemming from a dust echo. A dust echo evaluation score of 1 selects those infrared flares which have a reliable baseline in both bands before the flare, and for which the flare is evolving reasonably fast (rise time < 1000 days, fade time < 5000 days).

The review was performed using the web frontend, and each user rated the transients as either ‘interesting’, ‘maybe interesting’ or ‘boring’. The final selection comprised transients that were rated as ‘interesting’ by at least two users. An overview over the 37 transients that met this criterion can be found in the appendix in Table A.6.

6.9.3 New TDE Candidates

Among this final selection were 16 candidate TDEs that showed a strong dust echo, and which have not been published so far. Note that two of these events are already contained in the `XGBoost` TDE candidates presented in Section 6.8.1. The new dust-echo accompanied TDE candidates are highlighted in bold in Table A.6 in the appendix.

Two exemplary light curves are shown in Fig. 6.20, one bona-fide candidate TDE (*ZTF18abtnfnq*), and one more ambiguous accretion flare similar to *AT2019fdr*: *AT2020oio*.

One can now compare both populations, the dust-echo accompanied transients and those without. The rise and decay times resulting from the TDE fit for both can be seen in Fig. 6.21

The two populations are by far not distinct, but a slight trend towards longer rise and decay times for the transients accompanied by a dust echo can be discerned. This suggests that the fraction of ambiguous

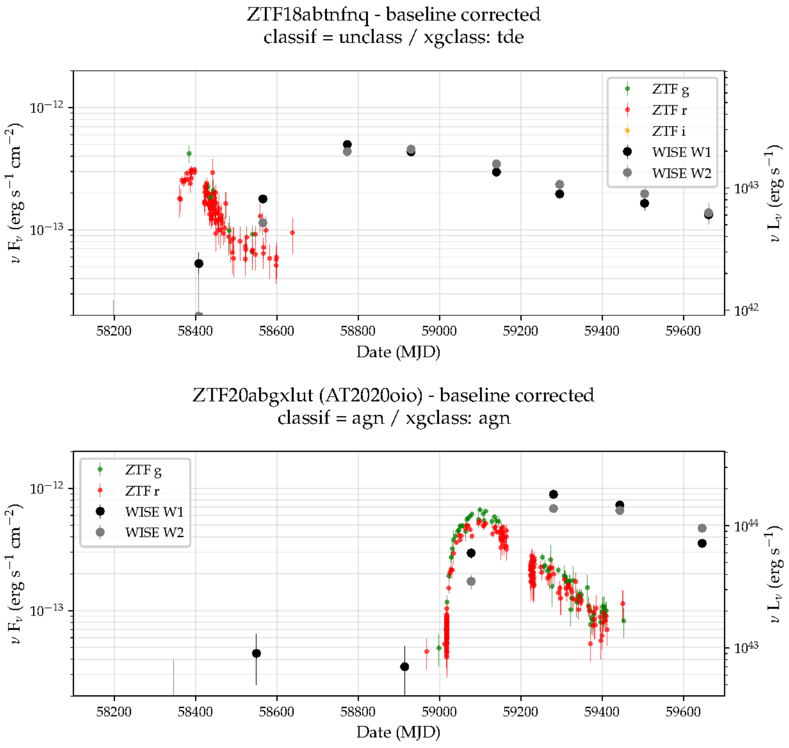
ZTF18abtnfnq - baseline corrected
classif = unclass / xgclass: tde

ZTF20abgxlut (AT2020oio) - baseline corrected
classif = agn / xgclass: agn

Figure 6.20: Two exemplary light curves from the dust echo section. Top: TDE candidate *ZTF18abtnfnq*. Bottom: Accretion flare candidate *AT2020oio*. Both have spectacular dust echoes. On the top, the delay between optical and infrared peak is roughly 400 days, while on the bottom it is about 180 days.

accretion flares is higher among the dust echo sample. Partially this might be a selection effect, as the 2D cuts introduced in Section 6.7 were custom-tailored to capture bona fide TDEs, while the dust-echo selection did not rely on these 2D cuts.

Comparison to Flux-Complete Sample

To check whether the number of new candidates from this work is compatible with expectations, I calculated a rough estimate on the expected number of TDEs when not limited by survey efficiency. To approximate the 'true' rate of TDEs accessible by ZTF, I assumed that in the magnitude bin between 17.5 and 18.0 all TDEs were in fact discovered. With this normalization, I calculated the flux-complete numbers for each magnitude bin, assuming the TDE population follows a distribution given by $N \propto f^{-\frac{3}{2}}$, where N is the expected number of TDEs and f is the TDE flux.

The resulting distribution is shown in Fig. 6.22, which displays the 66 known ZTF-detected TDEs used in this study (all of them fall into the time range of the nuclear sample) in blue. These are appended by the TDE candidates newly identified by this work (magenta), as well as the simulated flux-complete TDE sample (black dash-dotted line). This clearly shows that almost all candidates from this work could in fact be TDEs without violating the expected number of TDEs, especially in the two faintest magnitude bins.

This figure also highlights that even if all the TDE candidates truly were TDEs, the sample would still be far from complete at the faint end. For example, in the 19–19.5 magnitude bin, only half the number of expected TDEs is detected. This is not surprising, as the depth of ZTF survey observations does not exceed ∼ 20.5 mag, and only transients

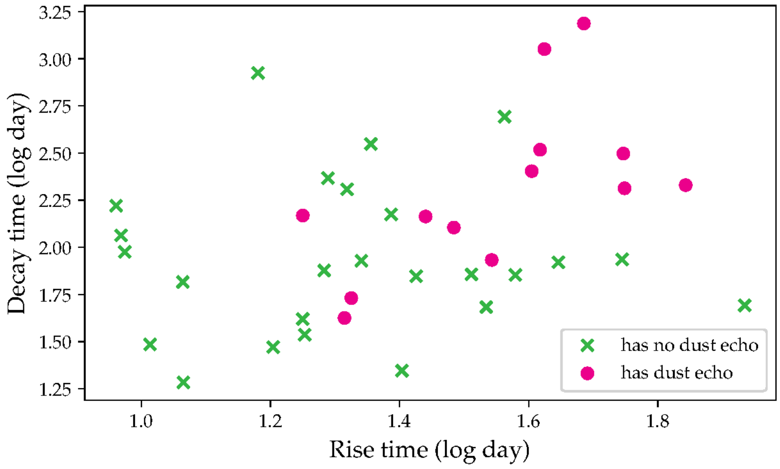

Figure 6.21: Rise and decay times for the transients accompanied by a dust echo versus those without one. Both axes are log day, and the dust-echo transients are shown as magenta dots, while the non dust-echo transients are shown as green crosses. On average, dust-echo accompanied transients evolve more slowly, suggesting a higher percentage of ambiguous nuclear transients among the dust-echo selected sample.

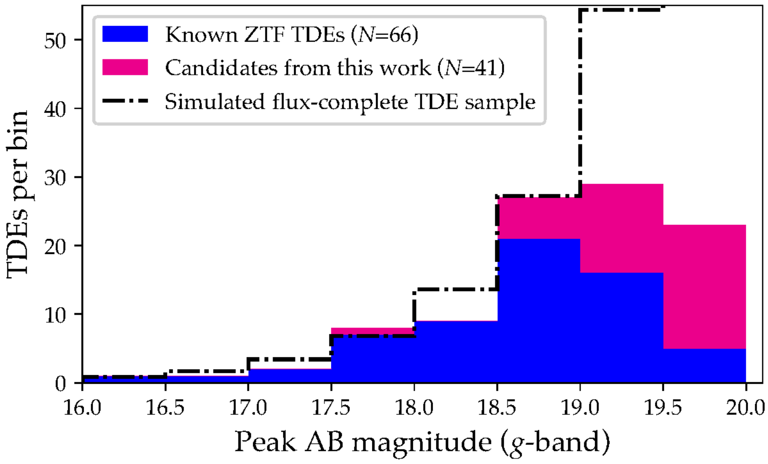

Figure 6.22: ZTF detected TDEs per g-band peak magnitude bin. The 66 known TDEs are shown in blue, while all the candidates from this work are shown in magenta. Furthermore, the flux-complete TDE sample (following a $f^{-\frac{3}{2}}$ distribution), normalized to the 17.5–18 magnitude bin, is shown as black dash-dotted line.

significantly brighter at peak will display a light curve that allows for reliable photometric identification.

6.10 Conclusion

Extracting previously unknown TDE candidates from the nuclear sample by using a trained algorithm aided by 2D cuts on the one side, and vetting promising dust-echo sources on the other side yielded in total 41 new candidates. This highlights that even dedicated programs are missing some candidates. If assuming that about half of the new candidates are in fact TDEs (in contrast to AGN activity and SNe, see Section 6.8.2 for a justification of that assumption), the luminosity function of TDEs might need to be adjusted. This is not exactly straightforward to do, as all the transients are over by now and the window of opportunity to classify them spectroscopically has closed. Nevertheless, the number of additional candidates and their magnitude distribution is in line with expectations when extrapolating from the ZTF-discovered TDEs in the 17.5–18.0 magnitude range.

One key takeaway from this study is that the contamination by AGN is the foremost problem in improving the quality of the photometric classifier. It seems that requiring exactly one coincident flare in two bands, and rejecting based on the reduced χ^2 of the two fits does not yield enough information, leaving a set of AGN misclassified as other transients. Therefore, it would be fruitful for future studies to add another classifier designed solely to reject stochastic AGN variability by identifying periodic behavior, e.g. based on Gaussian processes.

Outlook

It might be a promising avenue to obtain the missing host redshifts for the new transients identified in this study (of the 27 `XGBoost`-classified TDE candidates, 4 have a spectroscopic redshift, 12 a photometric redshift and 11 have none). There were 16 new dust-echo accompanied candidate TDEs and accretion flares, 14 of which are not already included in the 27 `XGBoost` candidates. Of these, 6 have a spectroscopic redshift, 4 have a photometric redshift and 4 have none. One could design a small program to acquire the 31 additional spectra needed to have reliable host redshifts for all the new transients.

Such a program will have to accommodate the fact that each additional TDE candidate will be intrinsically uncertain. Nevertheless, obtaining redshifts would allow to make the TDE nature of some candidates more plausible: The possible AGN nature of host galaxies would become apparent in their spectra, and in some cases alternative SN classifications might be ruled out based on the intrinsic luminosity of the respective event.

The ZTF TDE working group has spectroscopically identified 30 TDEs within a data-taking period of 2.6 years, translating to 11.5 discoveries per year. This study identified 27 additional candidates with the `XGBoost` classifier; when assuming at least 50 % of these are in fact TDEs, paired with a data-taking period of 4 years, this results in an additional 3.4 TDEs per year, an increase of 29 %. Adding 50 % of the candidate TDEs identified by their dust-echo results in an additional 1.75 TDEs per year.

One important remaining question concerns the differentiation between clear-cut TDEs and events which are somewhat ambiguous. The latter are in general more energetic, longer lived and occur in active galaxies. These are exemplified in Fig. 6.20, where *ZTF18abtnfnq* is a classical TDE candidate, and *ZTF20abgxlut* is more luminous, longer lived and situated in an AGN. One could hypothesize that both populations do in fact form a continuum; the classical TDEs on one end, and tidal disruptions happening in the environments of active galaxies on the other end, together constituting the class of 'accretion flares'.

Consequently, timely acquisition of spectroscopy is crucial in studying this possible continuum further; and the real-time application of the classifier developed in this work might be a promising avenue. Of course, such an application will have to deal with much younger and therefore shorter light curves and the fact that the transients are still evolving. Nevertheless, this can be done, as [328] have shown for the more general case of identifying transients worthy of early spectroscopic follow-up. Several of the features used by the classifier in this work will pose no problem, but the fit-based ones will need adjustments to allow the fitting of transients which are either young, or at least still evolving.

[328]: Miranda et al. (2022), *SNGuess: A method for the selection of young extragalactic transients*

Additionally, along the lines of [263], this study gave further evidence that the existence of a strong infrared flare is a good predictor for TDE-like behavior. We have successfully used an algorithm identifying periods of flaring activity which are then automatically checked if they could constitute a dust echo. Systematically using *WISE* data releases as soon as they are published to power an automatic pipeline searching for promising transients accompanied by dust echoes seems to be a promising avenue of research. This would also increase the fraction of spectroscopically classified transients belonging to the ambiguous 'accretion flare' category, possibly shedding more light on their nature.

It remains to be seen if dust-echo accompanied TDEs are in fact emitters of high-energy neutrinos. This work on *AT2019fdr* and the study additionally highlighting *AT2019aalc* [263] found evidence in favor of that hypothesis. On the other hand, when crossmatching the transients of that study with the full catalog of IceCube alerts—not only the alert neutrinos—[329] found no significant correlation. If the neutrino spectrum is hard enough, these two results do not contradict each other, though. Crossmatching the increased sample of dust-echo accompanied TDE candidates created in this work will shed more light on that question: The sample created here has a higher quality than what was used in [263], as the time range of observations is longer and there was a visual selection of TDE-like transients, rendering the sample more pure.

[329]: Necker (2023), *Search for High-Energy Neutrinos from TDE-like Flares with IceCube*

Conclusion and Outlook

The site of the Rubin Observatory in Chile. Image credit: Rubin Observatory, NSF and AURA.

The young field of neutrino astronomy is still in flux, with only a few established sources of high-energy neutrinos so far. This work gave an introduction to the field, followed by an overview of the predicted source classes, the instruments involved and the procedures and results of the high-energy neutrino follow-up program. The latter has now been running for over four years with significant contribution by the author, yielding several sources coincident with high-energy neutrinos. One of these, candidate Tidal Disruption Event *AT2019fdr* has been analyzed in detail in this work, including the compilation of an extensive set of multi-wavelength data from various instruments, modeling of its light curve, a study of its dust echo properties and a discussion of the chance coincidence of such an association.

This analysis has been supplemented by the ZTF Nuclear Sample, created by the author of this work. It comprises a unique dataset of nuclear transients observed with high-cadence in optical wavelengths over a period of several years. The sample has been tentatively classified with machine learning methods, utilizing an augmented version of the Bright Transient Survey as training sample. This classification, aided by an iterative selection of the 'TDE region', yielded 27 new candidate TDEs. Furthermore, a selection based on infrared dust echoes resulted in 16 previously unpublished TDE and accretion flare candidates (2 of which are also selected by the classifier). In total, 41 new TDE candidates were identified in this work.

Black Hole Masses for the Nuclear Sample

There are various further studies that could be done with the nuclear sample. The nature of accretion flares like *AT2019fdr* or e.g. *PS1-10adi* is still poorly understood. The TDE interpretation of *AT2019fdr*—though probably the best one—is still only one among two other (a SLSN Type II or an especially bright 'regular' AGN flare). A more extensive study of these objects is needed to provide a definite answer. A potential avenue of research is to obtain the missing host-galaxy spectra for all the accretion flare candidates found in the nuclear sample to estimate their corresponding black hole masses. One could use these to investigate if the host galaxy black hole masses of these events consistently fall below the Hills mass; if they do not, they cannot be TDEs. Also, as already mentioned in Section 6.10, spectra might help in updating the TDE luminosity function incorporating the new-found transients.

Real-time application of the Classifier

Furthermore, a real-time application of a modified version of the classifier would help in obtaining early spectroscopy, maybe additionally aided by the identification of potential dust echoes. The latter can be a powerful tracer of accretion flares, as was shown in this work. Such a program will also contribute to understanding the potential continuum of accretion flares ranging from classical TDEs to the somewhat more ambiguous events like *AT2019fdr*.

So far, we do not know for sure if Tidal Disruption Events or similar violent accretion events do in fact produce high-energy neutrinos. This work and the associations of a third event in [263] do stipulate that, but they hinge on the correctness of the error regions published by IceCube. Ultimately—due to the low-number statistics of the problem—only time will tell.

New Neutrino Detectors

Three important new instruments will allow shedding more light on the origin of high-energy neutrinos. Firstly, IceCube-Gen2 [330] is underway, with the goal of being fully operational in the middle of the 2030s. This extended detector will most likely improve upon the localization accuracy of IceCube, rendering associations with sources more secure. Additionally, the eightfold increase in volume when compared to IceCube will yield more events for which an association can be established.

[330]: Aartsen et al. (2021), *IceCube-Gen2: The Window to the Extreme Universe*

Secondly, the Cubic Kilometre Neutrino Telescope (KM3NeT) [331] is currently under construction in the Mediterranean. It is projected to have an angular resolution of $<0.2^\circ$ for neutrino energies above 10 TeV [332]. This is roughly half an order of magnitude better than the current IceCube accuracy, and will also aid in securing associations.

[331]: Adrián-Martínez et al. (2016), *Letter of intent for KM3NeT 2.0*

[332]: Aiello et al. (2019), *Sensitivity of the KM3NeT/ARCA neutrino telescope to point-like neutrino sources*

Rubin Observatory

Lastly—and much more imminent—first light of the Vera C. Rubin Observatory [175] is expected for November 2024[11]. With its 8.4 m primary mirror it is by far the largest survey telescope ever built, and will generate an order of magnitude more transient alerts when compared to ZTF. It will have an average sensitivity of 24.5 mag in the *r*-band compared to 20.6 for ZTF, which roughly translates to a 40-fold increase in sensitivity. This will result in a significantly higher number of candidate counterparts, with the majority being much fainter than ZTF-detected transients. The sheer volume and faintness of Rubin transients will render it impossible to obtain spectroscopy for the majority of the objects discovered by it.

[11] `https://www.lsst.org/about/project-status`

Therefore, making use of the Rubin Observatory in IceCube follow-up campaigns will necessitate further improvements in photometric typing. The work done on the nuclear sample might provide a stepping stone for such an undertaking. There are two potential obstacles, though.

Firstly, the lower cadence of the Rubin Observatory will lead to a much sparser sampling of light curves when compared to ZTF; this renders photometric typing even harder. In this vein, I have contributed to the HU/DESY participation in the Extended LSST Astronomical Time-series Classification Challenge (ELAsTiCC) [333]. This ongoing study for Rubin Observatory consists of photometrically typing simulated Rubin alerts, and the results so far give rise to optimism.

[333]: Narayan et al. (2023), *The Extended LSST Astronomical Time-series Classification Challenge (ELAsTiCC)*

Secondly, Rubin is located in the southern hemisphere, where IceCube's sensitivity is significantly reduced. It thus seems prudent to additionally keep on relying on smaller sky survey telescopes located in the northern hemisphere after the probable end of ZTF in 2024.

