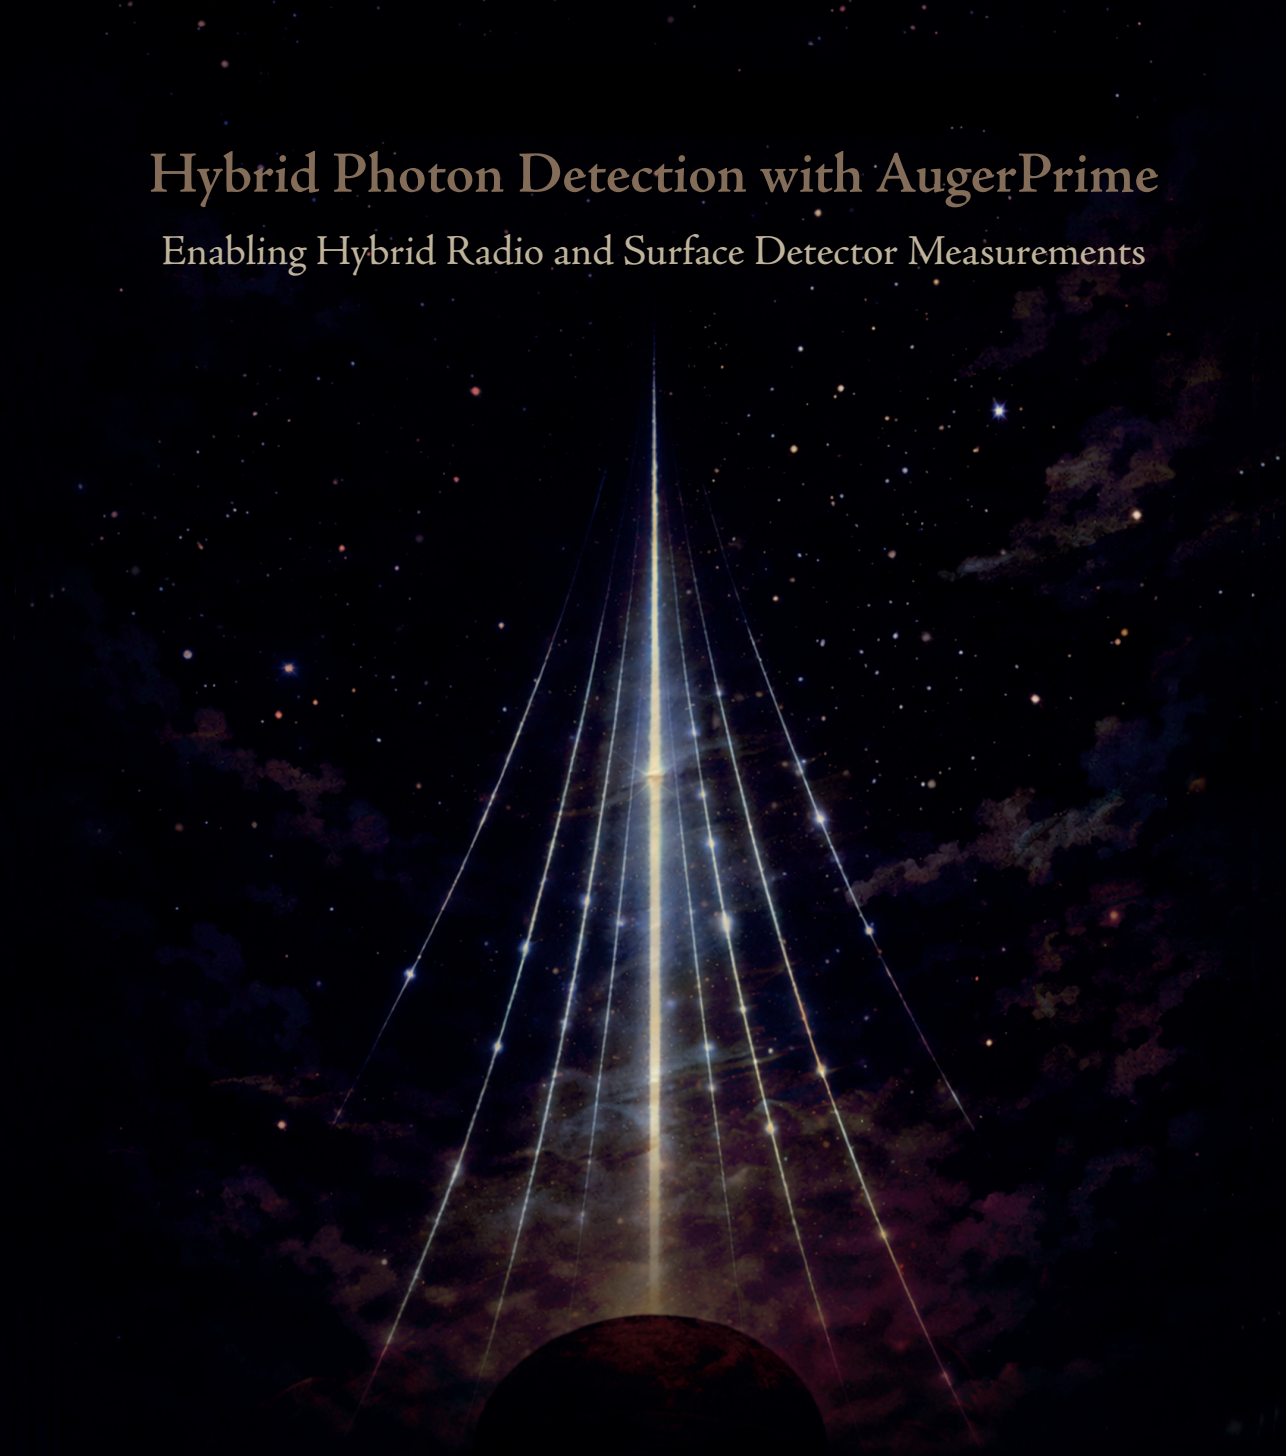
Hybrid Photon Detection with AugerPrime

Enabling Hybrid Radio and Surface Detector Measurements



I WAS JUST IN THE MODE, FOR LIGHT

Mohamed Ismaiel

Institute for Mathematics,
Astrophysics and Particle Physics

RADBOUD
UNIVERSITY
PRESS

Hybrid Photon Detection with AugerPrime

Enabling Hybrid Radio and Surface Detector Measurements

Mohamed Ismaiel



PIERRE
AUGER
OBSERVATORY

Mohamed Ismaiel
Hybrid Photon Detection with AugerPrime
Enabling Hybrid Radio and Surface Detector Measurements

Radboud Dissertation Series
ISSN: 2950-2772 (Online); 2950-2780 (Print)

Published by RADBOUD UNIVERSITY PRESS
Postbus 9100, 6500 HA Nijmegen, The Netherlands
www.radbouduniversitypress.nl

Design: Mohamed Ismaiel
Cover image: Mohamed Ismaiel
Production coordination & prepress: ProefschriftOntwerp.nl
Printing: ProefschriftMaken/Pumbo

ISBN: 9789465153476
DOI: 10.54195/9789465153476
Free download at: https://doi.org/10.54195/9789465153476

© 2026 Mohamed Ismaiel

RADBOUD
UNIVERSITY
PRESS

This is an Open Access book published under the terms of Creative Commons Attribution-Noncommercial-NoDerivatives International license (CC BY-NC-ND 4.0). This license allows reusers to copy and distribute the material in any medium or format in unadapted form only, for noncommercial purposes only, and only so long as attribution is given to the creator, see http://creativecommons.org/licenses/by-nc-nd/4.0/.

Hybrid Photon Detection with AugerPrime

Enabling Hybrid Radio and Surface Detector Measurements

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen
op gezag van de rector magnificus
prof. dr. J.M. Sanders,
volgens besluit van het college voor promoties
in het openbaar te verdedigen op
dinsdag 15 december 2026
om 12:30 uur precies

door

Mohamed Emam Ragab Ismaiel

geboren op 21 januari 1994
te Jeddah, Saoedi-Arabië

Promotor	Prof. dr. S.J. de Jong
Copromotor	Dr. H. Schoorlemmer
Manuscriptcommissie	Prof. dr. C.W.J.P. Timmermans (voorzitter)
	Prof. dr. A. Heijboer – Universiteit van Amsterdam
	Prof. dr. I. Mariș – Université libre de Bruxelles, België
	Dr. K.B. Mulrey
	Dr. J. Rautenberg – Bergische Universität Wuppertal, Duitsland

Contents

Introduction **3**

1 Ultra-High-Energy Cosmic Rays **5**
1.1 The Cosmic Ray Energy Spectrum 6
1.2 Extensive Air Showers 8
1.3 Cosmic Ray Composition 13
1.4 Ultra-High-Energy Photons 15
1.5 Searches for Ultra High Energy Photons 17

2 Pierre Auger Observatory and AugerPrime **23**
2.1 The Observatory 23
2.2 The AugerPrime Upgrade 26
2.3 The Radio Detector 28
2.4 The Surface Detector Triggers 29

3 The Radio Detector Monitoring and Commissioning **33**
3.1 The Deployment History of The Radio Detector 33
3.2 The Radio Detector Monitoring (MoRD) 34
3.3 Data Flow and Processing 36
3.4 Data Visualization 39
3.5 Evolution and Scaling 51
3.6 Using MoRD in the Commissioning Phase 54

4 Potential photon sensitivity of WCD and RD **61**
4.1 Threshold-Based Assessment of Hybrid Trigger Efficiency 61
4.2 Photon Reconstruction with WCD, Radio, and Hybrid Triggers 68
4.3 Exploring lower energies 75
4.4 Aperture for Ultra-High-Energy Photons 77

5 Data Reconstruction and Validation **81**
5.1 Water Cherenkov Detector Event Reconstruction 81
5.2 Radio Detector Event Reconstruction 83

5.2.1 $\overline{\text{Off}}\underline{\text{line}}$ Antenna Configuration . . . 86
5.3 Radio Detector Performance and Efficiency . . . 91

6 Data Quality Selection and Bias Assessment 97
6.1 Signal-Quality Selection . . . 97
6.1.1 Dual-Window Linear SNR . . . 99
6.1.2 Template Waveform Matching . . . 102
6.1.3 Polarization . . . 106
6.2 Events Quality Cuts . . . 110
6.2.1 Number of Radio Stations . . . 111
6.2.2 LDF Fit Quality . . . 113
6.2.3 Opening Angle . . . 115
6.2.4 Thunderstorm Flagging . . . 116
6.3 Angular Resolution . . . 119

7 Photon-hadron discrimination using WCD and RD 123
7.1 Simulation Reconstruction and Selection . . . 123
7.2 WCD and RD Observables . . . 126
7.3 Feature Engineering and Detrending . . . 128
7.4 Photon–hadron separation with $\boldsymbol{\delta'}$ . . . 130

8 Towards the search for ultra-high energy photons in data 135
8.1 The Hybrid Exposure . . . 135
8.2 Expected photon flux upper limits . . . 145
8.3 Validation: Burn Sample Proposal . . . 150

Summary and Outlook 155

Samenvatting 159

Appendix A: Simulation Sets 161

Appendix B: Aperture Derivation 163

Appendix C: OfflineB Configurations 165

Appendix D: Machine learning for photon–hadron discrimination 169

Appendix E: Data Management 180

Bibliography 183

Use of Generative AI 193

Acknowledgements **195**

Curriculum Vitae **197**

To my mother,

For years, I stayed close to you,
hesitant to leave because I wanted to be there for you.
But you were the one who encouraged me to go, to travel,
and to pursue the life you believed I was capable of.
You wanted to see the best of me, even when that meant
letting me go far from you.

This journey began because you believed I should not
give up my future to remain by your side.
I only wish you were here to see where that journey led.

And to my father,
who also left before he could see its end.

I carry you both with me to this final page.



Introduction

Since their discovery over a century ago, cosmic rays have been identified as relativistic subatomic particles—primarily protons and heavier atomic nuclei—that continuously bombard the Earth from outer space. These particles span an immense energy spectrum, with the most extreme examples reaching energies beyond 10^{20} eV, far exceeding the reach of human-made accelerators. While it is now established that these ultra-high-energy cosmic rays (UHECRs) are extragalactic in origin, the specific astrophysical sources and the mechanisms capable of accelerating particles to such macroscopic energies remain among the most significant mysteries in modern physics.

Cosmic rays reaching the Earth with energies beyond 10^{18} eV pose a fundamental challenge: their charged nature causes deflection by intervening magnetic fields, degrading directional information about their sources. Although the deflection is expected to be smaller at the highest energies, it can still be sufficient to hinder reliable source identification. Neutral particles (particularly photons) offer a more direct pathway to identifying cosmic accelerators. Gravitational deflections exist but are far below the instrumental pointing resolution of UHECR observatories. Ultra-high-energy photons serve a dual role: they act as messengers from source regions and provide critical tests of fundamental physics, including photon propagation limits and the expected cosmogenic flux from extragalactic cosmic ray interactions.

Current detection efforts face a significant limitation. Particle Detectors, the workhorse of ground-based cosmic ray observatories, perform well for near-vertical air showers but struggle with inclined geometries. At zenith angles exceeding 60°, atmospheric absorption eliminates most electromagnetic particles before they reach the surface. Hadronic air showers compensate through their muon content, but photon primaries (which generate few muons) produce signals too weak for reliable detection. Fluorescence detectors can directly observe the electromagnetic air shower development, but their limited duty cycle restricts the available exposure. This creates an observational blind spot precisely where radio detection techniques show greatest promise, motivating a hybrid approach combining radio antennas with surface detectors to achieve comprehensive sky coverage and enhanced photon sensitivity.

This thesis addresses this gap by leveraging the AugerPrime upgrade to the Pierre Auger Observatory, specifically the deployment of the Radio Detector array consisting of about 1660 antennas sensitive to 30–80 MHz air shower emissions. While the atmosphere absorbs electromagnetic cascade particles in inclined air showers, it remains largely transparent to radio frequency emission generated by these cascades. This allows the Radio Detector to capture the

electromagnetic component that particle detectors miss, specifically the Water Cherenkov Detectors in the case of the Auger Observatory, thereby providing a calorimetric measurement of air shower energy. By combining the muon-sensitive Water Cherenkov Detectors with the electromagnetically sensitive Radio Detectors, the hybrid measurement helps break the degeneracy between air-shower fluctuations and primary particle type, enabling robust photon–hadron separation.

This work establishes the sensitivity of this hybrid configuration to ultra-high-energy photons at large zenith angles and spans a large part of the experimental life-cycle: from commissioning the detector array through simulation studies and data quality selection, to developing a transparent hybrid photon–hadron separation strategy and proposing a framework for real-data analysis including exposure calculations and a burn-sample strategy for future validation.

Within the broader Pierre Auger Collaboration, this thesis makes several original contributions to the use of the AugerPrime Radio Detector for searches for ultra-high-energy photons. I developed the MoRD monitoring framework and associated commissioning diagnostics used to assess the status of the array and identify detector problems. I carried out simulation studies quantifying the photon sensitivity of the WCD, RD, and hybrid detector configurations, and validated the radio reconstruction on measured data, including investigations of detector and reconstruction effects such as the antenna-orientation issue. I developed station- and event-level quality-selection methods for reliable radio reconstruction and constructed and evaluated the hybrid photon–hadron separation observable based on the mismatch between the WCD and RD energy estimators. Finally, I developed the time-dependent exposure calculation and burn-sample framework for a future blinded search for photons in data. These contributions build on the detector hardware, standard reconstruction framework, simulations, and collaboration-wide data products developed within the Pierre Auger Collaboration.

The thesis is organized as follows: Chapter 1 provides an overview of ultra-high-energy cosmic rays, their energy spectrum, and the importance of photons in astroparticle physics. Chapter 2 describes the Pierre Auger Observatory and the AugerPrime upgrade, with emphasis on the Radio Detector. Chapter 3 details the monitoring and commissioning of the Radio Detector, including development of the monitoring tool. Chapter 4 investigates the sensitivity of the Water Cherenkov and Radio Detectors to ultra-high-energy photons using simulations. Chapter 5 presents the reconstruction and validation processes for data from both detector systems. Chapter 6 develops data quality selection criteria and assesses biases in the reconstruction of inclined air showers. Chapter 7 focuses on photon–hadron separation using hybrid observables, emphasizing a transparent cut-based selection validated on simulations. Chapter 8 outlines the framework for searching for ultra-high-energy photons in real data, including exposure calculations and a burn-sample strategy for validating the separation selection.

Chapter 1

Ultra-High-Energy Cosmic Rays

Ultra-high-energy cosmic rays (UHECRs), conventionally defined as cosmic rays with energies above 10^{18} eV, are the most energetic particles observed in nature. Their sources, acceleration mechanisms, and mass composition remain uncertain because of their very low flux and because charged particles are deflected by Galactic and extragalactic magnetic fields. This thesis focuses on ultra-high-energy photons, defined here as photons with energies $E_\gamma \geq 10^{17}$ eV; higher analysis-specific thresholds are stated explicitly where relevant [1].

UHE photons are expected in both astrophysical and exotic scenarios. In conventional bottom-up models, they are produced mainly through the decay of neutral pions generated in interactions of accelerated cosmic rays with matter or radiation fields. During UHECR propagation, photopion interactions of protons with cosmic background radiation can produce cosmogenic photons associated with the Greisen–Zatsepin–Kuzmin process [2]. In top-down scenarios, UHE photons may instead arise from the decay or annihilation of hypothetical super-heavy particles. Existing photon-flux limits therefore constrain both UHECR source and propagation models and scenarios beyond the Standard Model, including super-heavy dark matter and Lorentz-invariance violation [1, 3, 4].

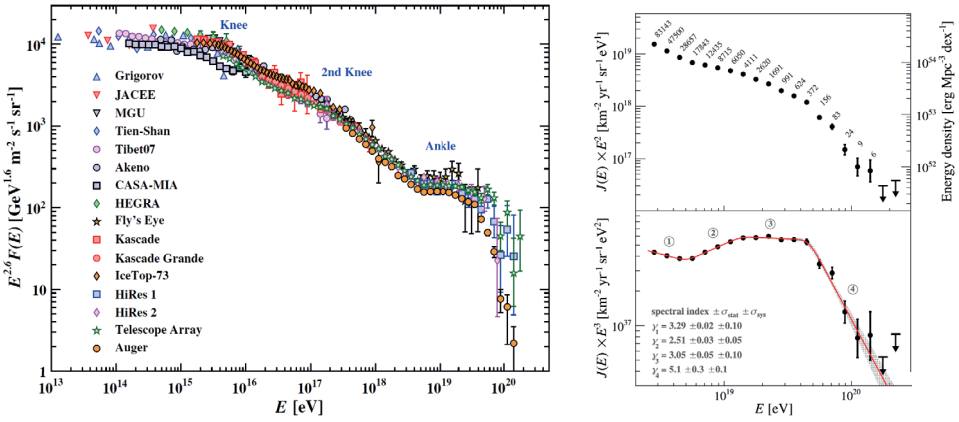
Because photons are electrically neutral, their arrival directions retain information about their production sites, although interactions with cosmic background radiation restrict the distance over which they can propagate. In the atmosphere, UHE photons produce predominantly electromagnetic, muon-poor air showers. Compared with hadron-induced showers of the same energy, they generally develop deeper in the atmosphere and exhibit different longitudinal profiles, ground-signal time structures, lateral distributions, and muon-to-electromagnetic signal ratios. At the highest energies, their development is additionally modified by the Landau–Pomeranchuk–Migdal effect and by conversion in the geomagnetic field. These characteristics form the experimental basis of the UHE-photon searches. In this thesis, ultra-high-energy photons are defined as photons with energies $E_\gamma \geq 10^{17}$ eV; when a particular search uses a higher threshold, that threshold is stated explicitly. Beyond their astrophysical significance, UHE photons may also provide probes of physics beyond the Standard Model [1].

