



THE UNIVERSITY

of ADELAIDE

Improved Neutrino Point Source Search Method for
IceCube's High Energy Starting Event Sample

Mark Gerald Aartsen
B.Sc. (Honours), Physics

A thesis submitted to the University of Adelaide in fulfilment
of the requirements for the degree of Doctor of Philosophy.

School of Physical Sciences
Department of Physics
December 2017

Contents

1	Introduction	1
1.1	Cosmic rays	1
1.2	Summary	7
2	Neutrinos and air shower particles	9
2.1	Neutrino production	9
2.2	Background particles from cosmic ray air showers	11
2.2.1	Atmospheric neutrinos from charged pions	13
2.2.2	Atmospheric neutrinos from kaons	16
2.2.3	Muons from cosmic ray air showers	18
2.2.4	Prompt atmospheric neutrinos from charmed mesons	20
2.2.5	Accompanying muons	21
2.3	Interaction at Earth	22
2.4	Propagation to the detector	25
2.4.1	Attenuation through the Earth	25
2.4.2	Neutrino oscillations	26
2.5	Summary	29
3	The IceCube Neutrino Observatory	30
3.1	The detector	32
3.2	Ice properties	33
3.3	Digital Optical Modules (DOMs)	35
3.4	Waveform unfolding	37

3.5	Reconstruction techniques	38
3.6	Summary	45
4	Neutrino contained event analyses	47
4.1	HESE filter	47
4.2	Event selection	51
4.3	Modelling the neutrino flux in IceCube	51
4.4	Current HESE results	55
4.4.1	Fit for a diffuse neutrino flux	55
4.4.2	Fit for point sources of neutrinos	59
4.5	Summary	61
5	Combined diffuse and point source analysis	63
5.1	Event reconstructions	65
5.2	Combined diffuse and point source likelihood	66
5.2.1	Derivation of unbinned Poisson likelihood	66
5.2.2	Application of the unbinned Poisson likelihood to Diffuse+PS analysis	68
5.2.3	Simplifying the Diffuse+PS likelihood	70
5.2.4	Comparing the Diffuse+PS likelihood to the previous point source likelihood	71
5.2.5	Gaussian priors	71
5.2.6	Test statistic	73
5.2.7	Probability of the test statistic under scrambling	74
5.3	Summary	75
6	Event selection, processing and generation	81
6.1	Event selection	82
6.2	Generation of randomly sampled skies	82
6.2.1	Scrambling method and sampling method	83

6.2.2	Resampling of event maps	84
6.2.3	Calculating rotation	85
6.3	Multinomial sampling and weighting of event sets	88
6.4	Example of a sampled sky	91
6.5	Summary	92
7	Statistical comparisons of the analyses	94
7.1	Direct comparisons between analyses	94
7.2	Definitions of power and significance	97
7.2.1	Power for a fixed direction	100
7.2.2	All-sky power	101
7.2.3	Scrambling and sampled maps	101
7.3	Investigations into bias	104
7.3.1	Tests of power for a fixed direction	106
7.3.2	Fixed direction tests of non-zero signal hypotheses	108
7.3.3	Tests of all-sky power	112
7.4	Reconciling these results	113
7.5	Diffuse+PS likelihood power on HESE sample	117
7.6	Summary	123
8	Sensitivity and discovery potential of the Diffuse+PS analysis	125
8.1	Definitions	125
8.2	Sensitivity and discovery potential towards a fixed direction	128
8.3	All-sky sensitivity and discovery potential	134
8.4	Summary	138
9	Results of the Diffuse+PS analysis	139
9.1	Results	139
10	Conclusion	150

Abstract

High Energy Neutrino Astronomy is a new field that has the potential to solve the mystery of the origin of cosmic ray particles, the highest energy particles that have been observed. The IceCube Neutrino Observatory has recently discovered the existence of a high energy astrophysical neutrino flux. We are able to calculate for our “track-like” sub-population of events a directional origin to within a square degree, yet it remains difficult to establish whether these neutrinos come from bright distinguishable sources or a sea of individual weak neutrino producers. The current goal of neutrino point source analyses is the observation of a bright neutrino source above expected background rates of neutrinos, in our case both from a uniform distribution of astrophysical neutrino sources and background neutrino events produced in cosmic ray interactions with our own atmosphere. This thesis presents a new form of point source analysis that tests the data under the model of the combination of multiple background hypotheses and a single point source hypothesis, where the predicted background distributions can be produced through simulation. We apply this analysis to the High Energy Starting Event sample produced by IceCube, and show how this analysis differs to previous analyses on these events. We find that the fit using a background in the likelihood that does not match the distribution of the events can result in a bias in the fitted strength of a point source, and that the relative power of the analysis compared to the previous point source analysis can depend on the point source location, where the power is seen to be comparable in the southern sky but improved in the northern sky. Our results from applying our new form of point source analysis to the data do not show strong evidence for a point source hypothesis, with p-values of 0.468 for all events in our sample and 0.922 for a subset of shower topology events in our sample.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

Mark Aartsen

June 2018

Acknowledgements

First and foremost, I would like to thank my supervisors Gary Hill, Ray Protheroe and Ben Whelan, whose support and guidance made this thesis possible, and without whom I could not have accomplished this work. I would also like to thank Sally, Rebecca and Elise for their continuing friendship, and all the Adelaide IceCubers including Alex K, Alex W and Natasha during my time at the University of Adelaide who all made the IceCube group feel like a team. I would like to thank Bruce, Roger and Gavin for their help and wisdom throughout my university education, all of you make a huge difference in the lives of aspiring physicists. I would also like to thank all of the High Energy Astrophysics Group at Adelaide, past and present, for making our group an incredible place where I felt truly welcome.

I would like to thank the world of IceCube outside Adelaide who, despite distance, managed to impart to me what I now know about this wonderful experiment, in particular Naoko, Chad, Claudio, Nathan, Jakob, Jake, Chris, Mike and Markus but also many, many others. In having such great collaborators you continue to motivate me to do better. A big thank you also to my new collaborators at the University of Canterbury who have allowed me to continue on this research path.

It goes without saying that my parents Liz and Gerald deserve a substantial amount of credit for this thesis, this work is a culmination of the path they opened for me and their endless encouragement and guidance have shaped me, so my infinite thanks go to them. My first steps on this path were walked with my Pake and Beppe: Harry and Martha, with all of the math books you provided; you both always made me feel proud of myself. A huge thanks also goes to Scott, I couldn't have asked for

a better brother. To the rest of my extended family which continues to grow, thank you all for your love and support.

A big thank you to my friends as well, especially Solly, is more than merited, you appear to have managed the exceptionally difficult tasks of keeping my head balanced and my chin up.

Emilie, my partner to the end of all things, I would like to thank most of all, for sharing all our adventures together and the burden of a PhD I struggled to carry alone.

I would like to acknowledge and dedicate my thesis to the late Associate Professor Raymond Protheroe, who supervised me through my honours year and introduced me to the field of neutrino astrophysics, and then introduced me to the IceCube detector for my PhD where he remained my supervisor. Sadly the decline of Ray's health led to him attending the University less and less throughout my PhD, and so our contact diminished over this time, and I regret not maintaining a stronger connection before his passing. For the time that Ray was able to supervise me, he remained a positive, friendly and insightful mentor, and those insights showed the depth of his knowledge, curiosity and love of physics. He remains dearly missed by the many who knew him, and my thoughts continue to go to Ray and his family.

Author's Note

This work was performed as a member of the IceCube Collaboration. Therefore it is important to note the contributions from other individuals to this work, and to show where contributions to the work were made by the author, and highlight the key components of the thesis.

The background material to the work of this thesis is presented by the author in Chapters 1 to 4. Chapters 1 and 2 summarise previous knowledge of cosmic rays and neutrinos, where Chapters 3 and 4 introduce the IceCube detector and the High Energy Starting Event selection. The reconstructions described in Chapter 3 and the event selection and analysis of Chapter 4 are the previous works of other members in the collaboration, which have been published separately.

The work of Chapters 5, 6, 7, 8 and 9 are the key chapters of this thesis and are the work of the author. The work used the existing reconstructed event point spread functions, and simulated event files which were used to produce background rate predictions, which were supplied by the collaboration. This work used new C++ programs to perform the new analysis of the data, written in collaboration with Ben Whelan.

Outline of this Thesis

This thesis concerns the detection of neutrinos of potentially astrophysical origin, and the characterisation of such an origin. Specifically, we wish to determine if a point source origin for a subset of these neutrinos is apparent when using all available information to characterise the various potential backgrounds which the events can originate from. As of yet, we have insufficient evidence to claim a discovery of a singular point source origin above expected statistical fluctuations of the diffuse backgrounds.

This work was conducted using the IceCube Neutrino Observatory, specifically looking at the High Energy Starting Event (HESE) catalogue. The work done within this thesis was performed as a member of the IceCube Collaboration.

In Chapter 1 the background for cosmic ray physics is discussed and the motivation for neutrino astronomy is outlined. In Chapter 2 the properties of the neutrino are discussed primarily concerning the mechanisms of its creation at various production sites, propagation to the detector and interaction within IceCube. The production of neutrinos and muons in air showers is also discussed in detail. In Chapter 3 the IceCube detector is introduced in addition to the means of determining neutrino properties with the detector. In Chapter 4 the event sample is introduced and the previous work on the sample is summarised. In Chapter 5 a new analysis technique for determining a point source is defined. In Chapter 6 we define event manipulations for the purpose of generating new pseudo-data sets. In Chapter 7 we compare the new technique to the old to assess biases found to emerge, and show how the sensitivity and discovery potential of the two techniques differ in

Chapter 8. In Chapter 9 the result of the new methodology on the event sample is shown, with concluding remarks in Chapter 10.

Chapter 1

Introduction

1.1 Cosmic rays

At the time of writing, it is currently 105 years since the balloon flights of Victor Hess in 1912 and the discovery of cosmic rays. Using several electroscopes which was the then-current means of measuring background radiation, Hess observed that the electroscopes had increasing readings at higher altitude, and therefore concluded that the radiation was extraterrestrial in nature, given his thorough studies of terrestrial radiation and the height to which it could penetrate, and his studies of the atmosphere and the material of the electroscope ruling out any atmospheric or instrumental origin [1]. By performing some flights at night and during an eclipse, Hess also concluded that the radiation was unlikely to be of solar origin [2], and this is true for cosmic rays at high energies, though the grounds for his conclusions were incorrect since solar cosmic rays arrive at the Earth isotropically [1]. Solar cosmic rays are relatively few, these particles are low energy and are associated with violent solar activity and hence variable in time [3].

In the time since Hess' discovery, numerous experiments have measured the cosmic ray spectrum and composition. We know that cosmic rays are primarily charged protons, and also heavier nuclei up to iron, and the relative abundances of these




Figure 1.1: The spectrum of cosmic rays as observed by multiple experiments over many decades in energy. Figure from [9].

nuclei change with energy. For cosmic rays from giga-electronvolts (GeV) up to peta-electronvolts (PeV), space- and balloon-based detectors (for example the Alpha Magnetic Spectrometer, designated AMS-02 aboard the International Space Station (ISS) [4], and the Cosmic Ray Energetics And Mass (CREAM) experiment, a balloon-based experiment flown in Antarctica [5]) measure cosmic ray particles directly with a variety of instruments to cover a large dynamic range, whereas at higher energies of up to zetta-electronvolts (ZeV) larger detectors (such as the Pierre Auger Observatory (PAO) [6] and the Telescope Array [7] [8]) measure the resulting particle cascade from the interactions of cosmic rays in the atmosphere.

The spectrum of cosmic rays is shown in Figure 1.1. The spectrum follows a power law with only small deviations. The first deviation, known as the knee, is an apparent steepening of the spectrum believed to be associated with the limitations of supernova shock acceleration (see Figure 1.2). The inability for an object to retain cosmic rays above a particular energy is due to the object's size and magnetic field strength; in approximate terms to retain cosmic rays the region in which the particles are accelerated must be of similar scale to the Larmor radius [10].

$$(r_L/1\text{pc}) = 1.08 \frac{(E/1\text{PeV})}{Z(B/1\mu\text{G})} \quad (1.1)$$

for a particle of charge Ze and energy E in a magnetic field \vec{B} normal to the particle's velocity. Equation 1.1 shows that heavier nuclei with greater charge Ze can be more effectively contained by a given source compared to particles with less charge such as protons with charge e . Therefore it is theorised that the knee shows a change in composition to progressively heavier nuclei from galactic sources.

The second deviation, known as the ankle, is a hardening of the spectrum suspected to be a transition toward extragalactic sources of cosmic rays. This conclusion is drawn from the acceleration limits of known sources in the galaxy, and that the galactic magnetic field itself is incapable of containing the highest energy cosmic rays. For the highest observed cosmic rays with energies $E > 10^{20}$ eV, the effect of the regular component of the galactic magnetic field of the Milky Way has been simulated showing that at $E/Z = 10^{20}$ eV, the deflection is of order 3° , and dependent on the path through the galaxy.

The observations of cosmic rays at the highest energies show a primarily isotropic distribution [12] [13], furthermore, over-densities appear away from the galactic plane. Therefore it is reasonable to conclude that it is highly likely the highest energy cosmic rays are of extragalactic origin. The deflection of protons over 50 Mpc through intergalactic magnetic fields with a correlation length less than 1 Mpc is less than 2° . If the sources of cosmic rays are nearby then the directions of the highest




Figure 1.2: Hillas plot showing the limitation of various sources to contain cosmic rays as they are accelerated based on their size and magnetic field strength. Supernova remnants (SNR) are not able to contain 10^{20} eV cosmic rays, and therefore their contribution to the cosmic ray spectrum cuts off at lower energies. Different capacities are shown for proton and iron nuclei, with iron nuclei having the potential to be accelerated to higher energies at the same sites. Figure sourced from [11], which is an updated version of the plot from [10].


energy cosmic rays are therefore indicative of the source positions. Attempts to find the origins of cosmic rays at these energies have yet to confirm any particular sources as origins. At present, a large scale anisotropy away from the galactic plane has been detected by the Pierre Auger collaboration using harmonic analysis at a significance of 5.2σ [14]. In addition, correlations with possible source candidates using the data from the Pierre Auger observatory show a post-trial p-value of 1.4% for Centaurus A and 1.3% with the selected Swift AGN catalogue [12], which while interesting is not sufficient to claim a related anisotropy. The TA observatory has claimed an anisotropy with a significance of 3.4σ for the most over-dense region on the sky, this region however has no obvious counterpart [13].

While some fraction of cosmic rays will escape their source regions, others will interact. Cosmic ray protons and nuclei interacting with other nuclei or photon fields will result in the production of mesons, predominantly charged and neutral pions, which decay to produce high energy neutrinos and gamma rays respectively. Having no charge, these gamma rays and neutrinos are not deflected by magnetic fields and therefore the direction toward their origin is retained when the particles are detected at Earth.

Observational evidence of cosmic ray acceleration in our galaxy has recently been shown through the detection of gamma rays; the signature “pion bump” of gamma rays due to neutral pion decay (see Equation 2.7) has been observed in two supernova remnants within our galaxy [15] as the observed gamma ray spectra from approximately 100 MeV to 100 GeV the two sources produce match the predicted spectra for pion decay and are inconsistent with leptonic origins. Future gamma ray telescopes may also be able to probe higher energy acceleration by observing gamma rays beyond our current observations to the range of 100 TeV, as a leptonic origin for a hard spectrum of gamma rays at such energies is disfavoured due to the softer spectrum of synchrotron and/or Inverse Compton (IC) cooling processes [16].

The range of cosmic rays is affected by local magnetic fields at lower energies, but at the highest energies pion photoproduction interactions of cosmic rays on the

Cosmic Microwave Background (CMB) photons results in the depletion of the high energy proton flux and the subsequent production of pions [9]:



where N is as given nucleon (in our case usually a proton), Δ is a short lived “Delta-resonance” which goes on to decay to both another nucleon and an outgoing pi-meson (π) (charge is omitted in Equation 1.2 as many different interactions of this type are possible). This process was theorised independently by Greisen [17] and Zatsepin & Kuz'min [18] and hence called the “GZK effect”. In the interaction approximately 20% of the cosmic ray energy is on average found in the nucleon, and 80% of the energy in the pion. This in addition to the density of CMB photons results in an effective “GZK Horizon” for cosmic ray protons with energies greater than $10^{19.5}$ eV of tens of mega-parsecs (Mpc). Heavier nuclei will undergo photo-disintegration on the CMB at much lower energies [19]. The fragments from photo-disintegration are lower energy nuclei that also constitute a small cosmic ray flux. The neutral and charged pions from the GZK interaction will decay to produce observable gamma rays and neutrinos respectively (See Equations 2.7 and 2.2). The steepening of the spectrum observed at the highest cosmic ray energies is indicative of the GZK effect, however it is also possible that this decrease in flux is in part due to the intrinsic limits of cosmic ray accelerators.

The range of gamma rays is limited by attenuation on background photon fields (e.g. local infra-red fields, radio backgrounds and the Cosmic Microwave Background) by photon-photon interactions [9]



The range of photons is shown in Figure 1.3.




Figure 1.3: The range of photons as a function of redshift. The shaded region shows the distance to sources that are invisible at the corresponding photon energies. At PeV energies the range of gamma rays is shorter than the distance to the Galactic Centre. At energies beyond this scale (> 10 PeV) the range increases, however no gamma rays have yet been observed at these energies. Figure from [20].

1.2 Summary

The study of cosmic rays covers a multitude of particles over many orders of magnitude, and the spectrum and composition can be inferred by the detection of these protons and nuclei at Earth using a variety of different instruments. However due to deflection by magnetic fields, the individual sources of cosmic rays are hard to determine from the detection of the cosmic rays themselves. The secondary particles from cosmic ray interactions, gamma rays and neutrinos, can be used to determine the source origins and their properties. While gamma rays can be used to test for cosmic ray acceleration this is not a trivial task, as gamma rays can be produced through both hadronic and leptonic processes, where we are most interested in hadronic processes, and in addition the range of gamma rays is limited with increasing photon energy.

In Chapter 2, we look at the properties of neutrinos as they relate to the ob-

servations of cosmic ray sources. Comparing neutrinos to other high energy cosmic messengers reveals the inherent advantages of using neutrinos for high energy astronomy: unlike cosmic ray nuclei, neutrinos trace back to their sources without deflection, unlike photons, neutrinos serve as an unambiguous tracer of hadronic acceleration sites, and uniquely neutrinos lack an observational horizon at all energies. A point source search for neutrinos therefore has the potential to definitively locate and characterise the dominant sources of cosmic rays throughout the universe which, after a century, remain unknown.

Chapter 2

Neutrinos and air shower particles

The neutrino was originally hypothesised by Wolfgang Pauli as a means of conserving momentum and energy in nuclear beta decay [21]. Pauli concluded that the particle would be of spin 1/2, electrically neutral, and of mass no greater than 0.01 times the mass of the proton. Current experiments put limits on the mass of the electron neutrino at $m_{\nu_e}c^2 < 2.05$ eV [22]. Neutrinos are in fact a family of leptons that are partners to the electron, muon and tau particles. Limits from cosmological data put tighter constraints on the mass of all flavours of neutrinos with results in the region $\sum m_\nu < 1$ eV [23].

Neutrinos interact rarely lacking electromagnetic charge and colour charge, interacting only via the weak force which is mediated through W^\pm and Z bosons. The interactions of neutrinos with matter as a means of detection are discussed in Section 2.3. This property of a rarely interacting particle gives us the ability to trace cosmic ray interactions beyond the effective horizons of both cosmic rays and gamma rays.


2.1 Neutrino production

Using neutrinos to perform astronomy has been discussed for over a half-century. Early publications by K. Greisen [24] and F. Reines [25] suggest the creation of

kilo-ton neutrino detectors for cosmic neutrinos and atmospheric neutrinos. The possibilities of neutrino astronomy are also summarised in J. Bahcall’s early paper [26] which discusses the possibility of using solar neutrinos to study the structure of the sun and details of the ongoing fusion reactions, but goes on to discuss the possibility of detecting neutrinos from proton-proton interactions in the atmosphere (atmospheric neutrinos) and also from the sites of cosmic ray acceleration (astrophysical neutrinos), primarily from “Strong Radio Sources”.

These strong radio sources we would now consider to include Active Galactic Nuclei and their relativistic jets which have the necessary magnetic fields and energy budgets to be potential cosmic ray acceleration sites, and therefore astrophysical neutrino production sites. Becker [27] provides an excellent summary of cosmic ray sources and their energy budgets.

The underlying motivation of neutrino astronomy is that the same proton-proton interactions in our atmosphere [26]



can be the same mechanism responsible for the production of neutrinos at cosmic ray sites, in addition to sources where $p\gamma$ interactions are also a possibility. There is a difference in the physics of atmospheric neutrino production compared to astrophysical neutrino production as the interacting cloud or other matter target is likely to be far less dense than our atmosphere, and therefore the mesons and resulting muons go on to decay without re-interaction losses. As a result, charged pi mesons (π^\pm) are considered the dominant channel for astrophysical neutrino production as they are the most readily produced mesons in pp collisions.

Charged pions decay via the following decay channel:

$$\pi^\pm \rightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu)
 \tag{2.2}$$

for over 99.98% of all decays. The decay of the resulting muon

$$\mu^\pm \rightarrow \nu_e(\bar{\nu}_e) + e^\pm + \bar{\nu}_\mu(\nu_\mu) \quad (2.3)$$

results in a production flavour ratio of $(\nu_e : \nu_\mu : \nu_\tau) = (\bar{\nu}_e : \bar{\nu}_\mu : \bar{\nu}_\tau) = (1 : 2 : 0)$ [27] assuming π^+ and π^- are produced in equal amounts and the source is not sufficiently dense that there are “muon-damping effects” [28] [29]. The flavour ratio $(1 : 2 : 0)$ is however not what we expect to observe at Earth due to the cosmological distances involved; neutrino oscillations (see Section 2.4.2) will change the flavour ratio $(1 : 2 : 0)$ to a mixed state of $(1 : 1 : 1)$. The energy of the resulting neutrinos is dependent on the energies of pions from $p + p$ interactions, typically $E_\pi/E_p \sim 1/5$. From the kinematics of pion and muon decay we expect roughly equal energy to be distributed among the final leptons, so neutrinos are typically of energy $E_\nu/E_p \sim 1/20$.

Precursors to IceCube, such as the Antarctic Muon And Neutrino Detector Array (AMANDA), set limits but were unable to confirm the presence of a high energy astrophysical neutrino flux at the energies associated with cosmic ray acceleration [30]. The eventual discovery of an astrophysical flux of neutrinos came from the gigaton scale IceCube Neutrino Observatory discussed in Chapter 3, using the same High Energy Starting Event (HESE) sample under investigation in this thesis.

2.2 Background particles from cosmic ray air showers

For the HESE analysis and the majority of astrophysical neutrino searches in IceCube, the background events come solely from the interaction of cosmic rays within Earth’s atmosphere.

Cosmic ray interactions result in numerous daughter particles, many of which are unstable with short decay times. One must take special relativity into account when determining the effective decay time for an observer on the surface of the earth

or in the atmosphere, as the time of decay for a particle at rest in the frame of an observer differs to that of a frame where the particle is in motion, due to relativistic time dilation. The effective decay time t of a particle is

$$t = \gamma\tau \quad (2.4)$$

where τ is the mean decay time for the particle at rest, and γ is the Lorentz factor

$$\gamma = \frac{1}{\sqrt{1 - \beta^2}} \quad (2.5)$$

where $\beta = v/c$ is the particle speed relative to the speed of light in a vacuum. The Lorentz factor can be used to relate total particle energy E to particle mass m :

$$E = \gamma mc^2 \quad (2.6)$$

In air showers we often consider particles with kinetic energies much greater than the particle rest mass energy, resulting in a large Lorentz factor and therefore substantial time dilation.

In the same vein of Heitler's model of electromagnetic air showers [31], a simple model of the interaction of cosmic rays in the atmosphere [32] divides the energy of the cosmic ray primary into the production of many pions; two-thirds of these pions will be charged (π^\pm) and one-third will be neutral (π^0). Neutral pions (π^0) have a mean lifetime 8.52×10^{-17} s [33]. This is short enough that for all but the most energetic cosmic rays, neutral pions are considered to decay rapidly and produce gamma rays:

$$\begin{aligned} \pi^0 &\rightarrow \gamma + \gamma \quad (98.82\%) \\ \pi^0 &\rightarrow e^+ + e^- + \gamma \quad (1.17\%) \end{aligned} \quad (2.7)$$

These neutral pions are the primary contribution to the electromagnetic component of cosmic ray air showers [34] with the charged hadrons feeding the hadronic com-

ponent. Both models share key features: the number of air shower particles grow at each interaction length until a critical energy where the particles that make up the shower are no longer sustained: in the case of electromagnetic showers where particle energies are too low to initiate pair production or bremsstrahlung; for hadronic showers the point at which it is more probable for the pion to decay than re-interact with air nuclei.

Neither the electrons nor gamma rays from neutral pion decays are able to penetrate into the deep ice to produce light in IceCube. What is observed is the muons and neutrinos resulting from the decay of charged pions and more massive mesons produced in air showers.

These simple models show a direct means to arrive at air shower details of particle number as a function of shower depth and describe the development of air showers, however the precise details of atmospheric neutrinos and accompanying muons require extra consideration (for example, if charged pions only decayed below their critical energy, there would be no neutrinos produced from charged pions with energies greater than the critical energy of charged pions deep in the atmosphere which is approximately 20 GeV). As air shower physics cannot be treated purely analytically, understanding the fine details of air showers generally requires Monte Carlo sampling methods such as with the air-shower simulation program CORSIKA [35]. A detailed treatment of the atmospheric neutrino flux can be found in papers by Honda et al. [36] and Enberg et al. [37]. Here we detail an overview of the generation of TeV energy neutrinos in atmospheric air showers capable of interacting in IceCube.

2.2.1 Atmospheric neutrinos from charged pions

The first contribution to the flux of atmospheric muon neutrinos is from the decay of charged pions (π^\pm) (see Equation 2.2). The pion mass is comparable to that of the muon mass, so in the rest frame of the pion the majority of the pion energy is given to the muon rather than the neutrino, as this is the general rule for most

two-body decays [38]. The resulting kinematic transform to the lab frame gives the results

$$\langle E_\mu \rangle / E_\pi = 0.79 \quad (2.8)$$

and

$$\langle E_\nu \rangle / E_\pi = 0.21 \quad (2.9)$$

In addition, charged pions have a longer mean lifetime of $\tau = 2.6 \times 10^{-8}$ s compared to the rapidly decaying π^0 , which means we need to consider the probability of re-interaction with particles in the atmosphere in addition to decay.

To describe the flux of neutrinos from pions and other mesons it is necessary to discuss how pions and other mesons behave in the atmosphere. Hadronic fluxes such as pions in air showers approximately follow the form [38]

$$\frac{dN_i(E, X)}{dX} = - \left(\frac{1}{\lambda_i} + \frac{1}{d_i} \right) N_i(E, X) + \sum_j \int \frac{F_{ji}(E_i, E_j)}{E_i} \frac{N_j(E_j)}{\lambda_j} dE_j \quad (2.10)$$

where N_i is the number of hadrons of type i , E is the energy of the particles and X is the slant depth which measures the depth of the shower from the top of the atmosphere in g/cm^2 (for a vertically inclined shower, X is $0 \text{ g}/\text{cm}^2$ at the top of the atmosphere and X is around $1030 \text{ g}/\text{cm}^2$ at sea level). The first term $-(1/\lambda_i + 1/d_i)N_i(E, X)$ shows the rate of loss of the hadronic species i to decay and re-interaction while the second term $\sum_j \int [F_{ji}(E_i, E_j)/E_i][N_j(E_j)/\lambda_j]dE_j$ describes the addition of particles i at energy E_i from interactions of all parent particles in the shower j . The pion development in the vein of the Heitler model [32] can be thought of as a simplification of this, for instance both the initial proton (p) and charged pions (π^\pm) interacting to generate lower energy charged pions (π^\pm) and neutral pions (π^0). We wish to examine the left term with its values of d_i for the decay length and λ_i for the interaction length, both in units of g/cm^2 . From the equation, it is clear that whichever length is shorter will be the method most responsible for the disappearance of the particle.

We express the decay length d_i and interaction length λ_i in units of g/cm^2 . As a consequence of this the decay length is dependent on the critical energy [38] of a particle

$$\epsilon_i = \frac{m_i c^2}{c\tau_i} h_0 \quad (2.11)$$

where τ_i is the mean decay time, $m_i c^2$ is the rest mass of the particle and h_0 is the scale height that describes the change in vertical atmospheric depth (the column density from the top of the atmosphere measured in g/cm^2). The critical energy ϵ is used to determine the decay length:

$$d_i = \frac{E_i X \cos \theta}{\epsilon_i} \quad (2.12)$$

where E_i is the energy of the particle, X is the slant depth, and θ is the angle of the shower axis from zenith. We are most interested in interaction and decay within the first few interaction lengths where $X < 200 \text{ g}/\text{cm}^2$. Typical interaction lengths are of order $\lambda_i \sim 100 \text{ g}/\text{cm}^2$ for charged pions and kaons [3]. Due to the similarity of λ_i and X in this case, the critical energy can be used as a good reference of how an unstable particle will behave in an air shower: for energies sufficiently below the critical energy, decay is the more probable outcome; sufficiently above the critical energy, re-interaction with air nuclei is more probable. For $X < 200 \text{ g}/\text{cm}^2$ we find $h_0 \approx 6.4 \text{ km}$, and so we obtain a critical energy for charged pions π^\pm of $\epsilon_{\pi^\pm} = 115 \text{ GeV}$. Correspondingly, the critical energy of neutral pions is $\epsilon_{\pi^0} = 3.5 \times 10^{10} \text{ GeV}$, and so for the shower energies relevant to IceCube, neutral pions can be considered as decaying before any re-interaction.

With a critical energy of 115 GeV, at TeV energies and above that are of interest to our analysis we find that most charged pions will re-interact with the atmosphere rather than decay. The re-interaction of the charged pion with target nuclei in the atmosphere is similar to the interaction of the cosmic ray in that they are both interactions via the strong nuclear force, both leading to the production of more pions both neutral and charged. The neutral pions produced will decay and add

to the electromagnetic shower. After several generations, the charged pions will be of sufficiently low energy such that they are more likely to decay and in doing so produce neutrinos. The re-interaction of pions and kaons in the atmosphere results in an atmospheric neutrino spectrum approximately one power steeper than the input cosmic ray spectrum for primary cosmic rays with energies much larger than the pion and kaon critical energies [39].

2.2.2 Atmospheric neutrinos from kaons

Charged kaons (K^\pm) and neutral kaons (K^0) are mesons which are also produced in the atmosphere, but more rarely with a relative frequency of 10 – 15% to that of pions. Charged kaons (K^\pm) decay predominantly via the following channels [33]:

$$\begin{aligned}
 K^+ &\rightarrow \mu^+ + \nu_\mu \quad (63.55\%) \\
 K^+ &\rightarrow \pi^0 + e^+ + \nu_e \quad (5.07\%) \\
 K^+ &\rightarrow \pi^0 + \mu^+ + \nu_\mu \quad (3.53\%) \\
 K^+ &\rightarrow (\text{hadrons}) \quad (28\%)
 \end{aligned}
 \tag{2.13}$$

and K^- decay via the charge conjugate modes. Compared to pions, the dominant two-body decay $K^+ \rightarrow \mu^+ + \nu_\mu$ results in a higher energy neutrino as, unlike for pions, the mass of the kaon is much larger than the muon and the energy of the kaon is split approximately equal between the pion and neutrino. The mean lifetime is 1.24×10^{-8} s for K^\pm which is approximately half the lifetime of charged pions. For K^\pm we have $\epsilon_K \sim 850$ GeV.

Neutral kaons (K^0) are separated into 2 equal modes, long (K_L^0) and short (K_S^0). From K_L^0 we observe the decay channels:

$$\begin{aligned}
 K_L^0 &\rightarrow \pi^\pm + e^\mp + \nu_e \quad (40.55\%) \\
 K_L^0 &\rightarrow \pi^\pm + \mu^\mp + \nu_\mu \quad (27.04\%)
 \end{aligned}
 \tag{2.14}$$

and from K_S^0 we observe a very small contribution

$$K_S^0 \rightarrow \pi^\pm e^\mp \nu_e \text{ (0.07)\%} \quad (2.15)$$

The mean lifetime of the neutral kaon mode K_L^0 is 5.12×10^{-8} s, with a critical energy of 205 GeV. For K_S^0 , the mean lifetime is shorter at 8.95×10^{-9} s giving a critical energy of 1.02×10^5 GeV, resulting in a small semi-prompt component as the energy lies between the small pion and kaon critical energies and the larger critical energies of charmed D mesons (Section 2.2.4).

Charged pion (π^\pm) and charged kaon (K^\pm) neutrino production in the atmosphere can be described by the model: [39]

$$\begin{aligned} \frac{dN_\nu}{dE_\nu} = & \frac{\phi_N(E_\nu)}{(1 - Z_{NN})(\gamma + 1)} \\ & \times \left(\frac{Z_{N\pi}(1 - r_\pi)^\gamma}{1 + B_{\pi\nu} \cos \theta E_\nu / \epsilon_\pi} + 0.635 \left[\frac{Z_{NK}(1 - r_K)^\gamma}{1 + B_{K\nu} \cos \theta E_\nu / \epsilon_K} \right] \right) \end{aligned} \quad (2.16)$$

For fixed θ , we observe that asymptotically with increasing energy the fraction of neutrinos from charged kaons exceeds that of charged pions due to the larger critical energy of kaons, with kaons becoming the dominant source of neutrinos above $E_\nu > 100$ GeV, as well as the first source of electron neutrinos ν_e observable by IceCube.

At TeV energies and above most kaons will also re-interact with hadrons in the atmosphere - at energies of 100 TeV the probability of decay of a charged kaon compared to re-interaction is significantly less than 1% [40].

We term the neutrinos that result from from π and K mesons in air showers the conventional atmospheric neutrino flux, and this flux has been measured by IceCube both as muon neutrinos [41] [42] [43] and also as electron neutrinos in IceCube-DeepCore [44] and the full detector [43] [45].

2.2.3 Muons from cosmic ray air showers

In addition to neutrinos directly from pions and kaons, the muons produced from pion and kaon decay can also then decay to produce neutrinos (Equation 2.3), however the contribution to the total neutrino flux from muon decay at 100 GeV is less than 15% of the total contribution [39], and this decreases significantly with increasing energy for muons produced in air showers. With a soft energy threshold of 30 TeV due to a threshold on total observed photoelectrons observed by the IceCube detector (see Chapter 4), the HESE analysis does not see a flux of neutrinos from muons, and therefore this contribution is ignored in simulation. We however consider muons themselves, as the light from the passage of muons through matter form the dominant background for IceCube analyses.

Muons have a mean life time $\tau = 2.197 \times 10^{-6}$ s and have a mass of $105.7 \text{ MeV}/c^2$. At energies of 100 GeV this would allow muons to travel lengths of $\beta c \gamma \tau \sim 10^6$ m, comparable to the radius of the Earth. Muons are not seen to travel these distances through the Earth as they lose large amounts of energy through stochastic loss processes as well as continuous ionisation as they pass through matter. This results in a significant shortening of the distance in which the muons produced in air showers will decay. In the dense Antarctic ice these stochastic processes lead to the production of observable Cherenkov light. A crude model of muon energy loss as a function of slant depth is given by

$$-\frac{dE}{dX} = a + bE \quad (2.17)$$

where E is the muon energy, X is the column density traversed and a and b are constants.

For propagation of muons in ice we find values of $a = 0.259 \text{ GeV/mwe}$ and $b = 0.363 \times 10^{-3} \text{ mwe}^{-1}$. Here mwe is Meter Water Equivalent, 1 mwe is the equivalent of 100 g/cm^2 . We find that this formula holds for ice as a good approximation above $E_\mu = 10 \text{ GeV}$ and improves with muon energy [46]. Therefore, at 1 TeV, we obtain




Figure 2.1: Figure from [47]. Rates of atmospheric muons and muon neutrinos which pass preliminary IceCube filters and cuts. Analyses in IceCube require additional reconstructions and cuts to filter out the muon background to assess the smaller-rate neutrino distribution. In the analysis performed in [47], the goal is to measure the energy spectrum of atmospheric neutrinos, for which events are considered only with reconstructed zenith angles greater than 86° .

an average loss rate of about 0.7 GeV/mwe . From such a loss rate it is easy to tell that muons will not propagate much further than several kilometers through ice at the energies observed by IceCube, and similar loss rates are seen for rock and water. Muons will therefore lose the majority of their energy traversing through matter and then decay (see Eq. 2.3).

Muon losses are much smaller in air than water, though there is still a large probability of interaction with air particles, with a muon critical energy of $\epsilon_\mu = 1.0 \text{ GeV}$. Typically we say that a single muon loses approximately 2 GeV traversing the atmosphere from the height of the air shower, and for E_μ much greater than ϵ_μ both muon decay and energy loss can be neglected for air shower modelling.

Down-going muons from air showers form the largest number of background events in IceCube by many orders of magnitude (see Figure 2.1). Atmospheric muons themselves are selected against in the HESE analysis (see Chapter 4), however they remain a potential background at lower energies, and their soft spectrum reduces the efficiency of the veto rapidly when the charge threshold is reduced. Their stochastic energy deposition in the ice means that there is no consistent pattern of

light emission along the length of the muon track: as opposed to the relatively consistent ionisation pattern seen for low energy muons in, for example, Cherenkov water tanks, the energy deposition from high energy muons over large distances is probabilistic and “lumpy”. This needs to be taken into account to reconstruct the muon energy which is discussed in Section 3.5, and to determine the rate at which they pass a muon veto. The details of the muon veto are discussed in Section 4.1.

2.2.4 Prompt atmospheric neutrinos from charmed mesons

Prompt neutrinos, by definition, occur due to the prompt decay of mesons prior to possible re-interaction with the air nuclei. At sufficiently low primary cosmic ray energies, the decay of the resulting pions and kaons would constitute a prompt neutrino flux as these mesons would be more likely to decay before interacting. However, given the neutrino energies we can observe with IceCube in the HESE analysis, there is the potential of seeing a neutrino flux from the decay of charmed mesons. Charmed mesons have substantially larger critical energies than pions and kaons in our atmosphere, meaning they are substantially more likely that they will decay before re-interaction in air showers. This prompt flux has a harder spectrum than the conventional atmospheric flux as the conventional flux from the cosmic ray spectrum is steepened by a power law which occurs due to re-interaction. The prompt flux by contrast has no such steepening effect. This can make it somewhat more difficult to discriminate an astrophysical neutrino flux from a prompt flux of atmospheric neutrinos in IceCube, particularly in an up-going muon neutrino sample where neutrinos are unaccompanied by detectable particles and therefore muon veto techniques cannot be applied.