Overcoming these two obstacles is possible, and making use of forthcoming instruments will help in answering some of the open questions. Whatever those answers will be, one thing is certain: The young field of neutrino astronomy does have an auspicious future.

A Appendix

Table A.1: Summary of the 34 neutrino alerts followed up by ZTF until March 2023. '90 % area' indicates the rectangular localization uncertainty region as reported by IceCube. 'ZTF obs' indicates the area observed at least twice by ZTF, within the reported 90 % localization (accounting for chip gaps). 'Signalness' estimates the probability that the neutrino is of astrophysical origin, rather than caused by atmospheric background (see Section 2.4.2). The total followed-up area (corrected for chip gaps) is 205.02 deg^2.

Event	**R.A. (J2000) (deg)**	**Dec (J2000) (deg)**	**90% area (deg^2)**	**ZTF obs (deg^2)**	**Signal-ness**	**Reference**
IC190503A	$120.28^{+0.57}_{-0.77}$	$6.35^{+0.76}_{-0.70}$	1.94	1.37	36 %	[334, 335]
IC190619A	$343.26^{+4.08}_{-2.63}$	$10.73^{+1.51}_{-2.61}$	27.21	21.57	55 %	[336, 337]
IC190730A	$225.79^{+1.28}_{-1.43}$	$10.47^{+1.14}_{-0.89}$	5.41	4.52	67 %	[338, 339]
IC190922B	$5.76^{+1.19}_{-1.37}$	$-1.57^{+0.93}_{-0.82}$	4.48	4.09	51 %	[340–342]
IC191001A	$314.08^{+6.56}_{-2.26}$	$12.94^{+1.50}_{-1.47}$	25.53	23.06	59 %	[241, 242, 343]
IC200107A	$148.18^{+2.20}_{-1.83}$	$35.46^{+1.10}_{-1.22}$	7.62	6.28	–	[344, 345]
IC200109A	$164.49^{+4.94}_{-4.19}$	$11.87^{+1.16}_{-1.36}$	22.52	22.36	77 %	[346, 347]
IC200117A	$116.24^{+0.71}_{-1.24}$	$29.14^{+0.90}_{-0.78}$	2.86	2.66	38 %	[348–350]
IC200512A	$295.18^{+1.72}_{-2.26}$	$15.79^{+1.26}_{-1.29}$	9.77	9.26	32 %	[351, 352]
IC200530A	$255.37^{+2.48}_{-2.56}$	$26.61^{+2.33}_{-3.28}$	25.38	22.05	59 %	[226–228, 231]
IC200620A	$162.11^{+0.64}_{-0.95}$	$11.95^{+0.63}_{-0.48}$	1.73	1.24	32 %	[353, 354]
IC200916A	$109.78^{+1.08}_{-1.44}$	$14.36^{+0.88}_{-0.85}$	4.22	3.61	32 %	[355–357]
IC200926A	$96.46^{+0.73}_{-0.55}$	$-4.33^{+0.61}_{-0.76}$	1.75	1.29	44 %	[358, 359]
IC200929A	$29.53^{+0.53}_{-0.53}$	$3.47^{+0.71}_{-0.35}$	1.12	0.87	47 %	[360, 361]
IC201007A	$265.17^{+0.52}_{-0.52}$	$5.34^{+0.32}_{-0.23}$	0.57	0.55	88 %	[362, 363]
IC201021A	$260.82^{+1.73}_{-1.68}$	$14.55^{+1.35}_{-0.74}$	6.89	6.30	30 %	[233, 234]
IC201130A	$30.54^{+1.13}_{-1.31}$	$-12.10^{+1.15}_{-1.13}$	5.44	4.51	15 %	[364, 365]
IC201209A	$6.86^{+1.02}_{-1.22}$	$-9.25^{+0.99}_{-1.14}$	4.71	3.20	19 %	[366, 367]
IC201222A	$206.37^{+0.90}_{-0.80}$	$13.44^{+0.55}_{-0.38}$	1.54	1.40	53 %	[368, 369]
IC210210A	$206.06^{+1.40}_{-0.95}$	$4.78^{+0.62}_{-0.56}$	2.76	2.05	65 %	[370, 371]
IC210510A	$268.42^{+1.47}_{-1.60}$	$3.81^{+0.68}_{-0.64}$	4.04	3.67	28 %	[372, 373]
IC210629A	$340.75^{+1.11}_{-2.23}$	$12.94^{+0.91}_{-0.93}$	5.99	4.59	35 %	[235, 237, 238]
IC210811A	$270.79^{+1.07}_{-1.08}$	$25.28^{+0.79}_{-0.84}$	3.17	2.66	66 %	[374, 375]
IC210922A	$60.73^{+0.96}_{-0.66}$	$-4.18^{+0.42}_{-0.55}$	1.57	1.16	93 %	[376, 377]
IC220405A	$134.47^{+1.71}_{-1.72}$	$-1.27^{+1.45}_{-1.02}$	8.47	7.52	32 %	[378, 379]
IC220405B	$320.62^{+1.37}_{-1.13}$	$29.06^{+0.94}_{-0.68}$	3.54	3.41	36 %	[379, 380]
IC220501A	$311.57^{+0.82}_{-1.07}$	$18.68^{+1.08}_{-0.92}$	3.58	2.47	40 %	[381, 382]
IC220513A	$224.03^{+1.36}_{-1.27}$	$-1.34^{+0.74}_{-0.81}$	4.08	1.93	56 %	[383, 384]
IC220624A	$224.12^{+2.23}_{-1.95}$	$41.31^{+1.56}_{-1.56}$	9.80	8.09	61 %	[385, 386]
IC220822A	$273.08^{+2.47}_{-2.50}$	$21.54^{+0.94}_{-1.18}$	9.81	9.63	38 %	[387, 388]
IC220907A	$224.81^{+2.07}_{-1.95}$	$44.7^{+0.94}_{-1.06}$	5.72	5.41	46 %	[239, 240, 389]
IC221216A	$6.86^{+1.08}_{-2.06}$	$10.43^{+1.54}_{-1.07}$	8.05	5.68	41 %	[390, 391]
IC221223A	$350.54^{+0.67}_{-0.67}$	$34.71^{+0.67}_{-0.67}$	1.48	0.97	79 %	[392, 393]
IC230112A	$24.35^{+1.43}_{-1.71}$	$0.90^{+0.63}_{-1.26}$	5.93	5.57	28 %	[394, 395]

Table A.2: VLA measurements of *AT2019fdr* and the VLASS archival limit (first row). From [259].

MJD	**Date**	**Band** (GHz)	**Flux density** (µJy)
58032	2017-11-25	3.00	≤ 320
59033	2020-07-03	3.00	226 ± 13
59033	2020-07-03	10.00	86 ± 11
59105	2020-09-13	1.52	390 ± 34
59105	2020-09-13	3.00	211 ± 10
59105	2020-09-13	6.00	118 ± 6
59105	2020-09-13	10.00	86 ± 5
59160	2020-11-07	1.62	373 ± 20
59160	2020-11-07	3.00	209 ± 9
59160	2020-11-07	6.00	91 ± 6
59160	2020-11-07	10.00	39 ± 5

Table A.3: Archival *AT2019fdr* host measurements from *GALEX*, SDSS, UKIRT (no errors were available on the isoMag value) and *WISE*, used to construct the synthetic host model (see 5.2.3). From [259].

Band	$\lambda_{\mathbf{eff}}$	**AB magnitude**
GALEX FUV		$22.32^{+0.07}_{-0.10}$
GALEX NUV	$21.52^{+0.06}_{-0.10}$	
SDSS *u*	$20.91^{+0.06}_{-0.07}$	
SDSS *g*	$19.97^{+0.04}_{-0.04}$	
SDSS *r*	$19.00^{+0.02}_{-0.03}$	
SDSS *i*	$18.64^{+0.02}_{-0.03}$	
SDSS *z*	$18.36^{+0.02}_{-0.03}$	
UKIRT *J*	18.18	
WISE W1	$17.83^{+0.05}_{-0.05}$	
WISE W2	$17.78^{+0.03}_{-0.05}$	

Table A.4: `XGBoost` hyperparameter space searched and values used in the photometric classification of the ZTF nuclear sample. The value chosen for the training of the full model is shown in bold.

Parameter	**Grid values**
`learning_rate`	[**0.1**, 0.01, 0.001]
`gamma`	[0.01, **0.1**, 0.3, 0.5, 1, 1.5, 2]
`max_depth`	[2, 4, **7**, 10]
`colsample_bytree`	[**0.3**, 0.6, 0.8, 1.0]
`subsample`	[0.2, 0.4, 0.5, **0.6**, 0.7]
`reg_alpha`	[0, **0.5**, 1]
`reg_lambda`	[**1**, 1.5, 2, 3, 4.5]
`min_child_weight`	[**1**, 3, 5, 7]
`n_estimators`	[100, 250, 500, **1000**]

Table A.5: Curated final list of TDE candidates as classified by `XGBoost`, appended with redshifts and community classifications from Fritz, the GROWTH Marshal and TNS. **The time range of all light curve plots is 3 years, with the flux ranging between** 2×10^{-14} and $3 \times 10^{-11}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$. The red, green and yellow dots are ZTF g-, r- and i-band, while the *WISE* $W1$- and $W2$-bands are shown in black and grey. The 27 objects which were previously unclassified and which are compatible with a TDE interpretation after visual inspection and manual cross checks are shown in bold (i.e. the new candidates).

Light Curve	Transient	z	z type	Community classification	IAU name	Peak mag. (g-band)	Notes
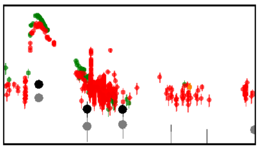	*ZTF17aaazdba*	0.022	spec.	TDE	*AT2019azh*	15.3	
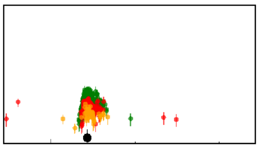	*ZTF18aakelin*	0.071	spec.	TDE	*AT2020ocn*	19.5	
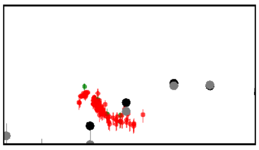	***ZTF18abtnfnq***	**0.131**	**phot.**			**20.3**	**Large dust echo**
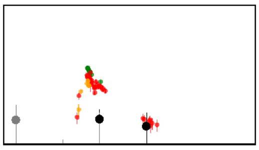	*ZTF18abxftqm*	0.108	phot.	TDE	*AT2018hco*	18.3	
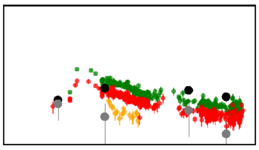	*ZTF18acaqdaa*	0.203	spec.	TDE	*AT2018iih*	18.3	
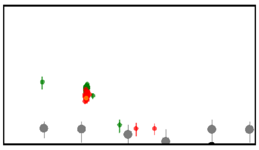	*ZTF18acbwomm*	0.137	phot.	AGN		19.2	
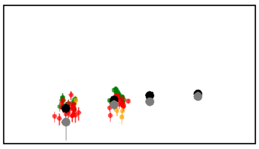	*ZTF18accttxu*				*AT2018lyq*	19.7	
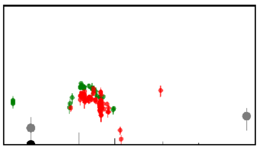	***ZTF18accvmgs***				***AT2018ibg***	**19.9**	
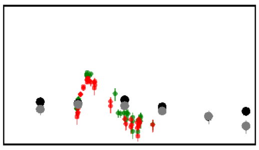	***ZTF18acetnxh***	**0.101**	**phot.**	**Blazar**	***AT2016fqa***	**18.6**	

Light Curve	Transient	z	z type	Community classification	IAU name	Peak mag. (g-band)	Notes
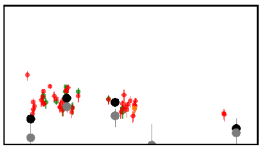	*ZTF18acqighv*			AGN	*AT2019afh*	19.3	
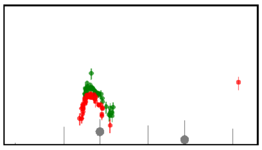	*ZTF18actaqdw*			TDE	*AT2018lni*	19.2	
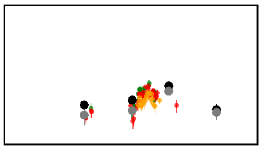	***ZTF19aaeukyu***				***AT2021aatf***	**19.2**	
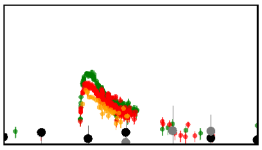	*ZTF19aaniqrr*			TDE	*AT2019cmw*	18.5	
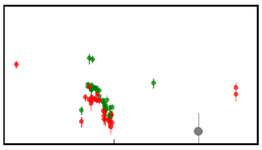	***ZTF19aanlttm***				***AT2019cyt***	**19.2**	
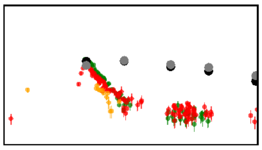	*ZTF19aapreis*			TDE	*AT2019dsg*	17.8	
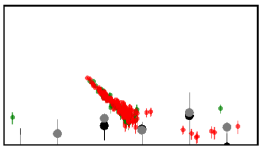	***ZTF19aaprhvf***				***AT2019ekr***	**19.1**	
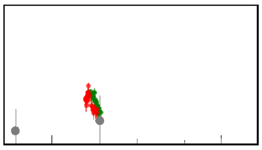	***ZTF19aayvxgy***					**19.7**	
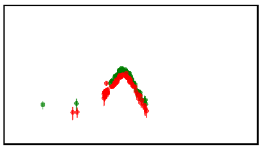	*ZTF19aaywayr*					18.2	Evolution too slow

Light Curve	Transient	z	z type	Community classification	IAU name	Peak mag. (*g*-band)	Notes
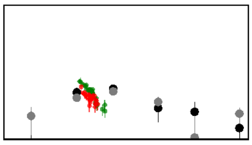	***ZTF19aazlrsy***			AGN	***AT2019hrg***	**19.2**	**Could also be AGN activity**
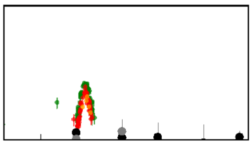	*ZTF19abhejal*	0.200	phot.	TDE	*AT2019mha*	19.3	
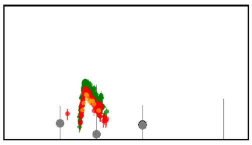	*ZTF19abhhjcc*	0.204	phot.	TDE	*AT2019meg*	19.2	
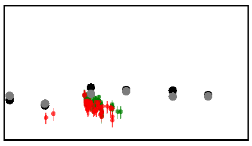	*ZTF19abicvxs*					19.9	
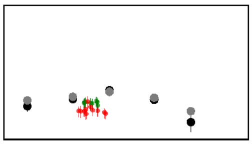	*ZTF19ablizhi*					19.8	
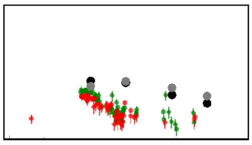	***ZTF19abqyosy***	**0.183**	**phot.**			**19.8**	**Large dust echo**
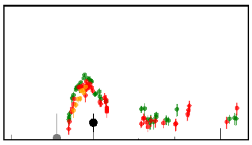	***ZTF19abrbskk***	**0.504**	**spec.**		***AT2019pcl***	**18.8**	**Very luminous**
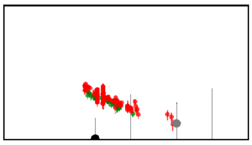	***ZTF19abzzuef***				***AT2019sez***	**19.5**	
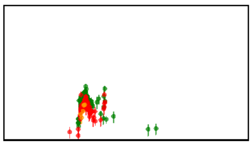	*ZTF19acanuza*	0.167	spec.	AGN		19.5	
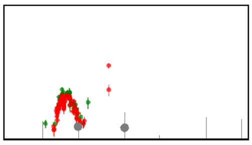	***ZTF19aclocml***					**19.7**	

Light Curve	Transient	z	z type	Community classification	IAU name	Peak mag. (*g*-band)	Notes
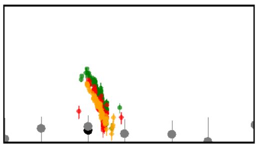	*ZTF19acspeuw*	0.134	phot.	TDE	*AT2019vcb*	18.4	
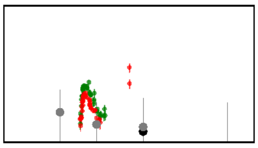	***ZTF19acssksf***			**AGN?**	***AT2019uyn***	**19.1**	
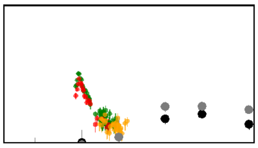	*ZTF20aabqihu*	0.070	phot.	TDE	*AT2020pj*	18.7	
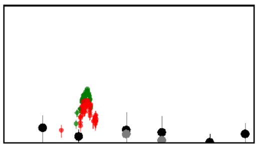	***ZTF20aavoumq***	**0.186**	**phot.**	**AGN?**		**19.5**	
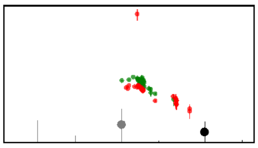	***ZTF20abbpxut***	**0.18**	**spec.**	**AGN?**	***AT2020kri***	**18.9**	
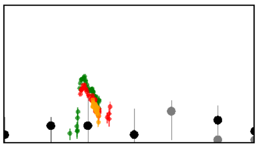	*ZTF20abefeab*	0.157	phot.	TDE	*AT2020mbq*	18.8	
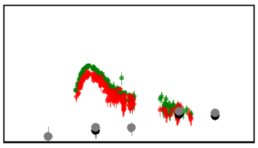	*ZTF20abfcszi*			TDE	*AT2020mot*	18.2	
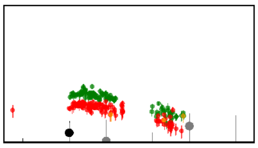	*ZTF20abjbkgq*	0.349	phot.	AGN	*AT2020qmx*	19.4	
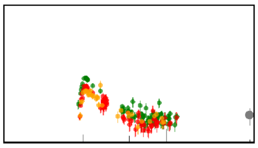	*ZTF20abjwvae*	0.200	phot.	TDE	*AT2020opy*	18.9	