1.1 The Cosmic Ray Energy Spectrum

Figure 1.1(a) displays the cosmic-ray energy spectrum above 10 TeV, measured by a range of ground-based observatories using extensive air shower (EAS) arrays. These experiments detect air showers — cascades of secondary particles produced by cosmic-ray interactions in the atmosphere — which span several kilometers at ground level, enabling sparse arrays to reconstruct their properties (see section 1.2). At lower energies (below ∼PeV), where the flux is sufficiently high, space-based instruments such as AMS and PAMELA complement these measurements through direct detection of cosmic-ray energy spectra and particle composition [5, 6]. The differential energy spectrum of cosmic rays can be, to first order, described by a broken power law:

$$J(E) \equiv \frac{d^4N}{dE\,dA\,dt\,d\Omega} \propto E^{-\gamma_i}$$
$$\text{for} \quad E_{\text{break},\,i-1} \leq E < E_{\text{break},\,i}, \quad i = 1, 2, \ldots, N+1,$$

where $J(E)$ is the differential intensity (number of particles per unit energy E, area A, time t, and solid angle Ω), and E is the particle energy (in eV). Thus $J(E)$ has units of $\text{eV}^{-1}\,\text{m}^{-2}\,\text{s}^{-1}\,\text{sr}^{-1}$. Beyond 1 PeV, four prominent features appear in the spectrum: the knee, the second knee, the



(a) (b)

Figure 1.1: (a) All-particle spectrum above 10 TeV, plotted as $F(E) = E^{2.6}J(E)$ to flatten the curve and highlight spectral-index changes (knee, second knee, ankle); vertical bars denote statistical uncertainties [7]. (b) Spectrum above 3 EeV from the Pierre Auger Observatory [8]: the top panel shows $E^2J(E)$ (energy density per decade), and the bottom panel shows $E^3J(E)$ to accentuate slope changes; shaded regions indicate fit uncertainties.

ankle, and the suppression. These features, observed as breaks or changes in slope, are typically linked to distinct astrophysical processes:

- **First Knee ($E \sim 4 \times 10^{15}$ eV)**
  The spectrum follows $J \propto E^{-\gamma}$ with $\gamma \approx 2.7$ up to about 4×10^{15} eV, so that a tenfold increase in E yields roughly a 500-fold decrease in flux. Above this energy, the slope

steepens to $\gamma \approx 3.1$, which is softer (i.e., $\gamma_2 > \gamma_1$) [9]. This bend might be attributed to the maximum energy achievable by galactic supernovae remnants for light particles such as protons [10]. In the range between $\sim 10\,\text{TeV}$ and $\sim 1\,\text{PeV}$ (just below the knee), the spectrum maintains this initial shallower slope.

- **Second Knee ($E \sim 8 \times 10^{16}\,\text{eV}$)**
  Located around $8 \times 10^{16}\,\text{eV}$, the second knee was identified by KASCADE-Grande [11] and confirmed by the Pierre Auger Observatory [8]. The spectral index steepens further to $\gamma \approx 3.3$, potentially marking the highest energies accessible to galactic accelerators for heavier nuclei or signaling the start of an extragalactic component.

- **Ankle ($E \sim 5 \times 10^{18}\,\text{eV}$)**
  Around $5 \times 10^{18}\,\text{eV}$, the spectrum hardens to a spectral index of approximately $\gamma \approx 2.5$–2.7. The ankle is commonly interpreted as a transition between different cosmic-ray source populations or components [12, 13]. Historically, it was also interpreted in the pair-production dip model as a propagation feature produced by extragalactic protons interacting with the CMB,

  $$p + \gamma_{\text{CMB}} \to p + e^+ + e^-,$$

  which requires an almost purely protonic composition above approximately $10^{18}\,\text{eV}$ [14]. However, measurements around the ankle indicate a mixed mass composition, including nuclei heavier than helium, while the measured fine structure of the energy spectrum is not reproduced by a proton-only propagation model. The classical proton-dip model is therefore excluded as the sole explanation of the ankle and of the UHECR flux observed at Earth [13, 8].

- **Suppression ($E \sim 5$–$6 \times 10^{19}\,\text{eV}$)**
  At the highest energies, the flux drops sharply. Two interpretations dominate:

  1. GZK Effect [2, 15]: Protons lose energy primarily via photopion production on the CMB, while heavy nuclei undergo photodisintegration [16].
  2. Maximum Acceleration Limit [17]: The flux cutoff may reflect the upper energy bound of extragalactic accelerators such as AGNs or GRBs.

  Both the Pierre Auger Observatory and the Telescope Array (TA) detect this suppression, though with slight differences in the energy threshold.

Figure 1.1(b) highlights the energy spectrum above 3 EeV measured by the Pierre Auger Observatory. In this energy range, the spectrum is described by successive power-law segments separated by the ankle, the instep, and the onset of the high-energy suppression. Region 1 covers energies below the ankle, while region 2 extends from the ankle to a further softening at approximately 13 EeV. This second break is known as the *instep*. At the instep, the spectral index changes from approximately 2.51 to 3.05. Region 3 extends from the instep to the stronger suppression near 50 EeV, above which the spectral index increases to approximately

5.1, defining region 4 [8]. These spectral features provide critical clues about cosmic ray origins and composition, as explored in section 1.3.

1.2 Extensive Air Showers

Extensive Air Showers (EAS) are cascades of secondary particles produced when a high-energy cosmic ray interacts with Earth's atmosphere, generating millions to billions of particles that spread over large areas as they propagate toward the ground [18, 19]. These air showers consist of electromagnetic (EM), muonic, and hadronic components. The EM and muonic components are sampled by ground-based particle detector arrays, while fluorescence telescopes detect the ultraviolet light emitted by atmospheric nitrogen along the shower track, together allowing inference of the primary cosmic ray's energy, composition, and arrival direction. The study of

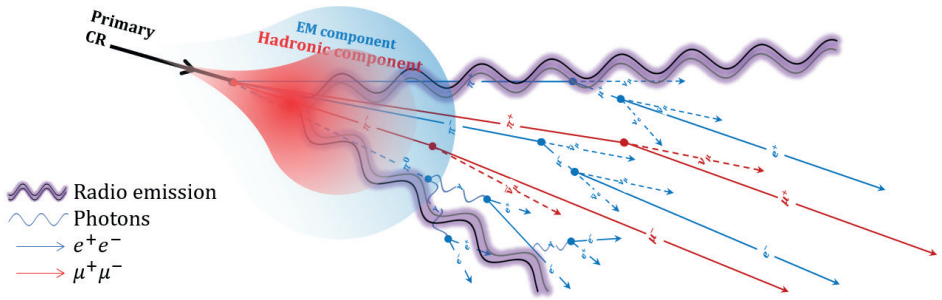


Figure 1.2: Schematic development of a typical EAS. A primary cosmic ray (black arrow) initiates a cascade in the atmosphere. Red region: hadronic core producing secondary particles. Blue region: EM component from photon-initiated cascades. Red arrows: muons from charged-particle decays. Purple wavy lines: coherent radio emission, detectable by antennas, dominated by geomagnetic and Askaryan mechanisms.

EAS is central to astroparticle physics for several reasons. First, their low flux makes UHECRs undetectable by space-based instruments, so EAS provide the only practical detection method. Second, they reveal particle interactions at energies far beyond those of terrestrial accelerators [20]. Finally, they probe potential new physics, such as tests of Lorentz invariance [4] and searches for exotic particles [21].

Particle Cascades

- **Hadronic Component:** The hadronic component forms the core of EAS, driven by nuclear interactions of the primary cosmic ray with atmospheric nuclei. These interactions produce hadrons, primarily pions and kaons, which either interact further or decay. Neutral pions (π^0) decay almost instantly into photons ($\pi^0 \to 2\gamma$), channeling approximately

one-third of the hadronic energy into the EM component per interaction generation [9]. Charged pions ($\pi^{\pm}$), with longer lifetimes, continue the hadronic cascade until they either interact or decay into muons and neutrinos ($\pi^+ \to \mu^+ + \nu_\mu$, $\pi^- \to \mu^- + \bar{\nu}_\mu$) [22]. The competition between decay and interaction depends on pion energy and atmospheric density, with a critical energy $E_{\pi,\mathrm{crit}} \approx 30\,\mathrm{GeV}$ at a density of $\sim 10^{-3}\,\mathrm{g\,cm^{-3}}$ [18].
Muons, due to their long lifetime and low interaction cross-section, penetrate the atmosphere with minimal attenuation and are therefore key observables for composition studies (see section 1.3). The Heitler-Matthews model, an analytical framework extending the Heitler splitting model to hadronic showers, predicts that the muon number scales as $N_\mu \propto E_0^{0.85}$, with heavier nuclei producing more muons by a factor of $A^{0.15}$ [22]. Despite this, current hadronic interaction models (e.g., EPOS-LHC, QGSJetII-04, Sibyll2.3d) predict 10 to 20% fewer muons than observed, the so-called *muon puzzle*, though this deficit has improved from earlier 30 to 60% discrepancies following new LHC forward-rapidity measurements [23].

- **Electromagnetic Component:** The EM component is initiated by photons from neutral pion decays, which produce electron-positron pairs via pair production. These pairs generate further photons through bremsstrahlung, amplifying the cascade [22]. The electron multiplicity at air-shower maximum is approximated by $N_e^{\max} \approx E_0/1.8\,\mathrm{GeV}$, where E_0 is the primary energy and 1.8 GeV represents the average energy per electron inferred from the Greisen cascade formula [22]. The longitudinal development of the air shower is modeled by the Gaisser-Hillas function:

$$N(X) = N_{\max}\left(\frac{X - X_0}{X_{\max} - X_0}\right)^{\frac{X_{\max}-X_0}{\lambda}} \exp\left(\frac{X_{\max} - X}{\lambda}\right) \tag{1.1}$$

where $N_{\max}$ is the maximum number of particles, $X_{\max}$ is the depth of air-shower maximum, X_0 is the depth of the first interaction, and λ is an interaction length parameter [24]. The lateral distribution of particles is described by the Nishimura-Kamata-Greisen (NKG) function:

$$\rho(r) = \frac{N_e}{r_M^2} \cdot C(s) \cdot \left(\frac{r}{r_M}\right)^{s-2} \cdot \left(1 + \frac{r}{r_M}\right)^{s-4.5} \tag{1.2}$$

where $\rho(r)$ is the particle density at distance r from the air-shower core, $r_M \approx 78\,\mathrm{m}$ is the Molière radius at sea level, and $C(s)$ is a normalization constant depending on air-shower age s [25]. The EM component deposits energy primarily through ionization, enabling nearly calorimetric energy measurements in fluorescence experiments:

$$E_{cal} = \alpha \int_0^\infty N(X) dX \tag{1.3}$$

where $\alpha \approx 2.2\,\mathrm{MeV/(g/cm^2)}$ is the average energy loss per particle [26]. An invisible energy correction, accounting for neutrinos and high-energy muons, is typically $\sim 10\%$ for proton primaries and slightly lower for heavier nuclei [27]. At ultra-high energies

($E > 10^{18}$ eV), the Landau-Pomeranchuk-Migdal (LPM) effect suppresses bremsstrahlung and pair production, causing deeper air-shower penetration [28]. Above 10^{19} eV, photons may interact with Earth's geomagnetic field, forming electron-positron pairs that initiate preshowers, altering the cascade [29].

Radio Emission

In this work, the radio signal is one of the primary observables: the air-shower energy is reconstructed from the measured radiation energy which will be used to search for photon-initiated events. A short overview of the underlying emission physics, therefore, sets the stage for the analysis that follows. At MHz frequencies, the signal is produced almost exclusively by the electromagnetic part of the air shower through two coherent mechanisms [30], sketched in Figure 1.3(a):

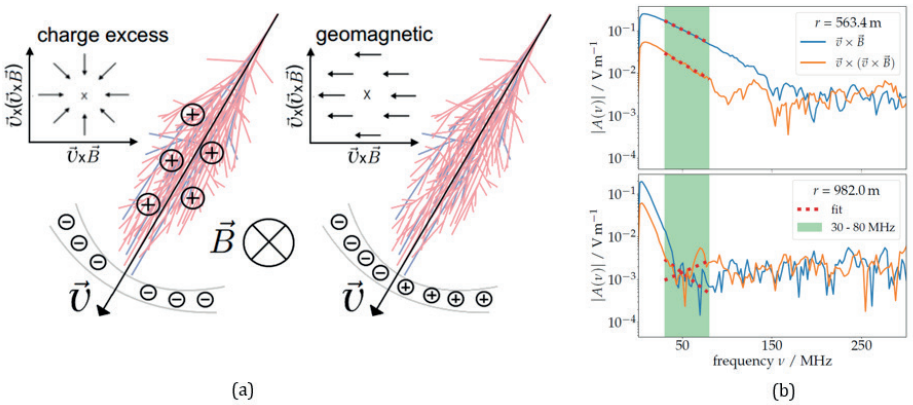


Figure 1.3: (a) Two emission mechanisms of coherent MHz pulses from extensive air showers and their polarization: charge excess emission (left) and the geomagnetic emission (right). (b) Spectra of two simulated pulses for observers situated on the $\vec{v} \times (\vec{v} \times \vec{B})$-axis at different distances to the MC air-shower axis. For each pulse, the frequency content in the $\vec{v}\times\vec{B}$ and $\vec{v}\times(\vec{v}\times\vec{B})$ directions is shown. Figures were obtained from [31, 32, 33].

- **Geomagnetic emission.** This mechanism occurs when electrons and positrons propagating along the air-shower axis—represented by the downward-directed velocity vector $\vec{v}$—are separated by the Earth's magnetic field $\vec{B}$. The Lorentz force

$$\vec{F} = q\,(\vec{v} \times \vec{B})$$

drives oppositely charged particles in opposite directions, creating a transverse current that evolves as the air shower develops and thus radiates. The resulting electric field is oriented[1]:

$$\vec{E}_{\text{geo}} \;\propto\; -\vec{v} \times \vec{B}.$$

[1]The minus sign follows from our choice that $\vec{v}$ points downward.

This geomagnetic component is typically the dominant contribution to the radio signal from air showers, and its linear polarization lies perpendicular to both $\vec{v}$ and $\vec{B}$. The geomagnetic emission is influenced by changes in particle numbers, making it crucial to observe a fully developed air shower for clear radio emission. Air showers that peak near or below the ground are not fully developed, so radio emissions are easier to detect in inclined air showers, with maxima farther from the ground. However, the geomagnetic emission disappears when the air shower's arrival direction is parallel to Earth's magnetic field [34].

- **Charge-Excess Emission:** Also known as the Askaryan effect [35], this secondary contribution arises from a net negative charge excess in the air-shower front (caused by processes like Compton scattering and positron annihilation). As the number of electrons N_e and their velocity $\vec{v}$ vary with time, the resulting time-varying current produces radiation described as:

$$\vec{E}_{\text{Ask}} \propto \frac{d}{dt}\big(N_e\,\vec{v}\big).$$

In terms of observed polarization, this effect often appears as

$$\vec{E}_{\text{ce}} \propto -\hat{r},$$

where $\hat{r}$ is the unit vector pointing from the air-shower axis to the observer, giving rise to a radial polarization pattern around the axis. While generally contributing around 10-20% of the total radio signal [36], it can be distinguished from the dominant geomagnetic component by its polarization characteristics.

The radio emission from air showers exhibits coherence at frequencies below $\sim 100\,\text{MHz}$. Coherence occurs when the wavelength of the emitted radiation exceeds the dimensions of the emission region, as shown in Figure 1.3(b). The spectral field strength typically peaks between 30-80 MHz and falls off at higher frequencies or far from the air-shower core as coherence is lost. Moreover, the radio signal strength observed at ground level decreases with increasing distance from the air-shower axis. A distinctive feature in the radio emission pattern is the formation of a Cherenkov ring [37]. When the refractive index of air (approximately $n \approx 1.0003$ at sea level and decreasing with altitude) is considered, coherent radio emission is enhanced at the Cherenkov angle θ_C given by:

$$\cos\theta_C = \frac{1}{n\beta} \tag{1.4}$$

where β is the particle velocity relative to the speed of light.
Gradients in temperature, pressure, and humidity change the air's refractive index, and those refractive-index variations [38], in turn, alter the radio-emission pattern. Moreover, the radio emission amplitude correlates strongly with the air-shower maximum depth $X_{\max}$. The peak radio field amplitude grows approximately linearly with the primary energy, while the total radiated energy scales quadratically [39].

Detection Techniques

Extensive air showers produced by UHECR can be detected using three complementary techniques: particle, optical, and radio detectors. Each method provides unique insights into the air-shower development, with different sensitivities, resolutions, and operational constraints.

Ground-Based Particle Detectors: These detectors sample the lateral distribution of secondary air-shower particles at the Earth's surface. Typically arranged as large arrays of many modular stations that allow for continuous operation in all weather conditions or even if a few individual units are offline.

- **Scintillation Detectors:** Plastic scintillators emit light when traversed by charged particles, and the resulting light is converted into an electrical signal by photomultiplier tubes or silicon photomultipliers. The signal amplitude provides an estimate of the energy deposited by the shower particles. Scintillation detectors form the surface array of the Telescope Array and are also used in AugerPrime, the Phase II upgrade of the Pierre Auger Observatory. In AugerPrime, a Scintillator Surface Detector is installed above each water-Cherenkov station. The different responses of the scintillator and water-Cherenkov detectors to the electromagnetic and muonic shower components enable improved separation of these components [40].

- **Water-Cherenkov Detectors (WCDs):** These detectors consist of purified water volumes viewed by photomultiplier tubes. Charged particles moving faster than the speed of light in water ($v > c/n$) emit Cherenkov radiation at an angle satisfying $\cos\theta_c = 1/(n\beta)$, with a light yield proportional to $1 - 1/(n^2\beta^2)$. WCDs are sensitive to both the muonic and electromagnetic components of an air shower. Muons, electrons, and positrons can produce Cherenkov light directly, while shower photons contribute after interactions such as pair production and Compton scattering generate charged secondary particles in the detector. These charged secondaries then emit the Cherenkov light recorded by the photomultiplier tubes. Thus, photons can contribute substantially to the WCD signal, although the detected optical light is still Cherenkov radiation produced by charged particles [41].

