A Search for Ultra-high Energy Cosmic Neutrinos: Data Analysis of the Antarctic Impulsive Transient Antenna, Third Flight

DISSERTATION

Presented in partial fulfilment of the requirements for the degree Doctor of Philosophy in the Graduate School of The Ohio State University

by

Sam Stafford, B.S., M.S.

Graduate Program in Physics

The Ohio State University

2017

Dissertation Committee:

Professor James Beatty, Adviser
Professor John Beacom
Professor Amy Connolly
Professor Richard Kass

Abstract

Ultra-high Energy (UHE) neutrinos represent an increasingly important messenger in astronomy and astrophysics. The Antarctic Impulsive Transient Antenna (ANITA) experiment campaign utilizes a balloon-borne phased antenna array to detect coherent Cherenkov radio-frequency pulses induced by UHE neutrinos interacting with the Antarctic ice. We analyzed the data from the third ANITA flight (ANITA-III) for evidence of Ultra-high energy neutrinos by augmenting interferometric methods used in analyses of previous ANITA flights. Continuous wave (CW) radio content from ground-based Antarctic habitations and orbiting geostationary communications satellites interferes with the detection and analysis of neutrino-induced radio signals; we developed circular polarization analysis methods to facilitate improved rejection of false positives induced by satellite CW. We also developed new methods of calculating signal-to-noise ratio (SNR) of event waveforms, and enhanced event localization by applying a probability distribution function (PDF) based on the measured resolution of our interferometry. We developed a final linear discriminant cut for rejecting thermal and anthropogenic signals by dividing the continent into equal-area bins and optimizing the cut to each individual bin, so as to obtain the strongest possible upper limit on cosmic neutrino flux.

to my dad, Joseph H. “Staff” Stafford

NJ8F

1934 - 1993

Acknowledgments

This work is possible in great part due to persons who showed greater confidence in me than I had in myself.

I acknowledge first and foremost my wife Lisa for her unending support and encouragement, without which I could not even have considered graduate school. Lisa's seemingly limitless belief in me, and her hard work as a businesswoman to provide financial support for our family, have been essential to my success. I also thank my daughters Ruth Volk and Dr. Emily Stafford for ongoing encouragement and validation. They have grown into fine adult women of whom I am exceptionally proud.

My late father, Joe Stafford, instilled in me a passion for learning and discovery, particularly in the areas of science and mathematics, as well as the intellectual fortitude to put that passion to action. Knowledge brings me joy, and I wish Dad were still here to share in it. My mother Sue Stafford has provided love and support during my graduate career, in times both good and hard.

I extend my sincere thanks to Professor Michael Stamatikos for his vigorous advocacy during my early graduate career. Mike introduced me to the scientific community, facilitated a research project that helped me build credentials to get admitted to graduate school, and sponsored travel for me during my first year. He also helped me obtain early graduate school funding through the Center for Cosmology and Astro-particle Physics (CCAPP).

Professor John Beacom welcomed me to Ohio State as a late-returning student and provided ongoing encouragement and advice, particularly during my early years in graduate school when my classwork seemed overwhelming. I thank Professor James Beatty for serving as my faculty advisor and for demonstrating patience and understanding during my adjustment, often very difficult, to the culture and climate of graduate school. Our conversations as fellow musicians provided frequent relief from the pressures. Rock on, Jim.

I am especially grateful to Professor Amy Connolly for her close guidance and assistance, particularly during the completion of my thesis this year. Thank you Amy for your patience, and for demanding more of me than I thought myself capable.

Numerous other colleagues have assisted me along the way, inspiring warm feelings of camaraderie and earning my full intellectual respect. Paul Schellin distinguished himself as an intellectually challenging academic peer and a true friend, especially during payload integration in Texas in the summer of 2014. Dr. Patrick Allison has continually challenged me throughout and has treated me with respect; his enthusiasm for discovery and astute intellect are ongoing inspirations. The friendship and engagement of my office partners brought a breath of fresh air to my academic endeavors during my last two years of graduate school: I appreciate Brian Clark for his friendly demeanor and for the example he has set by his positive attitude and work ethic; I've enjoyed working with Oindree Banerjee, a respected colleague, friend and fellow cat-fancier. My conversations with Dr. James Stapleton, about science, about graduate school, and about the trials of life in general helped me sustain through the ups and downs. I thank Dr. Brian Dailey for his friendship and advice, and for reminding me that this is all possible. I am grateful for the constructive advice and guidance of Jacob Gordon, Dr. Michael Sutherland, Dr. Carl Pfendner and Dr. Jordan Hanson. I thank Professor Richard Kass for filling in my committee on short notice. The ANITA collaboration, funded by NASA, was responsible for the majority of my graduate school financial support.

Finally, I take pride in my family's five-generation affiliation with The Ohio State University.

Vita

1959 .... Born, Columbus OH

1990 .... B.S. Engineering Physics, cum laude, The Ohio State University

2010-2012 .... Graduate Teaching Associate, The Ohio State University

2011 .... Hazel M. Brown Teaching Award, The Ohio State University

2012 .... M.S. Physics, The Ohio State University

2012-2016 .... Graduate Research Associate, The Ohio State University

Publications

Antarctic Surface Reflectivity Measurements from the ANITA-3 and HiCal-1 Experiments, P. W. Gorham et al., J. Astron. Inst. 06 (2017) no. 2, 1740002

Characteristics of Four Upward-pointing Cosmic-ray-like Events Observed with ANITA, P. W. Gorham *et al.*, Phys. Rev. Lett. 117 (2016) no. 7, 071101

The Prompt-afterglow Connection in Gamma-ray Bursts: A Comprehensive Statistical Analysis of Swift X-ray Light Curves, R. Margutti *et al.*, Mon.Not.Roy.Astron.Soc. 428 (2013) 729

Fields of Study

Major Field: Physics

Specialization: Experimental Astro-particle Physics, Ultra-high energy neutrinos

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Ultra-high energy (UHE) neutrinos represent a new and unique window into both particle physics and astrophysics. From a particle physics perspective, UHE neutrinos (UHEν) reveal details of interactions at energy levels far above those obtainable from synthetic sources; for astrophysics, they provide visibility into the most powerful and remote phenomena in the universe. In this work, we attempt to detect these neutrinos.

Throughout most of human history, observation of the universe was limited to the visible portion of the electromagnetic spectrum as perceived by the unaided human eye. The invention of the telescope about 400 years ago and subsequent refinements over centuries dramatically increased the sensitivity of optical observations, but our view of the cosmos remained limited to the optical regime until the development of infrared astronomy in the late 19th century. The mid-1900's witnessed the advent of radio astronomy; the subsequent development of ultraviolet, x-ray and gamma-ray instrumentation revealed a wide swath of the electromagnetic spectrum, spanning over twenty orders of magnitude in wavelength, from 10^{-20} m to 10^{2} m. The discovery of cosmic rays (high-energy charged particles from space) in 1910 was the first non-photonic astronomical observation, and marked the birth of multi-messenger astronomy and astro-particle physics. Then in the 1930's, Wolfgang Pauli proposed a massless neutral particle, later named the neutrino, to account for apparent conservation violations in beta decay. Neutrinos from the Sun were discovered in the 1960's[1], and a few years later a burst of neutrinos was detected from supernova SN1987A[2], introducing the neutrino as yet another class of cosmic messenger. By 2004, solar neutrinos were being used to generate crude images of the Sun[3]. The Antarctic Impulsive Transient Antenna (ANITA) campaign, the subject of this work, commenced in 2005, searching for radio-frequency (RF) impulses induced by UHE neutrinos interacting in ice.

1.1 The Neutrino

The Standard Model of particle physics includes a uncharged and nearly massless weakly-interacting particle, the *neutrino*, of which exist three *flavor* states, named the electron (ν_e), muon (ν_μ), and tau (ν_τ) neutrinos, in correspondence to the three lepton species. The neutrino is subject to neither electromagnetic interactions nor nuclear strong interactions; neutrino physics is thus restricted to the neutral-current (NC) and charged-current (CC) weak interactions between neutrinos, quarks and leptons (Figure 1.1). In the CC interaction, the neutrino exchanges a W boson with a nucleon and is transformed into its corresponding lepton. In the NC case, a Z boson is exchanged the nucleon but no lepton is produced. Identification of the outgoing lepton from a CC interaction is one strategy for detecting neutrinos; another is measuring element species changes from nucleon neutrino absorption, either directly or from subsequent decays. At ultra-high energies in matter, however, a particle cascade known as *shower* will occur. The shower can consist of hadrons and/or electromagnetic components (electrons, positrons, photons); the electromagnetic constituent is of primary interest in UHE neutrino detection.

An electromagnetic shower is initiated when a UHE neutrino traversing matter strikes an atomic nucleus, freeing high-energy electrons and nuclear fragments. Liberated electrons scatter off atoms, emitting photons through bremsstrahlung. These photons in turn scatter against atoms, creating e^+e^- pairs via pair-production, and freeing additional electrons via Compton scattering. Subsequent bremsstrahlung emissions and e^+e^- annihilations produce even more photons. A chain reaction ensues, and a conglomeration of e^+ and e^- begins to accumulate. The dimensions of the shower depend on the density of the medium and can vary from a few centimeters in solids to hundreds of meters in air. Compton electrons and positron annihilation contribute to a negative charge excess of about 20% in the conglomeration. As more particles are drawn into the shower, the energy per particle of course decreases. The cascade continues until the energy per $e^\pm$ falls to about 100 MeV; at this point ionization replaces bremsstrahlung as the primary stopping mechanism and the shower ceases to develop. The shower products emit photons by the Cherenkov[1] radiation process described below in Section 1.5.1. Detection via Cherenkov radiation is the prevailing strategy in UHE neutrino physics.

Neutrino masses are known to be exceedingly small ($< 10^{-6} m_e$)[5], and multiple mass states exist which do not correspond to the neutrino flavor eigenstates. *Neutrino oscillations* therefore occur, in which neutrinos change flavor over time at an energy-

---

[1] This transliteration from the Russian Cyrillic is used by Jackson[4]. Some writers use *Čerenkov*.

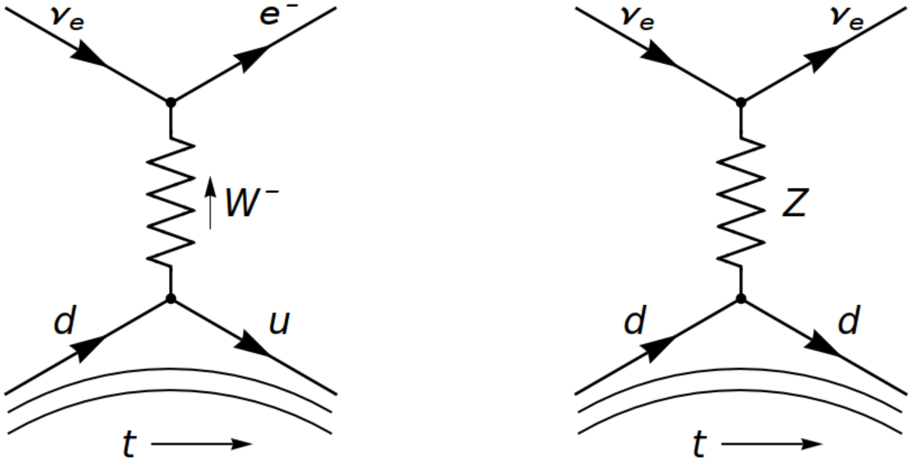

Figure 1.1: Examples of neutrino interactions. Left: deep inelastic scattering (a charged-current interaction). Right: elastic scattering (neutral-current).

dependent rate. Neutrino oscillations caused discrepancies in early measurements of the solar neutrino flux, as described in Section 1.3.

1.2 Origins and Motivations

Ultra-high energy neutrinos are expected to originate through an interaction sequence beginning with a photon-nucleon (γN) interaction such as [2]

$$p + \gamma \quad \longrightarrow \quad \Delta^+ \quad \longrightarrow \quad \begin{array}{c} p + \pi^0 \\ \text{or} \\ n + \pi^+ \end{array} \tag{1.1a}$$

---

[2] The interaction is actually a quark-photon ($q\gamma$) interaction and in general may occur with any nucleon in any atomic nucleus.

The weak decay of the π^+ produces neutrinos:[3]

$$\begin{aligned} \pi^+ &\rightarrow \mu^+ + \nu_\mu \\ \mu^+ &\rightarrow e^+ + \nu_e + \bar{\nu}_\mu \end{aligned} \tag{1.1b}$$

This may occur either within an accelerating engine such as a gamma-ray burst (GRB) or active galactic nucleus (AGN) where sufficiently energetic protons and photons are abundant, or via interactions with galactic and extragalactic photon backgrounds. A particularly important case of equation (1.1) is the Greisen-Zatsepin-Kuzmin (GZK) process, where ultra-high energy ($>10^{19}$ eV) cosmic-rays interact with the cosmic microwave background (CMB), typically within 50 Mpc of the CR source[6]. We therefore expect a distinct cutoff of earth-incident CR flux at about 10^{19} eV, and a complementary flux of GZK-produced UHEνs.

The total flux of UHE neutrinos at earth from the GZK process and source progenitors is on the order of $10^{-2}/\text{km}^2/\text{year}$[7]; therefore detection is challenging. At target-frame energies above the 6 PeV Glashow resonance, a νN interaction dominates, with cross-section $\sim 10^{-33}\,\text{cm}^2$[8], leading to an interaction rate in water of $\sim 10^{-2}/\text{km}^3/\text{year}$[7].

1.2.1 Cosmic Rays

The close association between UHE cosmic rays and UHE neutrino astrophysics merits brief discussion. Cosmic rays are energetic particles from space; they include all stable ($\tau > 10^6$ yr) charged particles and light nuclei up to and including iron[10], but their detailed composition and origin remain major questions. The high-energy spectrum of

---

[3] The π^0 decays electromagnetically and does not contribute to neutrino production.

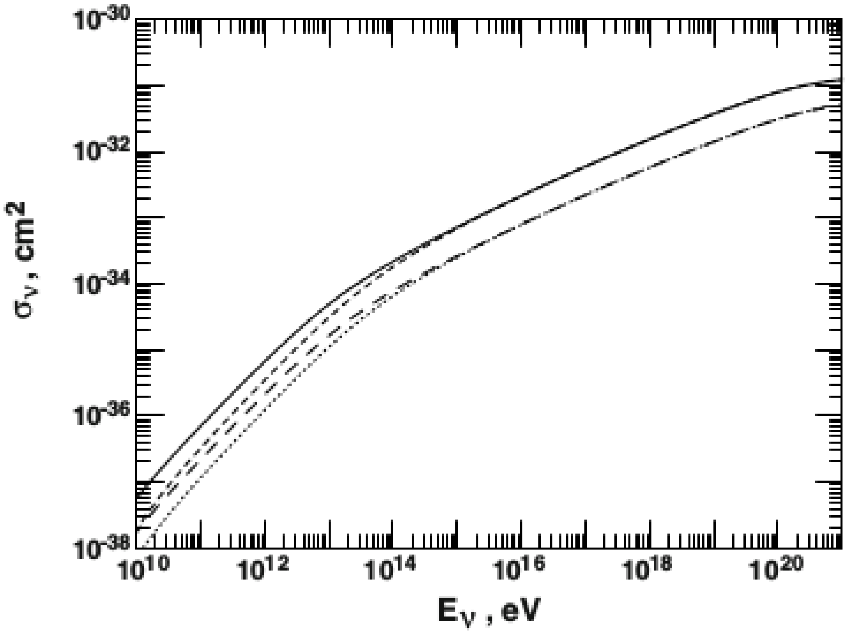

Figure 1.2: High-energy neutrino cross sections. Neutrino CC, antineutrino CC, neutrino NC and antineutrino NC cross sections are denoted respectively by the solid, short-dash, long-dash and solid lines. In the UHE regime ($>10^{18.5}$ eV, the CC cross section is $\sim 10^{-33}\,\mathrm{cm}^2$ dominating the NC by about 5:1. From [9].

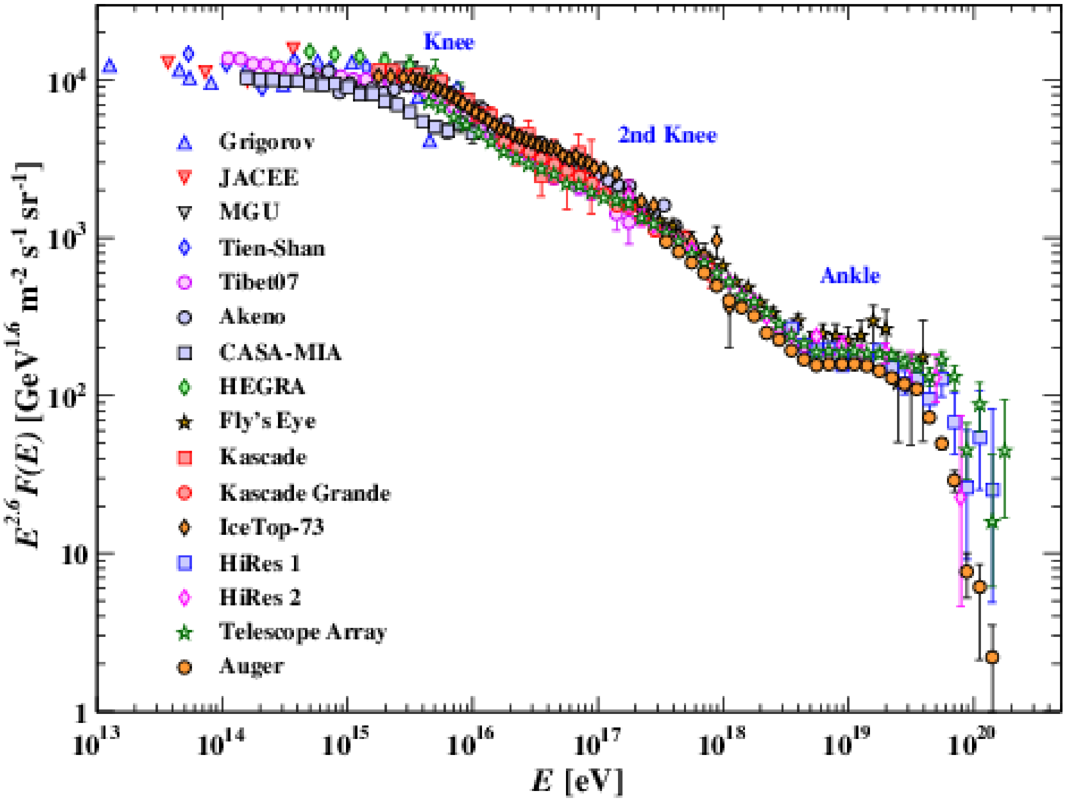

Figure 1.3: The cosmic ray spectrum, multiplied by $E^{2.6}$ to elucidate the slope changes. The rapid fall-off at energies above $10^{19.5}$eV may be due to the GZK effect, a neutrino-producing interaction with the CMB. From [10].

cosmic rays, shown in Figure 1.3, follows a broken power law in E. At energies below about 10^{18} eV, CR's are thought to be of galactic origin[10]. The spectral steepening near $10^{15.5}$ eV, called the *knee*, may be due to energy constraints at galactic sources, e.g. supernova remnants[10]. The flattening at $10^{18.5}$ eV, the *ankle*, is likely the takeover of a higher-energy, extragalactic constituent of CR's[10]. The final break, a sharp fall-off at $10^{19.5}$ eV, is consistent with the GZK process but also could indicate a cutoff in the source spectrum. Since the GZK process should produce neutrinos, the discovery a UHEν flux would be consistent with the hypothesis that GZK is, at least in part, responsible.

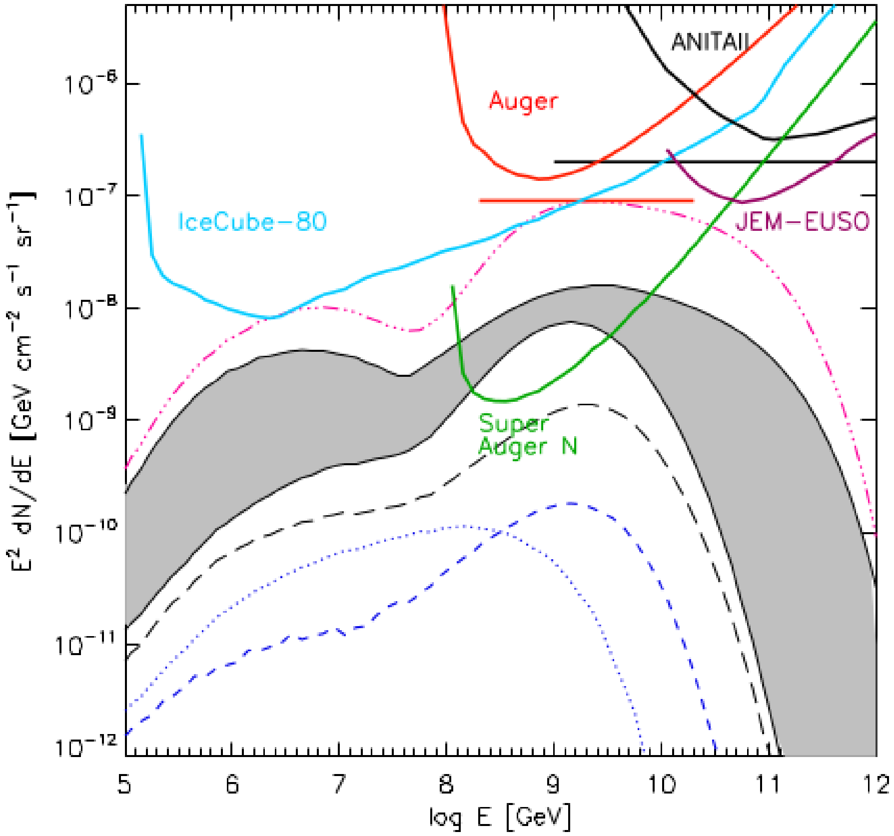

Figure 1.4: Models and limits of cosmic neutrino energy spectrum. From [11].

1.2.2 UHE Neutrino Flux

Various theoretical models have been formulated for predicting UHE neutrino flux at earth. Figure 1.4 shows flux predictions from Kotera et al.[11], and limits imposed by previous observations. The dashed and dotted blue lines represent pessimistic models assuming iron-rich CR composition, where the comparatively low energy per nucleon leads to lower neutrino production. The most optimistic model, represented by the pink dashed line, assumes a pure proton composition.

1.3 Ultra-high Energy Neutrino Astronomy

The energy of a particle arriving from space is related to the energy conditions at its source. Particles in the UHE regime thus provide insight into the highest-energy phenomena in the universe, such as Active Galactic Nuclei (AGN) and supernovae. Among the various messengers, the neutrino is of particular utility because of its uniquely long and straight path across the cosmos. Cosmic rays on the other hand are inhibited as long-distance messengers: Lorentz deflection in the extragalactic and galactic magnetic fields limits their pointing value; furthermore, the GZK process will attenuate CR's for center-of-momentum energies above $m_\Delta = 1232\,\mathrm{MeV}$. With a characteristic CMB photon energy $E_{CMB} = 0.7\,\mathrm{meV}$, the UHECR source-frame threshold energy E_{GZK} can be estimated by equating four-momenta:

$$m_\Delta^2 = (E_{GZK} + E_{CMB})^2 - (p_{GZK} - E_{CMB})^2 \approx m_p^2 + 4E_{GZK}E_{CMB}$$

$$E_{GZK} \approx \frac{m_\Delta^2 - m_p^2}{4E_{CMB}} \sim 10^{20}\,\mathrm{eV}. \tag{1.2}$$

A $p\gamma$ cross-section $\sim 10^{-4}\,\mathrm{mb}$ and CMB photon density $n \approx 500\,\mathrm{cm}^{-3}$ give an attenuation length of

$$\lambda_{GZK} = \frac{1}{\sigma n} \sim 2 \times 10^{26}\,\mathrm{cm} \sim 100\,\mathrm{Mpc}.$$

High-energy photons (gamma rays) are also attenuated by the CMB, through e^+e^- pair production with a threshold energy of $\sim 10^{15}\,\mathrm{eV}$. The $\gamma\gamma$ cross-section is of the same order of magnitude as that of $p\gamma$, and thus γ-rays above $10^{15}\,\mathrm{eV}$ are also attenuated within 100 Mpc of their sources[12]. The neutrino, in contrast, subject only to comparatively rare interactions via the weak nuclear force[4], can traverse great

---

[4] Photons and neutrinos are also both subject to the gravitational force; the effect on current UHEν astronomy is negligible.

distances undeflected and unattenuated, and in doing so points back to its source. Observational horizons thus exist for photons and cosmic rays, but not for neutrinos.

The neutrino's absence of electric charge precludes it from being directly accelerated to high energies (in supernovae, for example). UHE neutrinos therefore must be *created* either in the ultra-high energy environment of a source progenitor, or by interactions of UHE cosmic rays in interstellar or extragalactic space. The GZK prediction, and the detection of rare UHE cosmic rays above 10^{20} eV starting in the 1960's [9], are consistent with the expectation of a UHEν flux. UHEν's thus provide a unique window into ultra-high energy ($>$ PeV) astrophysical phenomena.

Another attraction of UHEν's to astrophysics is the fact that they possess much greater energy than is attainable in terrestrial laboratories such as the Large Hadron Collider (LHC), thereby providing a view into particle physics at energies inaccessible to the latter. UHEν energies are typically expressed in the target (earth) reference frame, while LHC energies are in the center-of-momentum (COM) frame. Since energy is not Lorentz-invariant, the two cannot be directly compared. Nevertheless, an earth-frame neutrino energy of 10^{19} eV corresponds to a COM energy of about 10^{14} eV, still exceeding the $< 10^{13}$ eV energy typical of the LHC. The extremely low flux at high energies represents a major challenge to UHE astronomy; this is exacerbated for the neutrino by its small interaction cross-section. But the upside of high energy is high individual detectability: the signatures of UHE particle interactions are themselves highly energetic and generally propagate much farther through a detector medium than those of lower-energy interactions. Thus, while decreasing incident flux demands a proportionally increasing volume of the detection medium, actual detectors can be deployed more sparsely over that volume.

The ANITA campaign carries this trade-off to its logical extreme by deploying a single instrument, a balloon-borne phased radio antenna array, over the $\sim 10^6\,\mathrm{km}^3$ Antarctic ice sheet in order to survey the highest-energy regime of cosmic neutrinos. A high-energy particle interacting in dense matter can induce a compact electromagnetic shower, leading to coherent, impulsive radio emission known as the Askaryan effect[13]. ANITA seeks to detect these rare signals and discern them from natural thermal backgrounds as well as from intrusive anthropogenic continuous-wave (CW) and pulse backgrounds. The ANITA collaboration to date has made four flights beginning in 2006, 2008, 2014 and 2016. This work is a survey for radio signatures of ultra-high-energy cosmic neutrinos through data analysis of the third flight of the ANITA campaign, and is described in the following chapters. First, however, we present some historical and scientific context.

1.4 Historical Background

The history of neutrino astronomy spans half a century. The first neutrinos were detected from nuclear reactors in 1956 by Reines and Cowan. They used a pair of water tanks, one of them doped with cadmium, to identify electron antineutrinos by detecting indirectly the inverse beta decay of hydrogen:

$$\bar{\nu}_e + p \rightarrow n + e^+. \tag{1.3}$$

A scintillator detected photons secondary to the β^{-1} decay: e^+e^- annihilation produced γ-ray pairs, while the cadmium atoms absorbed neutrons and γ-decayed[14]. Experimental work in neutrino astronomy began a decade later after John N. Bahcall calculated theoretical ν_e flux from fusion interactions in the Sun[15]. Raymond Davis, Jr. developed an experiment to detect this flux: in an underground laboratory at the Homestake mine site in South Dakota, a tank was constructed and filled with chlorine-containing fluid. Since the earth is opaque to cosmic rays but essentially transparent to ~MeV solar neutrinos, the underground site shielded the experiment from the CR background while admitting the neutrinos. Chlorine nuclei absorbed neutrinos and β-decayed, and the resulting argon atoms were periodically separated from the fluid and counted. The Homestake experiment, for which Davis shared the 2002 Nobel Prize in physics, indeed detected solar neutrinos[1], but measured only about a third of the flux predicted by Bahcall. This discrepancy was later attributed to neutrino flavor oscillations occurring en route from the Sun, as the Homestake experiment was sensitive only to the ν_e. The first detection of neutrinos from beyond the Solar system occurred in 1987 when Kamiokande-II (a water Cherenkov experiment in Japan) and other instruments detected a ~10 MeV neutrino flux from SN1987A, a supernova in the Large Magellanic Cloud. As predicted by long-held theory[16], the neutrino burst arrived a few hours before the visible light[17].

Later instruments began exploring the higher energy regimes. IceCube and its predecessor AMANDA deployed strings of photomultiplier tubes deep in the ice at the South Pole to detect the bluish Cherenkov light emitted by muons and shower products from ~PeV neutrinos interacting in the ice. This instrumented a large (~1 km^3 in the case of IceCube) detector volume. The ANITA campaign then extended this concept to neutrino energies of above 10^{18} eV by surveying ~10^6 km^3 of the Antarctic ice sheet from a high-altitude balloon and exploiting the coherent propagation of longer-wavelength Cherenkov radiation, i.e., the Askaryan effect, in the radio range.

1.5 Neutrino Detection by Cherenkov Radiation

The neutral charge that renders the neutrino an effective cosmological messenger also makes it very difficult to detect, but several techniques have nonetheless been developed. Some detectors employ utilize water as a detection medium; incident neutrinos induce inverse beta decays followed by e^-e^+ annihilations, and a photodetector such as a photomultiplier tube (PMT) detects photons from the annihilations. Another variation of this also exploits inverse beta decay, but detection is accomplished by identifying nuclei that have undergone secondary species changes or excitations, by counting them or detecting their subsequent decays, respectively. In the high-energy regime, however, the dominant neutrino detection methods utilize *Cherenkov radiation*, a phenomenon that occurs when a particle exceeds the phase velocity of light in a material.

1.5.1 Cherenkov Radiation

While the speed of light in a vacuum, $c \approx 3 \times 10^8\,\mathrm{m/s}$, is a constant of nature, light propagates through matter at $c_{medium} = c/n$, where $n > 1$ is the material's index of refraction. Thus, a high-velocity particle in matter may find itself accelerated beyond the phase velocity of light in that medium, in which case Cherenkov radiation occurs. The energy density per frequency ω per length x emitted by the particle is given by the Frank-Tamm formula [4]

$$\frac{dE}{dx d\omega} = \frac{q^2}{4\pi}\mu(\omega)\omega\left[1 - \frac{1}{\beta^2 n^2(\omega)}\right], \tag{1.4}$$

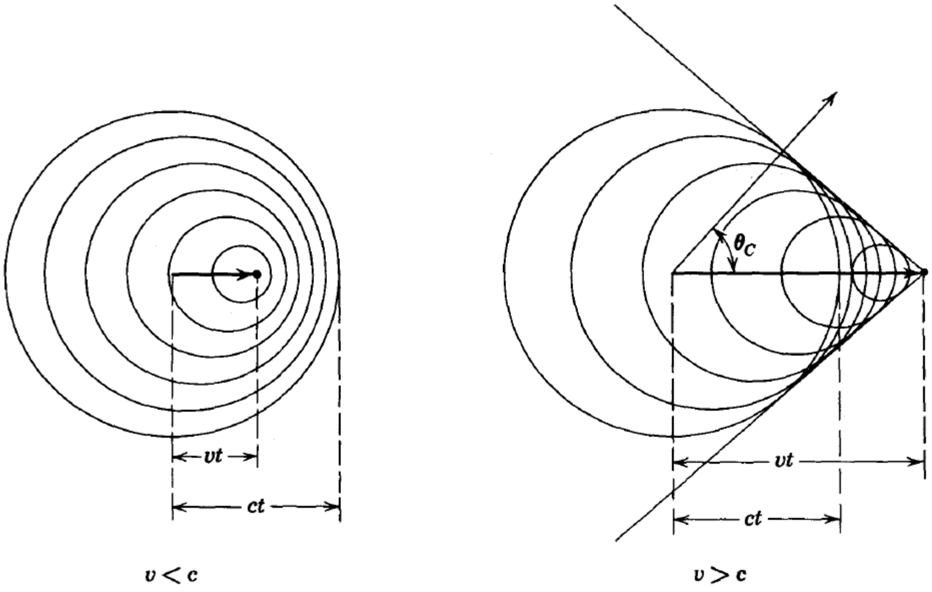

Figure 1.5: Geometry of Cherenkov radiation. For a particle traveling at $v < c_{medium}$ (left), wavefronts advance ahead of the particle, but for $v > c_{medium}$ (right), the particle outruns its own light, resulting in a shock-like front propagating along angle θ_c. Diagram from [4].

where $\beta = v/c$, and $\mu(\omega)$ and $n(\omega)$ are respectively the material's frequency-dependent permeability and index of refraction. In a process analogous to a sonic boom, the particle outruns its own radiation (Figure 1.5). The radiated power is concentrated in a shock front near the surface of a cone coaxial to the particle's path, with the cone angle given by

$$\cos\theta_c = \frac{1}{n\beta}. \tag{1.5}$$

In deep Antarctic ice, $n(\omega) \approx 1.6$ in the radio regime, corresponding to a Cherenkov angle of about 53°. The radiation is polarized normal to the cone surface[4]; this is useful during analysis to identify neutrino candidates.

In a typical neutrino experiment, an ultrarelativistic neutrino enters a dense medium such as liquid water, ice[18], or salt[19], and interacts with a nucleon leading to an electromagnetic shower, the constituents of which will emit Cherenkov radiation across a wide spectrum. Experiments such as IceCube and Super-K detect this radiation in and around the visible range, while ANITA and similar experiments survey a section of the RF spectrum.

1.5.2 The Askaryan Effect

Working in the radio spectrum enables a special case of Cherenkov radiation predicted by G. Askaryan[13] in 1962 and hence known as the *Askaryan effect.* Askaryan asserted that for wavelengths longer than the shower dimensions, the shower behaves as a point source, and constructive interference should result in a coherent pulse of Cherenkov radiation. Furthermore, a shower of charged particles in a dense medium such as a solid will be comparatively compact, on the order of a few centimeters, allowing for coherence at a much shorter wavelength than in air, where shower dimensions can reach hundreds of meters. A shower size of 10 cm, for instance, allows coherence at frequencies up to 1 GHz or more.

Although equation (1.4) is nearly linear in ω and thus predicts comparatively low energy output at low frequencies, detectability is enhanced by the concentration of energy on the shock front as well as by the quadratic dependence of coherent radiation energy density on the electric field. In ice, specifically, the electron shower is expected to have a characteristic size, or *Moliere radius* R_M, of about 10 cm[5]; thus coherence is expected at radio frequencies below about 1 GHz. This inspired the concept of the ANITA campaign, and was observed by the ANITA collaboration in 2007[18]. See Figure 1.6.

---

[5] R_M is the dimension of the shower itself, *not* the longitudinal propagation distance. The latter

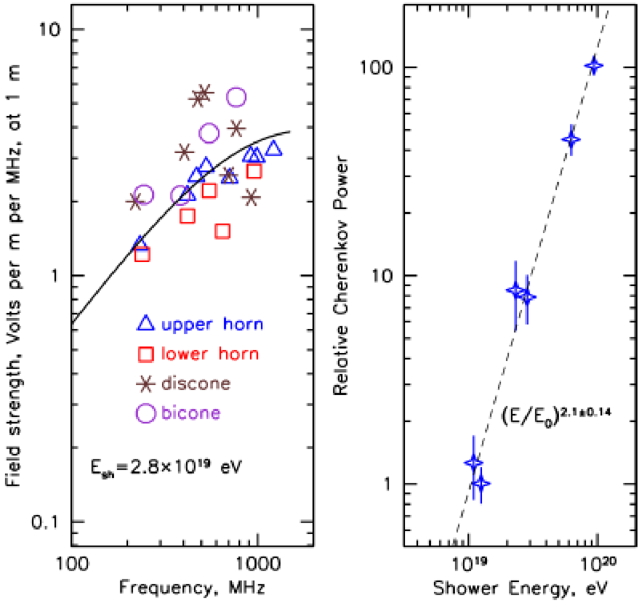

Figure 1.6: Cherenkov/Askaryan radiation in ice. The left plot shows the electric field strength vs. frequency for various antennas (the solid line is the theoretical expectation). On the right, the quadratic dependence of radiated power on shower energy reveals coherence. From [18].

1.5.3 Cherenkov Neutrino Experiments

A number of instruments have implemented the Cherenkov neutrino detection principle. Super-Kamiokande (Super-K), a water Cherenkov detector in Japan, searches for >MeV solar, atmospheric and reactor neutrinos. The detector medium is 50,000 metric tons of highly purified water, contained in an underground tank to shield from cosmic rays. An array of about 11,000 photomultiplier tubes (PMT's), deployed about the the tank perimeter, detects visible Cherenkov light from neutrino interactions[20]. Super-K confirmed atmospheric neutrino oscillations by comparing the flux of downgoing (from above) ν_μ to that of of upgoing (through-the-earth) atmospheric ν_μ; the longer flight path of upgoing ν_μ permits loss through oscillation to ν_e.

Cherenkov light detection by PMT's is also utilized in the IceCube Neutrino Observatory at the South pole; IceCube, however, explores a higher-energy (~PeV) astrophysical neutrino regime by using a larger detector volume and sparser PMT placement than Super-K. Over eighty vertical "strings", each equipped with sixty PMT assemblies, were lowered into the ice through boreholes to a depth of over 1km, spanning about a cubic kilometer of solid ice. Cherenkov light from neutrino interactions activates the PMT's in characteristic time and space distributions that allow identification of the induced electromagnetic shower, the daughter muon, or both. IceCube discovered the first high-energy neutrino flux in 2013 after analyzing three years of data [21]. Neutrinos of energy over 1PeV were observed.

Neutrino detection can be extended to yet higher energy regimes, employing even larger and sparser detectors. Since ice attenuates visible light over a few tens of meters, the preferred regime for UHE neutrino detection is radio, where attenuation

---

is $\sim$ 10m and, since $v_{shower} \approx c$, does not influence the regime of coherence lengths exceed 100 m and Askaryan coherence enhances signal power. The partially-constructed Askaryan Radio Array at the South pole will search for EeV neutrinos by surveying $\sim$200 km^3 of ice with an array of 37 antennas buried about 200 m beneath the surface, on a triangular grid at 2 km spacing. The instrument exploring the very highest neutrino energy regime to date is the ANITA, probing neutrino energies from 10^{18} eV to over 10^{20} eV [22].

1.6 The ANITA Campaign

ANITA uses the Antarctic continental ice sheet as its detector medium. By deploying a single detector on a balloon in the stratosphere at $>$35 km, it can at any one time survey a $\sim$600 km radius of ice surface. With an ice depth of more than 1 km, this represents a detector volume of up to $\sim 10^6$ km^3. A circular high-altitude wind pattern known as the Antarctic summer polar vortex enhances mission feasibility by carrying ANITA away from its RF-noisy launch point to quieter places, while simultaneously confining it to the Antarctic continent. The ANITA payload is a phased array of antennas mounted on a gondola with supporting power, control and data collection systems. Multiple antennas enable signal source localization by interferometry, and allow the use of coincidence algorithms to separate random thermal noise. Flights begin around the start of the austral summer in December and last from 20 to 60 days depending on wind conditions and instrument performance.

During the flight, data is collected and saved to on-board storage devices. Satellite telemetry enables ANITA collaborators to monitor instrument performance and manipulate instrument operating parameters during flight. At the end of the mission, the payload is jettisoned from the balloon and descends to earth by parachute. After recovery of the payload data storage, the raw data is distributed to collaborators who rewrite it into a user-friendly formats and make pre-analysis refinements such as time and temperature calibrations.

Most of ANITA's instrumented volume is near the horizon, so ANITA is most sensitive to down-going neutrinos at near-horizontal zenith angles and should therefore usually detect radiation from near the top of the Cherenkov cone, where the polarization is approximately vertical. This helps distinguish neutrino candidates from pulses induced by cosmic rays. The latter are due to geosynchrotron radiation; the roughly vertical orientation of the earth's magnetic field near the poles ensures that horizontal polarization dominates CR signals.

Analysis of ANITA data can be divided into two major components. First, events must be localized; this is done by interferometric methods described in Chapter 3. Events that originate from the continent are then subjected to analysis cuts to exclude non-impulsive events as well as anthropogenic background. After analysis cuts, surviving events are considered as neutrino candidates and used to establish an upper limit on cosmic neutrino flux.

* * * * * * *

In this work, we describe an analysis process for identifying neutrino candidate events from data of the ANITA-III flight. Chapter 2 describes the ANITA campaign in general, and the ANITA-III payload and flight specifically. My analysis methods are described in detail in Chapter 3. The development and optimization of our newly-developed circular polarization cuts and our enhanced linear discriminant cut merit detailed discussion, contained in Chapter 4. Chapter 5 contains a discussion of the actual analysis and its results. Conclusions are closing remarks are contained in Chapter 6.

The majority of computing pursuant to this work was performed using the resources of the Ohio Supercomputer Center (OSC)[23].

Chapter 2

The ANITA-III Payload and Flight

2.1 The ANITA instrument campaign

The ANITA campaign is a series of flights of a high-altitude balloon-borne phased antenna array, searching for the Askaryan radio signatures of ultra-high-energy neutrino interactions in the Antarctic ice sheet. The balloon missions are launched from the NASA Long Duration Balloon (LDB) facility near McMurdo Station, a United States base in West Antarctica on the coast of the Ross Sea. Launches occur around the beginning of the Austral summer in December, after the Antarctic summer polar vortex has formed up for the season. Flying at an altitude of about 35 km, the intent is to ride the vortex, circling the continent for up to 60 days. With a horizon distance of over 600 km, ANITA has a view of approximately $10^6\,\mathrm{km}^3$ of the Antarctic surface at any given time. At the end of the mission, the payload is jettisoned from the balloon and descends to earth by parachute. Experimental data is archived on-board during the flight, so recovery is critical.

Each ANITA instrument was comprised of three coaxial rings of radio antennas, with hardware for amplification, filtering, triggering, data storage, flight operations control, and telemetry. The ANITA-III payload is shown in Figure 2.1.

When RF radiation reaches an ANITA antenna, the resulting electrical signal is immediately filtered and amplified, then transmitted by coaxial cable to an instrument box on the flight deck. There the signal is filtered and amplified again, then split and fed to two components: an analog triggering circuit, which monitors incoming power levels; and a digitization circuit, which converts the analog signal into discrete time-stamped voltage samples. Since it is expected that a neutrino signal will trigger multiple antennas, the triggering circuitry and flight computer impose a timing coincidence algorithm to determine when a signal should be retained for storage. Signals satisfying the coincidence algorithm are formatted into *events* and forwarded to redundant on-board data storage systems. The flight computer also prioritizes the events for telemetry.

2.2 Instrument Overview

Figure 2.1 shows the ANITA-III payload. The forty-eight Seavey antennas are arranged in three coaxial circular arrays of 16 antennas each. The conical profile is necessary to accommodate a crane boom during transport and launch; the upper ring of sixteen antennas is therefore split into two mini-rings of eight. Since neutrino signals are expected from the ice, not the sky, each antenna is angled downward about 10°. Square black solar panels at the top and bottom of the payload provide electrical power for the instrument. The brains of the ANITA signal chain are contained in the instrument box, shown at the left side of the flight deck in Figure 2.1. The flight computer and supporting instrumentation for signal capture, triggering and data storage are contained there, and are described next.

2.2.1 Antennas and Signal Chain

The ANITA signal chain begins at the antennas. Since impulsive signals are broadband, and the characteristic 10 cm Moliere radius of the Askaryan shower imposes a coherence cutoff at ∼1 GHz, the antennas and signal chain hardware are designed for sensitivity across the 200-1200 MHz frequency band. The antennas (48 in the case of ANITA-III) are arranged outward-facing in vertically-displaced rings (Figure 2.1). They are manufactured to purpose by Seavey, Inc. in a quad-ridge horn configuration for directional sensitivity to horizontal and vertical polarizations (HPol and VPol), with a 6 dB falloff at about 45 deg from boresight. Figure 2.2 shows a typical antenna. The antenna gain response (Figure 2.3) is relatively flat across the band of interest, and gain variation among the antennas is less than 2 dB in most of the band. The angular beam response of a single antenna is shown in Figure 2.4. A pair of low noise filter-amplifier assemblies at each antenna provides ∼35 dB power gain to the two polarization channels. Band-pass filtering is applied at each amplification stage to eliminate the extraneous power of accumulated out-of-band noise and to suppress aliasing in subsequent digital signal processing. Coaxial cable transmits the signals to the instrument box, where they enter the Internal Radio Frequency Control Modules (iRFCMs). Each of the four iRFCM units services 24 input channels. Here the signals are amplified another 40 dB, and attenuators custom-fitted to each channel compen-

Figure 2.1: The ANITA-III payload during testing, Columbia Scientific Balloon Facility, Palestine, Texas, USA, August 2014. Photo by the author.

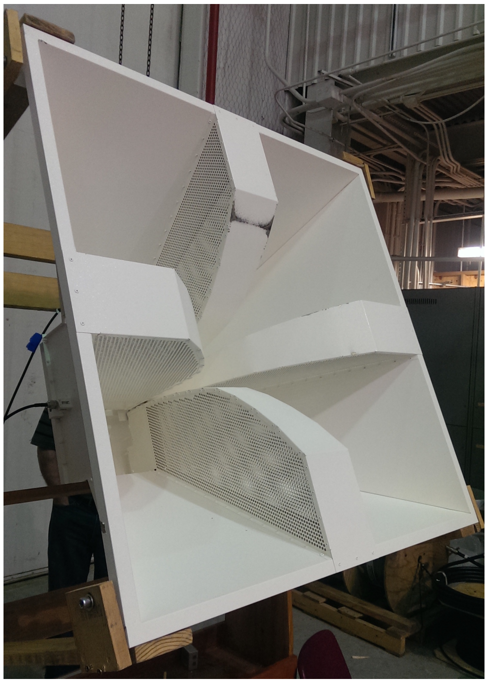

Figure 2.2: One of the 48 Seavey antennas used in ANITA-III. The quad-ridge horn design facilitates detection in horizontal and vertical polarization.

sate for gain variations between channels. The iRFCMs also contain bias tees which enable the signal cabling to transmit DC power to the amplifiers at the antennas. We assembled the iRFCMs and tested their frequency response during instrument integration at Palestine, Texas in the summer of 2014. Figure 2.5 shows the components of an iRFCM, and figure 2.6 shows a typical power spectrum.

Also contained in the instrument box is a Compact Peripheral Component Interface (cPCI) backplane interconnecting the flight computer and three electronic components: an array of SURF (Sampling Unit for Radio Frequencies) boards digitize the incoming signals; the SHORTs (SURF High-Occupancy RF Trigger) monitor power on each channel; and the TURF (Triggering Unit for Radio Frequencies) manages the higher-level logic of instrument triggering.