The prompt contribution comes primarily from the following decay channels:

$$\begin{aligned}
 D^+ &\rightarrow (\text{mesons}) + e^+ + \nu_e \quad (12.8\%) \\
 D^+ &\rightarrow (\text{mesons}) + \mu^+ + \nu_\mu \quad (13.0\%)
 \end{aligned}
 \tag{2.18}$$

along with corresponding charge conjugates for D^- decays, in addition to smaller contributions from D^0 decay channels:

$$\begin{aligned} D^0 &\rightarrow (\text{mesons}) + e^+ + \nu_e \quad (10\%) \\ D^0 &\rightarrow (\text{mesons}) + \mu^+ + \nu_\mu \quad (5.2\%) \end{aligned} \tag{2.19}$$

and D_s^+ decay channels:

$$\begin{aligned} D_s^+ &\rightarrow \tau^+ + \nu_\tau \quad (5.2\%) \\ D_s^+ &\rightarrow (\text{mesons}) + e^+ + \nu_e \quad (5.16\%) \\ D_s^+ &\rightarrow (\text{mesons}) + 2(e^+ + \nu_e) \quad (3.66\%) \end{aligned} \tag{2.20}$$

The decay times for the D^\pm , D^0 and D_s^\pm mesons of 1.0×10^{-12} s, 4.1×10^{-13} s, and 5.0×10^{-13} s respectively are short enough to ensure decay before interaction at TeV to PeV neutrino energies, with corresponding critical energies of 3.7×10^7 GeV, 9.7×10^7 GeV and 8.4×10^7 GeV respectively.

The neutrino flavour ratio observed from the prompt flux is approximately $(\nu_e : \nu_\mu : \nu_\tau) = (1 : 1 : 0)$, with only a small contribution to the ν_τ flux from D_s^+ .

The current upper limits of a prompt neutrino flux in contained events are placed at a level of 1.52 times the ERS [37] nominal flux model at the 90% confidence level [43]. Multiple analyses with the IceCube detector have also placed their own limits on the prompt flux contribution from up-going muons [42] [41] [48], with the latter setting an upper limit of $1.06 \times \text{ERS}$ at the 90% confidence level. The ERS model has been updated with recent data from the Large Hadron Collider (LHC) [49].

2.2.5 Accompanying muons

One possible way that conventional and prompt neutrino fluxes can be distinguished from astrophysical neutrinos is by the differing energy spectrum of each flux. However, there is another key difference between atmospheric and astrophysical neu-

trinos: atmospheric neutrinos are created in the Earth's atmosphere in cosmic ray air showers, and astrophysical neutrinos are not. For IceCube, this results in an extra means of discriminating astrophysical neutrinos from atmospheric neutrinos: for down-going neutrinos from the southern hemisphere, the presence of other detectable particles (namely muons) from the same air shower as atmospheric neutrinos. There are no expected accompanying particles associated with astrophysical neutrinos observable in IceCube on time with the neutrino and hence incoming muon veto techniques can be used to significantly reduce background.

The veto technique utilised for the High Energy Starting Event dataset of interest in this thesis utilises the principle that high energy muons continuously produce substantial amounts of Cherenkov light, and this light should be observed as the muons enter the detector. As such, a boundary of photo-detectors that make up the in-ice IceCube detector is selected as a veto region, and events which produce minimal light in this region but are of high energy within the detector are therefore most likely to be the result of neutrino interactions unaccompanied by muons. This technique removes both background muon events and a large fraction of down-going atmospheric neutrino events from cosmic ray air showers. Some signal events will be removed by the same technique due to chance background coincidences of light on the detector boundary.

2.3 Interaction at Earth

The neutrino interacts with other particles only via the weak force (with the exception of an unmeasurable gravitational force). Weak force interactions involve the exchange of a W^\pm or Z boson, which are called Charged Current (CC) and Neutral Current (NC) interactions respectively (see Figure 2.2).

Neutral current interactions result in an outgoing neutrino, carrying a fraction of the energy of the incoming neutrino. For neutral-current neutrino interactions observed within the detector, it is highly unlikely that the secondary neutrino from




Figure 2.2: Feynman diagrams showing typical interactions of neutrinos (ν_l) and anti-neutrinos ($\bar{\nu}_l$) of flavour l with standard matter quarks. Neutral Current (NC) interactions via a Z^0 boson which produce outgoing neutrinos ν_l ($\bar{\nu}_l$) are shown in the two left figures. Charged Current (CC) interactions via W^+ and W^- bosons which produce outgoing leptons l^- (l^+) are shown in the two right figures. Time is shown flowing left to right for each figure. Diagrams made using the TikZ-Feynman package [50].

that interaction will go on to interact on the scale of our experiment. Charged current interactions result in an outgoing charged lepton (e^\pm , μ^\pm , or τ^\pm) dependent on the incoming neutrino flavour.

For both CC and NC interaction types, the remaining fraction of the original neutrino energy not given to the outgoing lepton is imparted to the target nucleus, with sufficient energy to break apart the nucleus and produce jets with this type of interaction termed *Deep Inelastic Scattering* (DIS).

For IceCube and water-based neutrino observatories, detection of the outgoing charged particles in the ice can be performed indirectly by observing Cherenkov photons [52]. Cherenkov photons are produced by charged particles travelling through a dielectric medium (being electrically insulating rather than conductive) of refractive index n faster than the local speed of light c/n . The mechanism for this radiation is that in a dielectric an induced electromagnetic field from a charged particle will act to induce a small displacement of electrons bound to their atoms as the charge passes without inducing any excitation or ionisation (however close to the charged particle ionisation can still occur). For a particle traversing the medium at slow speed the displacement of the electron does not produce any observable radiation as wavelets from the displacement of neighbouring electrons interfere destructively; however, if a charged particle travels past at a speed greater than the local speed of




Figure 2.3: A schematic showing the principle of Cherenkov radiation. Spherical wavelets propagate at speed c/n within the medium shown emanating from points P_1 , P_2 and P_3 . These wavelets constructively interfere along BC corresponding to a Cherenkov angle θ . Figure from [51].

light, the wavelets from each electron along the path of the charged particle become coherent and constructively interfere [51].

Figure 2.3 shows the constructive interference of the wavelets. A charged particle travels a distance $AB = \beta c \Delta t$ where $\beta = v/c$ is the speed of the particle relative to the vacuum speed of light. Over the same time interval a wavelet of light propagates distance $AC = (c/n)\Delta t$. This gives the cherenkov angle relative to the path of the charged particle as

$$\cos(\theta) = \frac{AC}{AB} = \frac{(c/n)\Delta t}{\beta c \Delta t} = \frac{1}{\beta n} \quad (2.21)$$

This also imposes a threshold of $\theta = 0$, corresponding to particle speed $\beta = 1/n$, constructive interference does not occur for charged particles traversing a medium at speeds less than this.

The charged particles produced in neutrino interactions are typically highly relativistic (β very close to 1) and hence $\cos(\theta) = 1/n$. In south polar ice, at wavelength

$\lambda = 400$ nm the refractive index is $n(\lambda) \sim 1.32$ giving a Cherenkov angle of 41° [53]. Over the 300 nm to 600 nm wavelengths that IceCube is sensitive to, the spectrum of Cherenkov radiation has a $1/\lambda^2$ dependence.

Other methods to detect high energy neutrinos include the detection of radio Cherenkov emission. This is based on the principle that a charge asymmetry develops for showers in dense media, of about 20% the total particle number. The charge moves through the medium as a bundle on the scale of centimetres, which means for wavelengths greater than this size the radio Cherenkov emission over the charges is coherent. As the radiated power scales quadratically with the charge this results in a substantial amount of power being radiated at these wavelengths [54]. Several experiments including the Askaryan Radio Array (ARA) aim to detect this radio emission [55], as yet no detectors have observed an unambiguous ultra-high energy neutrino flux.

Air shower detection methods, for example those used at the Pierre Auger Observatory, are also being used to look for highly inclined neutrino interactions and the decays of up-going tau particles from Earth-skimming ν_τ charged-current interactions [56].

2.4 Propagation to the detector

2.4.1 Attenuation through the Earth

Neutrinos propagating through the Earth are notably diminished by Earth absorption. The cross section of neutrinos with nuclei rises approximately linearly with neutrino energy up to approximately 1 TeV, at which point the gauge-boson propagator restricts the momentum transfer, and the cross-section rises at a slower rate [57] (see Figure 2.4). This rise in cross-section with neutrino energy means that at higher energies the probability of the survival of neutrinos travelling long chords through the Earth is substantially smaller than at lower neutrino energies. The effect of Earth




Figure 2.4: Neutrino-nucleon (left) and antineutrino-nucleon (right) cross sections as a function of neutrino energy on stationary targets for the CTEQ3 distributions. Figures recreated from [57]. More recent estimates show a small relative decrease in the cross section at the highest energies [58].

absorption is a noticeable effect on all atmospheric and astrophysical fluxes as seen in Figures 5.1 - 5.8. In particular, the diffuse astrophysical flux prediction would appear isotropic within the detector were it not for the loss of neutrinos propagating through the Earth.

In addition to neutrino-nucleon interactions, we expect resonances of electron antineutrinos with electrons in matter [59]:

$$\bar{\nu}_e + e^- \rightarrow W^- \rightarrow (\text{Anything}) \quad (2.22)$$

at an energy of $E_\nu = 6.3\text{PeV}$, known as the Glashow resonance, which would result in an increase in the cross section and experimental effective area for electron neutrino interactions at that energy, to a factor of 300 times that of the charged-current neutrino-nucleon interaction [60].

2.4.2 Neutrino oscillations

In particle physics it is not immediately obvious that the neutrino should have mass, indeed in the Standard Model of particle physics all neutrinos are massless [61]. The existence of neutrino oscillations confirms that neutrinos must have differing mass states. The investigation of neutrino oscillations was motivated by the ap-

parent disappearance of atmospheric muon neutrinos in the Kamiokande neutrino detector [62], a detector designed to improve experimental limits on proton decay. Numerous advanced detectors able to carefully measure neutrino fluxes include the Sudbury Neutrino Observatory (SNO) which confirmed neutrino oscillations in solar neutrinos [63] by observing NC interactions on deuterium to quantify the interactions of all neutrino flavours against ν_e CC interactions. The Super-Kamiokande detector replaced the previous Kamiokande detector, and was designed to confirm the existence of atmospheric neutrino oscillations [64]. In the sub-GeV range, a significant deficit in the ratio of muon to electron neutrinos was confirmed.

The difference in neutrino masses gives rise to flavour oscillations as the neutrino flavour states do not correspond directly to the neutrino mass states. Neutrino flavour states can be expressed as a superposition of mass eigenstates

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (2.23)$$

where ν_1 , ν_2 , and ν_3 are the neutrino mass eigenstates and U is the unitary mixing matrix. To describe flavour oscillations it is easier to begin with two-flavour neutrino oscillations. The probability for a flavour state ν_α to be observed in flavour ν_β after travelling a distance L through a vacuum is given by the equation [61]

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2 2\theta \sin^2 \left(\frac{1.27 \Delta m^2 (\text{eV}^2) L (\text{km})}{E_\nu (\text{GeV})} \right) \quad (2.24)$$

where θ is termed the mixing angle, Δm is the mass difference between states, and E_ν is the energy of the neutrino.

In a full 3-flavour description of neutrino oscillations, oscillation probabilities are described by three mixing angles θ_{12} , θ_{13} , and θ_{23} , in addition to three mass-squared differences Δm_{12}^2 , Δm_{13}^2 , and Δm_{23}^2 , of which two are independent, and the Dirac phase δ . The mass hierarchy problem is to determine the ordering of the neutrino




Figure 2.5: Results from the IceCube DeepCore detector constraining two of the neutrino mixing parameters through the observation of atmospheric neutrinos. The 90% confidence level contours are shown here, other dedicated neutrino experiments however currently place better constraints on the parameters. Figure from [68].

masses given their mass-squared differences [65]. A more in-depth explanation of neutrino mixing is given in the Review of Particle Physics [33].

Vacuum neutrino oscillations account for the change in flavour of astrophysical neutrinos. Over cosmological distances the oscillation terms are averaged due to the scale of the source region and spectrum of neutrinos [66], which for the ratios expected from pion decay yields the expected $(\nu_e : \nu_\mu : \nu_\tau) = (1 : 1 : 1)$ ratio at Earth.

As can be seen in Equation 2.24 the effect of neutrino oscillations over a relatively small distance can be diminished with sufficiently large energy (P goes to 0 as the argument of the \sin^2 term goes to 0). The oscillations of atmospheric neutrinos in our analysis are limited to a chord through the Earth to the detector ($L < 2r_{\text{Earth}}$). At the neutrino energies of 30 TeV and above of concern in this analysis oscillations are therefore negligible, and do not need to be accounted for [67]. Within the IceCube DeepCore detector lower energy atmospheric neutrino oscillations are observed in IceCube and used to measure the neutrino mixing parameters [68] (see Figure 2.5).

2.5 Summary

In this chapter we have characterised neutrinos, which are separated into three flavour states, are of small mass and possesses a small cross-section with ordinary matter. The production of neutrinos from cosmic ray interactions both in source regions and in the atmosphere has been discussed, in addition to the characteristics of cosmic ray air showers in general, and how astrophysical neutrinos can be distinguished from the neutrinos and other particles produced in cosmic ray air showers in the atmosphere. In addition, we have discussed the detection mechanism for neutrino interactions, and the impact of neutrino oscillations on both astrophysical and atmospheric neutrinos. In Chapter 3 we discuss the IceCube Neutrino Observatory, and how this detection mechanism can be utilised within the glacial ice at the south pole.

Chapter 3

The IceCube Neutrino Observatory

High energy neutrino detectors typically aim to detect the radiation produced by charged secondary particles of neutrino interactions within a transparent medium. As these charged particles transit a dense transparent medium they emit Cherenkov radiation that can be detected as photons in the visible-to-UV light region by photomultipliers. This idea of using a gigaton-scale experiment with a transparent medium such as water or ice to detect neutrinos dates back as early as the late 1950s by Markov and Zheleznykh [69]. Low energy detectors comprising a smaller volume such as Super Kamiokande detect neutrinos with energies below 1 TeV interacting in water and producing light using a dense array of photomultipliers around the edge of a vast cylindrical well. Other low energy detection methods include Liquid Scintillator and Gas Counter techniques [61].

With increasing energies above GeV scales, sparser arrays of photomultipliers can be used to detect a relatively small fraction of the total Cherenkov light emitted by the interaction of secondary particles in order to infer the properties of the neutrino primary. Initial attempts at a sparse array in water include the DUMAND project [70], with current detectors including ANTARES [71], Baikal [72], and the




Figure 3.1: The IceCube Neutrino Observatory including surface and in-ice array. The differing colours on the surface correspond to the deployment seasons of the in-ice strings and IceTop tanks.

upcoming KM3Net [73]. Within ice, the AMANDA [74] detector preceded the IceCube detector at the same site beneath the Amundsen-Scott station at the south pole. AMANDA was used to characterise the atmospheric neutrino flux at high energy and set limits on the astrophysical diffuse flux and point sources in the northern sky, but was unable to make a clear detection of an astrophysical neutrino flux [75].

In this chapter we discuss the IceCube Neutrino Observatory located at the south pole, designed with the intent of measuring the interactions of high energy neutrinos in the south polar ice to discover an astrophysical component of the high energy neutrino flux. This chapter starts with a discussion of the properties of the detector itself including the instrumentation and properties of the surrounding ice, and includes the methods used to reconstruct the interactions of particles within the ice for events in the High Energy Starting Event (HESE) selection.

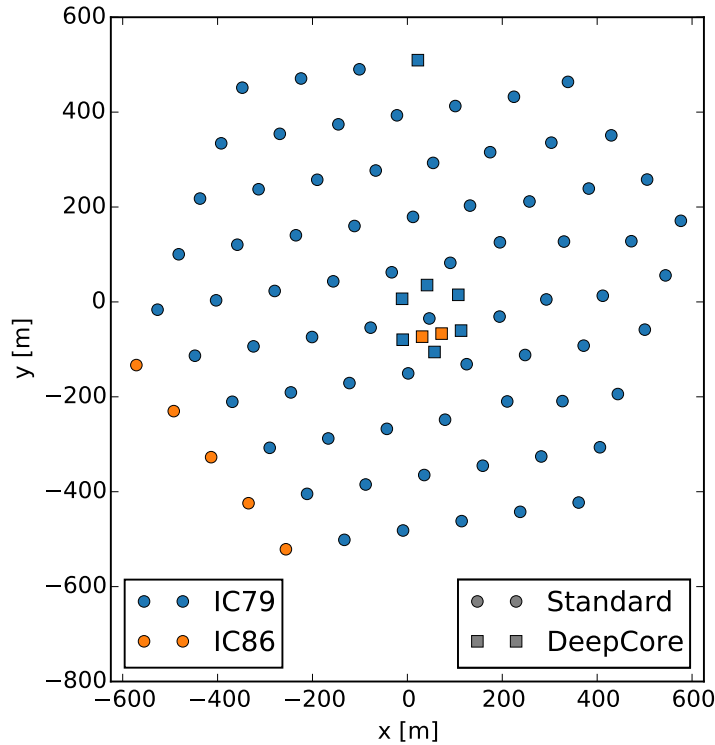


Figure 3.2: A top-down view of the IceCube strings relevant to the HESE analysis. Circles indicate standard in-ice strings, squares indicate DeepCore strings. The IC79 deployment covers the May 2010 - May 2011 season which is the first year of the HESE selection, and the IC86 deployment completes the current detector configuration which has been in operation from May 2011 onwards. IC86 adds 5 strings in the bottom-left corner and 2 additional in-fill strings to DeepCore.

3.1 The detector

The IceCube Neutrino Detector is an array designed to measure Cherenkov light yield of charged particles produced in neutrino interactions as they traverse the glacial ice at the south pole. IceCube consists of an in-ice array [76] of Digital Optical Modules (DOMs, discussed in Section 3.3), including a dense central array DeepCore, and an additional surface array called IceTop [77]. The in-ice array consists of 5160 DOMs which are deployed on 86 strings containing 60 DOMs on each. On standard (non-DeepCore) strings the DOMs are deployed between depths of 1450 to 2450 m beneath the surface, with a vertical spacing of 17 m between DOMs. The inter-string spacing between standard strings in the detector is 125 m

in a hexagonal pattern (Figure 3.2).

IceCube DeepCore is an additional dense infill array at the centre of the detector consisting of 8 of the 86 strings. DeepCore reduces the inter-string spacing substantially to about 70 m in an attempt to lower the energy threshold of IceCube to neutrino energies on the order of 10 GeV [78]. The detectors on the 8 DeepCore strings are deployed avoiding a layer of mineral dust deposition between 2000 to 2100 m, with 10 DOMs above this layer with a spacing of 10 m between DOMs, and 50 DOMs below with a spacing of 7 m between DOMs. The motivation of placing this array at the center of the detector is to use the surrounding instrumented volume as a veto of background muons while searching for low energy neutrino interactions in the clearest region of ice in the detector.

The IceTop surface array consists of 81 stations with 2 ice-Cherenkov tanks per station, and 2 DOMs in each tank. The stations generally align with the string configuration of the in-ice array below. The surface array is a cosmic ray air shower detector. This detector can be used independently for cosmic ray observation or combined with the in-ice array to collect coincident events. These coincident events are used for advanced composition studies in the “knee” region of the cosmic ray spectrum [77]. While the in-ice events from muons and neutrinos can consist of as little as tens of photoelectrons observed over as many DOMs in the in-ice array, IceTop detects the muons and electrons from air showers with each muon producing signals of approximately 130 photoelectrons in the corresponding DOMs.

3.2 Ice properties

Antarctic ice that is bubble free is the purest natural solid found on Earth, and is also the most transparent in the wavelengths of interest to us in the ultraviolet and visible wavelengths [80]. As such it is an ideal detection medium for observing neutrino interactions by their Cherenkov radiation. However, the optical properties of the ice are not homogeneous as a result of trace impurities (primarily mineral

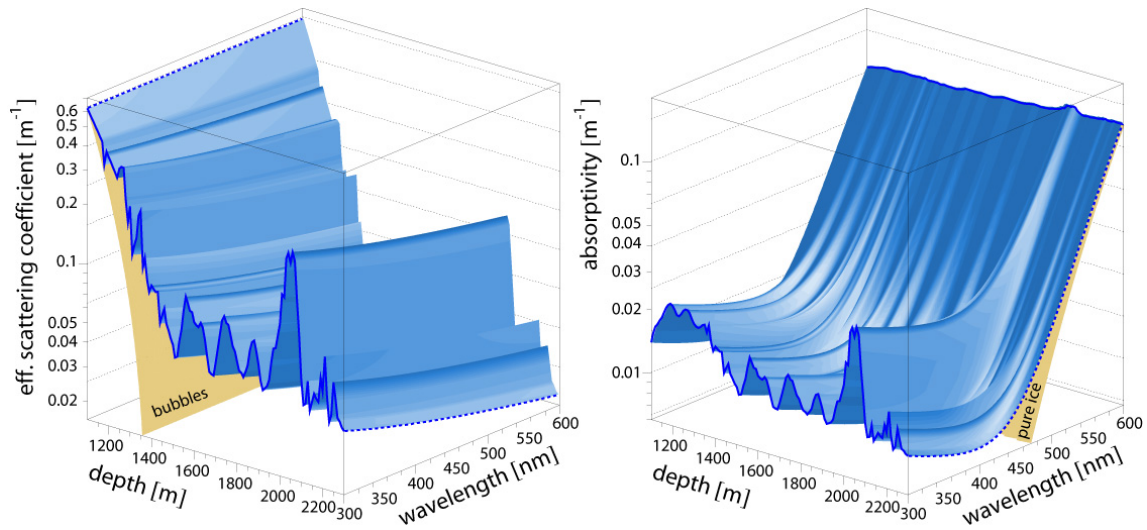


Figure 3.3: These images show the scattering (left) and absorption (right) of the ice over the range of observable wavelengths for the IceCube DOMs, smaller values here indicate better optical properties (larger distances to scattering and absorption). The surfaces in yellow indicate limits due to scattering off bubbles in the ice (left) and absorption in pure ice (right). The “dust layer” around 2000 m depth can clearly be seen as an increase in both scattering and absorption coefficients. Figures from [79].

dust and salts) that vary with depth which need to be accounted for to correctly reconstruct particle interactions in the ice.

The dust concentration throughout the ice is dependent on past variations in climate and glacial flow. In general we find relatively clear ice for depths above 2000 m, a layer of substantial dust deposition from 2000 to 2100 m which noticeably affects the visibility of Cherenkov light, and the clearest glacial ice is found in the layers beneath. Bubbles in the ice severely shorten the photon scattering length, and bubbles are prominent in the ice above 1300 m. The detector was deployed below this 1300 m depth to avoid this increase in scattering.

The change in scattering and absorption in the ice near the detector can be seen in Figure 3.3. Scattering of photons in the ice has the effect of changing the photon arrival timings by allowing different possible paths for photons to travel from their source to the PMT, adding complexity and additional uncertainty to event reconstruction. Absorption lowers the total photon statistics observed, requiring brighter events in order to observe them and resulting in less information available for event reconstruction. Both absorption and scattering need to be accurately

modelled in order to correctly reconstruct neutrino properties such as energy and direction from the observed light emission.

The scattering and absorption properties of the ice are measured by the DOMs themselves by using on-board LED flashers, in addition to laser dust-logging devices deployed on eight strings. Surrounding DOMs can measure the photon arrival distributions from the pulses of light emitted by the on-board LEDs, these arrival time distributions from multiple pulses throughout the detector are used to build a parameterised ice model, the best fit model giving the scattering and absorption coefficients [53].

The IceCube bore-holes were created using hot-water drilling techniques, the strings were then lowered into the holes and the melt water re-froze around the string. This re-freezing of the ice was non-uniform, with the edges freezing first before the central core. As a result, there is a higher concentration of bubbles in the core region of the hole ice. In addition to this, the DOMs do not necessarily freeze in the centre of the hole ice but prefer resting against the side of the hole, causing an asymmetry in the local ice around each DOM. Very recent ice models are capable of modelling these effects, at the cost of an increased complexity of the model.

3.3 Digital Optical Modules (DOMs)

The individual DOMs in the detector consist of a wide-angle 25 cm diameter Hamamatsu R7081-02 Photo-Multiplier Tube (PMT) [82] combined with onboard electronics encased in a spherical glass pressure sphere designed to withstand the pressure exerted by the ice. Unlike the AMANDA neutrino telescope which served as a predecessor to IceCube and sent analogue PMT responses to the surface to be digitised [83] the data collected by the DOMs are digitised and preliminary processing of the PMT response is performed with the DOM electronics before the data are sent to the surface. There are several benefits as a result of this preliminary processing: the time resolution is improved; and a larger dynamic range of the signal is able to




Figure 3.4: Side-view of a Digital Optical Module (DOM). Figure from [81].

be collected, i.e the amount of charge that can accurately be measured before saturation of the signal occurs. In addition, the correct clock times are maintained on each DOM using a local oscillator, and these times are synchronised with a master clock at the surface, allowing for precise measurement of photon arrival times.

A DOM is considered to have received a hit when the voltage reading on the PMT anode is 0.25 single photo-electrons (SPE) [81], where 1 SPE is the average single photo-electron pulse height. A hit causes the DOM to digitise the data from a 75 ns delay channel. Digitisation of the waveform is done in two different ways on the DOM mainboard. The first is with two dedicated Analog Transient Waveform Digitisers (ATWDs). The timing resolution of the ATWDs is 3.3 ns per sample, and 128 samples are stored per buffer, corresponding to a buffer size of 422 ns, and this buffer is recorded over 3 gain levels to capture both small and large signals. It takes a total of 29 μ s for an ATWD chip to digitize and report a waveform after capture, leading to a large downtime. To eliminate downtime, the second ATWD chip can collect data while the other transmits, but additional structure of the waveform after the trigger is collected by the second data collection method - the fast Analog-to-

Digital Converter (fADC) chip. This chip samples the data at a rate of 25 ns per sample with a buffer chosen to be $6.4 \mu\text{s}$ to exceed the likely timescale of events in IceCube, and has no downtime.

The reporting of the waveform does not occur unless a Local Coincidence (LC) condition is met, where a neighbour above or below the DOMs also receive hits. Neighbouring DOMs have dedicated copper links to transmit their trigger status. When a DOM receives a hit itself, it opens a receptive time window of the order $1 \mu\text{s}$ to receive a LC tag from its neighbours. Likewise, a quiescent DOM that has not received a hit but has received a tag from a neighbour will also open the same window with the possibility of receiving hits. Without a LC tag after a hit, to minimise downtime, the ATWD chip resets within 50 ns. Tagged hits occur as a small percentage of all PMT hits and are far less likely to be PMT noise. The detector primarily runs in *Soft Local Coincidence* (SLC) mode, where tagged LC hits are digitised and read out with ATWD and fADC data, while isolated hits contain no ATWD data and are reported as pulse height and time for 3 fADC bins around the peak, rather than full waveforms.

Within the IceCube Laboratory (ICL), computers at the surface termed the South Pole System (SPS), monitor and configure the DOMs remotely as well as processing the output data and waveforms to construct events. Events contain all the relevant data from the detector including isolated pulses and waveforms once a trigger condition is met. For the data in this analysis, the *Simple Multiplicity Trigger* (SMT8) was used. This trigger generates an event when 8 LC hits are observed within $5 \mu\text{s}$. When this trigger condition is met, all hits within a $20 \mu\text{s}$ window are stored as an event.

3.4 Waveform unfolding

To reconstruct events in IceCube, we need to know when and how many pulses were observed on each DOM over a buffer readout. Rather than working with the

digitised PMT response in terms of voltage offset, we wish to have this information converted to a pulse series for each DOM in the detector. This is the process of pulse extraction by means of unfolding: as we have measured the PMT response function for single photons, we can try to infer what the initial photons were that produced the complete PMT waveform. This is done by assuming that in normal ranges of operation the PMT electronics behave linearly, and therefore the waveform read by the PMT can be treated as a linear superposition of many single photon response functions with relative time offsets.

Such a system can be represented in matrix notation:

$$\vec{y} = \mathbf{B} \cdot \vec{\alpha} \quad (3.1)$$

Where \vec{y} is the observed digitised waveform, $\vec{\alpha}$ is the vector ultimately to be solved for containing the magnitude of photons at individual points in time received by the PMT, and $\mathbf{B} = [B_{ij}]$ is a set of basis functions that describe the response of the PMT. Since the data is recorded by both ATWD and FADC, we take into account both types of output. An initial solution can therefore be found by matrix inversion, however we wish to impose two additional constraints to the solution $\vec{\alpha}$. First, that all inputs be positive; second, that we avoid over-fitting the response and seek a minimal solution - that is to have as few pulses as possible while still accurately describing the waveform. This is achieved with a non-negative least squares fit and also maximising for the sparsity of the solution. Further details are discussed in [84].

3.5 Reconstruction techniques

In order to infer the properties of the interacting neutrino from the distribution of light observed in the detector, we need to apply reconstruction algorithms. Different algorithms are applied for different purposes: some are required to be fast in order to process the large quantities of data given by the detector, other algorithms generally




Figure 3.5: Figure showing the longitudinal shower size in photon count (top) and the ratio of photon count to input neutrino energy (bottom). Figure from [85].

trade off precision for computation time.

The high energy events observed by the IceCube detector are categorized primarily into two main topological types, track-type events and cascade-type events. Track-type events are due to the presence of a muon, either passing through the detector from a distant cosmic ray or neutrino interaction, or as part of a charged-current (CC) event producing a muon within the detector, either directly from a muon neutrino (ν_μ) interaction or as a possible decay of a tau lepton (τ) from a tau neutrino (ν_τ) interaction. Cascade-type events are produced from a particle shower that begins at the interaction vertex. The particle shower for cascade type events does not typically contain particles such as high energy muons that would extend the light emission to a track topology and so the light emission appears to have a single point-like origin on the scale of our detector. There are two subtypes of particle showers that apply to neutrino interactions at our event energies: hadronic showers and electromagnetic showers.

Cascade reconstruction

Hadronic and electromagnetic showers in the ice are short in extent, on the order of metres (top of Figure 3.5). The charged particles in the shower each travel a short distance in the ice yielding Cherenkov light. The total photon count scales linearly with the total energy of the shower: the number of photons per GeV of deposited energy remains constant (bottom of Figure 3.5). The fact that the light profile of the shower remains very consistent with energy scale over the distance scales that IceCube can resolve allows for a simple model of light deposition from such showers.

A model $\vec{\Lambda}$ of average light deposition over all observed DOM positions and time bins for an event of reference energy E_0 (usually 1 GeV) is tabulated from Monte Carlo simulation of cascades. These cascades are simulated over the geometry of the detector given a model of the ice properties. For a given event position, direction and time, one can look up the predicted DOM response in this table. The dimensionality of these tables are designed to be small in order to minimise the call time of the likelihood and the size of the tables in computer memory, and so higher order effects such as ice anisotropy [86] are not included.

Using this model we have the equation for the predicted photon rate in each time bin $\vec{\lambda}$ for our event energy E (in units of GeV):

$$\vec{\lambda} = \vec{\Lambda}E \quad (3.2)$$

Within the DOMs we not only expect light from the event but we also expect a uniform mean noise rate in each time bin $\vec{\rho}$ from a number of factors including natural radioactivity and scintillation in the DOM [82], therefore we need to include a noise term in each time bin

$$\vec{\lambda} = \vec{\Lambda}E + \vec{\rho} \quad (3.3)$$

The number of deposited photoelectrons observed for each time bin over all the DOMs \vec{k} should follow Poisson statistics of the predicted rate $\vec{\lambda}$, and therefore we

have the likelihood equation

$$\mathcal{L} = \prod_i \frac{\lambda_i^{k_i} e^{-\lambda_i}}{k_i!} \quad (3.4)$$

which relates the total probability of observing \vec{k} given $\vec{\lambda}$. Here, i denotes each time bin.

To infer the energy of the event, we substitute in $\lambda_i = \Lambda_i E + \rho_i$ and obtain

$$\mathcal{L} = \prod_i \frac{(\Lambda_i E + \rho_i)^{k_i} e^{-(\Lambda_i E + \rho_i)}}{k_i!} \quad (3.5)$$

Taking the logarithm, we obtain

$$\ln \mathcal{L} = \sum_i [k_i \ln(\Lambda_i E + \rho_i) - (\Lambda_i E + \rho_i) - \ln k_i!] \quad (3.6)$$

and then maximizing the equation with respect to E :

$$\frac{d \ln \mathcal{L}}{dE} = 0 \quad (3.7)$$

we obtain the result

$$\sum_i \left[\frac{k_i \Lambda_i}{\Lambda_i E + \rho_i} - \Lambda_i \right] = 0 \quad (3.8)$$

This cannot be solved for E algebraically due to the sum. In the case of $\vec{\rho} = \vec{0}$, we obtain

$$\sum_i \left[\frac{k_i}{E} - \Lambda_i \right] = 0 \quad (3.9)$$

and hence

$$E = \frac{\sum_i \Lambda_i}{\sum_i k_i} \quad (3.10)$$

which is the ratio of total predicted charge $\sum_i \Lambda_i$ for a cascade of energy E_0 to total observed charge $\sum_i k_i$. However equation 3.8 must be solved numerically if $\vec{\rho} \neq \vec{0}$, which is typically done using minimisation processes.

The reconstruction of direction also uses the same likelihood (Equation 3.5), however instead of maximising the likelihood by varying the energy of the event,

the overall likelihood is maximised while changing the model $\vec{\Lambda}$ of light emission by changing the interaction vertex position, shower direction and interaction time. Close to the shower, light will be preferentially emitted about the Cherenkov cone, however scattering in the ice will result in a spherical profile at large distances, therefore most of our directional information comes from DOMs located closest to the interaction vertex. The best fit direction is that which maximises the total likelihood, and the error can be deduced from the shape of the likelihood space scanned over all possible directions of the shower. The statistical median error (50% confidence interval about the maximum) is typically only a few degrees, but uncertainties in the ice parameters result in the uncertainty in the reconstructed direction being dominated by systematic errors resulting in a much larger angular uncertainty, which is determined on an event by event basis.

Track reconstruction

Track signatures in IceCube are produced by muons passing through the detector. When the muon energy is large enough (> 1 TeV) the muon loses its energy predominantly by stochastic processes including bremsstrahlung, pair production and photo-nuclear interactions. The light deposition from these losses is not averaged out over the scale of the detector but produces a “lumpy” pattern of light deposition that can vary greatly even for muons of the same energy.

To reconstruct events which contain muon components, including “pass-through” muons but also starting tracks and tau meson decays, we adapt the point model of light emission to a segmented line as a muon template. Each segment on this line has a freely scalable energy parameter to account for the stochastic nature of the energy loss pattern. The model $\vec{\Lambda}$ used before can therefore be thought of being broken up into multiple sub-showers $\vec{\Lambda}_1, \vec{\Lambda}_2, \dots, \vec{\Lambda}_N$ where the positions of each $\vec{\Lambda}_j$ are co-linear and equally spaced along the muon track. Then the predicted mean

rate of light deposition for our model is

$$\vec{\lambda} = \sum_j \vec{\Lambda}_j E_j + \vec{\rho} \quad (3.11)$$

or alternatively

$$\vec{\lambda} = \mathbf{\Lambda} \cdot \vec{E} + \vec{\rho} \quad (3.12)$$

noting that Λ is now a function of source j where light is emitted and also time bin i where light is observed and therefore a two-dimensional matrix $\mathbf{\Lambda} = [\Lambda_{ij}]$.

As with cascades (Equation 3.5), Equation 3.11 can be substituted into the Poisson likelihood (Equation 3.4):

$$\ln \mathcal{L} = \sum_i [k_i \ln(\sum_j \Lambda_{ij} E_j + \rho_i) - (\sum_j \Lambda_{ij} E_j + \rho_i) - \ln k_i!] \quad (3.13)$$

We wish to maximise Equation 3.13 for the best fit \vec{E} , a process we call the unfolding of the energy losses. The maximum of this equation can be found taking the derivatives with respect to each energy loss:

$$\frac{\partial \ln L}{\partial E_l} = 0 \quad (3.14)$$

This function in practice has as a first-order approximation:

$$\vec{k} - \vec{\rho} = \mathbf{\Lambda} \cdot \vec{E} \quad (3.15)$$

which is a solvable linear algebra problem for \vec{E} by matrix inversion. This is similar in form to Equation 3.1 and the problem of pulse extraction. Like pulse extraction we also wish to impose additional constraints, in this case when maximising Equation 3.13. The first is that each E_i be non-negative given a negative energy deposition is unphysical. Further regularisation terms may be used to favour solutions of smaller locally emitted light to a large but distant amount of light since both models produce small amounts of light in the detector, which can be a degenerate

solution particularly near the edges of the detector.

The HESE analysis uses deposited energy, inferred from the sum of the fitted losses for each segment, as an observable for the tracks in the sample as it is difficult to infer the neutrino energy from a starting track; not only due to the difficulty of reconstructing muon energies but also the ambiguity of the energy from the initial hadronic cascade overlapping the start of the muon track. For muon energies seen in this analysis the muons will not be calorimetric; only a fraction of the muon energy will be lost in the detector and the muon will exit the observable volume. Deposited energy is therefore a lower bound on the total event energy.

Directional reconstruction of a track-like event in the HESE selection is performed in the same way as a shower-like event, but here using the segmented track hypothesis. The hypothesis position, direction and time are varied along with the light output of each segment of the track, and the overall likelihood (Equation 3.13) is maximised. The directional reconstruction is far more accurate for a track event compared to a cascade as there is a large lever arm for the reconstruction from the linear deposition of the muon losses through the detector, making it easier to infer the direction of the neutrino.

Inference on neutrino energy

Absolute energy calibration can be performed by determining the in-situ DOM efficiency, and is performed in IceCube by studying minimum ionising muons. Minimum ionising muons in the context of the IceCube detector have energies of the order 100 GeV, produce little signal on the outer detector strings and stop within the fiducial volume [87]. The event selection of these muons looks for these properties in addition to high quality track reconstructions, and event directions of 45° - 70° with respect to zenith, as the resulting Cherenkov light emission should then be within the acceptance range of the PMTs. From this selection a correction of the nominal DOM efficiency is observed which is an excess of less than 5% on average over all track to DOM distances with a maximum of 9% at certain short distances [85].

Uncontained event analyses must infer the neutrino energy from a neutrino interaction that may have occurred at large distances from the detector for the vast majority of events. In a contained event analysis, we observe the neutrino interaction within the detector and so the entirety of the neutrino energy is observed as deposited energy in the detector for charged-current electron neutrino (ν_e) and most charged-current tau neutrino (ν_τ) interactions. A large amount of energy is observed in charged-current (CC) muon neutrino (ν_μ) interactions with some energy carried away by the muon, and NC interactions have energy from the secondary neutrino carried away that is invisible to IceCube. In addition, a small fraction of the energy observed will go into the generation of particles which are invisible to IceCube as they are too low energy to radiate Cherenkov light.