Optical Detectors: These instruments observe light emitted during the development of the air shower in the atmosphere, offering access to its longitudinal structure.

- **Fluorescence Detectors:** As air-shower particles excite nitrogen molecules in the atmosphere, UV fluorescence light (300–400 nm) is emitted isotropically. Telescopes capture this light, providing a near-calorimetric view of the air-shower development and enabling accurate measurements of $X_{\max}$ [12].

- **Imaging Atmospheric Cherenkov Telescopes (IACTs):** These detect Cherenkov light emitted in air by relativistic air-shower particles. The light forms a narrow cone and lands as a concentrated light pool on the ground (radius ∼120–150 m) [42]. IACTs

provide directional sensitivity and are widely used in gamma-ray astrophysics but are also valuable for high-energy cosmic-ray studies.

Radio Detectors: Radio detection exploits the coherent emission from extensive air showers, primarily due to geomagnetic emission. Most experiments operate in the 30–80 MHz band. The antenna is the key component in these radio detection setups because it serves as the primary sensor, and its design significantly affects the overall sensitivity. Examples of commonly used designs [43]:

- **Dipole Antennas (e.g., Butterfly):** These simple and low-cost antennas have broad angular acceptance and flat frequency response, making them ideal for large-scale arrays.
- **Log-Periodic Dipole Arrays (LPDAs):** Comprising multiple dipoles of varying lengths, LPDAs offer wide bandwidth and directionality, enhancing signal-to-noise performance.
- **Short Aperiodic Loaded Loop Antennas (SALLAs):** These non-resonant loops are mechanically robust and have a flat frequency response. While their sensitivity is slightly lower, their ease of calibration and durability make them suitable for large-scale deployments.

Each detection method offers distinct advantages and trade-offs in terms of resolution, coverage, and sensitivity. Table 1.1 presents a side-by-side comparison across key performance metrics. The synergy between particle, optical, and radio methods enables a more complete characterization of extensive air showers. Radio emission is generated almost entirely by the electromagnetic shower component, whereas a WCD records a mixture of electromagnetic and muonic signals. Combining the two measurements therefore provides sensitivity to the relative electromagnetic and muonic content of the shower. This complementarity is particularly useful in searches for photon-induced air showers, which are expected to be electromagnetically dominated and muon-poor.

1.3 Cosmic Ray Composition

The nuclear composition of cosmic rays is crucial for understanding their origins, propagation, and acceleration mechanisms. Composition is inferred from properties of extensive air showers (EAS, section 1.2), particularly the atmospheric depth of air-shower maximum ($X_{\max}$), where the number of secondary particles peaks. Below the ankle ($\sim 10^{15}$ eV), cosmic rays are dominated by light nuclei, primarily protons and helium. Above 10^{19} eV, measurements of the mean depth ($\langle X_{\max} \rangle$) and its distribution width ($\sigma(X_{\max})$) by the Pierre Auger Observatory show a decrease in $\langle X_{\max} \rangle$ and reduced fluctuations, indicating a shift toward heavier nuclei, such as intermediate and iron-like nuclei, which interact earlier in the atmosphere and produce more uniform air showers [49, 50]. As shown in Figure 1.4, Pierre Auger Observatory data suggest a lightening of composition between 10^{17} eV and $\sim 10^{18.3}$ eV, followed by a shift to heavier

Feature	Particle detectors	Radio detectors	Fluorescence detectors
Duty cycle	$\sim 100\%$ [41]	$\sim 100\%$ [44]	10–15% [45]
$X_{\max}$ **resolution**	$< 25\,\mathrm{g\,cm^{-2}}$ for $E > 2 \times 10^{19}\,\mathrm{eV}$ (indirect reconstruction) [46]	$< 15\,\mathrm{g\,cm^{-2}}$ near the upper end of $10^{17.5}$–$10^{18.8}\,\mathrm{eV}$ [39]	15–$20\,\mathrm{g\,cm^{-2}}$ for $E \gtrsim 10^{18}\,\mathrm{eV}$ [47]
Energy resolution	15–25% [41]	15–25% [44]	10–15% [45]
Measured component	Electromagnetic and muonic [47]	Electromagnetic [44]	Electromagnetic [47]
Angular resolution	$\sim 1^\circ$ [41]	$< 1^\circ$ [48]	0.5–1° [45]
Atmospheric dependence	Low [41]	Moderate [44]	High [45]
Operational constraints	All weather	Radio-frequency interference	Clear, moonless nights
Calibration complexity	Moderate [41]	High [48]	High [45]

Table 1.1: Comparison of particle, radio, and fluorescence detectors in extensive-air-shower observations. The quoted resolutions are representative values and depend on energy, event geometry, detector configuration, reconstruction method, and event-selection criteria.

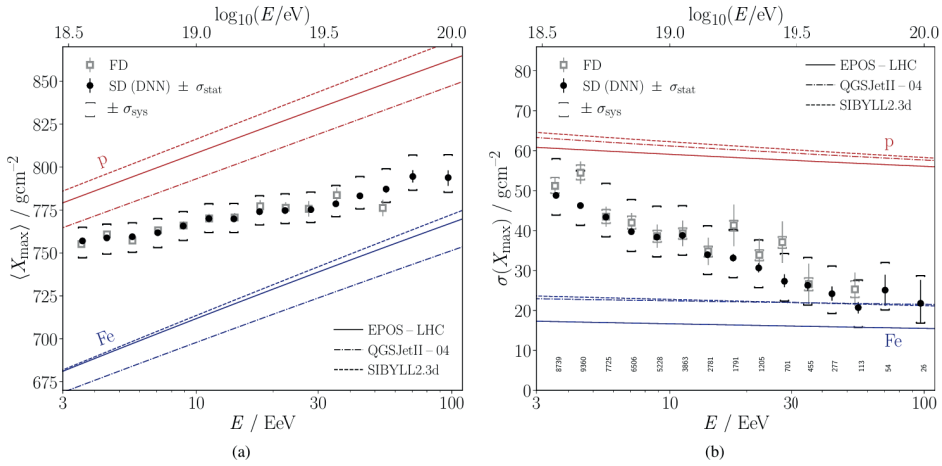


Figure 1.4: Energy evolution of (a) $\langle X_{\max} \rangle$ and (b) $\sigma(X_{\max})$ measured by the Pierre Auger Observatory, compared with model predictions for proton and iron primaries [51].

nuclei at higher energies [50]. The Telescope Array reports similar trends, though with larger uncertainties [52]. This composition shift impacts cosmogenic photon production (section 1.4), as heavier nuclei undergo photodisintegration in background radiation fields, reducing photon yields compared to pion production by lighter nuclei [53].

Composition analyses rely on hadronic interaction models (e.g., EPOS-LHC, QGSJetII-04, Sibyll2.3d), which extrapolate accelerator data to ultra-high energies,[2] where direct calibration is unavailable [55]. The trend toward heavier nuclei challenges single-component source models, suggesting either a source preference for accelerating heavy nuclei via rigidity-dependent mechanisms or significant photodisintegration during propagation [56].

For cosmogenic photon searches, composition governs secondary photon production rates and cosmogenic photons are expected to be approximately isotropically distributed. Sensitivity to cosmogenic photons set the benchmark for identifying photon-induced air showers [57]. Current upper limits on UHE photon fluxes rule out models predicting high photon fractions, but detecting even a small flux could reveal new physics [58].

1.4 Ultra-High-Energy Photons

Inferring cosmic-ray composition from extensive air showers (EAS, section 1.2) involves uncertainties arising from the dependence on hadronic interaction models (section 1.3). Ultra-high-energy photons provide a complementary probe, as photon-induced air showers are generally deeper and more muon-poor than hadronic showers. These differences provide observables for identifying photon candidates, as discussed in section 1.5. The production and propagation of UHE photons are described in section 1.4 and section 1.4, respectively.

[2]The relevant scale is the centre-of-mass energy, which exceeds LHC energies [54] by a factor of ~ 30.

Astrophysical Sources

High-energy photons are expected in two main scenarios:

- **Bottom-up scenario:** UHECRs are accelerated in astrophysical sources with intense magnetic fields. The Hillas criterion relates the maximum energy $E_{\max}$ to the source's size R and magnetic field B:

$$E_{\max} \sim \beta Z e B R \tag{1.5}$$

  where Z is the particle's charge, e is the elementary electric charge, and $\beta = v/c \leq 1$ [17]. However, dense photon fields or strong magnetic fields may render sources photon-opaque.

  - **Galactic Sources.** Galactic sources like the Crab pulsar or flaring magnetars emit photons up to the TeV–PeV range, but intense magnetic fields induce pair cascades, preventing UHE photon escape ($\geq 10^{18}\,\text{eV}$) [21]. Supernova remnants and stellar clusters may act as PeVatrons but are unlikely to reach EeV energies due to confinement and energy-loss constraints.

  - **Extragalactic Accelerators.** GRBs and AGNs are prime candidates for UHE photon production via pion production from proton-photon interactions (GZK effect):

$$p + \gamma \to \Delta^+ \to p + \pi^0 \to p + 2\gamma \quad (\text{BR} = 99.8\%[59]) \tag{1.6}$$

$$p + \gamma \to \Delta^+ \to p + \pi^0 \to p + \gamma + e^- + e^+ \tag{1.7}$$

    The energy threshold is:

$$E_p^{\text{th}} \approx \frac{m_\Delta^2}{2E_{\gamma_{\text{CMB}}}} \approx 5 \times 10^{19}\,\text{eV} \tag{1.8}$$

    The proton retains $\sim 80\%$ of its energy, while the neutral pion (π^0) carries $\sim 20\%$, decaying into two photons, each with $\sim 10\%$ of the initial proton energy ($E_\gamma \approx 5 \times 10^{18}\,\text{eV}$) [60]. GRBs produce photons via internal/external shocks in relativistic jets [61], while AGNs generate photons near supermassive black holes via pion production [62, 21]. These channels predict a photon flux of $\lesssim 0.1\,\text{km}^{-2}\,\text{yr}^{-1}$ above $10^{19}\,\text{eV}$ [63].

- **Top-down scenario:** UHE photons may arise from exotic processes, such as superheavy dark matter (SHDM) decay [64]. A heavy dark matter particle X may decay into photon pairs ($X \to 2\gamma$) with $E_\gamma \sim m_X/2$, or into a photon and other Standard Model particles ($X \to \gamma + \text{SM particles}$). For example, a particle with $m_X \approx 10^{15}\,\text{GeV}$ produces $\sim 10^{23}\,\text{eV}$ photons, which degrade to $\sim 10^{19}\,\text{eV}$ after propagating $\sim 100\,\text{Mpc}$ due to pair production with the CMB (section 1.4). Current experimental limits on UHE photon fluxes constrain these models.

Propagation and Energy Loss

UHE photons lose energy primarily through pair production with the cosmic microwave background (CMB):

$$\gamma + \gamma_{\mathrm{CMB}} \rightarrow e^{+} + e^{-} \tag{1.9}$$

The photon energy decreases exponentially with distance:

$$E_{\gamma}(d) = E_{\gamma,0} e^{-d/\lambda} \tag{1.10}$$

where $E_{\gamma,0}$ is the initial energy, d is the distance traveled, and λ is the attenuation length, given by $\lambda \propto (n_{\mathrm{CMB}} \cdot \sigma_{\gamma\gamma})^{-1}$, with $n_{\mathrm{CMB}} \sim 400\,\mathrm{cm}^{-3}$ and $\sigma_{\gamma\gamma} \sim 10^{-29}\,\mathrm{cm}^{2}$. At $10^{19}\,\mathrm{eV}$, the mean free path is $\sim 10\,\mathrm{Mpc}$, limiting the UHE photon horizon to $\sim 100\,\mathrm{Mpc}(z \lesssim 0.03)$ [65]. This rapid energy attenuation and the resulting limited propagation distance of UHE photons are illustrated in Figure 1.5. Given that photons travel in straight lines (ignoring gravitational

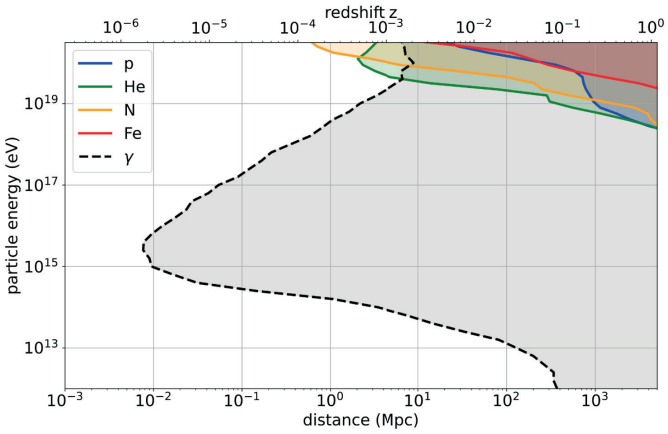


Figure 1.5: Energy-loss horizons for cosmic-ray primaries vs. distance (lower axis) and redshift (upper axis). Dashed curve: UHE photon attenuation. Solid lines: energy-loss horizons for protons, helium, nitrogen, and iron, with shaded regions indicating dominant energy-loss mechanisms (e.g., photodisintegration, photopion production). Produced by H. Schoorlemmer using CRPropa 3 [65].

lensing), their short attenuation length restricts detections to nearby sources.
The absence of detected UHE photons places stringent constraints on both astrophysical and exotic source models, enhancing our understanding of high-energy astrophysics and fundamental physics.

1.5 Searches for Ultra High Energy Photons

Unlike hadronic air showers described in section 1.2, photon-induced air showers are dominated by electromagnetic processes, where photonuclear reactions playing a secondary role at the

highest energies.

Electromagnetic Cascade Development The cascade begins when a UHE photon undergoes pair production in the atmosphere ($\gamma + \text{nucleus} \rightarrow e^+ + e^-$). These electron-positron pairs emit bremsstrahlung photons ($e^\pm \rightarrow e^\pm + \gamma$), which produce additional pairs, creating a multiplicative cycle. This EM cascade approximately doubles in particle count with each radiation length ($X_0 \approx 37\,\text{g/cm}^2$ in air) until individual particle energies fall below the critical energy ($E_c \sim 80\,\text{MeV}$), where ionization losses halt further multiplication [66].
Key features of photon-induced air showers include:

- **Longitudinal Profile:** The relatively low EM cross section at ultra-high energies delays initial development, resulting in deeper atmospheric penetration. At ultra-high energies, photon-induced air showers develop deeper in the atmosphere than hadron-induced showers. For primary energies around $E \sim 10^{19}\,\text{eV}$, photon showers can reach $X_{\text{max}} \sim 1000\,\text{g/cm}^2$, compared with typical values of $\sim 750\,\text{g/cm}^2$ for proton primaries and $\sim 650\,\text{g/cm}^2$ for iron nuclei at comparable energies [67].

- **Muon Deficiency:** Without significant pion production, photon-induced EAS generate only $\sim$ 10–30% of the muons observed in proton-induced air showers [68], a key discriminant for hybrid detectors.

- **Lateral Spread:** Photon-induced air showers are predominantly electromagnetic and have a much smaller muon content than hadron-induced showers. Since muons are highly penetrating and can contribute substantially to the signal at large distances from the shower core, their suppression in photon showers reduces the far-core ground signal. As a result, photon-induced showers exhibit a steeper, narrower lateral distribution function at ground level than hadron-induced showers of comparable energy [69, 66].

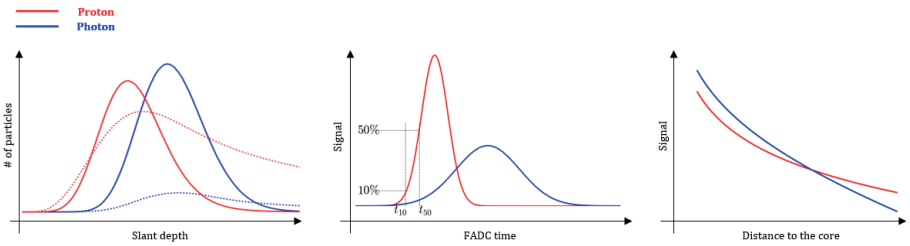


Figure 1.6: Comparison of a proton- and photon-induced air shower. Two air showers with the same zenith angle and primary energy are simulated, one initiated by a proton (red) and the other by a photon (blue). Left: Number of particles as a function of slant depth. Solid curves show $e^\pm$; dotted curves show $\mu^\pm$ (multiplied by a factor of 100 to be on a comparable scale as the $e^\pm$). Middle: Time structure of the signals in a Water-Cherenkov Detector (WCD); t_{10} and t_{50} mark the times at which the signal reaches 10% and 50% of its peak, respectively. Right: Lateral distribution of the WCD signals as a function of distance to the air-shower core.

Photonuclear Interactions at Ultra-High Energies Above $\sim 10^{19}$ eV, photonuclear interactions ($\gamma + \text{nucleus} \rightarrow \text{hadrons}$) introduce a hadronic component, producing $\pi^{\pm}$ that decay to muons and π^0 that feed back into the EM cascade. However, with a cross section ~ 1 mb orders of magnitude smaller than EM processes, these interactions contribute $< 10\%$ of the total muon content; the other $\approx 90\%$ of the air-shower energy remains in the purely electromagnetic cascade [68].

Key Observables for Photon Identification UHE experiments distinguish photon candidates using[3] well established classifiers illustrated in Figure 1.6:

- **Depth of Air-Shower Maximum:** Measured via air-fluorescence telescopes or radio antennas, photon $X_{\max}$ exceeds proton-induced air showers by ~ 200 g/cm^2 at 10^{19} eV.
- **Muon-to-Electromagnetic Signal Ratio:** Surface detector arrays exploit the lower muon content relative to the electromagnetic component in photon-induced air showers compared to hadronic primaries.
- **Lateral Distribution Function:** Photon-induced air showers exhibit steeper signal attenuation with distance from the air-shower core [67].
- **Rise Time of Ground Signals:** Photon-induced air showers are muon-poor and therefore have a reduced prompt muonic contribution to the surface-detector signal. Their ground signals are instead dominated by the more temporally dispersed electromagnetic component, leading to larger rise times than for hadron-induced air showers of comparable energy [70].
- **Energy Reconstruction Consistency:** Hybrid detectors that combine surface arrays with calorimetric observables, such as fluorescence or radio emission, provide a useful consistency check for photon-induced air showers. Surface-detector energy reconstruction is typically calibrated and optimized for hadron-induced air showers, where muons contribute significantly to the ground signal. Because photon-induced showers are muon-poor, the surface signal at a given calorimetric energy is reduced, leading to a possible tension between the calorimetric and surface-detector energy estimates.

Diffuse Flux Searches

The Pierre Auger Observatory, as the largest cosmic ray detector, is sensitive to a wide energy range, providing unprecedented exposure above 10^{17} eV utilizing different detection systems (details in chapter 2).
In Figure 1.7, the expected fluxes of UHE photons based on different assumptions are shown. These include predictions for two scenarios where the sources consist only of protons [71, 72], as well as a scenario with a mix of particles at the sources [73]. Although experimental results

[3]Not only, but widely used.