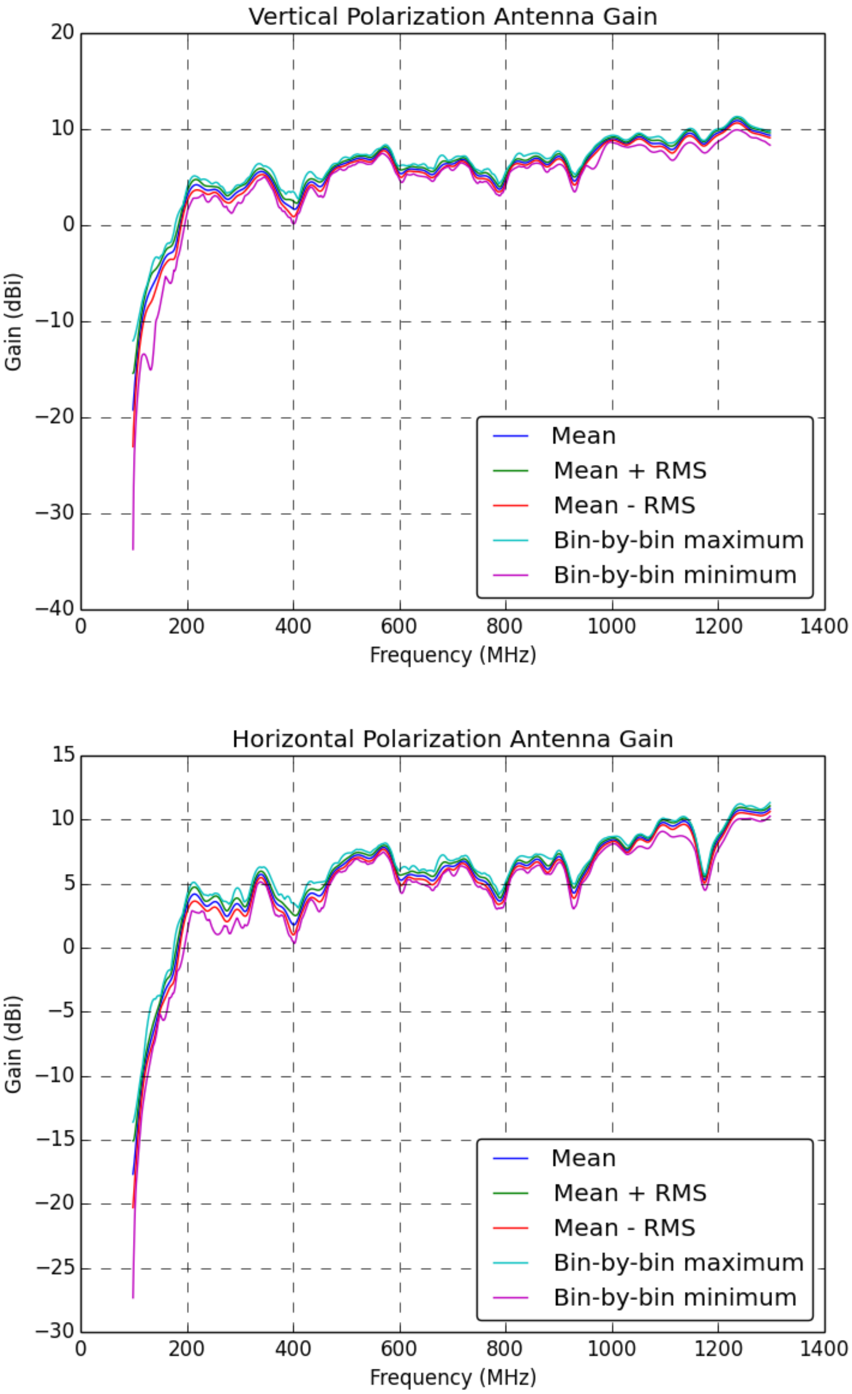

Figure 2.3: On-axis S12 gain response of the Seavey antennas, with calibration pulser frequency content and cabling response deconvolved. From [24]

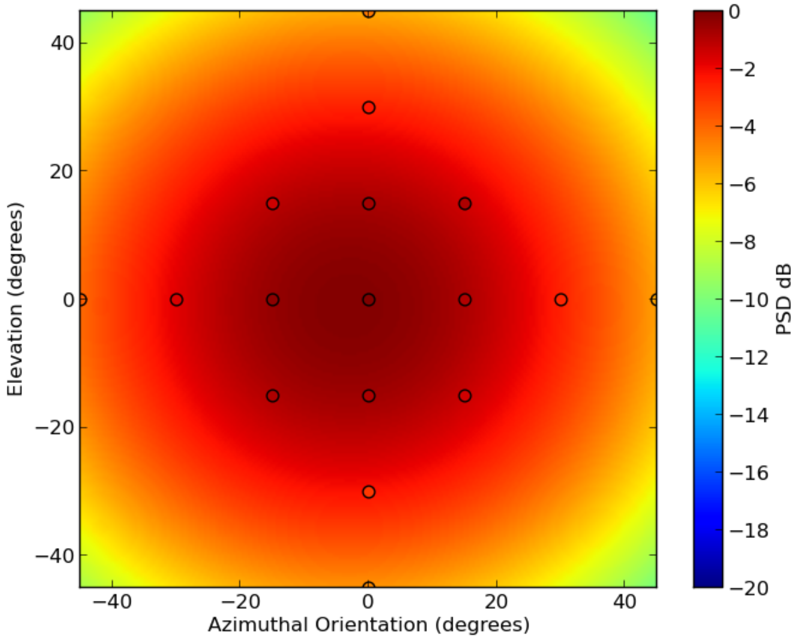

Figure 2.4: Typical hPol angular beam response pattern of Seavey antenna, 200-400MHz. Note the 6 dB fall-off at 45°; the pattern is similar across all in-band frequencies and polarizations. From [24]

After exiting the iRFCM, each signal is split two-fold. One branch goes to the SHORTS for analog processing to drive channel-level triggering, the other to the SURFs for digital sampling. These processes are described next.

2.2.2 Digital Sampling

For analysis, we need data in the form of lists of time-stamped voltages from each antenna/polarization. Each of the twelve SURF boards does this for eight channels. The actual analog-to-digital conversion is done by the Large Analogue Bandwidth Recorder And Digitizer with Ordered Readout (LABRADOR) integrated circuit (third generation), commonly called LAB3[26]. It is a switched capacitor array circuit that stores a 100 ns history consisting of 260 9-bit samples made at a rate of

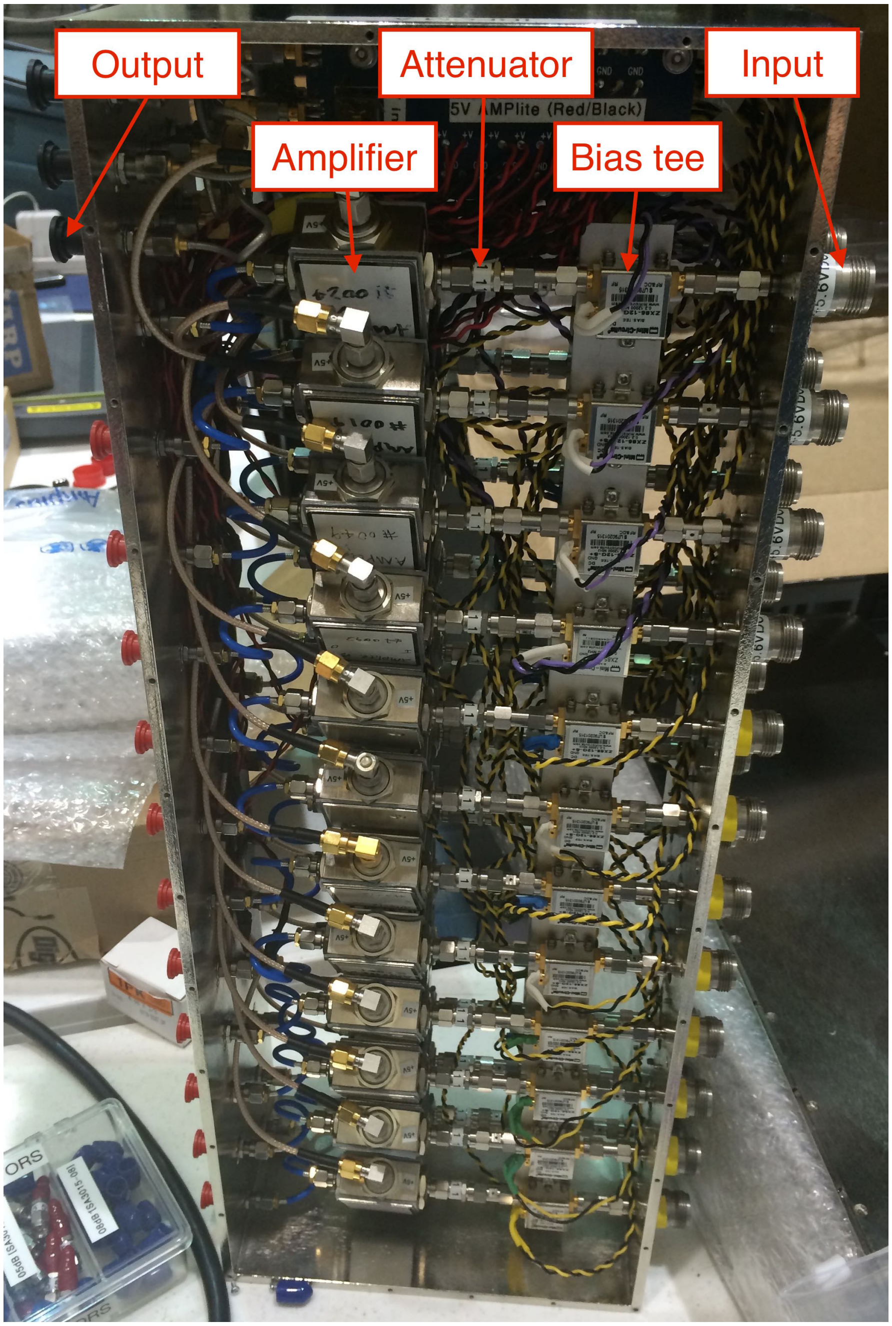

Figure 2.5: An Internal Radio Frequency Control Module (iRFCM) during assembly. From [25]

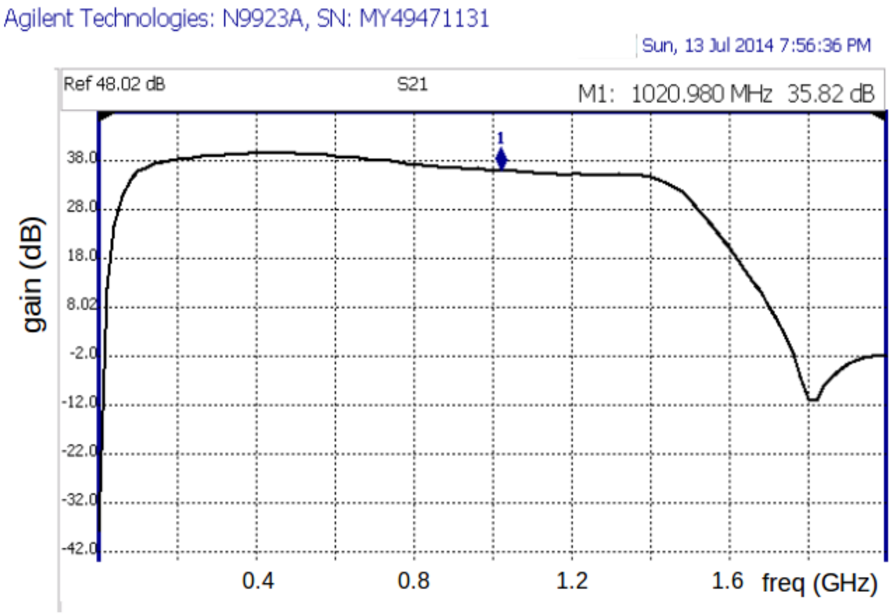

Figure 2.6: Frequency power response spectrum of a typical IRFCM channel, measured with a network analyzer. Note the flat response from 200 MHz to 1.4 GHz.

about 2.6 GHz, which can be read out on demand. Since the LAB3 has a recovery time of about 50 μs after readout, four LAB3 chips are used on each SURF to enable event buffering and increase livetime; each SURF therefore has an on-board clock which is used to calibrate timing between its four LAB3s.

2.2.3 Triggering

The trigger decides when to store data received by the instrument. With 48 antennas, each sampling two polarizations at about 2.6 GHz, a 50-day trip would yield about

$$(2.6\,\mathrm{GHz})(48)(2)(50\,\mathrm{days})(86,400\,\mathrm{s/day}) \approx 10^{18}\,\mathrm{bits} \approx 10^{5}\,\mathrm{TB}$$

of data. In reality, hardware limitations including power and bandwidth constrain the data processing rate of the instrument, and the recovery time of the LAB chips limits ANITA-III's maximum event processing rate to less than 50 Hz. We must therefore decide carefully when to read the LAB3s; this is the job of the triggering system.

On ANITA-III, a three-level triggering system was used. The signal content of each antenna channel[1] is continuously monitored by the SHORT, a tunnel diode circuit which functions as a square-law voltage meter, i.e., a power meter, integrated over a ∼10ns interval. The output of the SHORT is applied to a discriminator on the SURF; if it exceeds the discriminator threshold (typically $\sim 2.3\sigma$ above the expected thermal noise power) a level-one (L1) trigger is raised for the channel.

Fluctuations in thermal noise from the ice and from the instrument itself will induce frequent L1-triggers, but thermal noise is not correlated across antennas. Since a real signal from the ice should trigger multiple antennas, the second- and third-level (L2 and L3) triggers are based on coincidence of L1 triggers across multiple antennas within a specified time interval. An L1 trigger at any channel starts a timer on the SURF. The occurrence of another L1 trigger from an antenna in the same phi-sector within this time window will raise an L2 trigger. The length of the time window depends on the ring (top, middle or bottom) of the antenna that opened the window: the time window opened by a bottom-ring L1 trigger is longer (16 ns as opposed to 4 ns for the top ring), such that upward pointing events are favored for L2-triggering. L2 triggers are forwarded to the TURF board; here, a Field-Programmable Gate Array (FPGA) monitors for coincidence between L2 triggers in adjacent phi-sectors. If such a coincidence occurs within an 8 ns time window, then an L3 trigger is raised, and the content of the LAB3's is read out and forwarded for storage.

---

[1] Horizontal and vertical polarizations are always treated separately.

Triggers originating by the above process are referred to as RF triggers. The vast majority of these are false positives induced by thermal noise fluctuations or anthropogenic (CW and impulsive) sources. In order to discern signal from background it is helpful to have a sample of the RF environment independent of signal chain triggering. To obtain such *minimum-bias* events, the ANITA instrument is triggered intentionally at regular intervals. One of the on-board GPS units, described later, sends pulses to the SURF at 1 Hz to induce minimum-bias triggers, and the flight computer does the same. The former events are termed "PPS1" triggers, the latter "Soft" triggers.

2.2.4 Data storage

Three data storage arrays were prepared for the ANITA-III payload. A pair of "Helium drives", HDST Ultrastar helium hard disk units, provided about 6 TB of 2-fold redundant storage. Another storage component, the RIFRAF module, was an array of 48 100-GB USB "thumb" drives in a purpose-built enclosure which provided power, drive selection, and heat sinking. We programmed the drive selection switching on the RIFRAF using a MicroChip PIC microprocessor. We also assembled and tested the RIFRAF module before the flight. Another data storage system was developed by collaborators at National Taiwan University (NTU). The RIFRAF was not deployed on the payload due to power and throughput concerns, and the NTU module failed during the flight. The Helium drives, however, functioned properly throughout the flight and were successfully recovered thereafter.

Data stored during the flight included event waveforms, GPS information, and housekeeping data for monitoring instrument health and performance. In order to limit file sizes and reduce the risk of data loss, the data was divided up into *runs*, each containing information from up to two hours of operation.

2.2.5 GPS

Payload position and orientation must be precisely known in order to localize event sources. Since no meaningful analysis can occur without this data, three GPS systems were deployed on ANITA-III. Two Ashtech ADU5 GPS antenna arrays provide UTC time-stamped payload position (latitude, longitude, altitude) and orientation (heading, pitch, roll). The third GPS unit, an Ashtech G12, delivers time-stamped position and velocity information. The GPS data is sent to the flight computer, which forwards it to storage at 1-second intervals.

2.2.6 Other Payload Components

Flight Computer

The flight computer, a GE Intelligent Platforms XCR14 designed for use in extreme conditions, plugs into the cPCI backplane for close interaction with the SURF and TURF components. The flight software, described in detail in [25], is responsible for triggering, collecting housekeeping data, acquiring and storing event data, managing power and control systems, prioritizing and sending telemetry, and processing operator commands. This is accomplished by a set of daemon programs (Figure 2.7) running on a Linux operating system.

Power System

The ANITA payload requires about 500 W for nominal operation, all of which is

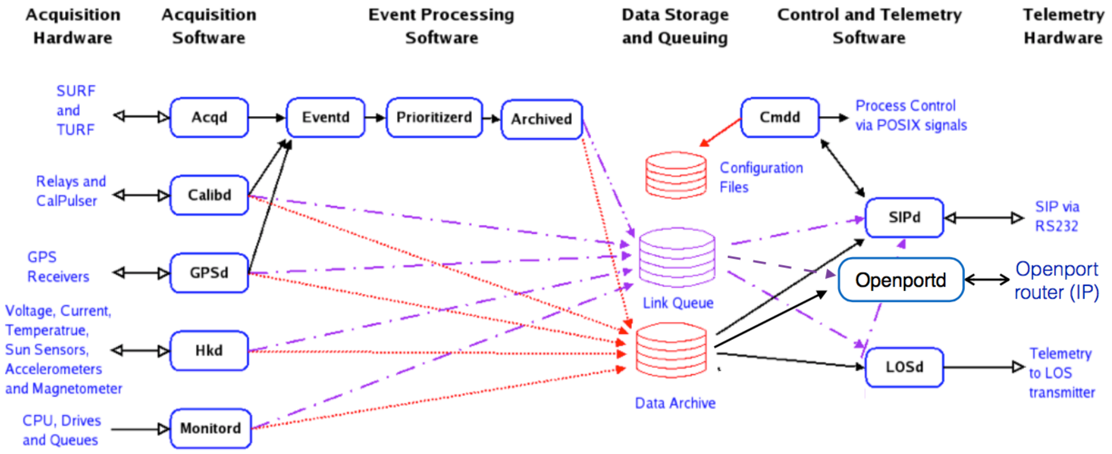

Figure 2.7: ANITA-III flight software modular structure and data flow.
Figure by R. Nichol

provided by the solar panels below the bottom antenna ring. ANITA flies in 24-hour daylight, but the ∼20° variation in Sun angle throughout the Antarctic day results in varying power input to the panels; a 12-cell, 2800 Wh-capacity battery array is used to ensure consistent power availability.

Housekeeping Sensors

Temperature sensors at various locations inside and outside the instrument box, as well as voltage and current sensors in the power systems and other components, allow collaborators to monitor the operating condition of the instrument. An array of sun sensors served as a potential backup to the GPS system for obtaining payload orientation.

Support Instrument Package

Aviation functions such as ballast release during ascent and payload jettison at the end of the mission, as well as various monitoring tasks, are the responsibility of NASA CSBF and are implemented in the Support Instrument Package (SIP), the enclosure shown in Figure 2.1 on the right side of the flight deck.

Anita Low Frequency Antenna

The serendipitous detection of cosmic rays by ANITA-I motivated the ANITA Low Frequency Antenna (ALFA). Sensitive to frequencies of 25 to 80 MHz, its purpose is to observe low-frequency emission from cosmic ray air showers.

2.2.7 Telemetry and Control

Communication with the payload is needed throughout the flight in order to monitor instrument health, manipulate instrument operating parameters, and observe a small sample of events. The flight computer compiled housekeeping information including triggering rates, temperatures at various points on the payload, and voltages and currents throughout the power systems and instrument circuitry. The GPS units delivered timestamped and payload position and orientation. Collaborators continually monitored the telemetry to assess instrument health during the flight.

Due to bandwidth constraints, only about 0.1% of triggered events can be telemetered; therefore an event prioritizer program was developed and deployed on the flight computer[25]. Using interferometry provided by a Graphics Processing Unit (GPU) program, events with impulsive characteristics are assigned higher telemetry priority than events resembling thermal or other backgrounds. Events are then placed in priority queues, from which the telemetry daemons select and send them.

Telemetry channels were used to control the instrument. Collaborators used a commanding system to manipulate parameters such as triggering thresholds and phi-sector masking, as well as to restart software or hardware components following failures.

Telemetry was transmitted by various channels depending on availability. Line-of-sight (LOS) telemetry provided the fastest data rate but was limited to when the payload was in line-of-sight (over 600 km at flight altitude) of McMurdo Station. The Iridium Openport network provided a 15-20 kbps data rate and was utilized throughout most of the flight. The Tracking and Data Relay Satellite System (TDRSS) provided a slower 6 kbps channel but was only available during limited time windows. A very slow Iridium service was also available, allowing a very limited subset of information to be sent when the other telemetry channels were unavailable.

2.2.8 Evolution of the ANITA Instrument

The design principles of the ANITA payloads were similar, albeit with a few important differences. ANITA-I had 32 antennas; for the ANITA-II and -III flights, the number was increased to 40 and 48, respectively, to improve coverage and sensitivity.

ANITA-I was sensitive to the circular polarizations. Since neutrino-induced pulses are expected to be linearly-polarized, the ANITA-I trigger required significant power in both circular polarizations. This allowed the ANITA-I trigger to be sensitive ultimately to linearly-polarized signals, but blind to the *direction* of the linear polarization. In ANITA-II, the trigger was changed to be sensitive only to vertical polarization because neutrino signals arriving at payload should be near the top of the Cherenkov cone, and hence predominantly vertically polarized. After the ANITA-II flight, however, collaborators analyzing ANITA-I data discovered horizontally-polarized radio impulses from cosmic-ray air showers emitting geosynchrotron radiation in the mostly-vertical polar geomagnetic field[27]. This serendipitous discovery brought an additional science motive to the campaign, but also augmented the need for discriminating between vertically-polarized neutrino signatures and horizontally-polarized CR ones. On ANITA-III, therefore, triggering was enabled in both horizontal and vertical polarizations.

In ANITA-I and -II, inbound signals from the antenna channels were each split into three frequency bands in order to identify broadband characteristics of impulsive signals and to help reject narrowband CW signals. Later experiments, however, determined that greater sensitivity could be achieved by triggering solely on integrated power over a ∼10 ns time interval, across the entire ANITA frequency band. ANITA-III, therefore, did not employ the frequency splitting strategy.

2.3 The ANITA-III flight

The ANITA-III mission launched from Williams Field at McMurdo Station, Antarctica on December 15, 2014. After completing about one-and-a-half circles around the continent, the payload drifted near the Antarctic seacoast, compelling termination of the mission on January 8, about 100 km from Davis Station (Figure 2.8).

2.3.1 In-flight monitoring and support

The performance of the ANITA-III payload was carefully monitored starting several hours before the launch and continuing throughout the duration of the flight. Collaborators worked seven-hour shifts, with a one-hour overlap between shifts to facilitate

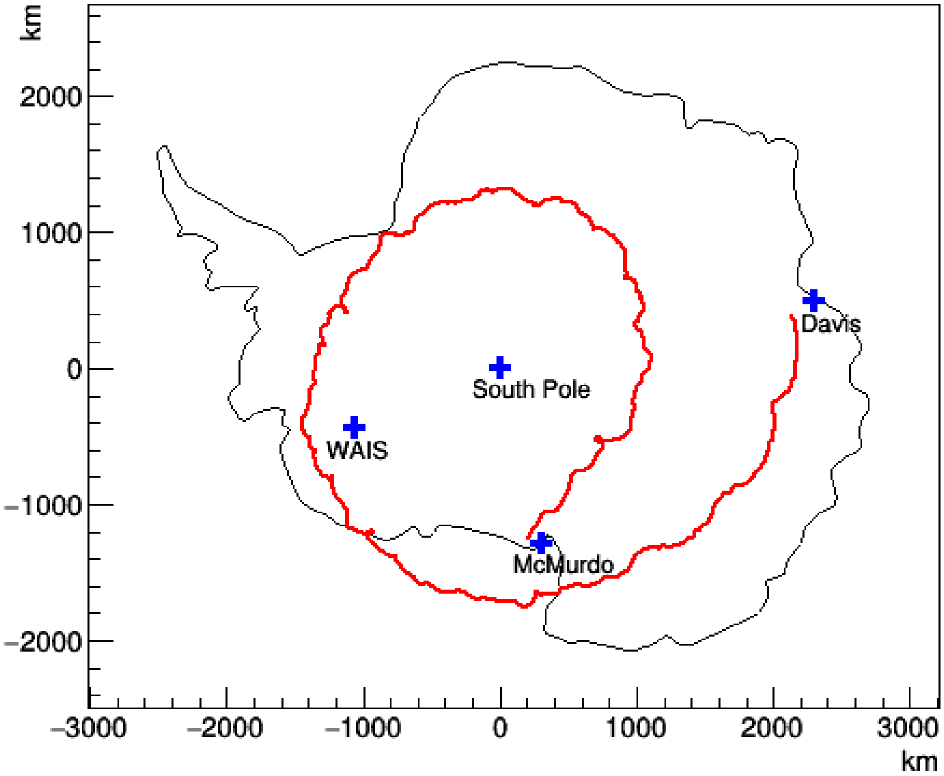

Figure 2.8: The Anita-III flight path.

communication of issues and priorities. During each shift, two persons ("shifters") were responsible for monitoring the payload, checking instrument parameters including event triggering rates, voltages and currents in various system components, and temperatures from sensors at multiple locations on the payload. The performance of the GPS units and telemetry channels were also carefully watched. Every hour, one of the assigned shifters filled out a log sheet, recording the relevant parameters and adding comments on particular issues or events that occurred during the shift. In addition to the shifters, one expert level person was assigned to each shift to provide second-level support and to direct any commands sent to the payload.

2.3.2 Limitations and Problems

The 22-day flight duration was a disappointment in light of the 60-day expectation, resulting in a >50% sensitivity loss. There were also problems with ϕ-sector masking, intended to suppress portions of the payload when subjected to excessive background. As a result, the event rate often spiked from its ideal value of about 30 Hz to over 80 Hz, leading to event buffer overflows at the SURFs, and in turn to frequent payload deadtime. Much of the shifters' attention was therefore focused on the monitoring the event rate and sending commands to the payload to manipulate ϕ-sector masking parameters.

In addition to instrument problems, ANITA-III was inundated with powerful CW background from geostationary satellites throughout the mission. As shown later in event sky maps of Figure 3.8, the analemmae of these satellites, distorted by changes in payload position during flight, are clearly visible. Furthermore, reflections from the satellites can be seen in a less distinct but similar pattern coming from the continent, just below the horizon. Likely satellite reflection events must be cut during analysis. This can be done based solely on event source location, albeit at the cost of a portion of the viewfield. But it also turns out that the satellite events are strong in circular polarization, and can be identified based on circular polarization parameters. This is discussed in more detail in Chapter 3.

2.3.3 Calibration pulsers

Calibration pulsers generate impulsive radio signals, dubbed *cal-pulses*, intended to trigger ANITA when the payload is within line-of-sight. Cal-pulses are used during flight to monitor instrument performance, and in post-flight analysis to assess pointing resolution and acceptance efficiency. For easy identification, the pulses are generated at regular time intervals, during a fixed time window of about 1 μs within each second. One calibration pulser was deployed at the Long Duration Ballon (LDB) facility at McMurdo and was in operation shortly before the launch, and later during the flight as the payload passed near McMurdo after completing one circuit around the continent[28]. Another pulser, located at the WAIS-divide research station in West Antarctica, operated during a two-day payload flyby, generating about 118,000 carefully-timed 6 kV pulses at one-second intervals[29]. Each pulser was activated when the payload approached line-of-sight distance. The RF triggers from these two pulsers were the major source for pointing calibration and acceptance efficiency in this work. Figure 2.9 shows interferometric localizations of WAIS pulses detected during the flyby. In Figure 2.10, the upper panel shows the payload distance from WAIS during the flyby. The lower left panel shows events from the flyby with nanosecond timestamps bewteen zero and 2000 ns. The timestamps span a 500 μs time window

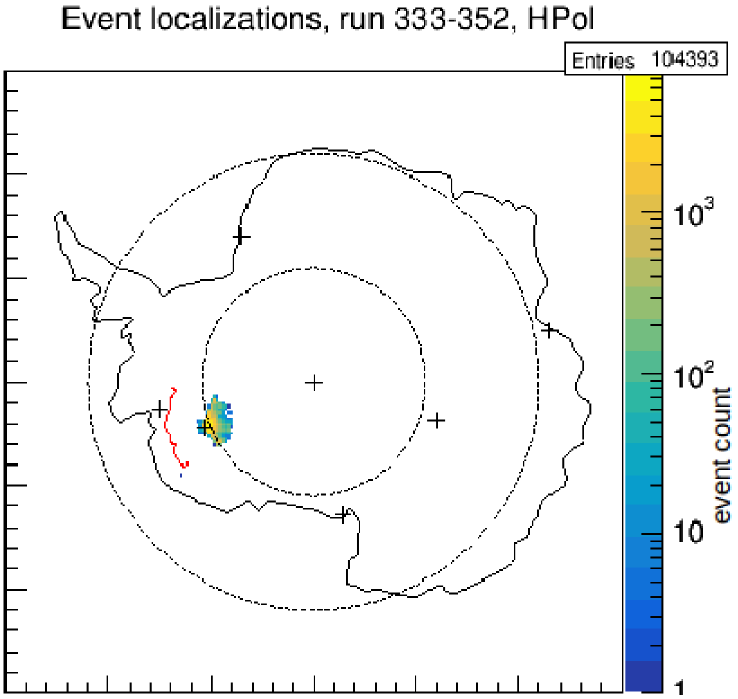

Figure 2.9: Localizations of WAIS pulser events on the Antarctic continent during the ANITA-III flight. The flight path during the flyby is shown in red.

due to the varying payload-WAIS distance. The lower right plot in Figure 2.10 shows the same event timestamps after adjustment for payload distance from WAIS. The distance-adjusted timestamps fall within a $2\,\mu$s time interval: the dearth of events outside the $2\,\mu$s window shows that non-WAIS events are not likely to fall within that window. WAIS pulses are thus easily distinguished by their nanosecond timestamps.

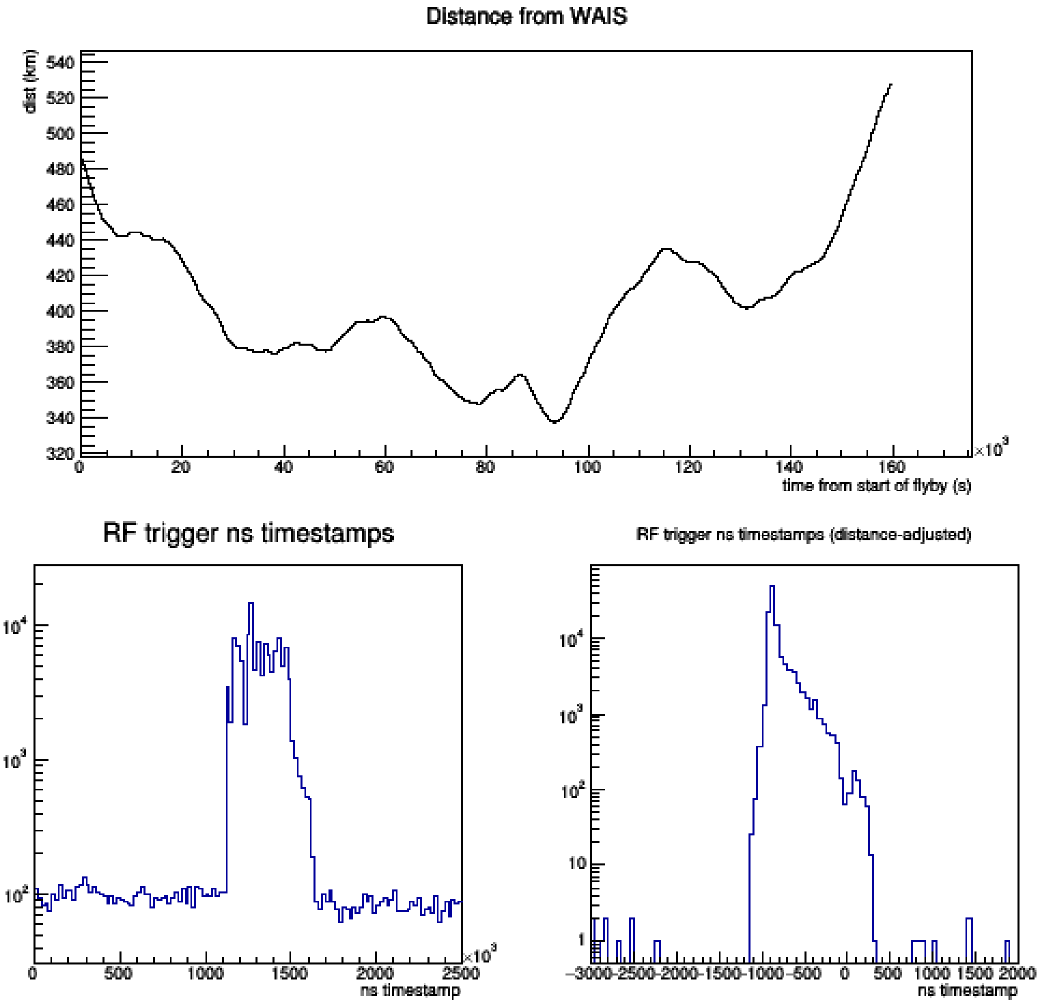

Figure 2.10: Top: payload distance from WAIS pulser during flyby. Bottom left: raw nanosecond timestamps of events during the WAIS flyby. Bottom right: distance-adjusted nanosecond timestamps of WAIS cal-pulses: all distance-adjusted pulses fall within a 2 μs time window.

2.3.4 Flight Termination and Payload Data Recovery

On January 8, 2015, the mission was terminated because the payload was drifting toward the seacoast. Drifting over the open ocean is not only prohibited under flight protocol, but also carries the risk that the payload could exit the polar vortex and be permanently lost. Since payload loss would mean complete loss of the on-board data storage and in turn substantial failure of the entire mission, the Collaboration decided to terminate the flight. The payload was jettisoned from the balloon, and descended to earth by parachute. The payload gondola was damaged in the fall, as expected, but the instrumentation remained largely intact. Fortunately, the flight terminated near Davis Base, an Australian outpost. Friendly relations between the U.S. and Australia facilitated cooperation, and the on-board data storage was recovered a couple of months later by a team deployed from Davis.

After recovery, the two Helium drives were shipped to University of Hawaii, Manoa, where they were successfully activated, and the data was distributed to team members for reformatting and calibration. The ANITA-III payload had recorded over 80 million events during the flight. A map of the Antarctic continent, showing a sampling of events that localized to the continent, is shown in Figure 2.11.

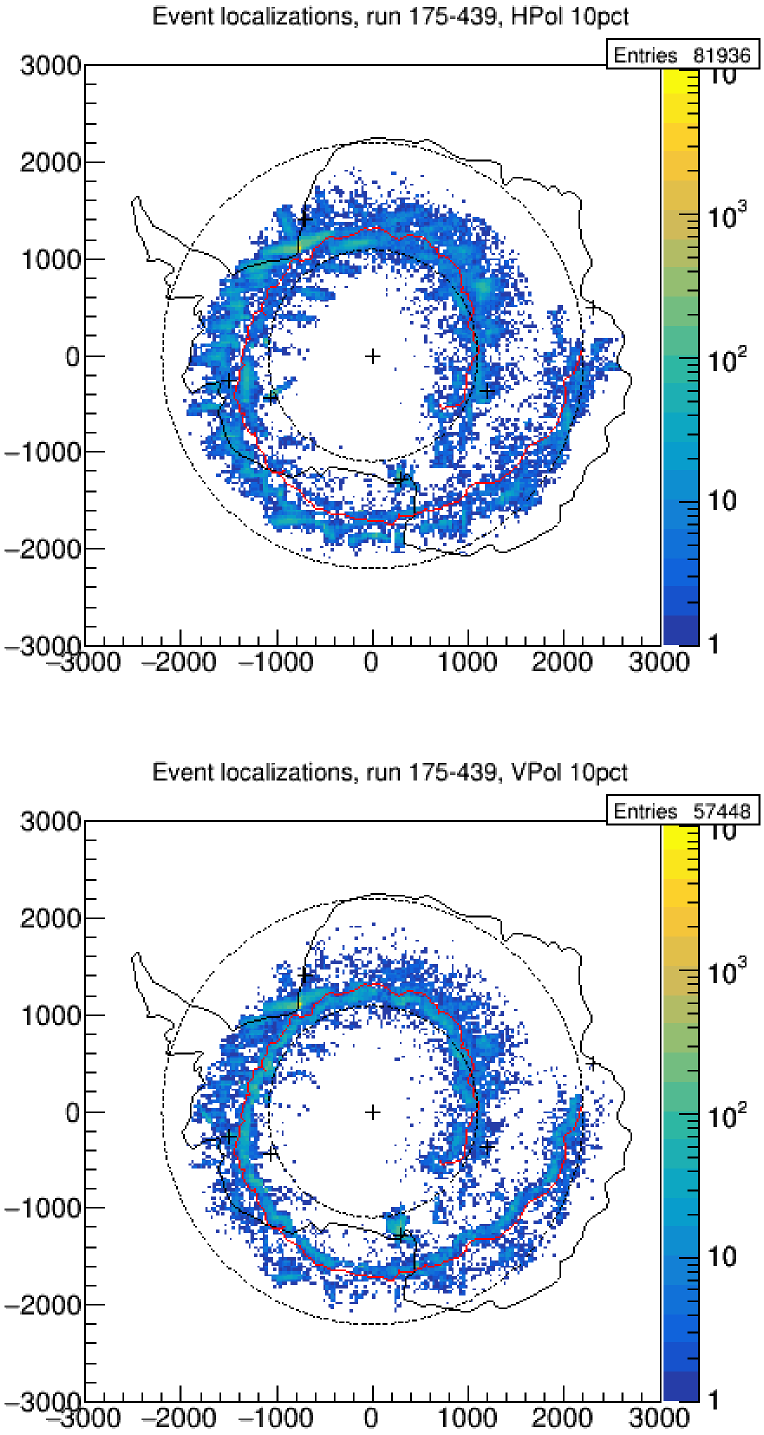

Figure 2.11: A sample of ANITA-III events that localized to the Antarctic continent.

Chapter 3

Analysis Methods

The purpose of this analysis is to derive a list of neutrino candidate events and/or a neutrino flux limit from the ANITA-III instrument data. ANITA data consists of event waveforms, gps data and payload housekeeping data. Shortly after the flight, collaborators created an analysis database from the original flight data. Time and temperature calibrations were applied to event waveforms, and gps and housekeeping data were organized [25]. Starting with this analysis database, we performed interferometry and event reconstruction on events using components of the `AnitaTools`, an analysis software package developed and used by the ANITA collaboration. We removed unusable events using quality cuts, and subjected the remaining reconstructed events to analysis cuts in order to eliminate non-impulsive events and anthropogenic triggers. Surviving events were examined as possible neutrino candidates, and statistics from the analysis were used to develop a process for establishing a flux limit.

This chapter describes in detail the methods used in analysis. We present first a description of the ANITA-III data format and organization.S The interferometry process used for event reconstruction is then discussed in detail, due to its critical role in analysis. This is followed by a discussion of the software employed in analysis. The analysis process itself is described, including the application of quality cuts and the motivations for choosing and optimizing the analysis cuts.

3.1 The ANITA-III Data

ANITA-III stored all of its data to on-board storage systems. After payload recovery, this data was retrieved and archived for analysis. Since the on-board data was optimized to throughput and capacity constraints and is not practical for analysis, it was reformatted post-flight into a 2.5 TB analysis database in the ROOT data analysis framework [30][1]. Data from instrument triggers was organized into two related ROOT trees[2]. A header tree (`headTree`) contains event number, trigger time, phi-sector masks and other single-valued data items from the event. A corresponding detail tree (`eventTree`) contains 108 uncalibrated time-domain voltage waveforms, corresponding to the 96 antenna/polarization channels and the twelve SURF clocks. Additional data trees were extracted containing payload position and orientation from the GPS systems, and data from on-board temperature, voltage and current sensors. Antenna phase center locations, required for interferometry, were approximated from pre-flight photogrammetry[31], then numerically optimized and embedded within `AnitaTools`.

Housekeeping and GPS data were delivered to the flight computer at various intervals, independently of triggering. Associating this data with events therefore requires

---

[1] ROOT is a prevailing analysis utility both in the Anita Collaboration and in the general particle physics community, and is particularly suited to managing large datasets.

[2] In ROOT, a *tree* (class `TTree`) is roughly analogous to a database table

care, especially with respect to GPS position and orientation which are critical to signal localization. Also, all three GPS units were not always working during the flight, although the ADU5A unit worked properly the vast majority of the time. The Collaboration created an event-specific database of payload position and orientation by interpolating event timestamps and calculating weighted averages of readings from the ADU5A and ADU5B units.

3.2 Interferometry Methods

Two methods of pulse reconstruction interferometry are described in Romero-Wolf et al.[32]: the correlation interferometer and the adding interferometer. In the correlation method, the delays between signal arrival times across multiple antenna pairings are used to fix a source direction. The adding interferometer seeks to maximize the total coherent power contained in a signal by varying an assumed source direction, calculating the relative arrival times at each antenna, and calculating the coherent sum of the input waveforms. While these two methods yield similar information, cross-correlation analysis is generally preferred due to its better handling of waveforms of varying gain and noise figure [32], and thus is employed in the `AnitaTools`.

Our interferometry exploits the fact that a directional plane-wave signal arrives at each antenna at a different time.[3] This arrival time delay Δt between any two antennas defines a circle on the 4π sr sky, as shown in Figure 3.1. The circle is centered on the direction of the baseline vector $\mathbf{b}$ connecting the two antennas; its radius, in the

---

[3] Typical source distance is over 10^4 m, compared to the $\sim$5 m dimensions of the ANITA payload; the plane-wave approximation of a spherical wavefront carries a path length error of $<10^{-3}$ m; the shortest wavelength in ANITA's regime is 0.3 m.

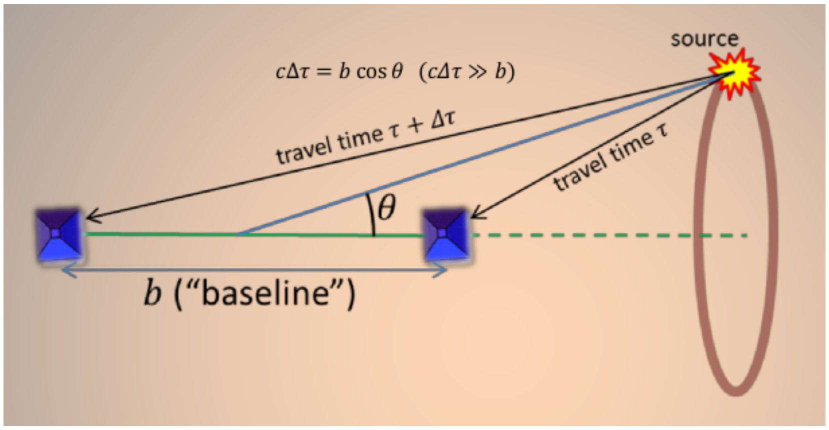

Figure 3.1: Circle on the sky defined by the time delay between two antennas. The difference in signal arrival times at two antennas defines the angle θ between the arrival direction of the signal and the baseline connecting the two antenna phase centers. This angle in turn constrains the signal's point of origin to a circle on the sky.

form of an angle θ_b from the baseline, satisfies

$$|\mathbf{b}|\cos\theta_b = c\Delta t, \tag{3.1}$$

where c is the speed of light. Due to the geometry of the ANITA payload, a strong impulsive signal will be detected by six to fifteen antennas pointed in the general direction of the signal. If n is the number of antennas detecting the signal, then we can generate $N = \frac{1}{2}n(n-1)$ pairs, measure the arrival delays of each pair, and calculate the location of the each of the N corresponding circles. Ideally, all of these circles will intersect at a unique point on the sky, thus localizing the signal's origin. In reality, however, the determination of the delay Δt between waveform arrivals is not trivial, and merits discussion.

3.2.1 Measuring delay by cross-correlation

Naïvely, one imagines simply measuring the time delay between maxima of the respective impulsive signals, but noise and bandwidth preclude this trivial means. Noise creates competing peaks, particularly at low SNR. Moreover, dispersion associated with finite bandwidth and nonflat in-band response temporally smears the impulsive signal, transforming it into a series of multiple peaks of similar magnitude, the "best" of which is not easily identified or even defined (Figure 3.2). To estimate the delay, we use the *cross-correlation function* (CCF) of two time-domain signals $x(t)$ and $y(t)$:

$$(x \star y)(\Delta t) = \frac{1}{M} \int_{-\infty}^{\infty} x(t) y(t + \Delta t) dt \tag{3.2a}$$

$$\text{with} \qquad M = \sqrt{\int_{-\infty}^{\infty} x^2 dt \int_{-\infty}^{\infty} y^2 dt} \tag{3.2b}$$

and Δt denoting the time delay between the two signals. Evidently, (3.2a) is essentially an inner product and thus a measure of similarity; the Δt of maximum correlation is an estimate of the time delay between signal arrivals at the two antennas.

Equations (3.2) are in fact an idealization in that they assume continuous sampling over infinite time. ANITA, however, samples data discretely over a finite time interval, and the exact time-spacing between individual samples varies over the duration of the waveform, and across antennas. But if the waveforms are interpolated onto an equally-spaced timestamp sequence, we can use a discrete reformulation of (3.2a) for the CCF of two discrete arrays $\mathbf{x}$ and $\mathbf{y}$ of size n, offset by $j \in \mathbb{Z}$:

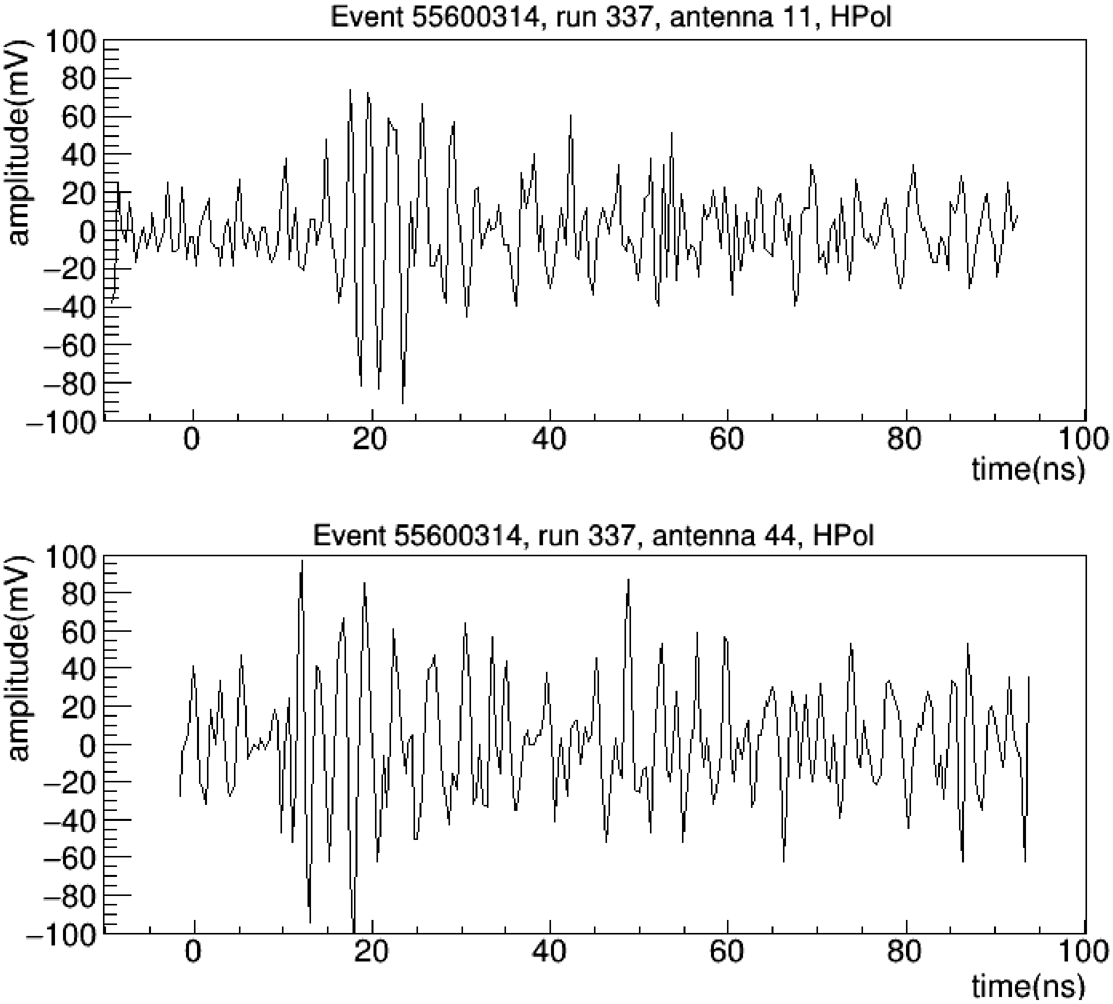

Figure 3.2: Waveforms from a calibration pulse as received at two ANITA-III antennas. The pulses are visible at about 20 ns (top) and 17 ns (bottom), but ringing and noise preclude identifying a definitive single peak in either waveform.

$$(\mathbf{x} \star \mathbf{y})_j = \frac{1}{M} \sum_{k=0}^{n-1} x_k y_{k+j} \qquad (1-n) \leq j \leq (n-1), \tag{3.3a}$$

$$\text{with} \qquad M = \sqrt{\left(\sum_{k=0}^{n-1} x_k^2\right)\left(\sum_{k=0}^{n-1} y_k^2\right)}. \tag{3.3b}$$

Note that in some terms of (3.3a), the index value $k+j$ falls outside the array bounds $[0, n-1]$. Depending on the application, terms containing these outlying indices may simply set to zero; or the index can be interpreted `mod`$[n]$, resulting in the *circular* or *wrap-around* correlation.