Light Curve	Transient	z	z type	Community classification	IAU name	Peak mag. (*g*-band)	Notes
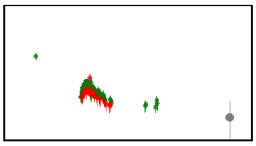	***ZTF20abkstjk***					**19.2**	
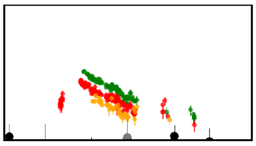	*ZTF20abnorit*	0.296	phot.	TDE	*AT2020ysg*	18.6	
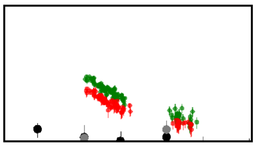	*ZTF20abowque*	0.333	phot.	TDE	*AT2020qhs*	18.9	
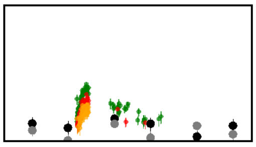	***ZTF20abqjebg***	**0.113**	**phot.**			**19.3**	
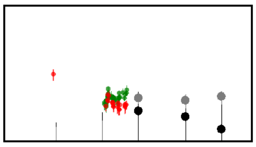	*ZTF20absxcuv*					18.6	
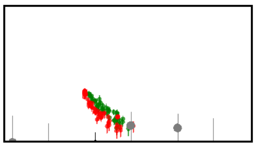	***ZTF20acfxnyr***	**0.127**	**phot.**			**19.8**	
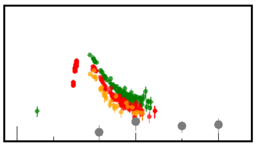	*ZTF20achpcvt*	0.033	spec.	TDE	*AT2020vwl*	17.6	
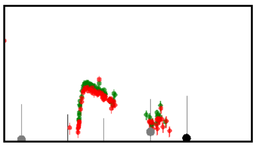	*ZTF20achupkw*	0.325	phot.	IIn	*SN2020vws*	19.2	
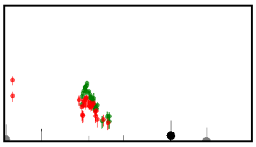	***ZTF20aclgfji***	**0.134**	**spec.**		***AT2020ygl***	**19.3**	
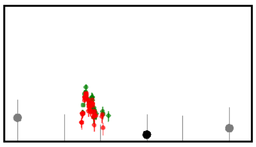	*ZTF20acllkua*	0.241	phot.			20.2	Evolving too fast

Light Curve	Transient	z	z type	Community classification	IAU name	Peak mag. (g-band)	Notes
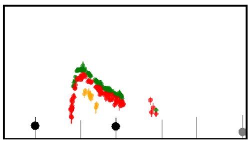	*ZTF20acnznms*	0.203	phot.	SLSN-II	*SN2020yue*	18.5	
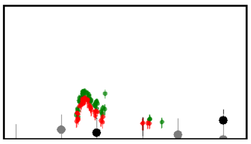	***ZTF20acpmkls***	**0.072**	**spec.**			**19.9**	
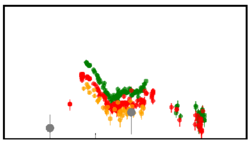	*ZTF20acwytxn*	0.319	phot.	TDE	*AT2020acka*	18.1	
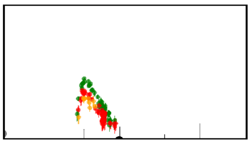	*ZTF20acyydkh*	0.127	phot.	TDE	*AT2021ack*	19.1	
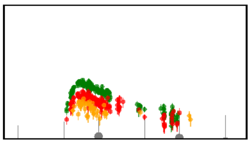	*ZTF21aabgjcz*			TDE	*AT2020aexc*	19.2	
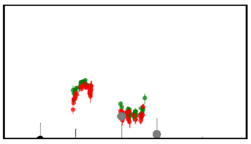	*ZTF21aabiipy*	0.278	phot.	TDE	*AT2021lo*	19.2	
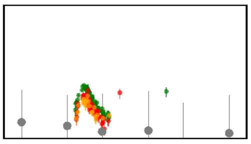	*ZTF21aaeoitd*	0.151	phot.	TDE	*AT2021jsg*	19.5	
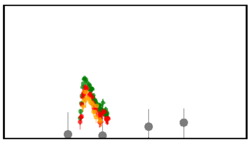	*ZTF21aakfqwq*	0.117	phot.	TDE	*AT2021crk*	19.1	
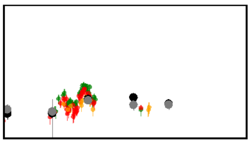	*ZTF21aalydww*	0.189	phot.			19.4	
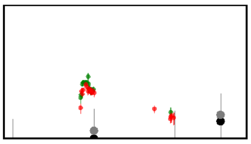	***ZTF21aanubdr***					**19.3**	**Little data**

Light Curve	Transient	z	z type	Community classification	IAU name	Peak mag. (*g*-band)	Notes
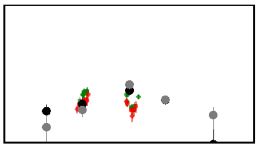	*ZTF21aaqgbln*					19.5	
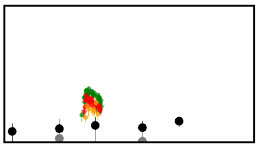	***ZTF21aaxtlvc***	**0.161**	**phot.**		***AT2021ldl***	**19.6**	
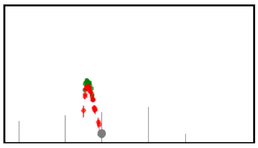	***ZTF21aaxtpty***	**0.126**	**phot.**		***AT2021kqp***	**19.1**	
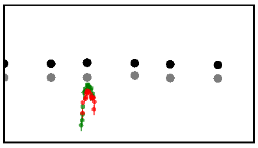	*ZTF21aaxxyha*	0.166	spec.		*AT2021kuk*	19.3	
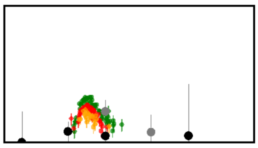	***ZTF21aazenvp***	**0.213**	**phot.**		***AT2021ovg***	**20.0**	
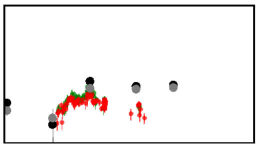	*ZTF21aazewul*				*AT2021uwm*	19.3	
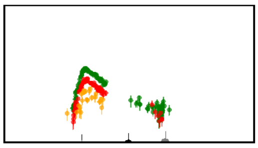	*ZTF21abcgnqn*	0.047	spec.	TDE	*AT2021nwa*	18.4	
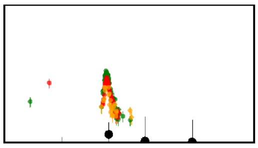	*ZTF21abjrysr*			TDE	*AT2021sdu*	11.6	
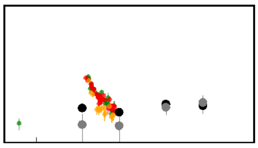	***ZTF21abkqvdo***	**0.096**	**phot.**		***AT2021swi***	**18.8**	
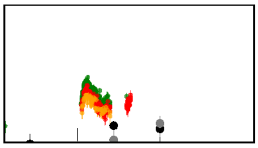	*ZTF21abqhkjd*	0.131	phot.	TDE	*AT2021uqv*	18.9	

Light Curve	Transient	z	z type	Community classification	IAU name	Peak mag. (g-band)	Notes
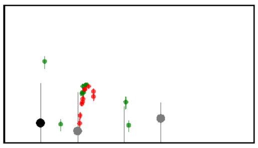	*ZTF21absltzp*	0.086	phot.		*AT2021vvw*	19.3	
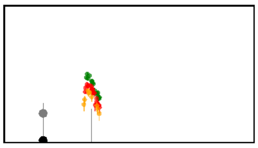	*ZTF21acafvhf*			TDE	*AT2021yte*	18.7	
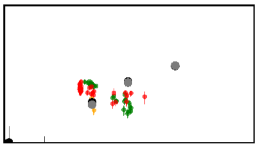	*ZTF21acdsrxc*	0.081	spec.			19.2	
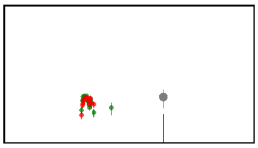	*ZTF21achnqdo*				*AT2021aczv*	19.9	
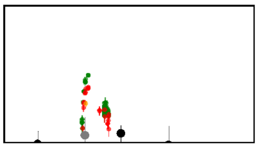	***ZTF21acojhgu***	**0.290**	**phot.**		***AT2021aees***	**18.7**	
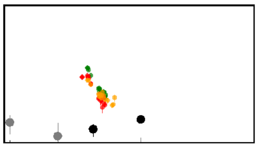	*ZTF22aaabovl*	0.137	phot.	TDE	*AT2022aee*	18.3	
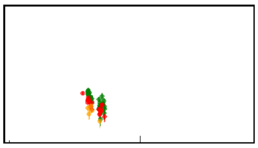	***ZTF22aaabqko***	**0.087**	**phot.**		***AT2022emf***	**19.7**	
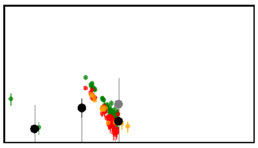	*ZTF22aaaedas*	0.186	phot.	TDE	*AT2022rz*	18.9	
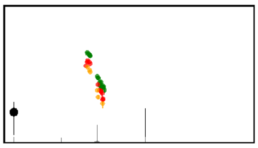	*ZTF22aaahtqz*	0.038	spec.	TDE	*AT2022bdw*	17.4	

Light Curve	Transient	z	z type	Community classification	IAU name	Peak mag. (g-band)	Notes
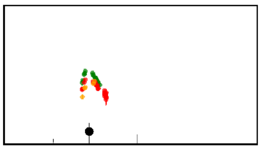	*ZTF22aabimec*	0.210	phot.	Ic-SLSN	*AT2022csn*	18.5	
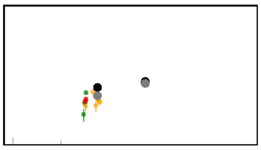	*ZTF22aacgcwv*	0.072	spec.	TDE	*AT2022dyt*	19.7	
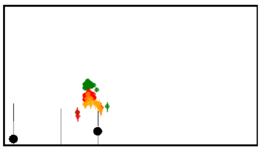	*ZTF22aaddwbo*	0.145	phot.	TDE	*AT2022gdw*	19.0	

Table A.6: Curated list of ZTF nuclear sample transients displaying strong *WISE*-detected infrared dust echoes. All transients listed here passed a visual selection by at least two persons. The 16 sources shown in bold are newly published in this work. Several sources were already published in [263], as a comparable metric to identify dust echoes was employed. **All light curve plots show a time range of 3 years, and the flux ranges between** 2×10^{-14} and 8×10^{-12} erg s^{-1} cm^{-2}. The red, green and yellow data points are ZTF measurements (*g*-, *r*- and *i*-band), while the black and grey datapoints are the *WISE W1* and *W2*-bands. Two candidates (*ZTF18abtnfnq* and *ZTF19abqyosy*) are already contained in the final `XGBoost` selection (Table A.5).

Light curve	Transient	z	z type	Community classification	IAU name	Peak mag. (*g*-band)	Notes
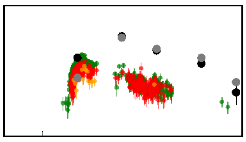	*ZTF18aajupnt*	0.037	spec.	CLAGN	*AT2018dyk*	18.58	[396, 397]
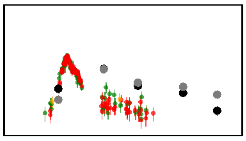	*ZTF18abjjkeo*	0.103	spec.	NLSy1 accr. flare	*AT2020hle*	18.65	[258]
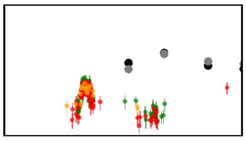	***ZTF18abnwufa***	**0.124**	**phot.**	**Unclass.**		**19.76**	
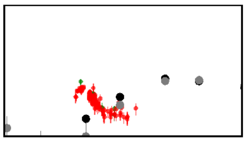	***ZTF18abtnfnq***	**0.131**	**phot.**	**Unclass.**		**20.26**	
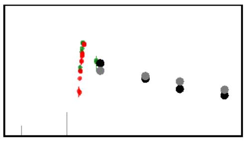	*ZTF18acgqweq*	0.167	spec.	AGN	*AT2018iql*	18.25	[263]
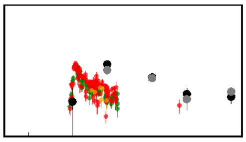	*ZTF18adbifqw*	0.121	phot.	Unclass.	*AT2018lhv*	19.14	[263]
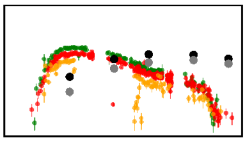	*ZTF19aailpwl*	0.374	spec.	various	*AT2019brs*	18.24	μTDE? [398]
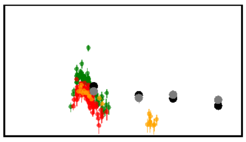	*ZTF19aakrnwh*	0.285	phot.	**Unclass.**	*AT2021aeuf*	18.26	[263]
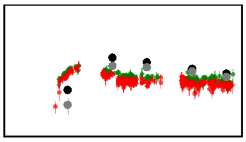	***ZTF19aamsgro***	**0.262**	**spec.**	**AGN**	***AT2019cyq***	**18.5**	**CLAGN?**

Light curve	Transient	z	z type	Community classification	IAU name	Peak mag. (*g*-band)	Notes
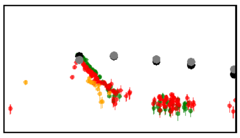	*ZTF19aapreis*	0.051	spec.	TDE	*AT2019dsg*	17.8	See Section 4.9.6
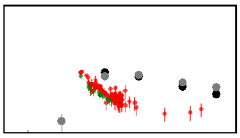	*ZTF19aasnaqa*	0.298	spec. by author	Unclass.	*AT2019gur*	19.74	[263]
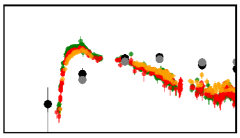	*ZTF19aatubsj*	0.266	spec.	accretion flare	*AT2019fdr*	18.12	See Chapter 5
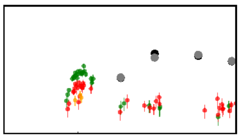	*ZTF19aaujlpo*	0.054	spec.	Unclass.	*AT2019idm*	19.21	[263], X-ray detected
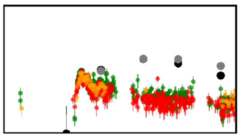	*ZTF19aavihif*	0.212	phot.	AGN	*AT2019hbh*	19.45	[263]
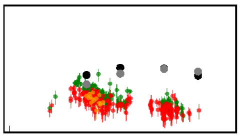	*ZTF19aavprgm*	0.183	spec. by author	Unclass.	*AT2019aami*	20.08	[263]
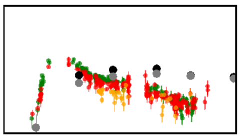	*ZTF19abhendr*			AGN	*AT2019mss*	18.88	[263]
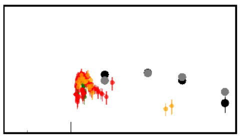	*ZTF19abiptrq*	0.1647	spec.	Unclass.	*AT2019nna*	19.36	[263]
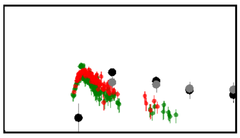	*ZTF19abkdlkl*	0.288	spec.	Unclass.	*AT2020afab*	19.76	[263]
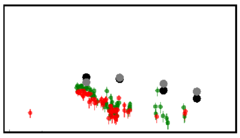	***ZTF19abqyosy***	**0.183**	**phot.**	**Unclass.**		**19.76**	

Light curve	Transient	z	z type	Community classification	IAU name	Peak mag. (*g*-band)	Notes
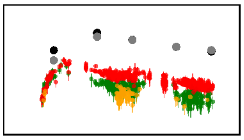	*ZTF19accdntg*	0.067	spec.	Unclass.	*AT2019thh*	19.24	[263]
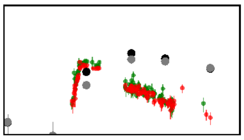	*ZTF20aaetsrw*	0.237	spec. by author	AGN	*AT2020atq*	18.93	[263]
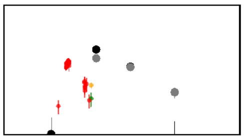	***ZTF20aaidhtp***	**0.148**	**spec.**	**Unclass.**		**19.84**	
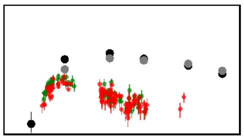	***ZTF20aaostow***	**0.184**	**spec.**	**AGN**		**19.9**	
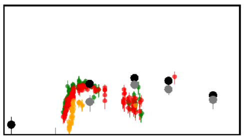	***ZTF20aaoxtxi***			**Unclass.**	***AT2020ima***	**19.92**	
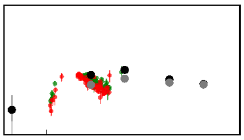	***ZTF20aapdqlk***	**0.487**	**spec. by author**	**Unclass.**		**19.62**	
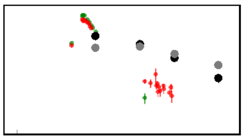	*ZTF20aasuiks*	0.159	spec.	CCSN	*SN2020edi*	16.59	[399], but large echo for a SN
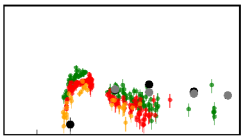	***ZTF20aauvhab***	**0.573**	**spec.**	**AGN**	***AT2020hip***	**19.27**	
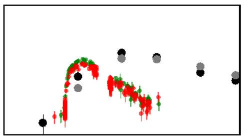	***ZTF20abgxlut***	**0.257**	**spec.**	**AGN**	***AT2020oio***	**18.85**	

Light curve	Transient	z	z type	Community classification	IAU name	Peak mag. (*g*-band)	Notes
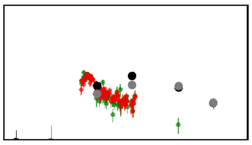	***ZTF20abhrmri***			**AGN**	***AT2020nnc***	**19.41**	
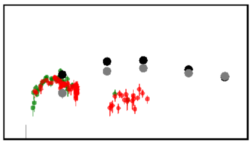	***ZTF20ablvwmh***	**0.337**	**phot.**	**AGN**	***AT2020xtj***	**19.36**	
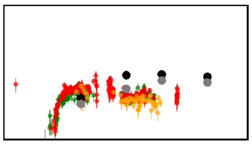	***ZTF20abxtsgg***			**Unclass.**		**19.94**	
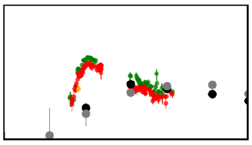	*ZTF20acbcfaa*	0.264	spec.	CCSN	*SN2020usa*	18.68	[399]
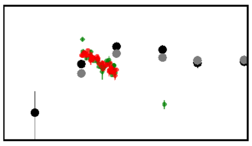	***ZTF20actrcji***			**Unclass.**	***AT2020abhp***	**17.77**	
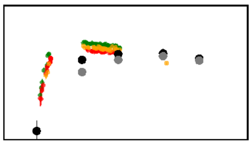	*ZTF20acvfraq*	0.26	spec.	AGN	*AT2020adpi*	17.9	[400]
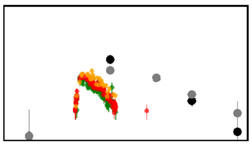	***ZTF20acyxxfo***	**0.268**	**phot.**	**Unclass.**	***AT2020aetz***	**19.6**	
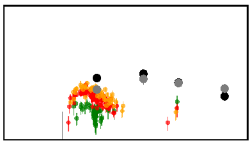	***ZTF21aaekxxf***	**0.26**	**phot.**	**Unclass.**	***AT2021esn***	**20.96**	
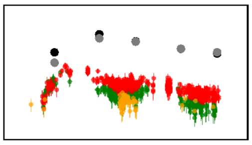	*ZTF21aawlhnk*	0.0506	spec.	Unclass.	*AT2019thh*	19.36	[263]

Bibliography

[1] W. Wall. *A History of Optical Telescopes in Astronomy*. Springer International Publishing, 2018 (cited on page 3).