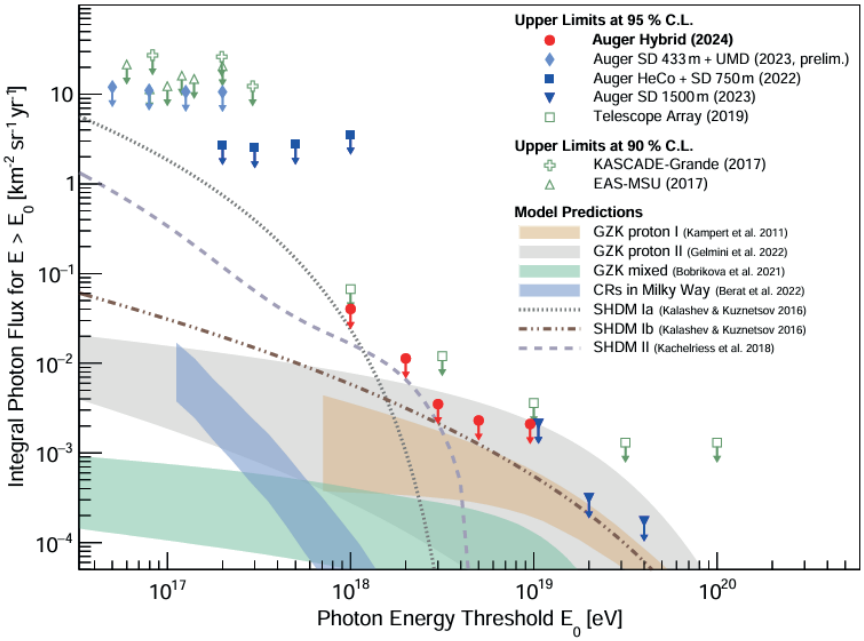


Figure 1.7: Upper limits on the photon fluxes as a function of energy [51].

are beginning to limit the optimistic predictions for photon fluxes from proton interactions at energies above 3×10^{18} eV, these predictions remain about 1 to 1.5 times higher than those from the mixed-composition model.

Previous upper limits on the photon flux have helped rule out certain models that do not involve acceleration, especially SHDM models, which aim to explain the origin of cosmic rays at very high energies (see, e.g., [74, 70]). These upper limits on photon flux have also been used to limit the mass and lifetime of SHDM particles [75, 76, 77].

Additionally, predictions for three different assumptions regarding SHDM decay are also shown. For the hadronic decay ($X \to q\bar{q}$), the expected photon fluxes are given for SHDM particles with masses within acceptable limits: M_X of 10^{18} GeV with a lifetime of 3×10^{21} years, 10^{12} GeV with a lifetime of 10^{23} years, and 10^{10} GeV with a lifetime of 3×10^{21} years [78].

Point Sources Searches

Searches for UHE photon point sources use two complementary approaches: (1) targeted cross-correlation of arrival directions with catalogs of candidate, or (2) blind searches for directional excesses without prior assumptions. Photons are uniquely valuable for such studies because—unlike charged nuclei—they propagate undeflected by magnetic fields, pointing directly back to their sources. The Pierre Auger Collaboration has applied both methods: cross-correlation tests constrain specific source classes [79], while blind searches probe anisotropies [80], limited only by the experiment's angular resolution.

A blind search for UHE photons ($10^{17.3}$ to $10^{18.5}$ eV) was conducted [1] using the Pierre Auger

Observatory hybrid data (January 2005 to September 2011). No excess was found surpass the expected background level. A targeted search on 364 pre-defined galactic and extragalactic sources (e.g., millisecond pulsars, Centaurus A) through December 2013 also found no significant excess, with the smallest penalized p-value[4] per target set indicating no detection at the 3σ level (10^{18} eV range).

Moreover, the Pierre Auger Observatory searched for UHE photons above 10 EeV coinciding with gravitational wave (GW) events from LIGO/Virgo's O1 to O3 runs [81]. Selected GW events, chosen for localization quality and proximity, were analyzed for photon signals within 1000 s and 1 day of the GW events. No coincidences were observed, yielding upper limits on UHE photon fluence of ~ 7 MeV cm^{-2} (1000 s) and ~ 35 MeV cm^{-2} (1 day). For the binary neutron star merger GW170817, less than 20% of its GW energy was converted into UHE photons above 40 EeV, marking the first constraints on UHE photon production from GW sources. Despite no significant detections, the results established new limits on UHE photon fluence and energy transfer from GW events, particularly compact object mergers like binary neutron stars and black holes.

[4]The probability of observing the detected events in the target's direction under the null hypothesis of no UHE photon emission, assuming events are due to background (e.g., hadrons), and accounting for expected background and statistical weights.

Chapter 2

Pierre Auger Observatory and AugerPrime

The Pierre Auger Observatory is currently the largest cosmic ray observatory studying UHECRs, located in Argentina at 1400 meters above sea level, covering an area of about 3000 km^2, and designed as a hybrid detector, combining surface and fluorescence detection techniques. This chapter provides an overview of the fundamental components of the observatory, followed by details of the AugerPrime upgrade. Key components of the upgrade include new detectors and enhanced systems. The chapter will focus on the Surface Detector[1] and Radio Detector, while providing brief descriptions of other components involved in the observatory.

2.1 The Observatory

The Pierre Auger Observatory's design, shown in Figure 2.1, before the AugerPrime upgrade, consists of two main detector components: the Fluorescence Detectors (FD) and the Surface Detector (SD). Together, they form a hybrid system capable of measuring both the longitudinal and lateral profiles of cosmic-ray air showers.

Fluorescence Detector

The FD system comprises 24 telescopes distributed across four sites (Los Leones, Los Morados, Loma Amarilla, and Coihueco) positioned along the perimeter of the observatory. Each telescope consists of a spherical mirror with a 3.6 m diameter and a camera of 440 photomultiplier tubes arranged in a 22×20 matrix. The FD operates during clear, moonless nights, achieving approximately a 15% duty cycle. Each telescope covers a view of 30° horizontally and 28.6° vertically. Three extra telescopes were later added to the Coihueco site to point at higher elevations and detect lower-energy air showers.

[1] In the original design report, the Surface Detectors (SD) referred exclusively to the array of Water Cherenkov Detectors (WCDs). However, following the upgrade, each water tank hosts multiple detection systems—including radio antennas and scintillators—making the term "Surface Detectors" less representative of just the WCD array. For consistency with existing literature, I will continue to use "SD" to refer to the Water Cherenkov Detector array as a whole, and "WCD" to refer to an individual tank unit

By observing the sky up to an elevation of 58°, they can capture air showers at energies as low as 10^{17} eV [82].

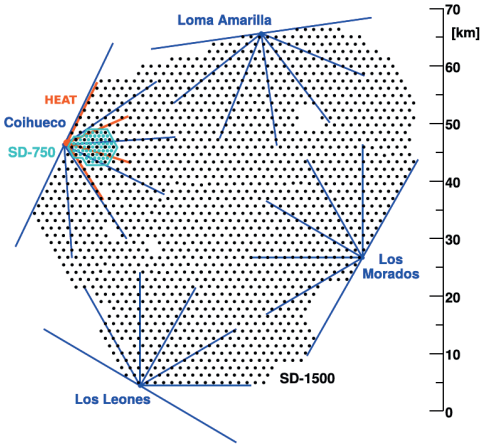


Figure 2.1: The layout of the Pierre Auger Observatory of SD and FD. The blue and orange regions indicate the fields of view for the four FD sites (adapted from [83]).

The FD measures the nitrogen fluorescence light emitted isotropically by air molecules excited by charged particles in the extensive air shower. This fluorescence emission, primarily in the 300–400 nm wavelength range, allows the longitudinal development of the air shower to be observed. The optical signal recorded by the telescopes also contains direct and scattered Cherenkov light. In the reconstruction, the fluorescence and Cherenkov contributions are modelled together, accounting for atmospheric transmission and the detector response, to obtain the longitudinal energy-deposit profile [84]. A Gaisser–Hillas function is then fitted to this profile. The position of its maximum defines $X_{\max}$, while its integral gives the calorimetric energy E_{cal}. The primary energy is obtained as $E_0 = E_{\text{cal}} + E_{\text{inv}}$, where the invisible energy E_{inv} accounts mainly for energy carried by muons and neutrinos that is not deposited in the atmosphere. This correction is not negligible. For vertical events, the Auger data-driven parametrisation corresponds to invisible-energy fractions of approximately 14%, 13%, and 11% at calorimetric energies of 1, 10, and 100 EeV, respectively [85]. The air-shower direction is determined from the timing of triggered camera pixels together with their pointing to define the air-shower–detector plane and the axis within it; in hybrid events, incorporating surface detector timing further constrains the geometry, yielding sub-degree accuracy in arrival direction. The FD's key performance characteristics included an energy resolution of $\sim 10\%$, angular resolution better than 1°, and the ability to determine the atmospheric depth of air-shower maximum with a resolution of about $20\,\text{g/cm}^2$. This precision in measuring $X_{\max}$ is crucial for studies of cosmic ray mass composition. The FD also provides a nearly calorimetric measurement of the primary cosmic ray energy, serving as the absolute energy calibration for the entire observatory. The FD is described in detail in [86].

Surface Detector

The SD has been deployed in a triangular grid across the observatory area. Each WCD unit comprises a cylindrical tank with a 3.6 m diameter and 1.2 m height, containing 12,000 liters of ultra-pure water in a sealed liner. The water volume is observed by three 9-inch photomultiplier tubes (PMT) mounted on the top surface of the liner, looking downward into the water.
When charged particles from air showers traverse the water at speeds greater than the local phase velocity of light, they emit Cherenkov radiation that is detected by the PMTs. Although the SD stations cannot directly identify individual particles or track their paths through the tank, the recorded light varies depending on the particle type and the distance it travels within the water. Electrons, positrons, and photons are largely absorbed before exiting the tank.
Each WCD operates as an autonomous detection station, powered by solar panels and batteries, with local electronics for signal processing, triggering, GPS timing, and communication. The detector design enables sensitivity to both the electromagnetic and muonic components of air showers, with nearly 100% duty cycle operation. The SD is described in detail in [87].
The SD includes multiple subarrays with different station spacings to extend the lower energy threshold of the observatory:

- **SD1500 (regular grid):** The main array consists of 1600 WCD arranged with 1500 m spacing, covering the entire 3000 km^2 area. This configuration provides full detection efficiency for hadronic-induced air showers with energies above 2.5 EeV and zenith angles less than 60°, with less sensitivity for inclined air showers above 65° where full efficiency is at 4 EeV.

- **SD750 (Infill):** A denser infill array consists of 61 WCDs covering $\sim 23.5\,km^2$ with 750 m station spacing. Located near the Coihueco FD site, this array lowers the full efficiency threshold to approximately 3×10^{17} eV, enabling studies in the energy region where the transition from galactic to extragalactic cosmic rays.

- **SD433:** A smaller, even denser array contains 19 WCDs with 433 m spacing was deployed to extend measurements to still lower energies, around 10^{17} eV. This array completed the energy coverage of the observatory, creating overlap with measurements from other experiments and provides crucial data in the ankle region of the cosmic ray spectrum.

The reconstruction of cosmic-ray air showers involves analyzing the spatial and temporal distribution of signals detected in triggered stations [88, 89].
The air-shower geometry is determined by fitting the arrival times of signals from the air-shower front's shape and particle density at different detector positions, establishing the air-shower axis direction and core position. The angular resolution for the Pierre Auger Observatory is reported to be better than 1.61° for events with more than three stations, and better than 0.91° for events with more than six stations. The lateral distribution of signals is fitted using an empirical function that models the signal strength as a function of distance from the air-shower axis. The signal at a reference distance of 1000 meters, typically for the SD1500 array, is referred to as

S1000 and is used as an energy estimator. S1000 is calibrated against FD energy measurements from hybrid events, establishing an energy scale with a systematic uncertainty of approximately 14% [90]. However, S1000 decreases with increasing zenith angle due to the attenuation of air-shower particles and geometric effects, as the path length through the atmosphere increases.

This reconstruction methodology enables the determination of primary cosmic ray properties with high precision across the operating energy range of the observatory. WCDs achieved energy resolution of approximately 12-16% (depending on energy and zenith angle), a core position resolution better than 50 m, and reliable event reconstruction for zenith angles up to 80°.

2.2 The AugerPrime Upgrade

The key scientific goals driving this upgrade are centered around understanding the suppression of the cosmic-ray flux above approximately 40 EeV and determining the corresponding mass composition.
The FD provides a direct measurement of $X_{\max}$, but its limited duty cycle restricts the available statistics at the highest energies. More recently, a deep neural network applied to the time traces recorded by the Phase-I WCD array, and cross-calibrated using hybrid events, has extended measurements of the $X_{\max}$ distributions and mass composition from 3 to 100 EeV, including the region above 40 EeV [51]. AugerPrime complements this measurement by providing additional event-by-event sensitivity to the electromagnetic and muonic components of air showers. The upgrade includes scintillator detectors, radio antennas, a small PMT in each WCD, upgraded electronics, and underground muon detectors [91].

The SD electronics has been upgraded to a higher-performance version. The pre-upgrade version features 6 channels, 40 MHz sampling, and 10-bit ADC resolution, used for high and low gain signals from the 3 PMTs in each WCD. In the new system, it is necessary to accommodate the new detectors, such as the Scintillator Surface Detectors (SSD), small PMT, Radio Detectors (RD), and Underground Muon Detectors (UMD). While the UMD is at present at only a few stations, the other detectors are located at nearly every station. The new system, called the Upgraded Unified Board (UUB), offers 10 channels, 120 MHz sampling, and 12-bit ADCs. The additional small PMT (sPMT), which has a smaller photocathode area and therefore collects less light than the three standard PMTs, extends the dynamic range of each WCD. The standard PMTs saturate for station signals above approximately 1000 VEM[2]. This is a station-signal threshold rather than a primary-energy threshold, and saturation can occur already in showers with energies of a few EeV when a station lies close to the shower core. The probability that an event contains at least one saturated station increases strongly with energy, reaching more than 40% above 30 EeV and approximately 80% near 100 EeV. The SPMT extends the

[2]A VEM denotes the average signal produced by a vertical-equivalent muon crossing the WCD.

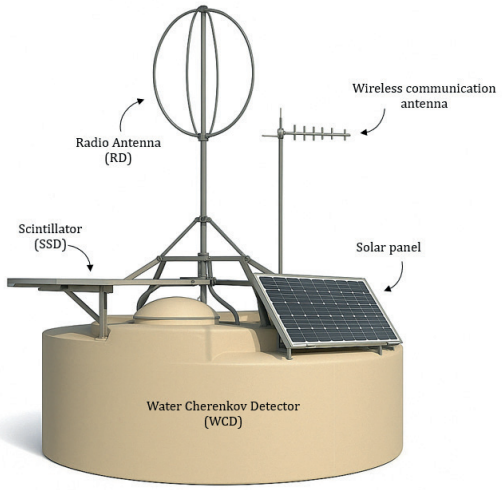


Figure 2.2: Schematic view of a surface detector station, with the radio antenna (RD) and the Scintillator Surface Detector (SSD) mounted above the Water Cherenkov Detector (WCD). Illustration by the author, adapted from on-site photography of an AugerPrime station.

measurable range to approximately 25 000 VEM, substantially reducing the occurrence of saturated signals [91, 92].

Each SSD unit is mounted on top of the WCD and consists of two plastic scintillator panels that cover a total area of $\sim 4\,\mathrm{m}^2$. The SSD enhances the WCD by offering a different response to electromagnetic particles compared to the WCD [93]. Its thin design makes it particularly sensitive to low-energy electrons and positrons. In contrast, WCDs are sensitive to both electromagnetic and muonic components. This complementary behavior allows the combination of SSD and WCD data to distinguish between the muonic and electromagnetic parts of an air shower. The SSD is crucial for distinguishing between different cosmic ray compositions and improving low zenith event ($\leq 45°$) reconstruction. At higher zenith angles, SSD measurements become less effective because the electromagnetic shower component is increasingly attenuated by the larger atmospheric slant depth and because the projected area of the horizontal scintillator decreases with increasing zenith angle. For this reason, the SSD–WCD combination is primarily used for vertical and moderately inclined showers, while the RD extends composition-sensitive measurements towards more inclined events.

Another subarray added is the UMD, located in the SD-750 and SD-433 arrays. The UMD is designed to directly measure the muon density in extensive air showers using scintillator detectors buried 2.3 meters underground. This depth effectively shields the detectors from most electromagnetic particles. The UMD will improve event-by-event mass composition reconstruction. The SD-750 and SD-433 arrays, together with the UMD, are part of the low-energy sensitivity enhancement.

To extend mass composition studies to inclined air showers, each SD1500 tank is also equipped with a radio antenna to form the RD array. In such inclined events, the electromagnetic parti-

cle component of the air shower is absorbed in the atmosphere before reaching the ground, so the WCD primarily detects muons. The radiofrequency signal originating from the air shower propagates mostly unperturbed through the atmosphere. The ratio of muonic to electromagnetic signals remains a key observable for inferring the primary particle's mass.
The SSD, UMD, and RD are operating in slave mode, meaning they only record data when the corresponding WCD is triggered.

2.3 The Radio Detector

Radio detection activities at the Pierre Auger Observatory began with the RAuger prototype experiment in 2006, continued with the Multi-Antenna eXperiment In Malargue Argentina (MAXIMA) installed in 2008 as an evolution of the initial setup to enable hybrid measurements with autonomous stations [94], followed by the deployment of the Auger Engineering Radio Array (AERA) in 2010, which expanded to 124 stations by 2013 and continued to develop to get to its final form in 2015 [95].
AERA was designed to explore and develop various aspects of radio detection techniques, such as testing different antenna designs, station electronics, and triggering methods. While primarily an engineering initiative supporting the Radio Detector component of the AugerPrime upgrade, it has produced numerous significant physics results, including precise $X_{\max}$ measurements, validation of radio emission models, and observations of inclined air showers, contributing substantially to cosmic-ray research [39, 96, 97]. AERA is located in the SD-750 region near Coihueco at the Pierre Auger Observatory, covering $\sim 17\,\mathrm{km}^2$ with 154 antenna stations spaced irregularly. Phase I (2010) included 24 LPDA-antenna stations with 144 m spacing. In 2013, Phase II added 100 Butterfly antenna stations with 250 m and 375 m spacing. Phase III introduced 25 more Butterfly stations, mainly spaced at 750 m near SD-750 stations.