The discrete circular cross-correlation function is related to the discrete Fourier transform (DFT) $\mathcal{F}$ by

$$\mathcal{F}\{\mathbf{x} \star \mathbf{y}\} = \mathcal{F}^*\{\mathbf{x}\}\mathcal{F}\{\mathbf{y}\}, \tag{3.4}$$

an analog to the convolution theorem. Since direct computation of the sums in (3.3a) requires $O(n^2)$ operations, implementations of cross-correlation usually work in the Fourier domain where the $O(n \log n)$ Fast Fourier Transform (FFT) algorithm can be used. By its nature, the FFT-based calculation delivers the circular cross-correlation, consistent with the assumption that the input waveforms are infinitely repeating. In our case, this assumption is manifestly untrue, and it is a common practice to pad the input waveforms with zero-valued bins to suppress this effect. Zero-padding lengthens the effective repetition period of the signal but also creates unsupported regions in the waveforms, contrary to typical physical reality where noise is expected. As a result, a constant normalization such as (3.2b) favors small values of the offset τ, where the two waveforms' support regions overlap strongly. The effect of variable overlap can be mitigated with an offset-dependent normalization M_τ, similar to (3.2b) but summed only over the overlapping support regions of the displaced waveforms [33]. sOur implementation of this method in `AnitaTools`, with enhancements to reduce computation time, is described in Appendix A.

Interferometry requires summing the cross-correlation functions of multiple antenna pairs. The ANITA antennas are highly directional (figure 2.4), so pair selection is limited to antennas within three ϕ-sectors of one another. Antenna channels containing voltages over 1500 mV, which saturate the ADC's in the digitizing circuitry, were omitted from pair selection. Once the antenna pairs are selected, correlated and summed, the resulting interferometric map is used to hypothesize the direction from which the signal originated.

3.2.2 Event Pointing and Localization

Ideally, an event's source direction is indicated by the intersection of the circles of equation (3.1) corresponding to multiple antenna pairs. In reality, these sky circles are not so well-defined. Thermal noise renders as fuzziness of the circles; CW contamination and dispersion-induced ringing lead to multiple interference fringes, i.e. "ghost circles". Furthermore, uncertainty in antenna phase-center positions and instrument timing means that the circles generated by the antenna pairs, already noisy and ambiguous, will not intersect precisely. Our best estimate of the true intersection point is obtained by adding up all the correlations from individual antenna pairs, and choosing the point on the sky with the largest correlation total. For each antenna pair, every point on the sky corresponds to a particular Δt, and consequently a particular baseline angle ϕ as in (3.1). To make the correlation map, we pixelate the sky and add up the correlation values of every eligible antenna pair at each pixel.

Positions on the celestial sphere are represented as an azimuth angle ϕ from the payload fore-to-aft line, and the elevation angle θ with respect to the payload xy-plane. For payload GPS orientation information, the (heading, pitch, roll) representation customary to aviation is used. For every (θ_{pl}, ϕ_{pl}) pixel in payload-centered coordinates, we calculate the corresponding baseline angle θ_b relative to each baseline vector $\mathbf{b}$; then the correlation values corresponding to the θ_b are calculated and accumulated. An impulsive signal should render as a narrow and pronounced peak in the correlation sky map. In practice, a coarse-grained correlation map is calculated first to save computation time. The correlation peaks in the coarse map are then used to make zoomed maps at finer resolution, from which the final interferometry peaks are obtained. The highest peak correlation peak is taken as the event localization on the sky; this is then used to reconstruct the original signal and, if possible, localize it to a point on the continent. The lesser peaks are used in analysis cuts.

3.2.3 Event Reconstruction and Error Bounds

After an event is localized on the sky, a reconstruction is attempted. Since the event source direction defines arrival delays between the antennas, the delays corresponding to the event source direction are applied to each event waveform, and the waveforms are added up. This *coherent reconstruction* is a hypothetical representation of the original signal, and is used to derive various waveform features, such as SNR and polarization, for use in analysis cuts. Figure 3.3 shows the correlation sky maps and the coherent sum reconstruction for a signal from the calibration pulser at WAIS Divide.

Using the event direction fixed on the payload-coordinate sky, the event source po-

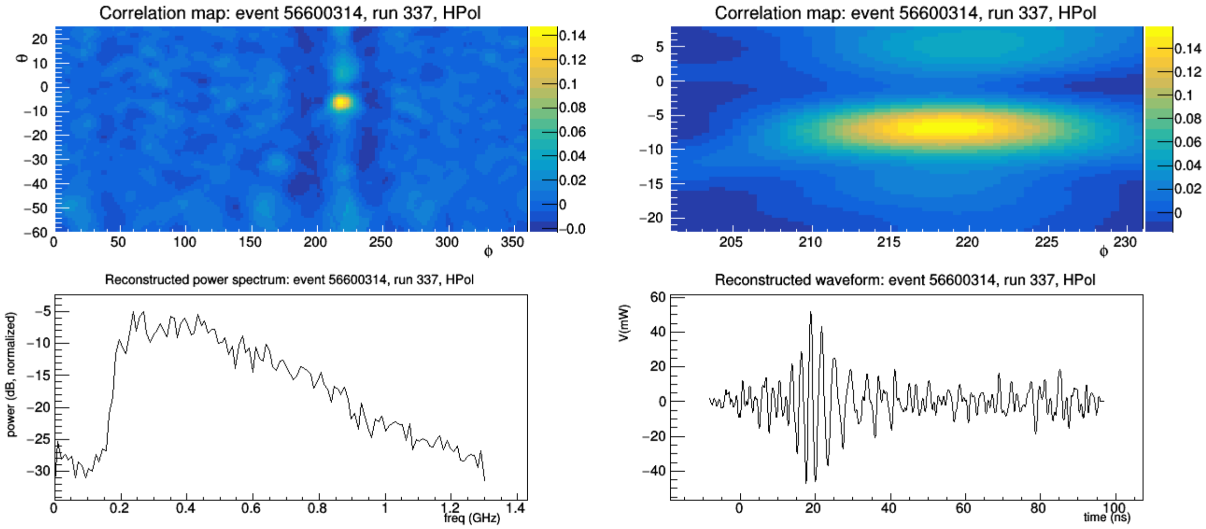
Correlation map: event 56600314, run 337, HPol

Correlation map: event 56600314, run 337, HPol

Reconstructed power spectrum: event 56600314, run 337, HPol

Reconstructed waveform: event 56600314, run 337, HPol

Figure 3.3: Top: Payload-coordinate correlation maps of a WAIS calibration pulse. The pulse is clearly localized to about $(220°, -6°)$. Bottom left: The frequency power spectrum of the coherent sum reconstruction. Bottom right: the reconstructed time domain waveform. The pulse arrives at $\sim$17 ns but dispersion smears it into a series of peaks spanning more than 10 ns.

sition is localized on the Antarctic continent if possible. The payload-relative localization delivered by the interferometry is transformed into earth-centered coordinates using the payload's position and orientation at the time of the event. A ray is traced from the known payload position to the ice surface, using 1km-resolution ice elevation data from BEDMAP2[34]. If the ray trace succeeds, the intersection of the tracing ray with the ice surface is taken as the nominal source location.

Once the event is localized to the continent, an error estimate is necessary. An empirical estimate of the error in payload coordinates was obtained using calibration pulses. A pointing probability distribution function (PDF) is obtained in payload coordinates using the interferometry peak (ϕ, θ) and the empirical pointing error RMS $(\delta\phi, \delta\theta)$ of calibration pulses (Figure 3.4, top panel). By projecting this distribution onto the continent, we obtain a PDF of the event's source location on the continent.

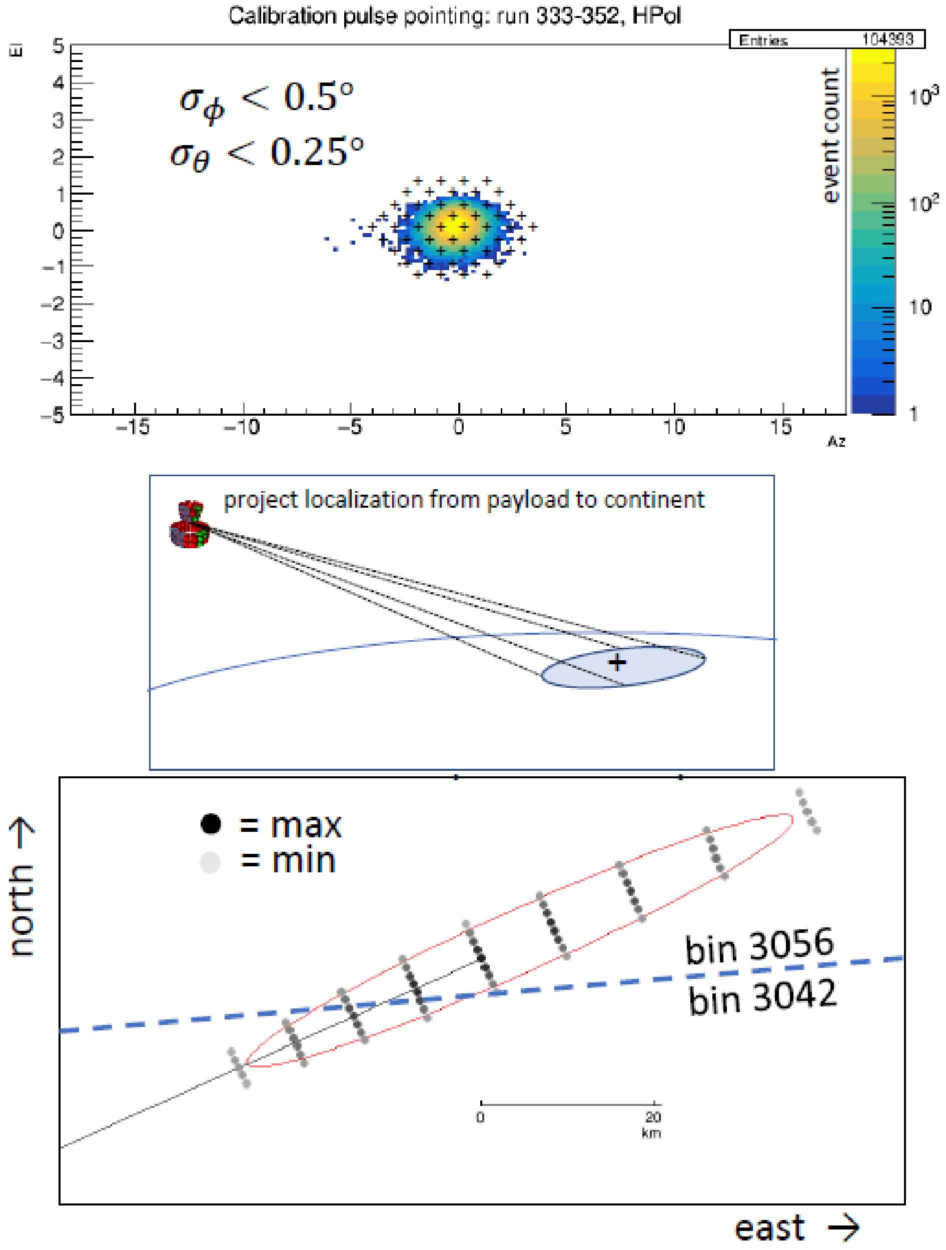

Figure 3.4: Pointing error and event localization. Top: the pointing error of WAIS pulses is used to obtain error ellipse parameters for defining a Gaussian PDF; the hexagonal grid is proportioned to the error ellipse. Middle: the localization grid is projected onto the continent. Bottom: for each point in the grid, the PDF is calculated and allocated to the Healpix bin in which it falls.

The size and shape of the error distribution are sensitive to the source's elevation angle θ_{pl}, especially for far-away events, which originate from near the horizon. In such cases, the error bound implied by a simple error ellipse may not be confined to the continent. Previous analyses such as B. Dailey[33] employed an analytical projection of the cone onto a pseudo-spherical earth, but this method assumes an equal probability distribution over the error cone, and moreover does not accommodate the irregular topography of the earth's surface. We sidestepped reliance on geometrical approximations by generating a hexagonal array of points in and around the pointing ellipse, and ray-tracing every such point to the continent. This method has the advantage of accommodating the contour of the earth's surface, and the resolution of the hexagonal grid can be optimized for precision and computational parsimony. Furthermore, we can easily allocate events to Healpix bins on the continent by applying a PDF to each point on the grid, and accumulating it in the bin to which the point is ray-traced. The PDF is a function of the angular distances $\delta\phi$ and $\delta\theta$ between the grid point and the interferometry peak (ϕ, θ). In this analysis, we used a 2-dimensional Gaussian PDF:

$$PDF(\delta\phi, \delta\theta) = \frac{1}{2\pi\sigma_\phi\sigma_\theta} \exp\left(\frac{\delta\phi^2}{2\sigma_\phi^2} - \frac{\delta\theta^2}{2\sigma_\theta^2}\right).$$

The width parameters σ_ϕ and σ_θ are determined empirically from calibration pulses, as described next.

3.2.4 Error PDF Determination

Calibration for error calculation was performed using calibration pulses from the transmitter systems WAIS Divide and LDB, described in Chapter 2. Signals from

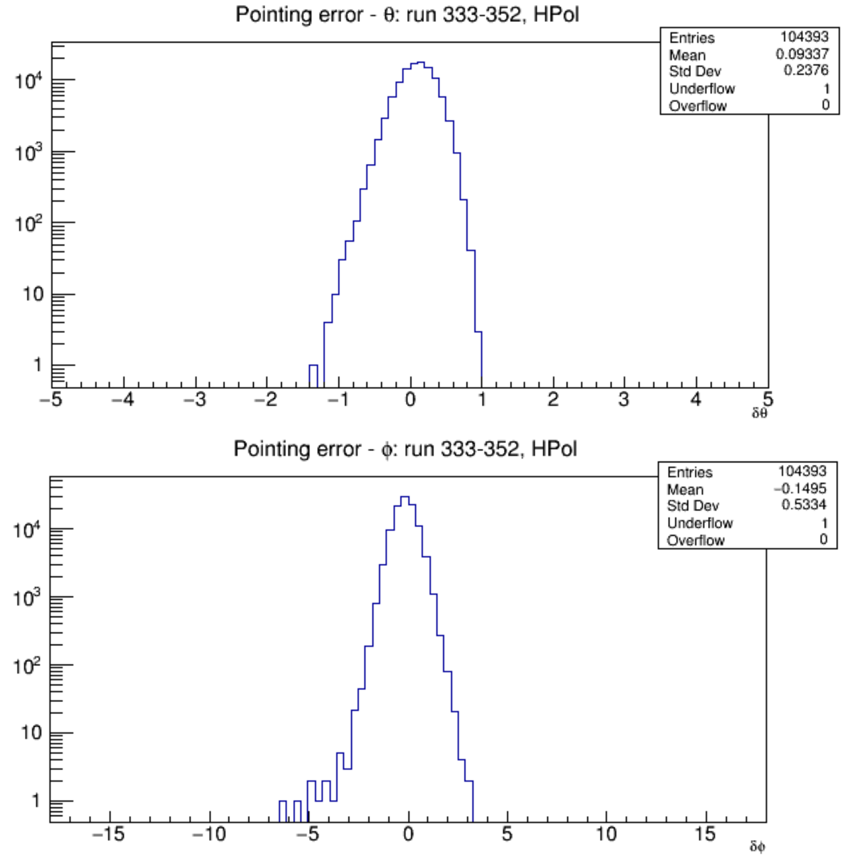

Figure 3.5: Distribution of calibration pulse azimuth pointing error $\delta\phi$ and elevation pointing error $\delta\theta$. The standard deviation of these distributions are used as parameters to the distribute a localization PDF on the continent.

the pulsers trigger the ANITA instrument and are recorded as events. Pointing uncertainties were obtained by comparing ANITA's localization of calibration events to the known locations of the pulsers. We ran interferometry and localization on batches of events from both pulses. Figure 3.5 shows the pointing results for a subset of WAIS-divide events. We used the standard deviations of the errors $\Delta\phi$ and $\Delta\theta$ from these runs as the σ_ϕ and σ_θ in the error PDF.

3.3 Background Separation

ANITA analysis must contend with three major types of background: anthropogenic continuous wave, anthropogenic impulses, and thermal noise. In ANITA's frequency band, continuous wave contamination originates primarily from satellites but can also come from aviation and ground-based radio communications. In previous analyses, CW has been mitigated by analyzing a waveform's frequency power spectrum, identifying peaks, and using notch filters to remove as much CW as possible while leaving intact any impulsive signal content. CW filtering is thus a major component of any ANITA data analysis and is discussed later in this chapter. Impulsive anthropogenic signals originate from transient electric discharges, most familiarly from switching functions of electrical equipment[4]. Since neutrino Askaryan events are impulsive, anthropogenic impulses are not easily distinguished from neutrino-induced events by their spectral characteristics. Instead, statistical methods are applied to the distribution of events on the continent to rule out those originating from areas of human habitation. The third kind of background is thermal noise, i.e., blackbody radiation from the ice. The thermal noise power received at the antenna follows a flat distribution across ANITA's frequency range, with a noise temperature of $\sim$250 K. As mentioned in chapter 2, coincidence algorithms were applied in the instrument hardware to screen out most thermal noise; but thermal fluctuations induce the majority of triggers in spite of this. Analysis cuts on correlation peak value, peak coherent signal power and polarization content are used to distinguish thermal from impulsive events. Following is a detailed discussion of these backgrounds and the methods used to identify mitigate them.

---

[4] Lightning induces radio pulses but virtually never occurs in Antarctica[35].

Any analysis must impose cuts to rule out background events as neutrino candidates, and should, when possible, employ waveform conditioning to mitigate misreconstructions. In this section, methods for contending with thermal, CW and impulsive backgrounds are discussed.

3.3.1 Continuous Wave background

Despite a comparatively low human presence, the Antarctic continent nonetheless abounds with anthropogenic radio, much in the form of narrow-band CW transmissions from radio installations and satellites. The effect of CW on ANITA is twofold. While non-impulsive, a CW event will often localize decisively to its actual source or to a prominent nearby interference fringe, where it could be mistaken for an impulse. Furthermore, strong CW may induce a trigger simply by delivering sufficient power to multiple antennas. Thus, CW must minimally be identified as background, and if possible, filtered out of the event waveforms to mitigate misreconstruction.

Continuous wave content is virtually always anthropogenic, and has two defining characteristics: it is narrow-band, and it is coherent. Carrier transmissions from radio stations are a familiar example. Even in the comparatively quiet radio environment of Antarctica, CW is pervasive and is a major obstacle to ANITA analysis. In ANITA's frequency band, the dominant CW sources are satellites, in particular, geostationary communications satellites. Communications from aircraft and ground bases also contribute. By virtue of its power and coherence, it can either trigger the ANITA payload outright, or it can dominate impulsive events and cause large localization errors, that is, misreconstructions. An example of CW-induced misreconstruction of an impulsive trigger is the calibration pulse shown in Figure 3.6. The actual event source is the peak at (250°, -6°), but this event localized to the brighter peak at (135°, +6°), i.e., the primary image of the CW source. The other peaks at 135° are the interference fringes from the CW; their arrangement is consistent with the vertical positioning of the ANITA antennas within each ϕ-sector. The CW comtamination is also apparent in the lower-left plot of Figure 3.6, in the form of two prominent peaks at 260 MHz and 380 MHz. A hint of modulated CW is also visible in the lower right plot of the time-domain reconstruction, starting at about 60 ns.

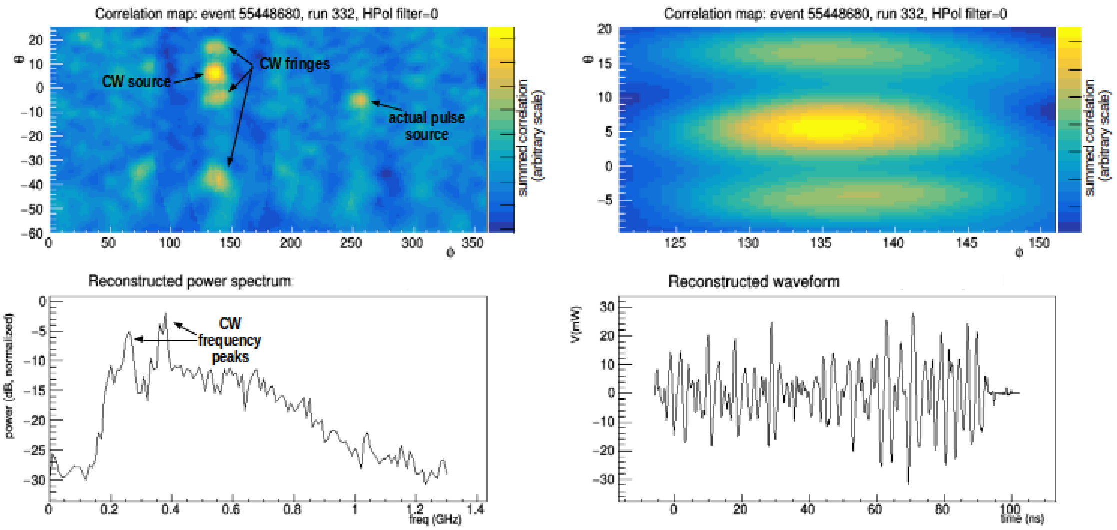

Figure 3.6: CW-induced misreconstruction of a WAIS calibration pulse. Interferometric map (top), reconstructed coherent power spectrum (bottom, right) and reconstructed coherent time-domain waveform (bottom right).

The goal of CW mitigation is two-fold. First, CW contamination should be removed from impulsive events if possible. This may prevent misreconstruction and allow productive analysis of the remaining, impulsive content in the waveform. Second, CW will continue to dominate some events despite attempts to remove it; such incorrigibles must be cut from the analysis. Two steps are therefore inserted into the analysis process to manage CW:

1. Before interferometry, identify CW and apply filtering to remove it, while preserving as much as the impulsive content as possible.
2. After interferometry, use analysis cuts to remove events that remained CW-influenced even after filtering.

Any CW mitigation process must first identify CW contamination in the waveforms. This is done by examining power spectra of individual events as well as in average throughout the flight. If a particular CW frequency is known to be prevalent throughout the flight, then a notch filter of fixed frequency and bandwidth can be applied uniformly to all waveforms in all events. This approach, however, fails to accommodate differences in antenna channel responses, and does not adapt to the changes in the CW environment throughout the flight. Adaptive filters have been developed to evaluate events individually, detect CW excesses and set filter notch frequencies and bandwidths accordingly. Various methods of adaptive filtering have been used previously; all of these work by comparing waveform power spectra to a reference power spectrum, typically representing the thermal background. Since thermal noise power density at the antennas is relatively flat across the ANITA's frequency spectrum, a thermal event, as received at the ANITA instrument box, should conform strongly to the inherent frequency-dependent gain of the signal chain. It is reasonable, then, to consider either a measured signal-chain power spectrum, or a sample of nonimpulsive trigger events from the flight, as a thermal baseline. To assess ambient background, both the flight computer and the instrument GPS hardware induced intentional triggers every second. These *minimum-bias* events can be compared to actual event power spectra to identify frequency peaks for filtering.

All digital filtering methods have side effects; disturbance of the phase/magnitude relationships can lead to artifacts that can resemble impulsive signals. Furthermore, filters will be noncausal unless care is taken to avoid this by, say, constraining transfer function poles or enforcing the Kramers-Kronig conditions. The goal of filtering in our case, however, is to improve pointing and acceptance efficiency of impulsive events. Since lack of causality does not necessarily detract from this objective, causality is not explicitly imposed in the filters used for analysis. Two adaptive methods, dubbed the sine subtraction filter and the geometric filter, both yielded improvements in ANITA-II analysis over the rectangular notch filters used in [36]. We therefore considered these two methods, discussed next, for use in our analysis.

3.3.2 Filtering

We considered two filtering methods. The first is a method developed by the ANITA Collaboration known as *sine subtraction*[37]. We also considered a "geometric" filtering method described in Dailey[33].

The sine subtraction method differs from most other filtering methods in that while most filters modify a waveform using a multiplicative transfer function, the sine subtraction method attempts to subtract CW contamination from the event waveforms. The algorithm finds a peak f_p in the frequency spectrum, which functions as a starting point for fitting a three-parameter sinusoid function S to the time-domain waveform,

$$S(t; A, f, \phi) = A \sin(2\pi f t - \phi).$$

The fitted function is then subtracted from the original waveform. This process of peak-finding, fitting and subtraction is repeated until the fraction of waveform power removed by the subtraction falls below a user-programmable threshold.

The second filtering method considered was the geometric filter. Previous collaborators noticed that existing digital notch-filtering methods led to abrupt changes in waveform phase near the notches, probably leading to artifacts in interferometry and reconstruction. The geometric filter attempts to mitigate these effects by identifying the phasor associated with the CW contamination and subtracting it out[33].

Adaptive methods were also developed for determining the locations and widths of the notches. CW in a waveform should render as a peak in the power spectrum, over and above the thermal background. In order to assess event contamination, a sample of thermal noise is required. This *baseline* thermal sample is then used as a reference for comparison against which the spectrum of an event waveform; large, narrow peaks in the difference are indicative of CW. As previously mentioned, the instrument hardware and the flight software induced intentional triggers every second. These events, however, are believed to be biased to absence of a hardware trigger and therefore may not be best suited for a rigorous comparison. Instead, We collected a sample of downgoing events that localized away from the sun and away from known locales of geostationary satellites, believing that such event would be free of ground-based CW. Since the phase of thermal noise is random, only the frequency power spectrum is of interest; we took an average power spectrum of this sample for each antenna/polarization channel. Despite our efforts to avoid CW contamination in the baseline sampling, a strong peak at 260 MHz was present, probably due to satellites. We removed CW peaks in the power spectrum by truncating them to the level of the second-highest peak in the power spectrum, then smoothing with a Gaussian kernel to remove sudden slope changes.

The performance of the two filters was similar; we ultimately selected the geometric filter for its slightly lower rate of misreconstructions of calibration pulses[33].

3.3.2.1 Post-interferometry removal of CW-influenced events

CW produces interference fringing by its nature, so CW-induced misreconstructions are often characterized by multiple correlation peaks: one corresponding to the actual CW source location and the others being ghost images due to fringe effects. An analysis cut is therefore made on the ratio of the second highest peak in the correlation map to the highest peak: a large value for this ratio implies a high likelihood of either CW contamination or a thermal trigger.

3.3.3 Thermal Triggers

The ANITA-III trigger worked by requiring surges in power coincident across multiple antennas. While in principle this should screen out uncorrelated thermal fluctuations, multiple antennas may indeed experience nearly simultaneous stochastic fluctuations, which must be tolerated in order to achieve acceptable instrument sensitivity. The low correlation of thermal events across antennas should be recognizable: we expect a correlation map lacking a distinct peak and a non-impulsive coherent sum waveform. Previous analyses such as Vieregg[36] employed a 2-parameter linear discriminant composed of the peak Hilbert envelope value of the reconstructed coherent sum waveform and the peak cross-correlation value from the interferometric sky map. A similar method is utilized in this work, and extended to other event characteristics.

Informed analysis of thermal triggers is facilitated by comparing the attributes of known impulsive events to known thermal events. While impulsive samples exist in the form of calibration pulser events from WAIS-divide, LDB McMurdo and Hi-Cal, obtaining a sample of known thermal triggers is less trivial. The **PPS1** and **Soft** triggers, intentional on-board triggers generated by payload hardware and software, provide a representative sampling of the ambient radio environment (including, importantly, prevailing narrow-band content). These samples, however, are taken in the *absence* of an RF trigger, and thus do not well represent the true background of interest: thermal fluctuations that *did* induce RF triggers. Vieregg used a sample of sky-localized, non-Solar events. A similar method is tried here, but such events are scarce in the ANITA-III dataset since large regions of the sky are dominated by satellite CW. Also, the representativeness of the sky sample is questionable due to the very low noise temperature of the sky ($\sim$10 K) as compared to the ice ($\sim$250 K); thus is motivated an attempt to obtain a thermal noise sample from the ice itself.

3.4 Analysis Process

The analysis process consists of the following steps:

1. Events are selected for interferometry. Non-RF triggers and calibration pulser events are cut. The final analysis traverses the 90% sample described above, implicitly cutting the 10% sample used in calibration and optimization.
2. Interferometry is performed on every selected event and correlation maps of the sky are constructed. Events are localized on the payload-relative sky and localization to the continent is attempted. Quality cuts are applied to remove

events affected by instrument aberrations or containing faulty or incomplete data. Interferometry results are stored in a database.

3. Uniformly-applied analysis cuts are performed. Waveforms that fail to reconstruct to the continent or that coincide with the position of the sun or satellites, or their respective reflections, are removed. Properties of the reconstructed waveforms are used to remove non-impulsive events.

4. The last analysis cut is intended to function as a final thermal noise rejecter, as well as to remove anthropogenic impulse events. The continent is divided into equal-area bins and a linear discriminant is optimized separately for each bin. This cut is an alternative to the clustering cuts performed in [36], and was motivated by the increased instrument sensitivity of ANITA-III and the noisier background at play during the mission, in comparison to previous flights.

5. Events surviving the cuts are used to make a list of weighted neutrino candidate events. An upper limit on the cosmic neutrino flux in the energy range of interest can be asserted, even if no event survives the analysis cuts.

3.4.1 Data Blinding

Before final analysis is attempted, the ANITA-III event data is blinded using a procedure described in [38]. To accomplish this, ∼10 randomly-chosen events were deliberately replaced by modified copies of WAIS-divide calibration pulser events, under the expectation that these should produce false positives in the final analysis. Since WAIS events are predominantly HPol and neutrino signals are expected to be VPol, the polarizations of the selected WAIS events were reversed such as to induce false positives in a VPol search. Ten percent of events were randomly selected and set aside for calibration and testing of the analysis software. The remaining ninety percent were preserved in order to function as the blinded dataset in the final analysis.

3.5 Quality cuts

Quality cuts eliminate waveforms with various problems that would prevent effective interferometry. The quality cuts used in this analysis are listed below.

SURF Saturation Events: Voltage levels above 1500 mV push the SURF out of its operating range, resulting in distortion to the waveform. [36] Events containing more than three waveforms with voltages over 1500 mV are therefore cut.

DC-offset Events: Waveforms with a large DC offset are thought to indicate digitization problems associated with high payload trigger rates [36]. Events with mean absolute value >100 mV in any channel are removed.

Short-trace Events: Events containing fewer that 240 samples in any waveform are cut.

Payload blast events: Events that trigger strongly in many phi-sectors are thought to result from on-board payload noise; we therefore cut events with at least 15 channels or nine phi-sectors with peak voltages >400 mV.

Nadir Noise Events: Events with excessive power in the lower antenna ring are thought to result from payload noise, and events appearing to originate from underneath the payload are outside the antennas' directional sensitivity. An event is cut if the maximum peak voltage in the top ring is less than half the maximum peak voltage in the bottom ring.

3.6 Analysis Cuts

Analysis cuts are applied after interferometry and quality cuts in order to remove non-impulsive and anthropogenic events. We started by applying analysis cuts similar to those used in [36] and [33]. We then applied newly-developed cuts based on events' circular polarization characteristics, primarily in order to reject satellite events.

As a general rule, the cuts use the values of selected event characteristics, e.g., correlation peak value, or reconstructed peak power, or waveform SNR(s) to remove anthropogenic, non-impulsive, or other extraneous events. The cuts used and their motivations are discussed here, while the actual parameter threshold values used to make the cuts in final analysis are discussed in the next chapter.

We applied analysis cuts in three stages. The first stage includes cuts for which the parameters are simple or well established. These cuts will remove a large proportion of events from the dataset and thus facilitate faster processing in the subsequent stages. The second stage consists primarily of cuts based on waveform characteristics, which require repeated test runs in order to optimize their parameters and thresholds. The final stage is a single cut based on a linear discriminant; this cut requires careful tuning and is discussed in detail later in this chapter.

3.6.1 Stage 1 Analysis Cuts

The first stage of cuts is comprised of very simple cuts such as to remove sun-induced events, events failing to localize to the Antarctic continent, events from the sky or from underneath the payload, and events that localize away from triggering phi-sectors.

Solar Reflection Cut. Events pointing directly back to the Sun or to known satellite locations are easily identified and removed; furthermore, such events come from above the horizon and are trivially cut on altitude angle[5]. But a significant number events localize to points near the Sun's reflection. While the general problem of a point source reflecting on a sphere is algebraically cumbersome, one can easily calculate the source direction of a presumed reflection event, given its arrival zenith angle θ_I, the payload altitude h_P, and the surface altitude at the reflection point, h_I. The geometry is shown in Figure 3.7. By the law of cosines,

$$(R_e + h_I)^2 = (R_e + h_P) + d^2 + 2d(R_e + h_P)\cos\theta_I, \tag{3.5}$$

where $R_e \approx 6398$km is the earth's curvature radius at latitude 80°S. Then,

$$d = -(R_e + h_P)\cos\theta_I - \sqrt{(R_e + h_P)^2\cos^2\theta_I - 2[R_e(h_P - h_I) + h_P^2 - h_I^2]}. \tag{3.6}$$

The central angle α is

$$\sin\alpha = \frac{d}{R_e}\sin\theta_I, \tag{3.7}$$

and the source direction polar angle at the reflection point is

$$\theta_R = 2\alpha - \theta_I + \pi, \tag{3.8}$$

---

[5] This applies to neutrino searches since neutrino events arrive only from the ice.

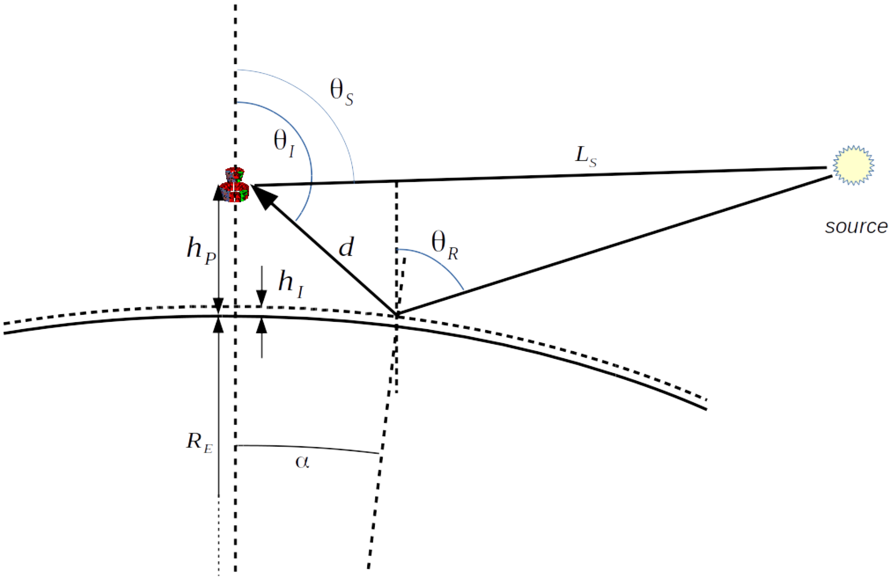

Figure 3.7: Geometry of reflection events. A reflection localized to elevation angle θ_R is traced to a source elevation angle θ_S.

then for a source assumed at distance L_S, the source zenith angle θ_S at the payload is obtained by

$$\sin(\theta_S - \theta_R) = \frac{d}{L_s} \sin(\theta_I - \theta_R) \tag{3.9}$$

Note that if $d \ll L_S$, $\theta_S \approx \theta_R$ and (3.8) suffices. Since d cannot exceed the flight-altitude horizon distance of ∼750 km, and $L_S \approx 36,000$ km for geostationary satellites, (3.8) is implemented in `AnitaTools`.

Localization to Continent Cut. Neutrino events are expected only from the ice, not from the sky. Events that fail to reconstruct to a location on the continent are therefore cut.

Elevation Angle Cut. Events originating from underneath the payload are cut by applying a threshold to the payload-coordinate elevation angle θ of the correlation map peak.

Triggering ϕ-sector Direction Cut. Impulsive events should trigger antennas that face the source. Events which localize far away from triggering ϕ-sector directions are therefore cut.

Calibration Pulser Cut. Events originating from the calibration pulsers at WAIS divide and LDB are removed based on their nanosecond trigger timestamps.

The stage 1 cuts remove about 95% of events from the dataset; this facilitates much faster processing of the remaining cuts.

3.6.2 Stage 2 Cuts

The second stage of analysis cuts are based on features of the interferometric maps and coherently-summed reconstructions of events.

Ratio of Highest Peaks Cut. Impulsive events should render as a distinct peak in the correlation map, and this peak should be much brighter than any other peaks in the map. CW, on the other hand, induces interference fringing which may result in other strong peaks in addition to the primary image. Thermal events, since they are incoherent, are expected to render many small peaks in the correlation map. The ratio of the second-highest peak in the map to the highest peak is calculated, and events for which this ratio exceeds a predetermined threshold are cut.

Correlation Peak Cut. An impulsive event should have a large correlation map peak value, compared to nonimpulsive events. Events with a correlation peak below a threshold value are cut.

Hilbert Peak Cut. Impulsive events are expected to have their power concentrated over a short time interval, and this should have a high peak power value in the time domain. The magnitude of the *analytic signal* of a waveform $V(t)$ is a measure of its instantaneous power. The analytic signal is a complex-valued time-domain function,

$$A(t) = V(t) + iH(t),$$

where $H(t)$ is the Hilbert transform of $V(t)$. The "Hilbert peak" is the peak value in the magnitude of the analytic signal of the reconstructed coherent sum waveform. Events with a Hilbert peak value below a threshold are cut.

3.6.3 Stage 3 Cuts

Stage three of analysis consists of six cuts. Three of these are optimized in order to establish the strongest possible neutrino flux limit. The other three are based on analysis of the allocation of events across Healpix bins on the Antarctic continent. The cuts are described here; the optimization process is described in detail in Chapter 4.

Circular Polarization Peak Proximity Cut. Neutrino-induced impulsive events are expected to be linearly polarized, whereas satellite-induced events are expected to be circularly polarized. Figures 3.8 and 3.9 show event localizations on the sky for the entire ANITA-III flight in the linear and circular polarizations, respectively. The smile-shaped arc in the upper center consists of events localized to the Sun; the arc is smeared by the changing latitude of the payload throught the flight. Solar reflection events below the horizon comprise a mirror-image of the sun arc; this feature is much stonger in the horizontal polarization due to the large incidence angle. The 90°-long horizontal stripe at 6°below the horizon, visible in the horizontal polarization but not the vertical, is calibration pulses from WAIS. The spot at (170°, -7°) in the horizontal polarization is events from the LDB cal pulser.

Concentrations of events tracing the analemmae of geostationary satellites are clearly visible in the north in Figure 3.8, just above the horizon. Apparent reflections from these satellites are also visible just below the northern horizon, albeit less distinctly than the direct satellite localizations. The visibility of the satellites is roughly equal in the horizontal and vertical polarizations: this is consistent with circularly-polarized signals. In Figure 3.9, however, the satellite traces appear strongly in the left-handed circular polarization (even stronger than in the linear polarizations), but are virtually invisible in the right-hand polarization. This shows that the satellite-induced events are intensely left-hand-circular polarized. Furthermore, this indicates that satellite events generally do not localize to the same point on the sky in left- vs. right-hand circular polarization. Linearly-polarized events thus appear to localize to *the same* point on the sky in both circular polarizations while circular-dominated events usually localize to *different* points in the two circular polarizations, suggesting that the mutual proximity of localizations in the two circular polarizations may be a powerful cut on satellite-induced triggers. We therefore cut events for which the distance between the left- and right-hand localizations exceeds a threshold value: we compare he location of the strongest interferometry peak in each circular polarization to that of the two strongest peaks in the opposite circular polarization (a total of three comparisons) and take the minimum separation of these three peak pairs as the cut variable. Events for which this separation *exceeds* a threshold value are considered to be circularly polarized, and are cut.

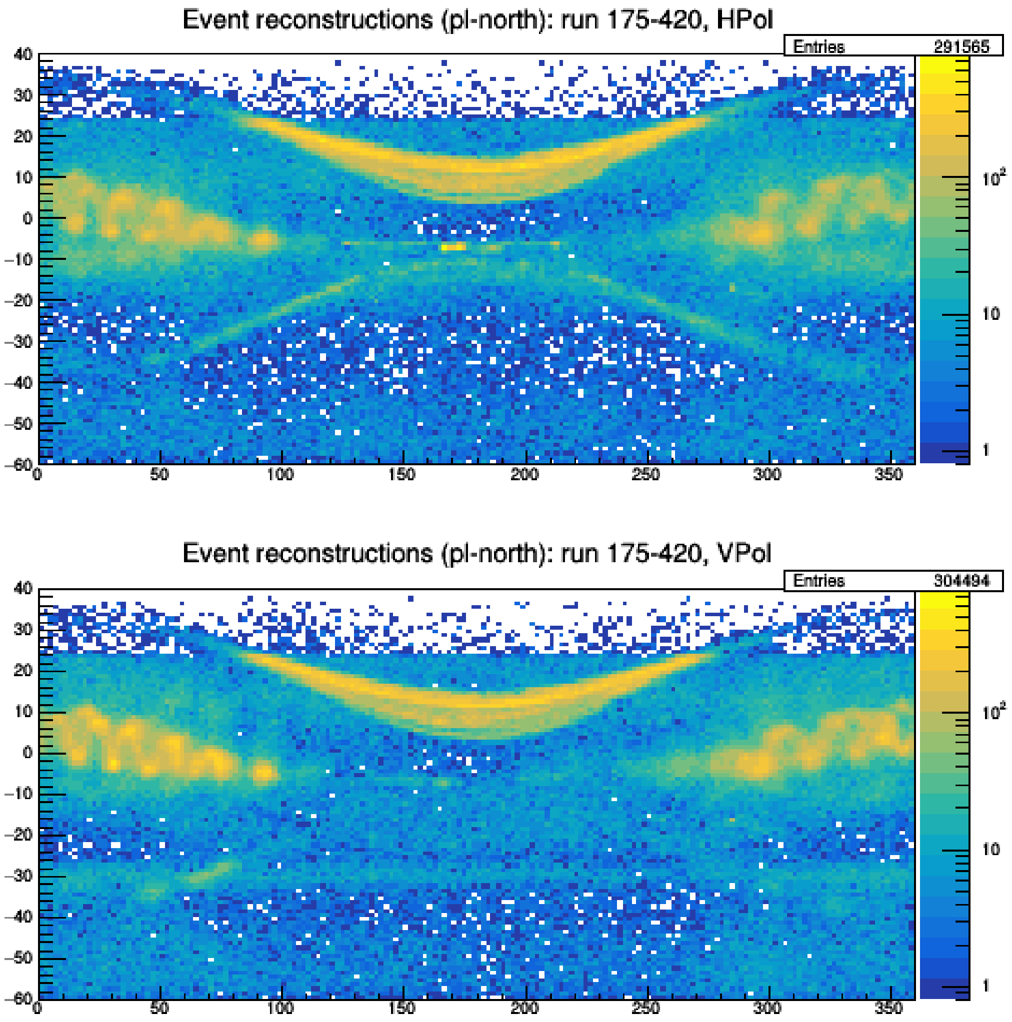

Figure 3.8: Event localizations across most of the ANITA-III flight, in the linear polarizations. The satellite analemmae traces and, to a lesser extent, their reflections, are clearly visible in both polarizations, consistent with circularly-polarized events.

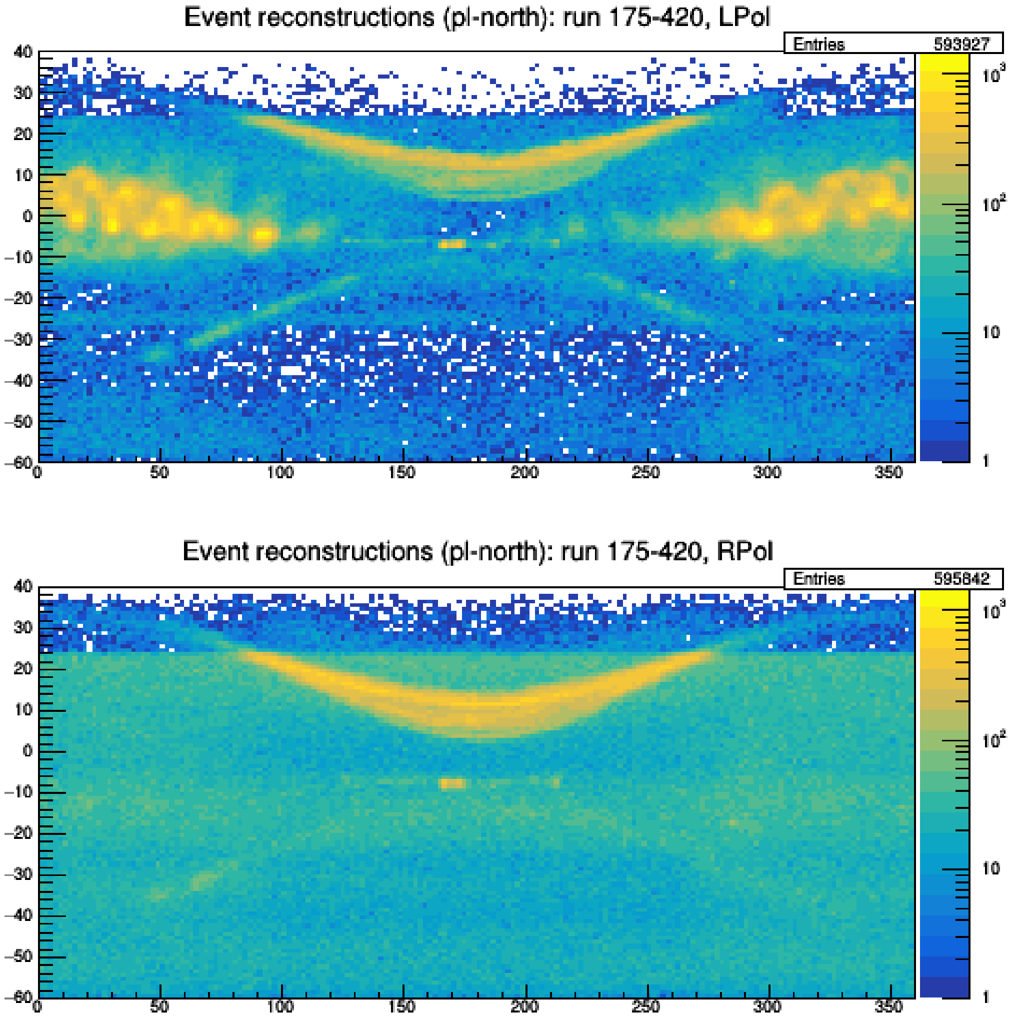

Figure 3.9: Event localizations across most of the ANITA-III flight, in the circular polarizations. The satellite analemmae are clearly visible in the left-hand polarization, but are virtually indistinguishable in the right-hand, a strong indicator of circularly-polarized events.