[2] C. M. Linton. *From Eudoxus to Einstein: A History of Mathematical Astronomy. A History of Mathematical Astronomy*. Cambridge University Press, 2004, p. 528 (cited on page 3).

[3] J. E. Beckman. *Multimessenger Astronomy*. Springer International Publishing, 2021 (cited on pages 3, 4).

[4] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, USA, 2007, p. 720 (cited on pages 5–7).

[5] E. Fermi. "Versuch einer Theorie der β-Strahlen. I". In: *Zeitschrift für Physik* 88.3-4 (1934), pp. 161–177. DOI: `10.1007/bf01351864` (cited on pages 5, 6).

[6] F. Reines and C. L. Cowan. "The Neutrino". In: *Nature* 178.4531 (1956), pp. 446–449. DOI: `10.1038/178446a0` (cited on pages 5–7).

[7] E. Rutherford. "Uranium radiation and the electrical conduction produced by it". In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 47.284 (1899), pp. 109–163. DOI: `10.1080/14786449908621245` (cited on page 5).

[8] M. P. Villard. "Sur le rayonnement du radium". In: *Comptes rendus de l'Académie des sciences* 130 (1900), pp. 1178–1179 (cited on page 5).

[9] C. Jensen. *Controversy and Consensus: Nuclear Beta Decay 1911–1934*. Ed. by F. Aaserud et al. Birkhäuser Basel, 2000 (cited on page 6).

[10] S. Bilenky. "Neutrino. History of a unique particle". In: *The European Physical Journal H* 38.3 (2012), pp. 345–404. DOI: `10.1140/epjh/e2012-20068-9` (cited on pages 6, 7).

[11] C. L. Cowan et al. "Detection of the Free Neutrino: A Confirmation". In: *Science* 124.3212 (1956), pp. 103–104. DOI: `10.1126/science.124.3212.103` (cited on page 6).

[12] G. Ekspong. *Nobel Lectures in Physics 1981 – 1990*. World Scientific, 1993 (cited on page 7).

[13] G. Danby et al. "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos". In: *Physical Review Letters* 9.1 (1962), pp. 36–44. DOI: `10.1103/physrevlett.9.36` (cited on page 7).

[14] SNO Collaboration, Q. R. Ahmad et al. "Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory". In: *Physical Review Letters* 89.1 (2002), p. 011301. DOI: `10.1103/physrevlett.89.011301` (cited on pages 7, 8).

[15] R. Davis et al. "Search for Neutrinos from the Sun". In: *Physical Review Letters* 20.21 (1968), pp. 1205–1209. DOI: `10.1103/physrevlett.20.1205` (cited on page 8).

[16] J. N. Bahcall and R. Davis. "Solar Neutrinos: A Scientific Puzzle". In: *Science* 191.4224 (1976), pp. 264–267. DOI: `10.1126/science.191.4224.264` (cited on page 8).

[17] V. Gribov and B. Pontecorvo. "Neutrino Astronomy and Lepton Charge". In: *Physics Letters B* 28.7 (1969), pp. 493–496. DOI: `10.1016/0370-2693(69)90525-5` (cited on page 8).

[18] Y. Fukuda et al. "Evidence for Oscillation of Atmospheric Neutrinos". In: *Physical Review Letters* 81.8 (1998), pp. 1562–1567. DOI: `10.1103/physrevlett.81.1562` (cited on page 8).

[19] DONUT Collaboration, K. Kodama et al. "Observation of Tau Neutrino Interactions". In: *Physics Letters B* 504.3 (2001), pp. 218–224. DOI: `10.1016/s0370-2693(01)00307-0` (cited on page 8).

[20] P. Fisher et al. "Neutrino Mass and Oscillation". In: *Annual Review of Nuclear and Particle Science* 49.1 (1999), pp. 481–527. DOI: `10.1146/annurev.nucl.49.1.481` (cited on page 8).

[21] The KATRIN Collaboration, and M. Aker et al. "Direct neutrino-mass measurement with sub-electronvolt sensitivity". In: *Nature Physics* 18.2 (2022), pp. 160–166. DOI: `10.1038/s41567-021-01463-1` (cited on pages 8, 9).

[22] V. Barger. *The Pysics of Neutrinos*. Princeton University Press, 2012, p. 240 (cited on page 9).

[23] A. Cooper-Sarkar et al. “The high energy neutrino cross-section in the Standard Model and its uncertainty”. In: *Journal of High Energy Physics* 2011.8 (2011). DOI: `10.1007/jhep08(2011)042` (cited on page 9).

[24] The Borexino Collaboration, M. Agostini et al. “Comprehensive measurement of pp-chain solar neutrinos”. In: *Nature* 562.7728 (2018), pp. 505–510. DOI: `10.1038/s41586-018-0624-y` (cited on page 10).

[25] W. Kunkel and B. Madore. “Supernova 1987A in the Large Magellanic Cloud”. In: *International Astronomical Union Circular* 4516 (1987) (cited on page 10).

[26] M. Spurio. *Probes of Multimessenger Astrophysics*. Springer International Publishing, 2018 (cited on pages 10, 17, 22).

[27] I. Arcavi. “Hydrogen-Rich Core-Collapse Supernovae”. In: *Handbook of Supernovae*. Springer International Publishing, 2017, pp. 239–276. DOI: `10.1007/978-3-319-21846-5_39` (cited on page 10).

[28] V. P. Utrobin et al. “Supernova 1987A: 3D Mixing and Light Curves for Explosion Models Based on Binary-merger Progenitors”. In: *The Astrophysical Journal* 914.1 (2021), p. 4. DOI: `10.3847/1538-4357/abf4c5` (cited on page 10).

[29] K. Hirata et al. “Observation of a neutrino burst from the supernova SN1987A”. In: *Physical Review Letters* 58.14 (1987), pp. 1490–1493. DOI: `10.1103/physrevlett.58.1490` (cited on page 10).

[30] R. M. Bionta et al. “Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud”. In: *Physical Review Letters* 58.14 (1987), pp. 1494–1496. DOI: `10.1103/physrevlett.58.1494` (cited on page 10).

[31] F. Halzen and K. Scholberg. “Neutrinos from Core-Collapse Supernovae and Their Detection”. In: *Handbook of Supernovae*. Springer International Publishing, 2017, pp. 1655–1670. DOI: `10.1007/978-3-319-21846-5_8` (cited on page 11).

[32] C. Grupen. *Astroparticle Physics*. Springer, 2005, p. 441 (cited on page 11).

[33] I. Iben and A. V. Tutukov. “Supernovae of type I as end products of the evolution of binaries with components of moderate initial mass”. In: *The Astrophysical Journal Supplement Series* 54 (1984), p. 335. DOI: `10.1086/190932` (cited on page 11).

[34] D. Maoz et al. “Observational Clues to the Progenitors of Type Ia Supernovae”. In: *Annual Review of Astronomy & Astrophysics* 52.1 (2014), pp. 107–170. DOI: `10.1146/annurev-astro-082812-141031` (cited on page 11).

[35] M. M. Phillips. “The absolute magnitudes of Type IA supernovae”. In: *The Astrophysical Journal* 413 (1993), p. L105. DOI: `10.1086/186970` (cited on page 11).

[36] A. W. Alsabti and P. Murdin. “Supernovae and Supernova Remnants: The Big Picture in Low Resolution”. In: *Handbook of Supernovae*. Springer International Publishing, 2017, pp. 3–28. DOI: `10.1007/978-3-319-21846-5_1` (cited on pages 11, 12).

[37] N. Sravan et al. “Progenitors of Type IIb Supernovae. II. Observable Properties”. In: *The Astrophysical Journal* 903.1 (2020), p. 70. DOI: `10.3847/1538-4357/abb8d5` (cited on page 12).

[38] H.-T. Janka et al. “Core-collapse supernovae: Reflections and directions”. In: *Progress of Theoretical and Experimental Physics* 2012.1 (2012). DOI: `10.1093/ptep/pts067` (cited on page 12).

[39] C. Lunardini. “Diffuse Neutrino Flux from Supernovae”. In: *Handbook of Supernovae*. Springer International Publishing, 2017, pp. 1637–1653. DOI: `10.1007/978-3-319-21846-5_6` (cited on page 13).

[40] H.-T. Janka. “Neutrino-Driven Explosions”. In: *Handbook of Supernovae*. Springer International Publishing, 2017, pp. 1095–1150. DOI: `10.1007/978-3-319-21846-5_109` (cited on page 13).

[41] R. Steinmaurer. “Erinnerungen on V. F. Hess, Den Entdecker der Kosmischen Strahlung, und an Die ersten Jahre des Betriebes des Hafelekar-Labors”. In: *Early History of Cosmic Ray Studies*. Springer Netherlands, 1985, pp. 17–31. DOI: `10.1007/978-94-009-5434-2_3` (cited on page 13).

[42] Particle Data Group, and R L Workman et al. “Review of Particle Physics”. In: *Progress of Theoretical and Experimental Physics* 2022.8 (2022). DOI: `10.1093/ptep/ptac097` (cited on pages 13, 14, 18).

[43] A. M. Hillas. *Cosmic Rays: Recent Progress and some Current Questions*. 2006. DOI: `10.48550/arxiv.astro-ph/0607109` (cited on page 14).

[44] A. M. Taylor. “Cosmic rays beyond the knees”. In: *Nature* 531.7592 (2016), pp. 43–44. DOI: `10.1038/531043a` (cited on page 14).

[45] F. Halzen. “Lectures on High-Energy Neutrino Astronomy”. In: *AIP Conference Proceedings*. AIP, 2006. DOI: `10.1063/1.2160978` (cited on page 14).

[46] F. M. Rieger. “Active Galactic Nuclei as Potential Sources of Ultra-High Energy Cosmic Rays”. In: *Universe* 8.11 (2022), p. 607. DOI: `10.3390/universe8110607` (cited on page 14).

[47] A. M. Hillas. “The Origin of Ultra-High-Energy Cosmic Rays”. In: *Annual Review of Astronomy & Astrophysics* 22.1 (1984), pp. 425–444. DOI: `10.1146/annurev.aa.22.090184.002233` (cited on page 14).

[48] M. S. Longair. *High Energy Astrophysics*. Cambridge University Press, 2011 (cited on pages 15–17).

[49] M. G. Baring. *Diffusive Shock Acceleration: the Fermi Mechanism*. 1997. DOI: `10.48550/arxiv.astro-ph/9711177` (cited on page 15).

[50] D. Urošević et al. “Particle acceleration in interstellar shocks”. In: *Astrophysics and Space Science* 364.10 (2019). DOI: `10.1007/s10509-019-3669-y` (cited on page 15).

[51] A. R. Bell. “The Acceleration of Cosmic Rays in Shock Fronts”. In: *Monthly Notices of the Royal Astronomical Society* 182.2 (1978), pp. 147–156. DOI: `10.1093/mnras/182.2.147` (cited on page 17).

[52] D. Fiorillo and M. Bustamante. “Bump-hunting in the diffuse flux of high-energy cosmic neutrinos”. In: *Proceedings of 38th International Cosmic Ray Conference — PoS(ICRC2023)*. Sissa Medialab, 2023. DOI: `10.22323/1.444.0999` (cited on page 18).

[53] D. F. Fiorillo et al. “Unified thermal model for photohadronic neutrino production in astrophysical sources”. In: *Journal of Cosmology and Astroparticle Physics* 2021.07 (2021), p. 028. DOI: `10.1088/1475-7516/2021/07/028` (cited on page 18).

[54] E. O. Ofek et al. “Precursors prior to Type IIn Supernova Explosions are common: Precursor Rates, Properaties and Correlations”. In: *The Astrophysical Journal* 789.2 (2014), p. 104. DOI: `10.1088/0004-637x/789/2/104` (cited on page 19).

[55] E. M. Schlegel. “A new subclass of type II supernovae?” In: *Monthly Notices of the Royal Astronomical Society* 244 (May 1990), pp. 269–271 (cited on page 19).

[56] F. Taddia et al. “Carnegie Supernova Project: Observations of Type IIn supernovae”. In: *Astronomy & Astrophysics* 555 (2013), A10. DOI: `10.1051/0004-6361/201321180` (cited on page 19).

[57] N. Kurahashi et al. “High-Energy Extragalactic Neutrino Astrophysics”. In: *Annual Review of Nuclear and Particle Science* 72.1 (2022), pp. 365–387. DOI: `10.1146/annurev-nucl-011122-061547` (cited on page 19).

[58] E. Waxman and B. Katz. “Shock Breakout Theory”. In: *Handbook of Supernovae*. Springer International Publishing, 2017, pp. 967–1015. DOI: `10.1007/978-3-319-21846-5_33` (cited on page 19).

[59] M. Petropoulou et al. “Point-source and diffuse high-energy neutrino emission from Type IIn supernovae”. In: *Monthly Notices of the Royal Astronomical Society* 470.2 (2017), pp. 1881–1893. DOI: `10.1093/mnras/stx1251` (cited on page 19).

[60] K. Murase. “New prospects for detecting high-energy neutrinos from nearby supernovae”. In: *Physical Review D* 97.8 (2018), p. 081301. DOI: `10.1103/physrevd.97.081301` (cited on page 19).

[61] R. W. Klebesadel et al. “Observations of Gamma-Ray Bursts of Cosmic Origin”. In: *The Astrophysical Journal* 182 (1973), p. L85. DOI: `10.1086/181225` (cited on page 19).

[62] G. Vedrenne. *Gamma-ray Bursts. The Brightest Explosions in the Universe*. Springer, 2009, p. 571 (cited on page 19).

[63] B. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Physical Review Letters* 119.16 (2017), p. 161101. DOI: `10.1103/physrevlett.119.161101` (cited on page 19).

[64] T. J. Galama et al. “An unusual supernova in the error box of the γ-ray burst of 25 April 1998”. In: *Nature* 395.6703 (1998), pp. 670–672. DOI: `10.1038/27150` (cited on page 20).

[65] J. M. Leblanc and J. R. Wilson. “A Numerical Example of the Collapse of a Rotating Magnetized Star”. In: *The Astrophysical Journal* 161 (1970), p. 541. DOI: `10.1086/150558` (cited on page 20).

[66] E. Pian and P. A. Mazzali. “Hydrogen-Poor Core-Collapse Supernovae”. In: *Handbook of Supernovae*. Springer International Publishing, 2017, pp. 277–292. DOI: `10.1007/978-3-319-21846-5_40` (cited on page 20).

[67] E. Waxman and J. Bahcall. “High Energy Neutrinos from Cosmological Gamma-Ray Burst Fireballs”. In: *Physical Review Letters* 78.12 (1997), pp. 2292–2295. DOI: `10.1103/physrevlett.78.2292` (cited on page 20).

[68] N. Senno et al. “Choked jets and low-luminosity gamma-ray bursts as hidden neutrino sources”. In: *Physical Review D* 93.8 (2016), p. 083003. DOI: `10.1103/physrevd.93.083003` (cited on page 20).

[69] P. Mészáros and E. Waxman. “TeV Neutrinos from Successful and Choked Gamma-Ray Bursts”. In: *Physical Review Letters* 87.17 (2001), p. 171102. DOI: `10.1103/physrevlett.87.171102` (cited on page 20).

[70] P. Padovani et al. “Active galactic nuclei: what’s in a name?” In: *The Astronomy & Astrophysics Review* 25.1 (2017). DOI: `10.1007/s00159-017-0102-9` (cited on page 21).

[71] C. K. Seyfert. “Nuclear Emission in Spiral Nebulae.” In: *The Astrophysical Journal* 97 (1943), p. 28. DOI: `10.1086/144488` (cited on page 21).

[72] M. Schmidt. “3C 273: A Star-Like Object with Large Red-Shift”. In: *Nature* 197.4872 (1963), pp. 1040–1040. DOI: `10.1038/1971040a0` (cited on page 21).

[73] P. Padovani. “On the two main classes of active galactic nuclei”. In: *Nature Astronomy* 1.8 (2017). DOI: `10.1038/s41550-017-0194` (cited on page 21).