The analysis used in this thesis uses deposited energy and direction from reconstructions to discriminate neutrinos from atmospheric and astrophysical sources, in addition to the directional reconstruction being used to test for point sources.

3.6 Summary

In this chapter the IceCube Neutrino Observatory has been introduced along side other neutrino experiments of differing scales and detection media, the properties of the south polar ice as the detection medium for IceCube were discussed, and we have introduced the fundamental unit that comprises IceCube, the Digital Optical Module (DOM), and the conditions by which the detection of photons by the DOM can result in a readout of data. The process of data unfolding, taking the voltage recording and summarising this as a series of input pulses has also been discussed, along with the methods of inferring back the properties of the particle interactions, specifically deposited energy and direction, through maximum likelihood techniques which compare the observed data to models of the light production from particle interactions in the ice. The inferences of these reconstructions on the properties of the neutrino are summarised, as these are dependent on the type of neutrino

interaction observed.

In the next chapter we discuss contained event analyses, focusing on the High Energy Starting Event analysis and its results, as this is the event selection considered in the remainder of this thesis.

Chapter 4

Neutrino contained event analyses

The event selection for the key analysis of this thesis, the combined diffuse and point source analysis defined in Chapter 5, uses the previous work of the High Energy Starting Event (HESE) analysis, which was first performed on two years of data from May 2010 to May 2012 [88]. From this analysis we obtain:

- The event selection using an active veto technique
- The reconstruction of the events in energy and direction
- As a result of the event selection, the determined shape of the signal and background distributions

This chapter discusses the motivation for the contained event analyses in IceCube including the precise details of the HESE analysis, and also includes a summary of the results of the data up to the 2014 season [89] which are the results of interest for this thesis. The results of six years of data have recently been published [90].

4.1 HESE filter

The High Energy Starting Event (HESE) filter was motivated by the discovery of two cascades with energies above 1 PeV in an earlier extremely high energy

(EHE) event sample designed to detect EeV neutrinos [91], however the idea of an active muon veto predates this discovery [92]. The selection of the events in this earlier EHE sample employed a directional track-based reconstruction “NPE-weighted LineFit” [93], and then performed a cut on the number of photoelectrons observed (NPE). This cut varied with the reconstructed LineFit direction:

$$\log_{10} \text{NPE} > \begin{cases} 4.8 & \cos \theta < 0.075 \\ 4.8 + 1.6 \sqrt{1 - \left(\frac{1.0 - \cos \theta}{0.925}\right)^2} & \cos \theta \geq 0.075 \end{cases}$$

Here θ is the zenith angle of the LineFit direction, with $\cos \theta < 0.075$. Atmospheric muons come exclusively from above the horizon due to their limited range through matter, but by chance can mis-reconstruct with horizontal or up-going directions ($\cos \theta < 0.075$). This cut on total photoelectrons and angle is able to substantially reduce the background of atmospheric muons in the sample by looking at the extremes of deposited charge. The cut enables a search for extraordinary events with track topologies but also allows for cascade events, which can most easily pass the selection criteria if they reconstruct as up-going with $\cos \theta < 0.075$ and a minimum light deposition of ~ 63000 NPE observed in all DOMs. Cascades, producing an almost point-like source of light, have directional reconstructions from LineFit that are essentially arbitrary, and so by chance can pass the cuts in this selection even if the true event direction would fall outside the desired angular range. This selection on the data between May 2010 and May 2012 yielded two cascades within the detector. These events reconstructed with zenith angles from LineFit satisfying ($\cos \theta < 0.075$), with light deposition slightly above the NPE cut, with 7.0×10^4 NPE and 9.6×10^4 NPE. Applying a full cascade reconstruction algorithm taking the ice model into account yielded reconstructed energies for the events of 1.04 ± 0.16 PeV and 1.14 ± 0.17 PeV respectively. The significance for the detection of these two events, which is assessed based on rates passing the selection criteria rather than a likelihood fit as used for the HESE analysis, looks at the probability of observing two events passing the selection criteria given an origin of an atmospheric only flux,




Figure 4.1: The HESE veto region uses the topmost layers of in-ice DOMs and the outer layer of strings as a filter against incoming muons. A layer of veto DOMs beneath the dust layer prevents horizontally-inclined down-going muons from “sneaking” into the bottom third of the detector that could otherwise pass the outer strings in the dust layer undetected. Images from [88].

with a nominal prompt component. The result of this was 2.9σ , and even with an over-scaling of the prompt neutrino flux to a previous IceCube upper limit of 3.6 times the Enberg flux, a significance of 2.3σ was still obtained, suggestive of an astrophysical origin.

This observation of two cascade events starting within the detector boundary motivated a contained event search using the outermost layers of the detector as a veto layer to reject incoming muon events. This muon veto leaves an internal detection volume with greatly reduced background, allowing us to probe the underlying flux that created the PeV cascades [88]. This veto allows the investigation of lower energy interactions due to a lower NPE threshold, and additionally does not rely on event reconstructions. This saves both initial computation time and the possibility of event mis-reconstructions affecting the selection. As discussed in section 2.2.5, a key property of this veto is not only to veto the atmospheric muon background but also the ability to reduce the down-going atmospheric neutrino background by means of vetoing the accompanying muons from the same air shower.

The HESE veto requires sufficient fiducial volume within the boundary later,

which is achieved with the deployment of IC79 onwards. Data are analysed beginning with IC79 in May 2010 to May 2011 and the completed detector IC86 from May 2011 onwards. In both cases, the conditions on the events that pass the HESE selection are that less than 3 photoelectrons of the first 250 observed photoelectrons appear on the veto layer; and the event must have 6000 photoelectrons (NPE) overall to reduce the background of entering muons - for a muon to produce this much light overall it is also highly likely to deposit light in the veto region upon entering.

The geometry of the veto layer (Figure 4.1) is similarly designed in both IC79 and IC86, the only difference is that the new detector boundary strings along one edge of the detector in IC86 are now selected over the previous boundary. To prevent horizontally inclined events the outermost strings are all selected as part of the veto. The top 90 m of the detector geometry are also selected as a measure to remove the substantial flux of down-going muons from air shower interactions: this is in part due to the increased probability of detecting an entering muon with additional DOMs, and in part because not all DOMs are deployed to precisely the same depth. The bottom of the detector is also selected as a veto and filters both up-going muons - while these are also neutrino induced events, they are not the target of this analysis - and more importantly, horizontal muons travelling below the detector which, by chance, can emit a stochastic burst of light directly beneath the detector, mimicking a high energy cascade, are also vetoed with this single layer. There is also an additional region selected beneath the “dust layer” at approximately two thirds of the detector’s depth, a region of strong light absorption caused by a layer of dust in the Antarctic ice as discussed in Section 3.2. This additional veto region is able to select against near-horizontal down-going muon events that could otherwise pass through the dust layer on the outer strings where the light is absorbed, and pass into the region of clear ice below the dust layer undetected by the outer strings to pass the HESE selection.

4.2 Event selection

A list of the 54 events in the 4-year HESE sample can be seen in Tables 4.1 and 4.2. Event 32 is removed from all analyses as it is a pair of coincident muons from independent cosmic ray air showers. Coincident muons are able to erroneously pass the veto condition, the first muon passes the outer veto and deposits the first 250 photoelectrons, and the second muon or muon bundle gives a subsequent large deposition of charge which fulfills the 6000 photoelectron criteria. As such, the event is identifiable as background and removed from the sample. In addition, Event 28 is removed from the point source sample due to an excess of hits observed on the IceTop surface array. If a down-going muon or neutrino originates from a cosmic ray air shower there is the possibility of an excess of hits observed on the IceTop surface array from the same air shower. If such hits on IceTop occur within a causal time window of the in-ice event this suggests that the in-ice event is likely associated with the same cosmic ray, and is therefore identifiable as background.

4.3 Modelling the neutrino flux in IceCube

Simulation of the events that are observed in the HESE analysis requires several levels of consideration. The backgrounds in this analysis are the diffuse flux of astrophysical neutrinos, and the resultant particles from the interaction of cosmic rays in the atmosphere. These cosmic ray air showers produce both muons and neutrinos that are seen within IceCube. The atmospheric muon rate in the HESE selection is suppressed at a level of 99.999%, and above an energy of 60 TeV of event deposited energy in the detector the passing probability over 4 years of data remains negligible compared to other components.

The baseline model of conventional atmospheric neutrinos used is the HKKMS2007 model [36] which is constrained in a previous IceCube analysis [94]. From the same

#ID	E (TeV)	Time (MJD)	Dec ($^{\circ}$)	RA ($^{\circ}$)	Res. ($^{\circ}$)	Topology	PS	E_{cut}
1	$47.6^{+6.5}_{-5.4}$	55351.3222143	-1.8	35.2	16.3	Shower	✓	
2	$117.0^{+15.4}_{-14.6}$	55351.4659661	-28.0	282.6	25.4	Shower	✓	✓
3	$78.7^{+10.8}_{-8.7}$	55451.0707482	-31.2	127.9	<1.4	Track	✓	✓
4	$165.4^{+19.8}_{-14.9}$	55477.3930984	-51.2	169.5	7.1	Shower	✓	✓
5	$71.4^{+9.0}_{-9.0}$	55512.5516311	-0.4	110.6	<1.2	Track	✓	✓
6	$28.4^{+2.7}_{-2.5}$	55567.6388127	-27.2	133.9	9.8	Shower	✓	
7	$34.3^{+3.5}_{-4.3}$	55571.2585362	-45.1	15.6	24.1	Shower	✓	
8	$32.6^{+10.3}_{-11.1}$	55608.8201315	-21.2	182.4	<1.3	Track	✓	
9	$63.2^{+7.1}_{-8.0}$	55685.6629713	33.6	151.3	16.5	Shower	✓	✓
10	$97.2^{+10.4}_{-12.4}$	55695.2730461	-29.4	5.0	8.1	Shower	✓	✓
11	$88.4^{+12.5}_{-10.7}$	55714.5909345	-8.9	155.3	16.7	Shower	✓	✓
12	$104.1^{+12.5}_{-13.2}$	55739.4411232	-52.8	296.1	9.8	Shower	✓	✓
13	$252.7^{+25.9}_{-21.6}$	55756.1129844	40.3	67.9	<1.2	Track	✓	✓
14	$1040.7^{+131.6}_{-144.4}$	55782.5161911	-27.9	265.6	13.2	Shower	✓	✓
15	$57.5^{+8.3}_{-7.8}$	55783.1854223	-49.7	287.3	19.7	Shower	✓	
16	$30.6^{+3.6}_{-3.5}$	55798.6271285	-22.6	192.1	19.4	Shower	✓	
17	$199.7^{+27.2}_{-26.8}$	55800.3755483	14.5	247.4	11.6	Shower	✓	✓
18	$31.5^{+4.6}_{-3.3}$	55923.5318204	-24.8	345.6	<1.3	Track	✓	
19	$71.5^{+7.0}_{-7.2}$	55925.7958619	-59.7	76.9	9.7	Shower	✓	✓
20	$1140.8^{+142.8}_{-132.8}$	55929.3986279	-67.2	38.3	10.7	Shower	✓	✓
21	$30.2^{+3.5}_{-3.3}$	55936.5416484	-24.0	9.0	20.9	Shower	✓	
22	$219.5^{+21.2}_{-24.4}$	55941.9757813	-22.1	293.7	12.1	Shower	✓	✓
23	$82.2^{+8.6}_{-8.4}$	55949.5693228	-13.2	208.7	<1.9	Track	✓	✓
24	$30.5^{+3.2}_{-2.6}$	55950.8474912	-15.1	282.2	15.5	Shower	✓	
25	$33.5^{+4.9}_{-5.0}$	55966.7422488	-14.5	286.0	46.3	Shower	✓	
26	$210.0^{+29.0}_{-25.8}$	55979.2551750	22.7	143.4	11.8	Shower	✓	✓
27	$60.2^{+5.6}_{-5.6}$	56008.6845644	-12.6	121.7	6.6	Shower	✓	
28	$46.1^{+5.7}_{-4.4}$	56048.5704209	-71.5	164.8	<1.3	Track		

Table 4.1: Table of HESE sample events in the first 2 years of observations. Columns shown are event ID, deposited EM-equivalent energy, the time of the event in Mean Julian Day, the event declination, right ascension and angular resolution, and event topology. The two rightmost columns show whether the event was used in the original 4-year analysis, and our selection for the combined diffuse and point source analysis which uses a $10^{4.8}$ GeV cut on deposited energy.

#ID	E (TeV)	Time (MJD)	(°)	RA (°)	Res. (°)	Topology	PS	E_{cut}
29	$32.7^{+3.2}_{-2.9}$	56108.2572046	41.0	298.1	7.4	Shower	✓	
30	$128.7^{+13.8}_{-12.5}$	56115.7283574	-82.7	103.2	8.0	Shower	✓	✓
31	$42.5^{+5.4}_{-5.7}$	56176.3914143	78.3	146.1	26.0	Shower	✓	
32	—	56211.7401231	—	—	—	Coinc.		
33	$384.7^{+46.4}_{-48.6}$	56221.3424023	7.8	292.5	13.5	Shower	✓	✓
34	$42.1^{+6.5}_{-6.3}$	56228.6055226	31.3	323.4	42.7	Shower	✓	
35	$2003.7^{+236.2}_{-261.5}$	56265.1338677	-55.8	208.4	15.9	Shower	✓	✓
36	$28.9^{+3.0}_{-2.6}$	56308.1642740	-3.0	257.7	11.7	Shower	✓	
37	$30.8^{+3.3}_{-3.5}$	56390.1887627	20.7	167.3	<1.2	Track	✓	
38	$200.5^{+16.4}_{-16.4}$	56470.1103795	14.0	93.3	<1.2	Track	✓	✓
39	$101.3^{+13.3}_{-11.6}$	56480.6617877	-17.9	106.2	14.2	Shower	✓	✓
40	$157.3^{+15.9}_{-16.7}$	56501.1641008	-48.5	143.9	11.7	Shower	✓	✓
41	$87.6^{+8.4}_{-10.0}$	56603.1116854	3.3	66.1	11.1	Shower	✓	✓
42	$76.3^{+10.3}_{-11.6}$	56613.2566890	-25.3	42.5	20.7	Shower	✓	✓
43	$46.5^{+5.9}_{-4.5}$	56628.5688531	-22.0	206.6	<1.3	Track	✓	
44	$84.6^{+7.4}_{-7.9}$	56671.8778763	0.0	336.7	<1.2	Track	✓	✓
45	$429.9^{+57.4}_{-49.1}$	56679.2044683	-86.3	219.0	<1.2	Track	✓	✓
46	$158.0^{+15.3}_{-16.6}$	56688.0702948	-22.3	150.5	7.6	Shower	✓	✓
47	$74.3^{+8.3}_{-7.2}$	56704.6001074	67.4	209.4	<1.2	Track	✓	✓
48	$104.7^{+13.5}_{-10.2}$	56705.9419933	-33.2	213.0	8.1	Shower	✓	✓
49	$59.9^{+8.3}_{-7.9}$	56722.4083554	-26.3	203.2	21.8	Shower	✓	
50	$22.2^{+2.3}_{-2.0}$	56737.2004652	59.3	168.6	8.2	Shower	✓	
51	$66.2^{+6.7}_{-6.1}$	56759.2159560	54.0	88.6	6.5	Shower	✓	✓
52	$158.1^{+16.3}_{-18.4}$	56763.5448147	-54.0	252.8	7.8	Shower	✓	✓
53	$27.6^{+2.6}_{-2.2}$	56767.0663034	-37.7	239.0	<1.2	Track	✓	
54	$54.5^{+5.1}_{-6.3}$	56769.0295975	6.0	170.5	11.6	Shower	✓	

Table 4.2: Table of HESE sample events in the last 2 years of observations, the 3rd year consists of events 29 to 37 and 4th year from 38 to 54. Columns shown are the event ID, deposited EM-equivalent energy, the time of the event in Mean Julian Day, the event declination, right ascension and angular resolution, and event topology. The two rightmost columns show whether the event was used in the original 4-year analysis, and our selection for the combined diffuse and point source analysis which uses a $10^{4.8}$ GeV cut on deposited energy.




Figure 4.2: Reconstructed zenith angles and energies of the 4-year HESE sample events. Error bars show the 68% confidence intervals on the reconstructed parameters. The Deposited EM-equivalent energy gives the light deposition as if all light were deposited due to electromagnetic showers. All deposited energies are therefore lower limits on the neutrino energy (see Section 3.5. Figure from [89].

reference the prompt atmospheric flux model used is the ERS model [37]. These models are adjusted to a specific cosmic ray flux model H3a which is a parameterisation of cosmic ray energies and masses which includes the change of cosmic ray flux at the “knee” [95].

To simulate the interactions of atmospheric and astrophysical neutrinos, the software “neutrino-generator” (NuGen), which is based on the program ANIS [96] used for the AMANDA detector, simulates the interaction of neutrinos produced isotropically at the surface of the Earth and then propagated to interact near the detector volume. An interaction weight is computed to account for the forcing of the neutrino interaction and the survival probability through the Earth, which is dependent on the Earth density which we approximate with the Preliminary Reference Earth Model [97].

To assess the atmospheric neutrino suppression rate, which is due to the presence of accompanying muons which trigger the veto condition, the CORSIKA software package [35] was used. CORSIKA is used to determine the number, type, energy,

position and direction of air shower particles observed in cosmic ray air showers. To estimate the veto rate only hadronic interactions at high energies needed to be considered, and the calculations of the atmospheric neutrino suppression factor took advantage of such simplifications. The suppression factor is then applied to the down-going atmospheric neutrino flux, with a maximum suppression of 90% to allow for a minimum 10% unsuppressed fraction to allow for uncertainties in hadronic interactions. The calculation of the suppression is detailed in [40].

Muons, both from the air showers and from CC ν_μ and ν_τ neutrino interactions, have periodic interactions simulated by the software PROPOSAL [98]. Light propagation from the interactions of all particles to the simulated DOM electronics is performed on GPUs using the software “PPC” and “CLSIM”.

4.4 Current HESE results

4.4.1 Fit for a diffuse neutrino flux

To assess the presence of a diffuse flux of astrophysical neutrinos in the data, the background-only hypothesis that the neutrinos in the sample arrived solely from atmospheric sources was tested against an alternative hypothesis that neutrinos could also be produced by a uniform distribution of neutrino sources with an $E^{-\gamma}$ spectrum. The expected background distributions were entirely due to particles produced in cosmic ray air showers, and consisted of atmospheric muons, atmospheric neutrinos due to π/K meson decay (conventional atmospheric neutrinos), and neutrinos due to charged meson decay (prompt atmospheric neutrinos).

The atmospheric muon rate in the event selection can be estimated by the data-driven process of muon tagging - a similar second internal region was selected within the existing detection volume as a new veto. The original outer veto was used to identify events that are not observed within the new inner veto. To determine the overall rate in the full experiment, the observed rate that pass the inner veto but




Figure 4.3: Fitted distributions (shaded regions) of possible event sources to observed data (black points with error bars) in deposited EM-equivalent energy. Figure from [89].

are observed (tagged) on the outer veto is measured. The rate observed within this small internal volume within the second veto is scaled to the internal volume of the full experiment, which is a factor of two times larger [88]. From this measurement we predict a total rate of 12.6 ± 5.1 muon events in 4 years of data, with very few muon events expected above 60 TeV.

These distributions were fit to the data (see Figures 4.3 and 4.4) for events over $10^{4.8}$ GeV (~ 60 TeV) over zenith angle, deposited energy and event topology (flavour). The distribution of the deposited energy, zenith angle and flavour of the events all suggest an astrophysical component (Figure 4.2). The spectrum of atmospheric neutrinos, both from π/K meson decay (conventional) and from charmed meson decay (prompt) is too steep to explain the data (Figure 4.3). The zenith distribution shows an excess from the southern hemisphere, which is heavily disfavoured due to the muon veto as we expect a large proportion of neutrinos from air showers to be accompanied by coincident muons that would be observed in the outer layers of the detector (Figure 4.4). Finally, the observed track to cascade ratio is consistent with a mixed distribution of neutrino flavours and differs substantially from that of conventional atmospheric neutrinos. Table 4.3 shows the relative expected rates of neutrino events in terms of track and cascade topology above $10^{4.8}$ GeV.

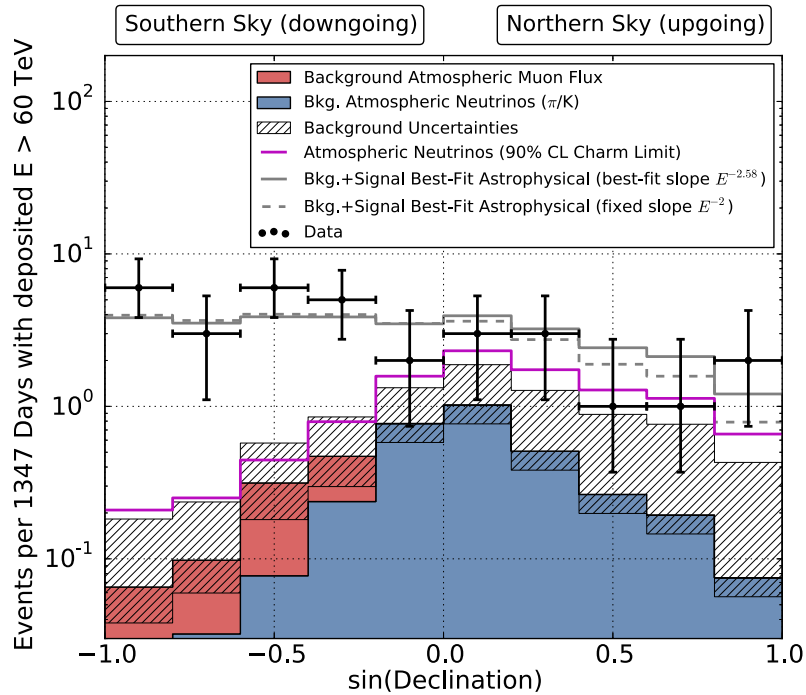


Figure 4.4: Fitted distributions (shaded regions) of possible event sources to observed data (black points with error bars) in zenith angle. Figure from [89].

	E^{-2}	$E^{-2.6}$	conventional	prompt
tracks	5.00	1.89	1.87	0.29
cascades	20.06	9.86	0.85	1.70
ratio	0.25	0.19	2.21	0.17

Table 4.3: Predicted rates of track events, cascade events and the track to cascade ratio above $10^{4.8}$ GeV, for the 4 years of HESE live-time. For the diffuse astrophysical flux both the nominal E^{-2} spectrum and 4-year best fit $E^{-2.6}$ spectrum are shown. For the E^{-2} astrophysical flux the normalisation is 1.0×10^{-8} $\text{GeV s}^{-1} \text{cm}^{-2} \text{str}^{-1}$ per neutrino flavour. The normalisation for the $E^{-2.6}$ flux is 1.0×10^{-18} $(E/100 \text{ TeV})^{-2.6}$ $\text{GeV}^{-1} \text{s}^{-1} \text{cm}^{-2} \text{str}^{-1}$ per neutrino flavour. For the conventional and prompt fluxes the nominal model predictions from the HKKMS07 and ERS models respectively are shown.

Of note is that unlike uncontained analyses with up-going neutrino events, the interaction vertex of all neutrinos including up going neutrinos is known in this sample. This property places a better upper limit on the neutrino energy for muon neutrino interactions than other analyses where a low energy muon seen at the detector can be the product of a much higher energy neutrino which interacted at a distance away from the detector.

With the initial 2 years of data (28 events), the significance of the excess above $10^{4.8}$ GeV above an atmospheric-only hypothesis was found to be 4.8σ on all events




Figure 4.5: Contours of confidence levels for the normalisation of astrophysical flux at 100 TeV (Φ_{astro}), and spectral index (γ_{astro}) for 3 and 4 years of data. The fit uses a prior on the prompt flux from previous upper limits [94]. This was not used in the original 3-year publication [99], the prior is used in both cases here for direct comparison. Figure from [89].

tested. The test for significance using all events is not an a priori result as the two highest energy events were already observed by the EHE selection. All normalisations except the atmospheric muon flux were allowed to float in the fit without priors, and in this case the astrophysical flux tested was a fixed E^{-2} astrophysical spectrum. The muon flux is constrained by a prior based on the predicted muon rate from data. As a conservative check due to the best fit prompt flux fitting to 0, the data was also fit with a pre-determined 90% upper limit on the charm flux, and find the test of an astrophysical flux in this case gives a 4.5σ significance. With an additional year of data collection, now totalling 3 years of data, the same calculation leads to a significance of 5.7σ , which includes a scaling of the prompt flux in the background-only fit to 3.6 times the existing prompt upper limit [99].

With 4 years of data the focus shifted from the confirmation, to the characterisation, of the diffuse astrophysical flux. To better resolve the ambiguity between the diffuse astrophysical flux and the prompt atmospheric neutrino flux a

prior was used to constrain the prompt normalisation based on the upper limit at 90% confidence from [94]. With this prior incorporated in the fit, a parameter scan over astrophysical flux normalisation and spectral index was performed and the best fit likelihood obtained for each point in the space. Performing this test with 4 years of data, better bounds on the nature of the astrophysical flux in total normalisation and spectrum were obtained, the results of which can be seen in Figure 4.5. The overall best fit including a variable spectral index is a flux of $\phi(E) = 2.2 \pm 0.7 \times 10^{-18} (E/100 \text{ TeV})^{-2.58} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$.

4.4.2 Fit for point sources of neutrinos

To fit for a point source of neutrinos, we use a likelihood. A likelihood is a probability function over the model parameter space, with the observables (data) held fixed, as opposed to a probability distribution function, which gives us the probability of the observables given a fixed model. The following point source likelihood \mathcal{L} , which assesses clustering of the data with the model of a point source and a single background distribution, was used to analyse the HESE data:

$$\ln \mathcal{L}(\vec{x}_s) = \sum_i^N \ln \left[\frac{n_s}{N} S_i(\vec{x}_s) + \left(1 - \frac{n_s}{N}\right) B_i \right] \quad (4.1)$$

Here, n_s is the number of model point source events out of N total events, S_i is the spatial point spread function (PSF) of event i , and this PSF is evaluated in the direction of the point source \vec{x}_s . These PSFs were determined using a grid scan of the sky for each event using reconstruction techniques discussed in 3.5, i.e. the error function on the event direction is derived from data, with additional widening added to account for systematic uncertainties.

A uniform background of $B_i = 1/4\pi$ is selected in this analysis as there is limited data with which to model the background from data alone [99].

The test statistic used is the ratio of the best fit signal hypothesis to the background-




Figure 4.6: All-sky scan of the test statistic using a uniform background in the point source likelihood, calculated on 4 years of HESE data. Cascades are marked with a + and tracks are marked with a \times . Figure from [89].

only hypothesis:

$$\text{TS} = -2 \ln \frac{\mathcal{L}(\hat{n}_s)}{\mathcal{L}(n_s = 0)} \quad (4.2)$$

and the post-trial p-value is determined by comparing the overall largest TS value on the sky, termed the “hotspot”, to the overall largest TS or hotspot of each of an ensemble of background-like skies, these background-like skies are produced by separately randomising the right ascension of each event. This method of determining the p-value is discussed further in Section 5.2.7. The post-trial p-value is the percentage of skies with the test statistic of the hotspot that are greater than that of the actual experimental data.

An important effect was found as a result of the point source search: the emphasis of muons with good angular resolutions in the point source likelihood map over the likelihood contribution of poorer resolution cascades. The result of this discrepancy is that a track on its own can appear as a stronger point source than a region of several cascades, independent of the fact that an isolated track on the sky can often be attributed to backgrounds (including the diffuse astrophysical background), whereas a cascade cluster could - albeit with poor resolution - indicate a source

location to which multiple events could be attributed.

It is important to note that the poor cascade resolutions are the result of the difficulties of resolving the event direction from the point-like source of light observed within the detector and are not fundamental limitations - in principle the cascade resolution could be known to high precision with improved knowledge of light propagation in the ice. The result is that these two topologies which are related to the neutrino flavour behave differently in the point source likelihood. To better utilise the cascade contributions, the point source likelihood is also calculated without tracks, this is termed a cascade-only point source search.

The results from this analysis on the full 4 years of HESE data give a post-trial p-value of 58% for all events, and 44% for the cascade-only subset of events.

To improve this analysis and future analyses of low-count point source data, an improved methodology is proposed in Chapter 5, investigated in Chapters 7 and 8, and applied to the data in Chapter 9.

4.5 Summary

In this chapter we discussed the High Energy Starting Event analysis, and the results from this analysis on the four years of data applicable to this thesis. We note that this analysis was motivated by the detection of two bright contained events, and discuss how the detector was utilised to select for these neutrino interactions within the detector toward lower energies while maintaining a high astrophysical purity. The model fits to the diffuse astrophysical flux show a statistically significant detection, with a steep spectral index of approximately $E^{-2.6}$. The fit for a point source, which uses a simple test of a uniform background in the likelihood, does not find a statistically significant “hotspot” of neutrino events.

In the next chapter we define the combined diffuse and point source (Diffuse+PS) analysis, which aims to improve upon the assumption of a uniform background in the likelihood by using simulation to describe the diffuse backgrounds in the fit for

a point source.

Chapter 5

Combined diffuse and point source analysis

The best-fit event directions and the likelihood scan over the sky of the previous point source analysis on the High Energy Starting Event (HESE) selection [89] can be seen in Figure 4.6 of Chapter 4. As there are too few events in the HESE sample to accurately estimate the background energy-zenith distribution from data, a uniform $1/4\pi$ background was used when assessing each event in the sample. As this uniform background contains no information of the energy, zenith and flavour distributions of the events, much of the available information associated with the HESE sample was not utilised in the formulation of the point source likelihood.

We wish to add to the point source search of neutrinos in the HESE sample features which can discriminate astrophysical point source signal events from background events. Neutrinos from astrophysical sources can be distinguished from atmospheric neutrino backgrounds by virtue of

- the energy of the neutrinos (the more energy an event has, the more likely it is to be astrophysical), and
- the incident zenith angle of the event (as zenith distributions of the atmospheric backgrounds are shaped by the veto and Earth absorption), and

- the topology of the events: for this analysis, whether the event has a cascade or track topology. A cascade only search, lacking tracks, therefore does not use topology as a discriminant.

A point source of astrophysical neutrinos can be distinguished from a background of diffuse astrophysical and atmospheric neutrinos by identifying a localised excess above the predicted background rate at a position on the sky.

We include the expected event distributions from atmospheric and astrophysical neutrinos in a combined fit. This fit is a combination of diffuse astrophysical neutrinos following the best fit diffuse spectrum of $E^{-2.6}$, an astrophysical point source also with an $E^{-2.6}$ spectrum, conventional atmospheric neutrinos and prompt atmospheric neutrinos. By including this information to analyse the expected distribution of events we aim to increase our power to observe a point source in the data.

The dependence of this analysis on topology served as a motivating feature, as it was thought that by fairly attributing tracks to atmospheric backgrounds we would avoid the issues of track and cascade resolution seen in the previous point source search (Section 4.4.2), where it was observed that isolated tracks that could reasonably be attributed to atmospheric backgrounds were being attributed a larger point source likelihood than clusters of cascades due to their superior angular resolution.

In order to correctly assess the presence of a point source on the sky, the assumed background distributions should be physically motivated. By using our simulated background distributions (Figures 5.1 - 5.8), we are assessing the point source strength fairly without under- or over-emphasising possible point source locations on the sky.

In comparing this new analysis to a test of a point source which uses a uniform background in the likelihood, it was observed that the assumption of uniform background results in a bias. This observed bias is discussed in depth in Chapter 7, but in essence it was observed that the zenith distribution on the sky results in an under-fitting of a point source in regions where the fitted background over-predicts the true background rate, and over-fits when the fitted background under-predicts

the true background rate, which is the case when fitting the HESE sample events with a uniform background, as the events do not follow a uniform zenith distribution. This has implications for the sensitivity and discovery potential of the analysis when assessing the “hotspot” on the sky.

The combined diffuse and point source analysis (here also referred to as the Diffuse+PS analysis) can be thought of as either an extension of the point source analysis to include a full description of the event backgrounds, or an extension of the fit for a diffuse astrophysical flux to include a point source component.

5.1 Event reconstructions

In the point source likelihood we incorporate the all-sky point spread function (PSF) of each event. The HESE sample events have been reconstructed with the maximum likelihood techniques from Chapter 3. To account for the uncertainty of scattering and absorption in the ice and the effect this has on angular reconstruction, the PSFs are convolved with an additional angular uncertainty determined for each event. This correction is typically of order 1° for track events and 10° for cascade events.

Such scans for the best fit reconstructed direction are performed over a finite grid. We use the HEALPix [100] software package to define the grid over the sky, giving an equal area binning of the event PSFs. Each bin on this grid is termed a “pixel”. These maximum likelihood scans over the sky can be directly converted to a per-pixel probability distribution over the sky which we call a point spread function or PSF. The PSF for each event i over all pixels j is denoted as ρ_{ij} , with $\sum_j^{N_{pix}} \rho_{ij} = 1$, which is the normalisation condition for each event i : the probability over all possible directions sums to 1.

5.2 Combined diffuse and point source likelihood

The combined diffuse and point source likelihood (also denoted as the Diffuse+PS likelihood), which considers all background diffuse components on the sky in addition to a point source contribution and tests their combined fit to the events, can be expressed as

$$\ln \mathcal{L} = \sum_{i=1}^{N_{\text{obs}}} \ln [n_s S_i(\vec{x}_s) + n_a A_i + n_c C_i + n_p P_i] - [n_s + n_a + n_c + n_p] - \frac{(n_c - \mu_{n_c})^2}{2\sigma_{n_c}^2} - \frac{(n_p - \mu_{n_p})^2}{2\sigma_{n_p}^2} \quad (5.1)$$

A full derivation and description of Equation 5.1 is outlined below.

5.2.1 Derivation of unbinned Poisson likelihood

Fundamentally, in this analysis we adapt the diffuse fit which uses a binned Poisson likelihood (denoted as \mathcal{L}), by taking the unbinned limit of the Poisson likelihood, and adding a point source term.

To derive the unbinned Poisson likelihood we begin with the binned Poisson likelihood \mathcal{L} , which is the product of Poisson probabilities of observing x_i events given a prediction of λ_i predicted events in bin i :

$$\mathcal{L} = \prod_{i=1}^{N_{\text{bins}}} \frac{\lambda_i^{x_i} e^{-\lambda_i}}{x_i!} \quad (5.2)$$

Taking the log of the Poisson likelihood, we obtain

$$\ln \mathcal{L} = \sum_{i=1}^{N_{\text{bins}}} x_i \ln \lambda_i - \lambda_i - \ln(x_i!) \quad (5.3)$$

where the $\ln(x_i!)$ term can be disregarded for the likelihood as it does not depend on the model parameters λ_i . In addition, $\ln(x_i!)$ will reduce to zero for $x_i = 0$ and $x_i = 1$.

If we take the case of infinitely small bins there must either be 0 or 1 events per bin, implying $x_i = 0$ in unoccupied bins and $x_i = 1$ for bins with an event in them. Therefore:

- If $x_i = 0$, the contribution to $\ln \mathcal{L}$ for bin i is: $-\lambda_i$
- If $x_i = 1$, the contribution to $\ln \mathcal{L}$ for bin i is: $\ln \lambda_i - \lambda_i$

Splitting the sum into cases where $x_i = 0$ and $x_i = 1$, we obtain

$$\ln \mathcal{L} = \sum_{i=1, x_i=1}^{N_{\text{bins}}} (\ln \lambda_i - \lambda_i) - \sum_{i=1, x_i=0}^{N_{\text{bins}}} \lambda_i \quad (5.4)$$

which can be rewritten as

$$\ln \mathcal{L} = \sum_{i=1, x_i=1}^{N_{\text{bins}}} \ln \lambda_i - \sum_{i=1, x_i=1}^{N_{\text{bins}}} \lambda_i - \sum_{i=1, x_i=0}^{N_{\text{bins}}} \lambda_i \quad (5.5)$$

and since $x_i = 0$ and $x_i = 1$ cover all cases, the right terms can be combined into a single sum:

$$\ln \mathcal{L} = \sum_{i=1, x_i=1}^{N_{\text{bins}}} \ln \lambda_j - \sum_{i=1}^{N_{\text{bins}}} \lambda_i \quad (5.6)$$

which we denote as λ_{total} :

$$\lambda_{\text{total}} = \sum_{i=1}^{N_{\text{bins}}} \lambda_i \quad (5.7)$$

The bins where $x_i = 1$ are the bins with observed events, hence we define a new index j for all cases $x_i = 1$, running from 1 to N_{obs} events. Rewritten in this form, Equation 5.6 becomes

$$\ln \mathcal{L} = \sum_{i=1}^{N_{\text{obs}}} \ln \lambda_i - \lambda_{\text{total}} \quad (5.8)$$

From here we expand the λ_i and λ_{total} terms to show how our underlying fluxes change the likelihood.

5.2.2 Application of the unbinned Poisson likelihood to Diffuse+PS analysis

The combined diffuse and point source likelihood is derived from the unbinned limit of a Poisson likelihood (eq. 5.8).

The term λ_i is the predicted rate for the i^{th} event, and λ_{total} is the total predicted rate of events. In determining λ_i we are assessing how frequently the event seen is produced by our model. From simulation, we can produce a complete model describing the rate of events given a point source direction \vec{x}_s , deposited energy E , event topology f and event direction j (which is discretised by HEALPix pixel). This distribution is denoted here as $\Lambda(\vec{x}_s, E, f)_j$. For an event which has a corresponding event deposited energy E and flavour f , we weight the i th event PSF ρ_{ij} by this distribution $\Lambda(\vec{x}_s, E, f)_j$ to determine the predicted event rate λ_i .

$$\lambda_i = \sum_{j=1}^{N_{\text{pix}}} \rho_{ij} \Lambda(\vec{x}_s, E, f)_j \quad (5.9)$$

This predicted event rate λ_i is used in our likelihood in Equation 5.8.

Expected event rate distributions

We incorporate the expected rate of neutrino events in flavour, energy, right ascension and declination from atmospheric and astrophysical origins into the fit. The distributions can be seen in Figures 5.1 - 5.2 for the conventional flux, Figures 5.3 - 5.4 for the prompt atmospheric flux, and Figures 5.5 - 5.8 for the astrophysical flux assuming a nominal E^{-2} flux (5.5 - 5.6) and best-fit $E^{-2.6}$ flux (5.7 - 5.8). On its own, Λ is defined as the linear combination of all fit components. Our model has four components in the fit, giving

$$\begin{aligned} \Lambda(\vec{x}_s, E, f)_j = & n_a \Phi_{\text{astro}}(E, f)_j + n_c \Phi_{\text{conv}}(E, f)_j \\ & + n_p \Phi_{\text{prompt}}(E, f)_j + n_s \Phi_{\text{ptsrc}}(\vec{x}_s, E, f)_j \end{aligned} \quad (5.10)$$

where all Φ terms are the now-normalised distributions of event rates over $(E, f)_j$ space, and all n terms are the relative event rates from each source and background: n_a and $\Phi_{\text{astro}}(E, f)_j$ are the event rate and shape of the diffuse astrophysical neutrino component, n_c and $\Phi_{\text{conv}}(E, f)_j$ are the event rate and shape of the conventional atmospheric neutrino component, n_p and $\Phi_{\text{prompt}}(E, f)_j$ are the event rate and shape of the prompt astrophysical neutrino component, and n_s and $\Phi_{\text{ptsrc}}(\vec{x}_s, E, f)_j$ are the event rate and shape of the astrophysical point source component.