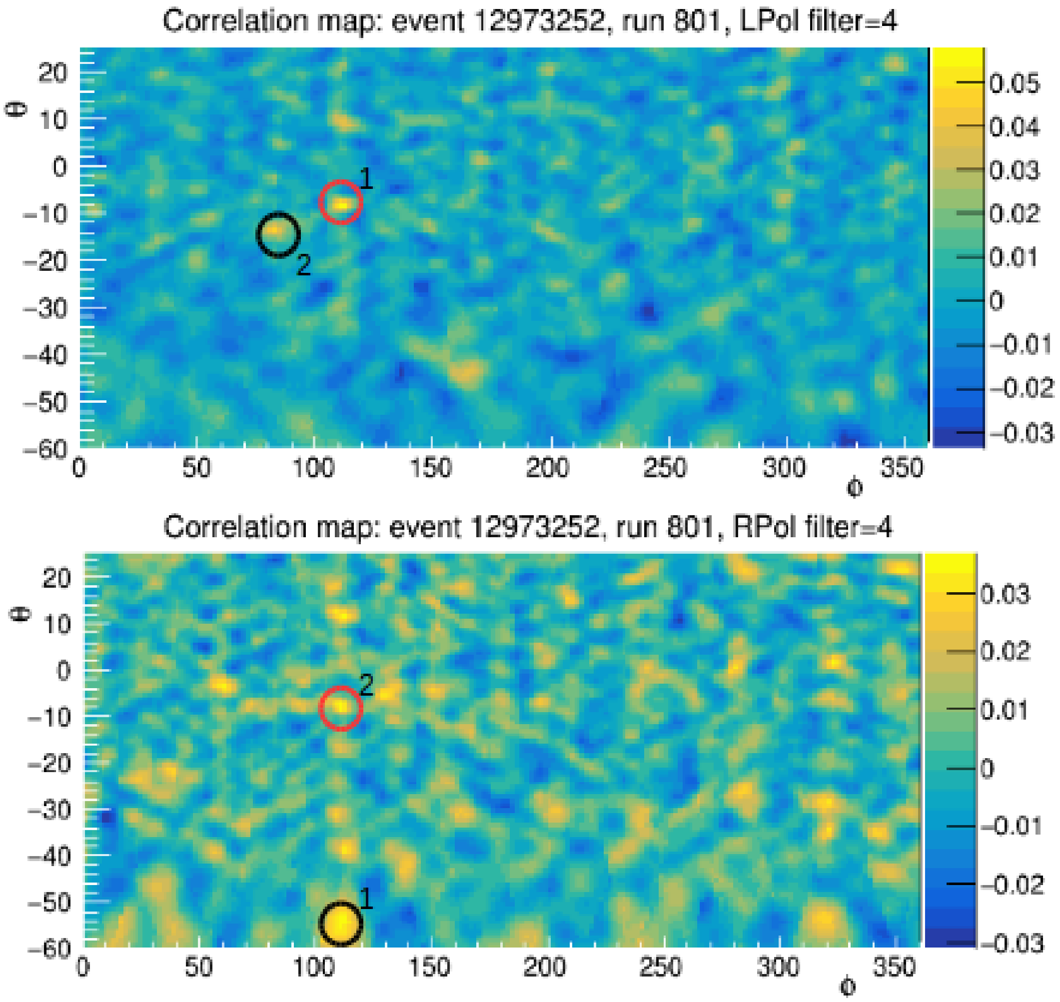

Figure 3.10: Peak location comparisons in the circular polarizations. The strongest peak in each circular polarization is compared to the two strongest peaks in the opposite circular polarization. The event in this example would pass the cut because the strongest left-polarized peak and the second strongest right-polarized peak are in close proximity to each other, indicating that this particular event is probably strongly linearly polarized.

Circular Polarization Strength Cut. Strongly linearly-polarized events should render localization peaks of roughly equal strength in the circular polarizations. In order to cut on this characteristic, we take the localization in linear polarization, and compare the values from two circular polarization correlation maps at that same location. Since the exact location of the correlation peak varies slightly across polarizations, even for strongly impulsive events, we consider the maximum correlation values within a small window of the two circular polarizations, constructed around the linear polarization localization point. The lesser of these two values, normalized with the linear polarization peak value, is used as the cut variable.

Low Binned Event Weight Cut. As described previously in section 3.2.3, events are allocated to bins on the Antarctic continent using a localization PDF based on payload pointing resolution. Thus, some events will be distributed across multiple bins and assigned a fractional weight in each bin. Events with weights that fall below a threshold value are cut.

Bin Rejection Cut. During the optimization process, bins on the continent will be subjected to selection criteria to ensure that a meaningful optimization can be made. The criteria are described in detail in the following Chapter. Events allocated to bins that fail the selection criteria are cut.

Bin-specific Linear Discriminant Cut. The last cut is made on a linear combination R of on the signal-to-noise ratio of the reconstructed coherent-sum waveform and the peak value in the correlation map. Events pass the cut if they fall to the right of the line

$$R_i = SNR + mP_{corr}, \tag{3.10}$$

where SNR is the signal-to-noise ratio of the coherent reconstructed waveform and

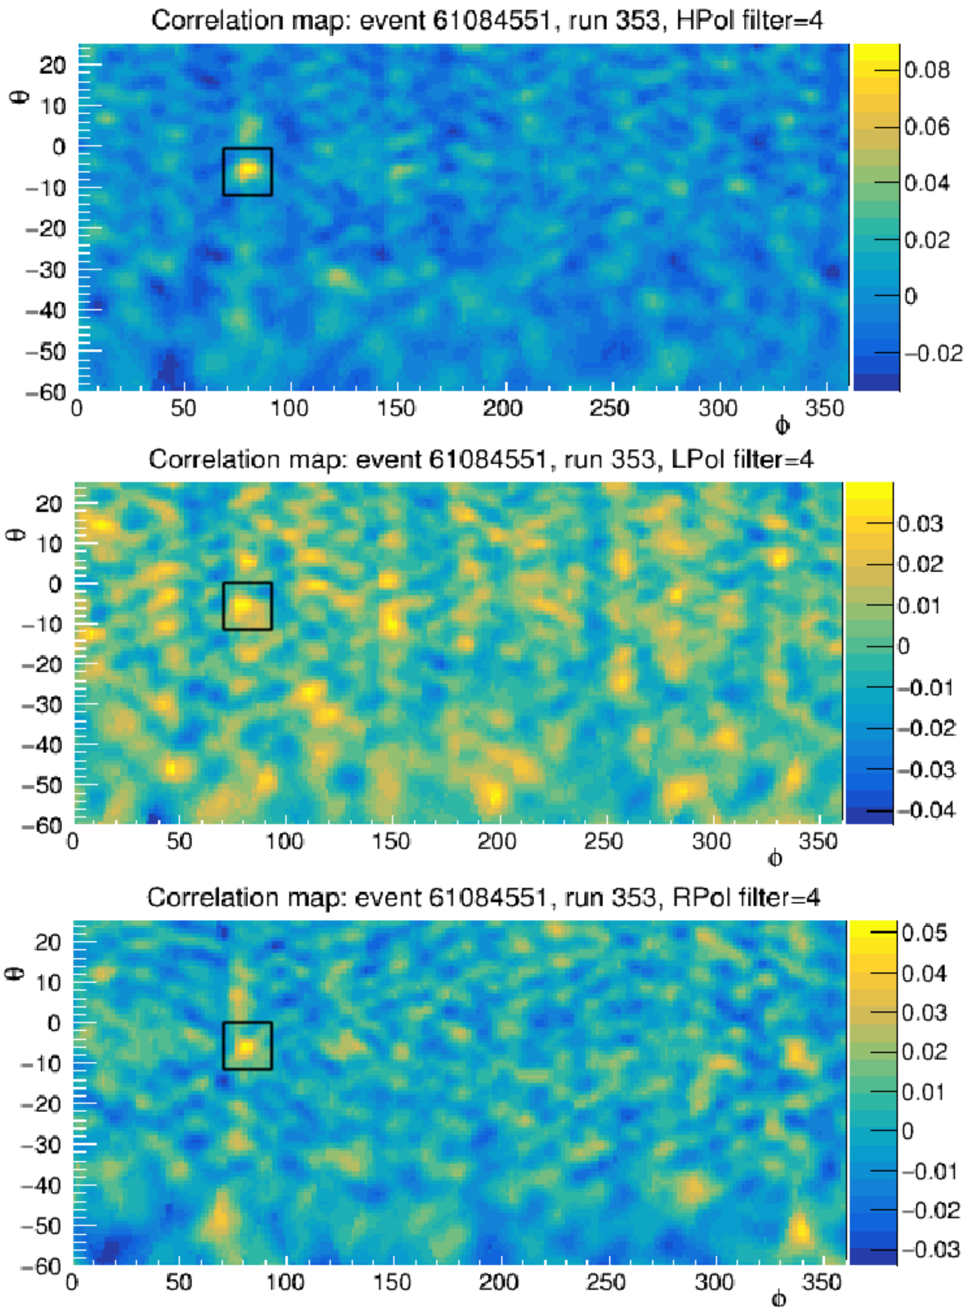

Figure 3.11: Correlation strength comparisons in the circular polarizations. This event, a WAIS calibration pulse, is strongly horizontally polarized, and therefore renders roughly comparable correlation strength in the two circular polarizations, in a window constructed around the HPol peak location. This event passes the circular polarization comparative strength cut.

P_{corr} is the correlartion peak value; that is,

$$R = SNR - mP_{corr} > R_i, \tag{3.11}$$

The slope and intercept parameters m and R_i are chosen by a process described in detail in Chapter 4. This cut is intended to enhance rejection of nonimpulsive events as well as anthropogenic impulsive signals.

Previous ANITA-II analyses[36] used a clustering algorithm to identify such events for cutting. Neutrino events are rare, so the probability of two neutrino events originating from closely-spaced locations on the ice is very small. The clustering algorithm therefore cut events that occurred in close proximity to other events. ANITA-III, however, flew in a noisier environment than that of -II, leading to concerns that clustering algorithms would exclude too large a proportion of the continent for effective analysis; essentially all events would be cut[33]. It was noticed also that for an ANITA-II 10% sample, the distribution of the values R in equation (4.1) tends to fall off exponentially after its peak, as shown in Figure 4.2. Dailey[33] took this exponential behavior to be a representation of background, and developed a new method for discriminating anthropogenic and thermal events. A spherical tesselation obtained from Healpix[39] was used to divide the surface of the earth into 3072 bins, each with an area of about $1.6 \times 10^5\,\mathrm{km}^2$. Every event was allocated to one or more bins by projecting its payload-frame localization ellipse onto the continent and weighting it proportionally to the fraction of the ellipse error falling in each bin. The linear discriminant cut threshold was then set individually for each bin to maximize the ability to discern a signal above this background. In this analysis, we use a similar process with a substantial enhancements; the optimization of the linear discriminant cut is described in detail in Chapter 4. As in Dailey[33], we allocate events to a Healpix tesselation of the continent; but instead of calculating fractional ellipse areas, we use the method described in Section 3.2.3. This method is computationally simple, and also allows us to develop a PDF of the events point of origin on the Antarctic continent.

3.7 `AnitaTools` Software Components

We performed our analysis using a software framework consisting of components developed by the ANITA Collaboration as well as original components developed as part of this Work. Throughout the ANITA campaign, collaborators have developed and evolved software for common analysis functions. A comprehensive discussion of the Collaboration software is not undertaken here, but a few components utilized heavily in this work merit brief discussion.

For use in ANITA-III and later missions[6], the Collaboration collected software components from previous flights, augmented them and refactored them into the `AnitaTools` package. The use of a C++ object-oriented class architecture provides layered data and process abstraction and thus facilitates modification and extension by collaborators. `AnitaTools` provides a mechanism for performing interferometry on Anita events, localizing them to the sky or continent, and reconstructing waveforms representing hypotheses of event source locations. Resultant pointing and localization hypotheses as well as reconstructed waveforms and their characteristics (e.g., SNR described in section and analytic waveform peak) are stored in a convenient format for subsequent cuts and statistical analysis.

---

[6] ANITA-IV launched in December/2016.

The container for an event waveform from an individual antenna/polarization channel is the `AnalysisWaveform` class; it upsamples the waveform and interpolates it on to a common time scale for analysis. Behind-the-scenes synchronization of time- and frequency-domain representations enables the user to manipulate and analyze a waveform in either domain, on equal footing.

The set of `AnalysisWaveform`s from a single event are organized in an instance of the `FilteredAnitaEvent` class. As its name suggests, this class also provides a filter interface; a `FilteredAnitaEvent` can be instantiated specifying a `FilterStrategy` (a list of zero or more filters) to be applied to the waveforms. Multiple filter types are also provided, including various causal and acausal notch filters, the sine-subtraction filter previously described in Section 3.3.2, and a hybrid filter for converting from the linear polarization basis to the circular basis. This facilitates experimentation with various filtering strategies for mitigating CW and noise. We augmented this by refactoring the geometric filter developed by Dailey[33] into the `AnitaTools`.

The `Analyzer` class is the interferometry engine; the user instantiates an `Analyzer`, then runs it using a `FilteredAnitaEvent` as input. Reconstructed waveforms are delivered in instances of `AnalysisWaveform`; their features and other results are stored in an `AnitaEventSummary`.

Simulated events are useful for optimizing the analysis cuts,as well as in establishing flux limit hypotheses. These were obtained from `icemc`, the prevailing simulation utility of the ANITA Collaboration. `icemc` starts with assumptions of diffuse cosmic UHEν flux, which can be chosen from a large selection of theoretical or toy models. After calculating an interaction probability, or *weight*, it then simulates the physics of the of neutrino interactions in the ice, the resulting shower, and the propagation of the ensuing Askaryan pulse to the payload[7]. At this point, with the power at the antennas known, the ANITA signal chain and triggering system response is simulated, and if an L3 trigger occurs, a data event is recorded. We used `icemc` events generated according to the maximum Kotera model[11].

3.7.1 Original Analysis Software Components

Our own original analysis software components are summarized in this section. These components are mostly scripts for executing `AnitaTools` components on batches of events. These are used both in optimizing the analysis cut parameters, and in performing the actual analysis: the analysis process is tested on a sample of ten percent of the ANITA-III data sample using various values for cut parameters. The cut parameters are tuned and optimized through multiple iterations of the analysis process; once tuning is complete, the results of the tuning are fed back into the analysis programs, and the final analysis is performed on a blinded dataset created from the remaining 90% of the data. Much of the work in this analysis involved running analysis steps on simulation events and in-flight calibration events in order to assess performance and determine analysis cut parameters. Components of our analysis software therefore allow for the selection and analysis of calibration pulses and simulated events. Calibration pulses from the WAIS and LDB pulsers are present in the ANITA-III flight dataset and are easily identified by trigger timing. Our analysis flow is a sequence of three main programs:

The first program, `runInterferometry`, reads events from the ANITA-III flight dataset, performs interferometry and localization, and applies quality cuts. Events are

---

[7] Probability-weighted event counts are calculated by adding up the event weights instead of counting the events discretely.

processed in the linear polarization (HPol, VPol) and the circular (LPol, RPol). User options in `runInterferometry` allow for the specification of various runtime options, including filtering and normalization methods (along with relevant parameters), input and output directories, and event selection constraints. Also, a single event can be selected for processing such that plots of correlation maps and event waveform reconstructions can be viewed in real time. The output consists of `AnitaEventSummary` objects (as defined in `AnitaTools`) and a few other values calculated from the waveform; header and gps information is tagged along for convenience in subsequent analysis.

The second program, `runAnalysisStage01`, reads results from `runInterferometry` and calculates the variable values used in the analysis cuts. The analysis cuts are actually applied in three stages, and are described in detail below Section 3.6. It also assigns events to Healpix bins on the continent (this is also described in Section 3.6). texttttrunAnalysisStage01 then produces an output summary tree containing the event numbers and the calculated cut variables. If requested, the stage 1 analysis cuts are enforced, such that only events passing the stage 1 cuts are written to the output. It also plots histograms of the cut variables. Optionally, `runAnalysisStage01` selects calibration pulses or simulation events. In this case, it evaluates and plots pointing error distribution of the localization errors, and the reports the means (μ_ϕ, μ_θ) and standard deviations $(\delta_\phi, \delta_\theta)$. Event acceptance efficiency for calibration pulses or simulations are also reported under this option. User options also exist to enable or disable the enforcement of quality and (stage 1) analysis cuts; it is also possible to apply cuts inversely, that is, retaining events that fail certain cuts instead of those that pass. This is useful in testing and cut tuning, as well as for identifying and investigating events that fail selected cuts. In all cases, maps of event localizations and plots of analysis cut variables are produced.

The next program in the sequence is `runAnalysisStage02`; this program enforces the stage 2 (and optionally, stage 3) analysis cuts and produces plots of the cut variables similar to those produced by `runAnalysisStage01`. The output of `runAnalysisStage02` is identical in format to that of `runAnalysisStage01`, but it contains only events that pass the stage 2 analysis cuts.

The last three cuts, the two circular polarization cuts and the linear discriminant, require detailed tuning. This is done using program `optimizeLDCut`. This program performs mock analyses using various values of the linear discriminant slope and intercept parameters, and performs the optimization process described in 3.6.3. The output consists of a spreadsheet for each slope tried. The number of Healpix bins that are successfully fitted and that meet pseudo-experiment p-value requirements are listed, as well as the expected background and the number of simulated events passing various stages of cuts. An example is shown in table E. This optimization is described in detail in the following chapter.

Chapter 4

Choosing the Circular Polarization and Linear Discriminant Cut Parameters

In this chapter we describe the process used to optimize the final three cuts of the analysis process, that is, the two circular polarization cuts and the linear discriminant. Our method is similar that that used by Dailey[33] to optimize the linear discriminant, but we made several major enhancements. First we optimized not only the linear discriminant intercept, but also its slope, as well as the two circular polarization cuts described in the previous Chapter. Second, instead of maximizing the acceptance of bins or the number of simulated events in accepted bins, we sought the strongest upper limit on neutrino flux that can be obtained from the cuts, while retaining about half or more of our total simulated events in the process.

4.1 Parameters to be optimized

As explained in Section 3.6.3, the two circular polarization cuts are based on the angular separation of the correlation peaks in left- and right-hand circular polarizations, and on the strengths of the circular polarization correlations near the point of linear polarization localization. We call these parameters respectively C_1 and C_2. The parameters for these cuts are simply thresholds on those values: the peak separation must fall under its threshold, and the circular polarization strengths must exceed a threshold. Regarding the linear discriminant cut, we recall eq. (3.11), which gives the criterion

$$R = SNR - mP_{corr} > R_i, \tag{4.1}$$

The parameter set that must be determined thus is $\{C_1, C_2, m, R_i\}$. A single set of values of the parameters $\{C_1,\ C_2,\ m\}$ will be chosen and applied across all Healpix bins; the linear discriminant intercept R_i, however, is to be optimized individually for each bin in order to accommodate the differing background levels across bins.

4.2 Optimization Process

Our optimization process is a four-level loop. The outer three levels are iterations over ranges of values of $\{C_1,\ C_2,\ m\}$. In the innermost loop, we take a particular $\{C_1,\ C_2,\ m\}$ and optimize against every retained Healpix bin to obtain the R_i specific to each bin. Our goal is to find the combination of parameters that generates the lowest possible limit on the neutrino flux, using scaled simulated event counts from `icemc` as a hypothetical representation of the flux. `icemc` delivers a set of events that follow an energy distribution form factor from the maximum Kotera model[11], but the scaling of this distribution is arbitrary. Thus, after optimizing our cut parameters, we will choose the scale factor that enables us to assert the lowest possible limit on cosmic neutrino flux. The bin-specific optimization of the linear discriminant intercept R_i is described next.

4.2.1 Optimization of the R_i Values

The optimization of the bin-specific R_i is based on an assumption that our 10% sample contains only *background*. This seemingly bold assertion is based on our expectation that the entire dataset will contain only ∼1 neutrino event; the 10% sample will thus contain only ∼0.1 neutrino events, which we will take to be zero. Using eq. (4.1), we calculate the value of R for every event in our sample and histogram the event count vs. R. We expect, as observed by Dailey[33], that for each continent bin there will be a peak value of R, beyond which the number density of events should decrease exponentially vs. R, as in Figure 4.2 top left. Any bin for which this is not the case should be explained if possible, and will in any case be omitted from the analysis. It is therefore necessary to assess each bin with respect to how closely it conforms to an exponential fall-off for each combination of $\{C_1, C_2, m\}$. To do this, take our histogram of 10%-sample events vs. R and make a two-parameter exponential fit to the descending portion of the histogram for every continent bin, and calculate the log-likelihood value L_b of the fit. Our histogram bin width in R is 0.5; in order to ensure a reasonable fit, we require that the five bins immediately after the peak-value bin be nonzero and decreasing; bins failing this criterion are omitted.

For bins surviving the retention criteria to this point, the fit is performed on the region of the histogram beginning two bins to the right of the peak-value bin, extending through the last nonzero bin; this is shown by the fitted line in the top-left panel of Figure 4.2. Next, we evaluate the quality of the fit by interpreting the fitted exponential function as a PDF of the R values of background in the Healpix bin of interest:

$$PDF_b(x) = y(x) = ae^{-ax} \qquad a > 0. \tag{4.2}$$

We perform a series of one thousand Monte Carlo pseudo-experiments on every bin, each consisting of a set of pseudo-events randomly thrown using eq. (4.2). The number of points thrown is equal to the number of events falling within the fitted range of the original histogram of the 10% sample. We generate the MC events by the inverse CDF method, in which uniformly-distributed random numbers on $[x_{min}, x_{max}]$ are generated and fed to the inverse CDF. Since the PDF is exponential, the normalized CDF is

$$CDF(x) = Y(x) = \int_0^x PDF(x')dx' = 1 - e^{-ax}, \tag{4.3}$$

which inverts to

$$X(y) = \frac{1}{a}\ln(1-y) = \ln(y') \qquad\qquad y' \in [e^{-ax_{max}}, e^{-ax_{min}}] \tag{4.4}$$

The linear discriminant values of the Monte-Carlo events are calculated using eq. (4.1), and a log-likelihood L_b of these values is calculated against the PDF of equation (4.2). Using this distribution of log-likelihoods as a hypothetical distribution, we calculate the p-value of the original log-likelihood L_b with respect to the former (Figure 4.3). If the original fit is a reasonable representation of the background, then this p-value should not fall near zero. Bins with p-values less than 0.05 are therefore omitted from the analysis.

For bins that deliver an acceptable p-value, we now optimize the rotated cut y-intercept R_i by analyzing the effect of the cut on the predicted background and on predicted neutrinos. Our background is obtained by integrating our exponential background PDF of equation (4.2) from the R_i to infinity:

$$b = a \int_{R_i}^{\infty} e^{-ax} dx \tag{4.5}$$

For an estimate of neutrinos, we sum the histogram of weighted simulation events from `icemc` (Figure 4.1), over the same limits (abusing notation for simplicity):

$$S = \sum_{R_i}^{\infty} S_i \tag{4.6}$$

Using this expectation of b background events and s neutrinos, the likelihood of detecting n events is

$$L(n|s) = \frac{(s+b)^n e^{-(s+b)}}{n!}. \tag{4.7}$$

Recalling our assumption that the 10% sample will contain zero signal, we take $b = n$. We then optimize the sensitivity of our cut by maximizing the ratio S/S_{up}. S is the number of simulated neutrinos passing the cut, from eq. (4.6). S_{up} is the lowest S that is excluded at 90% confidence if the number of observed events is equal to the expected background; that is,

$$\frac{\int_0^{S_{up}} L(b|s)ds}{\int_0^{\infty} L(b|s)ds} = \frac{\int_0^{S_{up}} (s+b)^b e^{-(s+b)} ds}{\int_0^{\infty} (s+b)^b e^{-(s+b)} ds} = U(S_{up}; b) = 1 - \alpha, \tag{4.8}$$

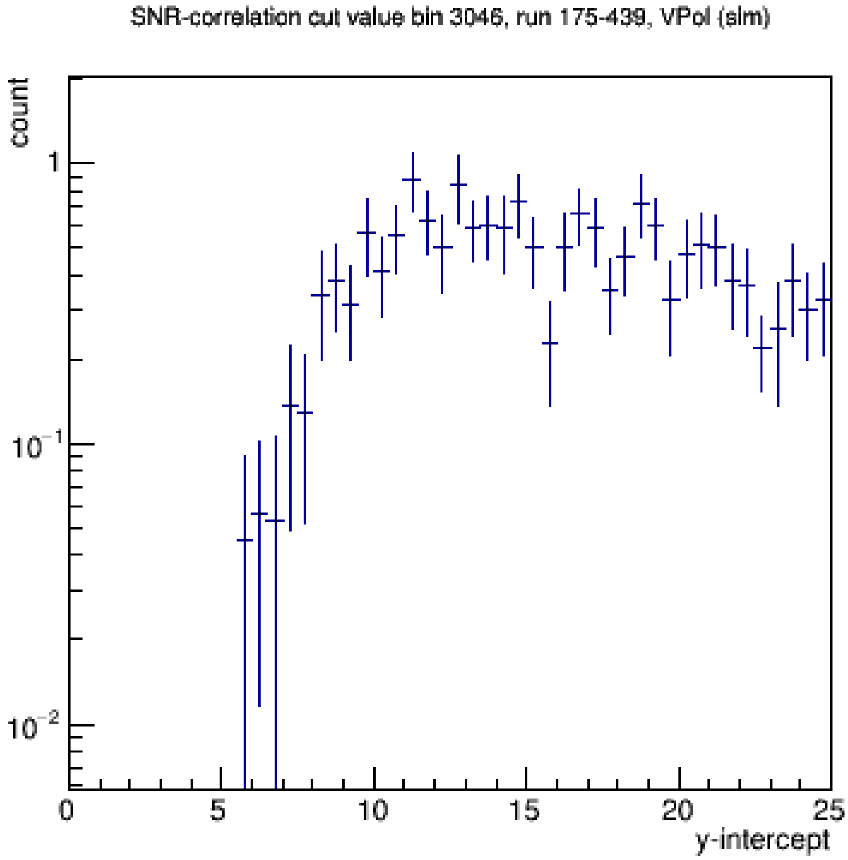

Figure 4.1: Distribution of simulated events S vs linear discriminant value R_i.

with $\alpha = 1.0 - 90\% = 0.1$ in our case. The integrals in (4.8) are calculated easily using a transformation shown in Appendix B. The value of S_{up} that satisfies (4.8) is then determined by an iterative process. Example plots of S and S_{up} are shown in the bottom-left panel of Figure 4.2. S/S_{up} is then optimized with respect to the cut intercept value as shown in Figure 4.2, top-right panel, and the intercept corresponding to the largest S/S_{up} is used in the final cut. The bottom-right panel of Figure 4.2 shows simulated events and 10% sample events with respect to the linear discriminant components, i.e., the correlation peak and SNR. The optimized linear discriminant is also shown.

4.2.2 Optimizing the Fixed Parameter Set $\{C_1, C_2, m\}$

At this point we have obtained optimizations of the linear discriminant intercept for various values of the common cut parameter set $\{m,\ C_1, C_2\}$, but it remains to be determined which of these parameter sets we should actually use in final analysis. We do this by optimizing the flux limit that can be asserted. For every parameter set, we generate a table like the one shown in Appendix E. Our distribution of simulated neutrino energies follows the Kotera model, but only up to a scaling constant. In order to determine the strength of the limit we can assert from our analysis, we must normalize this distribution. To do this, we optimize the scale factor through an iterative process. Taking the expected background of each bin as the mean μ of a Poisson distribution:

$$P(k) = \frac{e^{-\lambda}\lambda^k}{k!} \tag{4.9}$$

we multiply the count (summed weight) of simulated events passing our optimized cuts by the scale factor; then we calculate the Poisson CDF of that scaled count value.

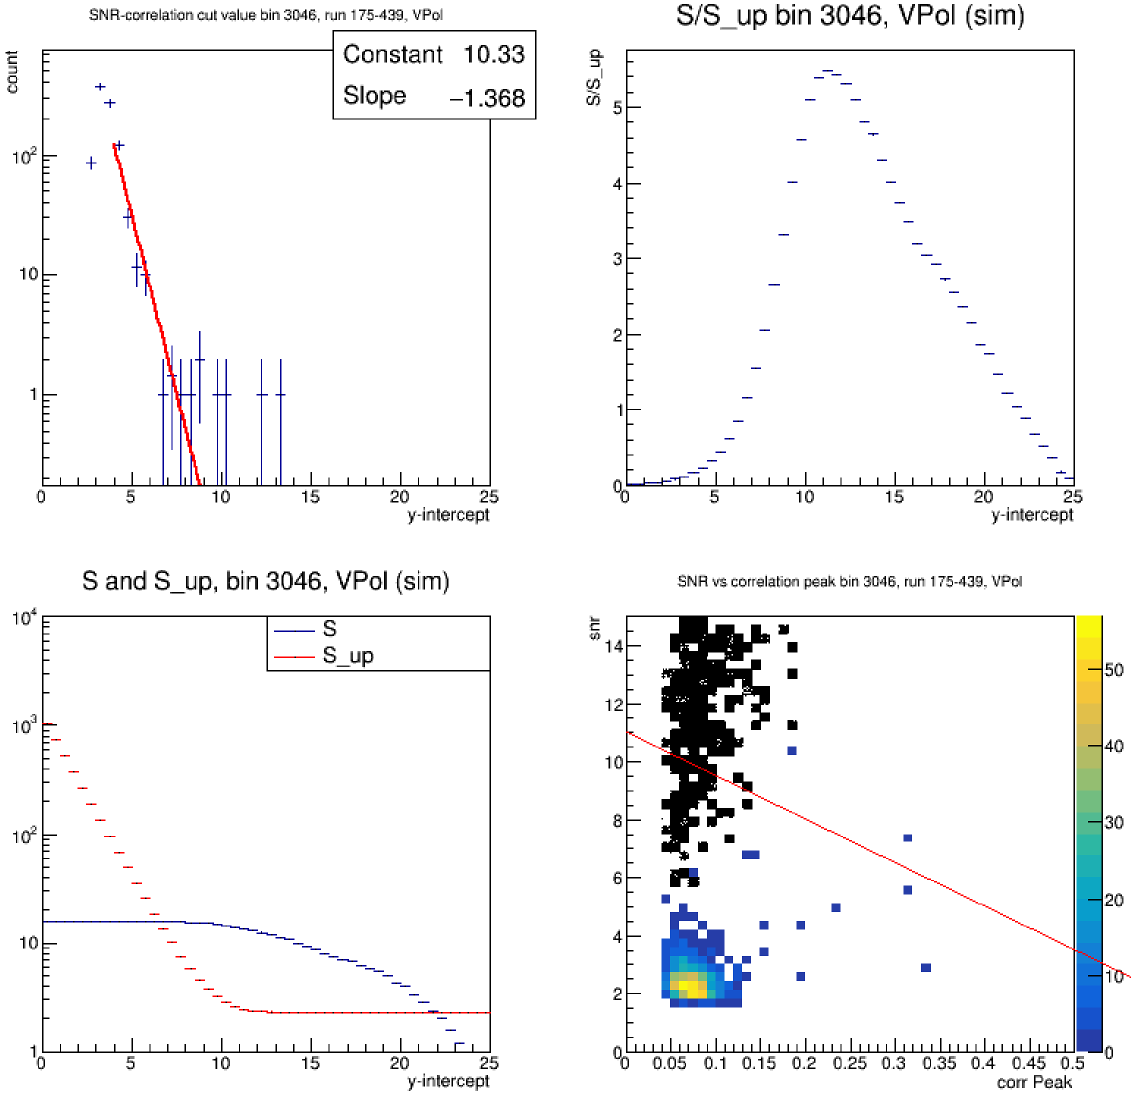

Figure 4.2: Optimization of the stage-3 cuts. This example is for healpix bin 3046 with common stage-3 cut parameters $\{m, C_1, C_2\} = \{-5.0, 5.0°, 0.02\}$. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i, and the value of the optimization parameter denominator S_up. Top right: The value of the optimization figure of merit, S/S_up vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

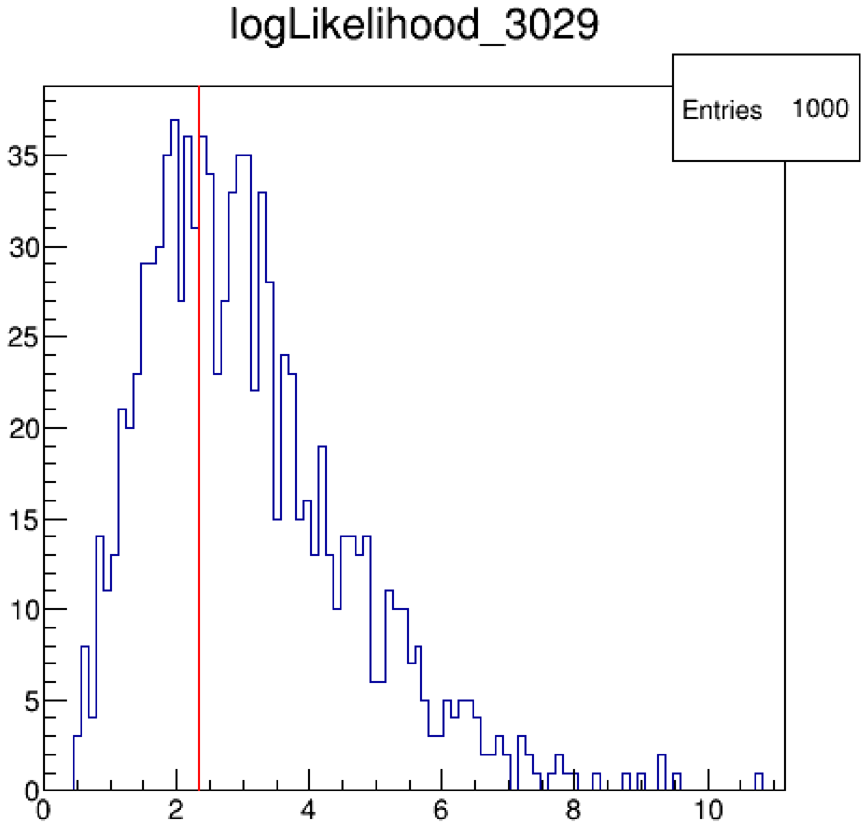

Figure 4.3: Log-likelihood assessment of exponential fit quality. For each Healbix bin that was fitted, the log-likelihood distribution from the pseudo-experiments is used to compute a p-value of the log likelihood L_b from the original exponential fit of the 10% sample. Bins with a p-value less than 0.05 are discarded from the analysis.

Since the events are weighted, the event count κ is a continuous variable; we therefore use a continuous analog to the Poisson distribution:

$$\mathcal{P}(\kappa;\lambda) = \frac{e^{-\lambda}\lambda^{\kappa}}{\Gamma(\kappa+1)}, \tag{4.10}$$

The corresponding CDF for each bin is

$$\mathcal{C}(\kappa;\lambda_i) = 1 - \gamma_r(\kappa+1,\lambda_i) = P_i. \tag{4.11}$$

γ_r represents the lower incomplete Gamma function; an explanation of equations (4.10) and (4.11) is given in Appendix C.

Having calculated the Poisson CDF for all bins, we multiply the CDF values of all N bins together to obtain a combined probability.

$$P_{total} = \prod_{i=1}^{N} P_i \tag{4.12}$$

We aim to assert a limit with 90 percent confidence, so we optimize our scale factor such that equation (4.12) delivers a value of about 10%. At this point, we have chosen a scale factor for every parameter set $\{m,\, C_1{,}C_2\}$ that we tried. To complete our optimization, we choose the parameter set that optimized to the *lowest* scale factor, and thus will allow us to assert the strongest limit.

4.2.3 Post-optimization bin rejection

Once the circular polarization and linear discriminant cut parameters have been optimized, we will apply a few additional criteria in order to Healpix bins that defy analysis due to low sensitivity or high background. These criteria are the basis of the bin rejection cut listed in Section 3.6.3. First, healpix bins with large numbers of background events (from the 10% sample) passing the optimized cuts are omitted. Next, recalling the exponential fit made on the linear discriminant values of the 10% sample, Healpix bins with very few weighted events in the fit range are omitted from the analysis. A threshold value is set and bins with a total event weight below this threshold are cut. Finally, bins with low cumulative sensitivity are omitted; the bins are sorted in ascending order of their sensitivity, that is, the number of weighted simulated events that would pass the all cuts, including the optimized linear discriminant cut. Starting with the lowest sensitivity bins, we accumulate the total sensitivity until a threshold value is reached. Bins that were accumulated before the threshold was reached are omitted. It is also to desirable retain as many simulated events as practical after bin rejection: to measure this, we consider the ratio of the total number of simulated events in all bins to the number in retained bins. We hope for a simulated event retention fraction of about 50%.

* * * * * * *

At this point, all of the cut parameters will have have been determined. In the next chapter, we list the actual cut parameters chosen and discuss the results of running the analysis on the 10-percent sample. This will prepare us to run the analysis on the blinded 90-percent dataset.

Chapter 5

Data Analysis

In the previous chapter, we described the analysis process, including the cuts to be applied during analysis. In this chapter, we show the actual cut parameter values used, as well as tables listing the cuts applied to the ANITA-III data, and the number of events failing each cut.

5.1 Analysis Cut Parameter Values

The parameter values used for the stage 1 cuts are shown in table 5.1

We set the Hilbert peak and correlation peak cut thresholds by choosing the largest reasonable value that would not reject WAIS pulses (Figure 5.1). The parameter values used for the stage 2 cuts are shown in table 5.2. The stage 3 analysis cuts were chosen by applying the iterative method described in Section 3.6.3. The

Analysis Cut Parameter Values - Stage 1

Solar reflection cut: Events localizing to within **5.0**° of the payload-coordinate solar reflection position, calculated using eq. (3.8), are cut.

Reconstruction to continent cut: Events are cut if the ray-trace of the payload-coordinate localization peak fails to intersect the continent. There is no threshold or paramater value for this cut.

Elevation angle cut: Events localizing to a payload-coordinate elevation angle of less than **-35**° or greater than **-6**° were cut.

Triggering ϕ-sector direction cut: Events localizing more than one ϕ-sector width (**22.5**°) away from the nearest triggering ϕ-sector were cut.

Calibration pulse cuts: There are no parameter values for the calibration pulse cuts; cal-pulse events from WAIS and LDB are reliably cut using preexisting built-in functions in `AnitaTools`.

Table 5.1: Cut parameter values used in the stage 1 analysis cuts.

Analysis Cut Parameter Values - Stage 2

2nd-to-1st peak ratio cut: Events in which the ratio of the second highest peak to the first highest peak in the correlation map exceeded **0.9** were cut.

Correlation peak value cut: Events in which the peak value in the correlation map fell below **0.04** were cut.

Hilbert envelope peak value cut: Events in which the peak value of the Hilbert envelope fell below **25** were cut.

Table 5.2: Cut parameter values used in the stage 3 analysis cuts.

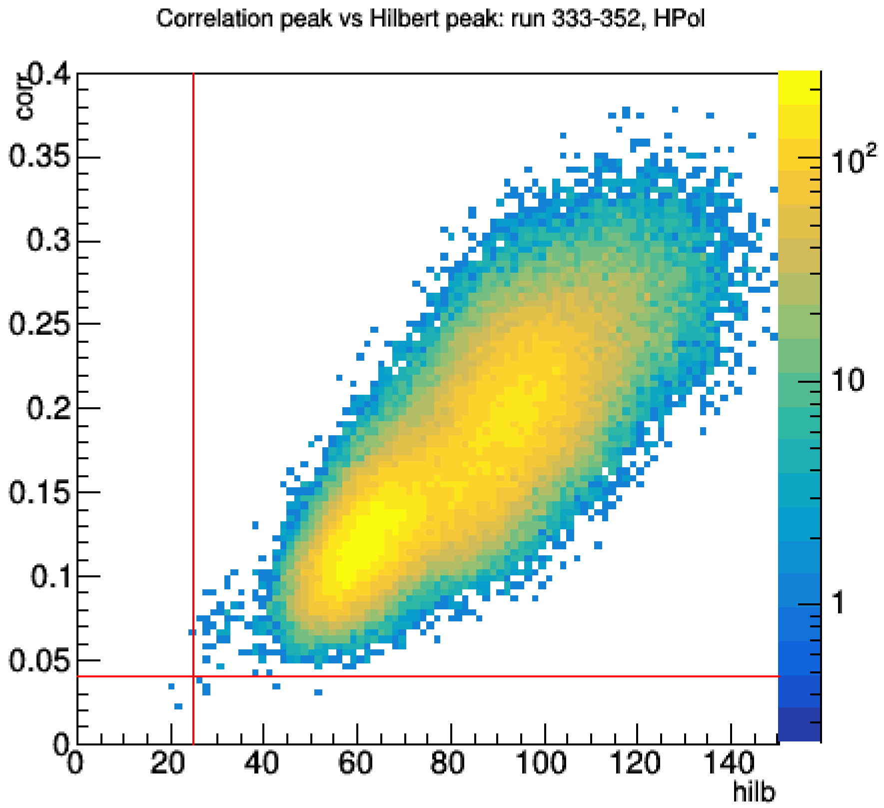

Figure 5.1: Plot of correlation peak and Hilbert peak values for WAIS calibration pulses. The cuts on these values were set as high as possible while rejecting very few WAIS pulses.

Analysis Cut Parameter Values - Stage 3

Circular polarization peak separation cut: Event for which the separation between the correlation peaks in left- and right-handed circular polarization (parameter C_1) exceeded **46°** were cut.

Circular polarization peak strength cut: Events in which the lesser circular polarization correlation peak strength, calculated in the proximity of the linear polarization peak (parameter C_2), was less than **0.015** were cut.

Linear discriminant cut: The linear discriminant slope (parameter m) chosen was **-6.0**. The intercept values (parameter R_i) were chosen individually for the Healpix bins that were included in the analysis and are shown in the table E.1.

Table 5.3: Cut parameter values used in the stage 3 analysis cuts.

actual values used for the stage 3 cuts are shown in table 5.3. In the parlance used in the text, this represents an optimized parameter set $\{C_1, C_2, m\} = \{46°, 0.015, -6.0\}$. The total number of simulated events surviving all cuts except the linear discriminant was 111.6; after the healpix bin rejections, 63.4 such events remain, for a retention ratio of 0.56. This is consistent with our desire to retain more than half of these events.

5.2 Results of the Analysis on the 10% Sample

We ran the analysis on the 10-percent sample of events from the ANITA-III flight. Quality cuts were applied prior to analysis to remove unusable events. The results of the quality cuts are shown in table 5.4.

Events that failed quality cuts were not included in the analysis. Figure 5.2 shows the

```
Quality cuts:
 polarization H
   total events processed        6316237
  cut-id       description               as first cut        as ordered cut        as last cut
                                       number  fraction     number  fraction     number  fraction
   0                 trigger type           0   0.00000          0   0.00000          0   0.00000
   1     L3 triggering phisectors     3217924   0.50947    3217924   0.50947      97764   0.01548
   2                payload blast          24   0.00000         24   0.00000          1   0.00000
   3          waveform saturation        2655   0.00042       1751   0.00028        294   0.00005
   4           waveform dc offset           0   0.00000          0   0.00000          0   0.00000
   5              short waveforms          77   0.00001         36   0.00001         35   0.00001
   6                  nadir noise       73234   0.01159      39542   0.00626      39542   0.00626
   7                   no trigger     3104268   0.49147          0   0.00000          0   0.00000
  total events cut:                                        3259277   0.51602
  surviving events:                                        3056960   0.48398
 polarization V
   total events processed        6316237
  cut-id       description               as first cut        as ordered cut        as last cut
                                       number  fraction     number  fraction     number  fraction
   0                 trigger type           0   0.00000          0   0.00000          0   0.00000
   1     L3 triggering phisectors     3062707   0.48489    3062707   0.48489      96930   0.01535
   2                payload blast         216   0.00003        216   0.00003         11   0.00000
   3          waveform saturation        2132   0.00034       1861   0.00029        512   0.00008
   4           waveform dc offset           0   0.00000          0   0.00000          0   0.00000
   5              short waveforms          77   0.00001         41   0.00001         41   0.00001
   6                  nadir noise       62292   0.00986      42901   0.00679      42901   0.00679
   7                   no trigger     2957773   0.46828          0   0.00000          0   0.00000
  total events cut:                                        3107726   0.49202
  surviving events:                                        3208511   0.50798
```

Table 5.4: Quality cut results. The number and fraction of events failing each cut is shown as if the cut were (1) made first; (2) made in the order shown; (3) made last.

```
Analysis cuts (stage 1):
 polarization H
   total events processed        3057121
  cut-id       description              as first cut        as ordered cut        as last cut
                                      number  fraction      number  fraction      number  fraction
   2            solar reflection      66370   0.02171      66370   0.02171       20014   0.00655
   6     reconstruct to continent    2362321   0.77273    2362295   0.77272        1928   0.00063
  11              elevation angle    2533149   0.82861     172600   0.05646       46405   0.01518
  29   hardware trigger direction    2193907   0.71764     316375   0.10349      316361   0.10348
  18            calibration pulse      11830   0.00387      11244   0.00368       11244   0.00368
  total events cut:                                       2928884   0.95805
  surviving events:                                        128237   0.04195
 polarization V
   total events processed        3208510
  cut-id       description              as first cut        as ordered cut        as last cut
                                      number  fraction      number  fraction      number  fraction
   2            solar reflection      28307   0.00882      28307   0.00882       10207   0.00318
   6     reconstruct to continent    2565087   0.79946    2565062   0.79946        1563   0.00049
  11              elevation angle    2780598   0.86663     216259   0.06740       53821   0.01677
  29   hardware trigger direction    2314227   0.72128     280554   0.08744      280545   0.08744
  18            calibration pulse        631   0.00020        381   0.00012         381   0.00012
  total events cut:                                       3090563   0.96324
  surviving events:                                        117947   0.03676
```

Table 5.5: Stage 1 analysis cut results. The number and fraction of events failing each cut is shown as if the cut were (1) made first; (2) made in the order shown; (3) made last.

Healpix binning of the Antarctic continent, with the total weights of events allocated to in each bin, before any analysis cuts were applied.

After quality cuts, the first stage of analysis cuts was applied, consisting of simple cuts expected to remove a large fraction of events. The results of the stage 1 analysis cuts are shown in Table 5.5.

The results of the second and third stage analysis cuts are shown in table 5.6.