[74] J. Thorne et al. *AGN Unification Diagram*. 2022. DOI: `10.5281/zenodo.6381013` (cited on page 21).

[75] The Event Horizon Telescope Collaboration, K. Akiyama et al. “First M87 Event Horizon Telescope Results. IV. Imaging the Central Supermassive Black Hole”. In: *The Astrophysical Journal Letters* 875.1 (2019), p. L4. DOI: `10.3847/2041-8213/ab0e85` (cited on page 22).

[76] V. Beckmann and C. Shrader. *Active Galactic Nuclei*. Wiley-VCH Verlag GmbH, 2012, p. 381 (cited on page 22).

[77] D. Eichler. “High-energy neutrino astronomy - A probe of galactic nuclei”. In: *The Astrophysical Journal* 232 (1979), p. 106. DOI: `10.1086/157269` (cited on page 22).

[78] P. Padovani. “The Microjansky and Nanojansky Radio Sky: Source Population and Multiwavelength Properties”. In: *Monthly Notices of the Royal Astronomical Society* 411.3 (2010), pp. 1547–1561. DOI: `10.1111/j.1365-2966.2010.17789.x` (cited on page 22).

[79] F. Combes. *Active Galactic Nuclei*. IOP Publishing, 2021 (cited on page 22).

[80] B. Khiali et al. *Particle Acceleration and gamma-ray emission due to magnetic reconnection around the core region of radio galaxies*. 2015. DOI: `10.48550/ARXIV.1504.07592` (cited on page 22).

[81] A. Reimer. “On the Physics of Hadronic Blazar Emission Models”. In: *Journal of Physics: Conference Series* 355 (2012), p. 012011. DOI: `10.1088/1742-6596/355/1/012011` (cited on page 22).

[82] S. Gao et al. “Modelling the coincident observation of a high-energy neutrino and a bright blazar flare”. In: *Nature Astronomy* 3.1 (2018), pp. 88–92. DOI: `10.1038/s41550-018-0610-1` (cited on pages 22, 25).

[83] M. J. Rees. “Tidal disruption of stars by black holes of $10^6 - 10^8$ solar masses in nearby galaxies”. In: *Nature* 333.6173 (1988), pp. 523–528. DOI: `10.1038/333523a0` (cited on page 23).

[84] S. Gezari. “Tidal Disruption Events”. In: *Annual Review of Astronomy and Astrophysics* 59.1 (2021), pp. 21–58. DOI: `10.1146/annurev-astro-111720-030029` (cited on page 23).

[85] J. G. Hills. “Possible power source of Seyfert galaxies and QSOs”. In: *Nature* 254.5498 (1975), pp. 295–298. DOI: `10.1038/254295a0` (cited on page 23).

[86] S. van Velzen et al. “Optical Discovery of Probable Stellar Tidal Disruption Flares”. In: *The Astrophysical Journal* 741.2 (2011), p. 73. DOI: `10.1088/0004-637x/741/2/73` (cited on page 23).

[87] E. Hammerstein et al. “The Final Season Reimagined: 30 Tidal Disruption Events from the ZTF-I Survey”. In: *The Astrophysical Journal* 942.1 (2022), p. 9. DOI: `10.3847/1538-4357/aca283` (cited on pages 23, 70, 101).

[88] S. Gezari et al. “PS1-10jh Continues to Follow the Fallback Accretion Rate of a Tidally Disrupted Star”. In: *The Astrophysical Journal* 815.1 (2015), p. L5. DOI: `10.1088/2041-8205/815/1/l5` (cited on page 23).

[89] K. Hayasaki. “Neutrinos from tidal disruption events”. In: *Nature Astronomy* 5.5 (2021), pp. 436–437. DOI: `10.1038/s41550-021-01309-z` (cited on pages 23, 24).

[90] I. Andreoni et al. “A very luminous jet from the disruption of a star by a massive black hole”. In: *Nature* 612.7940 (2022), pp. 430–434. DOI: `10.1038/s41586-022-05465-8` (cited on page 23).

[91] K. Murase et al. “High-energy Neutrino and Gamma-Ray Emission from Tidal Disruption Events”. In: *The Astrophysical Journal* 902.2 (2020), p. 108. DOI: `10.3847/1538-4357/abb3c0` (cited on page 24).

[92] R.-Y. Liu et al. “Neutrino emission from an off-axis jet driven by the tidal disruption event AT2019dsg”. In: *Physical Review D* 102.8 (2020), p. 083028. DOI: `10.1103/physrevd.102.083028` (cited on page 24).

[93] IceCube Collaboration, M. Aartsen et al. “Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector”. In: *Science* 342.6161 (2013). DOI: `10.1126/science.1242856` (cited on pages 24, 43).

[94] IceCube Collaboration, R. Abbasi et al. “Improved Characterization of the Astrophysical Muon–neutrino Flux with 9.5 Years of IceCube Data”. In: *The Astrophysical Journal* 928.1 (2022), p. 50. DOI: `10.3847/1538-4357/ac4d29` (cited on page 24).

[95] IceCube Collaboration, M. G. Aartsen et al. “First Observation of PeV-Energy Neutrinos with IceCube”. In: *Physical Review Letters* 111.2 (2013), p. 021103. DOI: `10.1103/physrevlett.111.021103` (cited on page 24).

[96] C. Kopper and E. Blaufuss. “IceCube-170922A - IceCube observation of a high-energy neutrino candidate event.” In: *GCN Circular* 21916 (Jan. 2017), p. 1 (cited on page 25).

[97] IceCube Collaboration, M. Aartsen et al. “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A”. In: *Science* 361.6398 (2018). DOI: `10.1126/science.aat1378` (cited on page 25).

[98] S. Paiano et al. “The Redshift of the BL Lac Object TXS 0506+056”. In: *The Astrophysical Journal* 854.2 (2018), p. L32. DOI: `10.3847/2041-8213/aaad5e` (cited on page 25).

[99] IceCube Collaboration, M. Aartsen et al. “Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert”. In: *Science* 361.6398 (2018), pp. 147–151. DOI: `10.1126/science.aat2890` (cited on page 25).

[100] N. L. Strotjohann et al. “Eddington bias for cosmic neutrino sources”. In: *Astronomy & Astrophysics* 622 (2019), p. L9. DOI: `10.1051/0004-6361/201834750` (cited on page 25).

[101] X. Rodrigues et al. “Leptohadronic Blazar Models Applied to the 2014–2015 Flare of TXS 0506+056”. In: *The Astrophysical Journal* 874.2 (2019), p. L29. DOI: `10.3847/2041-8213/ab1267` (cited on page 26).

[102] IceCube Collaboration, and R. Abbasi et al. “Evidence for neutrino emission from the nearby active galaxy NGC 1068”. In: *Science* 378.6619 (2022), pp. 538–543. DOI: `10.1126/science.abg3395` (cited on pages 26, 37).

[103] V. G. Rosas et al. “Thermal imaging of dust hiding the black hole in NGC 1068”. In: *Nature* 602.7897 (2022), pp. 403–407. DOI: `10.1038/s41586-021-04311-7` (cited on page 26).

[104] B. Eichmann and J. B. Tjus. “The Radio-Gamma Correlation in Starburst Galaxies”. In: *The Astrophysical Journal* 821.2 (2016), p. 87. DOI: `10.3847/0004-637x/821/2/87` (cited on page 26).

[105] G. Cecil et al. “Imaging spectrophotometry of ionized gas in NGC 1068. I - Kinematics of the narrow-line region”. In: *The Astrophysical Journal* 355 (1990), p. 70. DOI: 10.1086/168742 (cited on page 26).

[106] R. Silberberg and M. M. Shapiro. “Neutrinos as a Probe for the Nature of and Processes in Active Galactic Nuclei”. In: *International Cosmic Ray Conference*. Vol. 10. International Cosmic Ray Conference. Jan. 1979, p. 357 (cited on page 26).

[107] M. Ajello et al. “3FHL: The Third Catalog of Hard *Fermi*-LAT Sources”. In: *The Astrophysical Journal Supplement Series* 232.2 (2017), p. 18. DOI: 10.3847/1538-4365/aa8221 (cited on page 26).

[108] W. B. Atwood et al. “The Large Area Telescope on the *Fermi Gamma-Ray Space Telescope*”. In: *The Astrophysical Journal* 697.2 (2009), pp. 1071–1102. DOI: 10.1088/0004-637x/697/2/1071 (cited on pages 26, 75).

[109] IceCube Collaboration, M. Huber. “Searches for steady neutrino emission from 3FHL blazars using eight years of IceCube data from the Northern hemisphere”. In: *Proceedings of 36th International Cosmic Ray Conference — PoS(ICRC2019)*. Sissa Medialab, 2019. DOI: 10.22323/1.358.0916 (cited on page 26).

[110] G. Principe et al. “The first catalog of *Fermi*-LAT sources below 100 MeV”. In: *Astronomy & Astrophysics* 618 (2018), A22. DOI: 10.1051/0004-6361/201833116 (cited on page 26).

[111] IceCube Collaboration, R. Abbasi et al. “Search for Astrophysical Neutrinos from 1FLE Blazars with IceCube”. In: *The Astrophysical Journal* 938.1 (2022), p. 38. DOI: 10.3847/1538-4357/ac8de4 (cited on page 27).

[112] A. Plavin et al. “Observational Evidence for the Origin of High-energy Neutrinos in Parsec-scale Nuclei of Radio-bright Active Galaxies”. In: *The Astrophysical Journal* 894.2 (2020), p. 101. DOI: 10.3847/1538-4357/ab86bd (cited on page 27).

[113] R. Abbasi et al. *Search for correlations of high-energy neutrinos detected in IceCube with radio-bright AGN and gamma-ray emission from blazars*. 2023. DOI: 10.48550/a.2304.12675 (cited on page 27).

[114] E. Massaro et al. “VizieR Online Data Catalog: The Roma BZCAT - 5th edition (Massaro+, 2015)”. In: *VizieR Online Data Catalog*, VII/274 (Feb. 2016), pp. VII/274 (cited on page 27).

[115] IceCube Collaboration, C. Bellenghi et al. “Correlating high-energy IceCube neutrinos with 5BZCAT blazars and RFC sources”. In: (Sept. 2023). DOI: 10.48550/arxiv.2309.03115 (cited on page 27).

[116] IceCube Collaboration, R. Abbasi et al. “Search for neutrino emission from cores of active galactic nuclei”. In: *Physical Review D* 106.2 (2022), p. 022005. DOI: 10.1103/physrevd.106.022005 (cited on page 27).

[117] IceCube Collaboration, M. G. Aartsen et al. “An All-Sky Search for Three Flavors of Neutrinos from Gamma-Ray Bursts with the IceCube Neutrino Observatory”. In: *The Astrophysical Journal* 824.2 (2016), p. 115. DOI: 10.3847/0004-637x/824/2/115 (cited on page 28).

[118] IceCube Collaboration, R. Abbasi et al. “Searches for Neutrinos from Gamma-Ray Bursts Using the IceCube Neutrino Observatory”. In: *The Astrophysical Journal* 939.2 (2022), p. 116. DOI: 10.3847/1538-4357/ac9785 (cited on page 28).

[119] IceCube Collaboration, R. Abbasi et al. “Constraining High-energy Neutrino Emission from Supernovae with IceCube”. In: *The Astrophysical Journal Letters* 949.1 (2023), p. L12. DOI: 10.3847/2041-8213/acd2c9 (cited on page 28).

[120] O. Yaron and A. Gal-Yam. “WISeREP – An Interactive Supernova Data Repository”. In: *Publications of the Astronomical Society of the Pacific* 124.917 (2012), pp. 668–681. DOI: 10.1086/666656 (cited on page 28).

[121] J. Guillochon et al. “An Open Catalog for Supernova Data”. In: *The Astrophysical Journal* 835.1 (2017), p. 64. DOI: 10.3847/1538-4357/835/1/64 (cited on page 28).

[122] IceCube Collaboration, R. Stein. “Search for Neutrinos from Populations of Optical Transients”. In: *Proceedings of 36th International Cosmic Ray Conference — PoS(ICRC2019)*. Sissa Medialab, 2019. DOI: 10.22323/1.358.1016 (cited on page 28).

[123] C. Guépin et al. “High-energy neutrino transients and the future of multi-messenger astronomy”. In: *Nature Reviews Physics* 4.11 (2022), pp. 697–712. DOI: `10.1038/s42254-022-00504-9` (cited on page 29).

[124] I. Bartos et al. “The IceCube Pie Chart: Relative Source Contributions to the Cosmic Neutrino Flux”. In: *The Astrophysical Journal* 921.1 (2021), p. 45. DOI: `10.3847/1538-4357/ac1c7b` (cited on page 29).

[125] AMANDA Collaboration, E. Andrés et al. “The AMANDA neutrino telescope”. In: *Nuclear Physics B - Proceedings Supplements* 77.1-3 (1999), pp. 474–485. DOI: `10.1016/s0920-5632(99)00469-7` (cited on page 31).

[126] AMANDA Collaboration, E. Andrés et al. “The AMANDA neutrino telescope: principle of operation and first results”. In: *Astroparticle Physics* 13.1 (2000), pp. 1–20. DOI: `10.1016/s0927-6505(99)00092-4` (cited on page 31).

[127] P. A. Cherenkov. “Visible emission of clean liquids by action of gamma radiation”. In: *Doklady Akademii Nauk SSSR* 8 (1934), pp. 451–454 (cited on page 31).

[128] M. F. L’Annunziata. *Handbook of Radioactivity Analysis. Volume 1: Radiation Physics and Detectors. Radiation Physics and Detectors*. 4th ed. Elsevier Science & Technology Books, 2020 (cited on page 32).

[129] L. Filöp and T. Biró. “Cherenkov radiation spectrum”. In: *International Journal of Theoretical Physics* 31.1 (1992), pp. 61–74. DOI: `10.1007/bf00674341` (cited on page 32).

[130] S. G. Warren and R. E. Brandt. “Optical constants of ice from the ultraviolet to the microwave: A revised compilation”. In: *Journal of Geophysical Research: Atmospheres* 113.D14 (2008). DOI: `10.1029/2007jd009744` (cited on page 32).

[131] H. Iams and B. Salzberg. “The Secondary Emission Phototube”. In: *Proceedings of the IRE* 23.1 (1935), pp. 55–64. DOI: `10.1109/jrproc.1935.227243` (cited on page 32).

[132] T. Bednarski et al. “Calibration of photomultipliers gain used in the J-PET detector”. In: *Bio-Algorithms and Med-Systems* 10.1 (2014), pp. 13–17. DOI: `10.1515/bams-2013-0110` (cited on page 33).

[133] A. Einstein. “Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt”. In: *Annalen der Physik* 322.6 (1905), pp. 132–148. DOI: `10.1002/andp.19053220607` (cited on page 32).

[134] A. G. Wright. *The Photomultiplier Handbook*. Oxford University Press, 2017, p. 624 (cited on page 33).

[135] IceCube Collaboration, R. Abbasi et al. “Calibration and characterization of the IceCube photomultiplier tube”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 618.1-3 (2010), pp. 139–152. DOI: `10.1016/j.nima.2010.03.102` (cited on page 33).

[136] Icecube Collaboration, M. Aartsen et al. “The IceCube Neutrino Observatory: instrumentation and online systems”. In: *Journal of Instrumentation* 12.03 (2017), P03012–P03012. DOI: `10.1088/1748-0221/12/03/p03012` (cited on pages 33–36).

[137] M. Ahlers et al. “Probing particle physics with IceCube”. In: *The European Physical Journal C* 78.11 (2018). DOI: `10.1140/epjc/s10052-018-6369-9` (cited on pages 34, 35, 42, 43).

[138] T. Benson et al. “IceCube Enhanced Hot Water Drill functional description”. In: *Annals of Glaciology* 55.68 (2014), pp. 105–114. DOI: `10.3189/2014aog68a032` (cited on page 35).

[139] R. Abbasi et al. “IceTop: The surface component of IceCube”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 700 (2013), pp. 188–220. DOI: `10.1016/j.nima.2012.10.067` (cited on page 36).

[140] IceCube Collaboration, M. G. Aartsen et al. “Measurement of the cosmic ray energy spectrum with IceTop-73”. In: *Physical Review D* 88.4 (2013), p. 042004. DOI: `10.1103/physrevd.88.042004` (cited on page 36).

[141] IceCube Collaboration, M. Aartsen et al. “Cosmic ray spectrum and composition from PeV to EeV using 3 years of data from IceTop and IceCube”. In: *Physical Review D* 100.8 (2019), p. 082002. DOI: `10.1103/physrevd.100.082002` (cited on page 36).

[142] IceCube Collaboration, N. M. B. Amin. “Implementation of IceTop data in the IceCube Realtime Alert System”. In: *Journal of Physics: Conference Series* 2156.1 (2021), p. 012217. DOI: `10.1088/1742-6596/2156/1/012217` (cited on page 36).

[143] IceCube Collaboration, R. Abbasi et al. “The IceCube data acquisition system: Signal capture, digitization, and timestamping”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 601.3 (2009), pp. 294–316. DOI: `10.1016/j.nima.2009.01.001` (cited on pages 36, 37).

[144] IceCube Collaboration, M. G. Aartsen et al. “Energy reconstruction methods in the IceCube neutrino telescope”. In: *Journal of Instrumentation* 9.03 (2014), P03009–P03009. DOI: `10.1088/1748-0221/9/03/p03009` (cited on page 37).

[145] IceCube Collaboration, M. Aartsen et al. “The IceCube realtime alert system”. In: *Astroparticle Physics* 92 (2017), pp. 30–41. DOI: `10.1016/j.astropartphys.2017.05.002` (cited on pages 37, 42).

[146] AMANDA Collaboration, J. Ahrens et al. “Muon track reconstruction and data selection techniques in AMANDA”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 524.1-3 (2004), pp. 169–194. DOI: `10.1016/j.nima.2004.01.065` (cited on pages 37–40).

[147] IceCube Collaboration, R. Abbasi et al. “In-situ estimation of ice crystal properties at the South Pole using LED calibration data from the IceCube Neutrino Observatory”. In: (2022). DOI: `10.5194/tc-2022-174` (cited on page 39).

[148] IceCube Collaboration, R. Abbasi et al. “A muon-track reconstruction exploiting stochastic losses for large-scale Cherenkov detectors”. In: *Journal of Instrumentation* 16.08 (2021), P08034. DOI: `10.1088/1748-0221/16/08/p08034` (cited on page 40).