The Φ_{astro} , Φ_{conv} and Φ_{prompt} distributions take non-zero values over the whole sky, while Φ_{ptsrc} is only non-zero at the point source position \vec{x}_s and has a probability distribution function over deposited energy and event topology from the diffuse astrophysical component $\Phi_{\text{astro}}(E, f)_j$ at the zenith angle of the source, and this probability distribution is rescaled to a normalisation of one.

It is unphysical for a prediction of atmospheric or astrophysical neutrinos to be a negative contribution to the total neutrino flux. Therefore, a constraint is imposed that no parameter n can be negative, with a minimum of zero.

Total Predicted Rate λ_{total}

The fit of each model component to all events also must result in a total predicted rate of events across the sky from each source: conventional, prompt, astrophysical diffuse and point source. This total is the term λ_{total} in Eq. 5.8.

$$\lambda_{\text{total}} = n_a \sum_{E,f,j} \Phi_a + n_c \sum_{E,f,j} \Phi_c + n_p \sum_{E,f,j} \Phi_p + n_s \sum_{E,f,j} \Phi_s \quad (5.11)$$

and, since all Φ terms are normalised,

$$\lambda_{\text{total}} = n_a + n_c + n_p + n_s \quad (5.12)$$

If no constraints or priors are enforced on any flux in the Poisson fit, we find the fit normalises to the total event rate, $\lambda_{\text{total}} = N$, however, constraints are used here to result in plausible conventional and prompt rates and maintain consistency with

existing analyses.

5.2.3 Simplifying the Diffuse+PS likelihood

Putting all of these terms together, we can rewrite Equation 5.8 as

$$\begin{aligned} \ln \mathcal{L} = \sum_{i=1}^{N_{\text{obs}}} \ln & \left[\sum_j^{N_{\text{pix}}} (n_s \rho_{ij} \Phi_{\text{ptsrc}}(\vec{x}_s, E_i, f_i)_j + n_a \rho_{ij} \Phi_{\text{astro}}(E_i, f_i)_j \right. \\ & \left. + n_c \rho_{ij} \Phi_{\text{conv}}(E_i, f_i)_j + n_p \rho_{ij} \Phi_{\text{prompt}}(E_i, f_i)_j \right] \\ & - [n_s + n_a + n_c + n_p] \end{aligned} \quad (5.13)$$

We define new terms $S_i(\vec{x}_s)$, A_i , C_i and P_i as shorthand for the individual contributions to λ_i :

$$\sum_j^{N_{\text{pix}}} \rho_{ij} \Phi_{\text{ptsrc}}(\vec{x}_s, E_i, f_i)_j = S_i(\vec{x}_s, E_i, f_i, \rho_i) \equiv S_i(\vec{x}_s) \quad (5.14)$$

$$\sum_j^{N_{\text{pix}}} \rho_{ij} \Phi_{\text{astro}}(E_i, f_i)_j = A_i(E_i, f_i, \rho_i) \equiv A_i \quad (5.15)$$

$$\sum_j^{N_{\text{pix}}} \rho_{ij} \Phi_{\text{conv}}(E_i, f_i)_j = C_i(E_i, f_i, \rho_i) \equiv C_i \quad (5.16)$$

$$\sum_j^{N_{\text{pix}}} \rho_{ij} \Phi_{\text{prompt}}(E_i, f_i)_j = P_i(E_i, f_i, \rho_i) \equiv P_i \quad (5.17)$$

From this point we use the far-right terms as shorthand. With this notation the likelihood (without prior terms) becomes

$$\ln \mathcal{L} = \sum_{i=1}^{N_{\text{obs}}} \ln [n_s S_i(\vec{x}_s) + n_a A_i + n_c C_i + n_p P_i] - [n_s + n_a + n_c + n_p] \quad (5.18)$$

5.2.4 Comparing the Diffuse+PS likelihood to the previous point source likelihood

In the event that the parameters in the likelihood can freely float in the fit, we have a naturally imposed constraint from the Poisson likelihood that $n_s + n_a + n_c + n_p = N = N_{\text{obs}}$; that is, the predicted event rate will match the total rate when the likelihood is maximised (see Appendix A for proof) This means one of the parameters, we choose n_p here, can be expressed in terms of the others: $n_p = N - n_s - n_a - n_c$. Applying this, in addition to a slight rearranging of the equation, we obtain

$$\ln \mathcal{L} = \sum_{i=1}^{N_{\text{obs}}} \ln \left[\frac{n_s}{N} S_i(\vec{x}_s) + \frac{n_a}{N} A_i + \frac{n_c}{N} C_i + \left(1 - \frac{n_s + n_a + n_c}{N}\right) P_i \right] + N \ln N - N \quad (5.19)$$

The term $N \ln N$ results as we must remove $\ln N$ a total of N times due to the sum. As the far right term $N \ln N - N$ is constant, it can be dropped from the likelihood. This gives

$$\ln \mathcal{L} = \sum_{i=1}^{N_{\text{obs}}} \ln \left[\frac{n_s}{N} S_i(\vec{x}_s) + \frac{n_a}{N} A_i + \frac{n_c}{N} C_i + \left(1 - \frac{n_s + n_a + n_c}{N}\right) P_i \right] \quad (5.20)$$

which is very similar to Equation 4.1, a point source and a background term.

However, the property that $n_s + n_a + n_c + n_p = N$ no longer holds when any non-zero constraints or priors are placed on scaling parameters, unless they are constrained precisely to zero such as in the case of $n_s = 0$, as this is equivalent to having one fewer parameters in the fit.

5.2.5 Gaussian priors

With limited data, the sensitivity to a point source is lowered by having all parameters free in the fit. For example, events from a point source at the horizon can result in a fit which favours an atmospheric origin, with an atmospheric flux in excess of the constraints of other analyses. To avoid this we wish to constrain our fit with

a prior function on the conventional flux of 1 ± 0.3 times the nominal HKKMS07 conventional flux [36], and a prior of 1 ± 3.0 times the prompt ERS flux [37]. Both fluxes are adjusted in accordance with the H3a model (see Section 4.3). The priors used are taken from a high energy cascade analysis [101], and are conservative in the sense that tighter limits on these fluxes exist from other IceCube analyses - the purpose of these priors is simply to limit extreme fits. The astrophysical background flux (n_a) and the point source flux (n_s) are left as freely floating parameters in the fit (or equivalently they are said to have a uniform prior).

We can incorporate Gaussian prior probabilities into the combined diffuse and point source likelihood (Equation 5.18) to govern the normalisation of our atmospheric flux scaling terms n_c and n_p . The Poisson likelihood, \mathcal{L}' , including a prior for such a scaling term n therefore becomes

$$\mathcal{L}' = ae^{-\frac{(n-\mu)^2}{2\sigma^2}} \prod_{i=1}^{N_{\text{bins}}} \frac{\lambda_i^{x_i} e^{-\lambda_i}}{x_i!} \quad (5.21)$$

where a is a normalisation constant, and μ and σ are the mean and standard deviation of the prior term (e.g. for the conventional flux $\mu = \mu_{n_c}$ is 1 times the predicted event rate in the sample from the HKKMS07 conventional flux model, and $\sigma = \sigma_{n_c}$ is 0.3 times the same rate). Taking the logarithm of Equation 5.21, we obtain

$$\ln \mathcal{L}' = \ln a - \frac{(n-\mu)^2}{2\sigma^2} + \sum_{i=1}^{N_{\text{bins}}} [x \ln \lambda_i - \lambda_i + \ln x_i!] \quad (5.22)$$

and the term $\ln a$, being a constant term independent of fit parameters, can be left out of our formulation.

Including prior terms on both conventional and atmospheric fluxes in our analysis, the likelihood for our analysis becomes

$$\begin{aligned} \ln \mathcal{L} = & \sum_{i=1}^{N_{\text{obs}}} \ln [n_s S_i(\vec{x}_s) + n_a A_i + n_c C_i + n_p P_i] - [n_s + n_a + n_c + n_p] \\ & - \frac{(n_c - \mu_{n_c})^2}{2\sigma_{n_c}^2} - \frac{(n_p - \mu_{n_p})^2}{2\sigma_{n_p}^2} \end{aligned} \quad (5.23)$$

where μ_{n_c} and σ_{n_c} are the mean and width of the conventional flux prior and μ_{n_p} and σ_{n_p} are the mean and width of the prompt flux prior. The effect of these priors is that the likelihood becomes more negative, in other words a worse fit, when the fitted conventional and prompt fluxes deviate away from their nominal values.

5.2.6 Test statistic

The question we wish to address with this analysis is “Does a single point source explain any excess of events at a significant level above the background-only hypothesis?”

From this statement we can define a general test statistic [102]

$$\text{TS} = -2 \ln \frac{P(\text{data}|H_0)}{P(\text{data}|H_s)} \quad (5.24)$$

which takes the ratio of two hypotheses:

- H_0 : The null hypothesis that the data are the result only of our diffuse backgrounds of atmospheric neutrinos and a diffuse astrophysical flux.
- H_s : The alternative hypothesis that the data are the result of a single astrophysical point source with direction \vec{x}_s in addition to our diffuse backgrounds.

Given that the null hypothesis is the special case of the alternative hypothesis in our model where $n_s = 0$, the test statistic can be expressed as

$$\text{TS} = -2 \ln \frac{\mathcal{L}(n_s = 0)}{\mathcal{L}(\hat{n}_s)} \quad (5.25)$$

which compares the best fit likelihood $\mathcal{L}(\hat{n}_s)$ to the no-signal hypothesis $\mathcal{L}(n_s = 0)$. We find the best fit of n_s which maximises Equation 5.1, with n_s and all other fit parameters constrained to take positive values.

For our case with multiple fit parameters, this can be fully expressed as

$$\text{TS} = -2 \ln \frac{\mathcal{L}(n_s = 0, \hat{n}_a, \hat{n}_c, \hat{n}_p)}{\mathcal{L}(\hat{n}_s, \hat{n}_a, \hat{n}_c, \hat{n}_p)} \quad (5.26)$$

where the background model parameters n_a, n_c, n_p are optimised separately in both statistics denoted by \hat{n}_x and $\hat{\hat{n}}_x$ where x is a generic subscript $x = a, c$ or p .

An all-sky search for the maximum test statistic is performed by evaluating TS at each position on a HEALPix [100] grid covering the sky. HEALPix maps can be adjusted to different numbers of pixels dividing the sky giving different resolutions, the possible pixel counts are $N_{\text{pixels}} = 12N_{\text{side}}^2$ where N_{side} is a parameter that controls the step-size between different grids, and is restricted to the values $N_{\text{side}} = 2^k$ where k is an integer, therefore $N_{\text{side}} = 1, 2, 4, 8, \text{ etc.}$ For our scans we use an N_{side} value of 32, corresponding to 12,288 pixels with an average separation of 1.83° . The maximum value of TS is then compared to the maximum values of TS obtained from an ensemble of background-like skies to determine how frequently the observed value occurs, and in doing so we determine whether or not the result is a significant detection of a point source. ‘‘Scrambling’’ of the data is one method of generating background-like skies.

5.2.7 Probability of the test statistic under scrambling

The largest value of TS on the sky is used as the statistic to calculate the final significance of the data against background-only pseudo-experiments. As a conservative measure of significance, we determine the frequency of obtaining the value of TS observed in the data against the corresponding maximum TS in maps with the events randomised in right ascension to assess the rate of accidental clustering of the data. This technique is called ‘‘scrambling’’ of the data, a common technique when estimating the significance of a point source.

For a detector with constant uptime and a constant acceptance function on the sky as a function of time, the total rate of events as a function of declination should be independent of right ascension. The reasoning for this is that the sky will rotate overhead, and so no one value of right ascension is unique as over the course of a celestial day and for a given declination all right ascension values will be observed. IceCube is a special case of this: as the detector is located at the geographic south

pole, the right ascension of an event depends only on its azimuthal angle and not zenith. IceCube has highly uniform azimuthal acceptance with any irregularities averaged out quickly over short time-scales as the celestial sky rotates around the detector. This makes randomising each event's right ascension a practical way of generating background-only pseudo-experiments, and, as such, this was also used in the previous HESE analyses [88] [89] [99] and is a routinely used technique for the point source analyses of IceCube and other observatories.

Another potential method of generating background-only pseudo-experiments is to randomly sample background-only skies from the diffuse best fit hypothesis. We show in Chapter 8 that the differences in sensitivity between scrambling the data and testing against sampled skies are minimal, with noticeable differences seen at the celestial poles where scrambling in right ascension has minimal impact on event position. Therefore scrambling, being a more robust estimator of significance as it is less dependent on systematic uncertainties, is used to assess the final significance of the data, with simulated backgrounds being used within the likelihood fit to determine the test statistic.

5.3 Summary

In this chapter we defined the combined diffuse and point source (Diffuse+PS) analysis, deriving the likelihood from a Poisson distribution over infinitely many bins, and drawing on the energy, zenith angle and topology of each event to discriminate atmospheric backgrounds from events of an astrophysical origin. Priors were also added to the likelihood function to constrain the atmospheric neutrino background fit to within reasonable event rates. We also reintroduce the concept of the likelihood test as it pertains to this analysis, and discuss the scrambling method by which we determine the significance of the test statistic. The distributions determined from simulation of the neutrino backgrounds are also shown.

In the next chapter we discuss the data manipulation techniques used which are

required to perform statistical tests of the analysis.

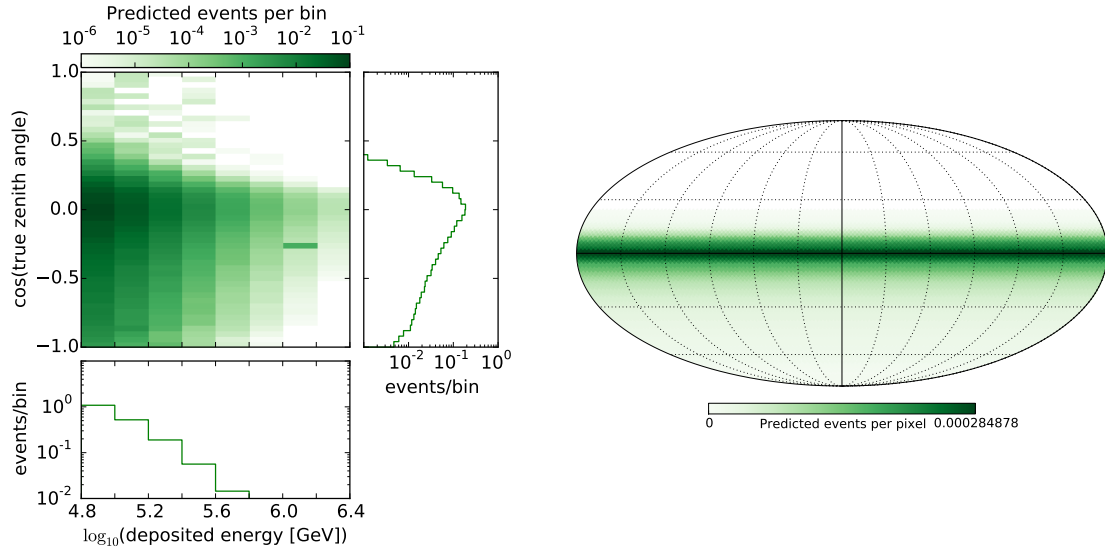


Figure 5.1: Conventional atmospheric neutrino track expectations. Left: The distribution of predicted track events in reconstructed deposited energy and true zenith angle, with corresponding 1D projections. Right: The track event distribution in detector zenith and azimuth using a HEALPix grid. Summing over all pixels on this grid for both tracks and cascades gives a combined normalisation of 1, such that the track to cascade ratio is preserved.

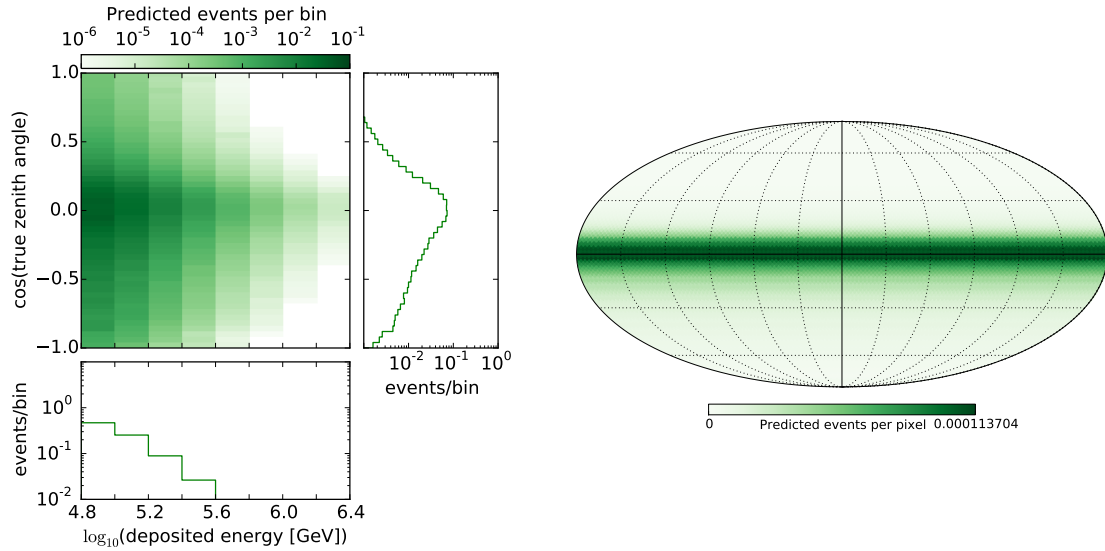


Figure 5.2: Conventional atmospheric neutrino cascade expectations. Left: The distribution of predicted cascade events in reconstructed deposited energy and true zenith angle, with corresponding 1D projections. Right: The cascade event distribution in detector zenith and azimuth using a HEALPix grid. Summing over all pixels on this grid for both tracks and cascades gives a combined normalisation of 1, such that the track to cascade ratio is preserved.

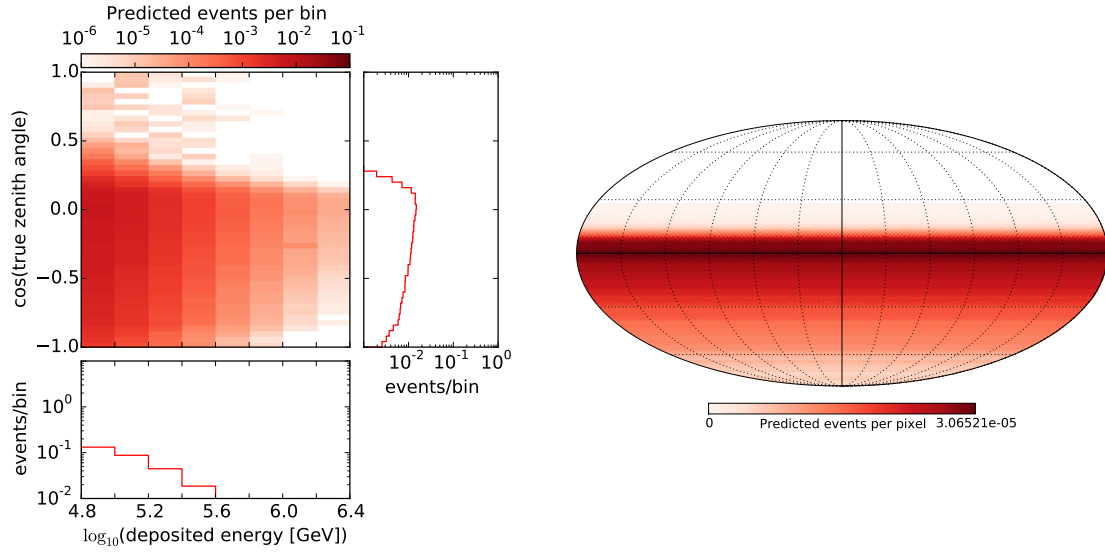


Figure 5.3: Prompt atmospheric neutrino track expectations. Left: The distribution of predicted track events in reconstructed deposited energy and true zenith angle, with corresponding 1D projections. Right: The track event distribution in detector zenith and azimuth using a HEALPix grid. Summing over all pixels on this grid for both tracks and cascades gives a combined normalisation of 1, such that the track to cascade ratio is preserved.

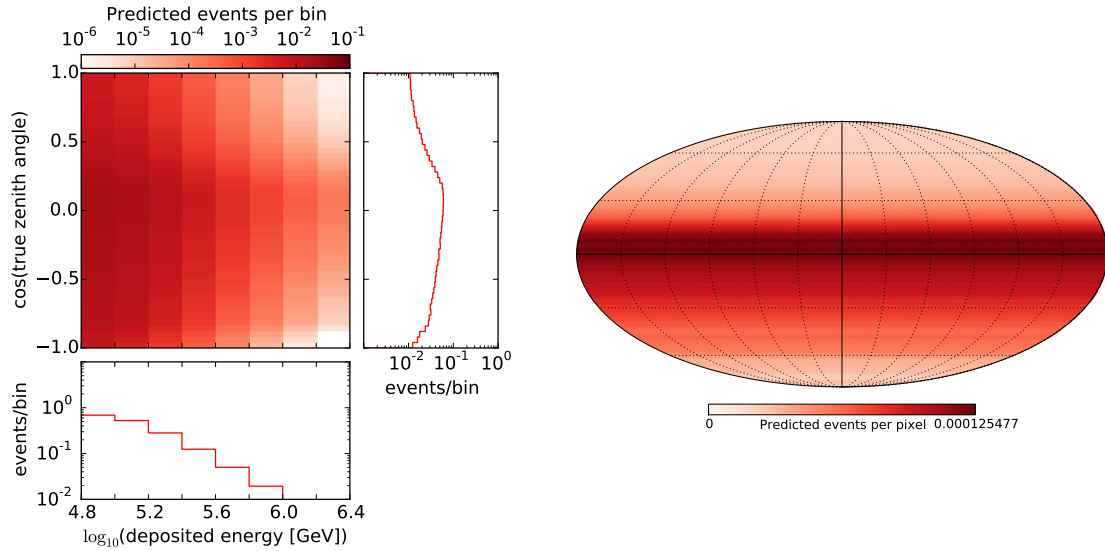


Figure 5.4: Prompt atmospheric neutrino cascade expectations. Left: The distribution of predicted cascade events in reconstructed deposited energy and true zenith angle, with corresponding 1D projections. Right: The cascade event distribution in detector zenith and azimuth using a HEALPix grid. Summing over all pixels on this grid for both tracks and cascades gives a combined normalisation of 1, such that the track to cascade ratio is preserved.

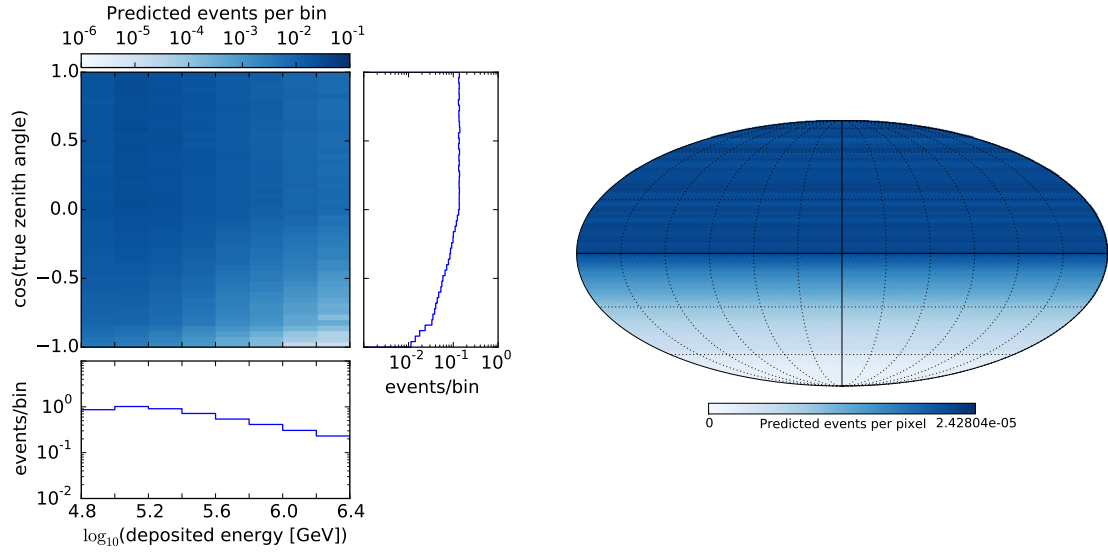


Figure 5.5: Diffuse astrophysical neutrino track expectations with a nominal E^{-2} spectrum. Left: The distribution of predicted track events in reconstructed deposited energy and true zenith angle, with corresponding 1D projections. Right: The track event distribution in detector zenith and azimuth using a HEALPix grid. Summing over all pixels on this grid for both tracks and cascades gives a combined normalisation of 1, such that the track to cascade ratio is preserved.

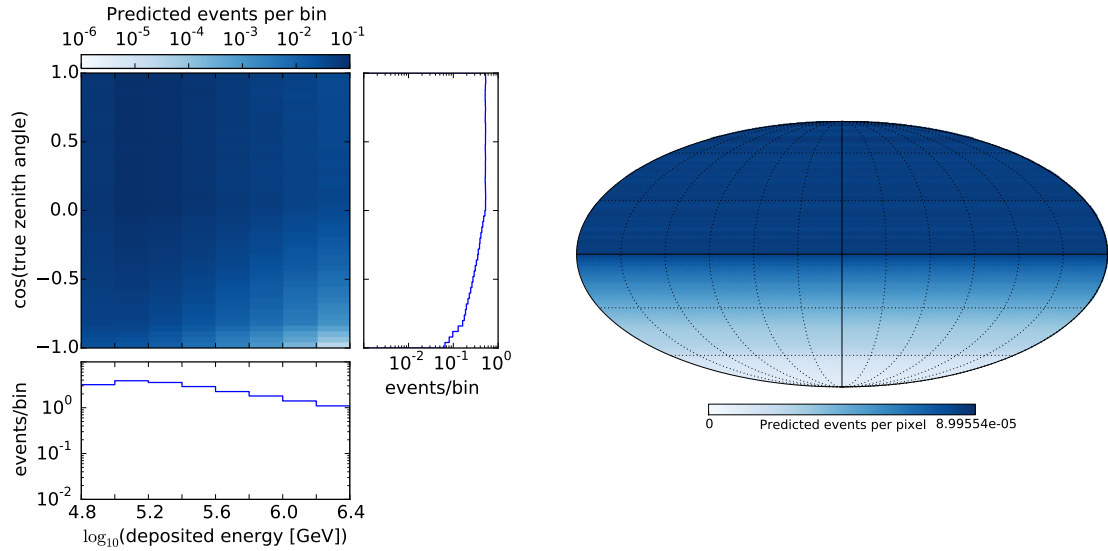


Figure 5.6: Diffuse astrophysical neutrino cascade expectations with a nominal E^{-2} spectrum. Left: The distribution of predicted cascade events in reconstructed deposited energy and true zenith angle, with corresponding 1D projections. Right: The cascade event distribution in detector zenith and azimuth using a HEALPix grid. Summing over all pixels on this grid for both tracks and cascades gives a combined normalisation of 1, such that the track to cascade ratio is preserved.

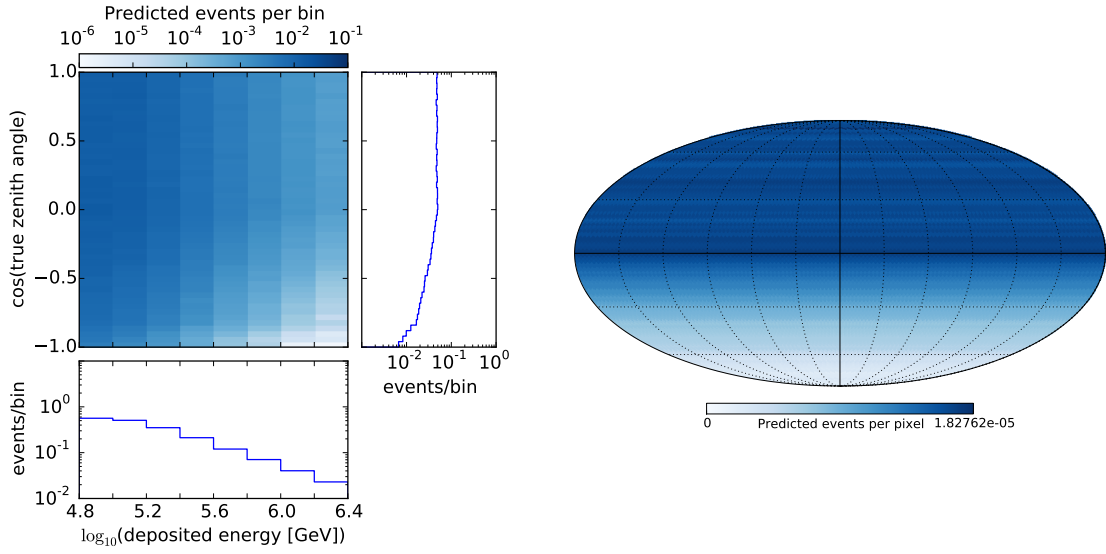


Figure 5.7: Diffuse astrophysical neutrino track expectations with a best-fit $E^{-2.6}$ spectrum. Left: The distribution of predicted track events in reconstructed deposited energy and true zenith angle, with corresponding 1D projections. Right: The track event distribution in detector zenith and azimuth using a HEALPix grid. Summing over all pixels on this grid for both tracks and cascades gives a combined normalisation of 1, such that the track to cascade ratio is preserved.

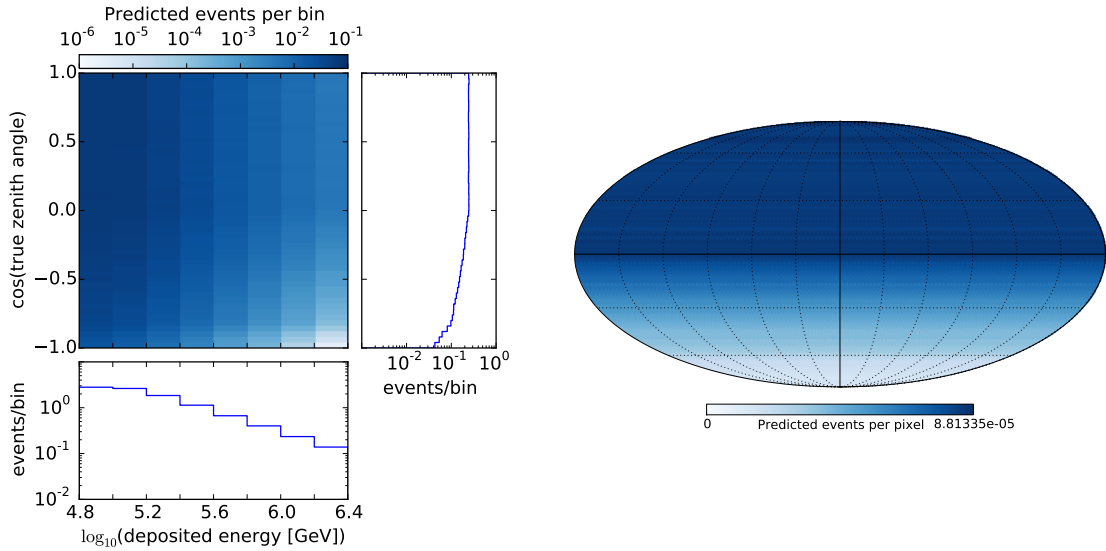


Figure 5.8: Diffuse astrophysical neutrino cascade expectations with a best-fit $E^{-2.6}$ spectrum. Left: The distribution of predicted cascade events in reconstructed deposited energy and true zenith angle, with corresponding 1D projections. Right: The cascade event distribution in detector zenith and azimuth using a HEALPix grid. Summing over all pixels on this grid for both tracks and cascades gives a combined normalisation of 1, such that the track to cascade ratio is preserved.

Chapter 6

Event selection, processing and generation

The events in the HESE analysis have already undergone a high level of computational analysis. Unlike many other analyses on IceCube data, the likelihood reconstruction techniques described in Chapter 3 were performed using a grid-scan, where the likelihood is maximised for each pixel over a many-pixel grid to determine the likelihood profile, giving a complete picture of the event direction and angular uncertainty. While computationally expensive, it was relatively straight forward to produce these high detail event point spread functions (PSFs) due to their small number. For this work, we use these detailed scans not just in the analysis itself but also when generating ensemble skies consisting of pseudo-events to test the sensitivity of the analysis.

The reason we reuse these detailed scans is that the number of ensemble skies in this work is very large, and therefore the ideal case of reconstructing each simulated event with such a scan is not practical. We have therefore chosen to reuse the PSFs of the event sample as representative of all event PSFs that could be seen in our ensemble skies. In this chapter we discuss the manipulation of the event PSFs needed to reuse them as pseudo-events, and the generation of ensemble skies.

6.1 Event selection

The fit to diffuse fluxes shown in Figures 4.3 and 4.4 use a low energy cutoff of $E_{cut} = 10^{4.8}$ GeV = 63.08 TeV, in addition to the 6000 PE total charge cut. This cut is made to remove the atmospheric muon background seen below this cut [89] and removes the uncertainty of the atmospheric simulation below this threshold. The precise selection of the 31 events remaining after the cut can be found in Tables 4.1 and 4.2.

6.2 Generation of randomly sampled skies

We wish to contrast the behaviour of the combined diffuse and point source fit against the fit using only a uniform background (see Equation 4.1). This comparison must be calculated for many possible observable skies and cannot simply be performed on the event sample we possess from the experimental data. For these calculations, we sample from our MC background distributions, which can be seen in Figures 5.1 - 5.8. The values we sample are the true event direction and reconstructed deposited energy of each event. The true event direction is sampled as we are unable to run the intensive directional reconstruction to determine the reconstructed event direction for each event, and so we use the point spread functions (PSFs) of existing events to represent angular uncertainties after sampling a true event zenith. However as each simulated event has a reconstructed deposited energy we can sample from the deposited event energy distribution and expect a representative event selection the same as if we had sampled from true neutrino energies. In addition, as we retain the original neutrino energy and angle we can also determine the effect of changing the spectrum on the distribution of reconstructed energies.

Our sensitivity studies rely on reusing the existing event point spread functions (PSFs), these PSFs are then used to generate representative maps of the data including signal events from the simulated signal and background distributions. Inherent

in this assumption is that the resolution of the HESE sample events is independent of the energy and angle of each event. Each such map contains the same number of events that appear in the experimental data sample, and the PSFs are sampled without replacement so that all events are used to produce each sampled map.

6.2.1 Scrambling method and sampling method

The scrambling method (see Section 5.2.7) used in this thesis is the same as used in the existing HESE 4-year point source search. Due to the near-continuous up-time, and the symmetry of the detector geometry in azimuth, we can perform the scrambling method by assigning a new random azimuth angle sampled uniformly from 0 to 2π radians to each event individually. The HEALPix implementation of interpolation is used to assign a new value to each pixel in the rotated event PSF given the corresponding location on the original event PSF with a differing right ascension. As one would expect, the zenith angle, energy and event topology of each event is unchanged under scrambling.

It is important to note that the final significance of the analysis was assessed only under scrambling. Sampled maps are used here only in the generation of example randomly sampled skies, which are used in our sensitivity and discovery potential calculations. For the tests of the analysis for a source in a fixed direction, these generated maps which include injected signal events are scrambled to generate a p-value. However, in the case of all-sky tests of the analysis where the sky is assessed for a “hotspot” using a grid search, (e.g. Figure 8.5), an ensemble of sampled maps which have contributions only from the background distributions are used to test significance as a proxy for scrambled maps, as sufficient statistics cannot be produced under scrambling due to computational limitations, such as in the case of all-sky sensitivity (Section 8.3). This same ensemble of background skies can be used for tests of the significance of each signal map by performing an identical likelihood test on each sky in the ensemble of background skies (for examples of this, see Figures 7.16 and 7.19). This makes the computation of significance relatively simple, as

only one ensemble of background skies is needed to assess the significance of all signal-injected skies, whereas to be assessed under scrambling, each signal-injected sky requires a unique ensemble of scrambled skies.

6.2.2 Resampling of event maps

Resampled maps use the true zenith and reconstructed deposited energy distributions from simulation to allocate random event directions and deposited energies on the sky. This is done to generate simulated skies that include a point source signal, and to assess the sensitivity using random background maps. For our final assessment of the Diffuse+PS likelihood, scrambled skies instead of sampled background skies were used for background estimation. This is because scrambling is a more robust statistical technique that is largely independent of any systematic uncertainties in our detector, and it was observed that the sensitivity of the analysis was not largely impacted by using scrambling relative to the simulated map ensembles.

We preserve the number of tracks and cascades from the data, and reuse the event PSFs from the data by rotating them to positions determined by random sampling of our possible event distributions (conventional neutrinos, prompt neutrinos, diffuse astrophysical neutrinos and neutrinos from a point source). While the possible true neutrino directions from our background distributions are potentially anywhere over the whole sky, the true directions of neutrinos from a point source are always the coordinates of the point source, however all events are later offset from their true directions (see Section 6.2.3) to account for the effects of reconstructing the events.

Simulated Event Direction and Energy

Simulated event direction and event energy are sampled from the two-dimensional zenith and energy distributions for conventional and prompt atmospheric neutrinos and astrophysical neutrinos, which return a true event zenith angle and reconstructed event energy within the detector. Due to the symmetry of the analysis in azimuthal angle for background distributions, the azimuthal angle for background

is sampled uniformly in the range of 0 to 2π radians. For a point source, we choose azimuthal angles corresponding to the point source right ascension. In addition, the energy of an event from a point source is taken from the astrophysical energy distribution evaluated at the corresponding zenith angle of the source. Having obtained a sampled energy and arrival direction, the PSF of the event is then calculated by a rotation of an existing, randomly-assigned event map to the direction of the sampled event direction. We term this sampled event direction $(\theta_{\text{MC}}, \phi_{\text{MC}})$, the term MC referring to the Monte Carlo method involving randomly sampling from a distribution. The sampled event directions from a point source have only a single specific direction - that of the point source.

6.2.3 Calculating rotation

We need to define a rotation that will place our event PSF with a correct offset from $(\theta_{\text{MC}}, \phi_{\text{MC}})$. The event PSFs define our uncertainty in the position of the event origins. As such, we do not expect that the PSF maximum or any other position will coincide with the MC event direction, but that there should be a distribution of events possible with PSF maximums distributed according to the width of the PSF itself.