5.2.1 Listing of 10%-sample Events Passing all Cuts

In the horizontal polarization, 5.3 weighted events from the 10% sample passed all cuts. In the vertical, 3.0 weighted events passed all cuts. The individual events and

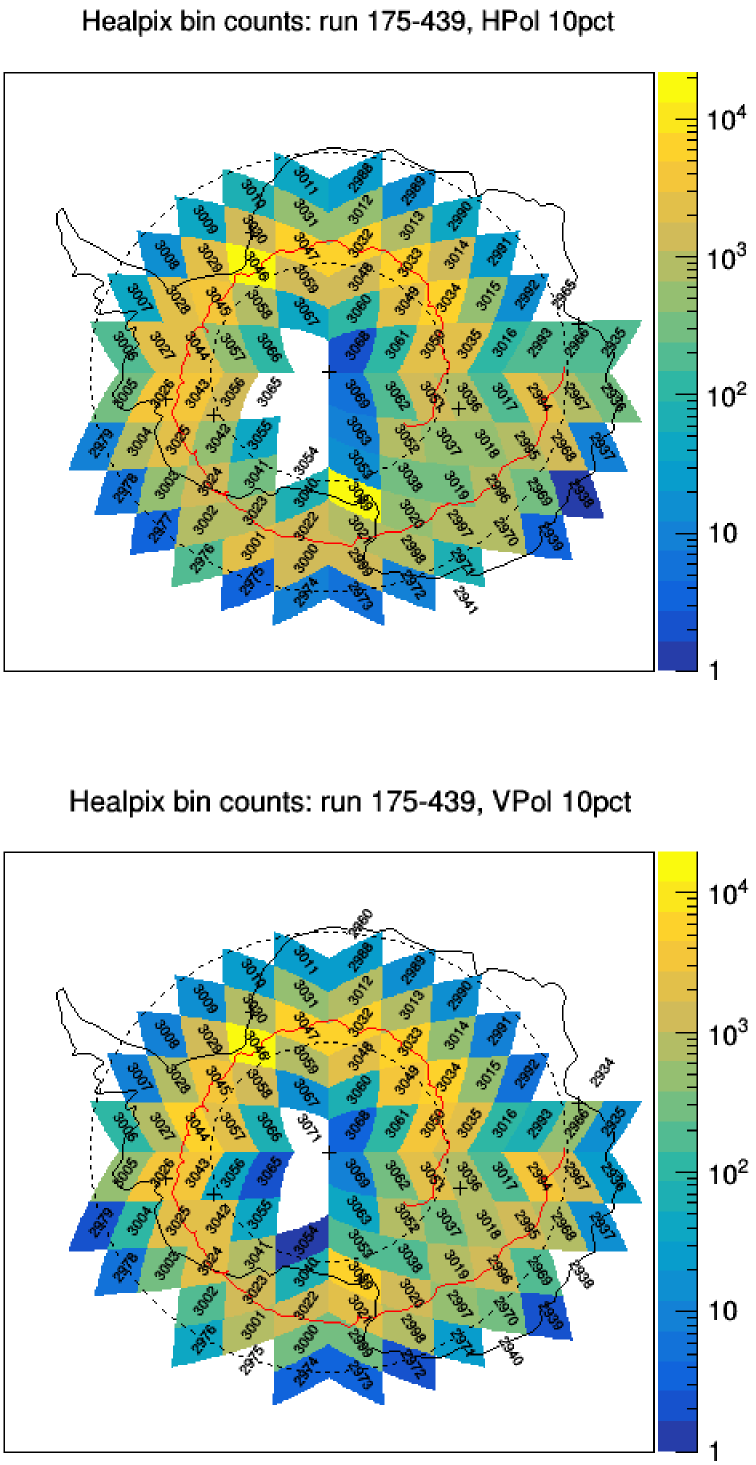

Figure 5.2: Allocation of event weights to bins on continent, before cuts. The numbers in the plot are the bin index numbers from the Healpix tesselation.

```
Analysis cuts (stage 2, 3):
 polarization H
   total events processed       124850.0
               description              as first cut        as ordered cut        as last cut
                                      number  fraction     number  fraction     number  fraction
          2nd/1st peak ratio         31487.0   0.25220    31487.0   0.25220     6002.0   0.04807
            correlation peak         20441.0   0.16372     7770.0   0.06223     1098.0   0.00879
                hilbert peak         19463.0   0.15589     6528.0   0.05229     1899.0   0.01521
        cpol peak separation         53085.0   0.42519    28367.0   0.22721    20947.0   0.16778
          cpol peak strength         30559.0   0.24477     5891.0   0.04718     5891.0   0.04718
        low event bin weight                                581.1   0.00465
        healpix bin rejected                              38436.0   0.30786
     linear discriminant cut                               5784.6   0.04633

  total events cut:                                      124844.7   0.99996
  surviving events:                                           5.3   0.00004

 polarization V
   total events processed       115338.0
               description              as first cut        as ordered cut        as last cut
                                      number  fraction     number  fraction     number  fraction
          2nd/1st peak ratio         41047.0   0.35588    41047.0   0.35588    12935.0   0.11215
            correlation peak         25455.0   0.22070    10046.0   0.08710      959.0   0.00831
                hilbert peak         29496.0   0.25574     8630.0   0.07482     2630.0   0.02280
        cpol peak separation         46650.0   0.40446    19772.0   0.17143    16108.0   0.13966
          cpol peak strength         29021.0   0.25162     4000.0   0.03468     4000.0   0.03468
        low event bin weight                                327.6   0.00284
        healpix bin rejected                              21073.9   0.18271
     linear discriminant cut                              10438.3   0.09050

  total events cut:                                      115335.0   0.99997
  surviving events:                                           3.0   0.00003
```

Table 5.6: Stage 2 and 3 analysis cut results. The number and fraction of events failing each cut is shown as if the cut were (1) made first; (2) made in the order shown; (3) made last.

```
Events surviving all cuts - 10% sample
HPol
 Event #    Run    bin    weight
 47429857   306    3058   0.8067
 51700777   320    3027   1.0000
 62273732   357    3024   1.0000
 69843935   383    2998   0.8247
 69969708   383    2998   1.0000
 70280225   385    2998   0.6432

VPol
 Event #    Run    bin    weight
 17715967   187    3052   1.0000
 35083936   254    3048   1.0000
 36478826   265    3048   1.0000
```

Table 5.7: Events from the 10-percent surviving all cuts.

their allocated weights are listed in Table 5.7

We expect these numbers to be smaller. Recalling that in the vertical polarization we expect a weighted neutrino event count ∼1.0 from the full flight dataset, the results of the 10% analysis are expected to deliver ∼0.1 events passing all cuts. In the horizontal polarization, we expect ∼0 neutrino events. The 3.0 VPol events and 5.3 HPol events surviving all cuts thus require additional consideration. We first examine the localizations of the surviving events, shown in figure 5.3. The colors represent the weighted event allocation to the Healpix bins. The red crosses represent the localizations of the events on the continent. At this point we discuss individually the events from the 10-percent sample that passed all cuts. Suggestions are made regarding possible causes and remedies. We acknowledge, however, that these arguments are anecdotal in nature and *do not* stand alone as bases for changes to our methods or cut parameters. They serve instead as motivations for further

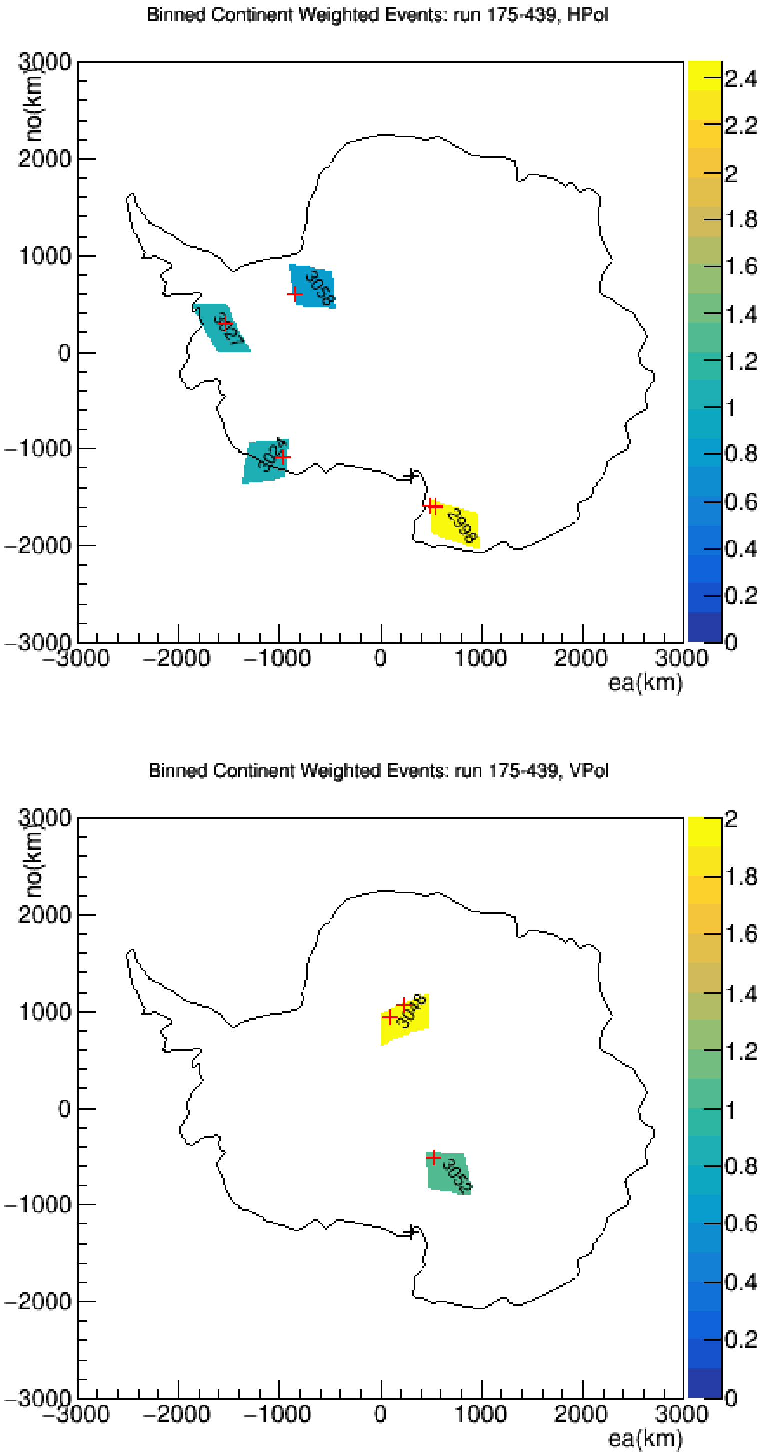

Figure 5.3: Events from the 10-percent sample surviving all cuts.

consideration.

5.2.2 Discussion of individual events passing the cuts

First, we consider bin 2998, containing a total event weight 2.45 in HPol comprised of three contributing events. All three originate about 200 km from McMurdo station. We consider the possibility that these are misreconstructions of anthropogenic pulses from McMurdo, but this is not obviously the case. Figures 5.4, 5.5, and 5.6 show the interferometry correlation maps of these three events. It is clear in from the figures that all three events are strongly impulsive; note that their correlation peaks are considerably stronger than the WAIS pulse example in Figure 3.3. Events of this strength are very unlikely to misreconstruct, so the localizations are probably accurate. While they do not originate directly from McMurdo, their source is close enough to McMurdo that human activity there is highly likely, and that a strong anthropogenic pulses from that location is highly plausible. We will consider omitting bin 2998 from the analysis.

Event 47429857 in bin 3058 is strong in horizontal polarization and in left-circular polarization (Figures 5.7 and 5.8), which is consistent with a satellite reflection. Furthermore, it appears to originate from the Filchner ice shelf, from which a strong reflection from the flat sea ice seems highly plausible. Further investigation into this event may reveal why it passed the circular polarization strength cut.

Event 51700777 in bin 3027 is HPol-dominated but weak in comparison to the events previously discussed. On initial examination it appears odd that this passed cuts, but the lower panel of Figure 5.9 shows that its Hilbert peak narrowly exceeds the

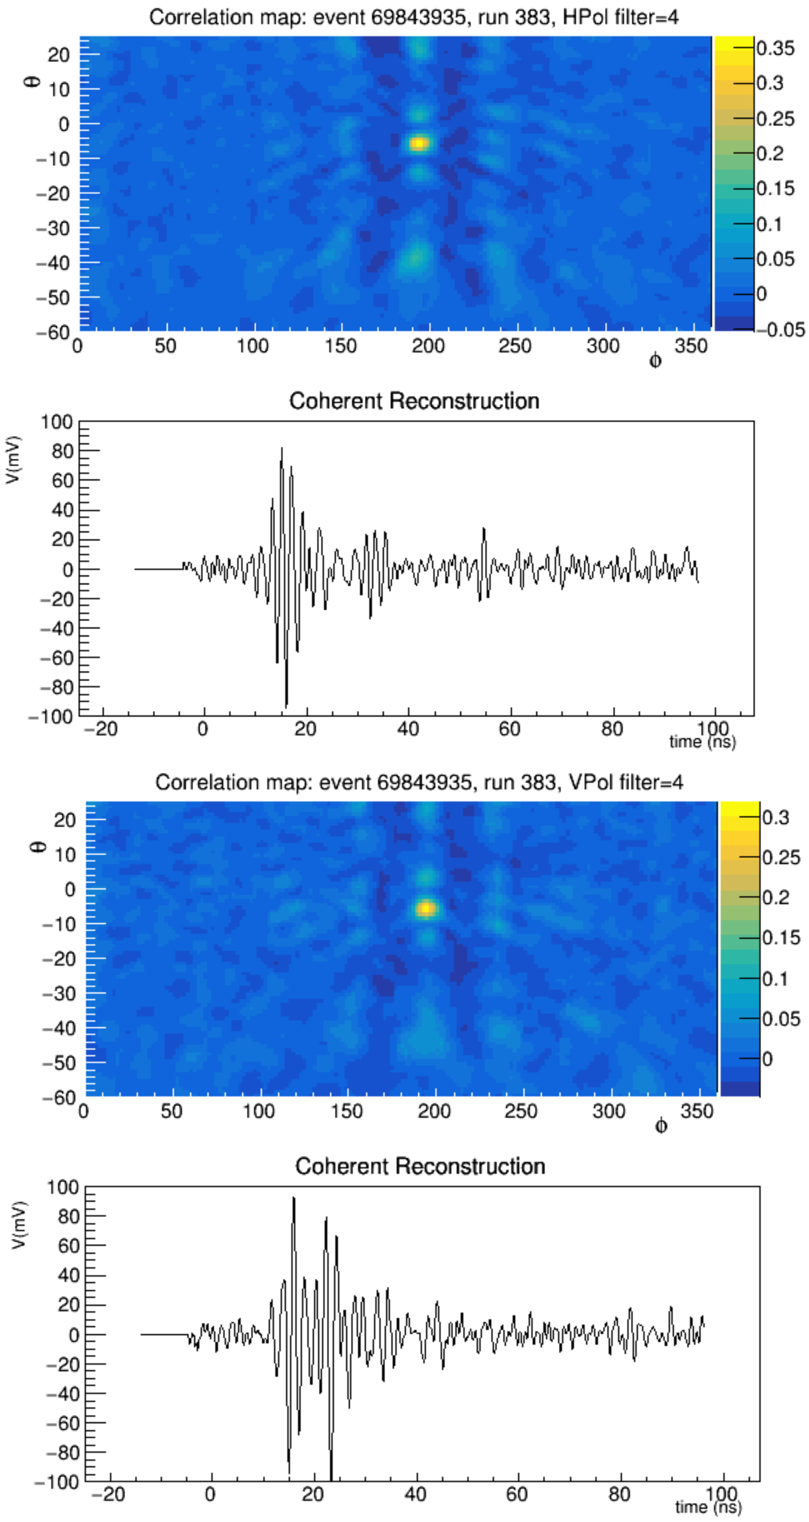

Figure 5.4: Reconstruction of event 69843935 in linear polarizations.

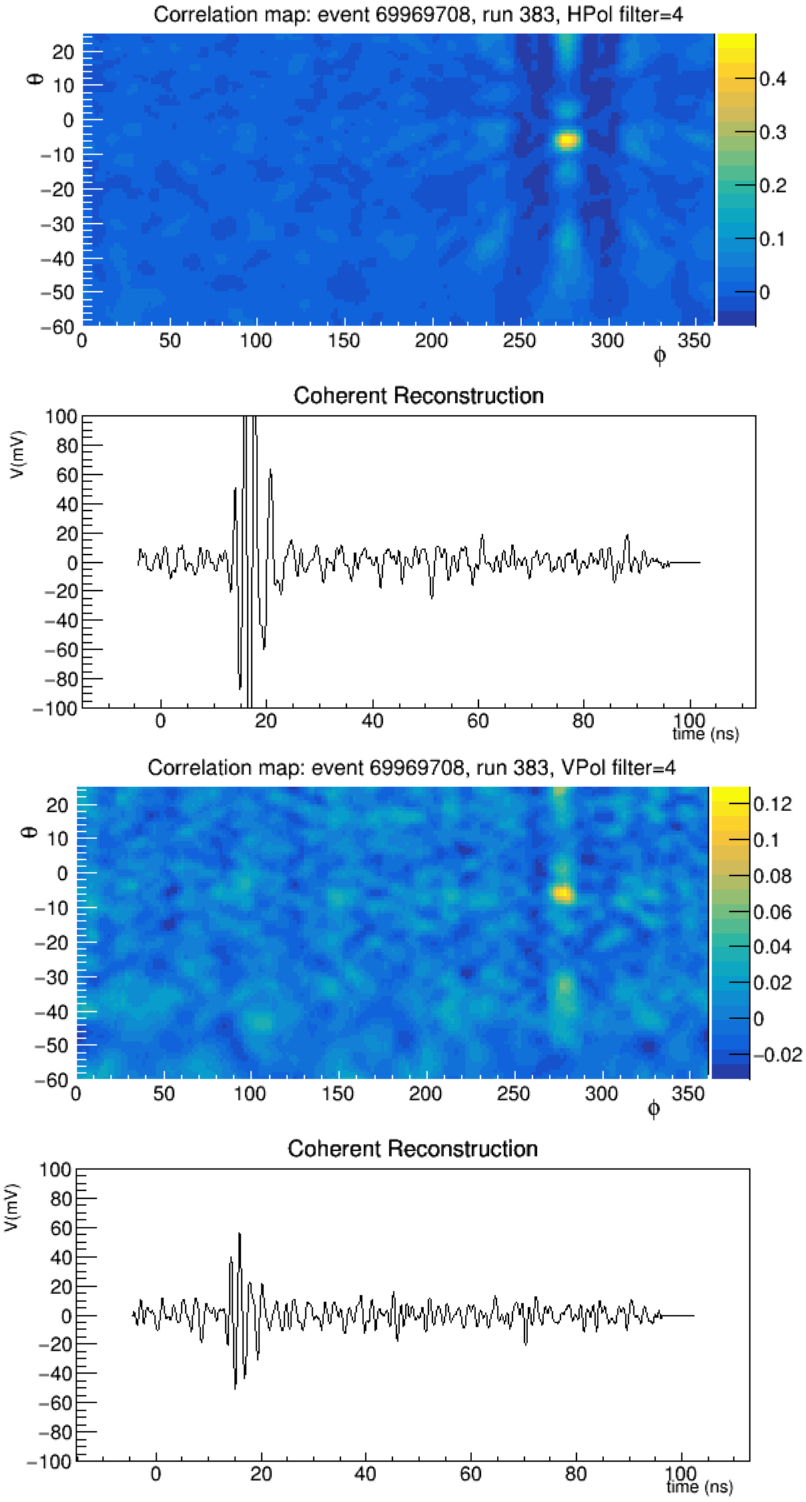
Correlation map: event 69969708, run 383, HPol filter=4

Coherent Reconstruction

Correlation map: event 69969708, run 383, VPol filter=4

Coherent Reconstruction

Figure 5.5: Reconstruction of event 69969708 in linear polarizations.

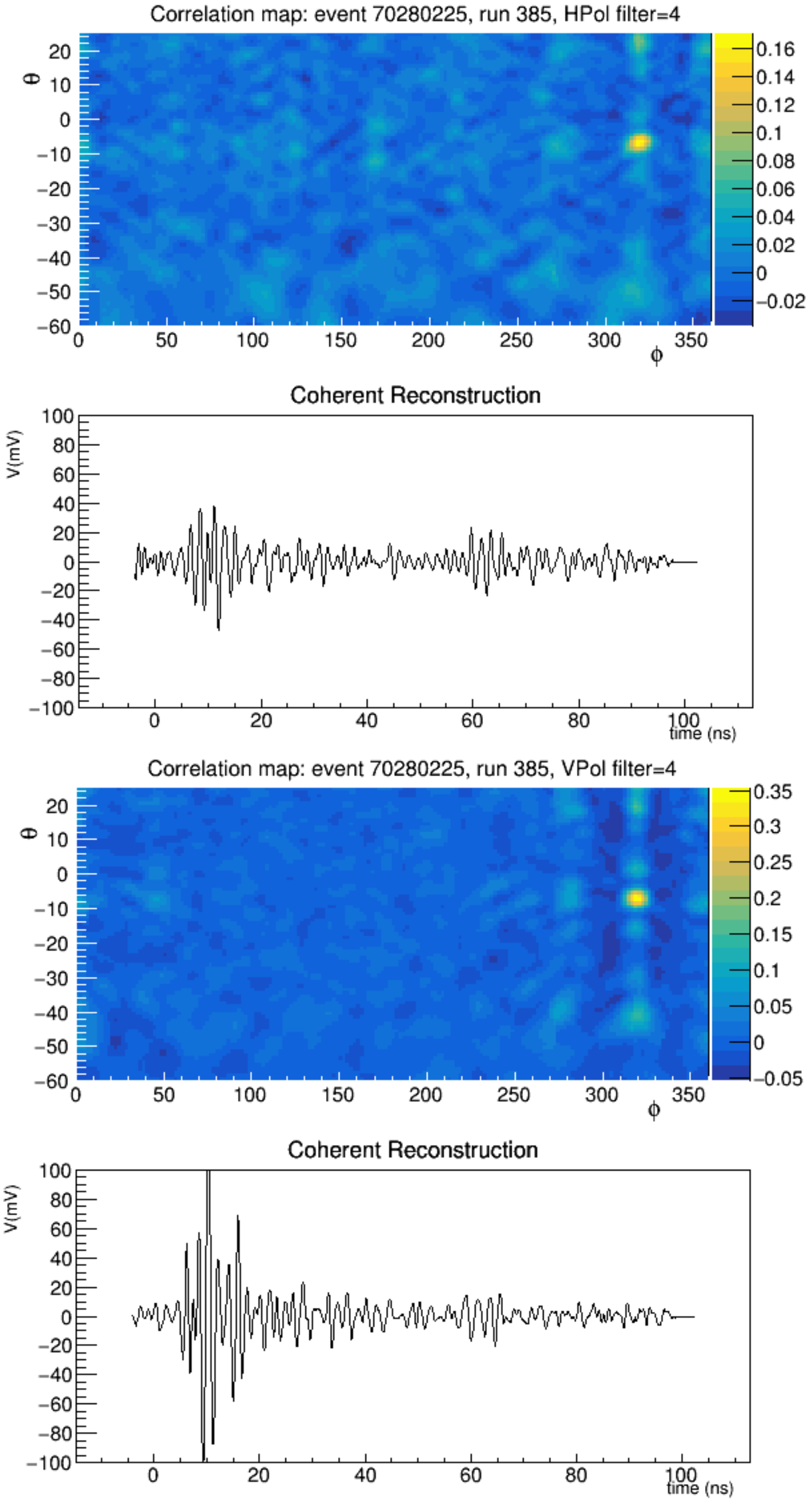

Figure 5.6: Reconstruction of event 70280225 in linear polarizations.

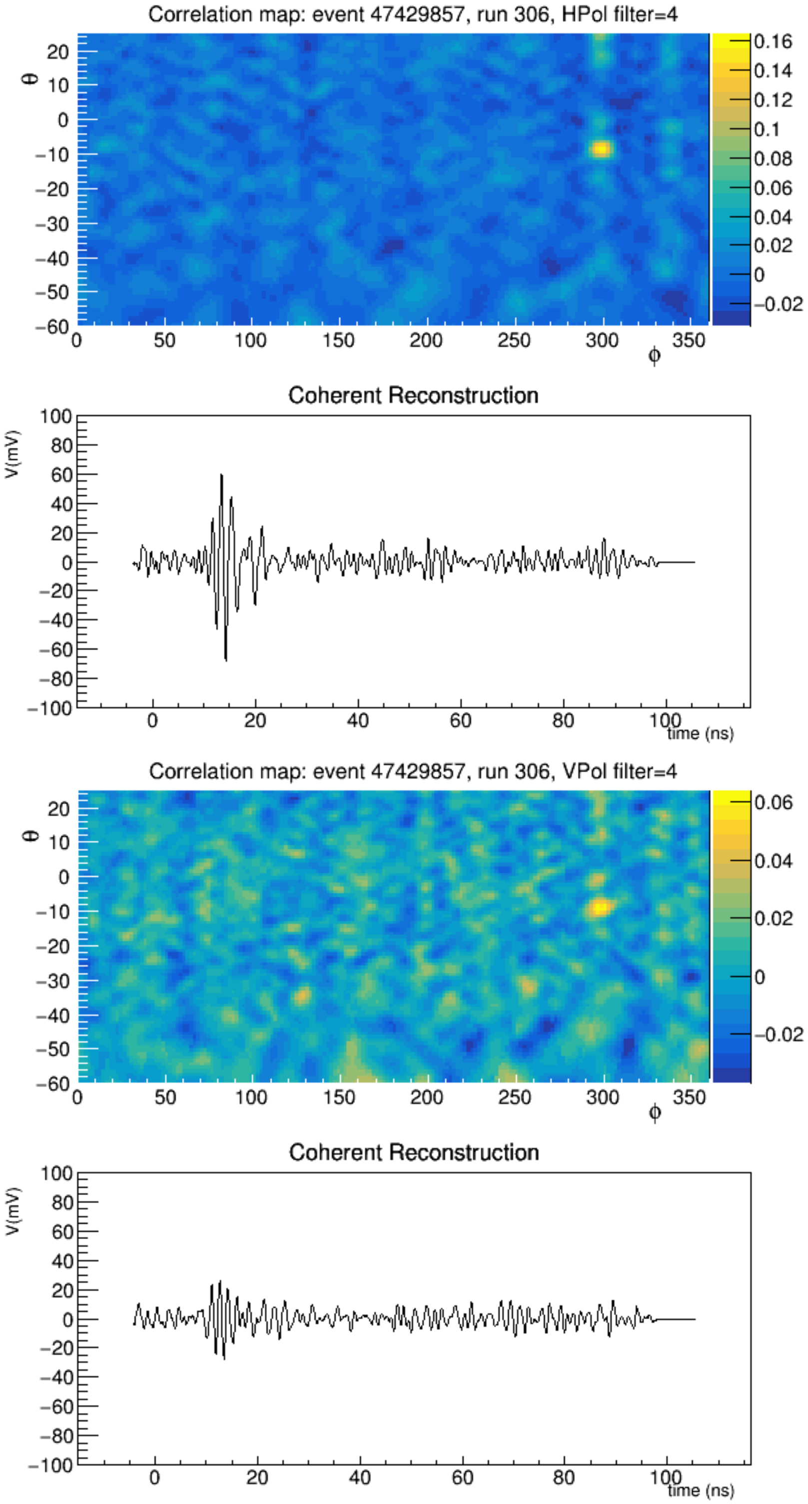

Figure 5.7: Reconstruction of event 47429857 in HPol. The hilbert envelope is shown in the lower panel because its peak narrowly exceeds the Hilbert cut threshold of 0.25

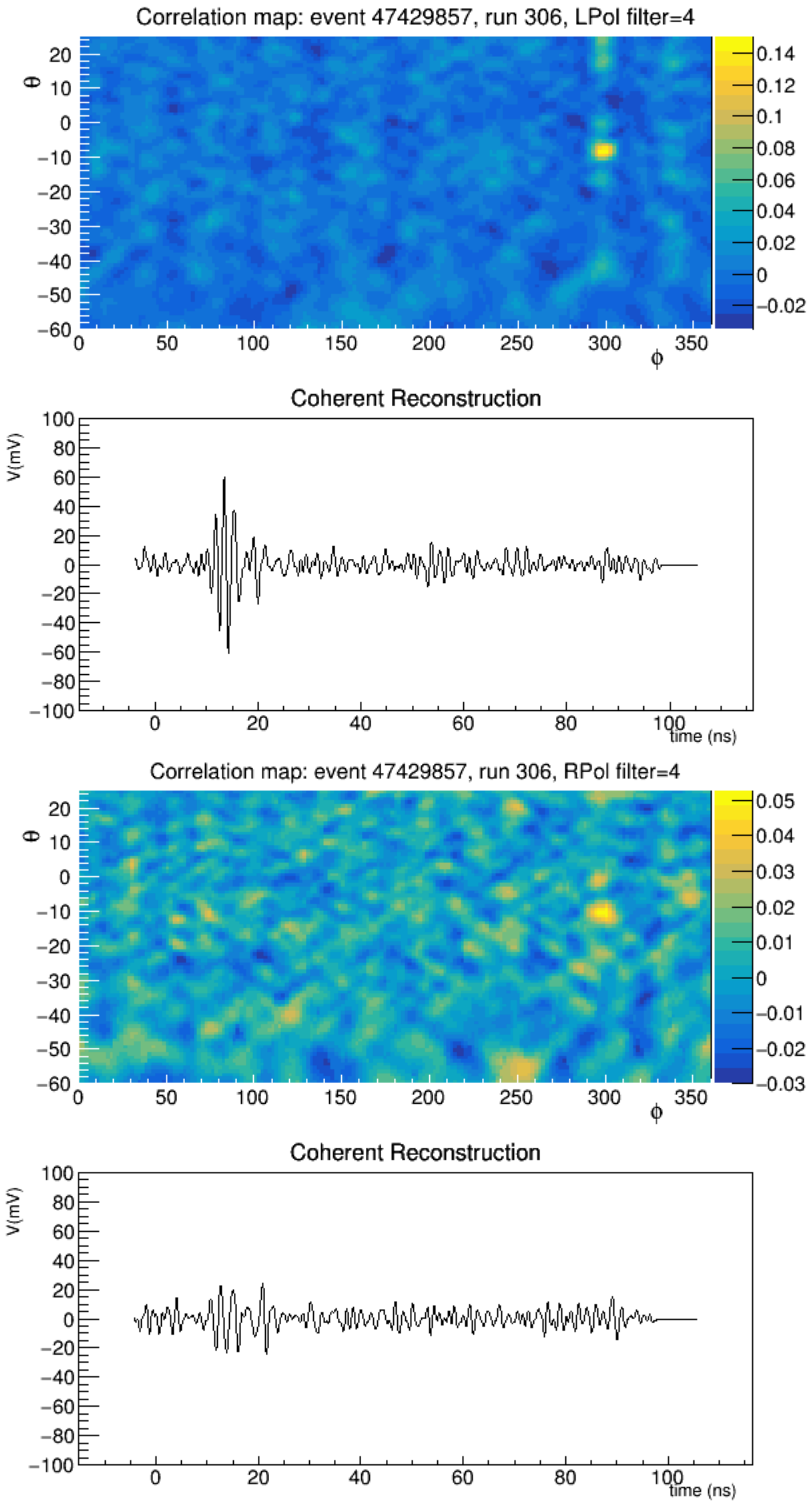

Figure 5.8: Reconstruction of event 47429857 in circular polarizations.

Hilbert cut parameter threshold of 0.25. The Hilbert and Correlation peak cuts were both set conservatively to avoid cutting too many WAIS pulse events; a tightening of the Hilbert cut should probably be considered.

Event 62273732 in bin 3024 is strong in both HPol and LPol, suggesting again the possibility of a satellite reflection. Viewing this event in payload-north coordinates may reveal if it originated from the north, where satellite reflections are common. If so, a cut on certain areas of the payload-north viewfield may be called for.

Event 17715967, bin 3052 is comparitively weak with respect to the other events that passed the cuts, as shown in the top and middle panels of Figure 5.12. However, as in the case of event 51700777, the Hilbert peak exceeds the cut threshold. This may bolster the case for strengthening the Hilbert peak cut.

Event 35083936, bin 3048 (Figure 5.13)is very strong and appears to be heavily CW contaminated. It passes the Hilbert peak cut by a wide margin such that tightening that cut would not help. This suggests that we consider a better mechanism for identifying CW. Filtering appears not to have sufficed in this case, so perhaps a cut can be developed in order to more consistently remove CW-contaminated events.

The last event passing all cuts, event 36478826 in bin 3048, appears anomalous. It appears from figure 5.14 that it should have not even have passed the Hilbert peak cut. Further investigation is necessary to determine if there is an abberation in the plot of the Hilbert envelope or if there exists a bug in the enforcement of the Hilbert peak cut.

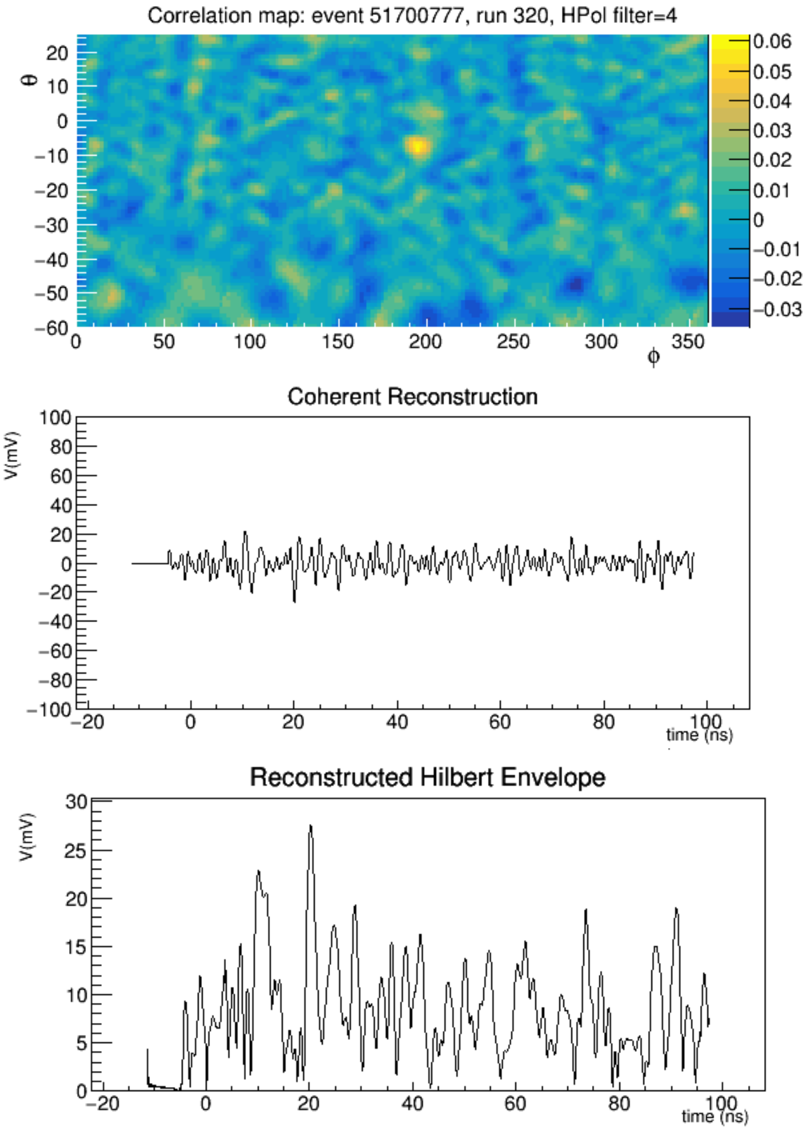

Figure 5.9: Reconstruction of event 51700777 in linear polarizations.

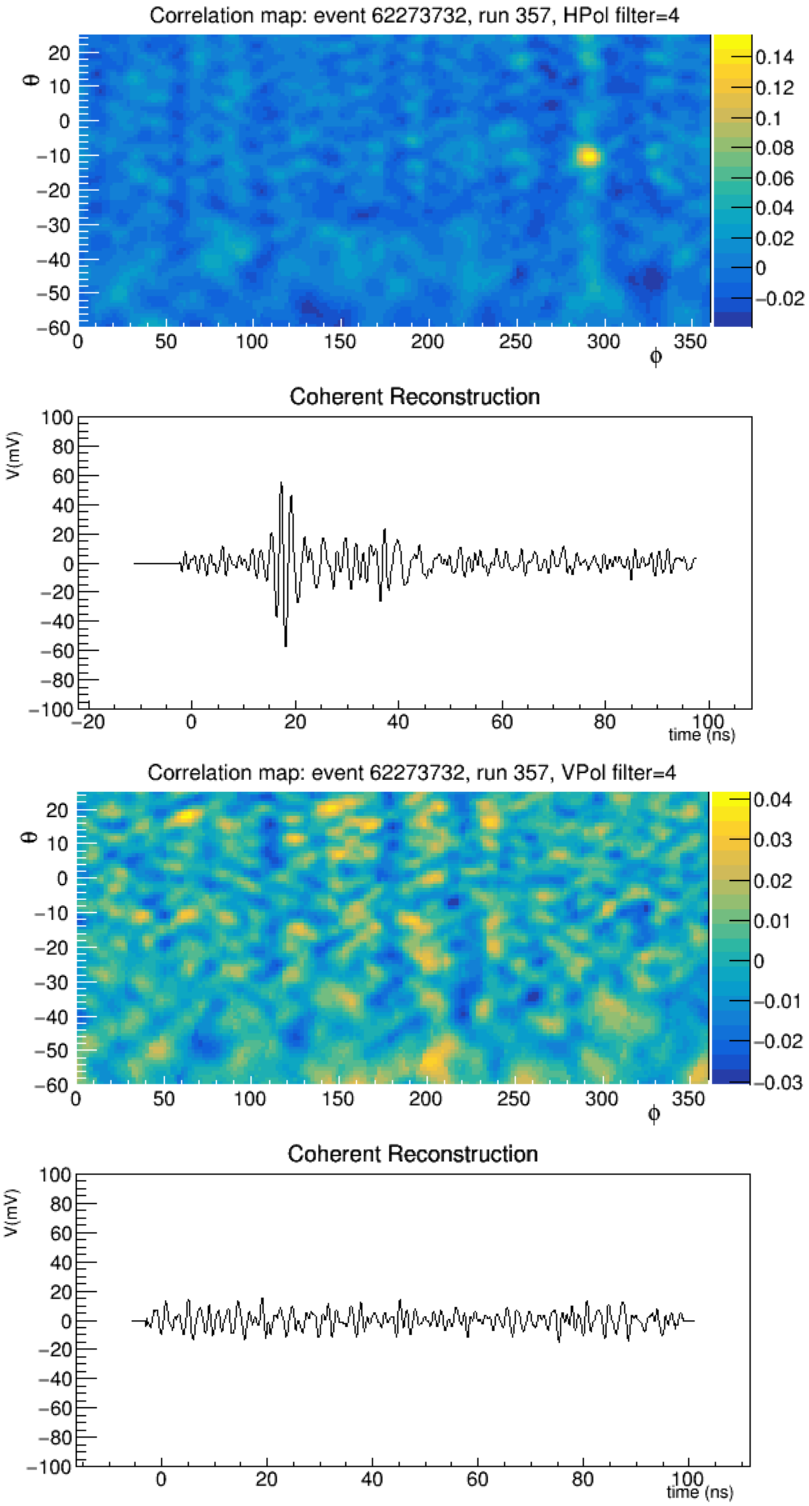

Figure 5.10: Reconstruction of event 62273732 in in linear polarizations.

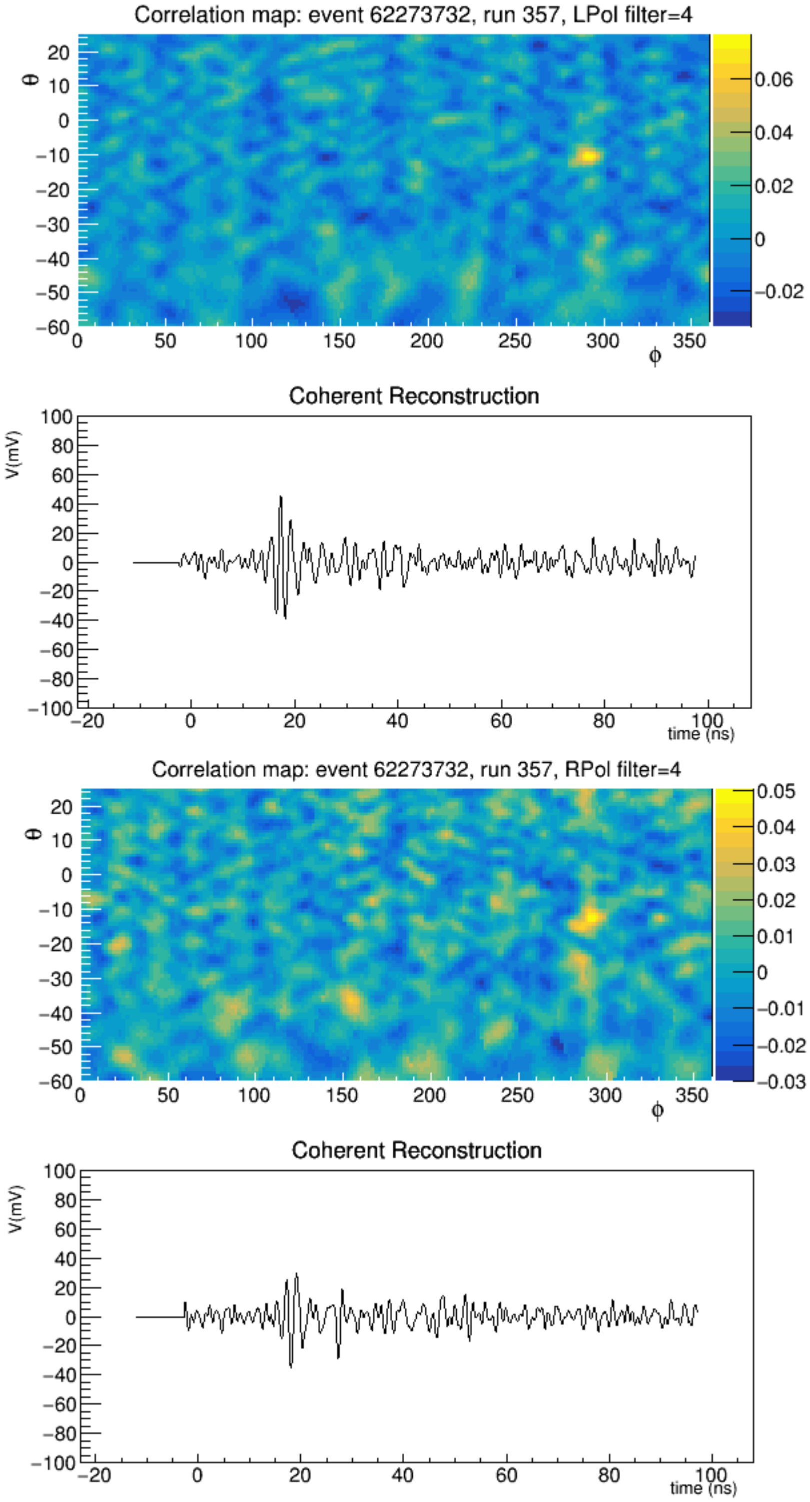

Figure 5.11: Reconstruction of event 62273732 in circular polarizations.

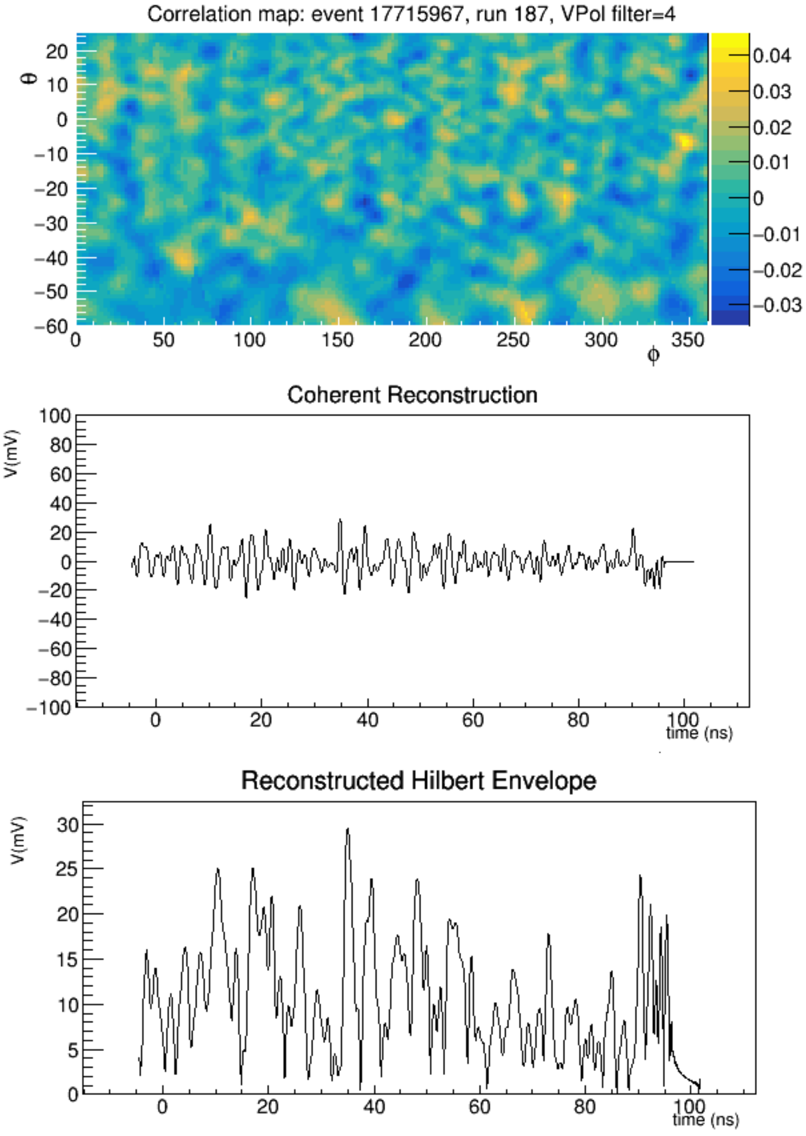

Figure 5.12: Reconstruction of event 17715967 in in linear polarizations.

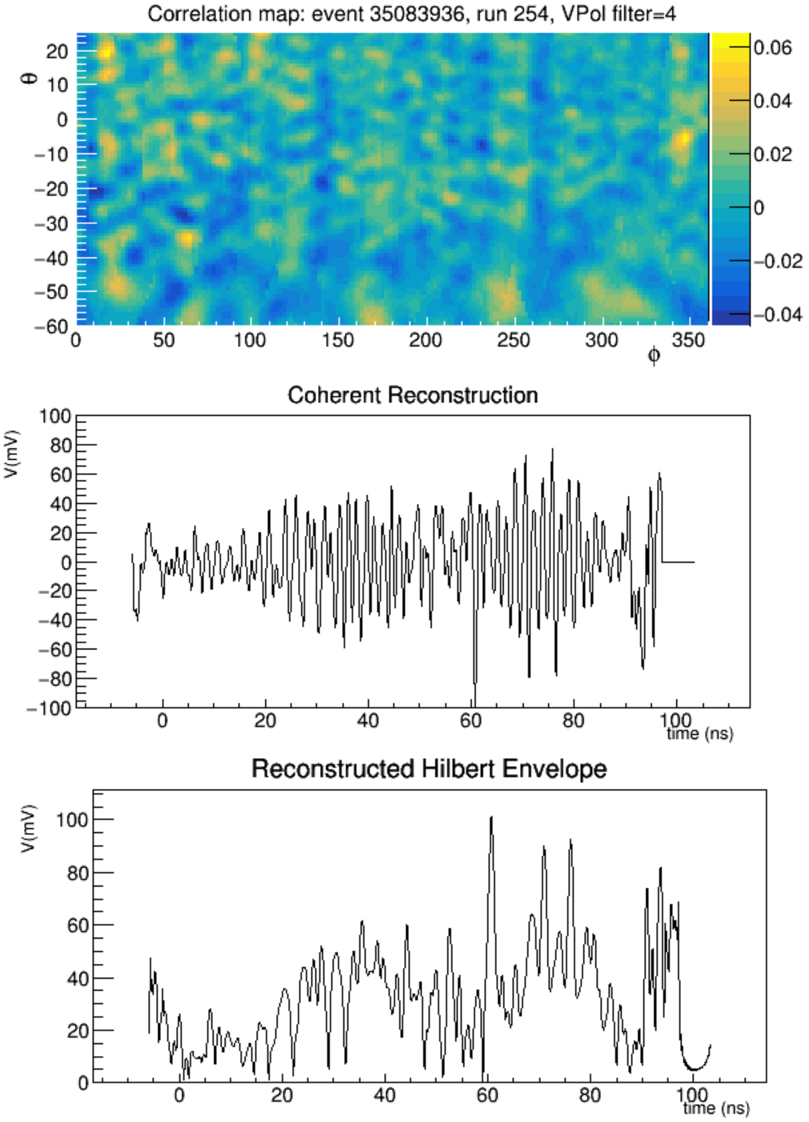

Figure 5.13: Reconstruction of event 35083936 in in linear polarizations.

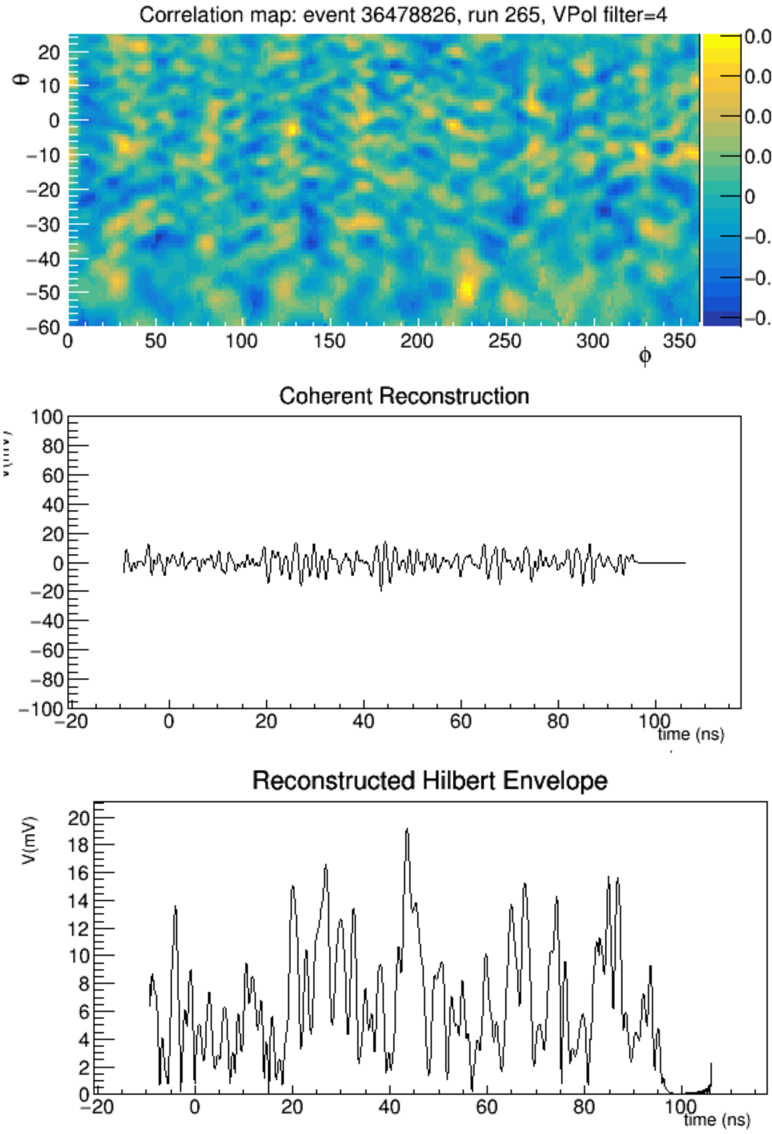

Figure 5.14: Reconstruction of event 36478826 in in linear polarizations.

5.2.3 Additional Comments

The 10% sample used in the analysis run was blinded, and therefore may contain inserted pulse events; this was an error in analysis the procedure. Therefore, the analysis should be rerun on an unblinded 10% sample. In any case, further analysis will be necessary to reconcile the high surviving event counts. The binned optimization of the linear discriminant cut was intended to obviate the need for a clustering algorithm such as was used in previous analyses, perhaps some analysis based on mutual event proximity would be helpful.