[149] N. Whitehorn et al. “Penalized splines for smooth representation of high-dimensional Monte Carlo datasets”. In: *Computer Physics Communications* 184.9 (2013), pp. 2214–2220. DOI: `10.1016/j.cpc.2013.04.008` (cited on page 40).

[150] K. M. Gorski et al. “HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere”. In: *The Astrophysical Journal* 622.2 (2005), pp. 759–771. DOI: `10.1086/427976` (cited on page 41).

[151] IceCube Collaboration, R. Abbasi et al. *IceCat-1: the IceCube Event Catalog of Alert Tracks*. 2023. DOI: `10.48550/arxiv.2304.01174` (cited on pages 41, 43, 44).

[152] E. Kankare et al. “Search for transient optical counterparts to high-energy IceCube neutrinos with Pan-STARRS1”. In: *Astronomy& Astrophysics* 626 (2019), A117. DOI: `10.1051/0004-6361/201935171` (cited on page 41).

[153] IceCube Collaboration, C. Lagunas Gualda et al. *Studies of systematic uncertainty effects on IceCube's real-time angular uncertainty*. 2021. DOI: `10.48550/arxiv.2107.08670` (cited on page 41).

[154] E. Blaufuss et al. “The Next Generation of IceCube Real-time Neutrino Alerts”. In: *Proceedings of 36th International Cosmic Ray Conference — PoS(ICRC2019)*. Sissa Medialab, 2019. DOI: `10.22323/1.358.1021` (cited on pages 42, 44).

[155] D. Chirkin and W. Rhode. “Propagating leptons through matter with Muon Monte Carlo (MMC)”. In: (July 2004). DOI: `10.48550/arxiv.hep-ph/0407075` (cited on page 42).

[156] IceCube Collaboration, M. G. Aartsen et al. “Measurement of the multi-TeV neutrino interaction cross-section with IceCube using Earth absorption”. In: *Nature* 551.7682 (2017), pp. 596–600. DOI: `10.1038/nature24459` (cited on pages 42, 43).

[157] IceCube Collaboration, R. Abbasi et al. “IceCube high-energy starting event sample: Description and flux characterization with 7.5 years of data”. In: *Physical Review D* 104.2 (2021), p. 022002. DOI: `10.1103/physrevd.104.022002` (cited on page 43).

[158] IceCube Collaboration, M. G. Aartsen et al. “Constraints on Ultrahigh-Energy Cosmic-Ray Sources from a Search for Neutrinos above 10 PeV with IceCube”. In: *Physical Review Letters* 117.24 (2016), p. 241101. DOI: `10.1103/physrevlett.117.241101` (cited on page 43).

[159] The IceCube, MAGIC and VERITAS Collaborations, M. G. Aartsen et al. “Very high-energy gamma-ray follow-up program using neutrino triggers from IceCube”. In: *Journal of Instrumentation* 11.11 (2016), P11009–P11009. doi: `10.1088/1748-0221/11/11/p11009` (cited on page 43).

[160] F. Zwicky. “Die Rotverschiebung von extragalaktischen Nebeln”. In: *Helvetica Physica Acta* 6 (Jan. 1933), pp. 110–127 (cited on page 45).

[161] W. Baade and F. Zwicky. “On Super-Novae”. In: *Proceedings of the National Academy of Sciences* 20.5 (1934), pp. 254–259. doi: `10.1073/pnas.20.5.254` (cited on page 45).

[162] R. G. Harrington. “The 48-inch Schmidt-type Telescope at Palomar Observatory”. In: *Publications of the Astronomical Society of the Pacific* 64 (1952), p. 275. doi: `10.1086/126494` (cited on page 45).

[163] S. H. Pravdo et al. “The Near-Earth Asteroid Tracking (NEAT) Program: An Automated System for Telescope Control, Wide-Field Imaging, and Object Detection”. In: *The Astronomical Journal* 117.3 (1999), pp. 1616–1633. doi: `10.1086/300769` (cited on page 45).

[164] C. Baltay et al. “The QUEST Large Area CCD Camera”. In: *Publications of the Astronomical Society of the Pacific* 119.861 (2007), pp. 1278–1294. doi: `10.1086/523899` (cited on page 45).

[165] N. M. Law et al. “The Palomar Transient Factory: System Overview, Performance, and First Results”. In: *Publications of the Astronomical Society of the Pacific* 121.886 (2009), pp. 1395–1408. doi: `10.1086/648598` (cited on page 45).

[166] E. C. Bellm. “Volumetric Survey Speed: A Figure of Merit for Transient Surveys”. In: *Publications of the Astronomical Society of the Pacific* 128.966 (2016), p. 084501. doi: `10.1088/1538-3873/128/966/084501` (cited on page 45).

[167] N. Blagorodnova et al. “The SED Machine: A Robotic Spectrograph for Fast Transient Classification”. In: *Publications of the Astronomical Society of the Pacific* 130.985 (2018), p. 035003. doi: `10.1088/1538-3873/aaa53f` (cited on page 46).

[168] B. Schmidt. “Ein lichtstarkes komafreies Spiegelsystem”. In: *Mitteilungen der Hamburger Sternwarte in Bergedorf* (1938) (cited on page 46).

[169] R. D. Cannon. “Schmidt Telescopes: Their Past, Present and Future”. In: *IAU Colloq. 148: The Future Utilisation of Schmidt Telescopes*. Ed. by J. Chapman et al. Vol. 84. Astronomical Society of the Pacific Conference Series. Jan. 1995, p. 8 (cited on page 47).

[170] *The HIPPARCOS and TYCHO catalogues. Astrometric and photometric star catalogues derived from the ESA HIPPARCOS Space Astrometry Mission*. Vol. 1200. ESA Special Publication. Jan. 1997 (cited on page 47).

[171] D. G. Koch et al. “Kepler Mission Design, Realized Photometric Performance, and Early Science”. In: *The Astrophysical Journal* 713.2 (2010), pp. L79–L86. doi: `10.1088/2041-8205/713/2/l79` (cited on page 47).

[172] R. Dekany et al. “The Zwicky Transient Facility: Observing System”. In: *Publications of the Astronomical Society of the Pacific* 132.1009 (2020), p. 038001. doi: `10.1088/1538-3873/ab4ca2` (cited on pages 47, 49–51).

[173] R. Dekany et al. “The Zwicky Transient Facility Camera”. In: *SPIE Proceedings*. Ed. by C. J. Evans et al. SPIE, 2016. doi: `10.1117/12.2234558` (cited on pages 47, 49).

[174] E. C. Bellm et al. “The Zwicky Transient Facility: System Overview, Performance, and First Results”. In: *Publications of the Astronomical Society of the Pacific* 131.995 (2019), p. 018002. doi: `10.1088/1538-3873/aaecbe` (cited on pages 47–49).

[175] Ž. Ivezić et al. “LSST: From Science Drivers to Reference Design and Anticipated Data Products”. In: *The Astrophysical Journal* 873.2 (2019), p. 111. doi: `10.3847/1538-4357/ab042c` (cited on pages 47, 60, 95, 128).

[176] R. Laher et al. “Processing Images from the Zwicky Transient Facility”. In: *Robotic Telescopes, Student Research and Education Proceedings, Vol 1, No 1*. Our Solar Siblings, 2018. doi: `10.32374/rtsre.2017.031` (cited on pages 47, 50).

[177] J. R. Janesick et al. “Scientific Charge-Coupled Devices”. In: *Optical Engineering* 26.8 (1987). doi: `10.1117/12.7974139` (cited on pages 47, 48).

[178] D. G. York et al. “The Sloan Digital Sky Survey: Technical Summary”. In: *The Astronomical Journal* 120.3 (2000), pp. 1579–1587. DOI: `10.1086/301513` (cited on page 49).

[179] N. Kaiser et al. “Pan-STARRS: A Large Synoptic Survey Telescope Array”. In: *SPIE Proceedings*. Ed. by J. A. Tyson and S. Wolff. SPIE, 2002. DOI: `10.1117/12.457365` (cited on page 49).

[180] *Gaia* Collaboration, T. Prusti et al. “The *Gaia* mission”. In: *Astronomy & Astrophysics* 595 (2016), A1. DOI: `10.1051/0004-6361/201629272` (cited on page 49).

[181] F. J. Masci et al. “The Zwicky Transient Facility: Data Processing, Products, and Archive”. In: *Publications of the Astronomical Society of the Pacific* 131.995 (2019), p. 018003. DOI: `10.1088/1538-3873/aae8ac` (cited on pages 50, 52, 54).

[182] E. C. Bellm et al. “The Zwicky Transient Facility: Surveys and Scheduler”. In: *Publications of the Astronomical Society of the Pacific* 131.1000 (2019), p. 068003. DOI: `10.1088/1538-3873/ab0c2a` (cited on pages 50, 55).

[183] F. J. Masci et al. *The ZTF Science Data System (ZDSD) – Explanatory Supplement*. 2020 (cited on pages 51, 53, 54).

[184] S. B. Howell. *Handbook of CCD Astronomy*. Cambridge University Press, 2006 (cited on page 51).

[185] *Gaia* Collaboration, A. G. A. Brown et al. “*Gaia* Data Release 1. Summary of the astrometric, photometric, and survey properties”. In: *Astronomy & Astrophysics* 595 (2016), A2. DOI: `10.1051/0004-6361/201629512` (cited on page 51).

[186] E. Bertin. “Automatic Astrometric and Photometric Calibration with SCAMP”. In: *Astronomical Data Analysis Software and Systems XV*. Vol. 351. Astronomical Society of the Pacific Conference Series, 2006, p. 112 (cited on page 52).

[187] E. Bertin and S. Arnouts. “SExtractor: Software for source extraction”. In: *Astronomy & Astrophysics Supplement Series* 117.2 (1996), pp. 393–404. DOI: `10.1051/aas:1996164` (cited on page 52).

[188] F. R. Chromey. *To Measure the Sky*. Cambridge University Press, 2016 (cited on page 52).

[189] P. B. Stetson. “DAOPHOT – A computer program for crowded-field stellar photometry”. In: *Publications of the Astronomical Society of the Pacific* 99 (1987), p. 191. DOI: `10.1086/131977` (cited on pages 52, 80).

[190] E. F. Schlafly et al. “Photometric Calibration of the First 1.5 Years of the Pan-STARRS1 Survey”. In: *The Astrophysical Journal* 756.2 (2012), p. 158. DOI: `10.1088/0004-637x/756/2/158` (cited on page 52).

[191] E. A. Magnier et al. “Pan-STARRS Photometric and Astrometric Calibration”. In: *The Astrophysical Journal Supplement Series* 251.1 (2020), p. 6. DOI: `10.3847/1538-4365/abb82a` (cited on page 52).

[192] K. Krisciunas. *A Brief History of Astronomical Brightness Determination Methods at Optical Wavelengths*. 2001. DOI: `10.48550/arxiv.astro-ph/0106313` (cited on page 53).

[193] D. Jones. “Norman Pogson and the Definition of Stellar Magnitude”. In: *Leaflet of the Astronomical Society of the Pacific* 10.469 (July 1968), pp. 145–152 (cited on page 53).

[194] J. B. Oke and J. E. Gunn. “Secondary standard stars for absolute spectrophotometry”. In: *The Astrophysical Journal* 266 (1983), p. 713. DOI: `10.1086/160817` (cited on page 53).

[195] J. L. Tonry et al. “The Pan-STARRS1 Photometric System”. In: *The Astrophysical Journal* 750.2 (2012), p. 99. DOI: `10.1088/0004-637x/750/2/99` (cited on page 53).

[196] E. Bertin. *SWarp: Resampling and Co-adding FITS Images Together*. Astrophysics Source Code Library, record ascl:1010.068. Oct. 2010 (cited on page 54).

[197] B. Zackay et al. “Proper Image Subtraction – Optimal Transient Detection, Photometry, and Hypothesis Testing”. In: *The Astrophysical Journal* 830.1 (2016), p. 27. DOI: `10.3847/0004-637x/830/1/27` (cited on page 54).

[198] A. Mahabal et al. “Machine Learning for the Zwicky Transient Facility”. In: *Publications of the Astronomical Society of the Pacific* 131.997 (2019), p. 038002. DOI: `10.1088/1538-3873/aaf3fa` (cited on pages 54, 55).

[199] Y. Tachibana and A. A. Miller. “A Morphological Classification Model to Identify Unresolved PanSTARRS1 Sources: Application in the ZTF Real-time Pipeline”. In: *Publications of the Astronomical Society of the Pacific* 130.994 (2018), p. 128001. DOI: `10.1088/1538-3873/aae3d9` (cited on pages 55, 96).

[200] M. T. Patterson et al. “The Zwicky Transient Facility Alert Distribution System”. In: *Publications of the Astronomical Society of the Pacific* 131.995 (2018), p. 018001. DOI: `10.1088/1538-3873/aae904` (cited on page 55).

[201] LIGO Scientific Collaboration, J. Aasi et al. “Advanced LIGO”. In: *Classical and Quantum Gravity* 32.7 (2015), p. 074001. DOI: `10.1088/0264-9381/32/7/074001` (cited on page 55).

[202] F. Acernese et al. “Advanced Virgo: A second-generation interferometric gravitational wave detector”. In: *Classical and Quantum Gravity* 32.2 (2014), p. 024001. DOI: `10.1088/0264-9381/32/2/024001` (cited on page 55).

[203] M. M. Kasliwal et al. “Kilonova Luminosity Function Constraints Based on Zwicky Transient Facility Searches for 13 Neutron Star Merger Triggers during O3”. In: *The Astrophysical Journal* 905.2 (2020), p. 145. DOI: `10.3847/1538-4357/abc335` (cited on pages 55, 61).

[204] S. Reusch et al. *planobs: Release 0.6.4*. 2023. DOI: `10.5281/zenodo.7774765` (cited on page 58).

[205] J. Nordin et al. “Transient processing and analysis using AMPEL: Alert management, photometry, and evaluation of light curves”. In: *Astronomy & Astrophysics* 631 (2019), A147. DOI: `10.1051/0004-6361/201935634` (cited on pages 60, 61).

[206] R. Stein et al. *nuztf: v2.6.5 Release*. 2023. DOI: `10.5281/zenodo.7710351` (cited on page 61).

[207] A. J. Drake et al. “First Results from the Catalina Real-Time Transient Survey”. In: *The Astrophysical Journal* 696.1 (2009), pp. 870–884. DOI: `10.1088/0004-637x/696/1/870` (cited on page 62).

[208] A. J. Drake et al. “The Catalina Survey Periodic Variable Star Catalog”. In: *The Astrophysical Journal Supplement Series10.1088/0004-637X/696/1/870* 213.1 (2014), p. 9. DOI: `10.1088/0067-0049/213/1/9` (cited on page 62).

[209] *Gaia* Collaboration, and A. G. A. Brown et al. “*Gaia* Data Release 2”. In: *Astronomy & Astrophysics* 616 (2018), A1. DOI: `10.1051/0004-6361/201833051` (cited on page 62).

[210] E. W. Flesch. *The Million Quasars (Milliquas) Catalogue, v8*. 2023. DOI: `10.48550/arxiv.2308.01505` (cited on pages 62, 108).

[211] C. P. Ahn et al. “The Tenth Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the SDS-III Apache Point Observatory Galactic Evolution Experiment”. In: *The Astrophysical Journal Supplement Series* 211.2 (2014), p. 17. DOI: `10.1088/0067-0049/211/2/17` (cited on page 62).

[212] E. L. Wright et al. “The Wide-Field Infrared Survey Explorer (WISE): Mission Description and Initial On-Orbit Performance”. In: *The Astronomical Journal* 140.6 (2010), pp. 1868–1881. DOI: `10.1088/0004-6256/140/6/1868` (cited on page 62).

[213] R. E. Hviding et al. “A New Infrared Criterion for Selecting Active Galactic Nuclei to Lower Luminosities”. In: *The Astronomical Journal* 163.5 (2022), p. 224. DOI: `10.3847/1538-3881/ac5e33` (cited on pages 62, 107, 108).

[214] S. Reusch. *fpbot: Release 1.1.2*. 2023. DOI: `10.5281/zenodo.7781291` (cited on page 63).

[215] H. Dembinski et al. *scikit-hep/iminuit*. 2023. DOI: `10.5281/zenodo.3949207` (cited on page 64).

[216] F. James and M. Roos. “Minuit – A system for function minimization and analysis of the parameter errors and correlations”. In: *Computer Physics Communications* 10.6 (1975), pp. 343–367. DOI: `10.1016/0010-4655(75)90039-9` (cited on page 64).

[217] A. A. Djupvik and J. Andersen. “The Nordic Optical Telescope”. In: *Astrophysics and Space Science Proceedings*. Springer Berlin Heidelberg, 2010, pp. 211–218. DOI: `10.1007/978-3-642-11250-8_21` (cited on page 64).

[218] D. Mancini et al. “Italian National Galileo Telescope (TNG) system description and tracking performance in the workshop”. In: *SPIE Proceedings*. Ed. by M. K. Masten and L. A. Stockum. SPIE, 1997. DOI: `10.1117/12.277175` (cited on page 65).

[219] I. M. Hook et al. “The Gemini–North Multi-Object Spectrograph: Performance in Imaging, Long-Slit, and Multi-Object Spectroscopic Modes”. In: *Publications of the Astronomical Society of the Pacific* 116.819 (2004), pp. 425–440. DOI: `10.1086/383624` (cited on page 65).

[220] R. W. Pogge et al. “The multi-object double spectrographs for the Large Binocular Telescope”. In: *SPIE Proceedings*. Ed. by I. S. McLean et al. SPIE, 2010. DOI: `10.1117/12.857215` (cited on page 65).

[221] J. B. Oke et al. “The Keck Low-Resolution Imaging Spectrometer”. In: *Publications of the Astronomical Society of the Pacific* 107 (1995), p. 375. DOI: `10.1086/133562` (cited on page 65).

[222] J. Cepa et al. “OSIRIS tunable imager and spectrograph for the GTC. Instrument status”. In: *SPIE Proceedings*. Ed. by M. Iye and A. F. M. Moorwood. SPIE, 2003. DOI: `10.1117/12.460913` (cited on page 65).

[223] D. Tody. “The IRAF Data Reduction And Analysis System”. In: *SPIE Proceedings*. Ed. by D. L. Crawford. SPIE, 1986. DOI: `10.1117/12.968154` (cited on page 65).