AERA has proved the science case for adding radio detection to the observatory. Despite the fact that the RD was not part of the original design of AugerPrime [91], extra funding was granted to develop a large-scale radio detector. The main component of the RD is a Short Aperiodic Loaded Loop Antenna (SALLA), illustrated in Figure 2.2. It consists of two aluminum loops mounted perpendicularly to form a dual-polarized system [98]. The entire structure is supported by a fiberglass mast. One loop is aligned in the geomagnetic north–south direction and the other in the east–west direction. The mast is stabilized by an aluminum frame connected to the WCD structure without touching the SSD. The SALLA is mechanically simple, cost-effective, and easy to maintain. With a diameter of 122 cm, it is optimized for the 30–80 MHz frequency range, providing a nearly uniform response with low dispersion.
Signal transmission is handled via two coaxial cables inside the mast, connected to the Low-Noise amplifiers (LNA) at the top of the antenna to the front-end digitizer electronics inside the WCD dome. Each of the two antenna channels is digitized using an analog-to-digital converter (ADC) with a sampling rate of 250 MHz and 12-bit resolution. Signals are amplified by 36 dB

and filtered with a bandpass between 30 and 80 MHz.
A Field Programmable Gate Array (FPGA) manages data exchange with the UUB, collecting 2048 samples when triggered by the WCD. Future plans include incorporating RD data into the trigger decision process, thereby going beyond slave mode operation. Finally, all data are transmitted via the wireless communication system used by the Surface Detector array. Additional monitoring data, such as measurements of the Galactic radio background, will also be sent as part of this stream. More technical details about the RD can be found in [98] and will be discussed in chapter 3.
The radio upgrade is meant to help achieve AugerPrime's goal of deriving the particle type of the incoming cosmic ray up to the highest energies. However, there is an important distinction in the methods used to determine the mass of cosmic rays between AERA, which is optimized for vertical air showers, and the RD, which targets horizontal air showers. For vertical air showers, a geometrical method is employed. This approach correlates the size of the air-shower footprint on the ground with the distance from the observer to $X_{\max}$. In contrast, for horizontal air showers, a different method is proposed [99]. It involves the combined use of radio antennas and WCD to measure the electron-to-muon ratio (e/μ) in air showers. This ratio serves as an indicator for determining the atomic mass as $\ln A$, thus providing a way to infer the mass composition of the primary cosmic ray.

2.4 The Surface Detector Triggers

The trigger system is a critical component of cosmic ray detection experiments to maximize signal over background. The SD trigger system is essential for understanding when signals are recorded by the WCD, which subsequently triggers the RD.

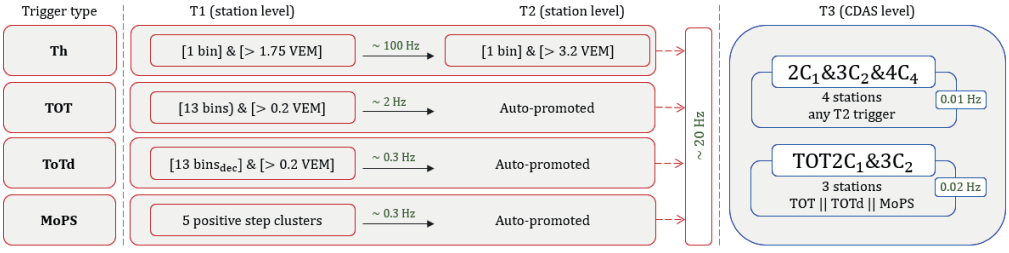


Figure 2.3: Schematic overview of the SD trigger hierarchy at the Pierre Auger Observatory, showing the progression from station-level to array-level triggers. T1 rate is dominated by the threshold trigger. The combined T2 rate is ~ 20 Hz. T3 rates are 0.01 Hz for $2C_1\&3C_2\&4C_4$ and 0.02 Hz for TOT $2C_1\&3C_2$. Adapted from [100].

The SD uses a hierarchical trigger system [100], as shown in Figure 2.3. This system includes station-level triggers (T1 and T2) and array-level triggers (T3, T4, and T5), coordinated by the Central Data Acquisition System (CDAS). When a detector generates a T2 trigger, it sends timestamps and trigger details to the CDAS, which buffers this information and waits for a

T3 trigger—formed by combining T2 signals from multiple nearby detectors to confirm the spatial–temporal pattern of a genuine extensive air shower (multiple tanks hit nearly simultaneously as the air-shower front sweeps across the array). Once a T3 trigger is formed, the CDAS collects all relevant data within a 30-μs window for event reconstruction. This hierarchical design addresses the bandwidth limitations of the array's wireless network.
The T3 trigger significantly reduces the rate of atmospheric muon background events while maximizing the detection of EAS events.
The T1 trigger consists of two primary modes, where meeting either of the modes is sufficient for an event to be checked for the possibility of higher-level triggers.

- **T1 Trigger:** The T1 trigger, implemented at individual SD stations are:
  - **Threshold Trigger (TH):** This mode requires each of the three PMTs in a station to exceed a signal threshold of 1.75 VEM. If only one or two PMTs are active, the threshold increases to 2.8 VEM or 2.0 VEM, respectively.
  - **Time-Over-Threshold (ToT):** The ToT mode requires at least 13 of the 25 ns time bins within a 3 μs window to exceed 0.2 VEM in at least two PMTs. It is sensitive to temporally extended signals, such as those from low-energy vertical air showers near the core or high-energy inclined air showers far from the core, and effectively filters out random muonic background.

  To further enhance performance, two additional T1 trigger modes have been introduced to further reduce the influence of atmospheric background muons and decrease the energy threshold of the array. The two modes, which build upon the ToT condition:

  - **Time-Over-Threshold Deconvolved (ToTd):** This mode deconvolves the FADC trace to suppress the exponential tail of diffusely reflected Cherenkov light before applying the ToT condition. It improves discrimination between EAS signals and muonic background by compressing the muon signal, which usually has a fast rise time in one or more bins before meeting the ToT condition. A minimum integrated signal of 0.5 VEM is required for the ToTd trigger.
  - **Multiplicity-of-Positive-Steps (MoPS):** The MoPS mode counts positive steps (successive signal increases) in the FADC trace within a 3 μs sliding window, requiring more than four steps above 5 $\times$ RMS noise (up to $\sim$ 0.5 VEM) in at least two PMTs. It enhances the recovery of low-energy signals, achieving rates below 2 Hz.

- **T2 Trigger:** The T2 trigger imposes stricter criteria on T1 events. T1-ToT, ToTd, and MoPS events are automatically promoted to T2-ToT, while a T1-TH event must exceed 3.2 VEM across three coincident PMTs (or 3.8 VEM for two PMTs, 4.5 VEM for one) to become T2-TH. This reduces the station event rate to approximately 23 Hz.

- **T3 Trigger:** Managed by CDAS, the T3 trigger integrates spatial and temporal data from T2-triggered stations. It operates in two modes, requiring only one to be satisfied [100]:

  - **ToT-ToT2C1 & 3C2:** Requires three stations with T2-ToT triggers, including at least one in the first crown[3] ($2C_1$) and one in the second crown ($3C_2$) around a central station, with timing within $(3 + 5C_n)\,\mu$s.
  - **2C1 & 3C2 & 4C4:** Requires a four-fold coincidence of any T2 triggers, with stations in the first, second, and up to the fourth crown ($4C_4$), favoring inclined air showers with sparse patterns.

  Upon T3 activation, all relevant station data within a $30\,\mu$s window are read out and stored.

- **T4 and T5 Triggers:** These offline triggers refine data quality:
  - **T4 Trigger:** Selects events based on geometric and timing compatibility, using either a 3-ToT compact configuration (three nearby T2-ToT stations) or the 4C1 configuration (three of six first-crown neighbors triggered).
  - **T5 Trigger:** Ensures air showers are well-contained within the array, with two levels: 6T5 (all six first-crown stations functional) and 5T5 (five functional). The 6T5 condition, critical for neutrino searches below 75°, enhances reconstruction accuracy.

Slave-Mode Triggers

The RD and UMD teams are independently developing autonomous trigger systems that may later be integrated into the station-level trigger. At present, the SSD, UMD, and RD operate in slave mode with respect to the WCD, which provides the station-level trigger. In this configuration, the RD is read out for every WCD T1 trigger, allowing radio traces to be recorded already at the lowest local trigger level. However, the corresponding event data are only transferred to CDAS and stored when CDAS requests a station read-out, typically after a higher-level array trigger has been formed from correlated T2 triggers in neighbouring stations.

[3] A concentric rings of detectors arranged around a central detector station.

Chapter 3

The Radio Detector Monitoring and Commissioning

The large-scale radio upgrade of the Pierre Auger Observatory motivated the development of a dedicated methodology to monitor and commission the deployment of more than 1,600 radio antennas. Given the issues that may arise from the hardware and software sides, it was important to have a bird's view of the array, pointing out the stations with issues, ensuring active commissioning as well as validating the radio array data-taking for further physics use. Understanding how the system was deployed and how it has evolved over time is crucial to developing a comprehensive understanding of its proper functioning.

3.1 The Deployment History of The Radio Detector

The deployment of the radio antennas began before the installation of the corresponding digitizers. Although an antenna could be mechanically installed at a station, its data became available only after the digitizer was installed and integrated into the standard Auger data stream. Consequently, a station could be monitored as part of the radio array only once both the antenna and digitizer were present. This distinction is important for interpreting the deployment history, since the number of installed antennas does not necessarily correspond to the number of stations providing usable radio data.

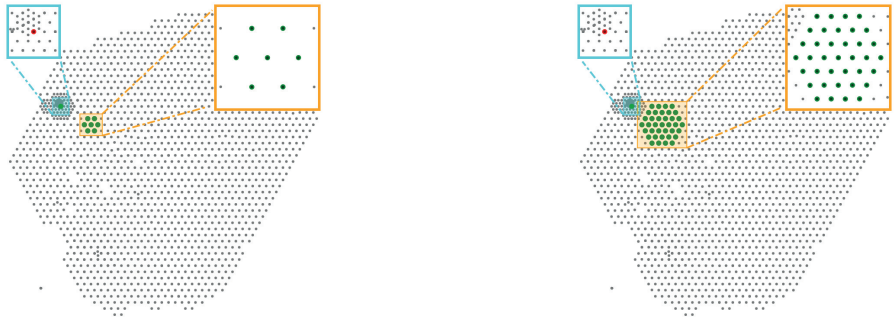
- In May 2019, the first RD station, Clais (LS 22), came online in the infill region. By November of the same year, six additional stations were equipped with digitizers: two in the infill region and four in the adjacent regular SD-1500 array, including one doublet station, laying the groundwork for future large-scale deployment.

- By September 2021, four additional stations were deployed in the regular grid, completing the hexagonal configuration of the upgraded stations. The engineering array (EA) now consisted of eight stations, as shown in Figure 3.1a. Seven of the stations formed a hexagonal pattern (with two stations forming a doublet) in the regular grid near the infill border, while one station was placed within the infill region. These locations were

deliberately chosen because of their high background noise levels, providing a worst-case environment for testing the array's performance limitations.

- Progress stalled due to delays in digitizer production due to the short supply of electronics components caused by the COVID-19 pandemic. However, in May 2023, a new batch of digitizers was received, allowing the expansion of the engineering array into what became known as the Super Hexagon, consisting of 40 stations Figure 3.1b. With an area $\sim 52\,\mathrm{km}^2$, the Super Hexagon became the largest ground coverage radio array in the world at that time.

- By October 2023, mass deployment was started and completed in November 2024. The final configuration included 1,530 stations in the regular grid, 49 in the infill, and 12 in the super-dense region (SD-433) within the infill. Figure 3.2 illustrates the current status of the deployment. In particular, a region in the southeast remains without radio upgrade due to unresolved land access restrictions imposed by local property owners.



(a) EA (b) Super Hexagon

Figure 3.1: Station configurations for (a) the Engineering Array (EA) in June 2023 and (b) the Super Hexagon in August 2024.

3.2 The Radio Detector Monitoring (MoRD)

During the deployment phase of the Super Hexagon, the need for an active monitoring system accessible to all users became evident. To address this, a dedicated website, **Mo**nitoring **R**adio **D**etector (MoRD), was developed, significantly improving the capability of AugerPrime upgrade commissioning. Its primary objective was to facilitate real-time access to data from newly deployed antennas, an aspect previously unavailable. The user-friendly interface of the website enables researchers and collaborators to track deployment progress in semi real-time (typically 12 hours delay), improving transparency, cross-validation, and routine maintenance. The approach of MoRD[1] is straightforward: collect the necessary data, process them efficiently,

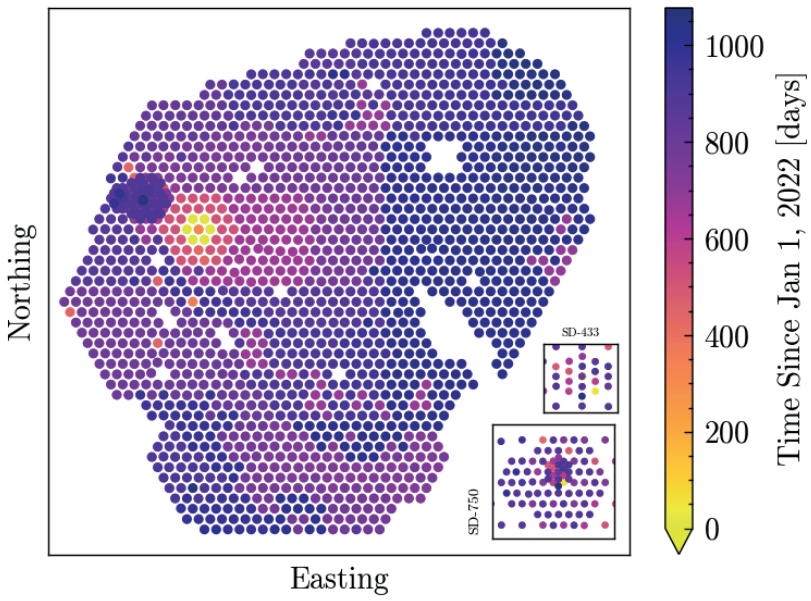


Figure 3.2: Spatial distribution of operated radio stations as of February 25, 2025. The color scale represents the number of days since January 1, 2022. The inset provides a detailed view of the super-dense (SD-433) and Infill (SD-750) regions, highlighting localized deployments.

and update the website. Automated scripts handle this process to keep updates regular while requiring only minimal maintenance.

The website workflow, shown in Figure 3.3, is structured as follows:

- Data Synchronization: The data flow script synchronizes raw data from remote servers (Lyon) to local servers for analysis and storage. The data copied from the Auger observatory in Argentina to Lyon everyday around 05:00 CET which typically include the two files of the previews 24-hours, for example, on February 20, the second of 18^{th} and the first 19^{th} data files are uploaded. Hence, the Lyon has a data latency of the order of one and a half days. However, during thunderstorms or technical problems, the files are not uploaded in time.

- Data Processing and Visualization: Once synchronized, the data are processed on the local cluster to generate figures and tables that provide insight into performance and deployment status. These visuals, along with the processed data, are published as static HTML pages on the MoRD website, allowing the AugerPrime team to access the information with minimal loading times.

[1]The website is hosted by the High Energy Physics (HEP) department at Radboud University and is accessible to users with Auger credentials.
Website: https://hef.ru.nl/mord/index.html
Repository: https://gitlab.iap.kit.edu/auger-observatory/monitoring/mord

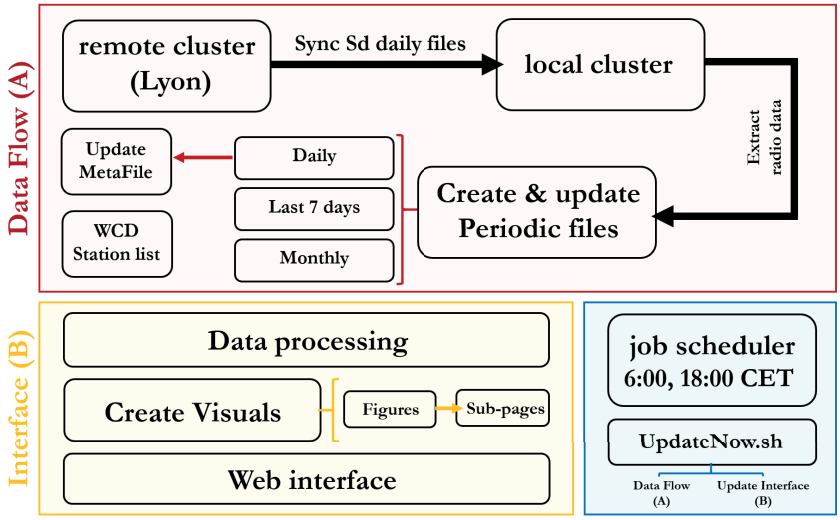


Figure 3.3: Workflow diagram for MoRD. The process begins with data synchronization from Lyon to a local cluster at Radboud University in Nijmegen then produce the periodic files (A). Then generate static web pages containing figures and tables for user interaction (B). The entire system is designed for full automation, ensuring that updates are carried out efficiently every 12 hours to reflect the most recent data.

- Automation: MoRD operates as a fully automated system, updating every 12 hours to keep the displayed data current. This reduces human error and minimizes the need for manual intervention, ensuring a smooth and reliable workflow for continuous monitoring. However, disruption can occur due to local cluster updates or corrupted data files.

3.3 Data Flow and Processing

The first step is retrieving the data from the Lyon Cluster, which contains the raw observational data that form the foundation for all subsequent analysis and dissemination. This acquisition process is automated and secured to minimize latency and ensure data integrity. The Pierre Auger Observatory stores data in two formats: **Sd** and **Ad**[2] files. The Sd data file contains all T3-triggered events processed by CDAS. Sd files contain 12 hours of data and are created every day with a one-day delay[3]. The ad files contain 24 hours of data and are created every day with a 7-day delay. For lower latency monitoring, Sd files are preferred.
The primary data stream is obtained by fetching only the radio data out of the Sd files using python library Uproot [101] to parse data. A typical Sd file is around 600 MB, while the extracted file containing only radio data is around 100 MB. Two data files are generated daily,

[2]Sd files include data from surface detectors, such as WCD and RD, and contain all triggered event instances. Ad files include FD data and merge events with matching timestamps, preventing event fragmentation, particularly for large footprint events.
[3]CET timezone

approximately 12 hours apart, triggering an update cycle on the website whenever a new file becomes available. To optimize data extraction, the processing code only selects events that include at least one antenna with valid data. In addition, to minimize data size, only tanks with an antenna are included in the dataset.
The generated file contains numerous parameters, but only a selected subset is essential for effective RD monitoring. The selected extracted data values are:

- Event ID
- Station ID (used to retrieve station name, position, and altitude)
- Time details of the event in GPS time format with subsecond information.
- Signal traces, which include the east-west (EW) and north-south (NS) channels labeled Ch0 and Ch1. Each signal is measured in ADC counts and has a length of 2048 samples in time.
- Trigger type.
- Digitizer firmware version: The firmware underwent several updates during the deployment with critical changes like where to place the radio signal (signal window).
- Signal parity[4]: Each sample in the signal trace has a parity bit assigned. The parity bit helps verify whether the sample data is healthy (valid) or corrupt (invalid). If a sample fails the parity check, it indicates that one or more bits in that sample have been flipped due to errors (e.g., transmission noise, memory faults).

During the early stages of deployment, the staff occasionally tested digitizers on different antennas before final installation. As a result, the data stream may include more stations than expected, some of which lack recorded signals due to these temporary setups.
For each new raw file that becomes available, named in the format sd_YYYY_MM_DD_HHhmm, the data flow script automatically performs the following tasks:

1. Daily Files: The script extracts radio data and saves it as sd_YYYY_MM_DD_HHhmm_rd, representing individual files.
2. Weekly File: A consolidated weekly file is generated at the beginning of each day, incorporating data from the previous seven days.
3. Monthly Files: Daily files are aggregated into a monthly file as soon as they are created, compiling all data for that month.