Finally, the optimization of the new cuts, that is, the circular polarization and linear discriminant cuts, is not a particularly mature process and can probably be improved. The cuts we chose on circular polarization are certainly not the only possibilities, and others should be considered and tried.

Chapter 6

Discussion and Conclusions

We developed a process for conducting a search for ultra-high energy cosmic neutrinos by analyzing data from the third flight of the Antarctic Impulsive Transient Antenna (ANITA-III). Starting with methods from analyses on the ANITA-II flight, we augmented those methods with the goal of identifying neutrino candidate events and asserting an upper limit on cosmic neutrino flux.

6.1 Enhancements to the analysis process

Our analysis process was based on the processes used in previous analyses, particularly those described in Dailey[33] and Vieregg[36]. We made several significant enhancements to those methods.

6.1.1 Circular Polarization Analysis

Most significantly, we performed analysis using the circular polarization representation of event waveforms. We discovered that high circular polarization content is a strongly correlated to satellite-influenced events. Knowing that Askaryan neutrino signals should be heavily linearly polarized, we developed two cuts based on circular polarization. We first considered the angular separation between the correlation map peaks derived from the left-and right-handed polarizations of the event waveforms because we expect that linearly polarized events should localize to roughly the same point on the sky when analyzed in the two circular polarizations. For our second cut, we localized the event in linear polarizations, then constructed a window around these localization points, imposed it on the circular polarization correlation map, and calculated the maximum correlation values in that window. We cut events for which these values fell below a threshold.

In an analysis of the 10% sample, as an ordered cut, the circular polarization peak separation cut removed 22% of the population of HPol events surviving the stage 1 cuts; in VPol, it removed 17%. The circular polarization strength cut, applied as an ordered cut, removed 4.7% and 3.4% of events surviving the stage 1 cuts in HPol and VPol, respectively. Both these cuts were also strong when considered as the last cut in the sequence with the peak separation cut removing 16.8% and 14.0% of events respectively in HPol and VPol, and the strength cut removing 4.7% and 3.7% in HPol and VPol, respectively.

6.1.2 Event Localization on the Antarctic Continent

Localization of the origination point of events on the continent was important in previous analyses in order to facilitate clustering algorithms for identification and removal of impulsive anthropogenic events. Instead of clustering, our analysis used a method similar to that described in [33], in which we localized events to Healpix bins on the Anatarctic continent, and optimized a linear discriminant individually for each bin. Since the error ellipse of an event can span multiple Healpix bins, careful weighting of events across these bins was essential. The previous analysis treated the error ellipse as a uniform PDF and allocated events to Healpix bins based on the fraction of the error ellipse falling in each bin.

We developed a method that allows flexibility in our choice of PDF. Using a PDF derived from the pointing performance of our interferometry, we made a grid of points around the payload-coordinate localization point, and ray-traced each of these points to the continent, weighting each point with the appropriate PDF value. This method has two advantages: first, it is computationally simple, employing straight-line ray-tracing functions from `AnitaTools`. Second, it allows any PDF to be used. We used a Gaussian PDF, with the standard deviations of calibration pulse pointing error in payload-coordinate azimuth ϕ and elevation θ as the PDF width parameters.

6.1.3 Cut Optimization

We extended the method of optimizing the linear discriminant cut. Dailey[33] chose a slope for the cut common across all Healpix bins, and optimized its y-intercept on individual bins. We, on the other hand, optimized all three of our stage 3 cuts, that is, the linear discriminant as well as our two newly-developed circular polarization cuts, as described in Chapter 4. Instead of optimizing on the number of simulated events passing cuts in accepted Healpix bins, we optimized out stage 3 cuts to the lowest flux limit that could be established.

6.2 Discussion of 10% Sample Analysis Results

We exercised and tuned the analysis process by running it against a 10-percent sample extracted from the full dataset. We expected ∼1.0 weighted background events to pass our analysis cuts, but in fact a combined 8.3 weighted events, comprised of nine (9) individual contributing events. passed in the horizontal and vertical polarizations. While it is not possible to make conclusions from the properties of nine events, consideration of these events led to some ideas for further investigation and improvements to the analysis process.

Three very strong impulsive HPol events were localized to points in bin 2998 within 200km of McMurdo Station. We do not know what these events are, but they were much too strong to be removed by waveform characteristic cuts, and are probably anthropogenic. We hoped that the optimized linear discriminant cut would remove such events; however, the optimization process is relatively immature and may be improved by further investigation.

Two comparatively weak events that passed cleared the Hilbert peak by a very narrow margin, but the Hilbert peak cut threshold was set loosely to begin with. More effort to optimize the Hilbert peak cut, and possibly the correlation peak cut as well, is recommended.

Two events that passed all cuts were strong in left-hand circular polarization. Our newly-developed circular polarization cuts were designed to remove such events, which we believe originate largely from satellites. We do not yet know from where in the sky these events originated, but this can be determined from the interferometry. In any case, we know that the areas of the northern sky near the horizon and the northern continent just below the horizon are the sources, respectively, of satellite signals and their reflections. We probably should have included a simple cut to exclude events from those regions.

Finally, due to a mistake in data preparation, we constructed our 10-percent sample from a blinded dataset which included inserted, modified calibration pulser events. Since the number of inserted events is on the order of 10, it is entirely likely that one or two of these events ended up in the 10-percent sample and, as would be expected, passed all cuts.

6.3 Opportunities for Improvement

6.3.1 Event Waveform Filtering

For CW mitigation, we used the geometric filter described in B. Dailey [33]. The performance of this filter was slightly better that the sine-subtraction method for ANITA-II events, and we determined that the performance of the geometric filter was on ANITA-III events was comparable to that of the sine-subtraction method, and nominally acceptable. We did not, however, undertake a deep analysis of filtering performance on ANITA-III. Furthermore, the ANITA Collaboration has made considerable improvements to the sine-subtraction method contemporaneously to the completion of this Work; that method has been optimized and compared side-by-side to the geometric filter and other methods. While the performance of the geometric and sine-subtraction filters remains comparable, the decision on which filter to use should be reconsidered seriously in subsequent analyses.

6.3.2 Signal-to-noise Ratio

As explained in Chapter 3, the computation of signal-to-noise ratio (SNR) for noisy, dispersed impulsive signals is challenging. Noise estimation of ANITA-III waveforms was complicated by the fact that our signals appear very early in ANITA-III's sampling window; obtaining a reliable noise sampling from before the signal arrival is difficult. The normalization technique we developed to deal with the differing start-up times of signals from the different antennas shows promise and merits more analysis and refinement. Moreover, the other SNR calculation technique we developed, using the Hilbert envelope, was not fully debugged in time to facilitate its use in the fiunal analysis. We think both of these developments have potential value and should be explored further.

6.3.3 Circular Polarization Analysis

The significance of circular polarization as a discriminator of satellite-influenced events occurred relatively late in the development of our analysis. The two circular polarization cuts that we developed are certainly not the only possible such cuts that could be imagined. For example, the simple fraction of circular polarization in a waveform is reasonable basis for a cut; we did not implement this due in part to concerns about the accuracy of the Stokes' parameter calculations in the `AnitaTools` package. These issues are currently being addressed by the Collaboration and reliable Stokes' parameter values are expected to be available. Further exploration of the circular polarization regime is strongly encouraged.

6.3.4 SNR-dependent Localization PDF

We allocated events to bins on the continent using a PDF based on pointing resolution data derived from calibration pulses. The standard deviations $(\sigma_\phi, \sigma_\theta)$ of calibration pulse pointing error were used as uniform width parameters in a Gaussian PDF. We did not consider the strength, that is the SNR, of the cal-pulses in this part of the analysis. Over 100,000 cal-pulses were obtained from WAIS alone, so assessing the SNR-dependence of our pointing resolution should not be difficult. This would allow us to develop a SNR-dependent PDF, which might lead to improvements in localization.

6.3.5 Payload Deadtime

The ANITA-III payload experienced frequent intermittent deadtime during flight due to spikes in the event rate, as mentioned in 2.3. Since these spikes were due in large part to malfunctions in instrument software, it is reasonable to consider cutting events that occurred during periods of high deadtime. Associating individual events with a deadtime rate, however, is not trivial. Deadtime information is delivered in housekeeping data packets, which were recorded at a rate of about 1Hz. These housekeeping data packets contain a timestamp in whole seconds. Some packets contain duplicate timestamps, so it is difficult to determine which packets to associate with some events. Also, there exist gaps in the stream of housekeeping data which make it difficult or impossible to obtain a meaningful deadtime fraction. We therefore recommend that succeeding analysts attempt to tackle this problem.

* * * * * * *

We believe that the most important contribution of this work is the determination that new cuts based on circular polarization analysis will enhance rejection of events contaminated by CW from geostationary satellites.

We are hopeful that this work will be refined and augmented by our successors in the collaboration. The circular polarization cuts which we developed are significant, but considerable opportunity exists for improvement. Improvements should also be made to the other analysis cut optimizations, as indicated previously in this chapter. After refinement, the optimization should be rerun on an unblinded 10% dataset. Then, a subsequent analysis on the blinded 90% dataset may lead to fruitful conclusions.

References

[1] R. Davis. *Phys. Rev. Lett.*, 12:303, 1964.

[2] W. D. Arnett, J. N. Bahcall, R. P. Kirshner, and S. E. Woosley. *Ann. Rev. Astron. Astrophys.*, 27:629–700, 1989.

[3] A. B. McDonald. *New Journal of Physics*, 6(1):121, 2004.

[4] J. D. Jackson. *Classical Electrodynamics.* John Wiley & Sons, Inc, 3 edition, 1999.

[5] R. Cahn and G. Goldhaber. *The Experimental Foundations of Particle Physics.* Cambridge University Press, 2009.

[6] K. Greisen. *Phys. Rev. Lett.*, 16:748–750, 1966.

[7] Ralph Engel, David Seckel, and Todor Stanev. *Phys. Rev. D*, 64:093010, Oct 2001.

[8] J. A. Formaggio and G. P. Zeller. *Rev. Mod. Phys.*, 84:1307–1341, 2012.

[9] T. Stanev. *High Energy Cosmic Rays.* Praxis Publishing Ltd., 2004.

[10] C. Patrignani et al. *Chin. Phys.*, C40(10):100001, 2016.

[11] K. Kotera, D. Allard, and A. V. Olinto. *JCAP*, 1010:013, 2010.

[12] F. A. Aharonian. *Very High Energy Cosmic Gamma Radiation: A Crucial Window on the Extreme Universe.* World Scientific, 2004.

[13] G. Askaryan. *Soviet Physics JETP-USSR*, 14 (2):441–443, 1962.

[14] F. Reines et al. *Phys. Rev.*, 117:159–173, Jan 1960.

[15] J. N. Bahcall. *Phys. Rev. Lett.*, 12:300–302, Mar 1964.

[16] S. A. Colgate and R. H. White. *Astrophys. J.*, 143:626, 1966.

[17] Y. Totsuka et al. *Nucl. Phys.*, A478:189–195, 1988.

[18] P. W. Gorham et al. *Phys. Rev. Lett.*, 99:171101, 2007.

[19] P. W. Gorham et al. *Phys. Rev.*, D72:023002, 2005.

[20] Y. Fukuda et al. *Nucl. Instrum. Meth.*, A501:418–462, 2003.

[21] M. G. Aartsen et al. *Phys. Rev.*, D88:112008, 2013.

[22] P. W. Gorham et al. *Astropart. Phys.*, 32:10–41, 2009.

[23] Ohio Supercomputer Center. Ohio supercomputer center. `http://osc.edu/ark:/19495/f5s1ph73`, 1987.

[24] B. Strutt and J. Avva. *ANITA Collaboration E-log*, 575, 2014.

[25] B. Strutt. PhD thesis, University College London, 2016.

[26] G. S. Varner, L. L. Ruckman, P. W. Gorham, J. W. Nam, R. J. Nichol, J. Cao, and M. Wilcox. *Nucl. Instrum. Meth.*, A583:447–460, 2007.

[27] S. Hoover et al. *Phys. Rev. Lett.*, 105:151101, 2010.

[28] S. Wissel and Saltzberg. *ANITA Collaboration E-log*, 617, 2015.

[29] S. Wissel and Saltzberg. *ANITA Collaboration E-log*, 595, 2014.

[30] R. Brun and F. Rademakers. In *Proceedings AIHENP'96 Workshop, Lausanne, Sep. 1996*, volume 389, pages 81–86, 1997.

[31] K. Belov et al. *ANITA Collaboration E-log*, 622, 2014.

[32] A. Romero-Wolf et al. *Astropart. Phys.*, 60:72–85, 2015.

[33] B. Dailey. PhD thesis, The Ohio State University, 2017.

[34] P. Fretwell et al. *The Cryosphere*, 7, 2013.

[35] Website Article:. *Where Lightning Strikes.* NASA, 2001 (accessed May 1, 2017). `https://science.nasa.gov/science-news/science-at-nasa/2001/ast05dec_1`.

[36] A. Vieregg. PhD thesis, University of California Los Angeles, 2010.

[37] A. Romero-Wolf. *ANITA Collaboration E-log*, 621, 2015.

[38] B. Strutt. http://www.hep.ucl.ac.uk/ strutt/meetingSlides/blindingSetup.pdf, 2016.

[39] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. *The Astrophysical Journal*, 622(2):759, 2005.

[40] S. Lea. *Mathematics for Physicists.* Thomson Brooks/Cole, 2004.

[41] R. Brun and F. Rademakers. *ROOT Data Analysis Framework User's Guide.* 2014.

[42] M. Galassi et al. *GNU Scientific Library Reference Manual (3rd Ed.).* 2009.

Appendix A

Normalization of cross-correlation using overlapping bins

Consider two arrays $\mathbf{x}_i$ and $\mathbf{y}_i$, both of length n and zero-padded such that their elements vanish unless $i \in [p, q]$, that is, [1]

$$\begin{aligned} \mathbf{x} = & \quad \{0, \quad 0, \quad ... \quad x_p, \quad x_{p+1}, \quad ... \quad x_{q-1}, \quad x_q, \quad ... \quad 0, \quad 0\} \\ \mathbf{y} = & \quad \{0, \quad 0, \quad ... \quad y_p, \quad y_{p+1}, \quad ... \quad y_{q-1}, \quad y_q, \quad ... \quad 0, \quad 0\}. \end{aligned} \tag{A.1}$$

While the domain of the circular cross-correlation is $j : j \in [(1 - n/2), (n/2 - 1)]$, its support is confined to $[(p - q), (q - p)]$. The overlapping normalization denominator is

$$M_j = \sqrt{X_j Y_j} \tag{A.2}$$

---

[1] The support intervals of $\mathbf{x}$ and $\mathbf{y}$ need not be identical as presented here; in the general case, $[p, q]$ is the interval of common support.

with

$$X_j = \sum_{i=p-j}^{q} x_i^2 \qquad Y_j = \sum_{i=p}^{q+j} y_i^2 \qquad j \in [p-q, 0] \tag{A.3a}$$

$$X_j = \sum_{i=p}^{q-j} x_i^2 \qquad Y_\tau = \sum_{i=p+j}^{q} y_i^2 \qquad j \in [1, q-p] \tag{A.3b}$$

The normalization has been used in ANITA-II analysis [33], but direct calculation of the sums in (A.3) for every M_τ is $O(n^2)$. For this ANITA-III analysis, we implemented this method in the AnitaTools suite while reducing the computation to $O(n)$ using the recurrence relations

$$\begin{aligned} X_j &= X_{j-1} + x_{p-j}^2 \qquad & Y_j &= Y_{j-1} + y_{q+j}^2 \qquad & j &\in [p-q, 0] \\ X_j &= X_{j+1} + x_{q-j}^2 \qquad & Y_j &= Y_{j+1} + y_{p+j}^2 \qquad & j &\in [1, q-p] \end{aligned} \tag{A.4}$$

to initial values from (A.3):

$$\begin{aligned} X_{p-q} &= x_q^2 \qquad & Y_{p-q} &= y_p^2 \\ X_{q-p} &= x_p^2 \qquad & Y_{q-p} &= y_q^2. \end{aligned} \tag{A.5}$$

The computation cost of (A.4) is negligible compared to the $O(n \log n)$ cost of the FFT.

Appendix B

Computation of S_{up}

In optimizing the linear discriminant cut as discussed in Chapter 4, it is necessary to determine S_{up}, the number of signal events, given background b, for which an observation of b is 90% excluded. Recalling equation (4.8), S_{up} is specified by requiring

$$\frac{\int_0^{S_{up}}(s+b)^b e^{-(s+b)}ds}{\int_0^{\infty}(s+b)^b e^{-(s+b)}ds} = U(S_{up};b) = 1-\alpha. \tag{B.1}$$

This equation is solved iteratively through repeated evaluation of the integrals in the left-hand side. Here, we show how (B.1) can be transformed to allow computation of $U(S_{up},b)$ through a simple code invoking built-in functions available in ROOT.

Note that with $t = s+b$,

$$\int_0^a (s+b)^b e^{-(s+b)}ds = \int_b^{(a+b)} t^b e^{-t}dt = \gamma(1+b,a+b) - \gamma(1+b,b), \tag{B.2}$$

where $\gamma(x,z) = \int_0^z t^{(x-1)}e^{-t}dt$ is the lower incomplete gamma function[40]. Noting that the ordinary gamma function $\Gamma(x) = \gamma(x, \infty)$, (4.8) becomes

$$U(S_{up}; b) = \frac{\gamma(1+b, S_{up}+b) - \gamma(1+b, b)}{\Gamma(1+b) - \gamma(1+b, b)}. \tag{B.3}$$

In terms of the regularized incomplete gamma function $\gamma_r(x, z) = \gamma(x, z)/\Gamma(x)$, which is available as a library function in ROOT[41], we have

$$U(S_{up}; b) = \frac{\gamma_r(1+b, S_{up}+b) - \gamma_r(1+b, b)}{1 - \gamma_r(1+b, b)}, \tag{B.4}$$

which can be coded in a simple function. S_{up} is then obtained by requiring

$$U(S_{up}; b) - (1 - \alpha) = 0, \tag{B.5}$$

and applying a root-finding utility such as those available in gsl[42] and ROOT[30].

Appendix C

Continuous Analog to the Poisson Distribution

Stochastic, uncorrelated discrete events occurring at a fixed mean frequency in space or time follow a Poisson distribution, provided that the frequency of occurrence is sufficiently small that the probably of multiple events within the smallest measurable interval is negligible. Given a mean event rate of λ within the measurement interval of interest, the probability of k events occurring within such an interval is

$$P(k) = \frac{e^{-\lambda}\lambda^k}{k!} \qquad k \in \mathbb{N} \tag{C.1}$$

Eq. (C.1), with a suitably chosen λ, is thus the PDF of the number of actual number of neutrino events expected in a particular bin on the Antarctic continent. The associated discrete CDF is

$$C(k) = \lambda^k \sum_{i=0}^{k} \frac{\lambda^i}{i!}. \tag{C.2}$$

But the optimization process of our linear discriminant cut (Chapter 4) utilizes simulated events modeled by the `icemc` program. These simulated events are weighted with an occurrence probability, so instead of counting discrete simulated events, we determine event counts by calculating the sum of their weights, a continuous variable. A discrete PDF, therefore, is not appropriate to analysis of simulated events; a continuous PDF must therefore be used.

We obtain a continuous analog to the Poisson distribution by making the substitutions

$$k \leftarrow \kappa, \quad k! \leftarrow \Gamma(\kappa+1) \qquad \kappa \in \mathbb{R},\ \kappa > 0, \tag{C.3}$$

where $\Gamma(x)$ is the complete Gamma function[40]. The continuous PDF analogous to the Poisson distribution is thus

$$\mathcal{P}(\kappa;\lambda) = \frac{e^{-\lambda}\lambda^{\kappa}}{\Gamma(\kappa+1)}, \tag{C.4}$$

and the corresponding CDF is

$$\mathcal{C}(\kappa;\lambda) = \int_0^{\kappa} \mathcal{P}(\kappa')d\kappa' = 1 - \gamma_r(\kappa+1,\lambda), \tag{C.5}$$

where $\gamma_r(x,y)$ is the regularized lower incomplete Gamma function previously described in Appendix B.

Appendix D

Alternative methods of SNR calculation

The signal-to-noise (SNR) ratio of a waveform is an important figure of merit used in making analysis cuts. Calculating SNR is challenging in our case because it is not possible to discern exactly the signal content from the noise. Dispersion in the instrument signal chain distorts the impulsive characteristic of the signal, and the nature of this distorted signal is known only approximately. A typical SNR calculation uses as the signal value a peak value from the waveform, or in some cases the difference between the maximum and minimum values in the waveform. Instrument dispersion also complicates the calculation of the noise level by pushing a portion of the signal power into later portions of the waveform. This is exacerbated if the trigger was influenced by CW, which by its nature is likely to continue long after the impulsive signal has subsided. Data analyses of previous ANITA missions mitigated this problem by using a sample of the waveform from *before* the arrival of the signal [33], [36]. In the case of ANITA-III, this was more difficult because the signal arrives earlier in the waveform sampling window than was the case in previous missions. In the `AnitaTools`, therefore, noise is by default sampled from the end of the waveform, but this is subject to noise overestimation for the reasons just mentioned.

We attempted to improve the SNR calculations. Instead of sampling the end of the waveform for noise, we sampled near the beginning. In order to mitigate problems arising from the comparatively early signal arrival, we take a very short (10ns) sample from before the signal arrival, accepting that this may be subject to larger fluctuations that a longer sample would be. Figure D.1 shows a typical coherent sum reconstruction of an ANITA-III calibration pulse. The increased waveform amplitude near the end of the sampling window, boxed in red is indicative of this problem. A noise sampling from the beginning of the waveform is boxed in blue.

The top panel in Figure D.1 shows a problem with sampling from the beginning of the waveform. When building the coherent sum reconstruction, a delay is applied to each waveform in accordance with the arrival direction of the signal. Therefore, data points very near the start of the waveforms may have zero values, and the number of these zero buckets varies across waveforms. As a result, early on in the waveform, only a portion of the waveforms actually contribute to the coherent sum. This is manifested in the approximately linear "ramp-up" of the noise amplitude of the waveform in the top panel, elucidated by the green line. This prompted us to consider an alternative normalization method to account for the smaller number of contributing antennas at the beginning of the waveform.

To compensate for a varying number of contributing antennas, one is inclined naïvely to normalize each data point in the coherent sum by dividing by the number of contributing antennas. That is, if we have a total of N antennas, but only n of these contributing an amplitude a_i to the kth data point in the waveform, the total normalized amplitude would be

$$A_k = \frac{1}{n} \sum_{i=1}^{n} a_i. \tag{D.1}$$

A problem with this method, however, is that it essentially treats the noise as adding coherently and scaling by n, when noise in fact adds incoherently and scales as $\sqrt{n}$. We mitigated this in the following way. If all antennas were contributing, the obvious normalization factor would be $1/N$, and the normalized amplitude at the kth data point would be (up to a scaling constant)

$$A_k = \frac{1}{N} \sum_{i=1}^{n} a_i. \tag{D.2}$$

Considering that we are adding up a total of N waveforms to make our coherent reconstruction, but only n_k waveforms actually contributing at the kth data point, we want to normalize to render a value as though all N waveforms were contributing. This to accommodate the $\sqrt{n}$ additive behavior of noise, we multiply eq. (D.2) by $\sqrt{N/n}$, so the normalized amplitude is

$$A_k = \frac{N}{n} \frac{1}{N} \sum_{i=1}^{n} a_i = \frac{1}{\sqrt{nN}} \sum_{i=1}^{n} a_i. \tag{D.3}$$

The effect of this can be seen in the lower panel of Figure D.1. The characteristic amplitude of the noise, shown by the green line, is much flatter than in the top panel. We implemented the modified normalization of eq. (D.3) as an option in the `AnitaTools`, and used it throughout our analysis. While noise estimates may be

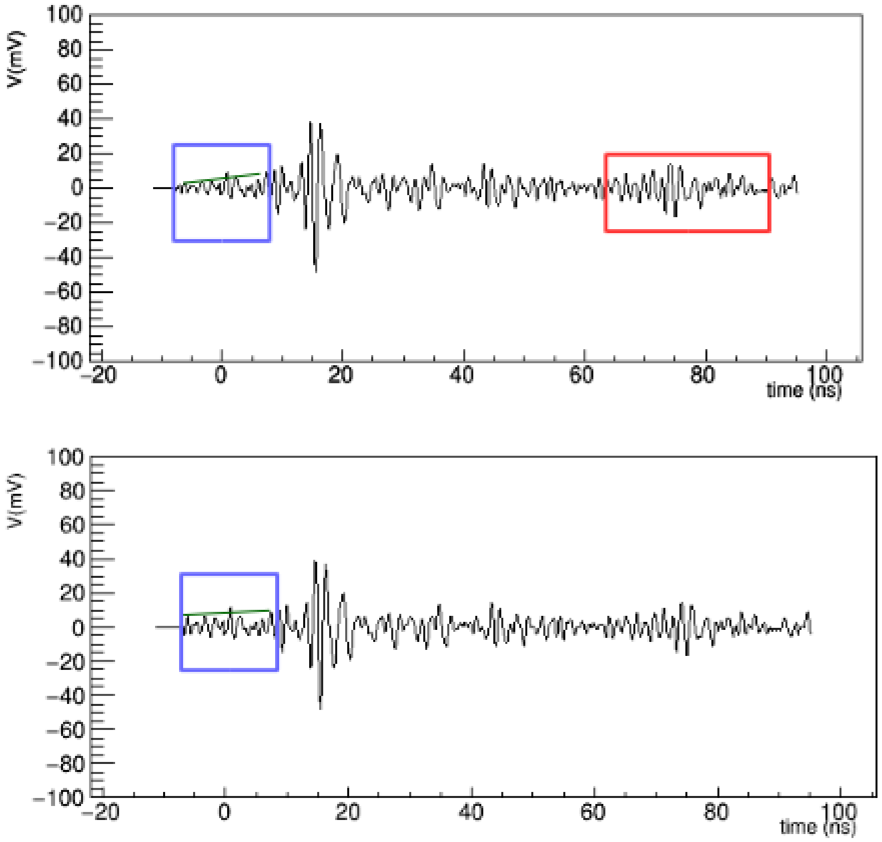

Figure D.1: Noise estimate considerations in coherent reconstructions of ANITA-III waveforms. Top: Near the end of the waveform is an increase in amplitude (red box), possibly due to signal dispersion or CW contamination, so we use instead a sample from before the signal arrival (blue box). However, a waveform normalized according to eq. (D.2) shows a quasi-linear "ramp-up" (green line) in the noise level as more antennas begin to contribute.
Bottom: by using eq. (D.3) instead, we obtain a much flatter noise amplitude (green line) despite the varying number of contributing antennas.

improved using (D.3), calculating signal level, mentioned previously, is also not trivial. Using peak voltage level or peak-to-peak voltage difference, as is the convention among ANITA analysts, is limited by the fact that it uses only one or two data points as input and may not accurately characterize total signal power. We developed an additional SNR calculation method that samples the waveform over a 20ns time window around the signal maximum, in hope of better characterizing overall signal power. A convenient representation of signal power arises from the complex-valued *analytic signal* of the coherently-summed waveform. The analytic signal employs the Hilbert transform $\mathcal{H}$ to obtain a "proper" imaginary part, i. e., one meeting the Cauchy-Riemann conditions[40], to combine with a real-valued time domain signal $\mathbf{x}(t)$:

$$\mathbf{A}(t) = \mathbf{x}(t) + i\mathcal{H}[\mathbf{x}](t) \tag{D.4}$$

The magnitude and phase of $\mathbf{A}(t)$ represent respectively instantaneous power and phase of the signal. In the collaboration, we refer to the magnitude of the analytic signal colloquially as the "Hilbert envelope".

The Hilbert envelope is desirable because it typically contains a single decisive maximum around which a sampling window can be constructed. We developed an SNR calculation using the RMS of the portion of the Hilbert envelope spanning from 7.5 ns before its maximum, to 12.5 ns thereafter, as shown in Figure D.2. For noise, we used a 10 ns window beginning at the earliest nonzero data point in the Hilbert envelope. We implemented this method in the `AnitaTools` package.

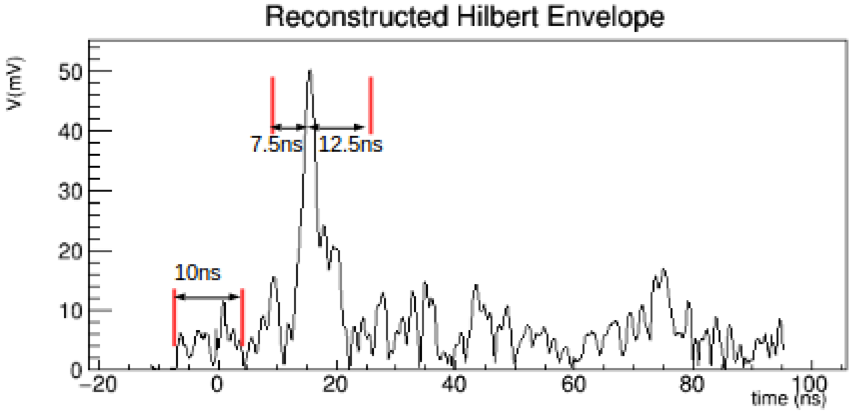

Figure D.2: SNR calculation using the Hilbert envelope. The signal power is estimated by integrating from 7.5ns before to 12.5ns after the time of maximum power. The noise is obtained from a 10ns sampling at the beginning of the waveform.

Appendix E

Sample Spreadsheet for Optimizing the Linear Discriminant Cut

This appendix shows an example of a spreadsheet used to optimize the linear discriminant cut. The first objective is to optimize the slope such that the maximum proportion of simulated events contained in retained Healpix bins. Once a slope is selected, the scale factor is optimized to deliver a total probability of 0.1.

Column Descriptions

1. **Bin number**: the index number of the Healpix bin on the Antarctic continent.
2. **Total events, sample**: the total number of events in the 10% sample that localized to this bin.
3. **Sim events pre-rotated cut**: the total number of simulated events localized to this bin that passed all cuts except for the optimized linear discriminant cut.
4. **At least 5 fit bins**: (1) means at least 5 descending bins were found to the right of the maximum in the plot of the linear discriminant value for events in

the 10% sample, so a fit was performed; otherwise, no fit was done and the bin was discarded from the analysis.

5. **Optimized cut intercept**: The optimized intercept value for the fitted bin.
6. **Bin fit p-value**: The p-value of the log-likelihood of the fit with respect to the pseudo-experiment log-likelihood distribution.
7. **p-value > 0.05**: (1) means that the p-value from the previous column is greater than 0.05, so the bin is retained; otherwise the bin is discarded from the analysis.
8. **Sim events passing**: the number of simulated events passing the linear discriminant cut, using the y-intercept from column 5.
9. **Expected background**: The background estimate for this bin, that is, the number of events from the 10% sample that pass the optimized linear discriminant cut.
10. **Sim events before cuts**: The total number of simulated events originating from this bin.
11. **Poisson CDF**: The Poisson CDF obtained using a mean μ equal to the expected background (column 9) and a number of events k equal to the number of events passing the optimized cut (column 8).

The final probability at the bottom of the table is obtained by multiplying all of the CDF values in column 11. We optimize the scale factor to obtain a total probability of 10%.

Optimization results: Slope=-30.0 simulation scale factor=0.100

1	2	3	4	5	6	7	8	9	10	11
bin number	total-events sample	sim-events pre-rotated-cut	at-least 5-fit-bins	optimized cut-intercept	bin-fit p-value	p-value >0.05	sim-events passing	expected background	sim-events before-cuts	Poisson CDF
3050	129.3414	5.1460	1	12	0.712713	1	0.4073	0.8238	70.6481	0.912236
3034	193.2236	5.0673	1	14	0.7	1	0.3666	0.1362	66.1795	0.751702
3032	150.1064	4.5643	1	13.5	0.97	1	0.3171	0.3886	64.6980	0.862606
3048	51.9940	4.1375	1	12	0.52	1	0.3167	0.4413	42.1035	0.875685
3019	86.9540	3.8628	1	14.5	0.593	1	0.2486	0.3529	33.6547	0.890353
2994	97.0000	3.5966	1	13.5	0.650651	1	0.2506	0.5359	41.5695	0.924623
3020	125.4621	3.3805	1	12.5	1	1	0.2677	0.1199	30.4317	0.811705
2995	128.6052	3.3267	1	15	0.820821	1	0.1972	1.9312	36.7775	0.998654
3043	245.4057	2.3783	1	13.5	0.013013	1	0.1638	0.6481	27.3822	0.969093
3053	20.3900	1.8085	1	10.5	1	1	0.1639	0.0437	14.2806	0.863248
3057	120.4200	1.6132	1	12	0.409409	1	0.1308	0.1297	17.3058	0.911690
2967	93.7215	1.4932	1	14.5	0.962	1	0.0955	0.3366	27.9514	0.965603
3059	29.4819	1.4916	1	11.5	1	1	0.1171	0.3716	14.0220	0.959558
3025	134.5427	1.2160	1	13.5	0.567	1	0.0861	0.2623	13.7412	0.962218
3041	69.3321	0.9744	1	15	0.035	1	0.0565	12.2334	9.9148	1.000000
3028	48.7696	0.2204	1	12	0.174174	1	0.0172	0.3406	6.2418	0.996420
3003	19.9514	0.0250	1	10.5	0.936	1	0.0023	1.1130	0.3433	0.999999
3029	62.0988	0.0118	1	12	0.777778	1	0.0012	0.0019	0.1433	0.998833
3022	232.6055	0.0072	1	9.5	0.197	1	0.0007	1045.6965	0.0163	1.000000
3000	33.5532	0.0000	1	0	0.969	1	0.0000	5431.4712	0.0000	1.000000
3051	85.0151	6.1792	1	16.5	0	0	0.3095	1.5786	69.5279	0.989254
3045	380.7355	2.0837	1	16.5	0	0	0.1009	19.8368	19.6556	1.000000
3046	3177.6577	1.4972	1	11	0	0	0.1334	0.0982	14.4331	0.902310
3039	873.5281	0.4637	1	13	0	0	0.0324	482.0551	4.1882	1.000000
3058	57.4521	0.0000	0	0	0	0	0.0000	0.0000	17.6843	1.000000
...										
3002	13.1698	0.0000	0	0	0	0	0.0000	0.0000	0.1452	1.000000
TOTAL		54.5451					37.8303		1473.6726	
ACCEPTED		44.3214					32.0682		517.4054	
FRACTION		0.8126					0.8477		0.3511	
24 bins fitted										
18 bins accepted										
total probability		0.2414								

Table E.1: Example spreadsheet for optimizing the linear discriminant. Some detail rows have been omitted in order to save space.

Appendix F

Plots From the Accepted Bins

This Appendix contains a set of figures for each Healpix bin that was accepted during the optimization of the Stage 3 analysis cuts. These plots are explained in Section 4.2.2, and an example of the plots can be found in Figure 4.2. Each figure consists of:

1. Plot of the histogram of event count from the 10-percent sample vs. linear discriminant value, including the exponential fit made on the descending portion of the histogram.
2. Plot S and S_{up} vs the linear discriminant cut threshold. S is the estimated signal passing the cut, calculated using simulated events obtained from `icemc`; S_{up} is the optimization parameter described in detail in Appendix B.
3. Plot of the ratio S/S_{up} vs. linear discriminant cut threshold; this is the figure that is maximized to obtain the cut threshold for each bin.

4. Two-dimensional histogram of 10-percent sample events and simulated events, and the line representing the optimized linear discriminant cut threshold. The y-intercept of the line corresponds to the maximum in the S/S_{up} plot of panel (2).

We optimized the parameter set $\{m,\ C_1,C_2\}$ for common values across all healpix bins: m is the slope of the linear discriminant cut; C_1 is the threshold (maximum) for the circular polarization peak separation cut; C_2 is the threshold (minimum) of the circular polarization strength cut. The optimized values chosen for the fixed cut parameters were $m = -6.0$, $C_1 = 46.0°$, $C_2 = 0.015$. We optimized the linear discriminant cut intercepts R_i individually to each Healpix bin. These cuts are explained in detail in Chapter 4.

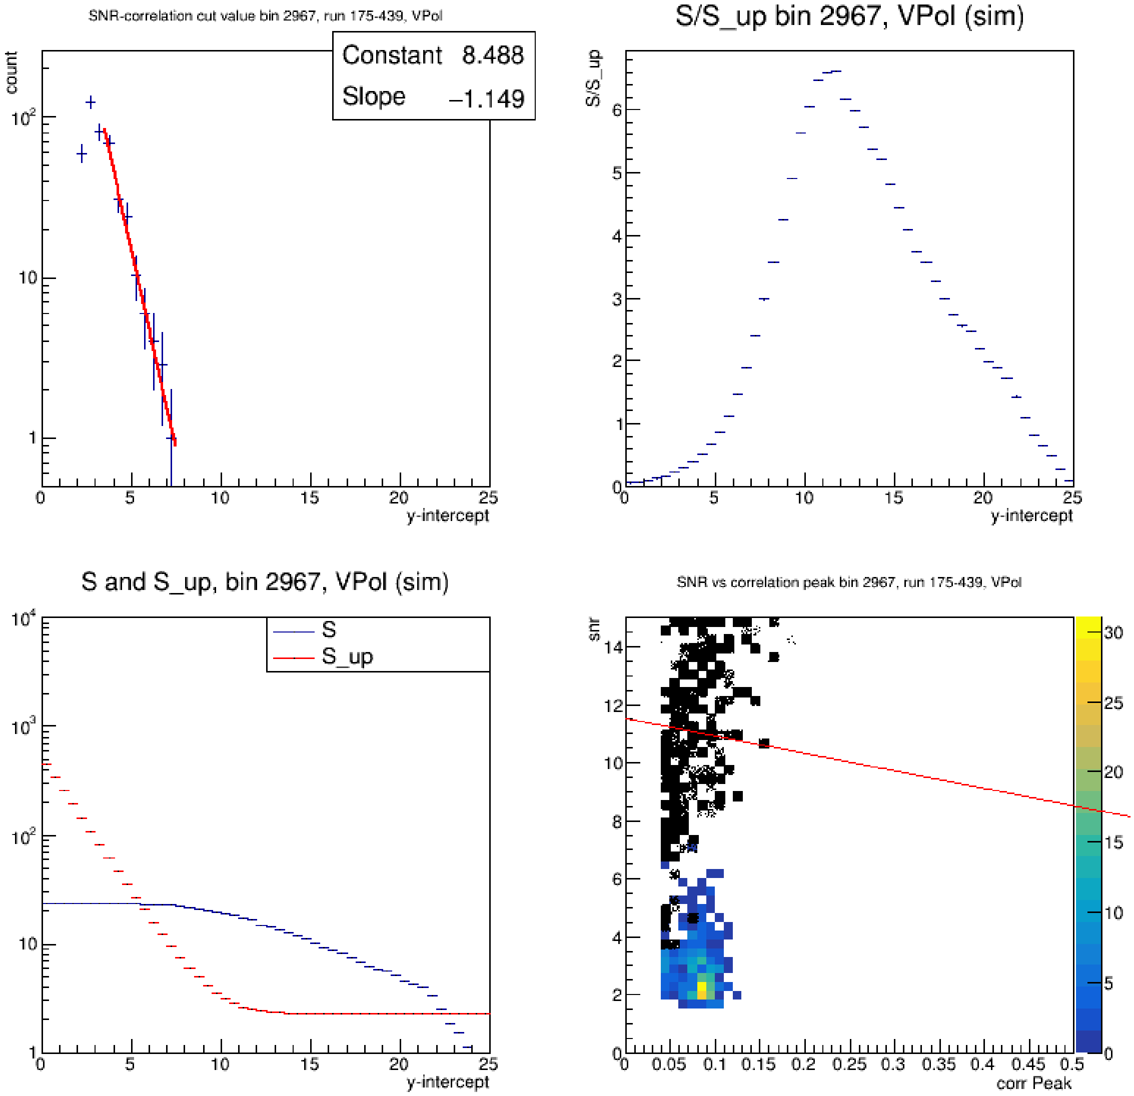

Figure F.1: Optimization of the linear discriminant cut, bin 2967. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

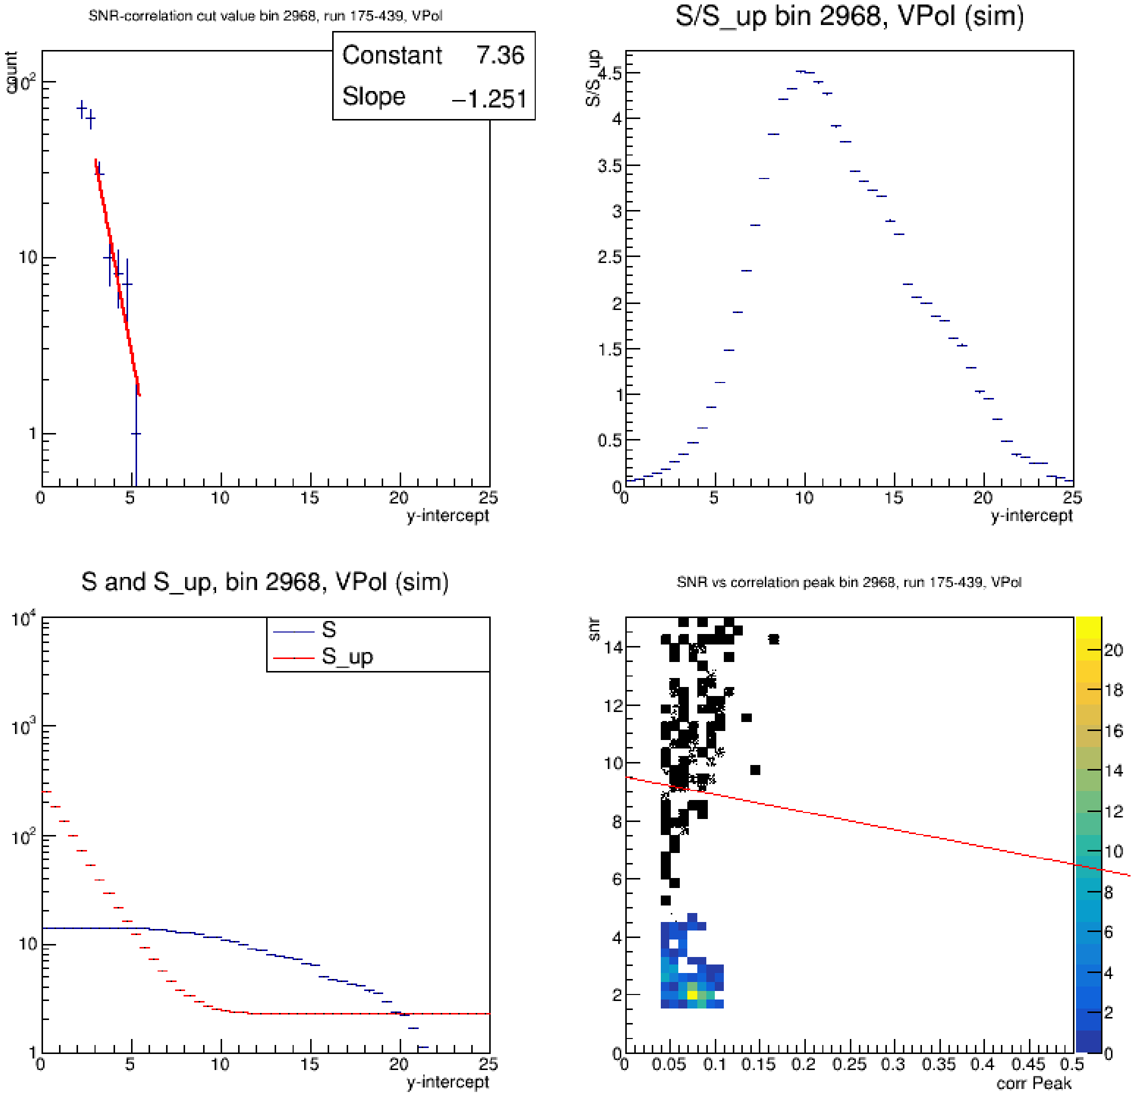

Figure F.2: Optimization of the linear discriminant cut, bin 2968. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 9.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

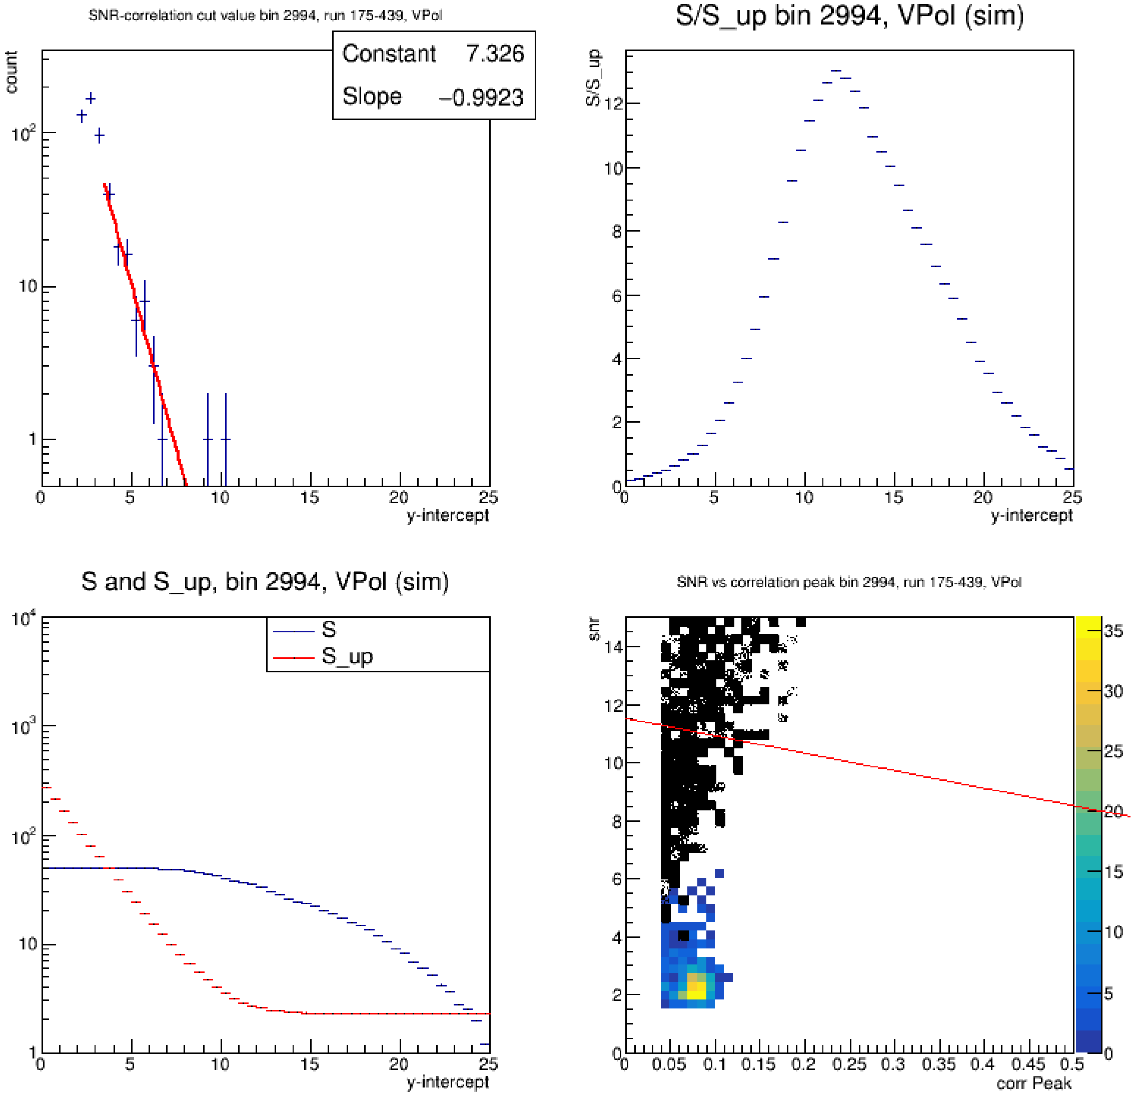

Figure F.3: Optimization of the linear discriminant cut, bin 2994. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