[224] J. X. Prochaska et al. *PypeIt: Release 1.0.0.* 2020. DOI: `10.5281/zenodo.3743493` (cited on page 65).

[225] R. Stein et al. “Neutrino Follow-Up with the Zwicky Transient Facility: Results from the first 24 Campaigns”. In: *Monthly Notices of the Royal Astronomical Society* (2023). DOI: `10.1093/mnras/stad767` (cited on pages 66, 67).

[226] R. Stein. “IceCube-200530A – IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 27865 (2020) (cited on pages 66, 73, 132).

[227] S. Reusch et al. “IceCube-200530A: Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 27872 (2020) (cited on pages 66, 132).

[228] S. Reusch et al. “IceCube-200530A – SN2020lls likely unrelated”. In: *GCN Circular* 27980 (2020) (cited on pages 66, 132).

[229] S. Blondin and J. L. Tonry. “Determining the Type, Redshift, and Age of a Supernova Spectrum”. In: *The Astrophysical Journal* 666.2 (2007), pp. 1024–1047. DOI: `10.1086/520494` (cited on page 66).

[230] J. Guillochon et al. “MOSFiT: Modular Open Source Fitter for Transients”. In: *The Astrophysical Journal Supplement Series* 236.1 (2018), p. 6. DOI: `10.3847/1538-4365/aab761` (cited on page 67).

[231] S. Reusch et al. “IceCube-200530A – SN2020lam classified as SN II without evidence of CSM interaction”. In: *GCN Circular* 27910 (2020) (cited on pages 67, 132).

[232] A. H. Harutyunyan et al. “ESC supernova spectroscopy of non-ESC targets”. In: *Astronomy & Astrophysics* 488.1 (2008), pp. 383–399. DOI: `10.1051/0004-6361:20078859` (cited on page 67).

[233] C. Lagunas Gualda. “IceCube-201021A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 28715 (2020) (cited on pages 68, 132).

[234] R. Stein et al. “IceCube-201021A: One candidate counterpart from the Zwicky Transient Facility”. In: *GCN Circular* 28757 (Oct. 2020), p. 1 (cited on pages 68, 132).

[235] M. Santander. “IceCube-210629A – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circular* 30342 (2021) (cited on pages 68, 132).

[236] P. Sánchez-Sáez et al. “Alert Classification for the ALeRCE Broker System: The Light Curve Classifier”. In: *The Astronomical Journal* 161.3 (2021), p. 141. DOI: `10.3847/1538-3881/abd5c1` (cited on page 68).

[237] J. Necker et al. “IceCube-210629A: One Candidate Counterpart from the Zwicky Transient Facility”. In: *GCN Circular* 30349 (2021) (cited on pages 68, 132).

[238] S. Reusch et al. “IceCube-210629A – Spectroscopic observation of AT2021osi”. In: *GCN Circular* 30419 (July 2021), p. 1 (cited on pages 68, 132).

[239] M. Lincetto. “IceCube-221223A – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GRB Coordinates Network* 33094 (2022) (cited on pages 69, 132).

[240] R. Stein et al. “IceCube-201021A: One candidate counterpart from the Zwicky Transient Facility”. In: *GRB Coordinates Network* 28757 (2020) (cited on pages 69, 132).

[241] R. Stein. “IceCube-191001A – IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 25913 (2019) (cited on pages 69, 132).

[242] R. Stein et al. “Candidate Counterparts to IceCube-191001A with ZTF”. In: *The Astronomer's Telegram* 13160 (Oct. 2019), p. 1 (cited on pages 69, 132).

[243] P. W. A. Roming et al. “The *Swift* Ultra-Violet/Optical Telescope”. In: *Space Science Reviews* 120.3-4 (2005), pp. 95–142. DOI: `10.1007/s11214-005-5095-4` (cited on page 70).

[244] N. Gehrels et al. “The *Swift* Gamma-Ray Burst Mission”. In: *The Astrophysical Journal* 611.2 (2004), pp. 1005–1020. DOI: `10.1086/422091` (cited on page 70).

[245] R. Stein et al. “A tidal disruption event coincident with a high-energy neutrino”. In: *Nature Astronomy* 5.5 (2021), pp. 510–518. DOI: `10.1038/s41550-020-01295-8` (cited on pages 70, 71).

[246] D. N. Burrows et al. “The *Swift* X-Ray Telescope”. In: *Space Science Reviews* 120.3-4 (2005), pp. 165–195. DOI: `10.1007/s11214-005-5097-2` (cited on page 70).

[247] J. T. L. Zwart et al. “The Arcminute Microkelvin Imager”. In: *Monthly Notices of the Royal Astronomical Society* 391.4 (2008), pp. 1545–1558. DOI: `10.1111/j.1365-2966.2008.13953.x` (cited on page 70).

[248] J. Hickish et al. “A digital correlator upgrade for the Arcminute MicroKelvin Imager”. In: *Monthly Notices of the Royal Astronomical Society* 475.4 (2018), pp. 5677–5687. DOI: `10.1093/mnras/sty074` (cited on page 70).

[249] J. Jonas. “The MeerKAT Radio Telescope”. In: *Proceedings of MeerKAT Science: On the Pathway to the SKA — PoS(MeerKAT2016)*. Sissa Medialab, 2018. DOI: `10.22323/1.277.0001` (cited on page 70).

[250] A. R. Thompson et al. “The Very Large Array”. In: *The Astrophysical Journal Supplement Series* 44 (1980), p. 151. DOI: `10.1086/190688` (cited on page 70).

[251] Y. Cendes et al. “Radio Observations of an Ordinary Outflow from the Tidal Disruption Event AT2019dsg”. In: *The Astrophysical Journal* 919.2 (2021), p. 127. DOI: `10.3847/1538-4357/ac110a` (cited on page 70).

[252] P. Mohan et al. “High-resolution VLBI Observations of and Modeling the Radio Emission from the Tidal Disruption Event AT2019dsg”. In: *The Astrophysical Journal* 927.1 (2022), p. 74. DOI: `10.3847/1538-4357/ac4cb2` (cited on page 70).

[253] S. Reusch. “Multi-messenger observations of Tidal Disruption Events”. In: *Proceedings of 27th European Cosmic Ray Symposium — PoS(ECRS)*. Sissa Medialab, 2023. DOI: `10.22323/1.423.0020` (cited on page 71).

[254] A. Mainzer et al. “Preliminary Results from NEOWISE: An Enhancement to the Wide-Field Infrared Survey Explorer for Solar System Science”. In: *The Astrophysical Journal* 731.1 (2011), p. 53. DOI: `10.1088/0004-637x/731/1/53` (cited on page 71).

[255] C. Fremling et al. “The Zwicky Transient Facility Bright Transient Survey. I. Spectroscopic Classification and the Redshift Completeness of Local Galaxy Catalogs”. In: *The Astrophysical Journal* 895.1 (2020), p. 32. DOI: `10.3847/1538-4357/ab8943` (cited on pages 71, 104).

[256] D. A. Perley et al. “The Zwicky Transient Facility Bright Transient Survey. II. A Public Statistical Sample for Exploring Supernova Demographics”. In: *The Astrophysical Journal* 904.1 (2020), p. 35. DOI: `10.3847/1538-4357/abbd98` (cited on pages 71, 104, 105).

[257] J. Nordin et al. “ZTF Transient Discovery Report for 2019-05-13”. In: *Transient Name Server Discovery Report* 2019-771 (May 2019), p. 1 (cited on page 73).

[258] S. Frederick et al. “A Family Tree of Optical Transients from Narrow-line Seyfert 1 Galaxies”. In: *The Astrophysical Journal* 920.1 (2021), p. 56. DOI: `10.3847/1538-4357/ac110f` (cited on pages 73, 74, 76, 87, 88, 143).

[259] S. Reusch et al. “Candidate Tidal Disruption Event AT2019fdr Coincident with a High-Energy Neutrino”. In: *Physical Review Letters* 128.22 (2022), p. 221101. DOI: `10.1103/physrevlett.128.221101` (cited on pages 74–79, 81, 83, 90, 133).

[260] J. C. Wilson et al. "A Wide-Field Infrared Camera for the Palomar 200-inch Telescope". In: *SPIE Proceedings*. Ed. by M. Iye and A. F. M. Moorwood. SPIE, 2003. DOI: `10.1117/12.460336` (cited on page 74).

[261] P. Predehl et al. "The eROSITA X-ray telescope on SRG". In: *Astronomy & Astrophysics* 647 (2021), A1. DOI: `10.1051/0004-6361/202039313` (cited on page 75).

[262] R. Sunyaev et al. "SRG X-ray orbital observatory". In: *Astronomy & Astrophysics* 656 (2021), A132. DOI: `10.1051/0004-6361/202141179` (cited on page 75).

[263] S. van Velzen et al. *Establishing accretion flares from massive black holes as a major source of high-energy neutrinos*. 2021. DOI: `10.48550/arxiv.2111.09391` (cited on pages 75, 80, 91–93, 99, 125, 128, 143–146).

[264] S. Abdollahi et al. "*Fermi* Large Area Telescope Fourth Source Catalog". In: *The Astrophysical Journal Supplement Series* 247.1 (2020), p. 33. DOI: `10.3847/1538-4365/ab6bcb` (cited on page 75).

[265] J. Ballet et al. *Fermi Large Area Telescope Fourth Source Catalog Data Release 2*. 2020. DOI: `10.48550/arxiv.2005.11208` (cited on page 75).

[266] S. Sazonov et al. "First tidal disruption events discovered by *SRG*/eROSITA: X-ray/optical properties and X-ray luminosity function at z<0.6". In: *Monthly Notices of the Royal Astronomical Society* 508.3 (2021), pp. 3820–3847. DOI: `10.1093/mnras/stab2843` (cited on page 75).

[267] UK *Swift* Science Data Centre. *Swift XRT Data Products Generator*. `https://www.swift.ac.uk/user_objects`. 2020. URL: `https://www.swift.ac.uk/user_objects` (cited on page 76).

[268] M. Arida and E. Sabol. *HEASARC WebPIMMS*. `https://heasarc.gsfc.nasa.gov/cgi-bin/Tools/w3pimms/w3pimms.pl`. 2020 (cited on page 76).

[269] A. A. Breeveld et al. "An Updated Ultraviolet Calibration for the *Swift*/UVOT". In: *AIP Conference Proceedings*. AIP, 2011. DOI: `10.1063/1.3621807` (cited on page 76).

[270] The CASA Team, B. Bean et al. "CASA, the Common Astronomy Software Applications for Radio Astronomy". In: *Publications of the Astronomical Society of the Pacific* 134.1041 (2022), p. 114501. DOI: `10.1088/1538-3873/ac9642` (cited on page 77).

[271] M. Lacy et al. "The Karl G. Jansky Very Large Array Sky Survey (VLASS). Science Case and Survey Design". In: *Publications of the Astronomical Society of the Pacific* 132.1009 (2020), p. 035001. DOI: `10.1088/1538-3873/ab63eb` (cited on page 77).

[272] K. De et al. "Palomar Gattini-IR: Survey Overview, Data Processing System, On-sky Performance and First Results". In: *Publications of the Astronomical Society of the Pacific* 132.1008 (2020), p. 025001. DOI: `10.1088/1538-3873/ab6069` (cited on page 78).

[273] M. F. Skrutskie et al. "The Two Micron All Sky Survey (2MASS)". In: *The Astronomical Journal* 131.2 (2006), pp. 1163–1183. DOI: `10.1086/498708` (cited on page 78).

[274] C. Y. Peng et al. "Detailed Structural Decomposition of Galaxy Images". In: *The Astronomical Journal* 124.1 (2002), pp. 266–293. DOI: `10.1086/340952` (cited on page 78).

[275] L. Bradley et al. *astropy/photutils: 1.0.0.* Version 1.0.0. Sept. 2020. DOI: `10.5281/zenodo.4044744`. URL: `https://doi.org/10.5281/zenodo.4044744` (cited on page 78).

[276] The Astropy Collaboration, A. M. Price-Whelan et al. "The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package*". In: *The Astrophysical Journal* 935.2 (2022), p. 167. DOI: `10.3847/1538-4357/ac7c74` (cited on page 78).

[277] A. W. Graham and S. P. Driver. "A Concise Reference to (Projected) Sérsic $R^{1/n}$ Quantities, Including Concentration, Profile Slopes, Petrosian Indices, and Kron Magnitudes". In: *Publications of the Astronomical Society of Australia* 22.2 (2005), pp. 118–127. DOI: `10.1071/as05001` (cited on page 78).

[278] F. Masci. *ICORE: Image Co-addition with Optional Resolution Enhancement*. 2013. DOI: `10.48550/arxiv.1301.2718` (cited on page 80).

[279] D. O. Jones et al. *PythonPhot: Simple DAOPHOT-type photometry in Python*. Astrophysics Source Code Library, record ascl:1501.010. Jan. 2015 (cited on page 80).

[280] S. van Velzen et al. "Optical-Ultraviolet Tidal Disruption Events". In: *Space Science Reviews* 216.8 (2020). DOI: `10.1007/s11214-020-00753-z` (cited on page 80).

[281] D. C. Martin et al. "The *Galaxy Evolution Explorer*: A Space Ultraviolet Survey Mission". In: *The Astrophysical Journal* 619.1 (2005), pp. L1–L6. DOI: `10.1086/426387` (cited on page 80).

[282] C. Million et al. "gPhoton: The *GALEX* Photon Data Archive". In: *The Astrophysical Journal* 833.2 (2016), p. 292. DOI: `10.3847/1538-4357/833/2/292` (cited on page 80).

[283] A. Lawrence et al. "The UKIRT Infrared Deep Sky Survey (UKIDSS)". In: *Monthly Notices of the Royal Astronomical Society* 379.4 (2007), pp. 1599–1617. DOI: `10.1111/j.1365-2966.2007.12040.x` (cited on page 80).

[284] B. D. Johnson et al. "Stellar Population Inference with Prospector". In: *The Astrophysical Journal Supplement Series* 254.2 (2021), p. 22. DOI: `10.3847/1538-4365/abef67` (cited on page 81).

[285] C. Conroy and J. E. Gunn. "The Propagation of Uncertainties in Stellar Population Synthesis Modeling. III. Model Calibration, Comparison and Evaluation". In: *The Astrophysical Journal* 712.2 (2010), pp. 833–857. DOI: `10.1088/0004-637x/712/2/833` (cited on page 81).

[286] D. Foreman-Mackey et al. *python-fsps: Python bindings to FSPS (v0.1.1)*. 2014. DOI: `10.5281/zenodo.12157` (cited on page 81).

[287] K. Barbary et al. *SNCosmo*. 2022. DOI: `10.5281/zenodo.592747` (cited on page 81).

[288] M. Newville et al. *lmfit/lmfit-py 1.0.2*. 2021. DOI: `10.5281/zenodo.4516651` (cited on page 82).

[289] K. Barbary. *Extinction V0.3.0*. 2016. DOI: `10.5281/zenodo.804967` (cited on page 82).

[290] D. Calzetti et al. "The Dust Content and Opacity of Actively Star-forming Galaxies". In: *The Astrophysical Journal* 533.2 (2000), pp. 682–695. DOI: `10.1086/308692` (cited on page 82).

[291] G. Schultz and W. Wiemer. "Interstellar reddening and IR-excesses of O and B stars". In: *Astronomy & Astrophysics* 43 (1975), pp. 133–139 (cited on page 82).

[292] C. Rodrigo and E. Solano. "The SVO Filter Profile Service". In: *XIV.0 Scientific Meeting (virtual) of the Spanish Astronomical Society*. July 2020, 182, p. 182 (cited on page 84).

[293] S. Dong et al. "ASASSN-15lh: A highly super-luminous supernova". In: *Science* 351.6270 (2016), pp. 257–260. DOI: `10.1126/science.aac9613` (cited on page 84).

[294] G. Leloudas et al. "The superluminous transient ASASSN-15lh as a tidal disruption event from a Kerr black hole". In: *Nature Astronomy* 1.1 (2016). DOI: `10.1038/s41550-016-0002` (cited on page 84).

[295] S. van Velzen et al. "Discovery of Transient Infrared Emission from Dust Heated by Stellar Tidal Disruption Flares". In: *The Astrophysical Journal* 829.1 (2016), p. 19. DOI: `10.3847/0004-637x/829/1/19` (cited on pages 85–87).

[296] I. Mann and H. Kimura. "Interstellar dust properties derived from mass density, mass distribution, and flux rates in the heliosphere". In: *Journal of Geophysical Research: Space Physics* 105.A5 (2000), pp. 10317–10328. DOI: `10.1029/1999ja900404` (cited on page 87).

[297] A. K. Leroy et al. "The CO-to-H_2 Conversion Factor from Infrared Dust Emission Across the Local Group". In: *The Astrophysical Journal* 737.1 (2011), p. 12. DOI: `10.1088/0004-637x/737/1/12` (cited on page 87).

[298] E. Kankare et al. "A population of highly energetic transient events in the centres of active galaxies". In: *Nature Astronomy* 1.12 (2017), pp. 865–871. DOI: `10.1038/s41550-017-0290-2` (cited on page 88).

[299] N. Jiang et al. "Infrared Echo and Late-stage Rebrightening of Nuclear Transient PS-10adi: Exploring the Torus with Tidal Disruption Events in Active Galactic Nuclei". In: *The Astrophysical Journal* 871.1 (2019), p. 15. DOI: `10.3847/1538-4357/aaf6b2` (cited on pages 88, 89).

[300] S. van Velzen et al. "Seventeen Tidal Disruption Events from the First Half of ZTF Survey Observations: Entering a New Era of Population Studies". In: *The Astrophysical Journal* 908.1 (2021), p. 4. DOI: `10.3847/1538-4357/abc258` (cited on pages 88, 96, 100, 101, 112).

[301] C. Ricci and B. Trakhtenbrot. *Changing-look Active Galactic Nuclei*. 2022. DOI: `10.48550/arxiv.2211.05132` (cited on page 89).

[302] R. Saxton et al. "X-Ray Properties of TDEs". In: *Space Science Reviews* 216.5 (2020). DOI: `10.1007/s11214-020-00708-4` (cited on page 89).