Using the existing HESE PSF maps, we can sample $(\theta_{\text{offset}}, \phi_{\text{offset}})$ from those PSFs by using the event PSF as a probability distribution with the result that regions closer to the maximum are sampled as offsets more frequently than those further away. This is analogous to sampling a deviation from a Gaussian or other probability distribution. The sampling in this case is discretised due to the size of the pixels. Thus with track-like events, the sampled offset is often the same direction as the most probable direction of the event given the large peaked probability within the central pixel corresponding to the best fit direction.

The rotation we wish to apply to each pixel in the event PSF is the rotation that will rotate the PSF from this offset direction $(\theta_{\text{offset}}, \phi_{\text{offset}})$ and move the distribution such that the value of the PSF at $(\theta_{\text{offset}}, \phi_{\text{offset}})$ is now the value at $(\theta_{\text{MC}}, \phi_{\text{MC}})$ (see

Figure 6.1).

We define two orthogonal rotation matrices to perform this rotation: one rotation in zenith R_θ and one rotation in azimuth R_ϕ .

R_θ is about a rotation axis perpendicular to both $(\theta_{\text{offset}}, \phi_{\text{offset}})$ and $(\theta_{\text{MC}}, \phi_{\text{offset}})$ (i.e. rotating directly vertically from $(\theta_{\text{offset}}, \phi_{\text{offset}})$ by $\theta = \theta_{\text{MC}} - \theta_{\text{offset}}$, and R_ϕ is a rotation about the \hat{z} axis of $\phi = \phi_{\text{MC}} - \phi_{\text{offset}}$.

$$R_\phi = \begin{bmatrix} \cos(\phi) & -\sin(\phi) & 0 \\ \sin(\phi) & \cos(\phi) & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$R_\theta(\phi = \phi_{\text{offset}}) = \begin{bmatrix} \cos(\theta) + \sin^2(\phi)(1 - \cos\theta) & -\sin(\phi)\cos(\phi)(1 - \cos(\theta)) & \cos(\phi)\sin(\theta) \\ -\cos(\phi)\sin(\phi)(1 - \cos(\theta)) & \cos(\theta) + \cos^2(\phi)(1 - \cos(\theta)) & \sin(\phi)\sin(\theta) \\ -\cos(\phi)\sin(\theta) & \sin(\theta)\sin(\phi) & \cos(\theta) \end{bmatrix}$$

These rotations are then applied to any unit vector on the sphere as

$$\vec{u}' = R_\phi R_\theta \vec{u} \quad (6.1)$$

The application of these rotations to any PSF given any start and end points has the property of preserving the orientation of the PSF with respect to the North-South axis: a vertical feature will remain vertical, a horizontal feature will remain horizontal. However, larger PSFs will appear - but not actually be - somewhat rotated away from this North-South axis due to their extent. This is particularly evident for the example event in Figure 6.1 which spans over 30° . We can see for this example event that the larger lobe goes from having a peak location beneath the offset position $(\theta_{\text{offset}}, \phi_{\text{offset}})$ before rotation, to a peak location above the MC position $(\theta_{\text{MC}}, \phi_{\text{MC}})$. This is due to the PSF being defined on the sphere: for instance, if the final locations were $(\theta_{\text{MC}}, \phi_{\text{MC}})$ the southern pole at $(90^\circ, 0^\circ)$, then all other points of the PSF would by definition be above this location.

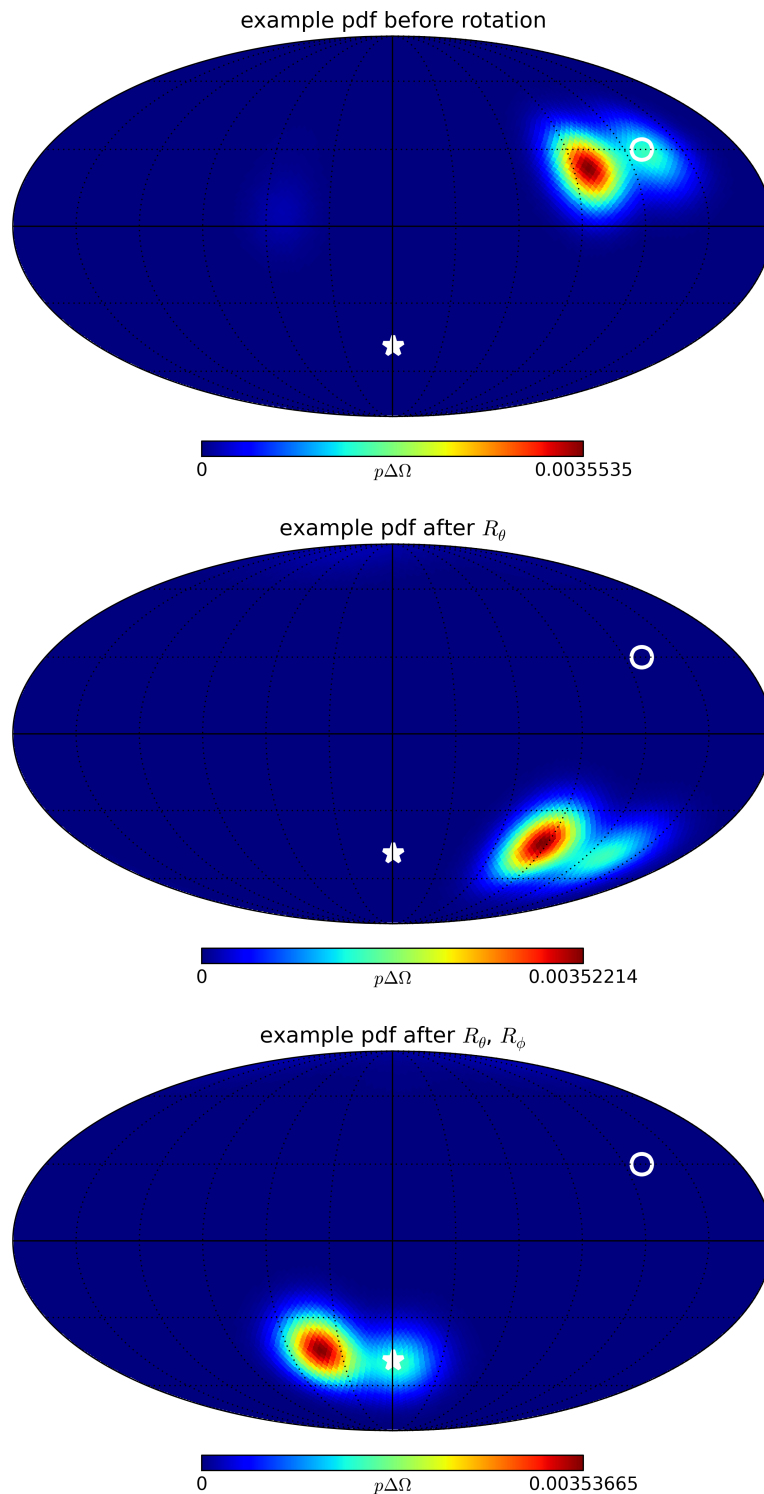


Figure 6.1: For an example PSF (in this case an artificial PSF with two lobes to demonstrate asymmetry), a view of the progress of rotation from start point (top) to end point (bottom) via rotations R_θ and R_ϕ . The circle shows the starting offset $(\theta_{\text{offset}}, \phi_{\text{offset}})$, the star the MC event position $(\theta_{\text{MC}}, \phi_{\text{MC}})$ for the PSF to be rotated to.

6.3 Multinomial sampling and weighting of event sets

So far, we have not yet defined how we are able to sample tracks and cascades from the differing backgrounds in this analysis. We achieve this through the process of multinomial sampling, which allows us to preserve the number of tracks and cascades. The skies generated from simulation contain events sampled with multinomial probability from our four possible sources: a sampled number of point source events (k_s), astrophysical background events (k_a), conventional atmospheric events (k_c) and prompt atmospheric events (k_p). Each of these are split by subscript “cas” for the number of cascades and “tr” for the number of tracks sampled for each source. A complete summary of these terms is shown in Table 6.1.

	point source	diffuse astro	conventional	prompt	sum
cascades	$k_{cas,s}$	$k_{cas,a}$	$k_{cas,c}$	$k_{cas,p}$	N_{cas}
tracks	$k_{tr,s}$	$k_{tr,a}$	$k_{tr,c}$	$k_{tr,p}$	N_{tr}

Table 6.1: Integer event counts for a sampled sky, these counts divide into eight values: four different source classes with two possible event types each.

The constraint we impose on the counts in Table 6.1 is that the total track count (N_{cas}) and cascade count (N_{tr}) are fixed to the observed rates of tracks and cascades in the HESE data set.

The sampled counts k are drawn from underlying rates λ . The values of λ for the astrophysical and atmospheric backgrounds come from the best fit flux values from the 4-year diffuse fit to the data in Chapter 4. We then partition the diffuse astrophysical flux such that the same fraction f of each of tracks and cascades are attributed instead to the point source. It then follows:

$$f = \frac{\lambda_{cas,s} + \lambda_{tr,s}}{\lambda_{cas,a} + \lambda_{tr,a}} = \frac{\lambda_{cas,s}}{\lambda_{cas,a}} = \frac{\lambda_{tr,s}}{\lambda_{tr,a}} \quad (6.2)$$

that is, the point source takes the same fraction of the astrophysical background

tracks and cascades. To remain physical, f must clearly be in the range $(0, 1)$. The relative rates can be expressed in terms of f as seen in Table 6.2, with $\sum_i \lambda_{\text{tr},i} = 1$ and $\sum_i \lambda_{\text{cas},i} = 1$ as these terms must independently govern the relative rates of tracks and cascades.

	point source	diffuse astro	conventional	prompt	sum
cascades	$f\lambda_{\text{cas},a}$	$(1-f)\lambda_{\text{cas},a}$	$\lambda_{\text{cas},c}$	$\lambda_{\text{cas},p}$	1
tracks	$f\lambda_{\text{tr},a}$	$(1-f)\lambda_{\text{tr},a}$	$\lambda_{\text{tr},c}$	$\lambda_{\text{tr},p}$	1

Table 6.2: Underlying rates for randomly sampled events that compose each example sky. The parameter f divides the astrophysical flux into a diffuse component and a point source component.

In general for a multinomial probability, we define relative rates λ_i over all possible source types i (as in whether the source type is conventional atmospheric neutrinos, prompt atmospheric neutrinos, astrophysical diffuse neutrinos or a point source of neutrinos), with $\sum_i \lambda_i = 1$. We define the observed integer counts for each source type i as k_i with $\sum_i k_i = N$ as the total count over all bins. The probability of observing k_i events given λ_i predicted events is then [103]

$$p(\vec{k}|\vec{\lambda}) = \frac{N! \sum_i \lambda_i^{k_i}}{\sum_i k_i!} \quad (6.3)$$

and this probability function is applied to our rates to independently sample track and cascade events, allocating them to the underlying flux distributions. With the number of events from each flux distribution given, each event is then sampled with an energy and zenith angle from the corresponding distribution. For example, if by multinomial sampling we obtain $k_{c,\text{tr}} = 2$, then 2 track events are drawn with deposited energy and zenith angle drawn from the energy and zenith distributions of conventional tracks.

Our ensemble maps with signal included are generated with specific values of the point source fraction $f = f_{\text{sample}} = 0.05, 0.1, 0.2, 0.3, 0.5$. However, the flux required to meet the conditions of sensitivity and discovery potential lies somewhere between these discrete values, so we need some means by which any ensemble of skies for

a specific fraction f_{test} can be tested. To obtain other fractions f_{test} between these values, we introduce to each map in the ensemble a weight. The weight is the ratio of probabilities of the new map being from the new value $f = f_{\text{test}}$ compared to the old value of $f = f_{\text{sample}}$, determined by the multinomial probability distribution function (Equation 6.3).

Specifically, the weight of a given map is given as the ratio of probabilities for the observed counts \vec{k}_{cas} and \vec{k}_{tr} , given the different rates governed by the sampled point source fraction f_{sample} to the new test point source fraction f_{test} . These new fractions result in new $\vec{\lambda}$ terms given as λ_{sample} and λ_{test} , and so our weights can be expressed as

$$w_{\text{cas}} = \frac{p(\vec{k}_{\text{cas}}|\vec{\lambda}_{\text{cas,test}})}{p(\vec{k}_{\text{cas}}|\vec{\lambda}_{\text{cas,sample}})} \quad (6.4)$$

and

$$w_{\text{all}} = \frac{p(\vec{k}_{\text{cas}}|\vec{\lambda}_{\text{cas,test}})p(\vec{k}_{\text{tr}}|\vec{\lambda}_{\text{tr,test}})}{p(\vec{k}_{\text{cas}}|\vec{\lambda}_{\text{cas,sample}})p(\vec{k}_{\text{tr}}|\vec{\lambda}_{\text{tr,sample}})} \quad (6.5)$$

where w_{all} is applicable for tests on the sky where all events are used in the likelihood, and w_{cas} is applicable for cascade-only tests of the sky where tracks are not considered in the likelihood. The probability p is given by Equation 6.3, with N in each case being the total number of tracks and cascades N_{tr} and N_{cas} for the corresponding expression of p .

As an example of how the weights are used, we might find that the condition for sensitivity (defined in Chapter 8) lies between an astrophysical fraction of $f_{\text{sample}} = 0.1$ and $f_{\text{sample}} = 0.2$. The value of the power (a term defined in Chapter 7) to meet the condition for sensitivity is $1 - \beta = 0.9$. The ensemble at $f_{\text{sample}} = 0.1$ results in a power of $1 - \beta = 0.86$ and the ensemble at $f_{\text{sample}} = 0.2$ results in a power of $1 - \beta = 0.95$.

In this case, we take the sampled maps at $f_{\text{sample}} = 0.1$ and re-weight each of them to values of f_{test} between 0.1 and 0.2. As we increase f_{test} , say to 0.12, maps which have a larger number of sampled point source events will receive a larger weight as they are more consistent with the larger point source fraction $f_{\text{test}} = 0.12$ than

with $f_{\text{sample}} = 0.1$, while events which have fewer point source events will be given a smaller weight as they are more consistent with $f_{\text{sample}} = 0.1$ than $f_{\text{test}} = 0.12$. The maps with more events from a point source are generally more likely to have a test statistic or p-value in the critical region (also defined in Chapter 8), and so their contribution to the power will increase with f_{test} , until the condition $1 - \beta = 0.9$ is met, in this case at $f_{\text{test}} = 0.135$.

Since the maps at $f_{\text{sample}} = 0.2$ are independent of those at 0.1, either can be re-weighted to the intermediate values. Therefore we can verify that the same value of $f_{\text{test}} = 0.135$ meets the condition of $1 - \beta = 0.9$ with this independent sample $f_{\text{sample}} = 0.2$. All results in this thesis showed good agreement in cases where such cross-verification was possible.

6.4 Example of a sampled sky

We show the test statistic scan on a sky using pseudo-events, which have been generated according to the procedure defined here, in Figure 6.2. The procedure for doing so is as follows:

- Background distribution rates are built from weighted neutrino simulation used in previous HESE diffuse flux tests.
- Multinomial sampling of these distributions, including a point source of relative flux f_{sample} , is used to generate a list of reconstructed event energies and true event directions, see Section 6.3.
- The counts of events from each source is stored so that a weighting to a new f_{test} can be calculated, also see Section 6.3.
- Each event is assigned an existing event PSF sampled without replacement, such that each new event corresponds to a unique PSF of the original data.
- Each of these event PSFs is rotated so that the offset direction sampled on the PSF now matches up with the sampled direction for each event, see Section

6.2.3.

- The Diffuse+PS likelihood (Chapter 5) is applied: at each pixel center on a HEALPix grid a point source is placed and the best fit is calculated using the new event PSFs.

These sampled skies are used to test the differences between the combined diffuse and point source fit and the uniform background fit (Equation 4.1), which are discussed in Chapters 7 and 8.

6.5 Summary

In this chapter we described our methodology to take input background and signal deposited energy and true zenith angle distributions and output ensembles of randomly sampled skymaps. This methodology requires the reuse of existing event PSFs, and allows for their reuse on any sampled event. This is an approximation to the ideal case in which each simulated event has its own corresponding reconstructed PSF. We also describe the method of multinomial sampling used to draw the event variables from the simulated distributions, and how this allows for a re-weighting of each sampled skymap to a different point source signal strength. An example of a generated sky tested with the Diffuse+PS likelihood is also shown.

In the next chapter we apply these techniques to perform statistical tests of the Combined Diffuse and Point Source (Diffuse+PS) analysis.

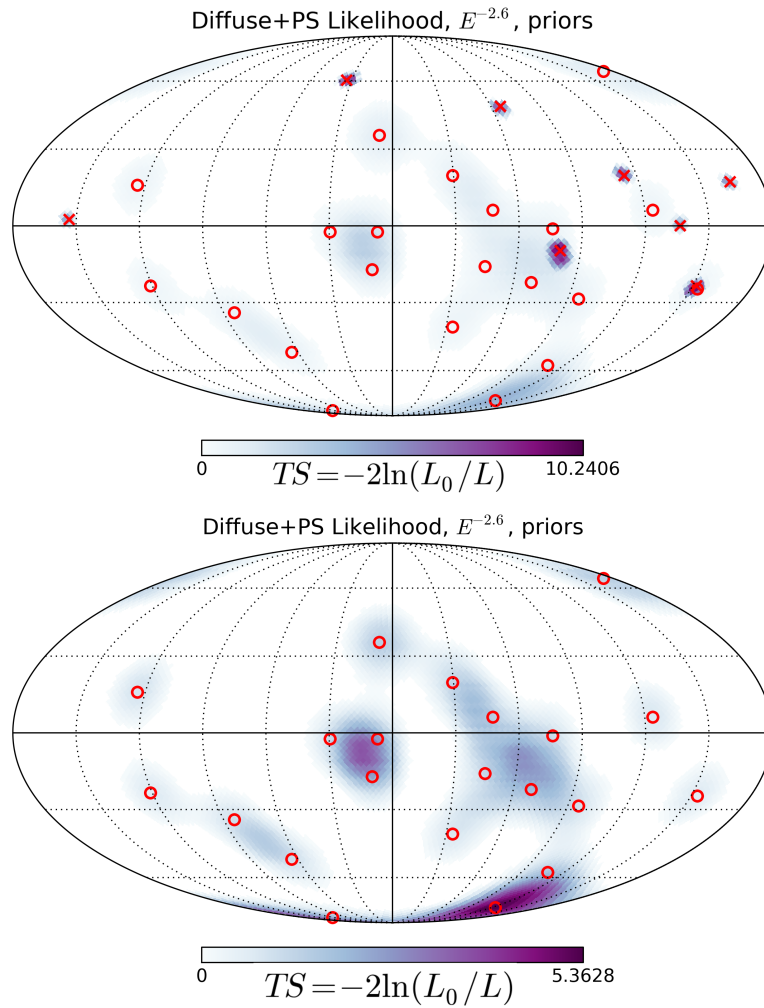


Figure 6.2: This example shows that we have the completed set of techniques to generate an example sky consisting of sampled events drawn from simulated background distributions, and to apply the Diffuse+PS likelihood to such sampled events. The combined diffuse and point source analysis is applied to this example sky which reuses the HESE sample event PSFs. Here, shower topology events are labelled by a red “o”, track topology events are labelled by a red “x”. Both plots show the result of the Diffuse+PS likelihood as shown by the test statistic TS in blue. Top: all events (track events and cascade events) are used in calculating the likelihood. Bottom: a cascade-only search where tracks are ignored in the likelihood.

Chapter 7

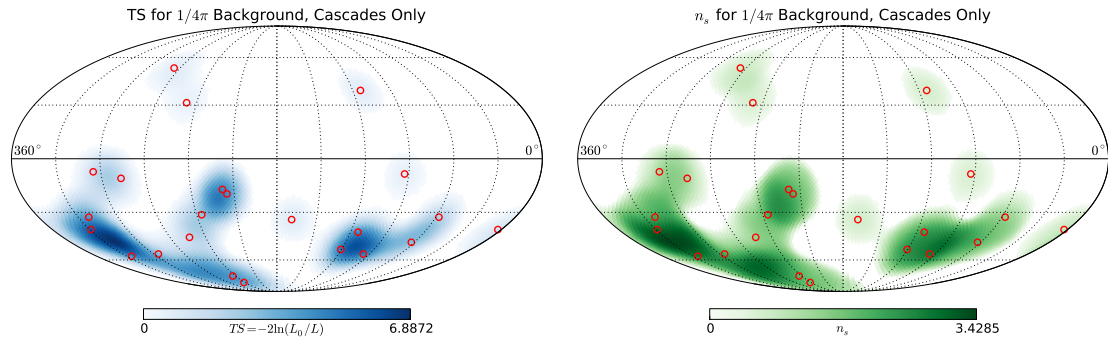
Statistical comparisons of the analyses

We wish to determine and assess the differences between the combined diffuse and point source fit, and the point source analysis with a uniform background. To do this, we test each analysis on generated HESE-like pseudo-experiments. These pseudo-experiments, like the data, have a reconstructed energy cut above $10^{4.8}$ GeV, and consist of simulated astrophysical and atmospheric neutrino events.

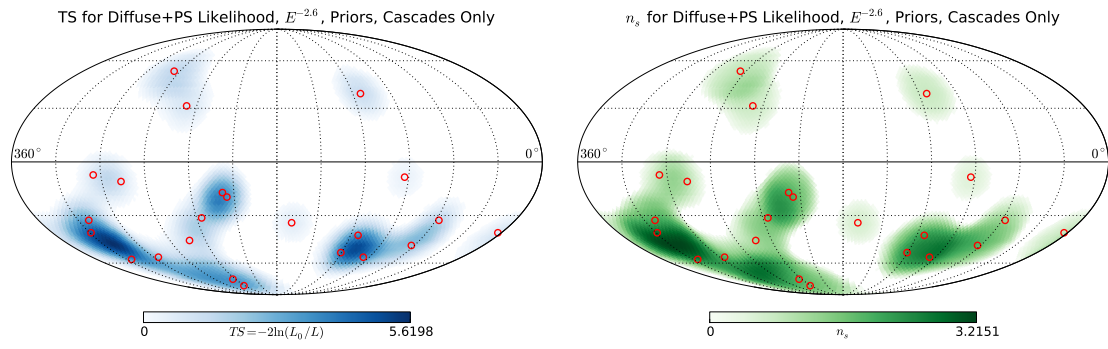
By adding events corresponding to a point source signal to the pseudo-experiments we are able to determine for each analysis the power to discern signal events above background. Differences in power between the two analyses emerge due to the choice of uniform background or the diffuse backgrounds in the likelihood used for the fit of a point source.

7.1 Direct comparisons between analyses

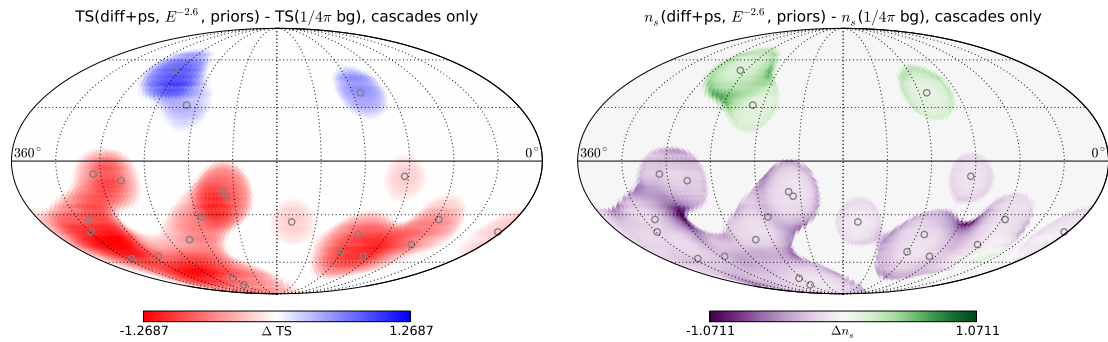
The uniform background likelihood (Eq. 4.1) treats events with equal weighting regardless of their energy and position in zenith, hence the best fit location of a point source will be the location that predicts the greatest number of high-resolution events. However, the astrophysical neutrino flux as observed by IceCube does not



(a) Likelihood map (left) and fitted signal strength (right) calculated with a uniform background likelihood (Equation 4.1).



(b) Likelihood map (left) and fitted signal strength (right) calculated with the Diffuse+PS likelihood (Equation 5.1).



(c) Difference between the Diffuse+PS likelihood (left) and fitted signal strength (right) compared to the uniform background case i.e. (b) - (a).

Figure 7.1: Comparisons between likelihood methods, in overall likelihood (left plots) and fitted signal strength (right plots). Here, shower topology events are labelled “o”, track topology events are not considered in this example. The differing emphasis of events at different declinations of the map due to the bias in fitted n_s is clearly visible in (c), with energy weighting having a relatively small impact. The most clear effect is that Earth absorption results in fewer events in the northern sky, and so this region is favoured with a larger value of n_s and TS for the best fit Diffuse+PS likelihood over the best fit uniform background likelihood.

produce a uniform distribution of events in the detector, it is a distribution over zenith angle that changes with energy and flavour. In considering the dependence on zenith angle alone, a point source could appear in a region of the sky where few background events are predicted: the strength of a point source in such a region would be substantially underestimated if the uniform background predicts a larger rate of background events in this region. In addition to a model of the zenith distribution, using the Diffuse+PS likelihood (Eq. 5.1) adds further information to determine the point source likelihood by comparing the energy and flavour rate from each component. These differences provide a small amount of additional power by discriminating astrophysical events from the small number of remaining atmospheric background events in our sample.

We can compare the results of likelihood scans to find the effect of the differing backgrounds on a given map. The events in Figure 7.1 were generated from a diffuse atmospheric and astrophysical background assumption with an $E^{-2.6}$ astrophysical spectrum, and we compare the results of the uniform background likelihood (Equation 4.1) to that of the Diffuse+PS likelihood (Equation 5.1).

We can see that the choice of likelihood affects the test statistic (TS) in Figure 7.1c, where for both tests a relative increase in the fit to signal strength between the tests appears to correlate with an increase to TS. In both analyses the fit to signal strength n_s compensates for the predicted background rate in the region: we know that in the northern sky, up-going events to the detector are diminished by Earth absorption (as can be seen in Figures 5.5 to 5.8) which the uniform background does not take into account. As a consequence the background rate in the northern sky is lower for the Diffuse+PS analysis compared to a uniform background, and this causes the fit for point source strength in this region to increase relative to the uniform background, resulting in a relative increase in the value of TS.

In contrast, the southern sky, which we see in IceCube as down-going events, is uniformly exposed to the astrophysical flux as it is almost unaffected by Earth absorption. Consequently the diffuse background fit correctly predicts more down-

going events than the uniform background does, and as a result the best fit n_s value will be smaller for the Diffuse+PS fit than the uniform background, with a relative decrease in TS between the two analyses.


This effect is shown for a one dimensional binned poisson fit in Figure 7.2. The same behaviour can be observed for a representative background (left) and a uniform background (right), that the more significant test statistic emerges when the corresponding background fits to fewer events at the source location.

Determining the effectiveness of the analysis however is not done on a map-to-map basis but by looking at the differences over a large ensemble of experiments, to do this requires statistical methods such as finding the power of a test.


7.2 Definitions of power and significance

The key goal of our analysis is to distinguish between two hypotheses: the null hypothesis that a given sky has no point source contribution, or alternatively that the sky does have a point source contribution. This separation can be accomplished through the use of hypothesis testing to discriminate between these hypotheses given the data available. This test is done using the test statistic from a likelihood ratio, where a larger test statistic will arise from our alternative hypothesis as a point source contribution begins to describe our data better than no point source contribution, as we increase the point source strength. Thus we will define a critical region in the tail of large values of the test statistic that allows us to reject the null hypothesis if the alternative hypothesis is true.

In simple terms, we wish to test an hypothesis H_0 against an alternative hypothesis H_1 . In testing for a point source, we generally define H_0 as the hypothesis that the point source signal strength is zero, that is that there is no contribution from a point source to the observed events, whereas H_1 is the alternative hypothesis, which in our case is the point source signal strength is the best fit \hat{n}_s from the likelihood fit. We have previously defined our test statistic TS in Equation 5.26 as the like-



(a) One dimensional examples of a two component fit - a background and signal distribution, for a correct background (left) and uniform background (right). The fit here is performed with signal injected in a region where the uniform background over-predicts the background rate, giving the correct background fit a larger signal fit and test statistic.



(b) One dimensional examples of a two component fit - a background and signal distribution, for a correct background (left) and uniform background (right). The fit here is performed with signal injected in a region where the uniform background under-predicts the background rate, giving the correct background fit a smaller signal fit and test statistic.

Figure 7.2: One dimensional simplification showing the impact of the background zenith distribution in a Poisson likelihood fit. If considering the test statistic and signal fit, the correct background does not always find a better result, but rather whichever background has a relatively lower prediction at the source location.




Figure 7.3: A graphic showing how significance α and power $1 - \beta$ are defined in our analysis; here α is the area in the light grey region, while $1 - \beta$ is the entire shaded area. In this example, as with our analysis, $TS < 0$ does not occur. In this instance, the region of acceptance is the region $0 \leq TS \leq 2$ and the critical region corresponds to $TS > 2$. In our specific case, the null hypothesis corresponds to the TS distribution seen with no injected signal, and here the alternative hypothesis includes some measure of signal from a point source.

likelihood comparison between the best fit signal \hat{n}_s and the null hypothesis $n_s = 0$, with $TS \geq 0$ by definition. We define our *region of acceptance* $0 \leq TS \leq TS_{\text{crit}}$ and the critical region $TS > TS_{\text{crit}}$. Values in the critical region suggest that H_0 is not true [103].

We can adjust the size of the critical region by changing TS_{crit} to obtain any possible *level of significance* α , the probability of observing a test statistic in the critical region if the null hypothesis is true:

$$P(TS > TS_{\text{crit}}|H_0) = \alpha \quad (7.1)$$

For instance, with a significance level of $\alpha = 0.01$, we have a probability of 0.01 of rejecting the null hypothesis even if the null hypothesis is true.

A test statistic TS is powerful if we can use it to discriminate between two hypotheses H_0 and H_1 . The *power* of a test is defined as the probability of TS belonging to the critical region if H_1 is true, that is the probability of correctly

rejecting the null hypothesis H_0 if H_1 is true:

$$P(\text{TS} > \text{TS}_{\text{crit}} | H_1) = 1 - \beta \quad (7.2)$$

This is shown graphically for a signal and background population in Figure 7.3. In addition these values can be interpreted as type-I and type-II errors: the probability α of observing a test statistic that rejects H_0 when H_0 is true is a type-I error, and the probability β of observing a test statistic that does not reject H_0 when it is not true is a type-II error.

There are two key types of tests that are performed in point source analyses: tests seeking power for a source in a fixed direction, and tests to maximise the all-sky power, in which the “hotspot”, the largest value of the test statistic on the sky, is evaluated.

7.2.1 Power for a fixed direction

To assess the power for a fixed direction \vec{x}_s , signal is injected at a given location in the sky \vec{x}_s as in Chapter 6. The value of TS is assessed specifically at the same location \vec{x}_s where signal has been injected. The value of TS is then assessed exactly the same way at the same location \vec{x}_s on all corresponding background maps to determine the probability of the test statistic at that point occurring by chance due to random fluctuations in the background.

This can be thought of as the sensitivity to a source when we want to assess a specific pre-trial location on the sky. This is the case when performing a predefined catalogue search (for an example see [104]), and is most often used as the point of comparison between different point source analyses. However, this method does not accurately represent the sensitivity to an all-sky search for a “hotspot” of neutrinos, which for our analysis is given by the all-sky power.

In Section 7.4 we show that the power for a source in a fixed direction is affected in an unusual way from the change in background distributions. The different back-

ground distributions which do not represent the data change when the likelihood fits to a singularity at zero, and therefore the change in background results in a change the power in this regime. Outside of the singularity at zero, the power appears unaffected.

7.2.2 All-sky power

All-sky power is defined in this analysis as a measure of how significant the location with the greatest TS value is when compared to randomised background skies. As when assessing power for a fixed direction, we inject signal events at a given location \vec{x}_s . However, for all-sky sensitivity our statistic is the largest value of TS evaluated over the entire HEALPix grid which covers the sky. This is repeated in the same way on background-only maps: the maximum value of TS is reported regardless of location. There is no constraint on which position of TS is considered on the signal or background maps.

This definition of power, not the power for a fixed direction, is the actual power of the analysis as it is performed, and tells us the flux required to see a point source in our all-sky test of a “hotspot” at a given significance.

7.2.3 Scrambling and sampled maps

All values of sensitivity and discovery in this thesis are generated numerically using randomly generated maps with no approximation or extrapolation of the test statistic distribution, as is sometimes done in point source analyses due to the large amount of computational time required to generate many background skies. The primary reason such approximations are not used in this analysis is that both the fixed-point and all-sky sensitivity of these analyses are not likely to follow a simple functional form. Power for a fixed direction is often approximated by a χ^2 distribution, however we expect deviations away from a χ^2 distribution due to the differing background in the analyses. When we had examined the all-sky test statistic dis-

tribution early on in the course of this work, we observed that bumps in the test statistic distribution that appeared to be due to the combination of tracks and cascades at different resolutions, and so a chi-squared distribution would not be likely to pick up this effect. We speculate that this change in shape may have been lessened with a larger data sample, as this effect is less apparent in Figure 7.19.

We have also found substantial deviations of the scrambled fixed-direction TS distribution between individual maps, and so this behaviour would be difficult to infer with any given approximation. These variations between event maps appear to be primarily due to the small event count: within a given declination band, only a few events will be responsible for the distribution of TS, therefore adding or removing a single event can noticeably affect the distribution.

We wish to assess power using sampled background maps in some circumstances, such as when computing the all-sky sensitivity and discovery potential, when doing so under scrambling would be many orders of magnitude more computationally intensive. To assess power using sampled background maps we directly compare the test statistic distribution of maps with signal included to the test statistic distribution of the background-only maps. We then determine TS_{crit} based on our predefined value of α , and assess the power, that is the probability of correctly rejecting the null hypothesis, by looking at the fraction of the signal distribution above this critical value.

We redefine power when assessing the power using scrambled maps due to the fact that the distribution of TS found by scrambling one experimental map is in no way guaranteed to be the same distribution as that of another experimental map. This is because under scrambling each experimental map defines its own background hypothesis H_0 : that there is no inherent clustering in right ascension but the zenith distribution is given by that of the events. This naturally affects both fixed direction and all-sky power.

The previous definition of power (Equation 7.2) fails to account for the fact that each map has a unique TS distribution under scrambling. Instead we must first

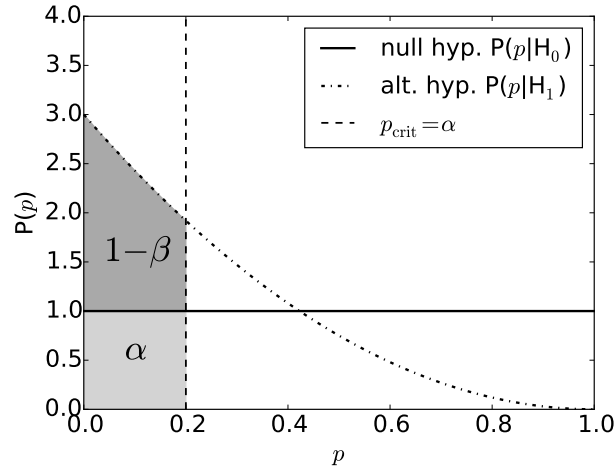


Figure 7.4: A graphic showing how significance α and power $1 - \beta$ are defined when using the p-value p ; here α is the area in the light grey region, while $1 - \beta$ is the entire shaded area. In our specific case, the null hypothesis corresponds to the p-value distribution seen with no injected signal, and is therefore a flat distribution. Here the alternative hypothesis includes some measure of signal from a point source, and is therefore more peaked towards p-values close to zero.

assess each map's p-value under scrambling $p(\text{TS}_{\text{obs}})$, so that different maps have a comparable measure. We define the p-value as the probability of observing a test statistic TS greater or equal to the observed statistic TS_{obs} if the null hypothesis - that there is no signal as given by the scrambled background maps - is true.

$$p(\text{TS}_{\text{obs}}) = P(\text{TS} \geq \text{TS}_{\text{obs}} | H_0) \quad (7.3)$$

In other words, the p-value is the fraction of scrambled skies in which the relevant test statistic is greater than that of the experimental map TS_{obs} . We wish to transform our definitions for the level of significance α and power $1 - \beta$ from TS-space to p-space. To define the level of significance α we want a definition of the critical p-value p_{crit} . To define this, first we suppose that all maps *do* have the same background test statistic distribution. Then by substituting TS_{crit} into Equation 7.3, we obtain

$$p_{\text{crit}} = p(\text{TS}_{\text{crit}}) = P(\text{TS} \geq \text{TS}_{\text{crit}} | H_0) \quad (7.4)$$

which is identical to the definition of the level of significance α in Equation 7.1.

Hence,

$$p_{\text{crit}} = \alpha \quad (7.5)$$

and thus α itself can be used in lieu of TS_{crit} to define the critical region in p-value space. From the definition of p in Eq. 7.3, it should be apparent that as the observed test statistic TS_{obs} increases, the corresponding value of p decreases, and so to reject skies where it is more likely that a point source is present, we define the critical region as $p < \alpha$. Therefore the equivalent definition of the power previously defined as Equation 7.2 instead becomes

$$1 - \beta = P(p \leq \alpha | H_1) \quad (7.6)$$

that is, the power of the analysis is the probability of observing a p-value of value less than the level of significance α given the alternative hypothesis H_1 . As the alternative hypothesis H_1 becomes increasingly distinct from the null hypothesis H_0 , we expect to observe a larger proportion of small p-values with a value less than the level of significance α . This is shown graphically in Figure 7.4.

7.3 Investigations into bias

Investigations into a possible bias in either analysis were motivated by early calculations which showed that an assumption of uniform background could result in greater power in a point source fit than the Diffuse+PS analysis, counter to our expectations at the time. Initial tests were performed to investigate this behaviour, these tests were designed to investigate the potential discrepancies that could arise due to:

- Event number - if there are too few events, perhaps there is not enough information to constrain the backgrounds in the fit.
- Energy - we observed a discrepancy in the predicted event rates compared to the data in the 10 – 60 TeV range. This precipitated the decision to use the




Figure 7.5: Backgrounds used in the likelihood to investigate bias. “Correct” and “Flat” backgrounds (left) show the difference of event rates in zenith angle of an E^{-2} diffuse flux and a uniform background. The “Inverted” background is added (right) as an extreme test of the background differing to the pseudo-experiment data; this background is the zenith distribution of the E^{-2} diffuse flux if the declination is inverted ($\delta \rightarrow -\delta$).