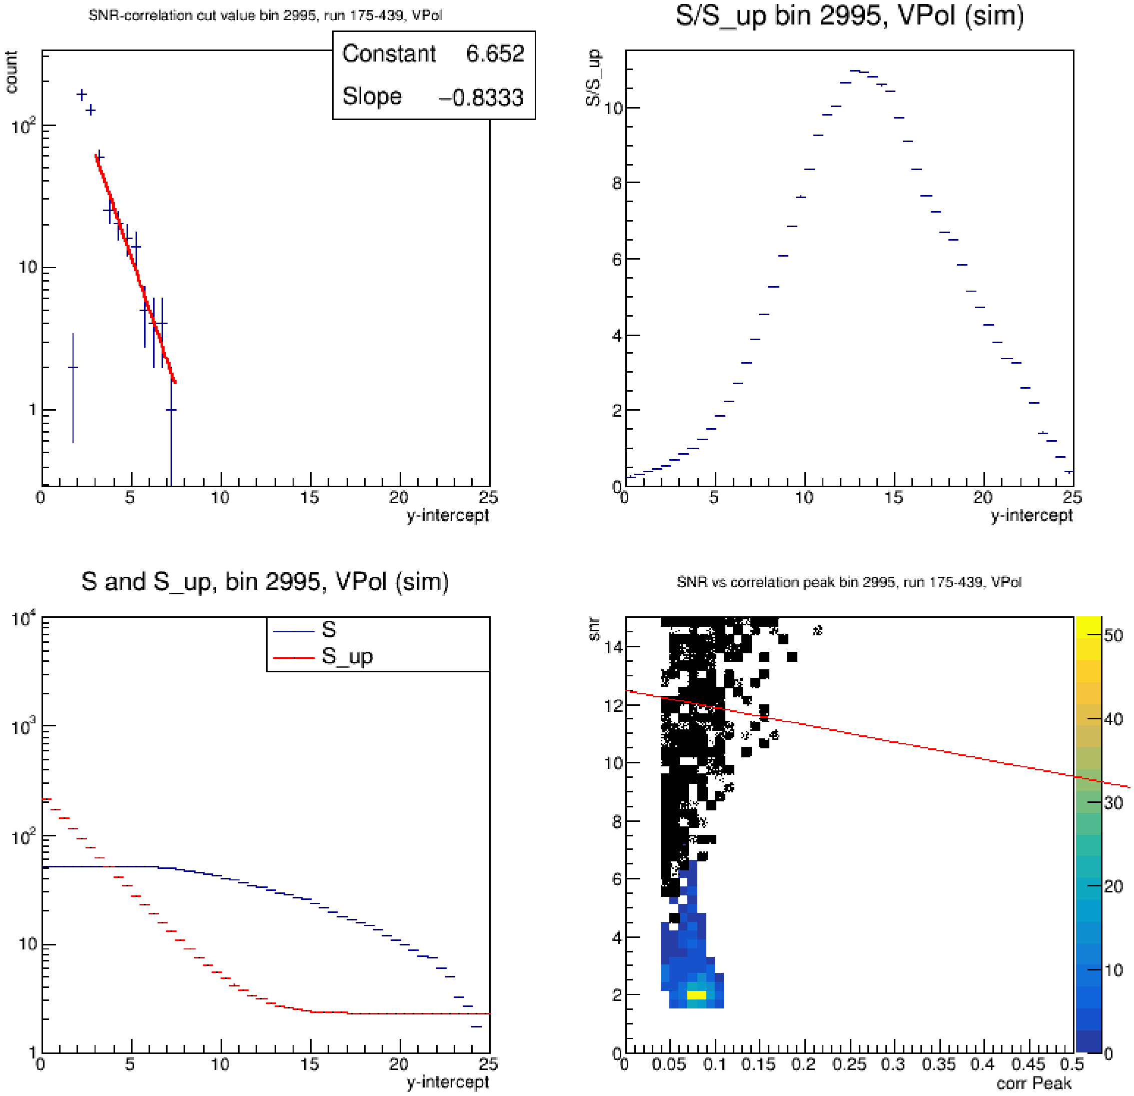

Figure F.4: Optimization of the linear discriminant cut, bin 2995. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 12.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

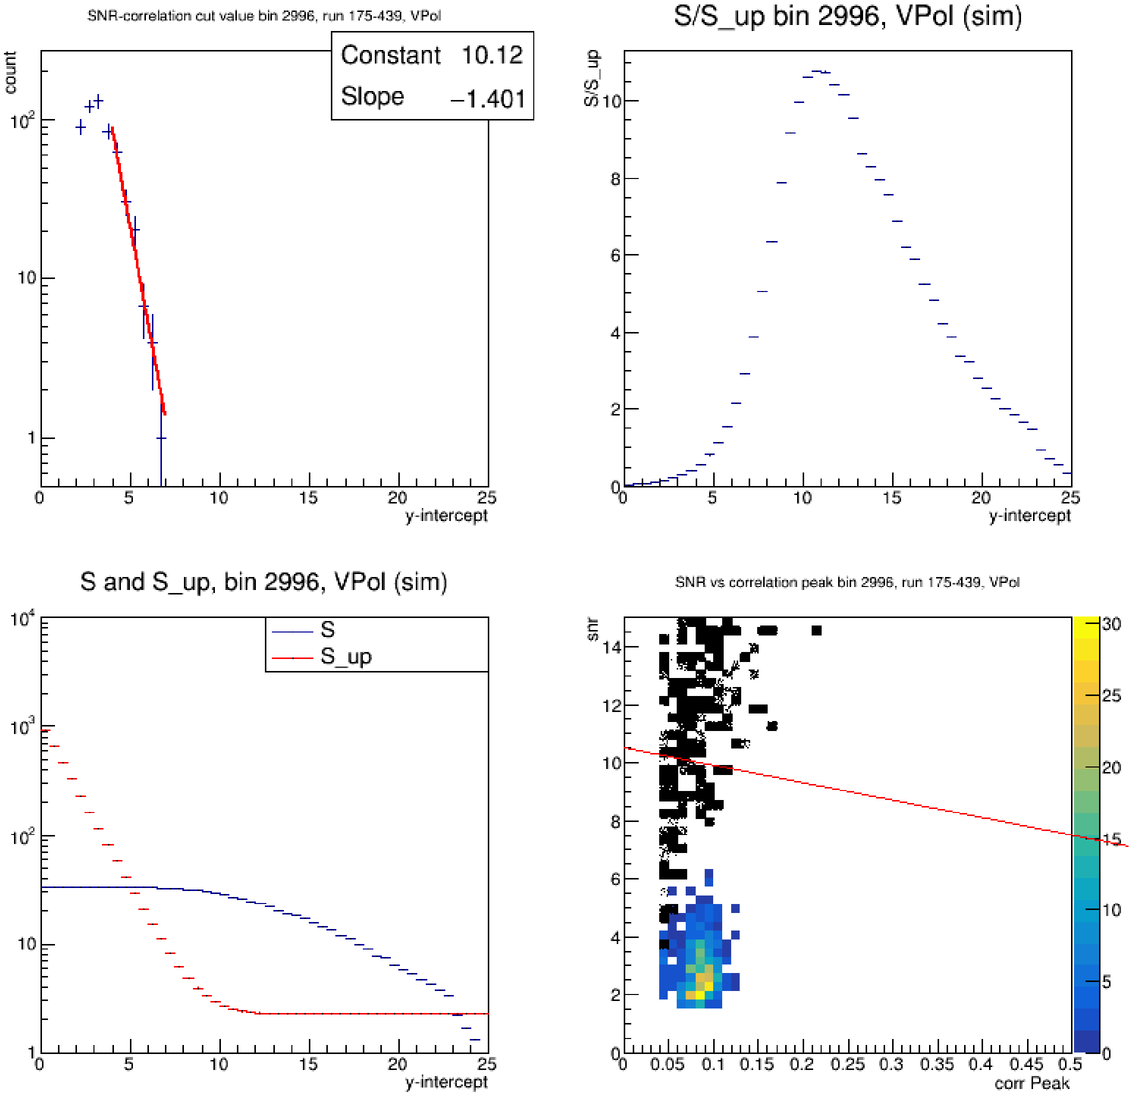

Figure F.5: Optimization of the linear discriminant cut, bin 2996. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 10.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

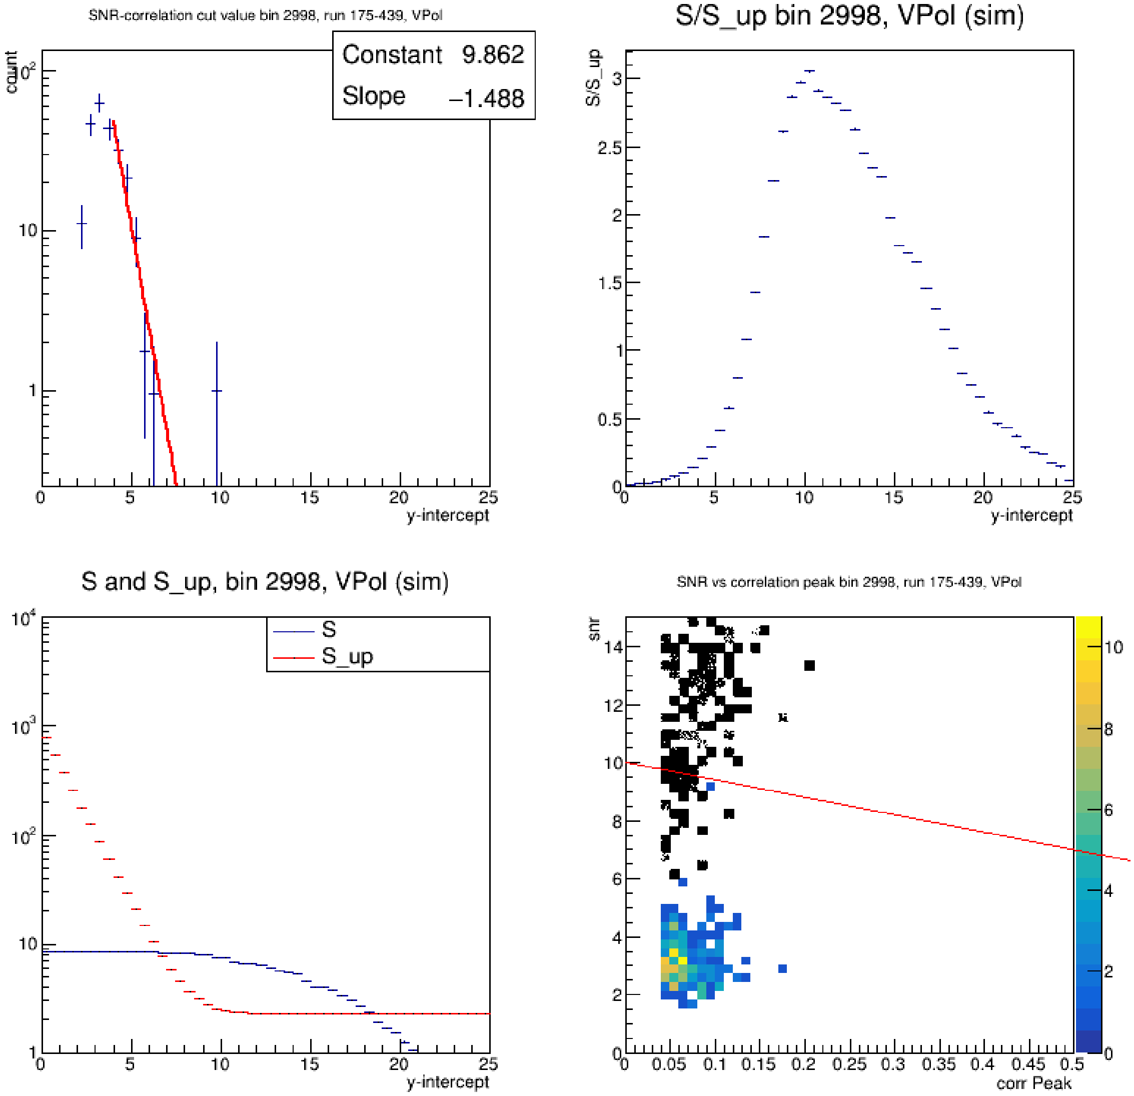

Figure F.6: Optimization of the linear discriminant cut, bin 2998. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 10.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

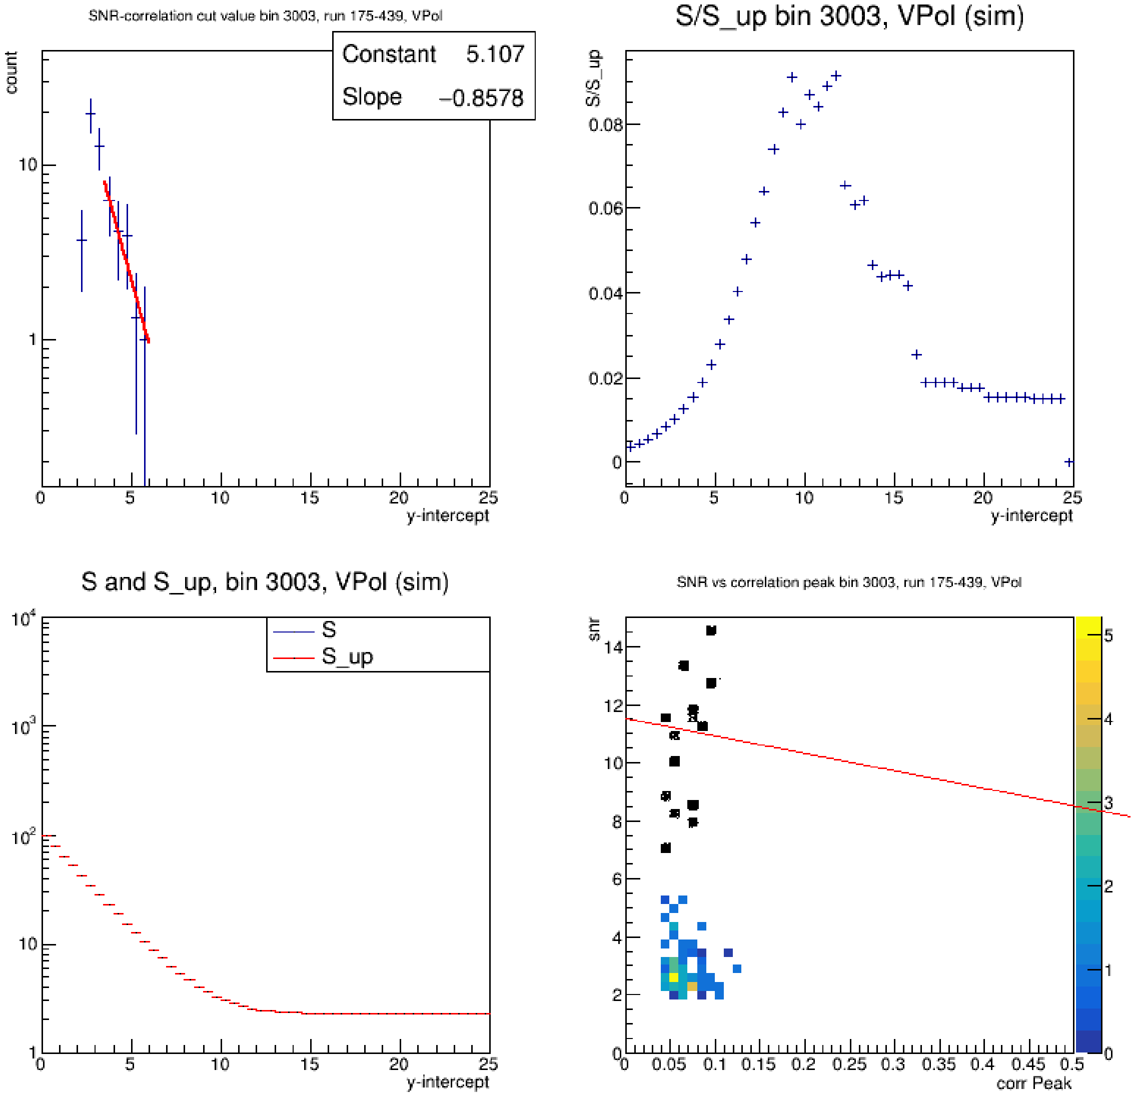

Figure F.7: Optimization of the linear discriminant cut, bin 3003. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

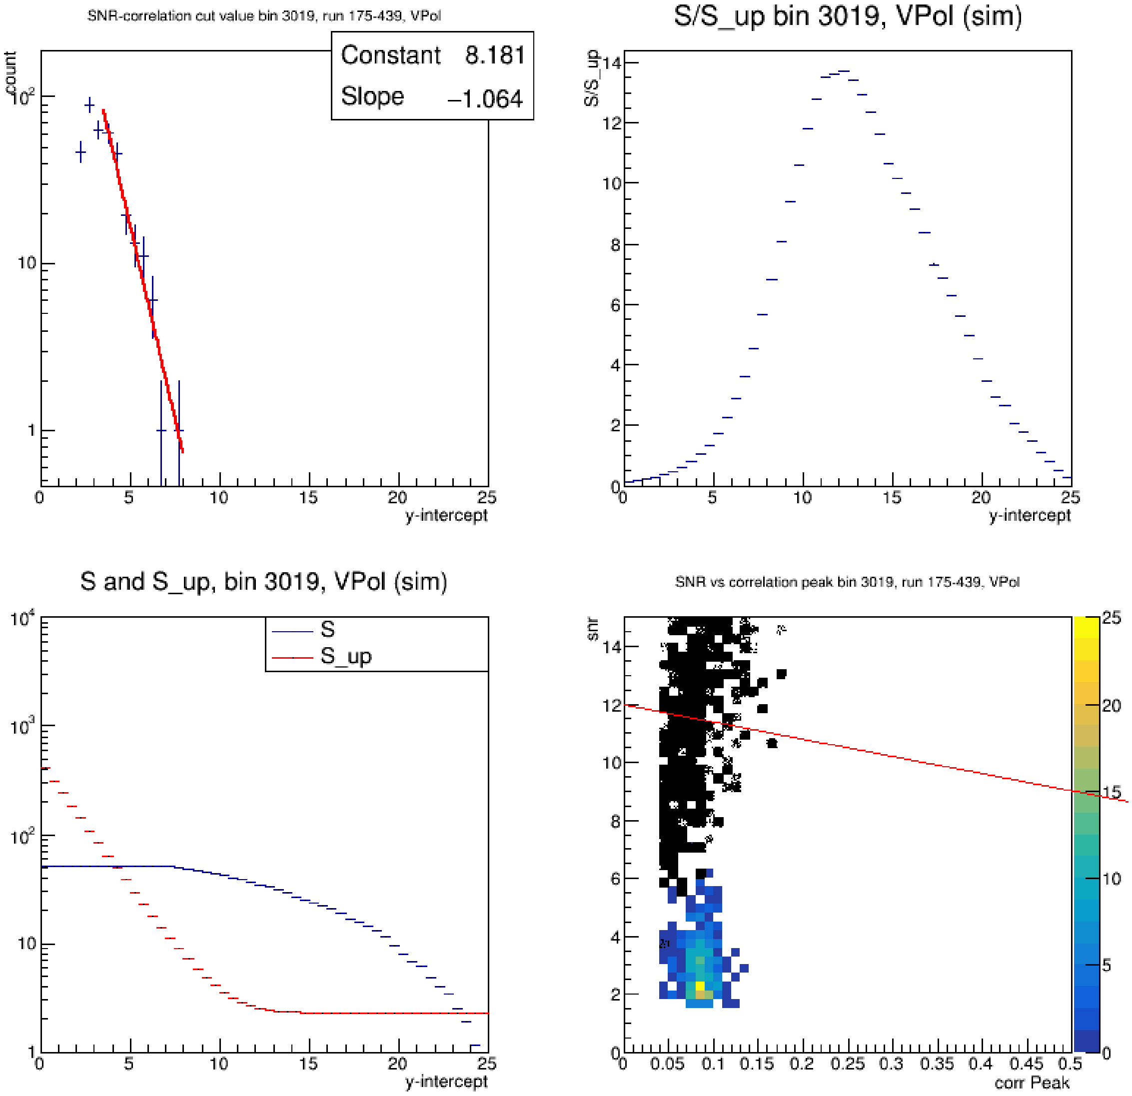

Figure F.8: Optimization of the linear discriminant cut, bin 3019. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 12.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

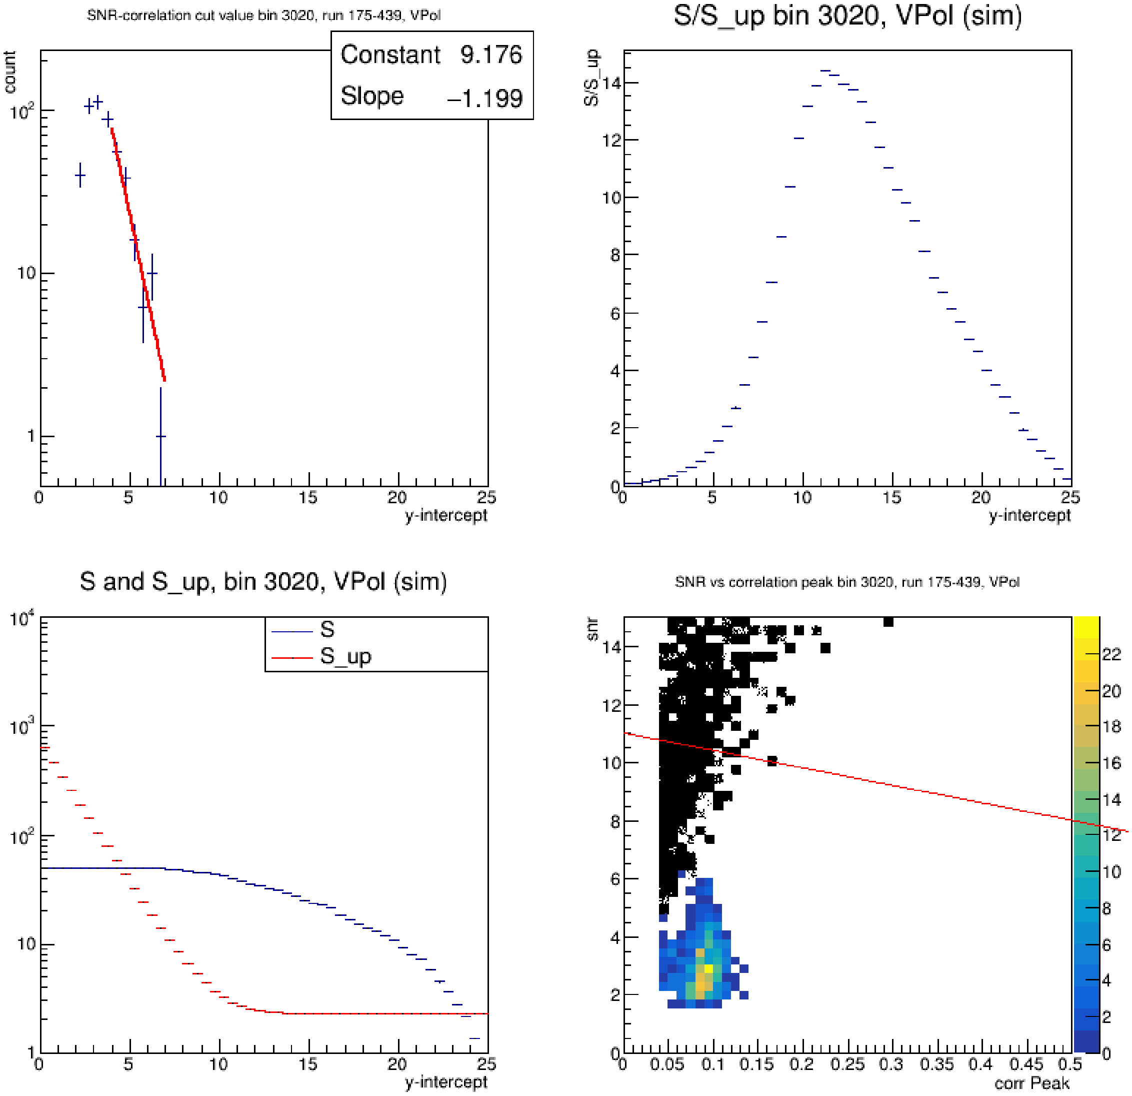

Figure F.9: Optimization of the linear discriminant cut, bin 3020. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

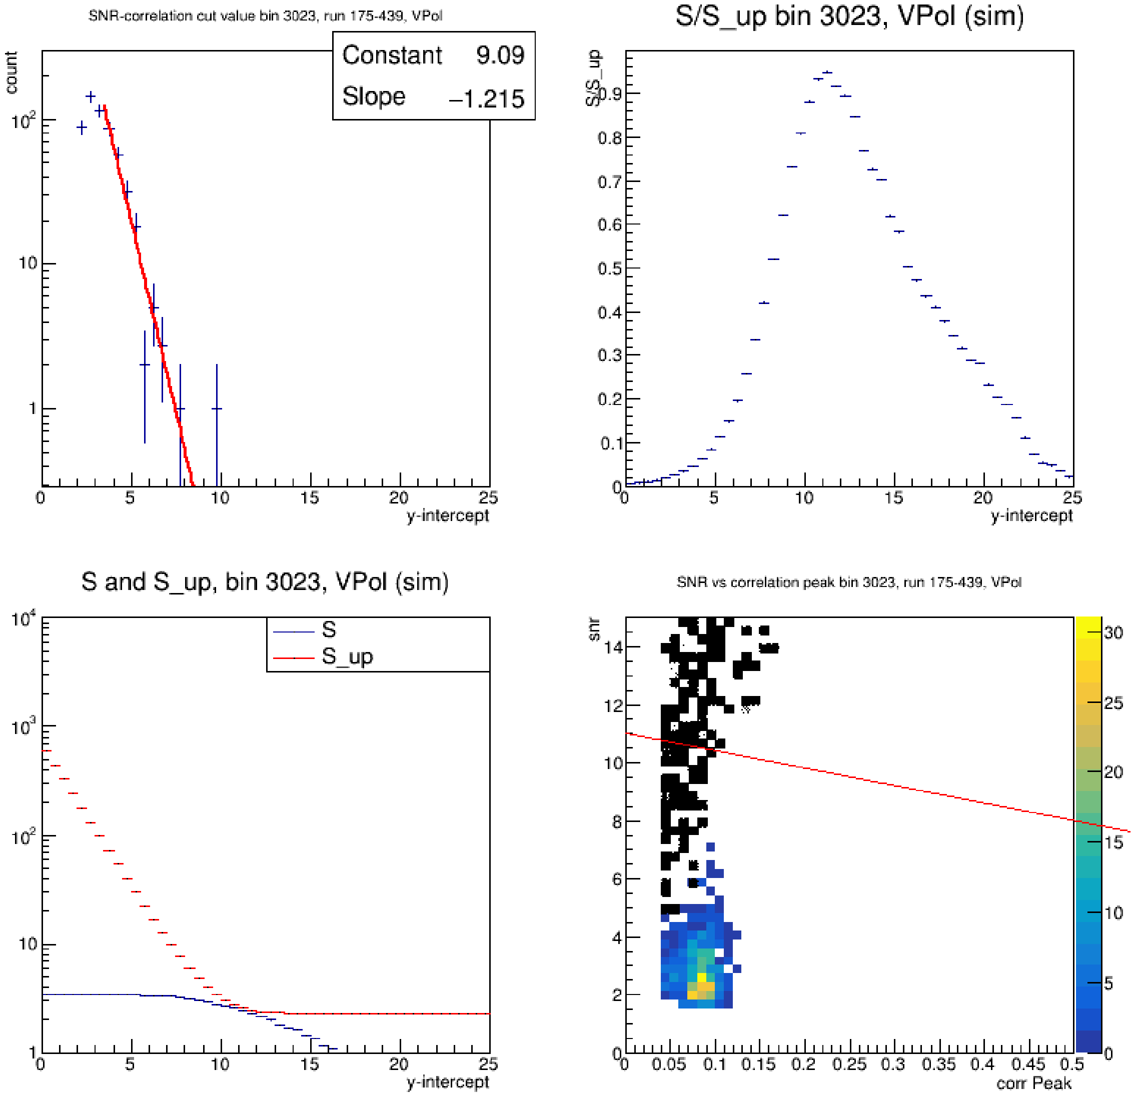

Figure F.10: Optimization of the linear discriminant cut, bin 3023. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

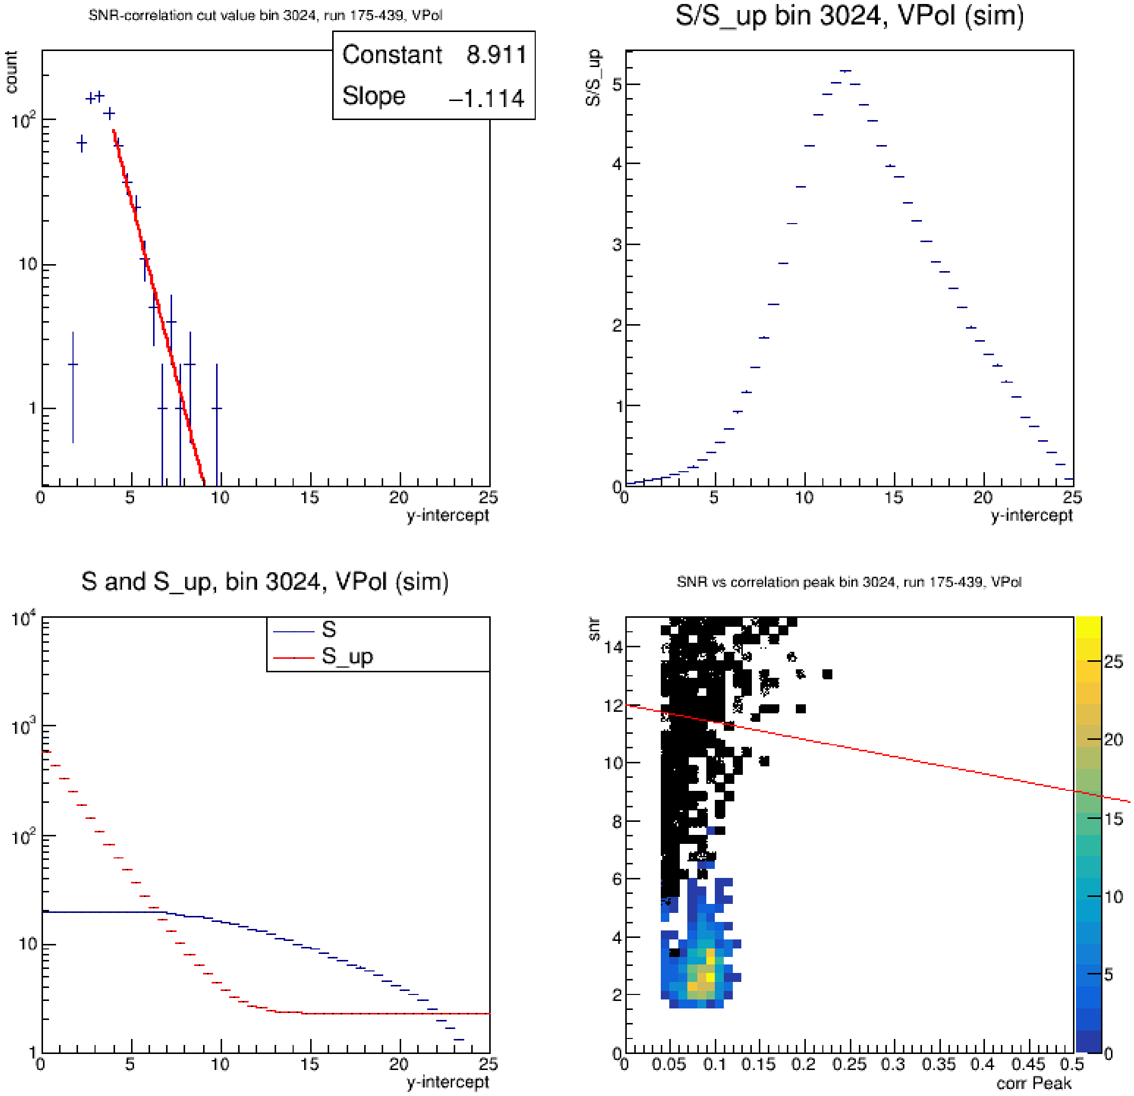

Figure F.11: Optimization of the linear discriminant cut, bin 3024. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 12.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

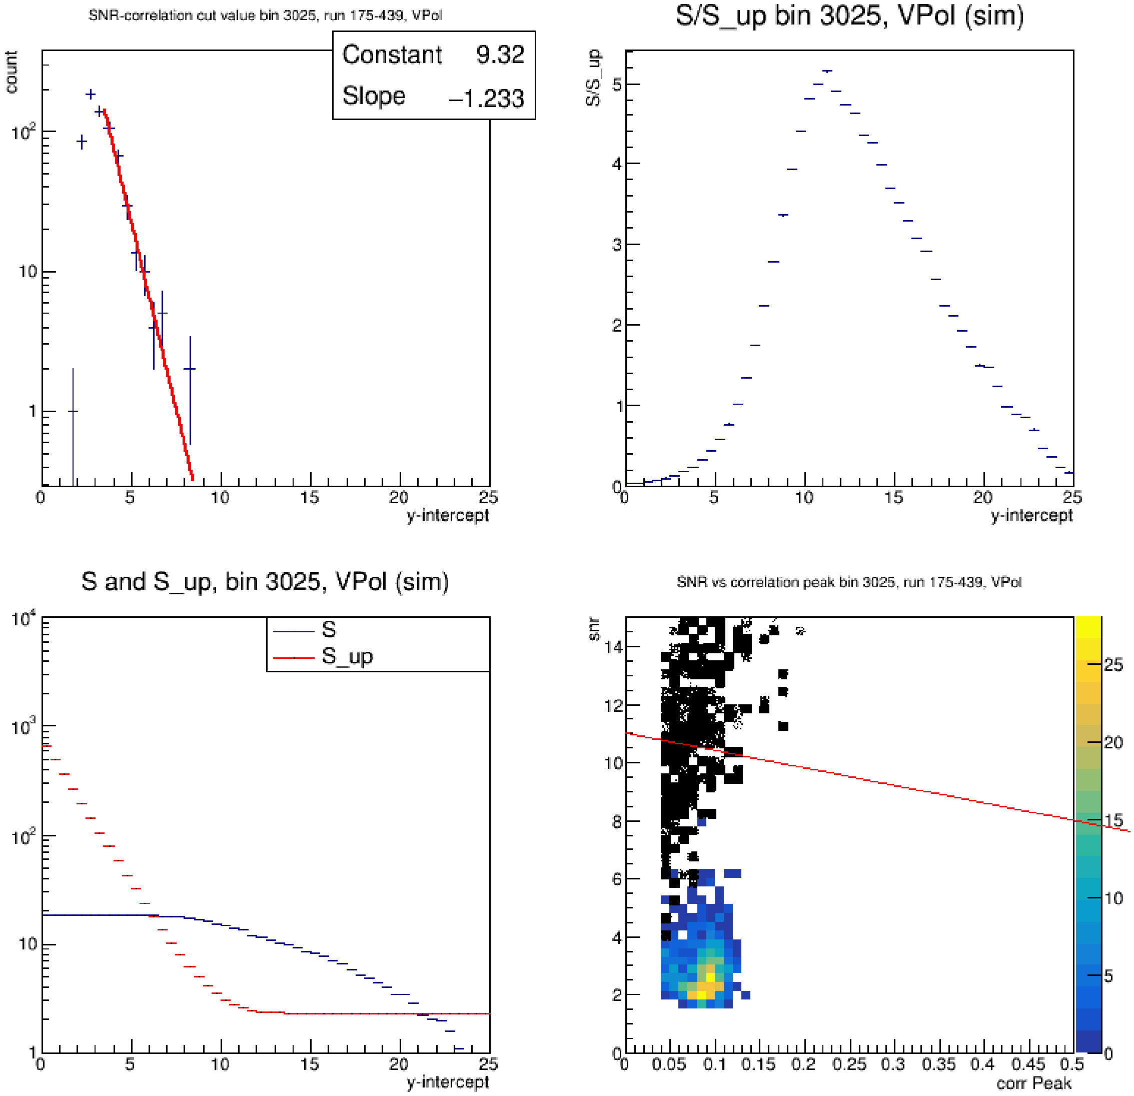
SNR-correlation cut value bin 3025, run 175-439, VPol

S/S_up bin 3025, VPol (sim)

S and S_up, bin 3025, VPol (sim)

SNR vs correlation peak bin 3025, run 175-439, VPol

Figure F.12: Optimization of the linear discriminant cut, bin 3025. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

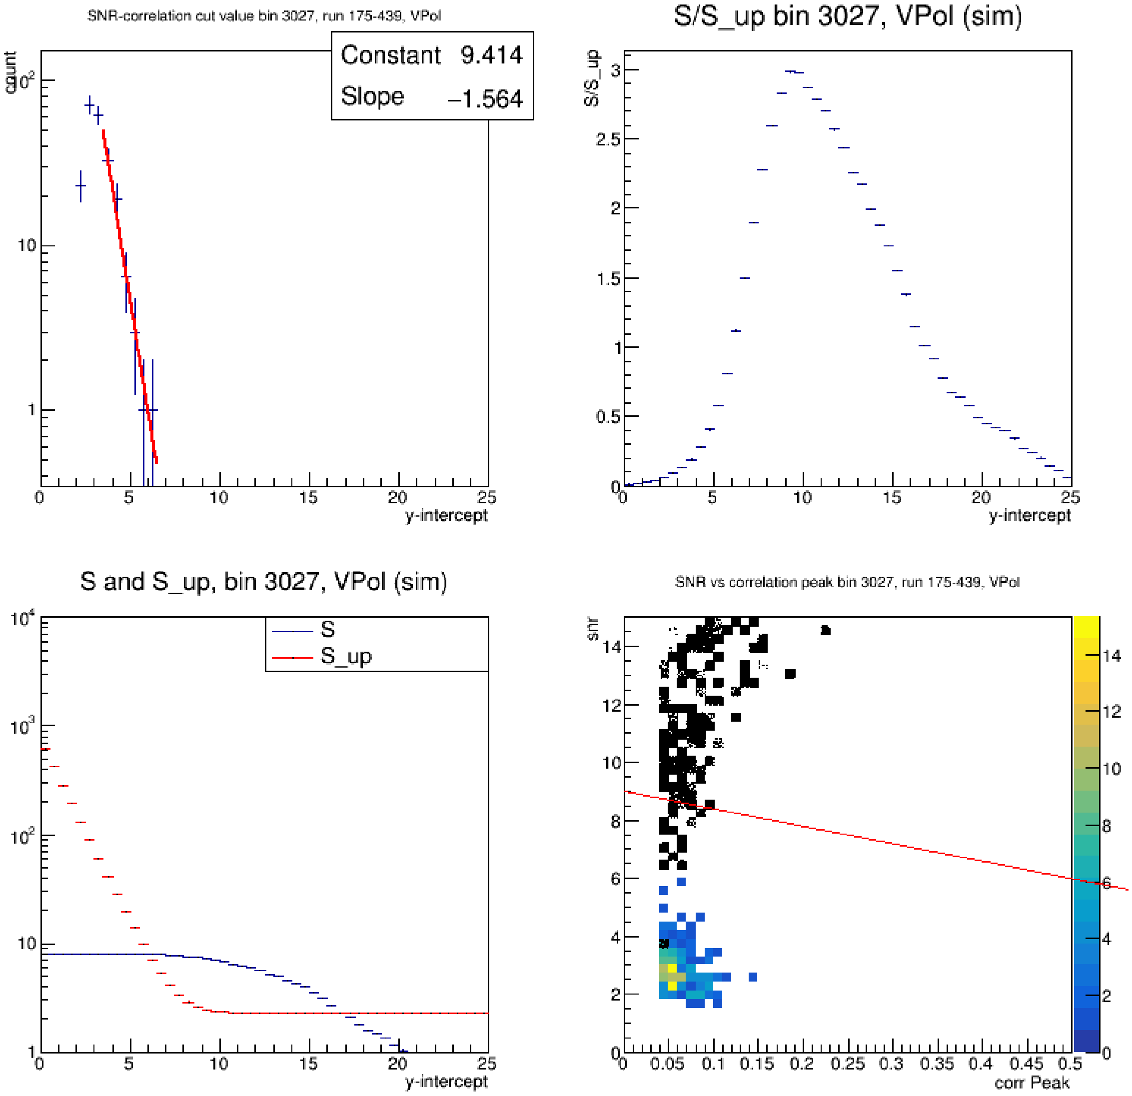

Figure F.13: Optimization of the linear discriminant cut, bin 3027. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 9.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

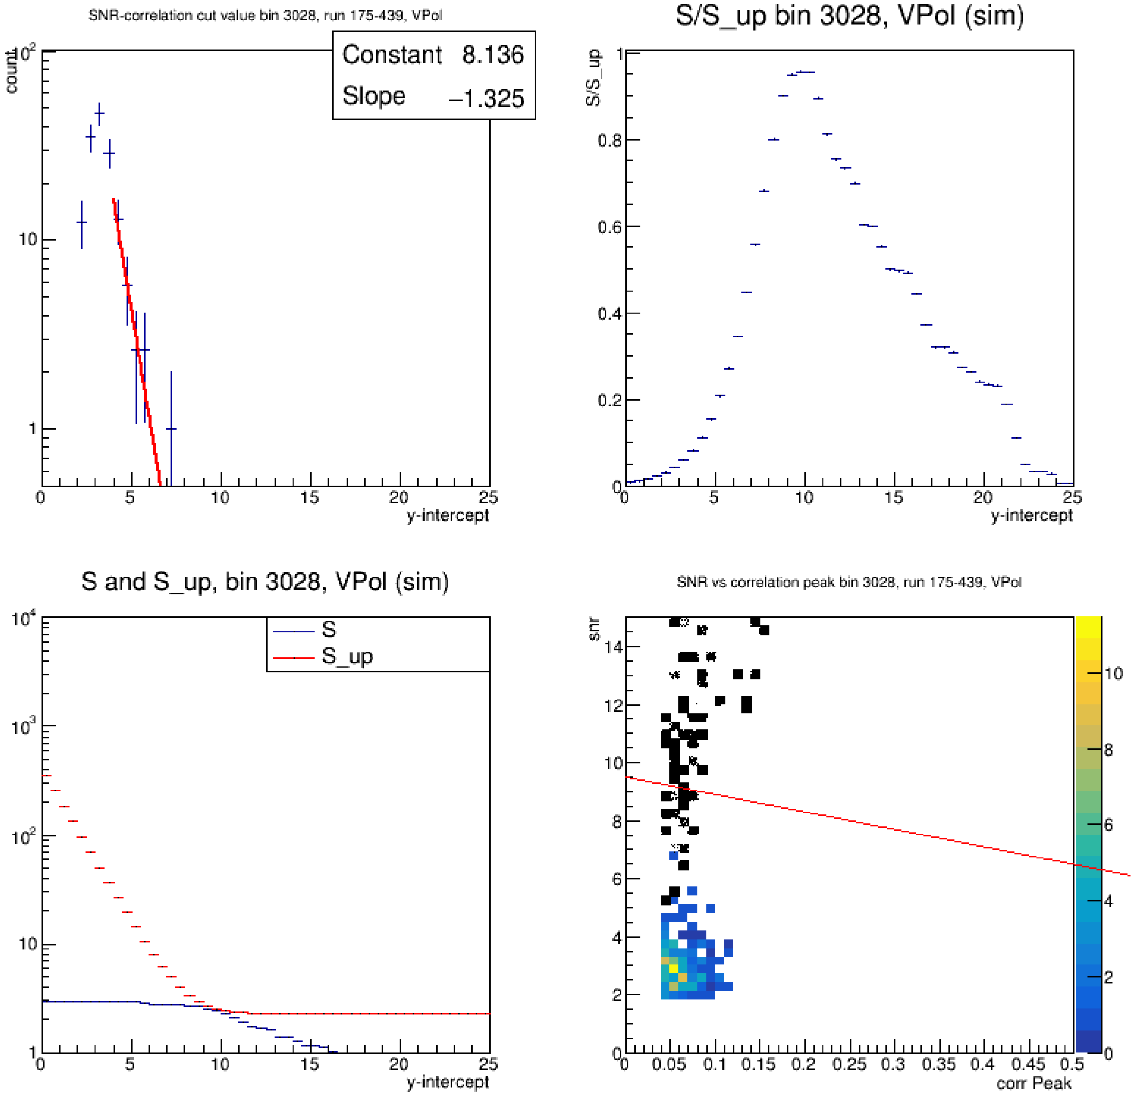

Figure F.14: Optimization of the linear discriminant cut, bin 3028. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 9.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

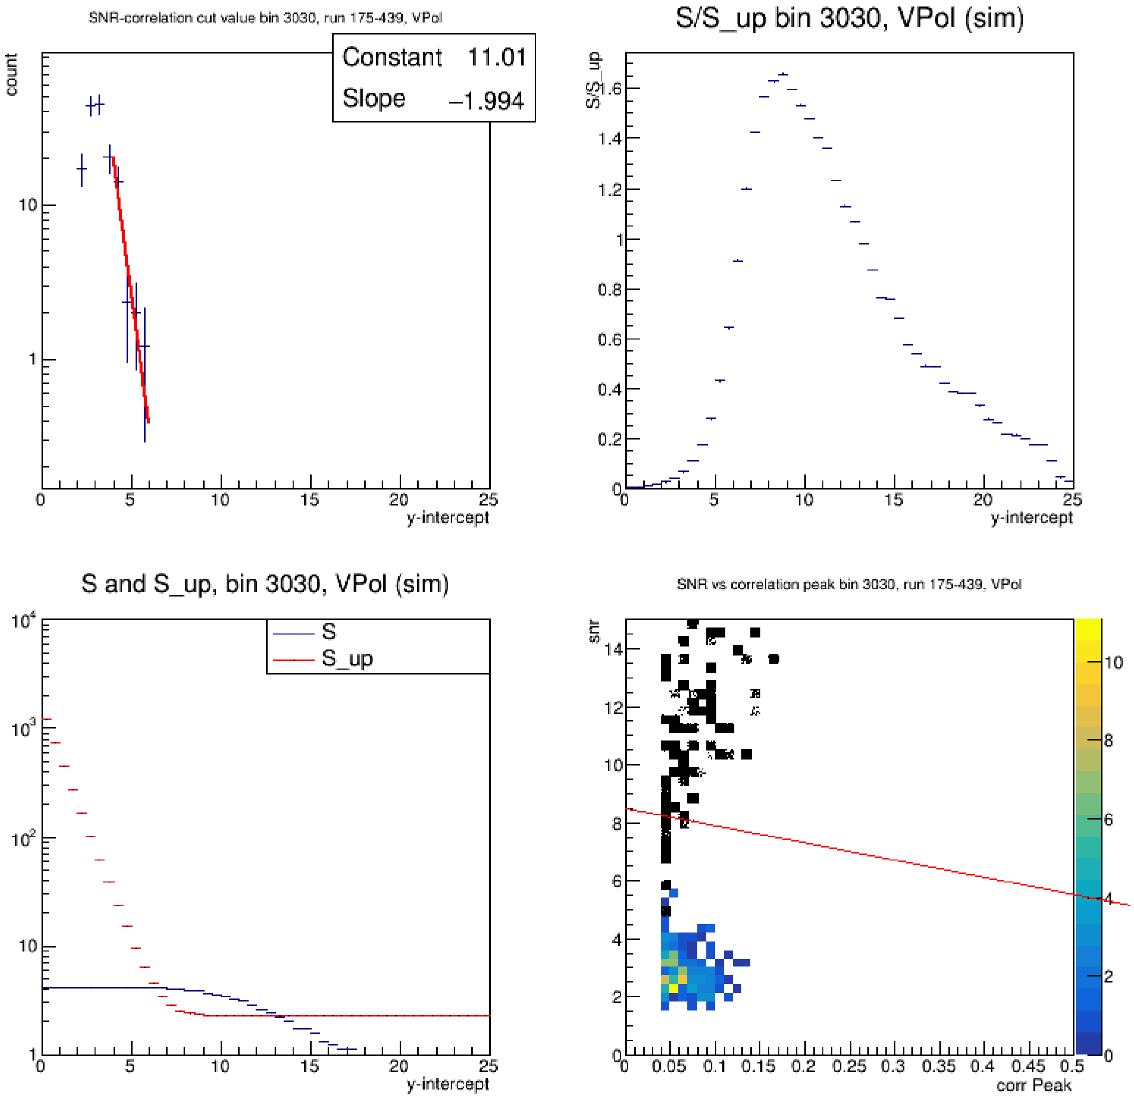

Figure F.15: Optimization of the linear discriminant cut, bin 3030. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 8.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