[303] K. M. Leighly. "A Comprehensive Spectral and Variability Study of Narrow-Line Seyfert 1 Galaxies Observed by *ASCA*. II. Spectral Analysis and Correlations". In: *The Astrophysical Journal Supplement Series* 125.2 (1999), pp. 317–348. DOI: `10.1086/313287` (cited on page 89).

[304] S. Vaughan et al. "X-ray spectral complexity in narrow-line Seyfert 1 galaxies". In: *Monthly Notices of the Royal Astronomical Society* 309.1 (1999), pp. 113–124. DOI: `10.1046/j.1365-8711.1999.02811.x` (cited on page 89).

[305] K. Barbary et al. "Discovery of an Unusual Optical Transient with the Hubble Space Telescope". In: *The Astrophysical Journal* 690.2 (2008), pp. 1358–1362. DOI: `10.1088/0004-637x/690/2/1358` (cited on page 89).

[306] B. T. Gänsicke et al. "SCP 06F6: A Carbon-Rich Extragalactic Transient at Redshift $z \simeq 0.14$". In: *The Astrophysical Journal* 697.2 (2009), pp. L129–L132. DOI: `10.1088/0004-637x/697/2/l129` (cited on page 89).

[307] E. C. Kool et al. "AT 2017gbl: A dust obscured TDE candidate in a luminous infrared galaxy". In: *Monthly Notices of the Royal Astronomical Society* 498.2 (2020), pp. 2167–2195. DOI: `10.1093/mnras/staa2351` (cited on page 89).

[308] S. Mattila et al. "A dust-enshrouded tidal disruption event with a resolved radio jet in a galaxy merger". In: *Science* 361.6401 (2018), pp. 482–485. DOI: `10.1126/science.aao4669` (cited on page 89).

[309] S. Reusch. *Multi-Messenger Observations of Tidal Disruption Events*. 2023. DOI: `10.48550/yfj.2307.00902` (cited on page 92).

[310] W. Winter and C. Lunardini. "Interpretation of the Observed Neutrino Emission from Three Tidal Disruption Events". In: *The Astrophysical Journal* 948.1 (2023), p. 42. DOI: `10.3847/1538-4357/acbe9e` (cited on page 92).

[311] D. A. Duev et al. "Real-bogus classification for the Zwicky Transient Facility using deep learning". In: *Monthly Notices of the Royal Astronomical Society* 489.3 (2019), pp. 3582–3590. DOI: `10.1093/mnras/stz2357` (cited on pages 96, 97).

[312] S. van Velzen et al. "The First Tidal Disruption Flare in ZTF: From Photometric Selection to Multi-wavelength Characterization". In: *The Astrophysical Journal* 872.2 (2019), p. 198. DOI: `10.3847/1538-4357/aafe0c` (cited on pages 96, 114).

[313] J. Necker and S. Mechbal. *timewise: v0.3.13 Release*. 2023. DOI: `10.5281/zenodo.8224743` (cited on page 98).

[314] G. Dálya et al. "GLADE: A galaxy catalogue for multimessenger searches in the advanced gravitational-wave detector era". In: *Monthly Notices of the Royal Astronomical Society* 479.2 (2018), pp. 2374–2381. DOI: `10.1093/mnras/sty1703` (cited on page 99).

[315] R. Zhou et al. "The clustering of DESI-like luminous red galaxies using photometric redshifts". In: *Monthly Notices of the Royal Astronomical Society* 501.3 (2020), pp. 3309–3331. DOI: `10.1093/mnras/staa3764` (cited on page 99).

[316] M. Bilicki et al. "Two Micron All Sky Survey Photometric Redshift Catalog: A Comprehensive Three-Dimensional Census of the Whole Sky". In: *The Astrophysical Journal Supplement Series* 210.1 (2013), p. 9. DOI: `10.1088/0067-0049/210/1/9` (cited on page 99).

[317] R. Beck et al. "PS1-STRM: Neural network source classification and photometric redshift catalogue for PS1 3π DR1". In: *Monthly Notices of the Royal Astronomical Society* 500.2 (2020), pp. 1633–1644. DOI: `10.1093/mnras/staa2587` (cited on page 99).

[318] M. M. Kasliwal et al. "The GROWTH Marshal: A Dynamic Science Portal for Time-domain Astronomy". In: *Publications of the Astronomical Society of the Pacific* 131.997 (2019), p. 038003. DOI: `10.1088/1538-3873/aafbc2` (cited on page 99).

[319] M. W. Coughlin et al. “A Data Science Platform to Enable Time-domain Astronomy”. In: *The Astrophysical Journal Supplement Series* 267.2 (2023), p. 31. DOI: `10.3847/1538-4365/acdee1` (cited on page 99).

[320] D. J. Schlegel et al. “Maps of Dust Infrared Emission for Use in Estimation of Reddening and Cosmic Microwave Background Radiation Foregrounds”. In: *The Astrophysical Journal* 500.2 (1998), pp. 525–553. DOI: `10.1086/305772` (cited on page 101).

[321] J. Guy et al. “SALT2: Using Distant Supernovae to Improve the Use of Type Ia Supernovae as Distance Indicators”. In: *Astronomy & Astrophysics* 466.1 (200710.1088/1538-3873/aafbc2), pp. 11–21. DOI: `10.1051/0004-6361:20066930` (cited on page 102).

[322] J. D. Scargle. “Studies in Astronomical Time Series Analysis. V. Bayesian Blocks, a New Method to Analyze Structure in Photon Counting Data”. In: *The Astrophysical Journal* 504.1 (1998), pp. 405–418. DOI: `10.1086/306064` (cited on page 103).

[323] J. D. Scargle et al. “Studies in Astronomical Time Series Analysis. VI. Bayesian Block Representations”. In: *The Astrophysical Journal* 764.2 (2013), p. 167. DOI: `10.1088/0004-637x/764/2/167` (cited on page 103).

[324] S. Reusch and A. Townsend. *ztfparsnip: Release v0.3.0*. 2023. DOI: `10.5281/zenodo.8251614` (cited on page 105).

[325] D. W. Hogg et al. *The K Correction*. 2002. DOI: `10.48550/arxiv.astro-ph/0210394` (cited on page 106).

[326] S. Mechbal et al. *Machine learning applications for the study of AGN physical properties using photometric observations*. 2023. DOI: `10.48550/arxiv.2303.18076` (cited on page 108).

[327] T. Chen and C. Guestrin. “XGBoost”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2016. DOI: `10.1145/2939672.2939785` (cited on page 109).

[328] N. Miranda et al. “SNGuess: A method for the selection of young extragalactic transients”. In: *Astronomy & Astrophysics* 665 (2022), A99. DOI: `10.1051/0004-6361/202243668` (cited on page 124).

[329] IceCube Collaboration, J. Necker. *Search for High-Energy Neutrinos from TDE-like Flares with IceCube*. 2023. DOI: `10.48550/arxiv.2307.15531` (cited on page 125).

[330] IceCube Collaboration, M. G. Aartsen et al. “IceCube-Gen2: The Window to the Extreme Universe”. In: *Journal of Physics G: Nuclear and Particle Physics* 48.6 (2021), p. 060501. DOI: `10.1088/1361-6471/abbd48` (cited on page 128).

[331] The KM3NeT Collaboration, S. Adrián-Martínez et al. “Letter of intent for KM3NeT 2.0”. In: *Journal of Physics G: Nuclear and Particle Physics* 43.8 (2016), p. 084001. DOI: `10.1088/0954-3899/43/8/084001` (cited on page 128).

[332] The KM3NeT Collaboration, S. Aiello et al. “Sensitivity of the KM3NeT/ARCA neutrino telescope to point-like neutrino sources”. In: *Astroparticle Physics* 111 (2019), pp. 100–110. DOI: `10.1016/j.astropartphys.2019.04.002` (cited on page 128).

[333] G. Narayan and ELAsTiCC Team. “The Extended LSST Astronomical Time-series Classification Challenge (ELAsTiCC)”. In: *American Astronomical Society Meeting Abstracts*. Vol. 55. American Astronomical Society Meeting Abstracts. Jan. 2023, 117.01, p. 117.01 (cited on page 129).

[334] E. Blaufuss. “IceCube-190503A – IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 24378 (2019) (cited on page 132).

[335] R. Stein et al. “Optical follow-up of IceCube-190503A with ZTF”. In: *The Astronomer’s Telegram* 12730 (May 2019), p. 1 (cited on page 132).

[336] E. Blaufuss. “IceCube-190619A – IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 24910 (2019) (cited on page 132).

[337] R. Stein et al. “Optical follow-up of IceCube-190619A with ZTF”. In: *The Astronomer’s Telegram* 12879 (June 2019), p. 1 (cited on page 132).

[338] R. Stein. “IceCube-190730A – IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 25225 (2019) (cited on page 132).

[339] R. Stein et al. “A candidate supernova coincident with IceCube-190730A from ZTF”. In: *The Astronomer's Telegram* 12974 (July 2019), p. 1 (cited on page 132).

[340] E. Blaufuss. “IceCube-190922B – IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 25806 (2019) (cited on page 132).

[341] R. Stein et al. “A candidate supernova coincident with IceCube-190922B from ZTF”. In: *The Astronomer's Telegram* 13125 (Sept. 2019), p. 1 (cited on page 132).

[342] R. Stein et al. “IceCube-190922B: Identification of a Candidate Supernova from the Zwicky Transient Facility”. In: *GCN Circular* 25824 (2019) (cited on page 132).

[343] R. Stein et al. “IceCube-191001A: Candidate Counterparts with the Zwicky Transient Facility”. In: *GCN Circular* 25929 (2019) (cited on page 132).

[344] R. Stein. “IceCube-200107A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 26655 (2020) (cited on page 132).

[345] R. Stein and S. Reusch. “IceCube-200107A: No candidates from the Zwicky Transient Facility”. In: *GCN Circular* 26667 (2020) (cited on page 132).

[346] R. Stein. “IceCube-200109A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 26696 (2020) (cited on page 132).

[347] S. Reusch and R. Stein. “IceCube-200109A: Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 26747 (2020) (cited on page 132).

[348] C. Lagunas Gualda. “IceCube-200117A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 26802 (2020) (cited on page 132).

[349] S. Reusch and R. Stein. “IceCube-200117A: Candidate Counterpart from the Zwicky Transient Facility”. In: *GCN Circular* 26813 (2020) (cited on page 132).

[350] S. Reusch and R. Stein. “IceCube-200117A: One Additional Candidate Counterpart from the Zwicky Transient Facility”. In: *GCN Circular* 26816 (2020) (cited on page 132).

[351] C. Lagunas Gualda. “IceCube-200512A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 27719 (2020) (cited on page 132).

[352] S. Reusch et al. “IceCube-200512A: No Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 27721 (2020) (cited on page 132).

[353] M. Santander. “IceCube-200620A – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circular* 27997 (2020) (cited on page 132).

[354] S. Reusch et al. “IceCube-200620A: One Candidate Counterpart from the Zwicky Transient Facility”. In: *GCN Circular* 28005 (2020) (cited on page 132).

[355] E. Blaufuss. “IceCube-200916A – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circular* 28433 (2020) (cited on page 132).

[356] S. Reusch et al. “IceCube-200916A: Two Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 28441 (2020) (cited on page 132).

[357] S. Reusch et al. “IceCube-200916A – ZTF20acaapwk/SN2020tno classified as SN Ia”. In: *GCN Circular* 28465 (2020) (cited on page 132).

[358] C. Lagunas Gualda. “IceCube-200926A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 28504 (2020) (cited on page 132).

[359] S. Reusch et al. “IceCube-200926A: No Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 28520 (2020) (cited on page 132).

[360] C. Lagunas Gualda. “IceCube-200929A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 28532 (2020) (cited on page 132).

[361] S. Weimann et al. “IceCube-200926A: No candidate counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 28551 (2020) (cited on page 132).

[362] M. Santander. “IceCube-201007A – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circular* 28575 (2020) (cited on page 132).

[363] S. Reusch et al. “IceCube-201007A2: No candidate counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 28609 (2020) (cited on page 132).

[364] C. Lagunas Gualda. “IceCube-201130A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 28969 (2020) (cited on page 132).

[365] S. Weimann et al. “IceCube-201130A: No Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 28989 (Nov. 2020), p. 1 (cited on page 132).

[366] C. Lagunas Gualda. “IceCube-201209A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 29012 (2020) (cited on page 132).

[367] S. Reusch et al. “IceCube-201209A: No candidate counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 29031 (2020) (cited on page 132).

[368] E. Blaufuss. “IceCube-201222A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 29120 (2020) (cited on page 132).

[369] R. Stein et al. “IceCube-201222A: No candidate counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 29172 (2020) (cited on page 132).

[370] C. Lagunas Gualda. “IceCube-2010210A: IceCube observation of a high-energy neutrino candidate event”. In: *GCN Circular* 29454 (2021) (cited on page 132).

[371] S. Reusch et al. “IceCube-210210A: Two Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 29461 (2021) (cited on page 132).

[372] M. Santander. “IceCube-210510A – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circular* 29976 (2021) (cited on page 132).

[373] R. Stein et al. “IceCube-210510A: No candidate counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 29999 (2021) (cited on page 132).

[374] M. Santander. “IceCube-210811A - IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circular* 30627 (2021) (cited on page 132).

[375] R. Stein et al. “IceCube-210811A: No Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 30644 (2021) (cited on page 132).

[376] M. Lincetto. “IceCube-210922A – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circular* 30862 (2021) (cited on page 132).

[377] S. Weimann et al. “IceCube-210922A: No Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 30870 (2021) (cited on page 132).

[378] M. Lincetto. “IceCube-220405A – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circula* 31833 (2022) (cited on page 132).

[379] S. Reusch et al. “IceCube-220405A: No Candidate Counterparts from the Zwicky Transient Facility”. In: *GCN Circular* 31843 (2022) (cited on page 132).

[380] E. Blaufuss. “IceCube-220405B – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circular* 31839 (2022) (cited on page 132).

[381] M. Santander. “IceCube-220501A - IceCube observation of a high-energy neutrino candidate track-like event”. In: *GRB Coordinates Network* 31986 (2022) (cited on page 132).

[382] R. Stein et al. “IceCube-220501A: No Candidate Counterparts from the Zwicky Transient Facility”. In: *GRB Coordinates Network* 31989 (2022) (cited on page 132).

[383] M. Santander. “IceCube-220513A – IceCube observation of a high-energy neutrino candidate track-like event”. In: *GCN Circular* 32037 (2022) (cited on page 132).

[384] S. Weimann et al. "IceCube-220513A: No Candidate Counterparts from the Zwicky Transient Facility". In: *GRB Coordinates Network* 32040 (2022) (cited on page 132).

[385] M. Santander. "IceCube-220624A – IceCube observation of a high-energy neutrino candidate track-like event". In: *GRB Coordinates Network* 32260 (2022) (cited on page 132).

[386] S. Reusch et al. "IceCube-220624A: No Candidate Counterparts from the Zwicky Transient Facility". In: *GRB Coordinates Network* 32270 (2022) (cited on page 132).

[387] M. Lincetto. "IceCube-220822A – IceCube observation of a high-energy neutrino candidate track-like event". In: *GRB Coordinates Network* 32475 (2022) (cited on page 132).

[388] J. Necker et al. "IceCube-220822A: One Candidate Counterpart from the Zwicky Transient Facility". In: *GRB Coordinates Network* 32480 (2022) (cited on page 132).

[389] R. Stein et al. "IceCube-220907A: Classification of AT2022oyn as a Type Ia supernova." In: *GCN Circular* 32573 (Jan. 2022), p. 1 (cited on page 132).

[390] M. Lincetto. "IceCube-221216A – IceCube observation of a high-energy neutrino candidate track-like event". In: *GRB Coordinates Network* 33065 (2022) (cited on page 132).

[391] S. Weimann et al. "IceCube-221216A: No Candidate Counterparts from the Zwicky Transient Facility". In: *GRB Coordinates Network* 33089 (2022) (cited on page 132).

[392] M. Santander. "IceCube-221223A – IceCube observation of a high-energy neutrino candidate track-like event". In: *GRB Coordinates Network* 33094 (2022) (cited on page 132).

[393] S. Reusch et al. "IceCube-221223A: No Candidate Counterparts from the Zwicky Transient Facility". In: *GRB Coordinates Network* 33134 (2023) (cited on page 132).

[394] M. Lincetto. "IceCube-230112A – IceCube observation of a high-energy neutrino candidate track-like event". In: *GRB Coordinates Network* 33161 (2023) (cited on page 132).

[395] J. Necker et al. "IceCube-230112A: No Candidate Counterparts from the Zwicky Transient Facility". In: *GRB Coordinates Network* 33174 (2023) (cited on page 132).

[396] S. Frederick et al. "A New Class of Changing-look LINERs". In: *The Astrophysical Journal* 883.1 (2019), p. 31. doi: `10.3847/1538-4357/ab3a38` (cited on page 143).

[397] S. Huang et al. "AT2018dyk revisited: A tidal disruption event candidate with prominent infrared echo and delayed X-ray emission in a LINER galaxy". In: *Monthly Notices of the Royal Astronomical Society* 525.3 (2023), pp. 4057–4064. doi: `10.1093/mnras/stad2541` (cited on page 143).

[398] Y. Yang et al. "Tidal Disruption on Stellar-mass Black Holes in Active Galactic Nuclei". In: *The Astrophysical Journal Letters* 933.2 (2022), p. L28. doi: `10.3847/2041-8213/ac7c0b` (cited on page 143).

[399] M. Thévenot et al. "Supernovae detected in NEOWISE-R 2021". In: *Transient Name Server AstroNote* 212 (Aug. 2021), pp. 1–212 (cited on pages 145, 146).

[400] J. T. Hinkle. *Mid-Infrared Echoes of Ambiguous Nuclear Transients Reveal High Dust Covering Fractions: Evidence for Dusty Tori*. 2022. doi: `10.48550/arxiv.2210.15681` (cited on page 146).

Selbständigkeitserklärung

Ich erkläre, dass ich diese Dissertation selbständig und nur unter Verwendung der von mir gemäß § 7 Abs. 3 der Promotionsordnung der Mathematisch-Naturwissenschaftlichen Fakultät – veröffentlicht im Amtlichen Mitteilungsblatt der Humboldt-Universität zu Berlin Nr. 42/2018 am 11.07.2018 – angegebenen Hilfsmittel angefertigt habe.

Simeon Reusch Berlin, den 14. Dezember 2023