63.08 TeV cut from the diffuse analysis. We also wanted to know what effect, if any, using energy has on the analysis.

- Additional fit parameters n_c and n_p for the atmospheric conventional and prompt fluxes - by having multiple backgrounds freely floating in the fit, there were two additional ways to explain point source events as background that could subtract from the point source significance.

To investigate each of these points, our test was then simplified to pseudo-experiment maps of 100 cascade events drawn from a single zenith distribution, in this case the zenith distribution of the E^{-2} astrophysical flux for this analysis, ignoring any contribution from the atmospheric fluxes for the purposes of these investigations. The energies of the events were also not used as a variable in these tests. We tested the statistical power of signal determination under varying background models in the likelihood including:

- “Flat” (the uniform $1/4\pi$ background),
- “Correct” (the same astrophysical E^{-2} distribution used to generate pseudo-experiment maps),

- and “Inverted” background (the same as the “Correct” E^{-2} background but flipped in declination).

These zenith distributions are shown in Figure 7.5.

To save computational time, the sampling method (see 6.2.1), that is comparing each result to an ensemble of sampled maps drawn from the background distribution, as opposed to scrambling each map to evaluate the significance, was used to determine the power of each test. In addition, rather than Poisson sampling from a rate of signal events, injected signal events are discretised for each sky: a specific number of integer events n_i are injected at the point source location. This allows for the testing of non-zero source event counts as null hypotheses as shown in Section 7.3.2.

7.3.1 Tests of power for a fixed direction

Initial tests of power were performed at significance levels of $\alpha = 0.1$ and $\alpha = 0.01$ under the background-only hypothesis of each case. An additional test at $\alpha = 0.5$ has been included to illustrate an effect of the test statistic distribution discussed later in Section 7.4. For the statistical power rejecting the background only hypothesis, $n_s = 0$, for a fixed direction (Figure 7.6), our findings were that for significances of $\alpha = 0.1$ and $\alpha = 0.01$ all tests have nearly equal power. An exception to this, where the power is dependent on the choice of background, can be seen at $\alpha = 0.1$ for the case of the “Inverted” background for a source at $\delta = +45^\circ$, and also at $\alpha = 0.5$ for all backgrounds, but most prominently at $\delta = +45^\circ$. We can see that at $+45^\circ$ the correct background has more power to reject the null hypothesis when α is large, but there are also hints that the flat and inverted backgrounds result in marginally more power at $\delta = -45^\circ$ than the correct background in the likelihood at large α .

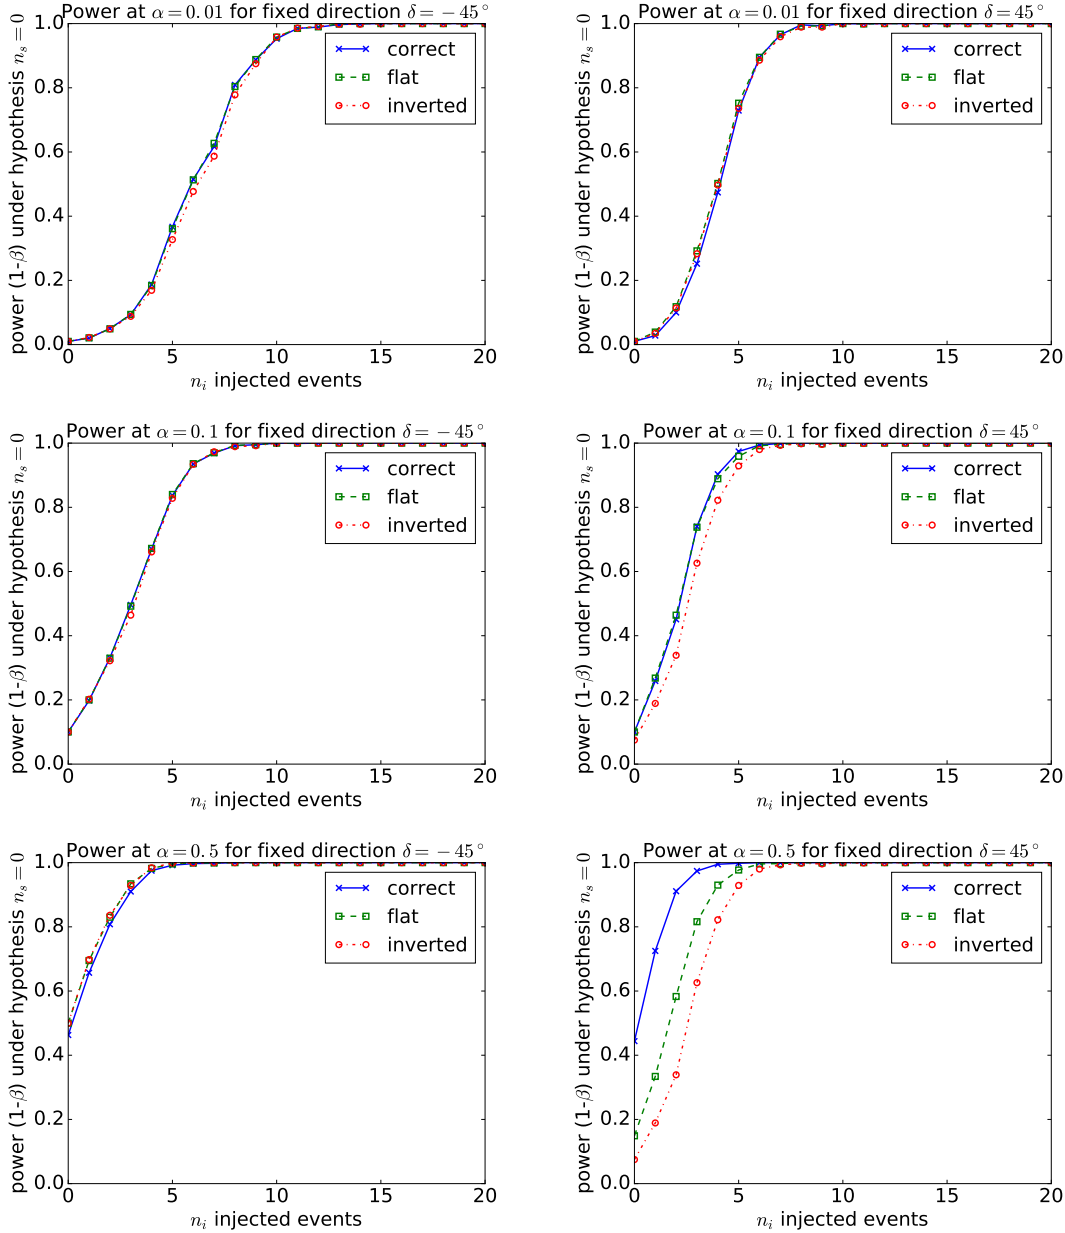


Figure 7.6: Testing power under the null hypothesis $n_s = 0$ for a fixed direction against injected signal with varying background hypotheses (coloured lines) at declinations of $\delta = -45^\circ$ (left) and $\delta = +45^\circ$ (right), with different critical regions $\alpha = 0.01$ (top), $\alpha = 0.1$ (centre), $\alpha = 0.5$ (bottom). All plots show that the power to reject the background-only hypothesis increases as signal events are injected, and most plots show that this power is near equal for each background tested in the likelihood. However, at large values of α , for example $\alpha = 0.5$ and $\delta = +45^\circ$ (bottom right), we observe that the choice of background does impact the power, in this case favouring the correct background model in the likelihood.

7.3.2 Fixed direction tests of non-zero signal hypotheses

Testing the best fit signal hypothesis $\mathcal{L}(\hat{n}_s)$ against the null hypothesis of background only $\mathcal{L}(n_s = 0)$ is only one possible hypothesis test of point source signal strength (see Equation 5.26). In the regime where the signal strength is substantial enough, we may wish to place a confidence interval on the possible values of n_s . In such a case, we would like to test non-zero hypotheses of signal strength against the best fit signal, where n_s is the hypothesis point source signal strength of fixed value, giving a test statistic

$$\text{TS} = -2 \ln \frac{\mathcal{L}(n_s)}{\mathcal{L}(\hat{n}_s)} \quad (7.7)$$

We can then assess the test statistic over the range of possible n_s hypotheses and determine if any method shows a bias in the interval of allowed n_s . As stated earlier the power $1 - \beta$ is governed by the excess of TS for the null hypothesis over the alternative hypothesis (see Figure 7.3). In this case, the null hypothesis is $n_s = n_i$, the actual number of injected events, so we assess the ensemble of test statistics TS for injected strength n_i against various signal strength hypotheses n_s , and look for the excess relative to the TS distribution for the null hypothesis $n_s = n_i$. A natural consequence of this is that we will always obtain $1 - \beta = \alpha$ for this test when $n_i = n_s$.

For clarity, let us start with the case we are already familiar with, that of a null hypothesis $n_s = 0$. Suppose we wish to test for this with an injected signal $n_i = 3$. We require for this two ensembles, that of our alternative hypothesis with $n_i = 3$, and that of our null hypothesis $n_i = n_s = 0$ (a background only ensemble). We look at the distribution of $\text{TS} = -2 \ln \mathcal{L}(n_s = 0) / \mathcal{L}(\hat{n}_s)$ for these two ensembles as in Figure 7.3 in order to determine the power $1 - \beta$ as the sum of all best-fit test statistics TS that are greater than the critical value TS_{crit} . The critical value is determined by the predefined value of α , since this determines the size of the tail in the null hypothesis TS distribution.

For tests of different non-zero null hypotheses, the same principle can apply. As

an example, we can select a null hypothesis $n_s = 5$. Suppose we wish to test for this with injected signal $n_i = 3$. We take two ensembles, one with $n_i = n_s = 5$ as the null test distribution, and one with $n_i = 3$ as the alternative test distribution. Testing with the null hypothesis $n_s = 5$:

$$\text{TS} = -2 \ln \mathcal{L}(n_s = 5) / \mathcal{L}(\hat{n}_s) \quad (7.8)$$

and applying this test on an ensemble with injected signal $n_i = 5$ we obtain the null test statistic distribution, and applying this same statistic to the alternative distribution for injected signal $n_i = 3$ we obtain the alternative test statistic distribution. Examples of this case are shown in Figure 7.7. We then determine the power of this test by looking at the excess of test statistics of the $n_i = 3$ distribution above the critical value determined by the $n_i = 5$ distribution.

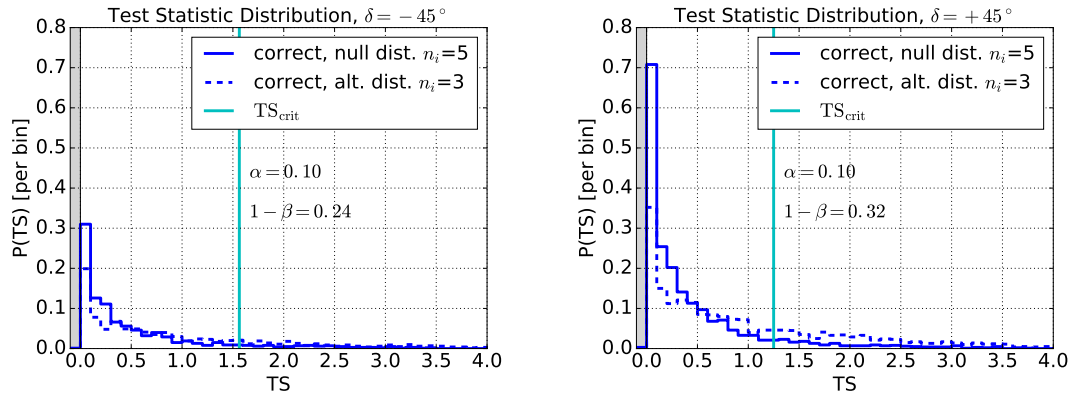
For any test we might perform, if there are any regions which satisfy the condition

$$1 - \beta < \alpha \quad (7.9)$$

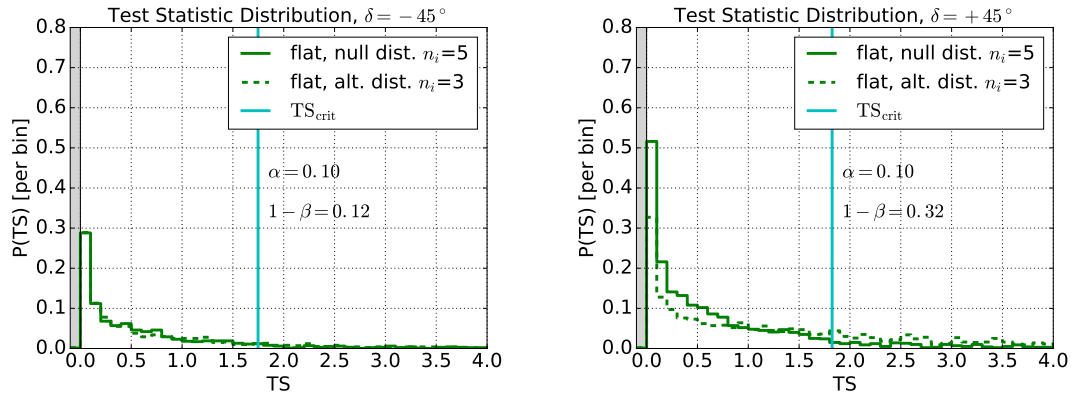
that is, if the statistical power of a test falls below the level of significance, then this specific test is called a *biased test*. This means that for such a region we are more likely to accept the null hypothesis when it is false than when it is true. This is an undesirable property for a test to have in this region, though the test may potentially be more powerful in other regions [103].

This bias manifests as a result of the fit \hat{n}_s being a biased estimator when the true background distribution is not used in the likelihood when testing for a point source. We can observe this bias as the fitted value of \hat{n}_s is systematically offset from the injected signal strength n_i . The fit of \hat{n}_s compensates for the estimated background excess or deficit relative to the true background rate. This property of n_s is shown in Figure 7.8.

Figure 7.9 and the corresponding vertical slices through these plots in Figure 7.10 show the effect of bias when testing signal strength for a fixed direction against



(a) Test statistic distributions testing the null hypothesis $n_s = 5$ for a fixed direction, using the “Correct” background in the likelihood, for a point source at $\delta = -45^\circ$ (left) and $\delta = +45^\circ$ (right).



(b) Test statistic distributions testing the null hypothesis $n_s = 5$ for a fixed direction, using the “Flat” background in the likelihood, for a point source at $\delta = -45^\circ$ (left) and $\delta = +45^\circ$ (right).

Figure 7.7: The distributions of the test statistic TS for the non-zero null hypothesis $n_s = 5$ as described in Section 7.3.2.

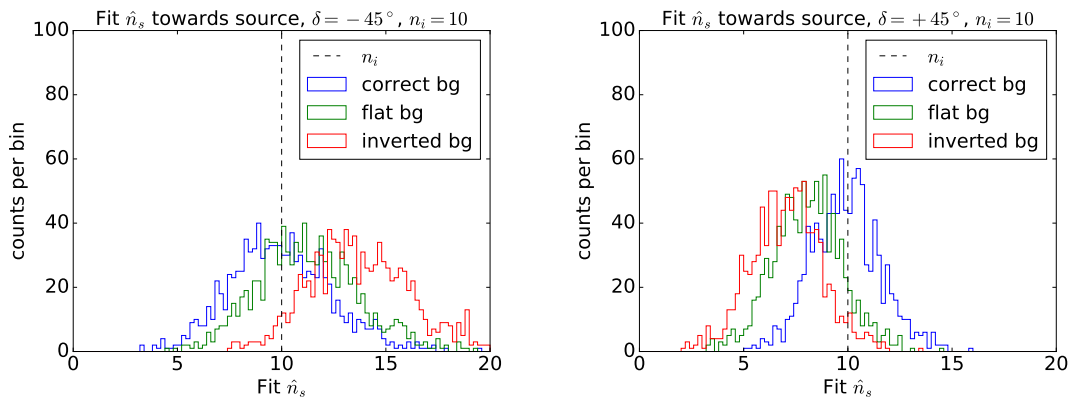




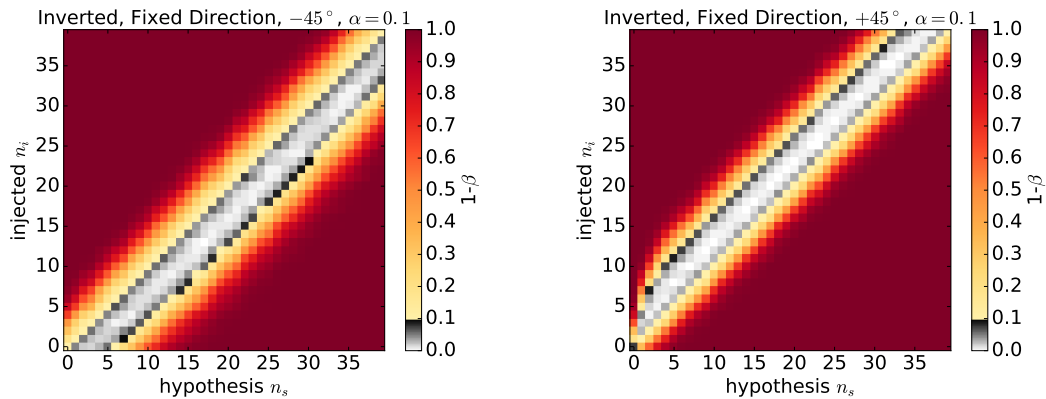
Figure 7.8: Plot showing the fit to n_s for a fixed injection of $n_i = 10$ at $\delta = -45^\circ$ (left) and $\delta = +45^\circ$ (right). It can be seen that the fit to n_s is centred around differing values dependent on the background model fit in the likelihood.



(a) The “Correct” background is tested in the likelihood for a point source at -45° (left) and at $+45^\circ$ (right) declination.



(b) As above for the “Uniform” background



(c) As above for the “Inverted” background.

Figure 7.9: 2D plots of the observed power $1 - \beta$ (colour scale) for fixed significance level $\alpha = 0.1$ testing different signal injection hypotheses n_s and given injected signal strengths n_i . Tests are performed using the background models in the likelihood as indicated: “Correct” (top), “Flat” (center) and “Inverted” (bottom). Power shown is for a fixed direction against sampled background maps. Regions in greyscale are biased regions where the power is less than the level of significance. The “Correct” background shows no bias, whereas the “Flat” and “Inverted” backgrounds appear to show biased regions due to under- or over-estimating the background rate at the point source, which can be observed as a change from the bias regions for a source at $\delta = -45^\circ$ appearing at $n_s > n_i$, switching to biased regions for $\delta = +45^\circ$ appearing at $n_s < n_i$.

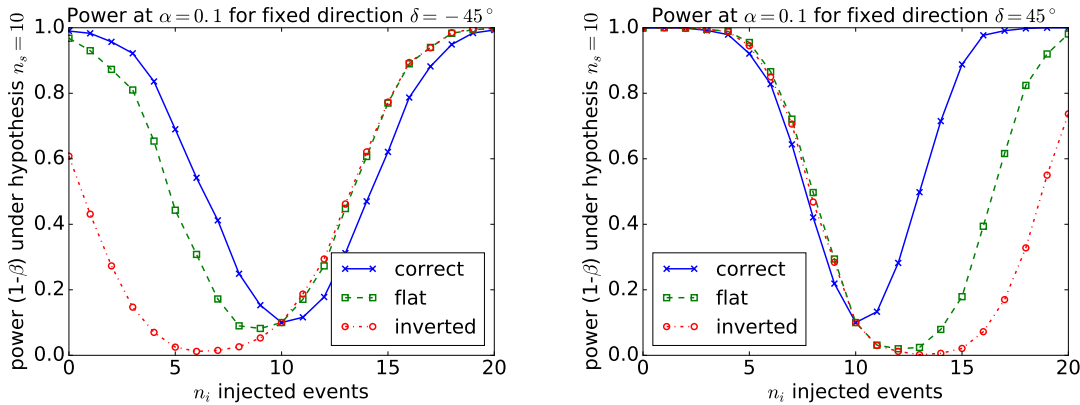


Figure 7.10: Vertical slices of Figure 7.9 which show the power $1 - \beta$ to reject a fixed hypothesis $n_s = 10$ for $\alpha = 0.1$ for a range of injected signal n_i . This is shown for an injected source at declinations of -45° (left) and $+45^\circ$ (right). Biased regions are where the power drops below the level of significance $\alpha = 0.1$. Where $n_s = 10$, the power $1 - \beta = \alpha = 0.1$ as both null and alternative distributions of TS are identical.

a background that does not accurately represent the data. We observe that all but the “Correct” background exhibit bias for non-zero n_s , and this biased region is reversed for the two different directions on the sky, echoing the dependency of fitted signal on the choice of background in the fit (Figure 7.1), specifically the balancing of n_s to compensate for any excess or deficit in the background rate.

7.3.3 Tests of all-sky power

The all-sky power involves a test of the point on the sky where the maximum value of TS is found. The all-sky power when testing the background-only hypothesis $n_s = 0$ was assessed, and the power as a function of injected signal can be seen in Figure 7.11. These results show that which choice of background used in the likelihood is more powerful between the E^{-2} flux and the uniform background depends on where a point source is injected.

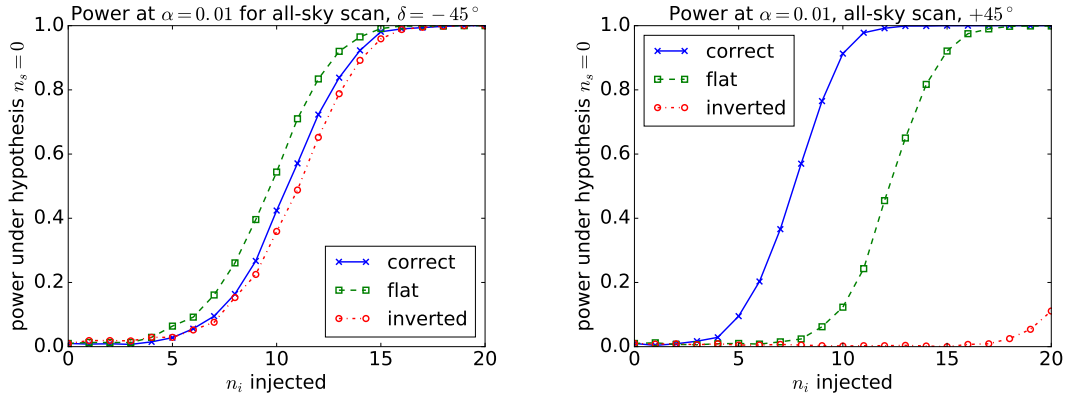
The power appears to be greater for the backgrounds which predict fewer events at the source location, and therefore the point source over-fits, with the “Correct” background predicting fewer events in the northern hemisphere ($\delta = +45^\circ$), and the “Flat” background predicting the fewer events in the southern hemisphere ($\delta = -45^\circ$). This is consistent with the relative fit to TS in the northern and southern

sky, in general where the test statistic fit by one background is greater, the power is also greater. However, it appears that the “Inverted” background is too poor a fit to the data, and consequently no additional power is found at either angle. Further investigation beyond the scope of this thesis may give insight into precisely when and why this is the case.

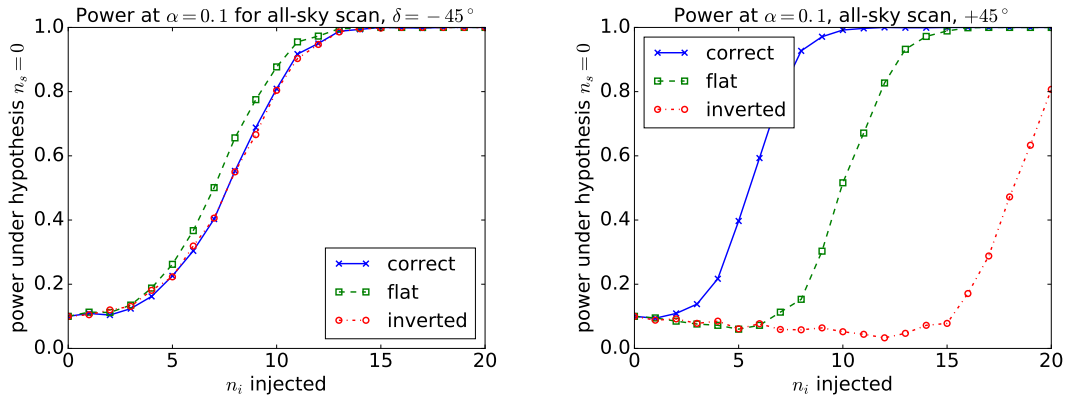
As was done in the test for a fixed direction on the sky, one could test for non-zero hypotheses of signal strength rather than our $n_s = 0$ test in a test of the all-sky hotspot. In principle however it is difficult to interpret tests other than $n_s = 0$ for the all-sky power due to the ambiguity of the injected point source direction and best fit directions, which is not an issue for a fixed direction test of power as the two are always identical. As such, the results of these tests are not presented here.

7.4 Reconciling these results

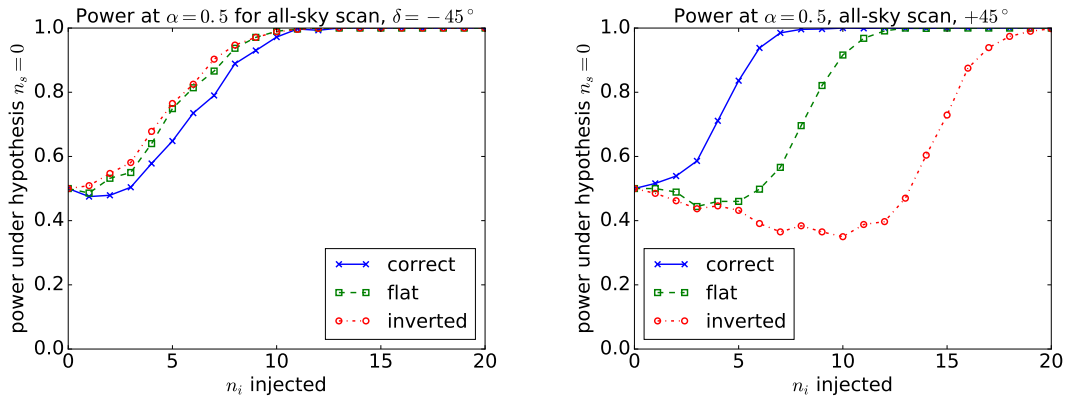
In Figure 7.6 we can see that for all point sources, at a significance level of $\alpha = 0.5$, the power is dependent on the choice of background fitted in the likelihood, and for a point source at $\delta = +45^\circ$ with significance $\alpha = 0.1$ the inverted background showed a difference in power compared to the other backgrounds. Meanwhile, with significance $\alpha = 0.01$ the power at a fixed point is unaffected by choice of background. Further investigation showed that this is due to the nature of the background distribution of the test statistic TS. Looking at Figure 7.1 it is obvious that there are large regions of the sky with $TS = 0$, and similar regions will be present if the events are scrambled in right ascension. As a result, assessing any given position on the sky and re-assessing the same position on the background distribution will result in a large percentage of the TS distribution to be at a value of 0. The effect these $TS = 0$ regions have on a fixed direction background TS distribution can be seen in Figure 7.12: the proportion of test statistics in the $TS = 0$ region depends on the choice of background in the likelihood: the proportion is larger where the predicted background is larger as it takes more signal events when the fitted background



(a) All-sky power against the hypothesis $n_s = 0$ at a significance of $\alpha = 0.01$; $-45^\circ\delta$ (left) and $+45^\circ\delta$ (right).



(b) As above at a significance of $\alpha = 0.1$; $-45^\circ\delta$ (left) and $+45^\circ\delta$ (right).



(c) As above at a significance of $\alpha = 0.5$; $-45^\circ\delta$ (left) and $+45^\circ\delta$ (right).

Figure 7.11: Test of all-sky power against $n_s = 0$ for various levels of injected signal. As expected, we see that unlike the power for a fixed direction, the all-sky power does depend on the choice of background at all tested levels of significance α . In this case, the uniform background test is marginally more powerful than the “Correct” background test in the southern sky, but much weaker in the northern sky.




Figure 7.12: The fixed direction test statistic TS distribution for background events sampled from the “Correct” background distribution, each TS value comes from a sampled map evaluated at -45° and $+45^\circ$ declination.

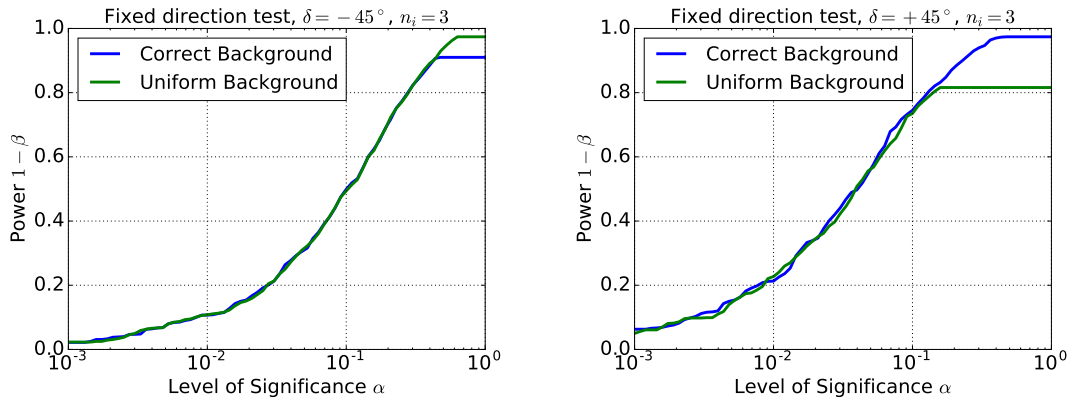


Figure 7.13: The fixed direction power $1 - \beta$ against the range of possible levels of significance α comparing the test of the “Correct” (astrophysical) background to the previous uniform background assumption for a signal injection of $n_i = 3$ events.

distribution is larger for a fit to a point source to be non-zero.

Viewing the power as a function of the level of required significance α (Figure 7.13) showed that the power for a fixed direction at large values of significance $\alpha \gtrsim 0.1$ depends on the choice of background in the likelihood. At the same time small values of significance $\alpha \lesssim 0.1$ show little difference between the choice of background. The distributions show that for small α , the power curves for different backgrounds almost precisely follow each other, while for large α the power of both analyses stay at their own corresponding constant value. This constant value plateau in both curves can be attributed to the corresponding value of α at which $TS_{\text{crit}} = 0$: $\alpha = P(TS > 0|H_0)$. At this value and all larger values of α , the power is precisely $1 - \beta = P(TS > 0|H_1)$. At $\alpha = 0.5$ - which is used in our definition of sensitivity




Figure 7.14: The all-sky test statistic TS distribution for background events sampled from the “Correct” background distribution, each TS value is the hotspot from a sampled background map.




Figure 7.15: The all-sky power $1 - \beta$ against the range of possible levels of significance α comparing the test of the “Correct” (astrophysical) background to the previous uniform background assumption for a signal injection of $n_i = 6$ events.

(see Table 8.1) - the value of TS_{crit} is 0 for our selected backgrounds and the power clearly diverges for the two assumptions of background.

In contrast, the all-sky test statistic TS is a continuous distribution (Figure 7.14). We must find that $TS > 0$ for the distribution of the all-sky statistic as we specifically select the maximum value of the test statistic on the sky (the “hotspot”), the only means by which $TS = 0$ could occur for such a test is by having no events in the sky, or to have no point on the sky at which the likelihood is improved by a point source.

The all-sky power (as seen in Figure 7.15) shows a continuously increasing power with increasing α as a consequence of the all-sky TS forming a continuous distribution. The choice of background can change the all-sky TS distribution by altering the

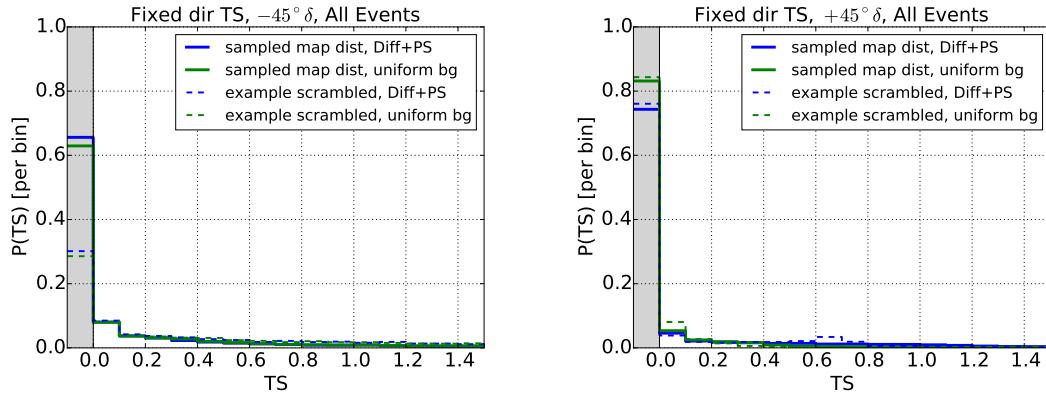
best fit position and test statistic (Eq. 5.26) as a result of under- or over-estimating the fit to signal strength n_s ; these results are consistent with Figure 7.11. Since the corresponding fit in each background map could appear at any location on the sky, this change in signal strength will not be compensated for in the same way as the fixed-point statistic, and for all values α we expect the power $1 - \beta$ of the analysis to be changed by choice of background.

7.5 Diffuse+PS likelihood power on HESE sample

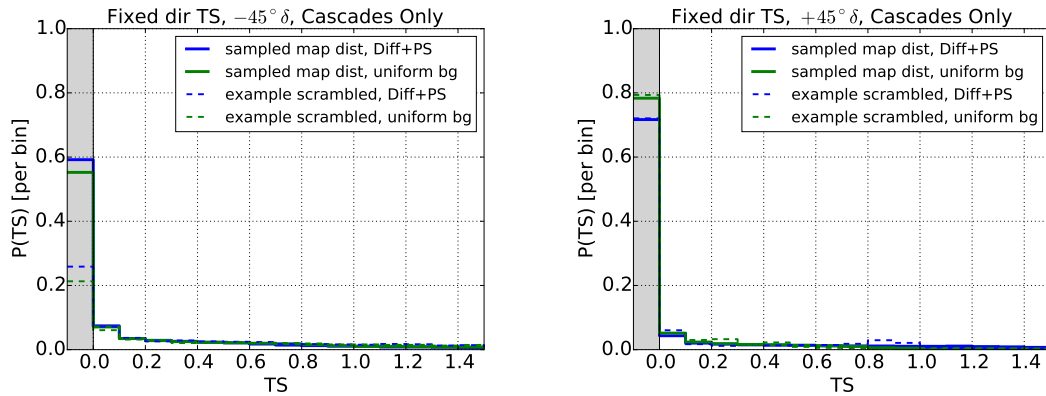
Having now investigated the specifics of the point source test and the effect of the background zenith distribution that is fit in the likelihood, we return to a complete test of the Diffuse+PS likelihood. To understand our eventual calculation of sensitivity and discovery potential it is first beneficial to re-evaluate the TS distributions and then show the effect of changes in the full likelihood on the power of the analysis. The previous tests only took into account the shape in zenith of the background and no additional information that is added by the full Diffuse+PS likelihood. In this chapter, we aim to accurately represent the power of the final analysis: to this end, we:

- sample events from our best description of the sky: an $E^{-2.6}$ spectrum for astrophysical events, and include atmospheric conventional and prompt neutrinos,
- match for each sky the number of tracks and cascades to that of the true event sample using the sampling method described in Chapter 6,
- and use the full diffuse likelihood with four fit components described in Chapter 5.

As a result a small amount of additional power is gained by discrimination of event energy and topology compared to those shown previously, and in most cases we show



(a) Fixed-direction test statistic (TS) distribution for all events (both tracks and cascades) at $-45^\circ\delta$ (left) and $+45^\circ\delta$.



(b) As above, for cascades only.

Figure 7.16: Distributions of the fixed direction test statistic TS for sampled backgrounds (solid lines) and an individual scrambled map seen in Figure 7.1 (dashed lines). The shaded regions correspond to the closed bin $[-0.1, 0]$. As $TS \geq 0$, this bin consists entirely of entries where $TS = 0$.


the results for both tracks and cascades as *all events* and for the subset of the data with no tracks included as *cascades only*. The choice to split the analysis in such a way stems from the decision to do the same in previous High Energy Starting Event analyses in IceCube.

The fixed direction test statistic distributions (Figure 7.16, solid lines) bear strong resemblance to the previous results of Figure 7.12 in that the frequency of $TS = 0$ between the backgrounds depends strongly on declination, however the rates here appear to be in greater agreement between the test of uniform background and test of the Diffuse+PS likelihood, with less variance in the uniform background. Given that the uniform background likelihood is unchanged and independent of en-


ergy and flavour, and the result holds for all events and cascades only, this result can only come about by either the change to an $E^{-2.6}$ diffuse astrophysical flux, the reduced number of events or the inclusion of atmospheric backgrounds.

In addition, the test of signal in the final data necessitates revisiting the effect of scrambling which was not fully assessed in our tests of power using the simplified background model. An example TS distribution for a single map under scrambling is shown as the dotted lines in Figure 7.16. For the majority of cases the distinction between the power under sampled backgrounds and under scrambled backgrounds is small; however there are subtle distinctions to point out. The first is that all events in a zenith band, which maintain their value of energy, are likely to influence each statistic in the background distribution, and so any outlying events such as high energy events can reduce the overall sensitivity. For example, there is significantly reduced point source sensitivity for sources close to the poles under scrambling - any event appearing at the exact poles must appear in the exact same way in all background skies because randomising the right ascension leaves the event unchanged. The second is again due to the fact that the p-value is distinct from the comparison purely of the signal test statistic TS to the null (background) distribution - and must be - as each ensemble sky defines its own background TS distribution. An example of the p-value distribution for a small injection of signal can be seen in Figure 7.17.

When determining the power for a source in a fixed direction under scrambling, the build-up at $TS = 0$ directly impacts the p-value. If, for our experimental map, we find $TS = 0$ at our chosen source position \vec{x}_s , the obtained p-value will be $p = 1$. Therefore when determining the power of our analyses under scrambling we find a distribution of p that consists of two regions: $p = 1$ where $TS = 0$ for the experimental map, and $p < 1$ where $TS \neq 0$ for the experimental map. As there is a large fraction of $TS = 0$ in the majority of scrambled TS distributions for small or no signal injected (see Figure 7.16), these two regions can be noticeably separate (Figure 7.17) and yield similar power curves (Figure 7.18) to the distribution under scrambling.




(a) Distribution of the p-value frequency for all events (both tracks and cascades) at $-45^\circ \delta$ (left) and $+45^\circ \delta$.




(b) As above, for cascades only.

Figure 7.17: An example of the p-value distributions for a relatively weak signal injection (fraction of the diffuse astrophysical flux $f = 0.1$).



(a) Power $1 - \beta$ against all events (both tracks and cascades) at $-45^\circ \delta$ (left) and $+45^\circ \delta$.