The different file formats are designed to accommodate the varying needs of website users, enabling them to access data based on their specific interests and requirements.
One of the initial challenges was ensuring that the website remained fast and responsive while

[4]More about signal parity: Parity bit

retaining all historical data, not just the most recent updates. To address this, a local comma-separated values file, named **MetaData**, was created to store data from each file as a single row. This file is updated with each new entry, eliminating the need to iterate over all previous files, thus improving efficiency.

Since this monitoring effort focuses on commissioning the RD deployment, only minimal data processing is required, with no cuts applied. Additional quantities are computed to provide a more detailed analysis.
For each station, each polarization is processed independently. The Root Mean Square (RMS) and the Discrete Fourier Transform (DFT) computed using the fast Fourier transform algorithm are evaluated using the following formulas:

- Root Mean Square:

$$\mathrm{RMS} = \sqrt{\frac{1}{N}\sum_{k=1}^{N} |a_k - \bar{a}|^2}, \quad \text{where} \quad \bar{a} = \frac{1}{N}\sum_{k=1}^{N} a_k$$

  a_k represents the k-th value among the total entries used for the baseline evaluation, where $N = 2048$. In addition, $\bar{a}$ denotes the mean of the distribution, representing the baseline value in ADC counts.

- DFT of the time domain:

$$A_j = \left| \sum_{k=0}^{N-1} T(k) \cdot e^{-\frac{2\pi \mathrm{i}}{N} jk} \right|$$

  A_j represents the amplitude of the frequency component at frequency index j, while $T(k)$ denotes the k-th sample of the time-domain signal, and i is the imaginary unit.

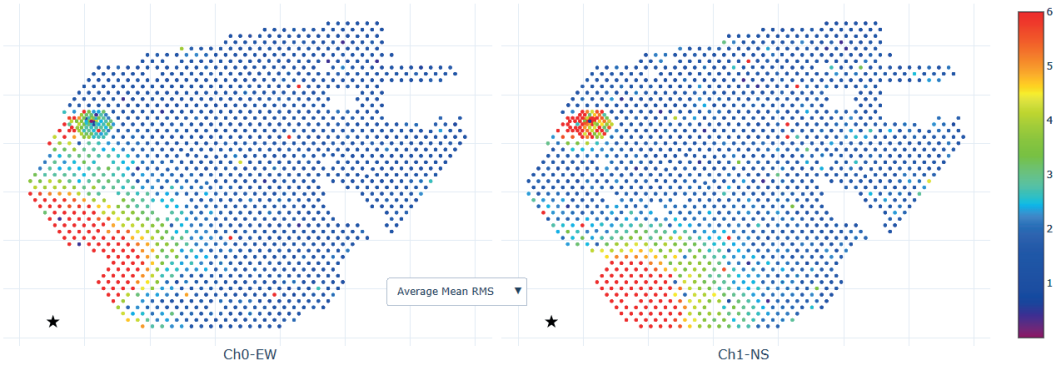


Figure 3.4: The color scale represents the average mean RMS for both polarizations over a 12-hour period. The high-noise region on the left in the EW polarization is caused by a nearby TV transmitter, indicated by a star symbol. Screenshot taken on May 21, 2024.

Since the samples fluctuate around their baseline, the RMS is typically a few ADC counts, with variations depending on the electronics used. It serves as a useful metric for assessing

the detector's operational stability. Controlled RMS values within a specific range suggest that there is no significant noise or major hardware issues present. However, RMS alone is insufficient to monitor the overall condition of the detectors. For example, a malfunctioning hardware may still produce a stable RMS value, similar to electronic noise, making it difficult to detect issues based only on RMS. Therefore, while RMS is valuable for identifying certain problems, a comprehensive diagnostic approach involving multiple variables is necessary, as implemented in MoRD is necessary.

The EW polarization in Figure 3.4 reveals a localized high-noise region, represented by a red cluster, which is attributed to a nearby TV transmitter. Similarly, the NS polarization also shows an area of elevated noise in a nearby region. Additionally, a secondary high-RMS region appears in the infill area on the top left of the array. For example, during the week of March 10–17 (Figure 3.11), the RMS mode for EW is approximately 1.5, while for NS, it is around 1.6. Similarly, the 95^{th}-quantile for EW is 8.02, whereas for NS, it is 13.4, indicating higher noise levels in the NS direction.

The strong TV transmitter serves as a reference point for diagnosing and characterizing antenna gain and orientation. The infill region experiences higher noise levels due to increased human activity and the presence of multiple subarrays, which contribute to radio frequency interference (RFI). Meanwhile, stable functioning stations are primarily affected by electronic noise, galactic signals, and persistent narrow-band radio interference, which forms a stable background for the detector.

3.4 Data Visualization

An effective visualization is essential for interpreting complex data, identifying trends, and making informed decisions, aligning with one of MoRD main objectives: ensuring accessibility for the entire collaboration. To facilitate efficient data access, the processed data are presented through static HTML pages. This decision was initially motivated by the project's origins as a small personal undertaking with limited access to dedicated web hosting, making simplicity, reliability, fast loading, and minimal server-side computation. When the website was introduced to the collaboration on 24 May 2023, it initially featured only three pages. However, as the deployment progressed, the demand for additional pages with more detailed information became evident.

Data visualization is a key component of the monitoring system, enabling effective tracking of deployment progress. By transforming complex datasets into intuitive visual formats, it facilitates analysis and informed decision-making. Visual representations allow users to quickly identify patterns, trends, and anomalies that may not be immediately apparent in the raw data. For example, time series plots reveal variations in signal strength over time, while heat maps highlight regions with consistently strong or weak signals, as shown in Figure 3.4. These visualizations provide immediate feedback on the performance and status of radio stations, which makes them essential for early issue detection and prompt intervention. Interactive dashboards

enable users to dynamically explore data. This interactivity enhances system monitoring, allowing for a more comprehensive analysis and enabling a quick response to emerging issues as they arise.
Visual tools are essential to identify specific issues within the deployment. For example, if an antenna's signal strength drops unexpectedly, visual trend analysis can help correlate this change with potential causes, such as environmental factors or equipment failures. Comparison of data from multiple antennas highlights discrepancies and inconsistencies, helping diagnoses and troubleshooting. This comparative analysis plays a crucial role in maintaining the overall integrity and performance of the system.
Developed entirely in Python, MoRD utilizes several interactive plotting libraries, although only two offer a high level of interactivity: Plotly [102] and Bokeh [103]. Plotly was selected for this project for two main reasons:

1. Framework Compatibility: Plotly is an agnostic framework that offers full-featured bindings for Python, JavaScript, and HTML, making it adaptable for various implementations.

2. Enhanced Interactivity: It provides a superior interactive experience, particularly for complex dashboards, ensuring seamless user interaction and data exploration.

The first version of MoRD was launched in May 2023. At that time, only eight stations were operational: the Engineering Array (EA) and one station in the infill array (Peteroa Jr.). Figure 3.5 shows the initial version of MoRD. The primary focus during this phase was to detect hardware failures, incorrect installations, and significant environmental problems. As a result, the initial implementation included readout rates and station spectra to aid in diagnostics. As new insights emerged, additional pages and parameters were incorporated to enhance the system's capabilities.
All MoRD plots are interactive and user-friendly, leveraging Plotly's extensive range of interactive features, including:

- Zooming and Panning: Users can zoom in and out to explore the data in greater detail, which is particularly useful for time- and frequency-domain signals where specific intervals need closer examination. The panning allows users to navigate seamlessly through large visuals by clicking and dragging.

- Hovering: Hovering over data points reveals detailed information, such as the station ID, without cluttering the plot. This feature is especially useful when comparing subplots.

- Clickable Legend: Users can click on data points to view additional details. Clicking on legend items filters the displayed data, allowing users to focus on specific subsets. Legends can also be toggled to show or hide traces, making it easier to compare different data series.

- Range Sliders: For time-series and continuous data, range sliders allow users to select and analyze specific time intervals, which aids trend analysis.

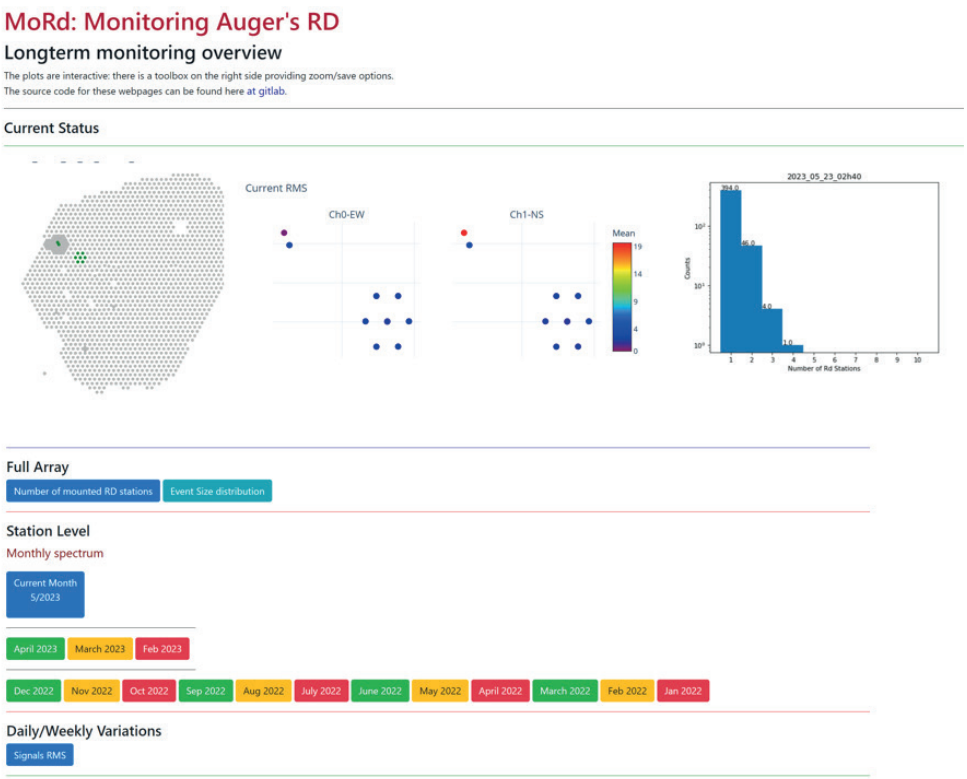


Figure 3.5: MoRD main page at its initial launch in May 2023. The contents are explained in the text.

- Data Selection and Filtering: Buttons and dropdown menus enable users to quickly adjust time frames, data ranges, and displayed information.
- Export and Save: Users can export plots as static images (PNG, JPEG) directly from the interactive interface for further analysis, documentation, or the pleasure of showing them to others.

Index page

As the array and user interest continue to grow, new pages are regularly added to the website. The index page begins with brief instructions on how to navigate the interface, along with links to the User Guide and the GitLab repository for further reference.
The website index is structured into a header page and five sections, each dedicated to a specific type of information. These sections are designed to present data starting with a general overview of the array, followed by sections detailing its monthly and weekly status updates. Additionally, a section is dedicated to recorded events, providing insights into significant occurrences. Each section includes buttons that lead to relevant sub-pages, ensuring users can easily access detailed

information as needed.
The index page header provides a quick overview of the array status of the last processed file. It consists of four figures:

- **Spatial distribution of stations in the array and operated Antennas**: Displays the spatial distribution of stations within the array, highlighting the total number of operated antennas (Figure 3.6-top left). It illustrates the operational status of the grid system, showing the allocation of operated and non-operated stations across different categories, including the main (SD-1500), infill (SD-750), super-dense (SD-433), and test grids. The map also indicates the number and locations of offline stations (stations not sending data). Some of these offline stations, marked as gray points, have RD antennas installed but are still missing digitizers. Others are awaiting landowner agreements before RD antennas and associated electronics can be installed. Additionally, some stations have a "black tank" status, meaning the WCD is not sending data.[5]

- **Event distribution by number of radio stations**: Displays the number of events based on the number of RD stations involved (Figure 3.6-top right). Blue bars represent the regular grid, while red bars correspond to the Infill region. Most events involve fewer stations, with the count decreasing as the station multiplicity increases.

- **RMS map for both polarizations**: The color scale represents the selected metric from the dropdown menu, offering two mapping options: Average Mean RMS or Average Median RMS (Figure 3.6-bottom). A black star marks the location of a nearby TV transmitter, a likely source of RFI. Higher values (red/yellow) cluster near the transmitter, indicating strong interference, while lower values (blue) dominate most of the array, reflecting lower noise levels.

Maps

Map-based station information

- Array overview: A station distribution map that allows for easy identification of stations within the grid. The map provides a visual overview of station placements and includes a search function to locate stations by name or ID.

- Array History with Accumulated Station Data by Date: Displays the operational status of stations over the past three days, with different colors representing various dates to indicate when each station was active. The spatial distribution provides insights into the chronological progress of operations within the grid as shown in Figure 3.7. A dropdown menu allows users to filter data by specific 12-hour intervals or view the accumulated data.

[5]Radio data only requested when WCD trigger.

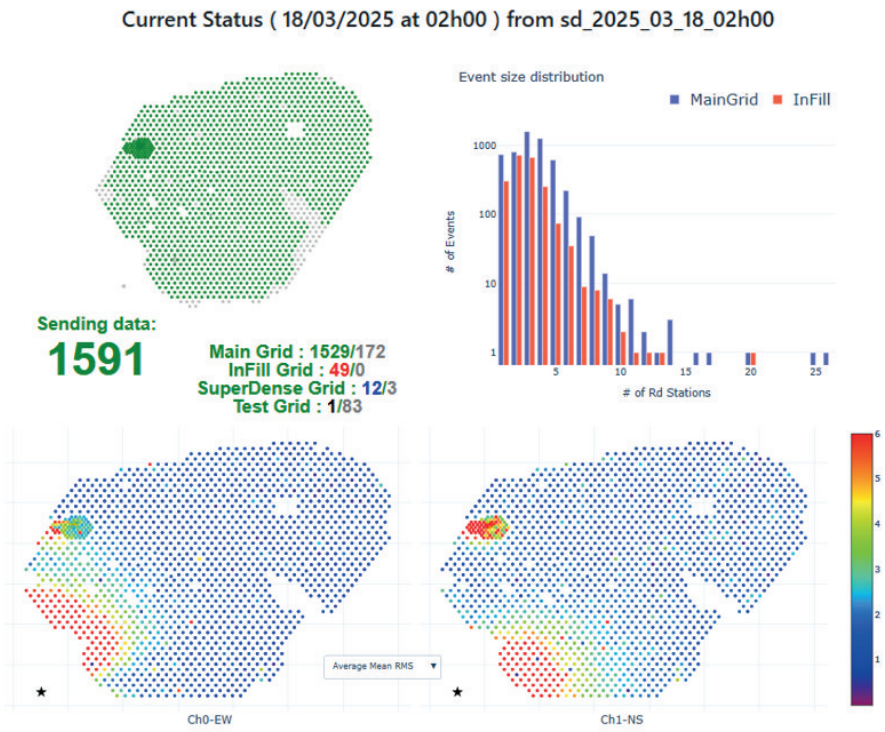


Figure 3.6: The current MoRD header page (updated on March 18, 2025) displays the time of the last update and its file name at the top.

- Trigger Volume: Shows the trigger volume for each file over approximately 12-hour intervals defined as the number of radio waveforms received. Each map represents the spatial distribution of triggered stations at different times, providing insights into their activity patterns. The infill region exhibits a higher trigger volume compared to other areas.

- Bad Periods and Status Map: This sub-page was developed by B. Pont[6], and the content is described in detail in [104].
  RD Bad Periods List: A downloadable list of RD bad periods is available, starting from January 1, 2022, with varying end dates specified in the files. These periods can be added to the RdStationRejector.xml in Offline mode to enable them as bad periods. Users must update the LastAllowedEventDateTime_RD in their application, which by default is set to the end date of the last committed bad period.
  Maintenance Overview Plot: A plot provides a status update of the RD array based on the last 7 days of data. Including roads map, antenna misalignment angles in color coding for easy recognition.
  Stations experiencing issues are marked with colored square markers, and a horizontal

[6]Repository: RDBadPeriods

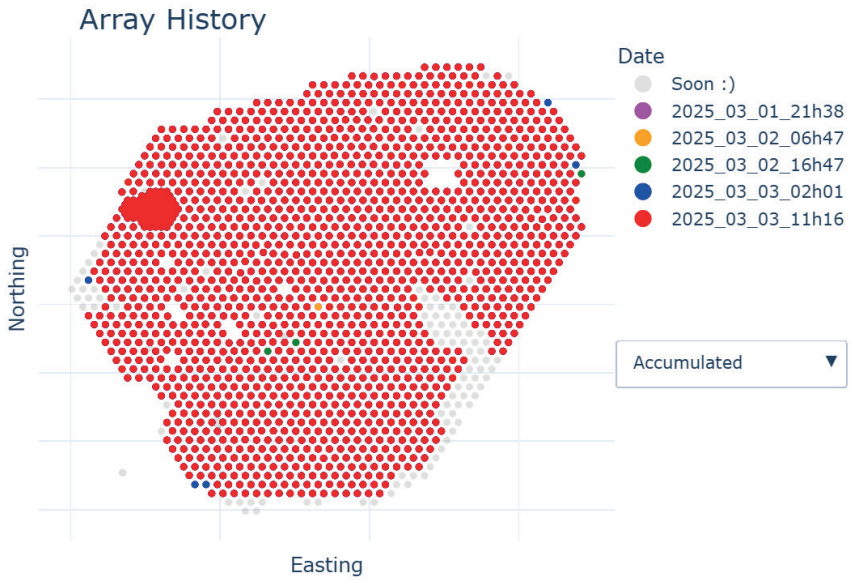


Figure 3.7: Visualization of the Array History, displaying station activity over time. Color-coded is according to its recorded status at different timestamps. Gray markers indicate offline stations, red indicate online station in the last processed data while colored markers represent active stations at specific time points. The dropdown menu allows users to view accumulated data or filter by individual timestamps.

bar beneath them provides a mini-timeline of problems, segmented in 3-day intervals for better visualization. Stations labeled as "variable alignment" have alignment arrows indicating shifts over time. Moreover, Monthly overview maps are available and provide a more stable status check compared to the 7-day plots, which can be affected by short-term noise conditions.

Spectrogram Plot: The spectrogram is generated using 20-minute time bins, displaying the median frequency spectrum over the last 7 days. Historical monthly spectrograms are also available for long-term monitoring. The processing workflow includes multiple steps:

- Data from Channel 0 (EW) and Channel 1 (NS) is analyzed, excluding bad periods and Infill regions to minimize noise interference.
- Raw spectra are normalized, and noise suppression is applied to remove narrow-band RFI, with a separate plot showing the removed RFI.
- Signal excess above the background is calculated using sidereal hour filtering, with deviation plots highlighting changes in specific frequency bands (e.g., 30-40 MHz for solar storm indicators).
- Frequency bins track known emitters (e.g., TV lines, beacon signals), and the final median spectrum displays both raw and filtered spectra.