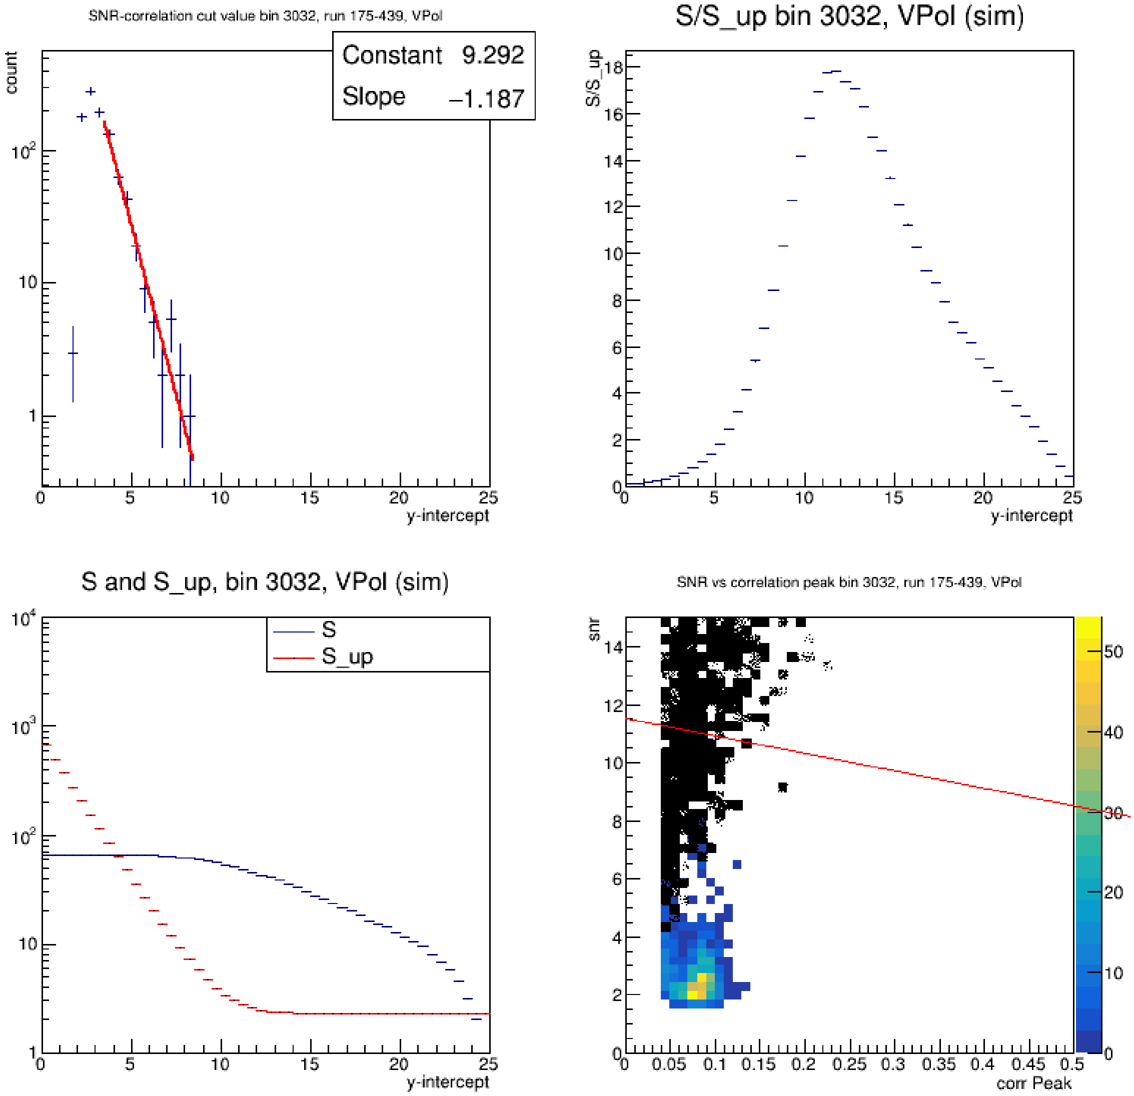

Figure F.16: Optimization of the linear discriminant cut, bin 3032. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

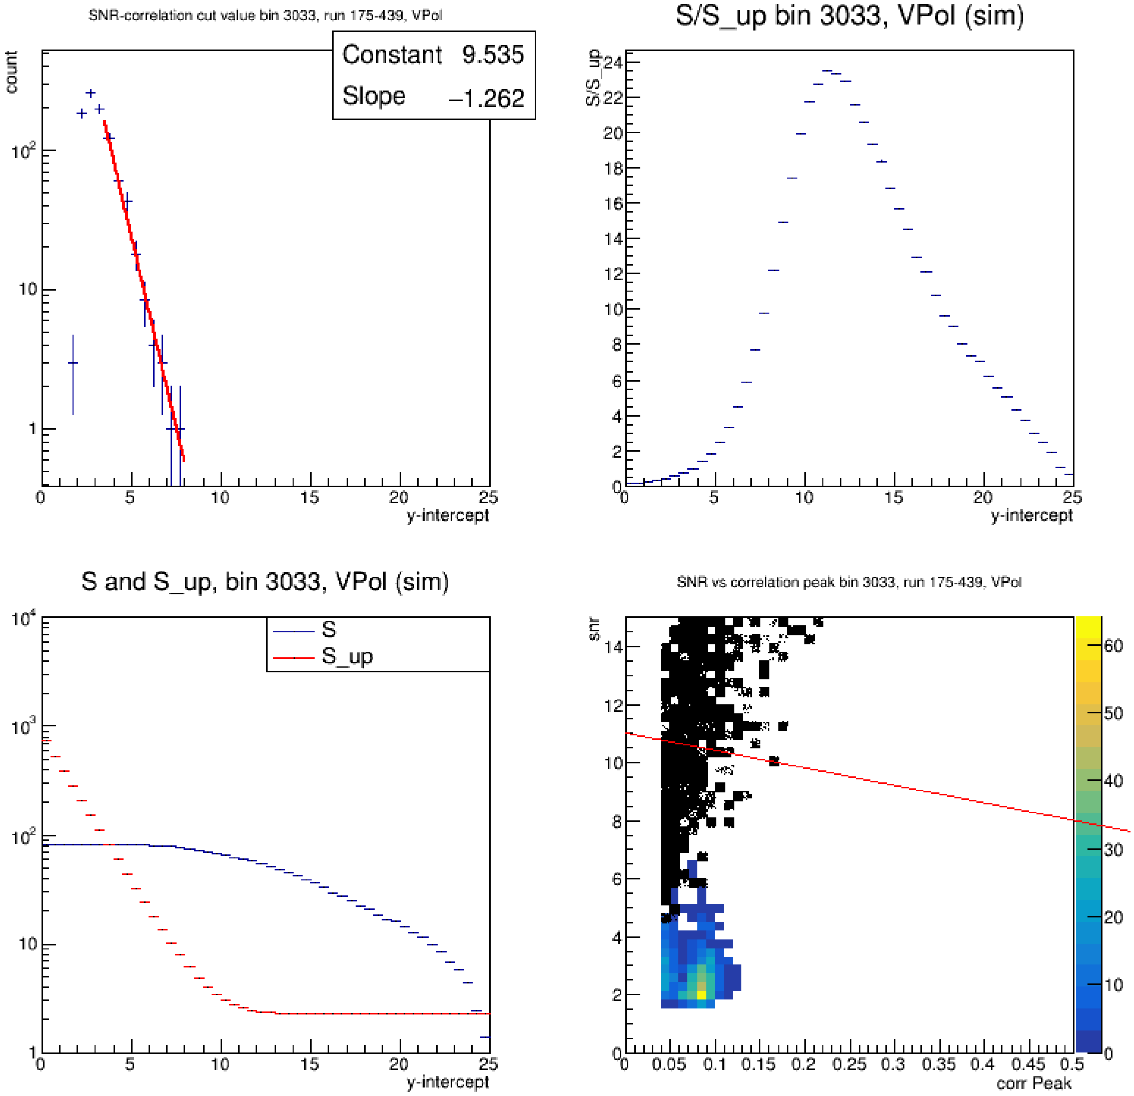

Figure F.17: Optimization of the linear discriminant cut, bin 3033. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

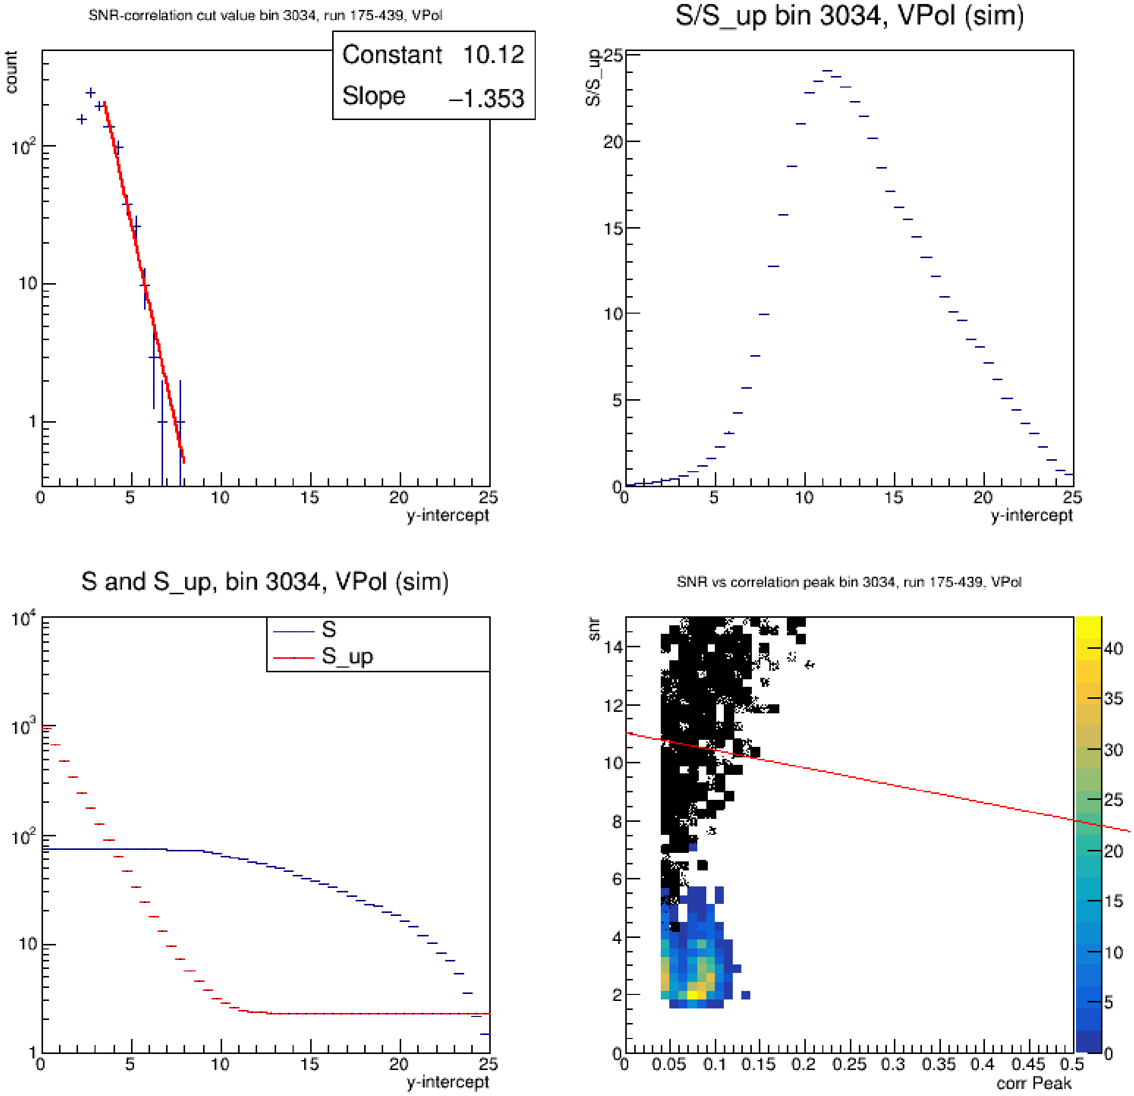

Figure F.18: Optimization of the linear discriminant cut, bin 3034. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

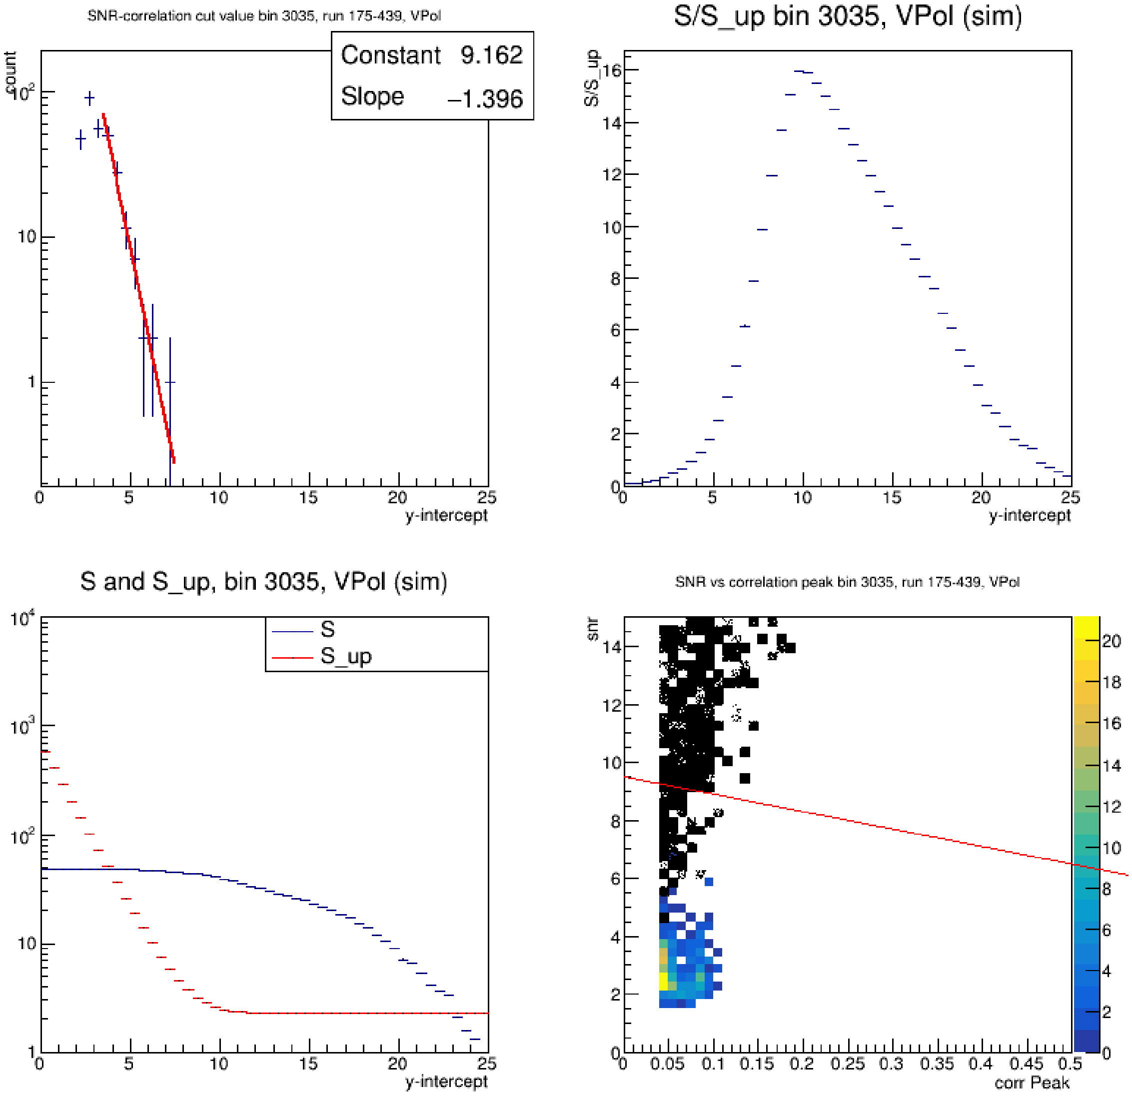

Figure F.19: Optimization of the linear discriminant cut, bin 3035. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 9.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

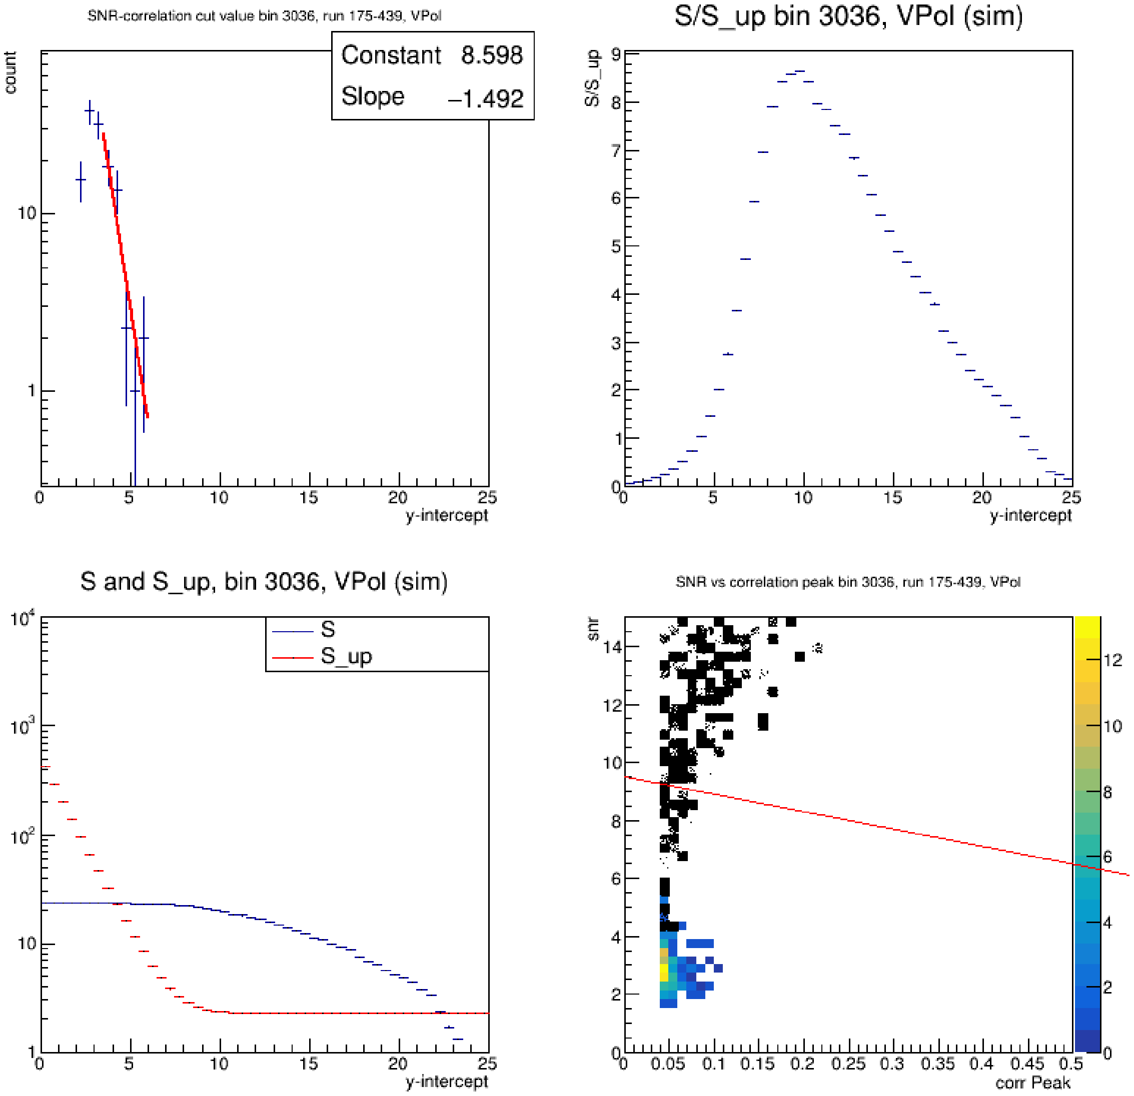

Figure F.20: Optimization of the linear discriminant cut, bin 3036. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 9.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

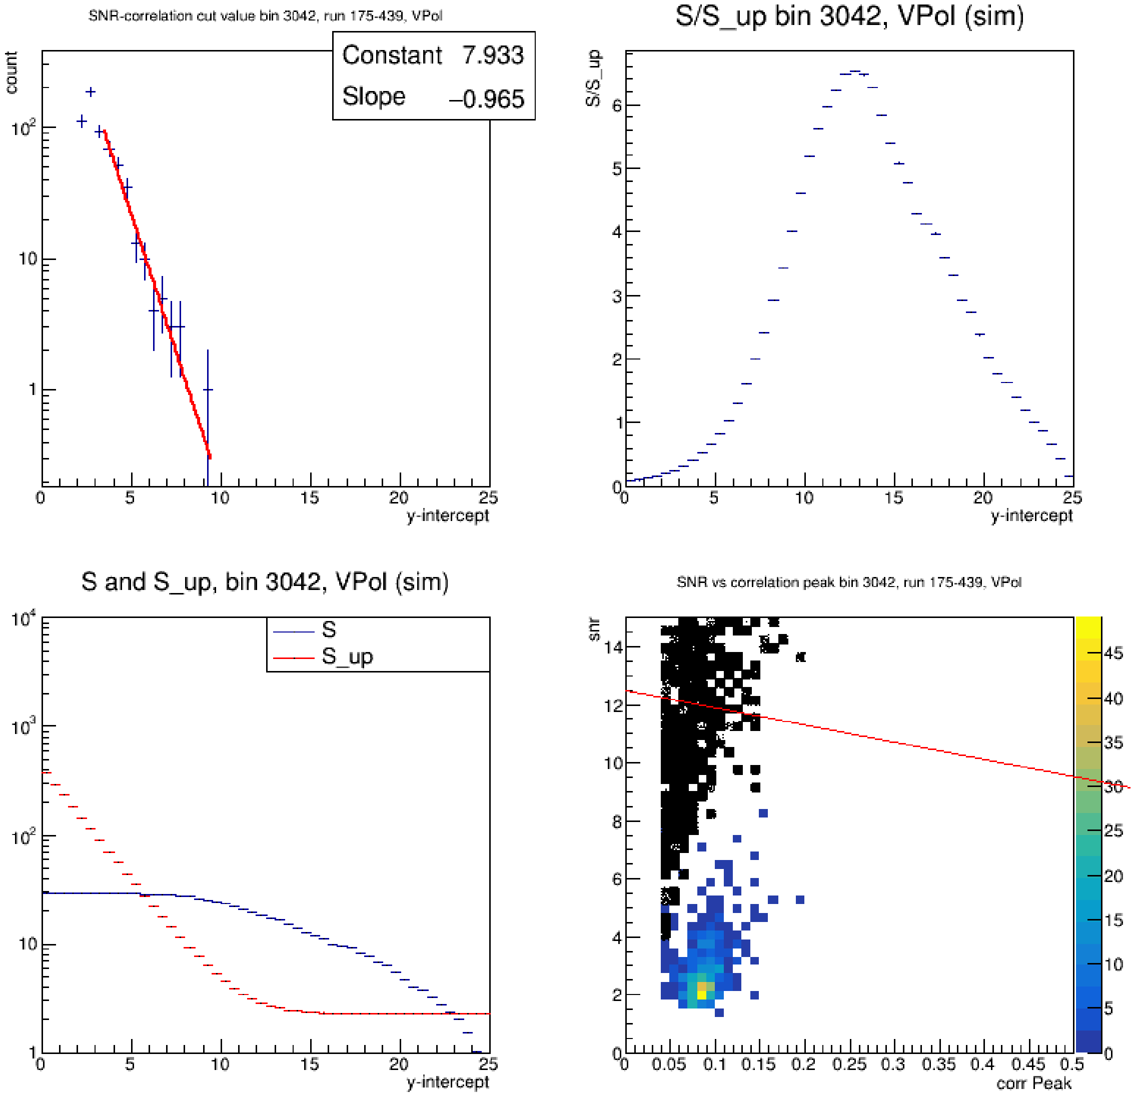

Figure F.21: Optimization of the linear discriminant cut, bin 3042. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 12.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

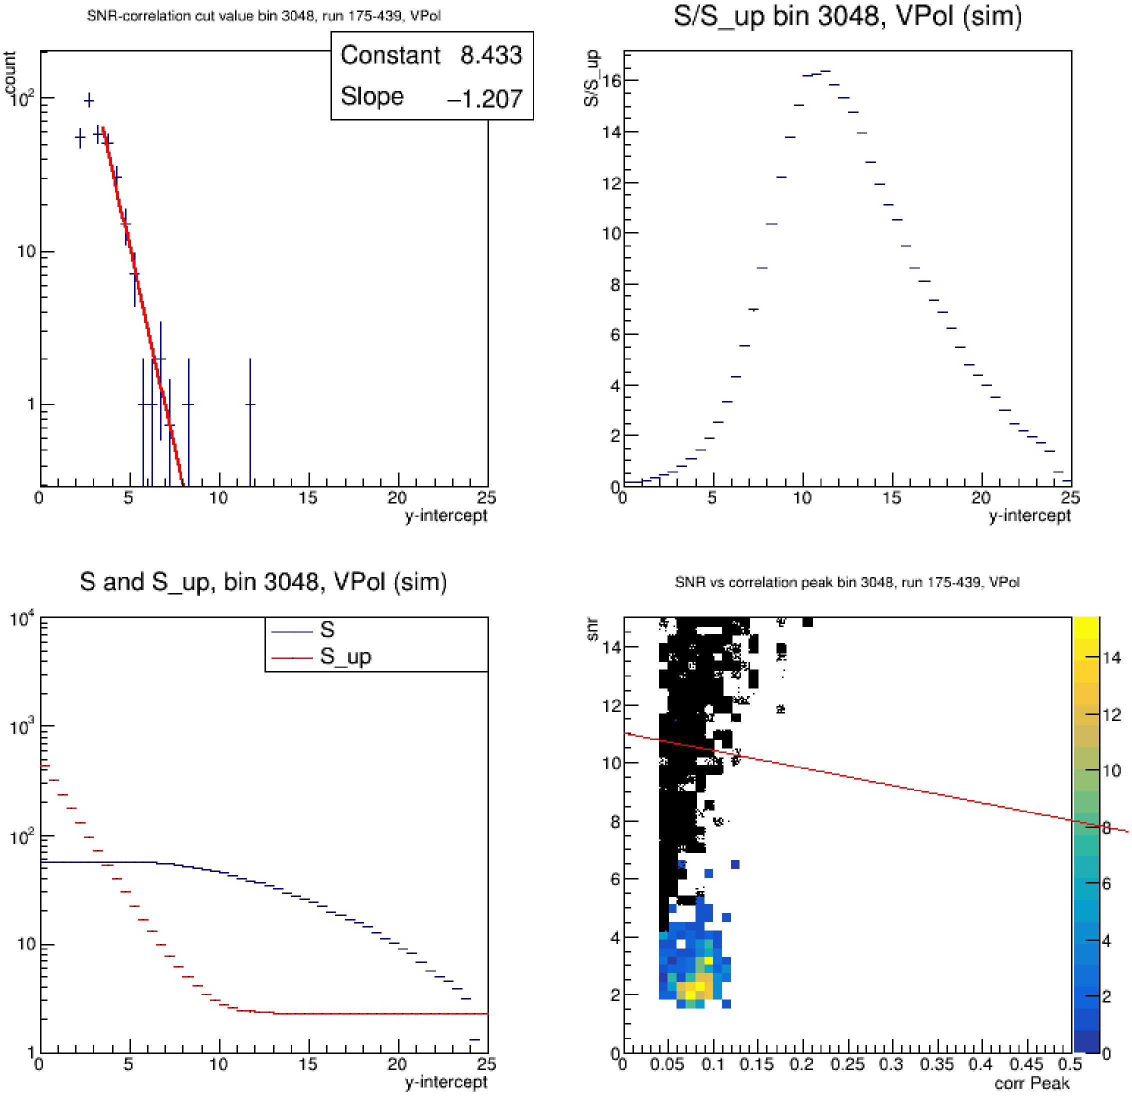

Figure F.22: Optimization of the linear discriminant cut, bin 3048. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

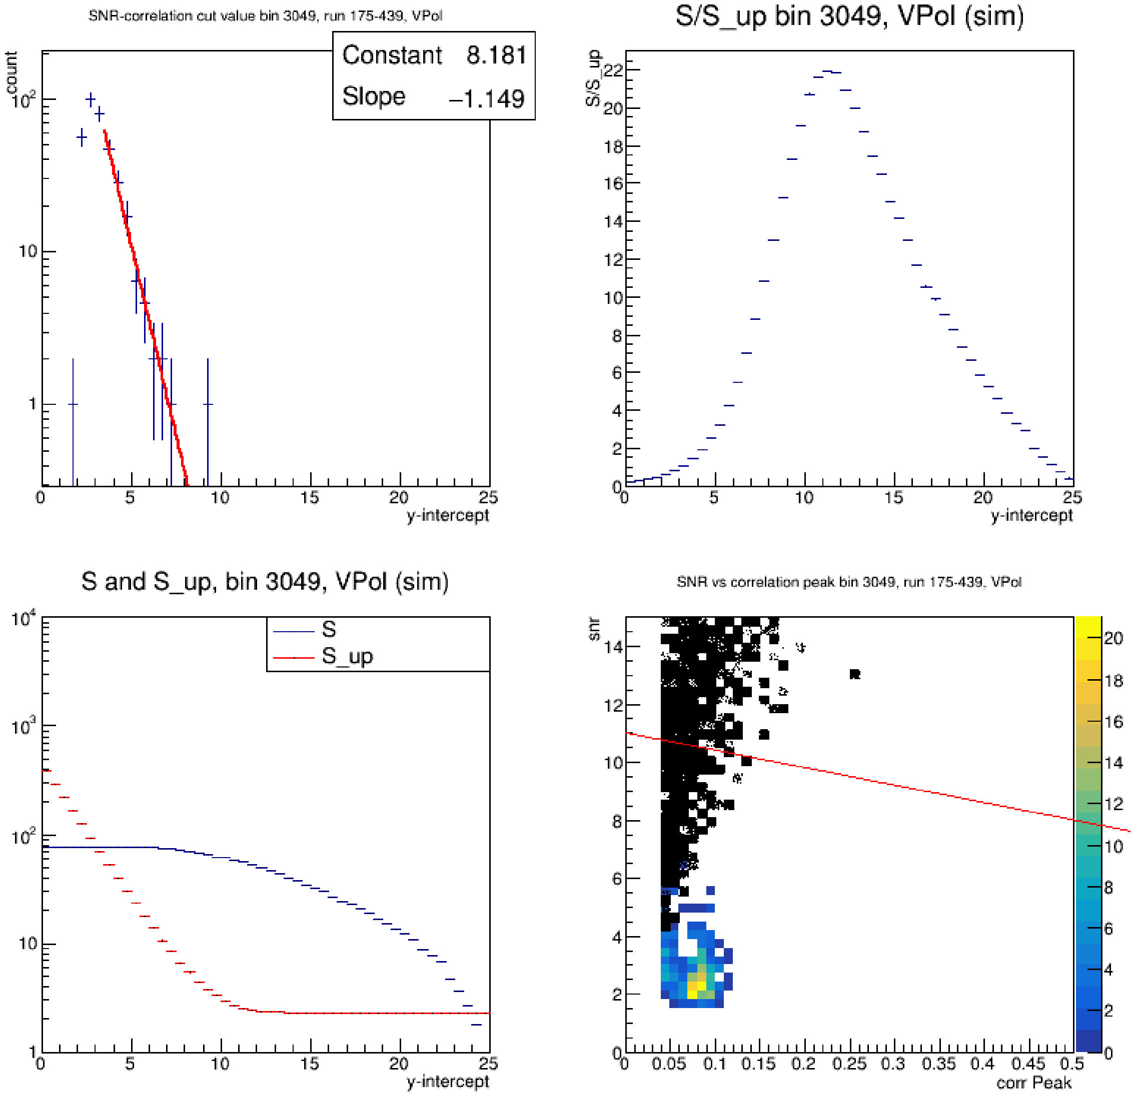

Figure F.23: Optimization of the linear discriminant cut, bin 3049. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 11.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

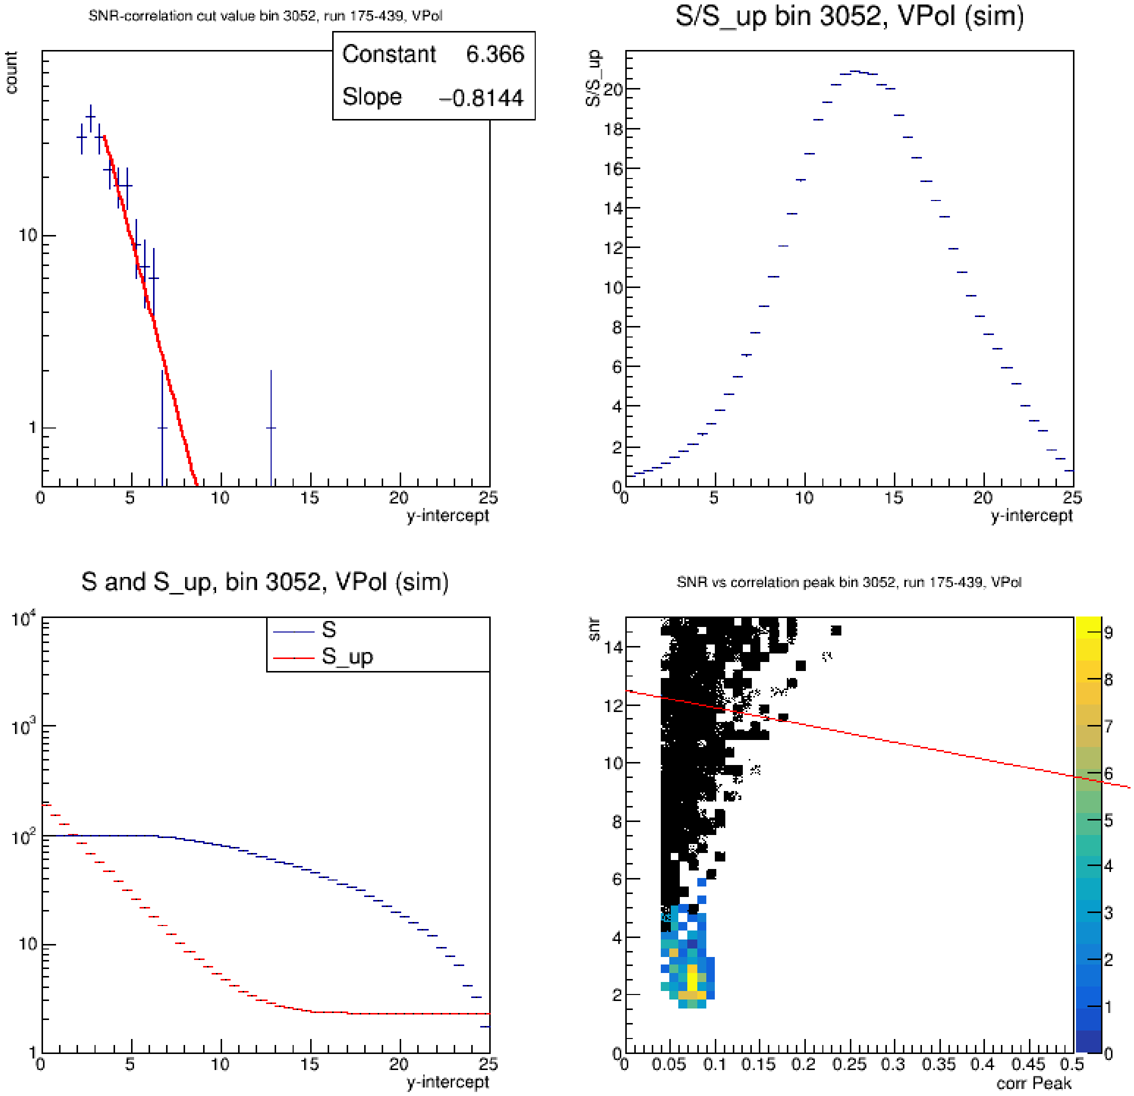

Figure F.24: Optimization of the linear discriminant cut, bin 3052. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 12.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

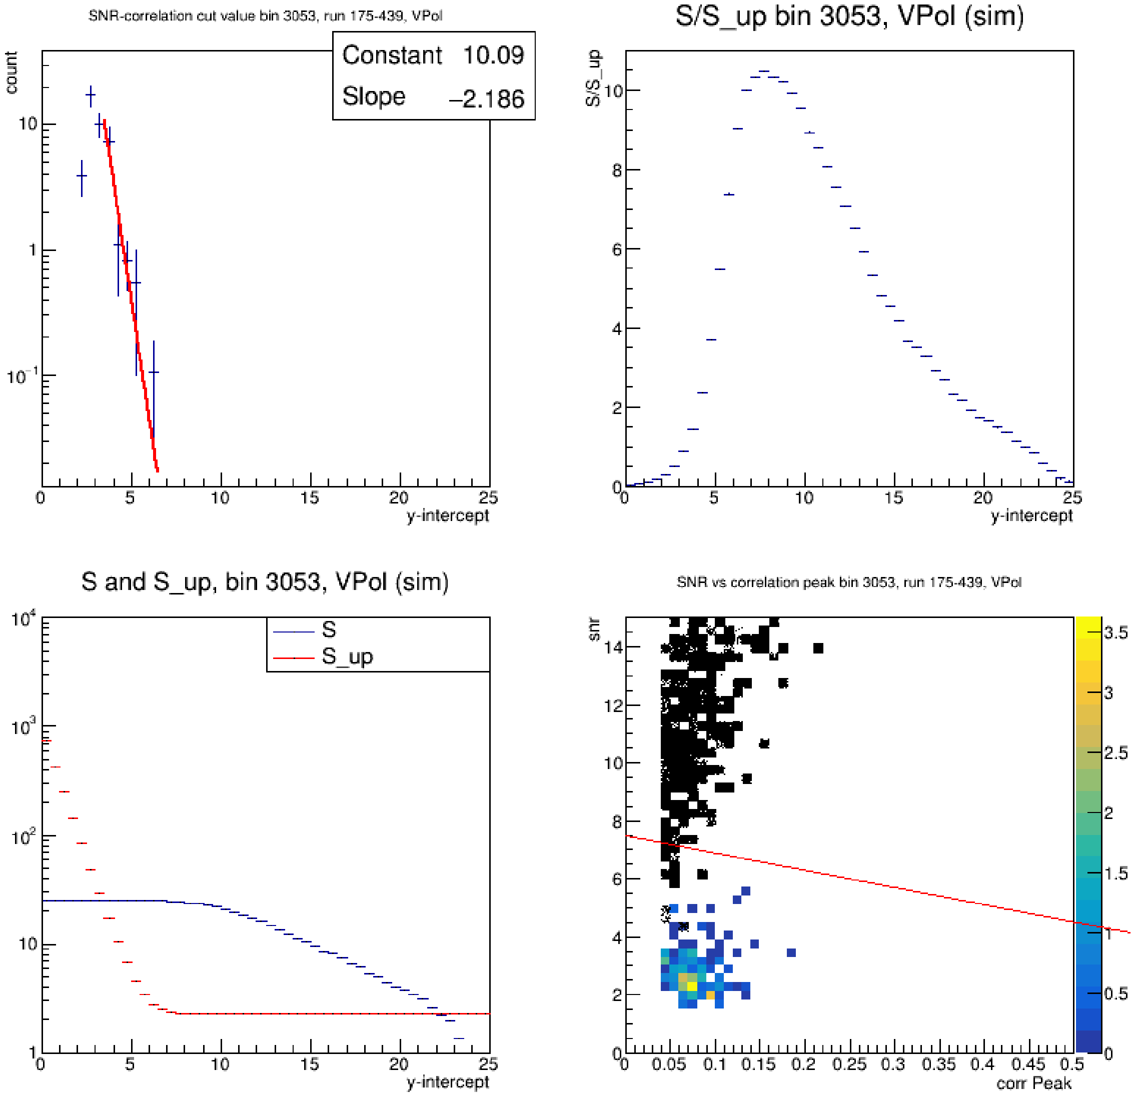

Figure F.25: Optimization of the linear discriminant cut, bin 3053. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 7.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

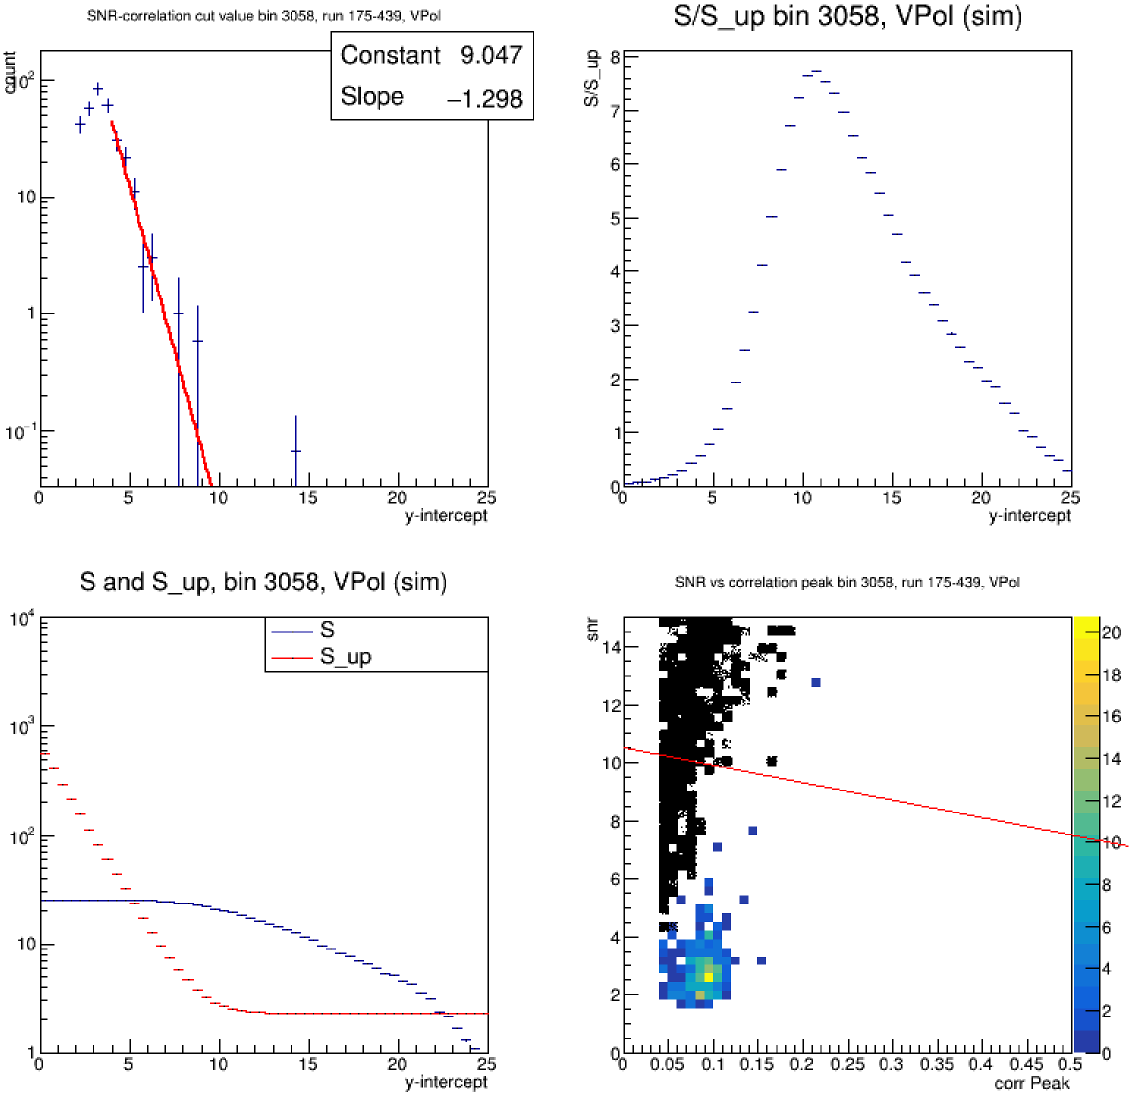

Figure F.26: Optimization of the linear discriminant cut, bin 3058. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 10.5. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.

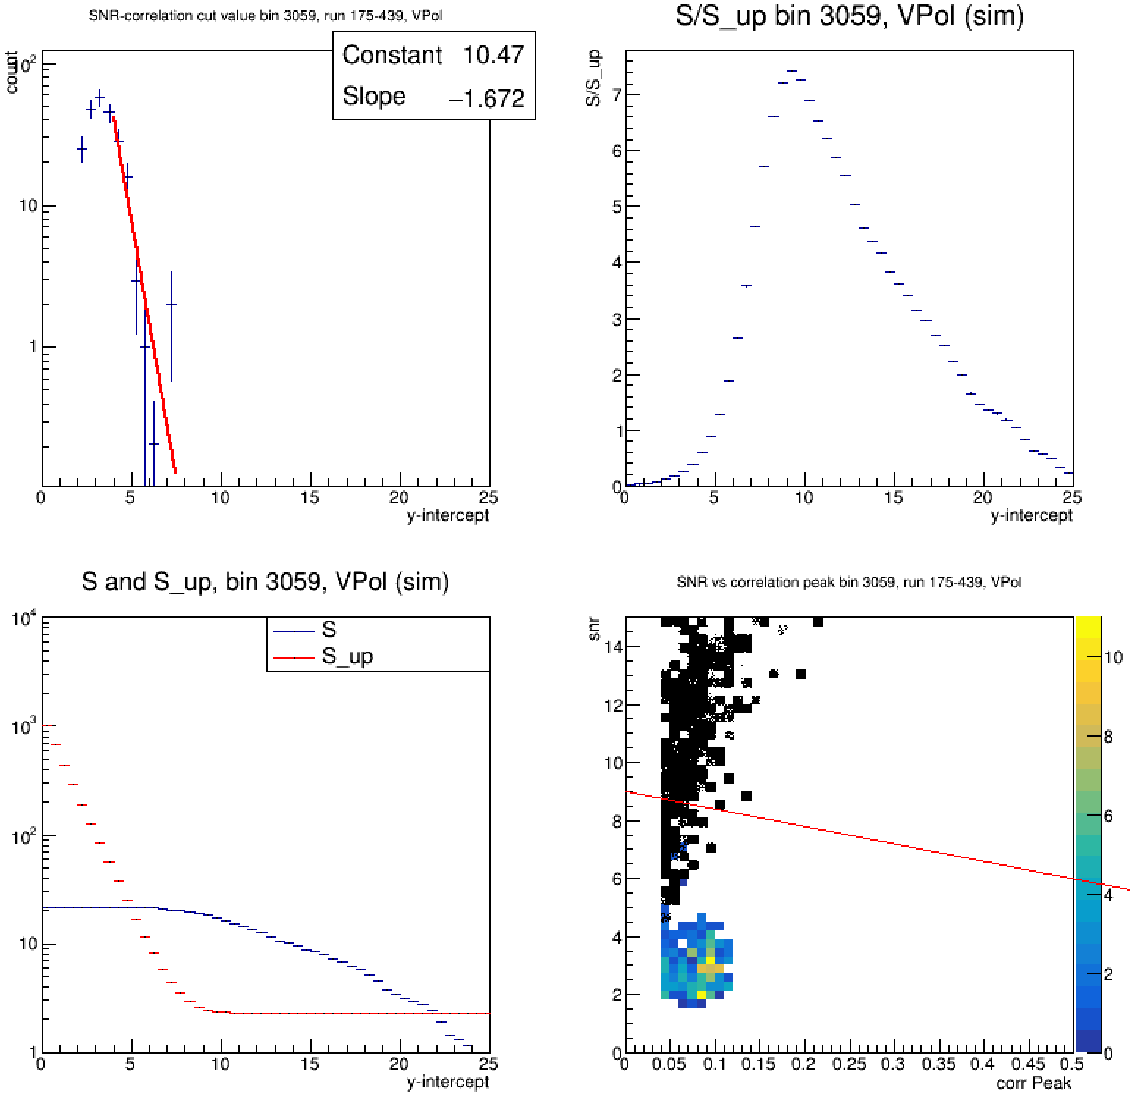

Figure F.27: Optimization of the linear discriminant cut, bin 3059. Top left: the exponential fit on the linear discriminant value of events from the 10% sample is used as a representation of the background in the bin of interest. Bottom left: the number S of simulated events passing cuts for values of the linear discriminant y-intercept R_i , and the value of the optimization parameter denominator S_{up}. Top right: The value of the optimization figure of merit, S/S_{up} vs. linear discriminant y-intercept; the optimized y-intercept in this case is 9.0. Bottom right: Simulated events (black) and 10% sample events (histogrammed in color) against correlation peak value and SNR. The red line is the linear discriminant.