(b) As above, for cascades only.

Figure 7.18: Fixed direction power $1 - \beta$ vs significance α assessed under scrambling (from the same data in Figure 7.17: an injected signal of $f = 0.1$).

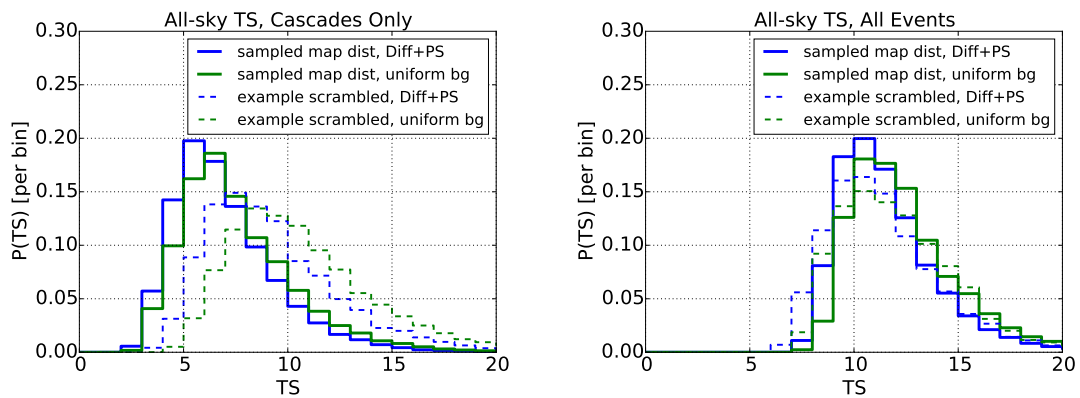


Figure 7.19: The all-sky distributions of TS, each statistic being the maximum value of TS on the entire sky. Solid lines show the distribution aggregated over many background-only sampled skies, and dashed lines show the distribution for the example map (Figure 7.1) under scrambling.




Figure 7.20: Plots of power $1 - \beta$ against significance α for the combined diffuse and point source analysis. The power is shown for an injected signal of $f = 0.2$ at varying declinations δ . Individual lines show the power under different conditions of the atmospheric flux parameters. Results at lower declinations $\delta < -15^\circ$ are consistent with that at $\delta = -15^\circ$ and are not shown.

For the all sky power the TS distribution, as mentioned previously (Section 7.4 and Figure 7.14), does not have the issue of the $TS = 0$ build-up. Example all-sky test statistic distributions for the final Diffuse+PS analysis are shown in Figure 7.19. This figure also shows examples of the distribution for a single randomised sky, showing that one should expect on a map-by-map basis some difference in the test statistic distribution when comparing sampled maps drawn from simulated backgrounds to that of the same sky randomised in right ascension.

To see the impact of priors on the atmospheric flux normalisations, we assessed the all-sky power $1 - \beta$ as a function of α with an $E^{-2.6}$ flux for different tests (Figure 7.20):

- using free atmospheric parameters n_c, n_p ,

- a prior on the atmospheric parameters in accordance with [101],
- and fixed atmospheric parameters at their nominal values (1 times the HKKMS07 atmospheric flux and no ERS prompt contribution).

The results from this test show a small but non-negligible difference in the power, particularly at low values of α which correspond to greater significance. From these results it is apparent that the use of a prior results in an increased power to distinguish a point source compared to keeping the atmospheric flux rates as free parameters. Earlier tests had shown that this improvement in power was likely: given the small event rate, an unconstrained atmospheric flux could increase beyond constraints from other analyses to explain a point source at the horizon, where the atmospheric fluxes predict the largest numbers of events. These earlier tests involved the use of 2 years of data rather than 4, so this effect was further emphasised as there were fewer events to constrain the fit. The most powerful test of a point source would therefore employ fixed atmospheric flux parameters n_c and n_p , however this is not physically motivated as the model conventional and prompt atmospheric flux contributions are not precisely known.

The distribution including all events is peaked at a larger value of TS than the cascade-only distribution due to the presence of track events: they increase the minimum possible likelihood and in the vast majority of cases define the “hotspot”. Larger values are therefore found when cascades overlap with the track events.

7.6 Summary

In this chapter several different tests of the Diffuse+PS analysis were shown. First the statistical terms of power and level of significance which we use to assess the test of the Diffuse+PS likelihood against the test of a uniform background in the likelihood. Then, to investigate the fundamental impact of the background zenith distribution chosen in the likelihood, we reduced the data sample and tests down to more fundamental units: a single background distribution, a selection of many

cascade-type events without track-type events, and removing any testing of the energy.

Then testing the simplified likelihood for a fixed direction we found that, due to the shape of the TS distribution and specifically the build-up of skies where $TS = 0$, we see a plateau of the power $1 - \beta$ where the level of significance α is large. This therefore has an impact on the calculation of sensitivity where $\alpha = 0.5$, as we will see in Chapter 8. Interestingly, we find that for sufficiently small values of the levels of significance α , the choice of background in the likelihood does not appear to influence the power of the test for a source in a fixed direction.

We also investigate the bias inherent in computing the likelihood with a background that does not represent the shape of the data, by looking at tests against non-zero point source signal strengths as the null hypothesis. This showed that under- and over-estimating the background event at the point source location does result in a biased test of signal strength. Finally, with this reduction of the data sample and tests we assessed the power of testing for a “hotspot”, the largest value of TS on the sky, and show that this by its nature does not have the issue of the $TS = 0$ build-up, but the choice of background in the likelihood does impact on the power of the analysis.

Returning to the full test of the Diffuse+PS likelihood, we find that the key results from the reduced tests continue to hold true for the full formalism, and also find that the inclusion of priors in the likelihood does improve the power of testing for a point source.

In the next chapter, we compute the required fluxes to meet the conditions of sensitivity and discovery potential of the full Diffuse+PS analysis and compare these to the test of a uniform background in the likelihood.

Chapter 8

Sensitivity and discovery potential of the Diffuse+PS analysis

In this chapter we evaluate the power of an analysis that uses the full Diffuse+PS likelihood formalism (Equation 5.23). We quantitatively evaluate specific levels of power for the Diffuse+PS fit compared to the assumption of using a uniform background in the likelihood specifically for this analysis with our final event sample. We find that the Diffuse+PS analysis is in all cases more powerful at determining point sources of neutrinos in the northern hemisphere, and is competitive in the southern hemisphere.

8.1 Definitions

The terms Sensitivity and Discovery Potential are measures of the threshold point source flux required to exceed a specific value of power $1 - \beta$ of an experiment at a specific level of significance α . These values are given in Table 8.1.

	α	$1 - \beta$
Sensitivity	0.5	0.9
Discovery (v.1)	0.0013	0.5
Discovery (v.2)	0.0013	0.9

Table 8.1: Values of the level of significance and power for the definitions of Sensitivity and Discovery Potential.

Sensitivity is defined here as it is defined in several other point source analyses (for example see [104]), Discovery Potential is defined here at a significance corresponding to a one-tailed test at 3σ , with differing power $1 - \beta$ (often discovery is quoted at 5σ , with $1 - \beta = 0.5$). We choose 3σ here so the necessary number of trials can be computed within a reasonable time-frame using our full computation of the TS distribution.

To find the point source flux from the injected event rate, we start with the prediction of the diffuse astrophysical flux. The HESE analysis uses weights for the live time of the 4-year analysis which can be recalculated depending on the spectral index γ , and corresponding to a flux of

$$\Phi(E, \gamma) = 3 \times 10^{-18} \left(\frac{E}{100\text{TeV}} \right)^{-\gamma} [\text{GeV}^{-1}\text{s}^{-1}\text{cm}^{-2}\text{str}^{-1}] \quad (8.1)$$

We define n_Φ as the number of events this flux predicts in our detector.

We call n_{exp} the expected rate of events in a small solid angle $\Delta\Omega$ at declination δ .

In regions where the expected rate of events from the diffuse flux is greater on average, for example in the southern hemisphere, a smaller point source flux is needed to result in the same number of events as from a point source elsewhere.

The point source flux corresponding to signal number for sensitivity n_{sens} is therefore

$$\Phi(E, \gamma) = 3 \times 10^{-18} \left(\frac{E}{100\text{TeV}} \right)^{-\gamma} \left(\frac{n_{\text{sens}}}{n_\Phi} \right) \left(\frac{n_\Phi}{n_{\text{exp}}} \right) \left(\frac{\Delta\Omega}{4\pi} \right) 4\pi [\text{GeV}^{-1}\text{s}^{-1}\text{cm}^{-2}] \quad (8.2)$$

and simplifying we obtain

$$\Phi(E, \gamma) = 3 \times 10^{-18} \left(\frac{E}{100\text{TeV}} \right)^{-\gamma} \left(\frac{n_{\text{sens}}}{n_{\text{exp}}} \right) \left(\frac{\Delta\Omega}{4\pi} \right) 4\pi [\text{GeV}^{-1}\text{s}^{-1}\text{cm}^{-2}] \quad (8.3)$$

We present sensitivity and discovery potential in terms of the point source flux capable of producing the mean rate of events observed. As they are both measures of the power of the test, we continue to have two independent definitions of sensitivity and discovery as in Chapter 7: namely fixed-point sensitivity and discovery potential (shown in Section 8.2) and all-sky sensitivity and discovery potential (8.3).

As we assess our final significance under scrambling, we prefer to use scrambling in all calculations, however sampled maps are used in place of scrambling where computation time limits our ability to complete the required calculations.

To assess the sensitivity and discovery potential of our analysis, we wish to find the signal strength that corresponds to our chosen figure of merit, for sensitivity this is a level of significance α of 0.5 and power $1 - \beta$ of 0.9, for discovery potential this is a level of significance α of 0.0017 and a power $1 - \beta$ of 0.5 (v.1) or 0.9 (v.2) (see Table 8.1). To interpolate between signal strengths, we must use the weighting scheme defined in Section 6.3. When assessing the power for a given signal strength, the power is given as the sum of the weights w_i for skies which belong to the critical region (either $\text{TS}_i > \text{TS}_{\text{crit}}$ or $p_i < \alpha$) over the total weights of all maps:

$$1 - \beta = \frac{\sum_i^{N_{\text{skies}}} w_i (\text{TS}_i > \text{TS}_{\text{crit}})}{\sum_i^{N_{\text{skies}}} w_i} \quad (8.4)$$

for the power assessed under sampling and

$$1 - \beta = \frac{\sum_i^{N_{\text{skies}}} w_i (p_i < \alpha)}{\sum_i^{N_{\text{skies}}} w_i} \quad (8.5)$$

for the power assessed under scrambling. Here, w_i is the weight for each sky for the given signal strength assessed given by Equation 6.4 or Equation 6.5, depending on whether all events or solely cascades are evaluated. In other words, the power is

given by the fraction of skies, with the correct weight for each sky, which exceed the level of significance, using either the sampling method or scrambling method to evaluate the significance of each skymap.

8.2 Sensitivity and discovery potential towards a fixed direction

In this section the calculations are shown for the required point source flux to meet the predefined levels of significance and power given in Table 8.1. In figures of the sensitivity and discovery potential in this chapter (Figures 8.1 to 8.3 and Figures 8.5 to 8.7)), the same curves are shown twice in different configurations. First, the top two plots compare the same test (Diffuse+PS or the test of a uniform background in the likelihood) on the different event selections: all events including tracks, and the cascade-only test. Secondly, the bottom two plots compare the same curves in the top two plots but now between the different tests (Diffuse+PS and uniform background) allowing for direct comparisons of the two likelihoods on the same data samples.

In this chapter, for the Diffuse+PS analysis, an $E^{-2.6}$ astrophysical point source signal is injected amongst sampled background and the point source is tested for in the likelihood against an $E^{-2.6}$ diffuse astrophysical component in addition to atmospheric backgrounds constrained by priors in the likelihood.

The sensitivity for a fixed direction point source search is shown in Figure 8.1 indicate a difference in apparent sensitivity between the methods; however we know that this can be due to the change in power by virtue of the ambiguity of the test statistic distribution at $TS = 0$.

As we expect from the plots of power against significance (e.g Figure 7.13), we find the discovery potential (v.1 and v.2) for a fixed direction (Figures 8.2 and 8.3 respectively) show little to no difference between the Diffuse+PS likelihood and the uniform background likelihood. This also suggests that adding the energy and


flavour information to the fit does not significantly impact this analysis. This is not particularly surprising as the astrophysical background appears to produce many more events than the atmospheric backgrounds in this analysis; as such, a point source cannot be easily distinguished from the background on the basis of energy or flavour.

The general shape of the curves, that of a smaller flux in the southern hemisphere and a larger flux in the northern hemisphere, is dominated by the effect of Earth absorption for the signal flux. A much larger point source flux from the northern hemisphere is required to produce a number of signal events compared a point source in the southern hemisphere. This is particularly true in this high energy dataset where the effect of Earth absorption is significant due to the increase in the neutrino cross section with energy.


The tests of the analysis for a fixed direction here use scrambled background maps to assess the significance of each map with injected signal. This can be done with the tests for a fixed direction as this is computationally much easier than for an all-sky scan. As such we observe that the tests closest to the poles (at $\delta = -75^\circ$ and $\delta = +75^\circ$) show a larger flux is required to achieve the same power relative to the rest of the southern hemisphere and equator as expected.

Figure 8.4 shows the ratios between the required fluxes when testing the likelihood using sampling and testing the likelihood using scrambling to calculate sensitivity and discovery potential for a fixed direction. The sensitivity and discovery potential when using scrambling, as expected, is diminished most at the poles, and scrambling visibly affects both all-event and cascade-only analyses due to the size of the cascade PSFs on the sky which are important in both event selections (as a track which aligns with one or more cascades is more significant than a lone track). What is also observed is that discovery potential appears to be affected more than sensitivity, which can be interpreted due to the “scrambling in” of more point source signal events into the background distribution when the signal flux is larger.

It is important to note that although sampling would give an improved sensitivity




(a) Sensitivity testing a fixed direction, comparing all event and cascade-only tests, for the Diffuse+PS analysis (left) and uniform background assumption (right).




(b) Sensitivity testing a fixed direction, comparing the two likelihoods, for all events (left) and cascades only (right).

Figure 8.1: Calculations of sensitivity for a source in a fixed direction, first comparing event selections (top), then comparing the Diffuse+PS likelihood and likelihood testing a uniform background (bottom). The sensitivity was assessed under scrambling.

at the poles and improved discovery potential in general, we choose to use scrambling as the final estimator of the p-value due to the robust nature of scrambling against both systematic variations and changes in the model fluxes (such as spectral index) which are otherwise not accounted for in this analysis.

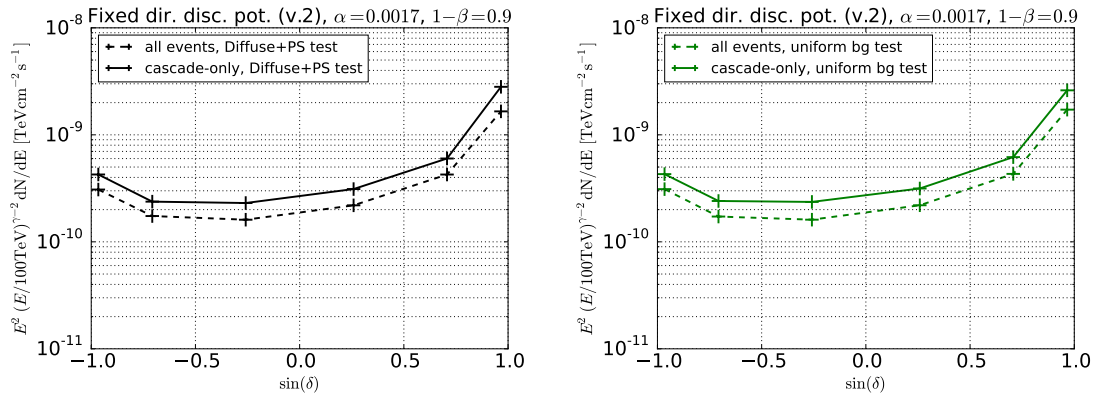


(a) Discovery potential (v.1) testing a fixed direction, comparing all event and cascade-only tests, for the Diffuse+PS analysis (left) and uniform background assumption (right).




(b) Discovery potential (v.1) testing a fixed direction, comparing the two likelihoods, for all events (left) and cascades only (right).

Figure 8.2: Calculations of discovery potential (v.1) for a source in a fixed direction, first comparing event selections (top), then comparing the Diffuse+PS likelihood and likelihood testing a uniform background (bottom). The discovery potential was assessed under scrambling.



(a) Discovery potential (v.2) testing a fixed direction, comparing all event and cascade-only tests, for the Diffuse+PS analysis (left) and uniform background assumption (right).



(b) Discovery potential (v.2) testing a fixed direction, comparing the two likelihoods, for all events (left) and cascades only (right).

Figure 8.3: Calculations of discovery potential (v.2) for a source in a fixed direction, first comparing event selections (top), then comparing the Diffuse+PS likelihood and likelihood testing a uniform background (bottom). The discovery potential was assessed under scrambling.

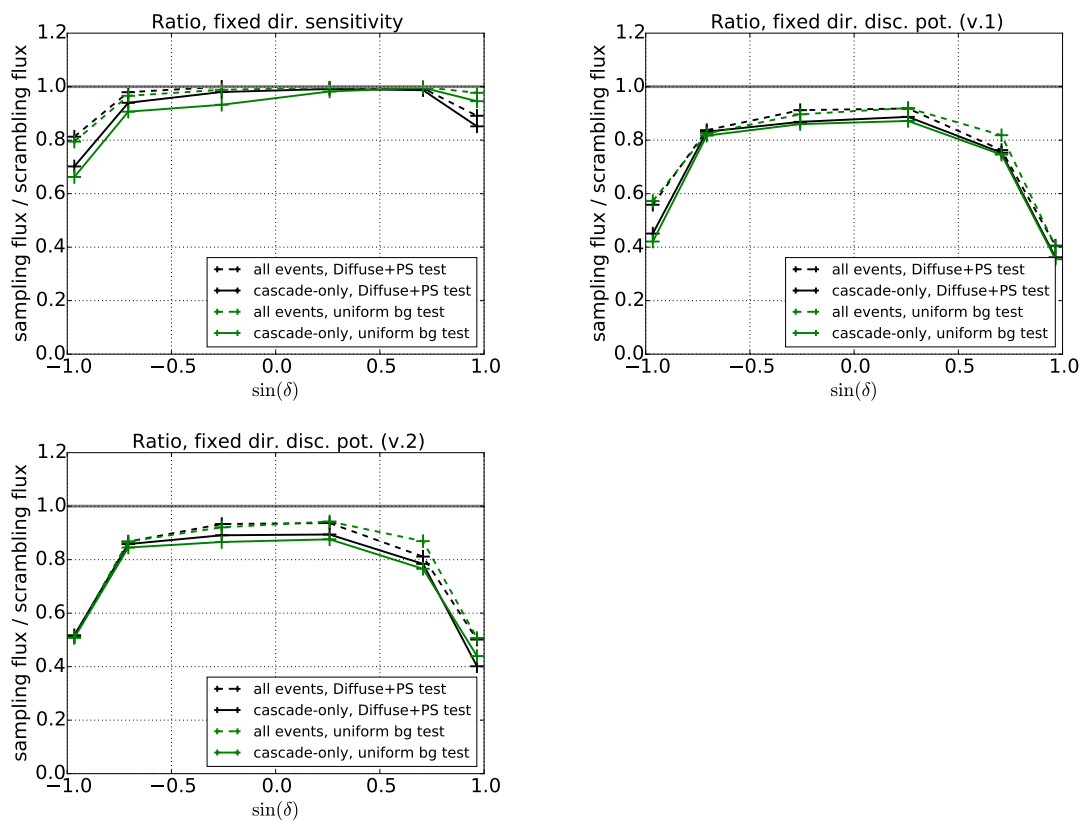
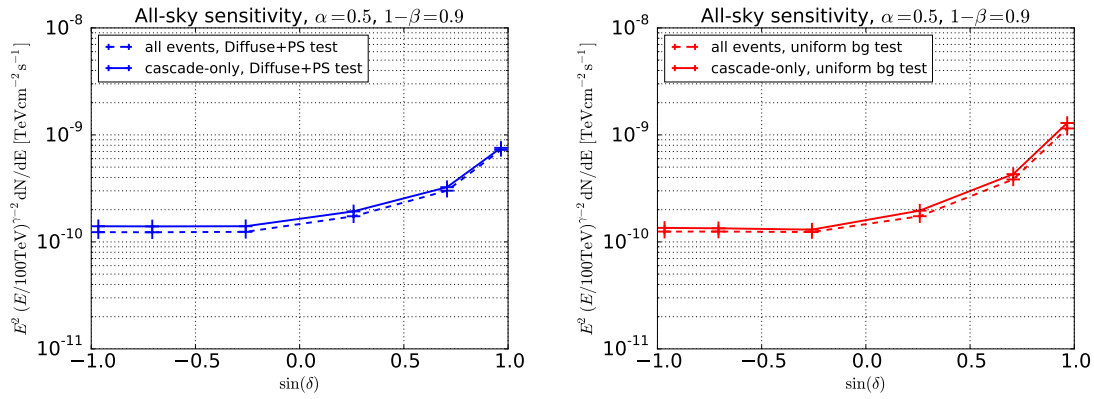


Figure 8.4: The ratio of the required flux using sampling over using scrambling to meet sensitivity (upper left) and discovery potential (v.1 upper right, v.2 lower left), for tests of the likelihoods for a source in a fixed direction.

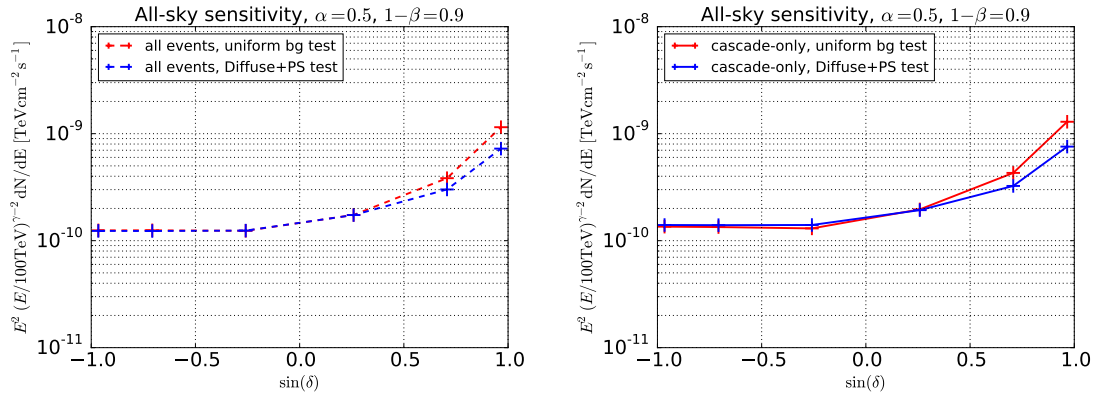
8.3 All-sky sensitivity and discovery potential

The calculation of the all-sky sensitivity and discovery potential are the key measures of the performance of this analysis as it is applied to the data in Chapter 9. The all-sky sensitivity, testing an $E^{-2.6}$ astrophysical point source component, using the full Diffuse+PS likelihood assuming $E^{-2.6}$ astrophysical fluxes with priors on the atmospheric backgrounds, is compared to the test using the uniform background in the likelihood, which be seen in Figure 8.5. As it is far more computationally expensive to assess the all-sky significance under scrambling as each randomised sky must be again scanned over all points on the HEALPix grid, the all-sky sensitivity has been assessed instead using sampled background skies. This also allows us to compute the discovery potential for this analysis which is shown in Figures 8.6 and 8.7. As a consequence of using sampled backgrounds, the sensitivity and discovery potential shown may be enhanced at the poles relative to the result one might see under scrambling.

The results for all-sky sensitivity show a general trend that the northern hemisphere sensitivity is improved by the combined diffuse and point source analysis, and the sensitivity is approximately the same as the test of the uniform background in the southern hemisphere. The results for the all-sky discovery potential show this same trend at the correspondingly higher level of significance α . The Diffuse+PS analysis appears to be a test that is more sensitive to the detection of a point source in the northern sky, due to identifying the relatively low frequency of events in that area, and appears to have minimal change on the test in the southern hemisphere.




(a) All-sky sensitivity, comparing all event and cascade-only tests, for the Diffuse+PS analysis (left) and uniform background assumption (right).




(b) All-sky sensitivity, comparing the two likelihoods, for all events (left) and cascades only (right).

Figure 8.5: Calculations of the all-sky sensitivity, first comparing event selections (top), then comparing the Diffuse+PS likelihood and likelihood testing a uniform background (bottom). The calculation uses sampled background maps.




(a) All-sky discovery potential (v.1), comparing all event and cascade-only tests, for the Diffuse+PS analysis (left) and uniform background assumption (right).

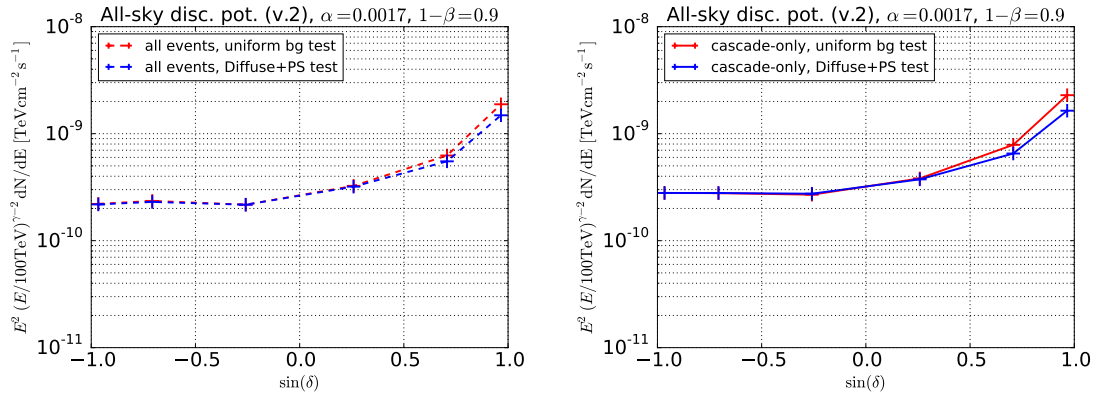


(b) All-sky discovery potential (v.1), comparing the two likelihoods, for all events (left) and cascades only (right).

Figure 8.6: Calculations of the all-sky discovery potential (v.1), first comparing event selections (top), then comparing the Diffuse+PS likelihood and likelihood testing a uniform background (bottom). The calculation uses sampled background maps.



(a) All-sky discovery potential (v.2), comparing all event and cascade-only tests, for the Diffuse+PS analysis (left) and uniform background assumption (right).



(b) All-sky discovery potential (v.2), comparing the two likelihoods, for all events (left) and cascades only (right).

Figure 8.7: Calculations of the all-sky discovery potential (v.2), first comparing event selections (top), then comparing the Diffuse+PS likelihood and likelihood testing a uniform background (bottom). The calculation uses sampled background maps.

8.4 Summary

In this chapter we have defined sensitivity and discovery potential, showed how the fluxes required to meet these definitions can be calculated, and the corresponding fluxes for both the Diffuse+PS analysis and test of the uniform background have been compared. We find that the results in all cases are consistent with the key results of Chapter 7, that the sensitivity for a fixed direction is influenced by the shape of the test statistic distribution, while the discovery potential remains consistent for both tests. We find that for the all-sky sensitivity and discovery potential that, on average, the Diffuse+PS likelihood has resulted in an improvement in that the required flux for this test is lower, but only in the northern sky where other IceCube point source analyses can be used to test for the presence of a point source.

In the next chapter we show the results of the Diffuse+PS analysis on the final data sample, and determine whether this test shows any significant result.

Chapter 9

Results of the Diffuse+PS analysis

The combined diffuse and point source technique was applied to the four years of HESE data. We perform separate analyses for the set of all events passing our selection criteria in Chapter 4, and for the subset of shower topology events that also pass the criteria (a cascade-only search).

9.1 Results

For the all-event point source search we find a “hotspot” centered on event 13, a clear track topology event with 250 TeV deposited energy in the detector. The significance comes from a partial alignment with event 51, a 66 TeV cascade also in the northern celestial sky. The hotspot coordinate is located at $(\delta = 40^\circ, \alpha = 68^\circ)$, with a best fit n_s of 1.71 events (see Figures 9.1 and 9.3). We find that the hotspot, with test statistic 10.93, was not significant under trials as 46.8% of scrambled skies have test statistics greater than this value.


For the cascade-only point source search, we find a “hotspot” close to the galactic centre region. The past analyses of the HESE data have also found that the cascade-only hotspot favours this region in the data. Our coordinate for the maximum of the test statistic is $(\delta = -26^\circ, \alpha = 281^\circ)$, with a best fit n_s of 2.44 events (see Figures 9.2, 9.3). We find that this hotspot, with test statistic 4.05, was also not significant

with 92.2% of scrambled skies having a test statistic greater than this value.


From these results we conclude that the HESE sample above 63 TeV with 4 years of data does not show evidence of a point source of neutrinos.

We can also observe the behaviour of the diffuse fit, that is the fit to n_a , n_c and n_p in addition to n_s for each point on the sky as we see events being attributed to the point source hypothesis and the remaining background. We find that, as one might expect, the point source can be thought of as absorbing events that would otherwise be predominantly attributed to a particular background, and so with these events instead being attributed to a point source hypothesis, that particular background fits to a smaller value (see Figures 9.3, 9.4, 9.5, 9.6). This effect is clearly observed in the fit of the diffuse astrophysical and conventional atmospheric fluxes over the sky, but of these, the fit of the diffuse astrophysical flux is changed the most by the presence of a point source. The fit of a prompt component over the sky is disfavoured for a point source in almost all directions with the exception of the south celestial pole, and then only when considering both track and shower events.

The results from the application of the Diffuse+PS likelihood can be compared to the fit using the uniform background in the likelihood for our subset of the HESE data. The corresponding likelihood scans for the uniform background fit are shown in Figures 9.7 and 9.8. We observe that for the all event scan the hotspots differ based on the choice of likelihood, and as might be expected this therefore results in different p-values. For the cascade only scan, the hotspots correspond to the same event cluster, with p-values that are also similar. The difference between the fits of the Diffuse+PS likelihood and the uniform background are shown in Figure 9.9. We see from these difference maps that the differences between the likelihoods are similar to what we observed from randomly sampled skies: the Diffuse+PS analysis favours the northern sky relative to the uniform background test due to the reduced background rate.




(a) Plot of the test statistic TS for the combined diffuse and point source analysis on the all-event sky in equatorial coordinates. Shower topology events are labelled by “+”, track topology events are labelled as “x”. The “hotspot” of this map with the largest value of TS is highlighted with a red square.




(b) The distribution of largest TS for each realised sky when all events are randomised in right ascension.

Figure 9.1: Results for all-event sky applying the Diffuse+PS likelihood. The hotspot on this map (a) is not considered significant when the right ascension of each event is randomised many times and the hottest spot of the resulting maps is considered (b).



(a) Plot of the test statistic TS for the combined diffuse and point source analysis on the cascade-only sky in equatorial coordinates. Shower topology events are labelled with a “+”. The “hotspot” of this map with the largest value of TS is highlighted with a red square.



(b) The distribution of largest TS for each realised sky when all cascade events are randomised in right ascension.

Figure 9.2: Results for cascade-only sky with the Diffuse+PS likelihood. The hotspot on this map (a) is not considered significant when the right ascension of each event is randomised many times and the hottest spot of the resulting maps is considered (b).

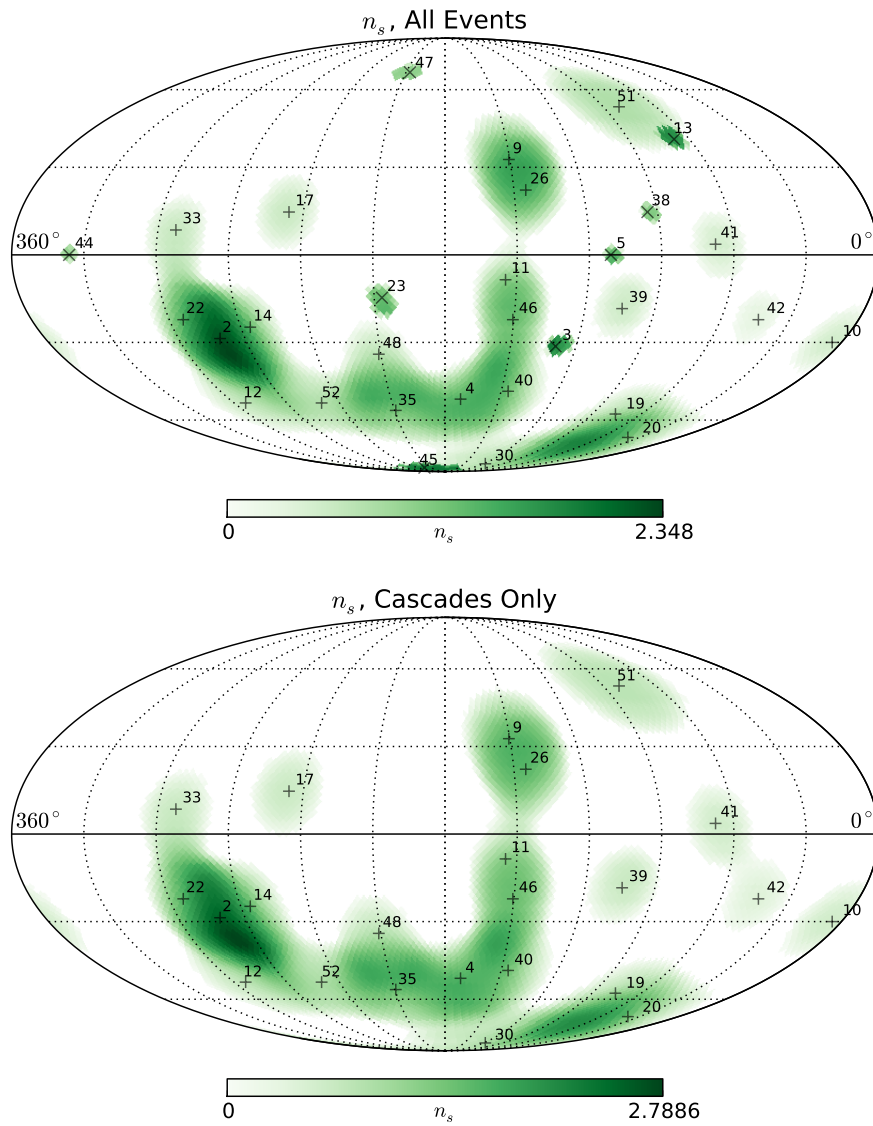


Figure 9.3: Distribution of the fit to n_s , the number of events attributed to a point source at the HEALPix grid locations shown, for all events (top) and cascades only (bottom).

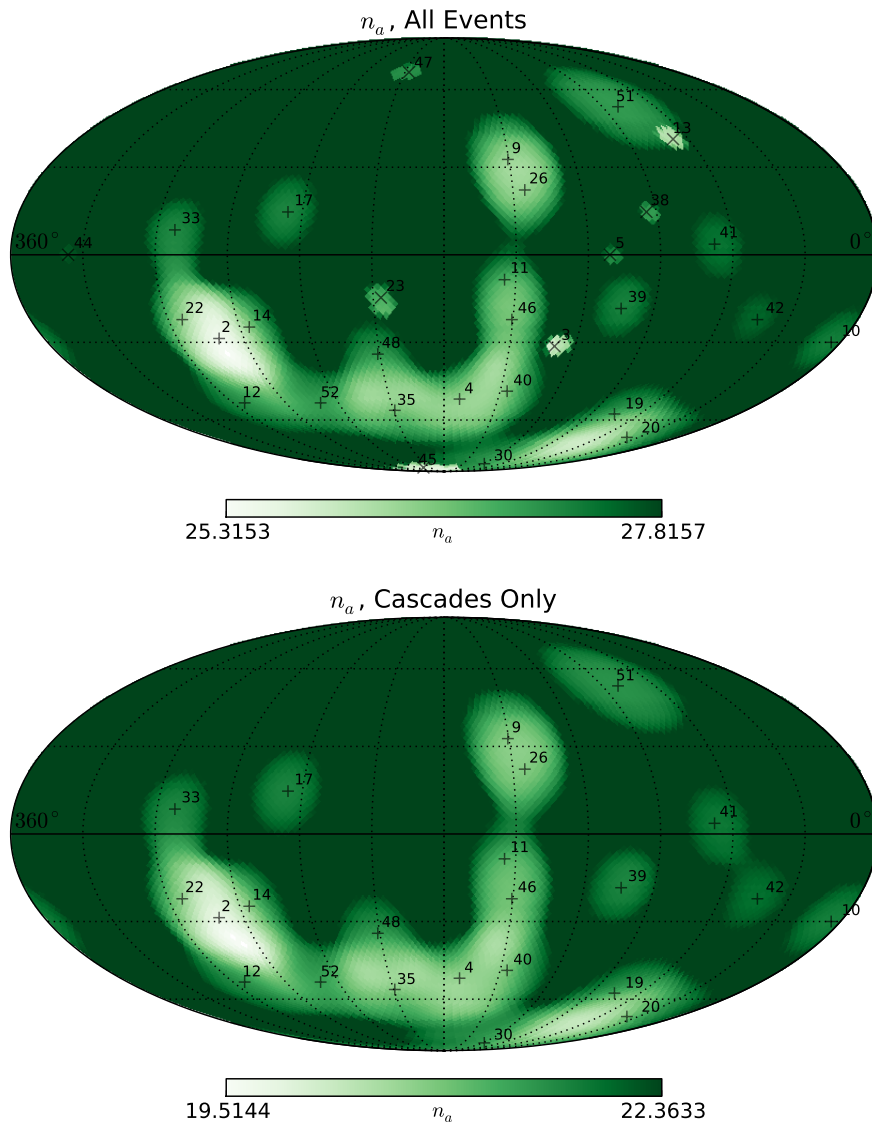


Figure 9.4: Distribution of the fit to n_a , the number of events attributed to the diffuse astrophysical background at the HEALPix grid locations shown, for all events (top) and cascades only (bottom).