Full-Array View

Provides performance metrics and insights for the entire array.

- Number of stations and readouts over time: The evolution of station count and activity from 2022 onward illustrated in Figure 3.8. The top section tracks the number of stations, while the middle section displays the readout count over time. Different grid types—regular grid (SD1500), infill grid (SD750), super-dense (SD433) grid, and test grid (used to test new features) are represented by distinct colors. Some drops in station count are observed, primarily due to non-operational WCD tanks (black tanks). The

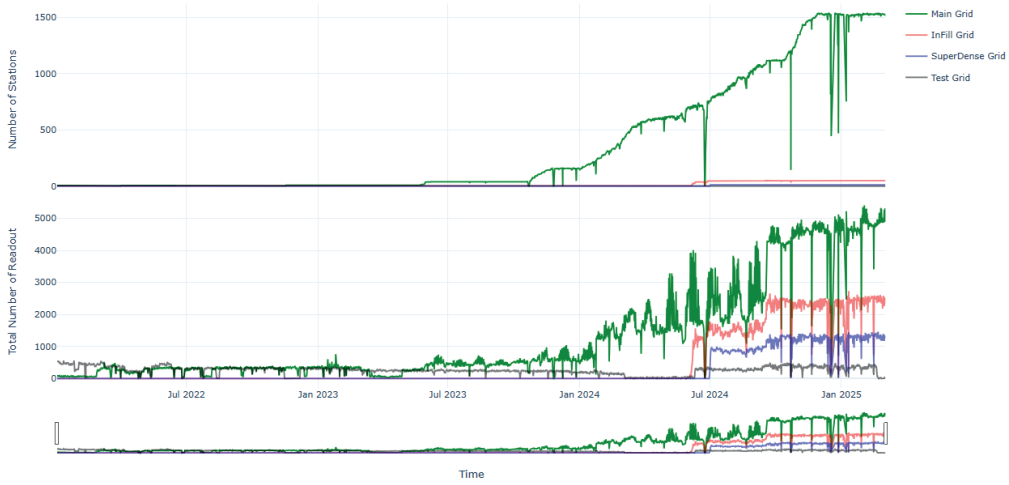


Figure 3.8: A MoRD page titled "Number of Operated RD Stations" presents the deployment history of the Auger Observatory's main and sub-grids, along with the number of triggers (readouts) per grid for all stations. The Infill grid became operational around June 2024, followed by the super-dense grid.

middle section displays readout activity, featuring a distinct, recognizable pattern shown in Figure 3.9. The FD's operation relies on sky conditions, particularly the moon cycle, externally triggering the WCD during clear, dark nights. This results in higher readout rates at night compared to daytime. Daytime files (ending with 00hxx or before 12h) consistently show lower rates than nighttime files (after 12h), reflecting FD's increased activity under optimal night sky conditions, which drives noticeable spikes in readout events.

- Station list: Displays a list of all stations that have been online at any point in time, including each station's name, grid classification, and timestamps of its first and last appearance sending radio data.

- Rd Trigger: Two operational stations were running with RD trigger firmware as of April 2024, but the RD trigger was deactivated in May 2024. Despite this deactivation, the

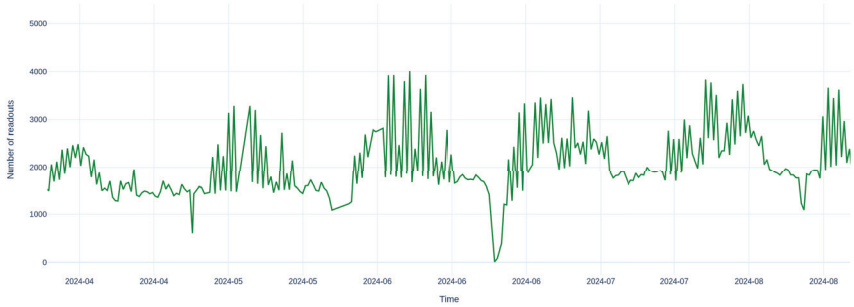


Figure 3.9: Screenshot from the MoRD monitoring framework showing the total number of readouts for the regular grid between March and August 2024. The recurring variations in the readout rate are associated with FD operation under favorable night-sky conditions, as discussed in the text.

page was developed and continues to operate to capture any future instances of RD trigger activation. The page displays all radio self-trigger instances, whether initiated solely by radio or in combination with WCD (particle) triggers. A table is provided with detailed information on specific triggers, including their types, station names, compatibility, and characteristics. The lower section of the page highlights RD trigger pulses, emphasizing key trigger events across both channels.

- Accumulated event size (Sd-1500, SD-750): Displays the accumulated frequency of events based on the number of radio stations in both the regular grid and infill grid since January 1, 2022.

Monthly spectrum

Provides a monthly view of the array's spectra.

- Spectral viewer: Provides a monitoring interface that displays frequency spectra and deployment status. Figure 3.10 shows a screenshot of the page.
  - A search bar allows users to filter specific radio stations.
  - A grid map visualizes the spatial distribution of stations, where each point represents a station (the station marker colors are arbitrary).
  - A legend interface shows all stations' names supporting manual search (e.g., using browser ctrl+F).

  Users can select multiple stations to visualize their spectra either by using the search bar or clicking on station legends. This is one of the core pages of MoRD, as it allows collaborators to check the spectra of any station compared to another station.

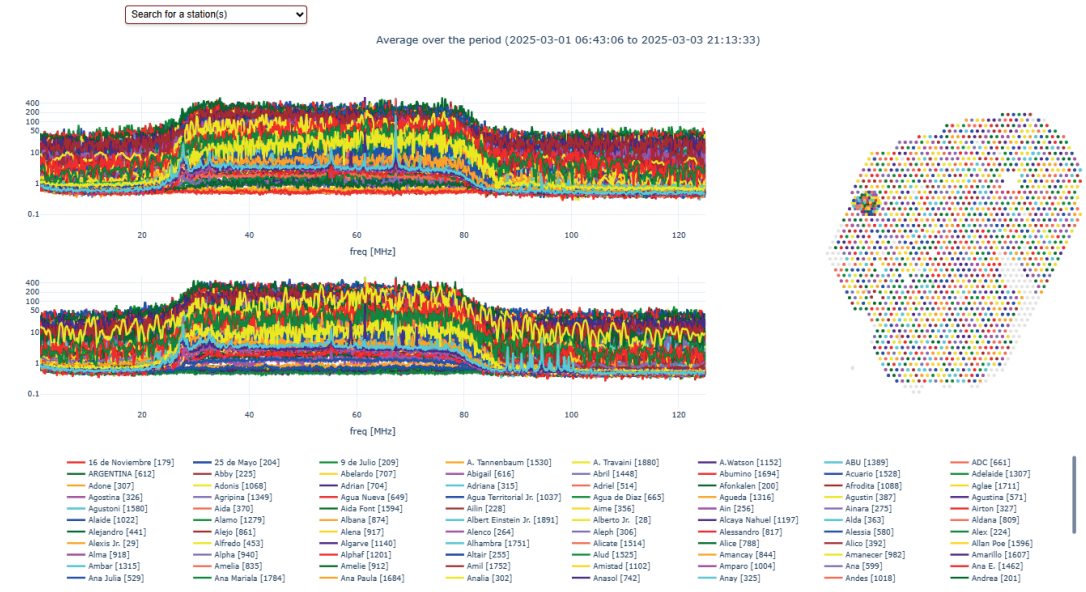


Figure 3.10: regular grid Spectral Viewer Page of MoRD. The top and bottom figures display Ch0 and Ch1, respectively.

- Current month frequency spectrum and time traces overview: Displays both the average frequency spectrum and time traces for various stations during the current month. Each station is color-coded and labeled, with a corresponding map showing the spatial distribution of the stations. Users can search, select, and compare different stations at various locations, providing a flexible and detailed analysis.

- Historical data access for frequency and time: This section provides a navigation interface to access historical data by month, categorized by frequency and time. This allows collaborators to check any station at any given time, making it a valuable tool to track the behavior of the array over time without the need to reprocess all data.

Weekly/Daily Variations

Provides a weekly and daily view of the array's spectra.

- Weekly frequency spectrum and time traces overview: Same feature as monthly spectral viewer but only for the last 7 days of available data.

- RMS values: The sub-page presents two figures:

  – The distribution of RMS values for both polarizations indicating the median, mode, and the 95^{th}-quantile. Given the range of RMS values, as discussed previously, each sub-plot is equipped with a slide-bar shown in Figure 3.11.

– The average RMS value per 12-hour interval for both channels over the past seven days. This allows users to identify days affected by thunderstorms and the regions of the array impacted. Thunderstorms typically saturate the channels gain, leading to high RMS values, which can be observed in the data.

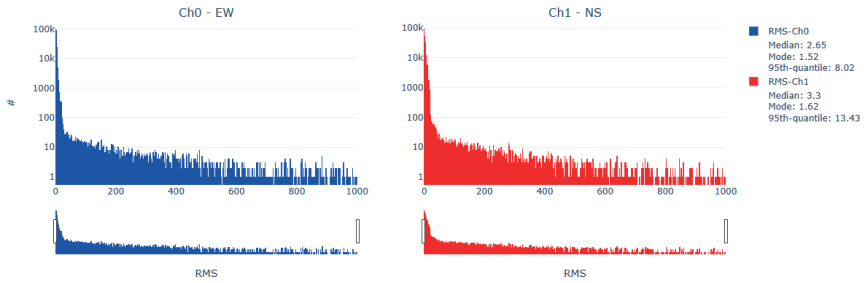


Figure 3.11: regular grid RMS for the Last 7 Days Page of MoRD. RMS values are calculated as described in 3.3.

- Latest file traces and frequency spectrum overview: Presents the average latest file trace for the stations as shown in Figure 3.12, excluding the Infill grid. This page is designed for collaborators who check MoRD daily and are more focused on daily variations rather than weekly or monthly averages. It includes both time-traces and frequency-domain data ranging from 20 to 100 MHz.

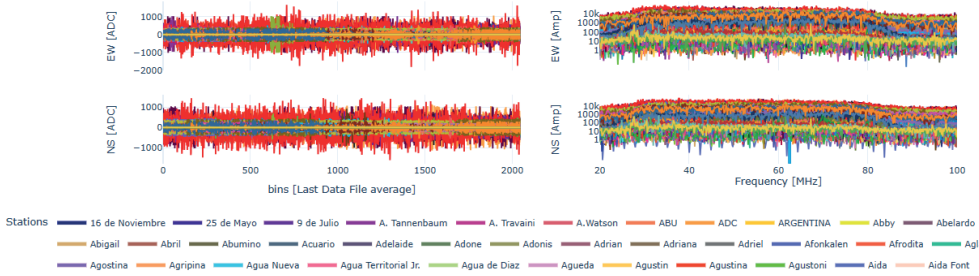


Figure 3.12: regular grid Stations' Traces for the Last 12 hours Page of MoRD.

- Spectral viewer for individual station: As the number of stations and the size of data increased exponentially, some users were less interested in the comparison feature, which allowed users to search for stations based on monthly or weekly viewership. Instead, they preferred to focus on individual station spectra. To address this, a dedicated page was developed where users can search for a station by name or ID and view its recent spectral data. Figure 3.13 provides an example of this functionality, showing the frequency spectrum analysis for the station "Jaco" (ID: 1742). This page displays both the weekly average and the most recent trace spectra for two channels (Ch0 and Ch1), and it also

indicates the number of traces available for this station over the past 7 days. In total, Jaco has 254 recorded traces.

The Latest Events

Provides an overview of events with radio data from the past three days for main and infill grids.

- Events with high multiplicity (regular grid, Infill grid): Displays all triggered events involving at least 15 stations in case of the regular grid and at least 10 stations for the infill. Users can navigate through the recorded events or download them using the provided buttons for each event.

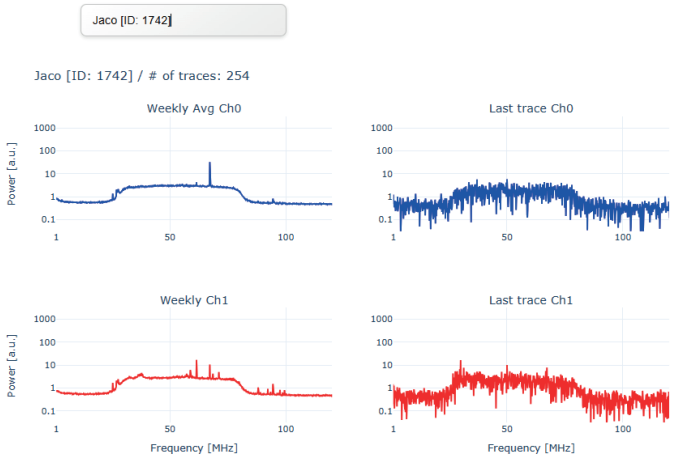


Figure 3.13: Frequency Spectrum for Jaco (ID: 1742) via individual station spectral viewer. The weekly spectrum is an average of 254 traces.

- Events with strong signals (regular grid, Infill grid): Displays all triggered events where at least 8 stations for regular grid and at least 5 for infill detected a strong signal. The signal selection method uses Signal-to-Noise Ratio (SNR), defined as the ratio of peak amplitude in the signal window (samples 1030-1070) to the RMS noise in the baseline (first 500 samples). Traces are selected when either channel (Ch0 or Ch1) shows an SNR above 4, ensuring only data with sufficiently strong signals relative to background noise. An example is shown in Figure 3.14.

Example of reconstructed events

A list that is manually updated, displaying selected events grouped by their energy and arrival direction. The list showcases notable and exceptional events which was particularly useful in the early stages for validating RD operation.

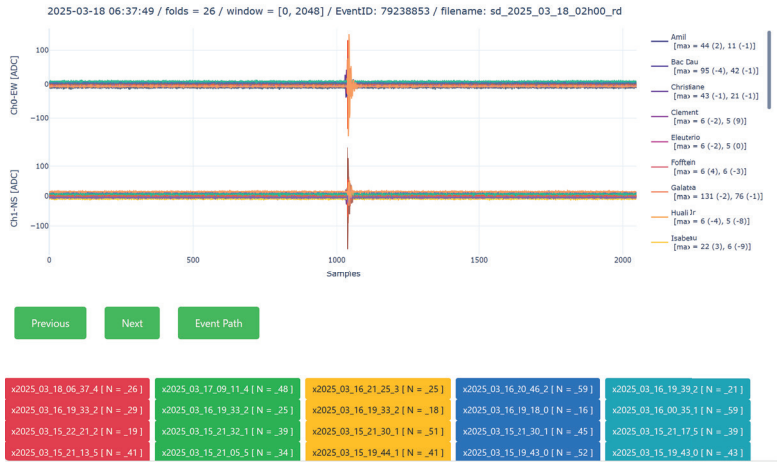


Figure 3.14: Events with at least 8 stations detecting signal-like within the signal window. This is the regular grid page in MoRD. The button colors are assigned arbitrarily for visual clarity.

Selected featured event until December 2024

EventID	Energy [EeV]	Zenith [deg]	Azimuth [deg]	# WCD tanks	# Rd antennae	View Event	Notes
75451494	20.55	59.24	233.16	15.0	3.0	View	Almost vertical event reconstruction led by a station in the cherenkov cone
76398224	13.5	79.36	143.3	22.0	43.0	View	Boom event in the infill
75363629	0.54	73.92	169.02	3.0	3.0	View	High quality low energy
74870325	21.32	85.68	260.65	61.0	29.0	View	High quality Southeren event
73570233	74.25	74.67	217.18	52.0	11.0	View	highest energy event measured during Superhexagon
77651647	15.51	88.89	241.26	66.0	60.0	View	Highest station multiplicity yet (May, 2025); the longest you can get in Pierre Auger Observatory
75377292	38.42	84.59	156.89	66.0	51.0	View	Legacy event for have a good luck
74870325	21.1	85.8	260.6	61.0	29.0	View	Legacy event for Rd deadtime

Figure 3.15: MoRD featured event page where the 'view' button opens the event display.

Figure 3.15 shows part of the page content where users can sort the selected events based on Energy, Zenith, Azimuth, or number of stations.

A dedicated event browser was developed to display all relevant WCD and RD data. Figure 3.16 shows an example of this browser's dashboard, which supports all MoRD features such as zooming and selecting. At the top of the display, the event ID, event time, and raw data file are shown. The interface is organized into three columns:

1. **Left column:** It contains six sub-figures. The first two sub-figures show the raw ADC signals recorded by both channels. The next three sub-figures show the reconstructed electric field signals in the Cartesian coordinates E_x, E_y, and E_z. The final sub-figure displays the signals transformed into the air-shower plane, where the $\vec{v}\times\vec{B}$ and $\vec{v}\times\vec{v}\times\vec{B}$ components appear as thick and thin lines, respectively. All these sub-figures display only the signal window (rather than the entire trace).

2. **Middle column:** This section presents a station map. The green and gray markers

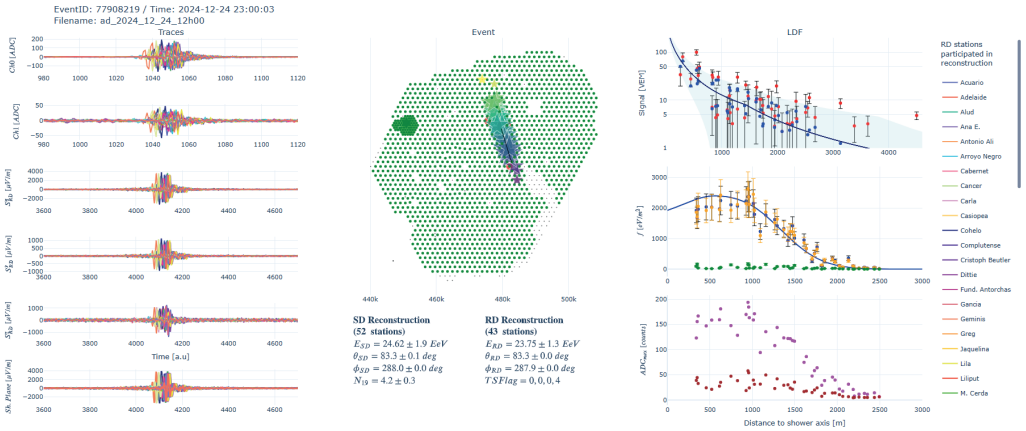


Figure 3.16: Event ID: 77908219 displayed using the MoRD Event Viewer.

represent online and offline stations at the time of the event, respectively, while stars indicate stations with a radio signal (colored by signal time). Red dots mark stations that took part in the WCD reconstruction but had weak or noisy signals, so they did not contribute to the RD reconstruction. A black line indicates the reconstructed arrival direction, and the WCD and RD cores are also marked on the map. Below the map is a summary of the WCD and RD reconstruction parameters: the number of stations used, energy, direction, muon content, and a thunderstorm flag for the radio reconstruction.