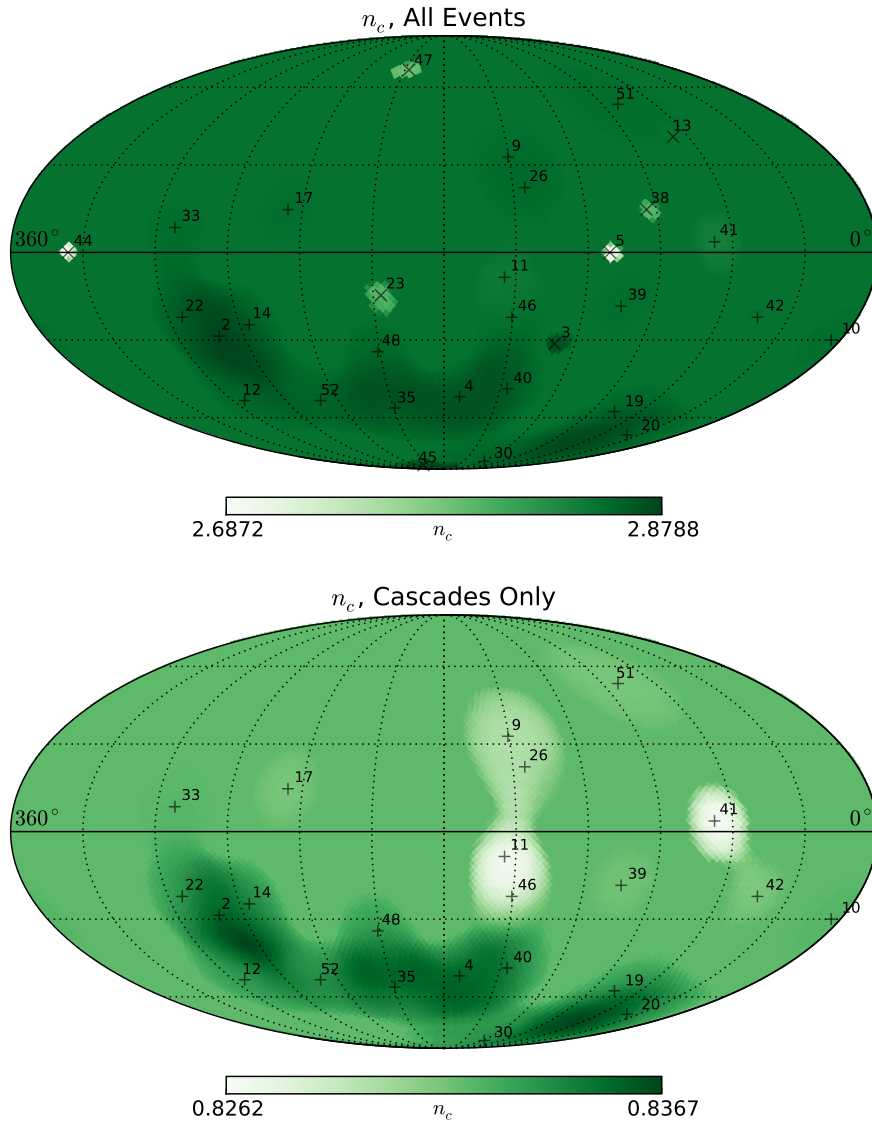


Figure 9.5: Distribution of the fit to n_c , the number of events attributed to the conventional atmospheric flux at the HEALPix grid locations shown. This parameter is constrained with a Gaussian prior, see Chapter 5. Locations where the fit is smaller can be attributed to the point source absorbing events favoured by the conventional flux hypothesis. Note that the net change to the fit here are about $\sim 10\%$ of the conventional flux for the all-sky scan and about $\sim 1\%$ for cascades. This is likely due to the conventional flux being predominantly track topology events, therefore a track being attributed instead to a point source has a larger impact on the fit than a cascade.

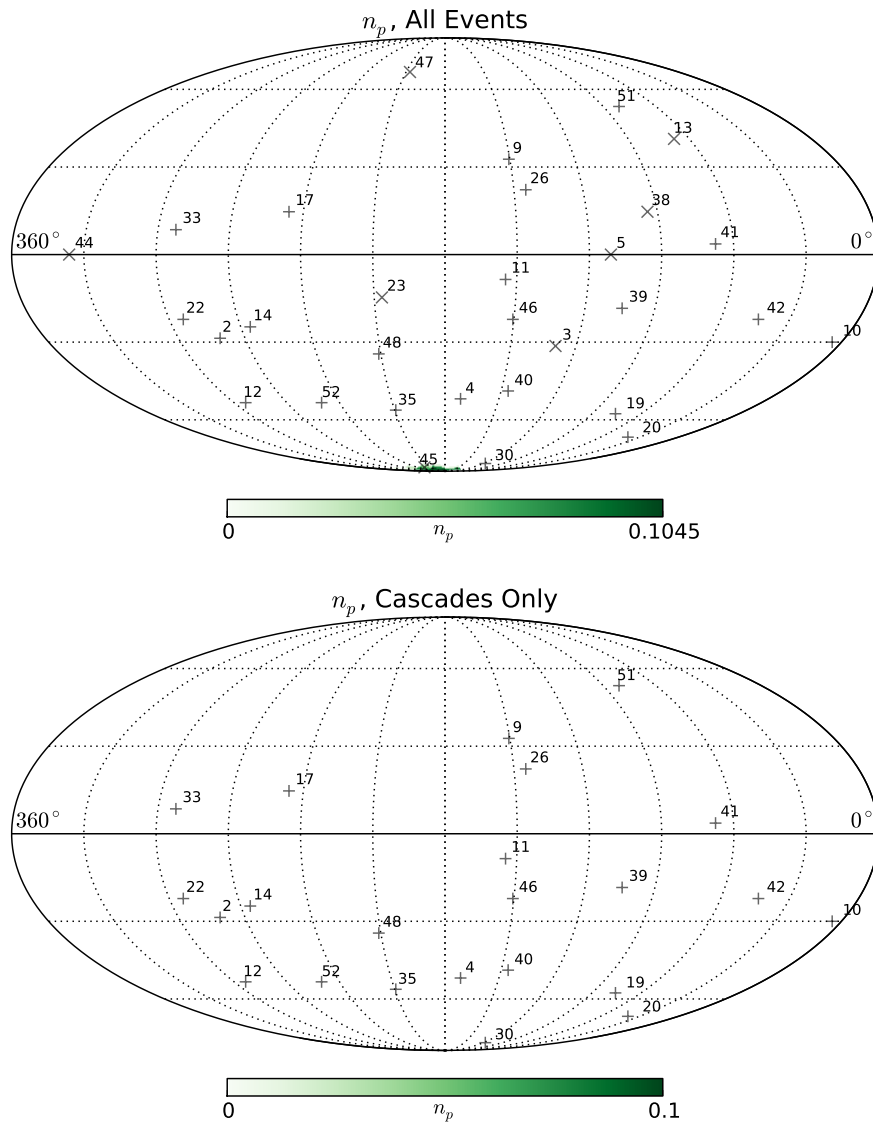




Figure 9.6: Distribution of the fit to n_p , the number of events attributed to the prompt atmospheric flux at the HEALPix grid locations shown. This parameter is also constrained with a Gaussian prior. We find that the fit to prompt appears disfavoured in this scan: the only non-zero fit to the prompt flux is seen in the all-event scan (top) at the south celestial pole (near event 45). In all other locations, including the entirety of the cascade-only scan, the prompt flux has a best fit of 0.




(a) Plot of the test statistic TS for all events (cascades and tracks) under the test of a uniform background in the likelihood. Shower topology events are labelled by “+”, track topology events are labelled as “x”. The “hotspot” of this map with the largest value of TS is highlighted with a blue square.




(b) The distribution of largest TS for each realised sky when all events are randomised in right ascension.

Figure 9.7: Results for all-event sky under the test of a uniform background in the likelihood. The hotspot is located at the south celestial pole (near event 45) and is not significant under scrambling.



(a) Plot of the test statistic TS for cascade events only under the test of a uniform background. Shower topology events are labelled with a “+”. The “hotspot” of this map with the largest value of TS is highlighted with a blue square.



(b) The distribution of largest TS for each realised sky when all cascade events are randomised in right ascension.

Figure 9.8: Results for cascade-only sky under the test of a uniform background. The hotspot is not significant under scrambling.

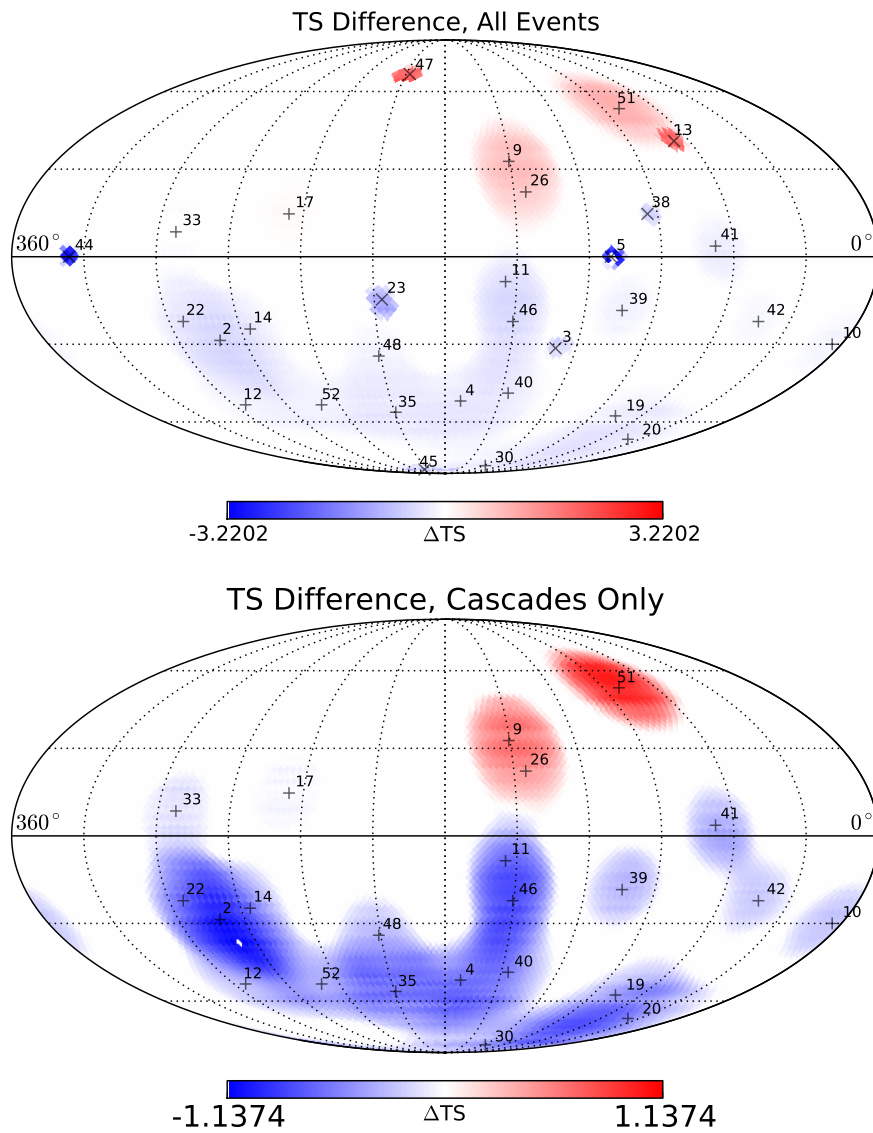


Figure 9.9: Difference in the test statistic $\Delta TS = TS_{\text{diff+ps}} - TS_{\text{uniform}}$. Red hues show a greater test statistic for the Diffuse+PS analysis, blue for the uniform background. The feature around event 5 is a real feature of the test statistic difference between the maps. Top: all events. Bottom: cascade-only sky.

Chapter 10

Conclusion

As shown in Chapter 9, we did not observe a significant fit to a point source component using the combined diffuse and point source analysis likelihood when determining the significance from scrambled trials, with p-values of 0.468 and 0.922 for the all-event and cascade-only event analyses respectively.

The specific methodology used here, where the diffuse background components, estimated from simulated distributions of the background fluxes, are fit separately in addition to a point source, has not previously been used on IceCube data and is to our knowledge a new technique in the field of high energy neutrino astronomy. During the course of this work, some similar analyses which utilised best-fit simulated background distributions in IceCube have been shown on different datasets, for instance lower energy contained cascades [105] and up-going muon neutrinos [106]. In these analyses the different background components are not fit to, instead a global best fit background is assumed. This is because the event number in these analyses is sufficiently large such that a point source contribution is unlikely to affect the fit of the individual background distributions. The specific test of this thesis, which incorporates the fit of several background components, is applicable to any point source data set where the observed data rate is small enough for the fit of events to a point source signal to have an effect on the fit of the remaining background

distributions, and where the separate background distributions from simulation are well known. In this analysis the fit was performed not only in energy and zenith but also in neutrino flavour, however in theory any observable which discriminates point source signal events from background could be introduced into such an analysis.

Through studies on generated example skies, drawn from simulated backgrounds, the effects of assessing a point source with a background distribution in the likelihood that does not match the true event distribution of the data were determined. It was found that fitting with a background that does not describe the data causes a biased fit to a point source, this bias resulting in a fit that favours a point source hypothesis in regions where the background rate fitted in the likelihood is lowest, and giving an incorrect estimation of the signal strength and its error interval as shown in Chapter 7. The results of the all-sky power shows that the biased fit to signal strength, which correlates with the test statistic, also correlates in general with the power of the analysis, such that the use of the correct background in the likelihood results in a slight reduction of power in the southern celestial sky compared to the uniform background, and an increase in power in the northern celestial sky.

The method of scrambling produces a robust post-trial p-value regardless of any induced bias in the fit for a source, and in this way different choices of the background distribution in the likelihood, while they may favour different source locations, will not give a false reporting of a point source.

The sensitivity and discovery potential of the Diffuse+PS analysis, shown in Chapter 8, which uses the correct background in the likelihood in addition to energy and flavour, shows approximately equal sensitivity and discovery potential in the southern celestial sky compared to the use of the uniform background in the likelihood, while the Diffuse+PS analysis does show an improvement in the sensitivity and discovery potential in the northern celestial sky. This result is similar to that described above for tests of the all sky power, as one would expect as the power is directly related to the sensitivity and discovery potential of an analysis, however we note the slight discrepancy in power in the northern sky is now largely

indistinguishable in the sensitivity and discovery potential. This is possibly due to the inclusion of energy and flavour to the fit, however the exact cause is unknown.

As the selection is dominated by astrophysical neutrino events, the addition of energy and flavour do not have a large impact on the power to distinguish a point source from background in this analysis. Were the effects of energy and flavour more able to distinguish a point source, the biases inherent in the choice of background zenith distribution in the likelihood would have been less apparent. Our investigation into bias stemmed from curiosity as to why the Diffuse+PS analysis was not uniformly more powerful over the full zenith range.

There are many future developments that could lead to improvements for this analysis. Additional sensitivity could be achieved by adding the remainder of the event sample below $10^{4.8}$ GeV to the point source search. To do this would require the inclusion of an atmospheric muon background component in the fit as there is a non-negligible contribution of this flux to the sample below this threshold.

The methodology is further applicable to other potential analyses on the HESE data and on other data sets with a small event rate. We find from our studies of the discovery potential in a fixed direction that it is unlikely a search which looks for a source in a fixed direction in the sky, such as a catalogue search, will have a substantially different result between this analysis and an assumed uniform background, due to the equivalent power for a source in a fixed direction being largely independent on the choice of background for sufficiently small values of α . However, other types of searches for sources which are not in a singular fixed direction, such as galactic plane analyses and all-sky extended source searches, should see a difference and potentially an improvement in the result from the application of a combined diffuse and point source style likelihood.

Appendix A

Model parameters in a Poisson likelihood

This note is relevant to the Diffuse+PS likelihood as mentioned in Chapter 5, specifically showing that

$$n_s + n_a + n_c + n_p = N = N_{\text{obs}} \quad (\text{A.1})$$

that is, the total fitted event count matches the observed event count, in an unconstrained Poisson fit. This is a well known result, for full details see Baker and Cousins [107], who show that this result holds when the fit result is unconstrained, i.e. if the parameters are given complete freedom to scale the fitted models. Here I outline my general illustration of the proof.

First we show a trivial result to illustrate the concept on a per-bin basis, however this result is not sufficient to show that the counts will match in the Diffuse+PS likelihood. Consider a Poisson process over n bins which results in observations k_i in bins $i = 1 \dots n$. The Poisson likelihood of the underlying mean rate μ_i in each bin resulting in the observations seen is

$$\mathcal{L}(\vec{\mu}|\vec{k}) = \prod_{i=1}^n \frac{\mu_i^{k_i} e^{-\mu_i}}{k_i!} \quad (\text{A.2})$$

Finding the most probable value of μ_i in each bin i is trivial, first taking the logarithm of both sides:

$$\ln \mathcal{L}(\vec{\mu}|\vec{k}) = \sum_{i=1}^n (k_i \ln \mu_i - \mu_i + \text{const.}) \quad (\text{A.3})$$

and then differentiating with respect to μ_i :

$$\frac{\partial \ln \mathcal{L}}{\partial \mu_i} = \frac{k_i}{\mu_i} - 1 \quad (\text{A.4})$$

and, since the most probable value will be when

$$\frac{\partial \ln \mathcal{L}}{\partial \mu_i} = 0 \quad (\text{A.5})$$

we obtain the result

$$\mu_i = k_i \quad (\text{A.6})$$

that is, the predicted value μ_i which has the highest probability to result in the observed count in each bin, is equal to the observed count.

Now we show a result that is applicable for the full Diffuse+PS likelihood. The models tested in this thesis are such that the relative rates between bins for each model are constrained, and only the overall rate over all bins is allowed to vary in the fit. First we consider two such models, one predicting rates of $a\mu_{a_i}$ in each bin, and the other predicting $b\mu_{b_i}$ in each bin, with both a and b variables free to fit in the likelihood while μ_{b_i} and μ_{a_i} terms, which can differ from bin to bin, are treated as constants.

First these terms which replace μ_i are substituted into the Poisson likelihood:

$$\mathcal{L}(a, b|k_i) = \prod_{i=1}^n \frac{(a\mu_{a_i} + b\mu_{b_i})^{k_i} e^{-(a\mu_{a_i} + b\mu_{b_i})}}{k_i!} \quad (\text{A.7})$$

Again taking the logarithm of both sides we obtain

$$\ln \mathcal{L}(a, b | k_i) = \sum_i^n (k_i \ln(a\mu_{a_i} + b\mu_{b_i}) - (a\mu_{a_i} + b\mu_{b_i}) + \text{const.}) \quad (\text{A.8})$$

and now we perform partial differentiation with respect to the variables a and b to obtain the formulae:

$$\frac{\partial \ln \mathcal{L}}{\partial a} = \sum_{i=1}^n \left(\frac{k_i \mu_{a_i}}{a\mu_{a_i} + b\mu_{b_i}} - \mu_{a_i} \right) \quad (\text{A.9})$$

and

$$\frac{\partial \ln \mathcal{L}}{\partial b} = \sum_{i=1}^n \left(\frac{k_i \mu_{b_i}}{a\mu_{a_i} + b\mu_{b_i}} - \mu_{b_i} \right) \quad (\text{A.10})$$

and as the maximum likelihood will occur when all partial differential equations equate to 0 we substitute this in, giving

$$\sum_{i=1}^n \left(\frac{k_i \mu_{a_i}}{a\mu_{a_i} + b\mu_{b_i}} - \mu_{a_i} \right) = 0 \quad (\text{A.11})$$

and

$$\sum_{i=1}^n \left(\frac{k_i \mu_{b_i}}{a\mu_{a_i} + b\mu_{b_i}} - \mu_{b_i} \right) = 0 \quad (\text{A.12})$$

Multiplying (A.11) by a and (A.12) by b , and then taking the sum of both, we obtain

$$\sum_{i=1}^n \left(k_i \frac{a\mu_{a_i} + b\mu_{b_i}}{a\mu_{a_i} + b\mu_{b_i}} - a\mu_{a_i} - b\mu_{b_i} \right) = 0 \quad (\text{A.13})$$

which simplifies to

$$\sum_{i=1}^n (k_i - a\mu_{a_i} - b\mu_{b_i}) = 0 \quad (\text{A.14})$$

or

$$\sum_{i=1}^n k_i = \sum_{i=1}^n (a\mu_{a_i} + b\mu_{b_i}) \quad (\text{A.15})$$

that is, the total predicted event rate over all bins matches the observed event count

over all bins when the likelihood is maximised. Thus, if

$$\sum_{i=1}^n \mu_{a_i} = \sum_{i=1}^n \mu_{b_i} = 1 \quad (\text{A.16})$$

and we term

$$\sum_{i=1}^n k_i = N \quad (\text{A.17})$$

then we obtain

$$a + b = N \quad (\text{A.18})$$

as desired.

This same procedure can be expanded from two models with variables a and b to the fit of an arbitrary number of models which are scaled by variables c, d, e, \dots : the partial derivative is taken for each model parameter, the resultant equation is multiplied by this factor and the sum taken of all partial differential equations, one for each model, yielding a similar result as equation A.15. Therefore, we can conclude that the equation

$$n_s + n_a + n_c + n_p = N = N_{\text{obs}} \quad (\text{A.19})$$

holds if all parameters are not constrained in the fit. We note that the case of the Diffuse+PS likelihood, where prior functions on the fit variables have been introduced, does therefore not obey this condition, and so simplifications of the likelihood that would be possible if Equation A.19 held cannot be applied. This poses no issue for the minimisation of the likelihood.

References

- [1] Mario Bertolotti. *Celestial Messengers: Cosmic Rays: The Story of a Scientific Adventure*. Springer Science & Business Media, 2012.
- [2] Viktor F. Hess. *Observations of the Penetrating Radiation on Seven Balloon Flights*. In A.M. Hillas (editor), *Cosmic Rays*, The Commonwealth and International Library: Selected Readings in Physics, pages 139 – 147. Pergamon, 1972.
- [3] Thomas K. Gaisser, Ralph Engel, and Eliza Resconi. *Cosmic Rays and Particle Physics, Second Edition*. Cambridge University Press, 2016.
- [4] M. Aguilar et al. *Precision Measurement of the Proton Flux in Primary Cosmic Rays from Rigidity 1 GV to 1.8 TV with the Alpha Magnetic Spectrometer on the International Space Station*. Phys. Rev. Lett., 114:171103, Apr 2015.
- [5] H.S. Ahn et al. *The Cosmic Ray Energetics And Mass (CREAM) instrument*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 579(3):1034 – 1053, 2007.
- [6] *The Pierre Auger Cosmic Ray Observatory*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 798(Supplement C):172 – 213, 2015.
- [7] T. Abu-Zayyad et al. *The surface detector array of the Telescope Array experiment*. Nuclear Instruments and Methods in Physics Research Section A:

- Accelerators, Spectrometers, Detectors and Associated Equipment, 689:87 – 97, 2012.
- [8] H. Tokuno et al. *New air fluorescence detectors employed in the Telescope Array experiment*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 676:54 – 65, 2012.
- [9] James J. Beatty and Stefan Westerhoff. *The Highest-Energy Cosmic Rays*. Annual Review of Nuclear and Particle Science, 59:319–345, 2009.
- [10] A. M. Hillas. *The Origin of Ultra-High-Energy Cosmic-Rays*. Annual Review of Astronomy and Astrophysics, 22:425–444, 1984.
- [11] Luis Anchordoqui, Thomas Paul, Stephen Reucroft, and John Swain. *Ultrahigh Energy Cosmic Rays: The State of the Art before the Auger Observatory*. International Journal of Modern Physics A, 18(13):2229–2366, 2003.
- [12] A. Aab et al. *Searches for Anisotropies in the Arrival Directions of the Highest Energy Cosmic Rays Detected by the Pierre Auger Observatory*. The Astrophysical Journal, 804(1):15, 2015.
- [13] R. U. Abbasi et al. *Indications of Intermediate-scale Anisotropy of Cosmic Rays with Energy Greater Than 57 EeV in the Northern Sky Measured with the Surface Detector of the Telescope Array Experiment*. The Astrophysical Journal Letters, 790(2):L21, 2014.
- [14] A. Aab et al. *Observation of a large-scale anisotropy in the arrival directions of cosmic rays above 8×10^{18} eV*. Science, 357(6357):1266–1270, 2017.
- [15] M. Ackermann et al. *Detection of the Characteristic Pion-Decay Signature in Supernova Remnants*. Science, 339(6121):807–811, 2013.
- [16] Ye Liu et al. *Expectation on Observation of Supernova Remnants with the LHAASO Project*. The Astrophysical Journal, 826(1):63, 2016.

-
- [17] Kenneth Greisen. *End to the Cosmic-Ray Spectrum?* Phys. Rev. Lett., 16:748–750, Apr 1966.
- [18] G. T. Zatsepin and V. A. Kuz'min. *Upper Limit of the Spectrum of Cosmic Rays.* Soviet Journal of Experimental and Theoretical Physics Letters, 4:78, August 1966.
- [19] D Harari, S Mollerach, and E Roulet. *On the ultrahigh energy cosmic ray horizon.* Journal of Cosmology and Astroparticle Physics, 2006(11):012, 2006.
- [20] John G. Learned and Karl Mannheim. *High Energy Neutrino Astrophysics.* Annual Review of Nuclear and Particle Science, 50(1):679–749, 2000.
- [21] W. Pauli. *Open letter to the group of radioactive people at the Gauverein meeting in Tbingen, 1930.*
- [22] V. N. Aseev et al. *Upper limit on the electron antineutrino mass from the Troitsk experiment.* Phys. Rev. D, 84:112003, Dec 2011.
- [23] K.N. Abazajian et al. *Cosmological and astrophysical neutrino mass measurements.* Astroparticle Physics, 35(4):177 – 184, 2011.
- [24] K. Greisen. *Cosmic Ray Showers.* Annual Review of Nuclear Science, 10(1):63–108, 1960.
- [25] F. Reines. *Neutrino Interactions.* Annual Review of Nuclear Science, 10(1):1–26, 1960.
- [26] John N. Bahcall. *Observational Neutrino Astronomy.* Science, 147(3654):115–120, 1965.
- [27] Julia K. Becker. *High-energy neutrinos in the context of multimessenger astrophysics.* Physics Reports, 458(45):173 – 246, 2008.
- [28] Tamar Kashti and Eli Waxman. *Astrophysical Neutrinos: Flavor Ratios Depend on Energy.* Phys. Rev. Lett., 95:181101, Oct 2005.

- [29] Jörg P. Rachen and P. Mészáros. *Photohadronic neutrinos from transients in astrophysical sources*. Phys. Rev. D, 58:123005, Nov 1998.
- [30] T. De Young and Ice Cube Collaboration. *Results from seven years of AMANDA-II*. In *Journal of Physics Conference Series*, volume 136 of *Journal of Physics Conference Series*, page 022046. November 2008.
- [31] Walter Heitler. *The Quantum Theory of Radiation*. Oxford University Press, third edition edition, 1954.
- [32] J. Matthews. *A Heitler model of extensive air showers*. Astroparticle Physics, 22(56):387 – 397, 2005.
- [33] J. Beringer et al. *Review of Particle Physics*. Phys. Rev. D, 86:010001, Jul 2012.
- [34] Ralph Engel, Dieter Heck, and Tanguy Pierog. *Extensive Air Showers and Hadronic Interactions at High Energy*. Annual Review of Nuclear and Particle Science, 61(1):467–489, 2011.
- [35] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw. *CORSIKA: a Monte Carlo code to simulate extensive air showers*. February 1998.
- [36] M. Honda, T. Kajita, K. Kasahara, S. Midorikawa, and T. Sanuki. *Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data*. Phys. Rev. D, 75:043006, Feb 2007.
- [37] Rikard Enberg, Mary Hall Reno, and Ina Sarcevic. *Prompt neutrino fluxes from atmospheric charm*. Phys. Rev. D, 78:043005, Aug 2008.
- [38] Thomas K. Gaisser. *Cosmic Rays and Particle Physics*. Cambridge University Press, 1990.
- [39] Thomas K Gaisser. *Atmospheric Neutrino Fluxes*. Physica Scripta, 2005(T121):51, 2005.

- [40] Thomas K. Gaisser, Kyle Jero, Albrecht Karle, and Jakob van Santen. *Generalized self-veto probability for atmospheric neutrinos*. Phys. Rev. D, 90:023009, Jul 2014.
- [41] M. G. Aartsen et al. *Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere Using Six Years of IceCube Data*. The Astrophysical Journal, 833(1):3, 2016.
- [42] M. G. Aartsen et al. *Evidence for Astrophysical Muon Neutrinos from the Northern Sky with IceCube*. Phys. Rev. Lett., 115:081102, Aug 2015.
- [43] M. G. Aartsen et al. *Atmospheric and astrophysical neutrinos above 1 TeV interacting in IceCube*. Phys. Rev. D, 91(2):022001, January 2015.
- [44] M. G. Aartsen et al. *Measurement of the Atmospheric ν_e Flux in IceCube*. Phys. Rev. Lett., 110:151105, Apr 2013.
- [45] M. G. Aartsen et al. *Measurement of the Atmospheric ν_e Spectrum with IceCube*. Phys. Rev. D, 91:122004, Jun 2015.
- [46] D. Chirkin and W. Rhode. *Propagating leptons through matter with Muon Monte Carlo (MMC)*. ArXiv High Energy Physics - Phenomenology e-prints, July 2004.
- [47] M. G. Aartsen et al. *Measurement of the ν_μ energy spectrum with IceCube-79*. The European Physical Journal C, 77(10):692, Oct 2017.
- [48] Anne Schukraft. *A view of prompt atmospheric neutrinos with IceCube*. Nuclear Physics B - Proceedings Supplements, 237238:266 – 268, 2013. Proceedings of the Neutrino Oscillation Workshop.
- [49] Jeong, Yu Seon et al. *Prompt atmospheric neutrino flux from the various QCD models*. EPJ Web Conf., 141:07002, 2017.
- [50] Joshua P. Ellis. *TikZ-Feynman: Feynman diagrams with TikZ*. Computer Physics Communications, 210(Supplement C):103 – 123, 2017.

-
- [51] J V Jelley. *Cerenkov radiation and its applications*. British Journal of Applied Physics, 6(7):227, 1955.
- [52] Pavel A Cherenkov. *Visible emission of clean liquids by action of γ radiation*. Doklady Akademii Nauk SSSR, 2:451, 1934.
- [53] M.G. Aartsen et al. *Measurement of South Pole ice transparency with the IceCube LED calibration system*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 711:73 – 89, 2013.
- [54] P. W. Gorham et al. *Observations of the Askaryan Effect in Ice*. Phys. Rev. Lett., 99:171101, Oct 2007.
- [55] P. Allison et al. *Performance of two Askaryan Radio Array stations and first results in the search for ultrahigh energy neutrinos*. Phys. Rev. D, 93:082003, Apr 2016.
- [56] Pierre Auger Collaboration. *Ultrahigh Energy Neutrinos at the Pierre Auger Observatory*. ArXiv e-prints, April 2013.
- [57] Raj Gandhi, Chris Quigg, Mary Hall Reno, and Ina Sarcevic. *Ultrahigh-energy neutrino interactions*. Astroparticle Physics, 5(2):81 – 110, 1996.
- [58] Amanda Cooper-Sarkar and Subir Sarkar. *Predictions for high energy neutrino cross-sections from the ZEUS global PDF fits*. Journal of High Energy Physics, 2008(01):075, 2008.
- [59] Sheldon L. Glashow. *Resonant Scattering of Antineutrinos*. Phys. Rev., 118:316–317, Apr 1960.
- [60] S. Sahu and B. Zhang. *On the non-detection of Glashow resonance in IceCube*. ArXiv e-prints, December 2016.
- [61] Takaaki Kajita. *Atmospheric neutrinos*. New Journal of Physics, 6(1):194, 2004.

- [62] K.S. Hirata et al. *Experimental study of the atmospheric neutrino flux*. Physics Letters B, 205(2):416 – 420, 1988.
- [63] Q. R. Ahmad et al. *Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by 8B Solar Neutrinos at the Sudbury Neutrino Observatory*. Phys. Rev. Lett., 87:071301, Jul 2001.
- [64] Christopher W. Walter, , for the Super-Kamiokande collaboration, and for the Super-Kamiokande collaboration. *The Super-Kamiokande Experiment*, pages 19–43. World Scientific, 2012.
- [65] The IceCube-PINGU Collaboration. *Letter of Intent: The Precision IceCube Next Generation Upgrade (PINGU)*. ArXiv e-prints, January 2014.
- [66] J. G. Learned and S. Pakvasa. *Detecting ν_τ oscillations at PeV energies*. Astroparticle Physics, 3:267–274, May 1995.
- [67] J. van Santen. *Neutrino Interactions in IceCube above 1 TeV: Constraints on Atmospheric Charmed-Meson Production and Investigation of the Astrophysical Neutrino Flux with 2 Years of IceCube Data taken 2010 to 2012*. Ph.D. thesis, University of Wisconsin, Madison, 2014.
- [68] M.G. Aartsen et al. *Neutrino oscillation studies with IceCube-DeepCore*. Nuclear Physics B, 908:161 – 177, 2016. Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015.
- [69] M. A. Markov. *On High Energy Neutrino Physics*. In *Annual International Conference on High Energy Physics*. 1960.
- [70] A. Roberts. *The birth of high-energy neutrino astronomy: A personal history of the DUMAND project*. Reviews of Modern Physics, 64:259–312, January 1992.

- [71] M. Ageron et al. *ANTARES: The first undersea neutrino telescope*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 656(1):11 – 38, 2011.
- [72] I.A Belolaptikov et al. *The Baikal underwater neutrino telescope: Design, performance, and first results*. Astroparticle Physics, 7(3):263 – 282, 1997.
- [73] S Adrian-Martinez et al. *Letter of intent for KM3NeT 2.0*. Journal of Physics G: Nuclear and Particle Physics, 43(8):084001, 2016.
- [74] E. Andres et al. *The AMANDA neutrino telescope: principle of operation and first results*. Astroparticle Physics, 13(1):1 – 20, 2000.
- [75] F. Halzen. *Astroparticle physics with high energy neutrinos: from AMANDA to IceCube*. The European Physical Journal C - Particles and Fields, 46(3):669–687, 2006.
- [76] A. Achterberg et al. *First year performance of the IceCube neutrino telescope*. Astroparticle Physics, 26(3):155 – 173, 2006.
- [77] R. Abbasi et al. *IceTop: The surface component of IceCube. The IceCube Collaboration*. Nuclear Instruments and Methods in Physics Research A, 700:188–220, February 2013.
- [78] R. Abbasi et al. *The design and performance of IceCube DeepCore*. Astroparticle Physics, 35(10):615 – 624, 2012.
- [79] M. Ackermann et al. *Optical properties of deep glacial ice at the South Pole*. Journal of Geophysical Research: Atmospheres, 111(D13), 2006. D13203.
- [80] MG Aartsen et al. *South Pole glacial climate reconstruction from multi-borehole laser particulate stratigraphy*. Journal of Glaciology, 59(218):1117–1128, 2013.

- [81] R. Abbasi et al. *The IceCube data acquisition system: Signal capture, digitization, and timestamping*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 601(3):294 – 316, 2009.
- [82] R. Abbasi et al. *Calibration and characterization of the IceCube photomultiplier tube*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 618(13):139 – 152, 2010.
- [83] M. Ackermann et al. *The IceCube prototype string in Amanda*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 556(1):169 – 181, 2006.
- [84] N. Whitehorn. *A Search for High-Energy Neutrino Emission from Gamma-Ray Bursts*. Ph.D. thesis, The University of Wisconsin - Madison, 2012.
- [85] M. G. Aartsen et al. *Energy reconstruction methods in the IceCube neutrino telescope*. Journal of Instrumentation, 9:P03009, March 2014.
- [86] Dmitry Chirkin. *Evidence of optical anisotropy of the South Pole ice*. In *33 rd International Cosmic Ray Conference, The Astroparticle Physics Conference*. 2013.
- [87] J. Feintzeig. *Searches for Point-like Sources of Astrophysical Neutrinos with the IceCube Neutrino Observatory*. Ph.D. thesis, The University of Wisconsin - Madison, 2014.
- [88] IceCube Collaboration. *Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector*. Science, 342:1242856, November 2013.
- [89] Claudio Kopper, William Giang, and Naoko Kurahashi. *Observation of Astrophysical Neutrinos in Four Years of IceCube Data*. PoS, ICRC2015:1081, 2016.

-
- [90] C. Kopper. *Observation of Astrophysical Neutrinos in Six Years of IceCube Data*. PoS, ICRC2017:981, 2017.
- [91] M. G. Aartsen et al. *First Observation of PeV-Energy Neutrinos with IceCube*. Physical Review Letters, 111(2):021103, July 2013.
- [92] Stefan Schönert, Thomas K. Gaisser, Elisa Resconi, and Olaf Schulz. *Vetoing atmospheric neutrinos in a high energy neutrino telescope*. Phys. Rev. D, 79:043009, Feb 2009.
- [93] R. Abbasi et al. *First search for extremely high energy cosmogenic neutrinos with the IceCube Neutrino Observatory*. Phys. Rev. D, 82:072003, Oct 2010.
- [94] M. G. Aartsen et al. *Search for a diffuse flux of astrophysical muon neutrinos with the IceCube 59-string configuration*. Phys. Rev. D, 89:062007, Mar 2014.
- [95] Thomas K. Gaisser. *Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio*. Astroparticle Physics, 35(12):801 – 806, 2012.
- [96] A. Gazizov and M. Kowalski. *ANIS: High energy neutrino generator for neutrino telescopes*. Computer Physics Communications, 172(3):203 – 213, 2005.
- [97] Adam M. Dziewonski and Don L. Anderson. *Preliminary reference Earth model*. Physics of the Earth and Planetary Interiors, 25(4):297 – 356, 1981.
- [98] J.-H. Koehne et al. *PROPOSAL: A tool for propagation of charged leptons*. Computer Physics Communications, 184(9):2070 – 2090, 2013.
- [99] M. G. Aartsen et al. *Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data*. Physical Review Letters, 113(10):101101, September 2014.
- [100] K. M. Grski et al. *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere*. The Astrophysical Journal, 622(2):759, 2005.

-
- [101] Hans Martin Niederhausen, Mariola Lesiak-Bzdak, and A. Stoessl. *High energy astrophysical neutrino flux characteristics for neutrino-induced cascades using IC79 and IC86-string IceCube configurations*. PoS, ICRC2015:1109, 2016.
- [102] J. Braun et al. *Methods for point source analysis in high energy neutrino telescopes*. *Astroparticle Physics*, 29:299–305, May 2008.
- [103] Frederick James. *Statistical Methods in Experimental Physics*. World Scientific Publishing, 2nd edition edition, 2006.
- [104] M. G. Aartsen et al. *All-sky Search for Time-integrated Neutrino Emission from Astrophysical Sources with 7 yr of IceCube Data*. *The Astrophysical Journal*, 835(2):151, 2017.
- [105] M. G. Aartsen et al. *Search for Astrophysical Sources of Neutrinos Using Cascade Events in IceCube*. *The Astrophysical Journal*, 846(2):136, 2017.
- [106] René Reimann. *Search for point-like sources in the astrophysical muon neutrino flux with IceCube*. PoS, ICRC2017:997, 2017.
- [107] Steve Baker and Robert D. Cousins. *Clarification of the use of CHI-square and likelihood functions in fits to histograms*. *Nuclear Instruments and Methods in Physics Research*, 221(2):437 – 442, 1984.