3. **Right column:** It features three plots. The first shows the WCD lateral distribution function (LDF), where red points represent measured signals and blue points represent horizontal-template-fit predictions [89] in VEM units. The fitted NKG function is drawn as a blue line. The second plot displays the radio LDF: green, orange, and blue points represent the signal fluence (in eV/m^2) in $\vec{v}\times\vec{B}$, $\vec{v}\times\vec{v}\times\vec{B}$, and the derived geomagnetic component, respectively, with a solid blue line showing the fit. The third plot shows the recorded signals (in ADC counts) of both channels as a function of their distance from the reconstructed core. A legend on the right allows users to select or hide specific stations or reconstruction components.

3.5 Evolution and Scaling

As the array expanded, both in the number of stations and in the data volume, several challenges arose in maintaining the functionality and performance of the website. The primary issue stemmed from the lack of a dedicated hosting server, as the project initially relied on static preloaded HTML pages. Since MoRD began as a local initiative and later gained collaboration support, scalability became a key concern.

One of the main challenges was visualizing and managing data from an increasing number of stations, growing from dozens to hundreds, and eventually surpassing 1,600 stations while keep-

ing website generation time within reasonable limits. The total generation time depended on both the amount of trace data processed and the complexity of the visualizations. The average number of readouts per station remained relatively constant, as shown in Figure 3.17 for four regular grid stations. In particular, station 734, located in the Infill region, exhibited a higher number of readouts due to its noisy background, a trend that was more pronounced when the array was relatively small.
To overcome this challenge, a database system was introduced that allowed new data to be stacked incrementally rather than reprocessing the entire dataset with each update. Moreover,

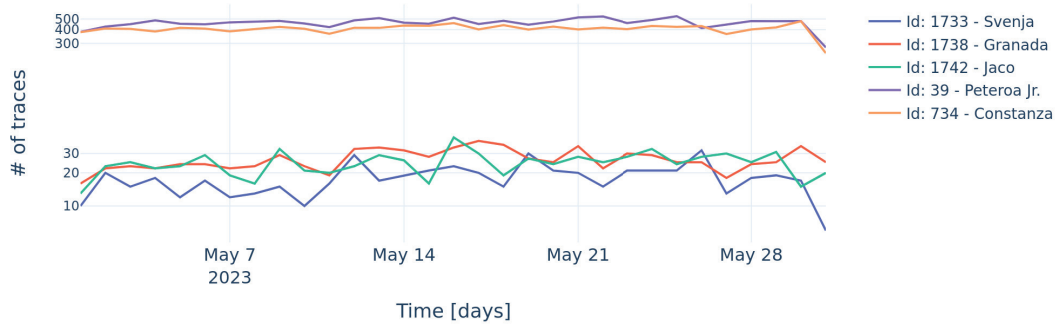


Figure 3.17: Daily number of triggers per station during May 2023. The regular grid station Constanza is located near the Infill area, while Peteroa Jr is an Infill station. regular grid stations outside the Infill area show fewer triggers. Overall, the trigger count remains stable over time.

the raw data file includes all stations, covering both the regular grid and various sub-grids within the Auger detector. Each grid is processed independently to ensure a tailored analysis, allowing for a more efficient handling of its unique characteristics and requirements. To improve performance, each update processes only the latest data, rather than reprocessing the entire dataset from scratch. Figure 3.18 illustrates the relationship between the number of regular grid stations and the computational time required to generate monthly spectrum grid view web pages. The x-axis represents the number of regular grid stations, while the y-axis denotes the generation time in seconds. Each data point is color-coded based on the number of readouts. The generation time increases exponentially as the number of stations grows, primarily due to the increased computational complexity associated with processing larger datasets, including data aggregation, formatting, and rendering operations. With the full array of stations operational, approximately 25 minutes is required to generate the page. This presents a scalability challenge, as further grid expansion or an increase in readout rates could result in prohibitively long update times. The station's monthly spectra and time traces were initially displayed in a grid format with five columns. The number of rows was calculated by dividing the total number of stations by five and taking the floor of the result. This grid-based format was visually

effective and manageable in the early stages when the number of operating stations increased at a steady and gradual pace. Users could easily navigate the subplots and compare station performance without difficulty.

However, a significant challenge arose at the beginning of 2024 with the mass deployment of new stations. By December 2023, the number of operational stations had reached 157, but by the end of January 2024, this number jumped to 256. This rapid expansion resulted in a grid with over 50 rows of subplots, making the original layout impractical for handling the expanded dataset. Navigating through such a large number of rows became cumbersome, and users struggled to compare stations effectively. The increasing station count also introduced more distinct spectral features and variations in performance, further complicating the analysis. Furthermore, when the number of operating stations exceeded 1,000, the page failed to load properly, as at the end of each month, it had to process more than 3.5×10^5 traces.

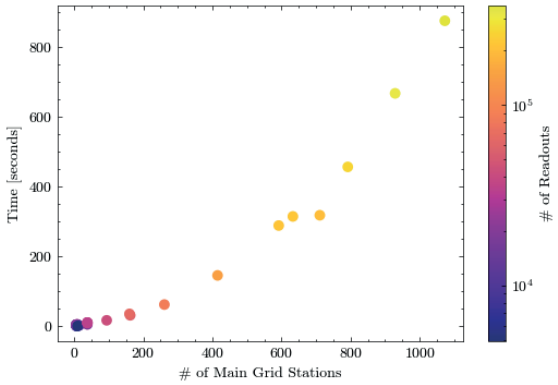


Figure 3.18: Computation time required to generate the regular grid spectra as a function of the number of regular grid stations.

Solution:

First, the workflow was optimized by switching to daily file processing instead of reprocessing the entire dataset. This adjustment significantly reduced the computational time to just 7 minutes by processing only the newest data, rather than reanalyzing the entire dataset from scratch. Table 3.1 summarizes the computational time for the MoRD pipeline. Although initialization is relatively quick, tasks such as array visualization, event size distribution, and monthly data processing require significantly more time.

Further optimization of the HTML format itself, particularly for preloading all data, proved challenging. Therefore, all stations' monthly average spectra are saved into a single plot that supports all necessary dynamic features, such as selecting, hiding, or searching. This plot is saved in JSON format, which can be called and fetched within the displayed HTML page. This approach significantly improves performance and scalability. Unlike dynamically generated Plotly figures, which require back-end computations each time they are loaded, the JSON-based method allows the browser to instantly retrieve and display data, reducing computational over-

Process	Time (s)
Imports and Paths	19.6
Array_View	284.6
Array_View_Focused	288.5
EventSizeDistributionLatest	290.7
UpdateStationList	349.5
RdTrigger	0.1
GetTriggerVolume	2.3
EventSizeDistribution	14.0
TriggerRatePlusRDstations	45.4
Weekly Data Ready	64.4
WeeklyFreqTracesOne	75.3
WeeklyTimeTracesOne	107.7
Weekly RMS	162.6
LastFileTracesTrain	178.3
Monthly Data Ready	264.3
FreqTracesGrid	408.2
FreqTracesOnePlot	419.2
TimeTracesGrid	638.0
TimeTracesOnePlot	670.4
MainGrid Events	305.1
InFill Events	417.7
Pages Automation	0.5
Total Time	**1682.6**

Table 3.1: Processing Times Summary (Wed 26 Feb 2025).

head. The JSON structure is also lightweight and easily parsed, ensuring smooth interaction without delays. In addition, a spatial map subplot displays the geographical distribution of the selected stations within the array.
This feature is essential for comparing station spectra, highlighting regional variations, and facilitating spatial analysis.

3.6 Using MoRD in the Commissioning Phase

MoRD was initially designed for real-time deployment monitoring, but quickly proved valuable beyond its original scope. As the number of stations increased, unexpected challenges emerged, including hardware failures, misconfigurations, and environmental interference. Given the array's size, it became essential to accurately pinpoint the location of any issues.
MoRD capability to process and visualize data in near real-time transformed it into a powerful diagnostic tool, allowing the early detection of operational anomalies. Rather than relying solely on post-deployment quality checks, researchers could proactively assist local technical staff by addressing issues as they arose. The success of the experiment depends on the precise configuration of essential components, such as the LNAs and antenna orientation. During deployment, MoRD was the key tool in detecting and resolving critical issues, such as swapped antenna arms, malfunctioning LNA, and significant RFI that compromised data collection, ul-

timately optimizing the RD array.

RD commissioning is outside the scope of this thesis; however, it is important to describe the role of MoRD in supporting the commissioning team in their tasks. MoRD raw data can be used to construct spectrograms that quickly diagnose the status of RD stations. Since there is a strong TV emitter near the array, as shown in Figure 3.4 during the Super Hexagon phase, this signal was used as an external reference to indirectly check the antenna response, especially in the southern and western parts of the array, where the signal is strong enough for spectral analysis. By moving from the time domain to the frequency domain, we can inspect the spectral content and calculate the average frequency spectrum for each station separately for the EW and NS channels. The strength of the TV line is then estimated by taking the maximum spectral amplitude in the frequency interval 66.9–67.5 MHz, corresponding to the selected FFT bins 549–553. This gives one amplitude value for each channel, A_{EW} and A_{NS}. The relative line strengths are calculated as

$$R_{\mathrm{EW}} = \frac{A_{\mathrm{EW}}}{A_{\mathrm{EW}} + A_{\mathrm{NS}}}, \qquad R_{\mathrm{NS}} = \frac{A_{\mathrm{NS}}}{A_{\mathrm{EW}} + A_{\mathrm{NS}}}. \tag{3.1}$$

Here, A_{EW} and A_{NS} are the maximum amplitudes found in the 66.9–67.5 MHz interval in the EW and NS spectra, respectively. This normalization compares the relative contribution of the two channels to the observed TV line and provides insight into the directional and polarization-dependent response of the antenna.

The result shown in Figure 3.19 for June 2023, when the Super Hexagon was operational, highlights several key features on the plot:

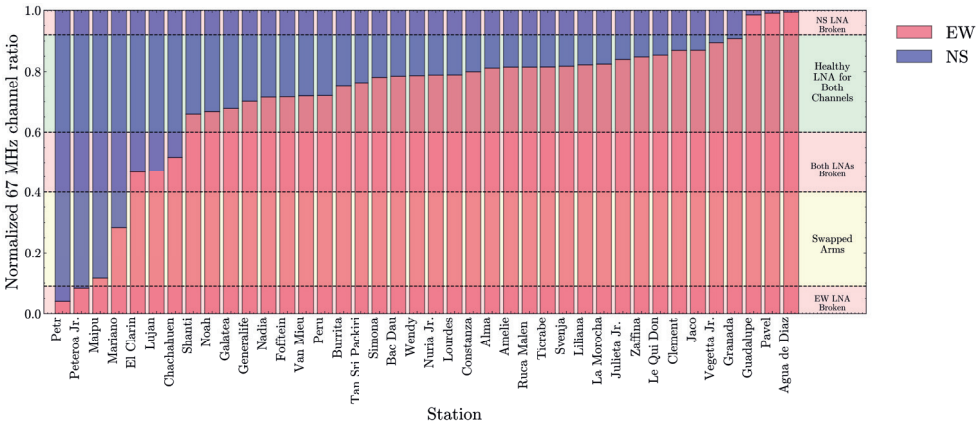


Figure 3.19: June 2023 data of the relative 67 MHz line strength between the two polarizations for various stations, displaying the contributions from EW (red) and NS (blue) polarizations. For each station, the maximum spectral amplitude is extracted in the 66.9–67.5 MHz frequency interval for both channels and normalized by the sum of the two channel amplitudes. The annotations on the right indicate the health status of the LNAs, identifying cases where the NS or EW LNA is broken, both LNAs are broken, or the antenna arms are swapped.

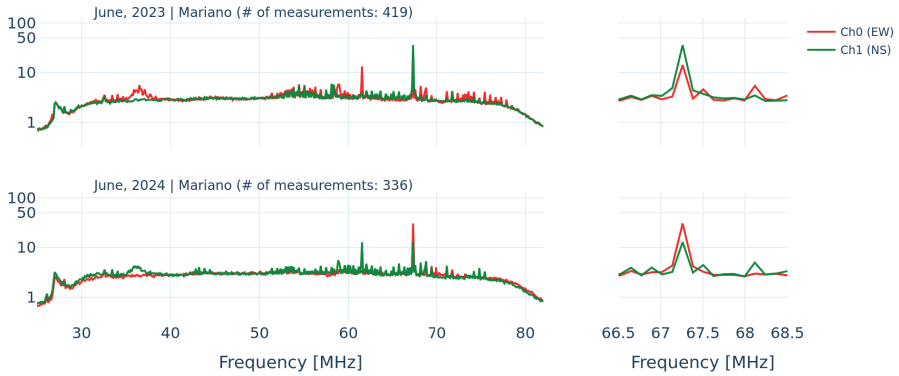


Figure 3.20: Comparison of frequency spectra recorded at Mariano station in June 2023 (top) and June 2024 (bottom). The left panels show the full frequency range, while the right panels zoom in on the 67-MHz region. The comparison shows the polarization-channel configuration before and after correction of the swapped antenna-arm cables.

LNA status

When the power/gain of one polarization direction is minimal, it suggests a failure in the signal chain, such as an LNA failure. Stations with fractional contributions to the 67-MHz line of less than 10% in the EW direction typically have a broken EW LNA. For example, stations Petr and Peteroa Jr. exhibit this behavior. Similarly, in the NS direction, when the line power is dominant in the EW direction, such that the power in NS is less than 10% of the total power, it suggests a broken NS LNA. For example, stations Guadalupe, Pavel, and Agua de Diaz show this pattern. Moreover, when the power of the TV line is competing between both channels, it indicates that the LNA is broken for both directions, as neither channel is receiving the correct power.

Swapped Arms

Another issue arose during the early stages when the antenna arm cables were switched mistakenly during deployment, causing a polarization direction change. Since the TV transmitter is located south-west of the array, we expect a stronger signal in the EW polarization direction. Therefore, when the contribution of the 67-MHz line in the EW channel at the Super Hexagon location is less than 40-50%, it indicates a swapped arm issue. For example, the Mariano station showed this issue, as shown in Figure 3.20. When comparing the June 2023 data with the June 2024 data, it is clear that the cables were fixed, and the two channels now function as expected. This is an easy fix, as the commissioning team simply directed the deployment team to the affected station, where the issue was resolved by switching the cables. The station also lacks the 67-MHz line and exhibits lower gain compared to its peers.

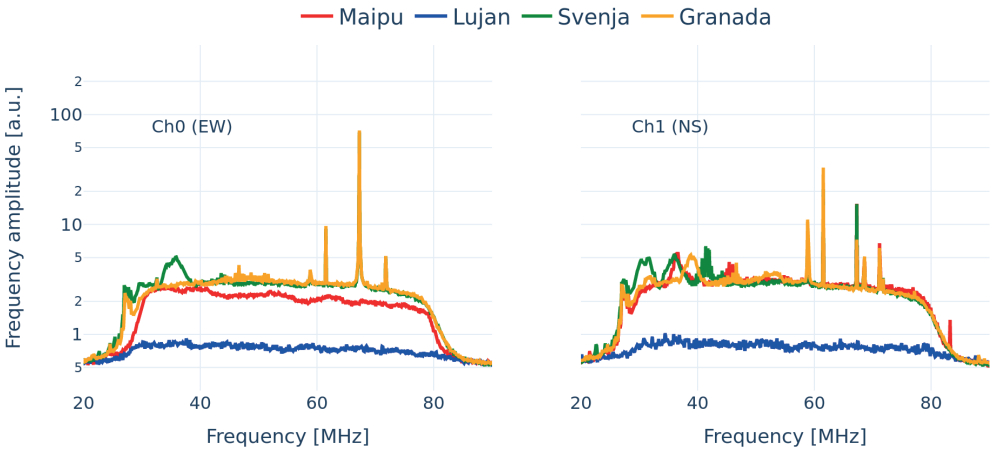


Figure 3.21: Examples of frequency spectra from RD stations in June 2023 illustrating different LNA conditions. Svenja and Granada show normal LNA operation, Maipu shows a degraded or failed LNA in one polarization channel, and Lujan exhibits failures affecting both polarization channels.

Healthy LNAs

When the 67-MHz line shows more power in the EW direction with some power in the NS direction, it is a sign of healthy operational LNAs. Figure 3.21 shows a subset of stations where the behavior of a broken LNA is clearly visible, with minimal gain. However, the presence of healthy LNAs does not always guarantee that the system is operating as it should. For example, station Maipu appears normal in the 67-MHz spectral analysis. The spectral analysis of the Maipu station in Figure 3.22 revealed a significant issue in June 2023. The overall intensity of Ch0 in 2023 was noticeably lower compared to Ch1, indicating a systematic reduction in signal strength. This behavior strongly suggests a failure of the LNA responsible for the EW polarization.

Another example of a low gain issue is shown in Figure 3.23, where using the MoRDspectral viewer to compare the spectra of two neighboring stations, Netherlands and Poland, we observe that the Netherlands station still has the 67-MHz line but exhibits lower gain in the NS polarization.

Maipu and Netherlands are examples of partial LNA failures, which do not necessarily result in a complete loss of signal, but can lead to a reduction in gain, causing weaker spectral intensity and selective loss of specific frequency components, such as the 67-MHz feature. The fact that the overall spectral shape appears normal, but with lower power in Ch0, further supports the hypothesis that the LNA was not entirely nonfunctional, but rather operating at reduced efficiency.

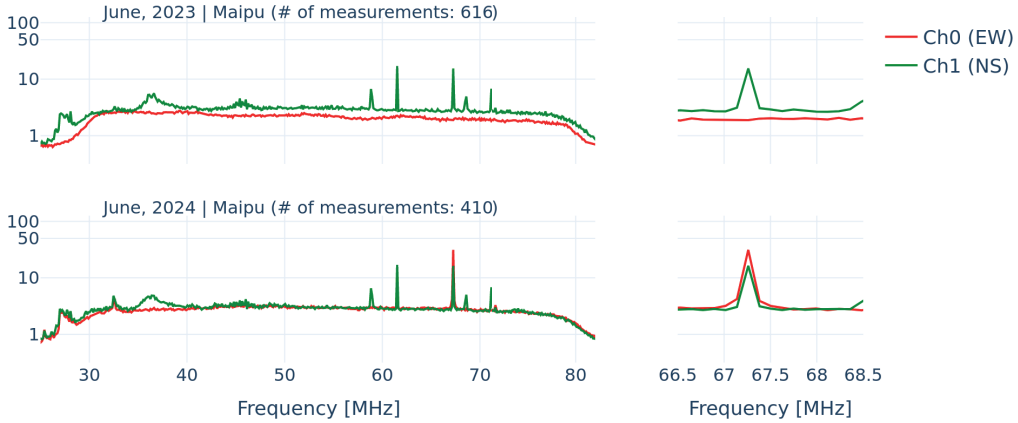


Figure 3.22: Frequency spectra recorded at the Maipu station in June 2023 (top) and June 2024 (bottom), illustrating the detector response before and after replacement of the EW-polarization LNA. The restored response of the EW channel is visible in the 2024 data.

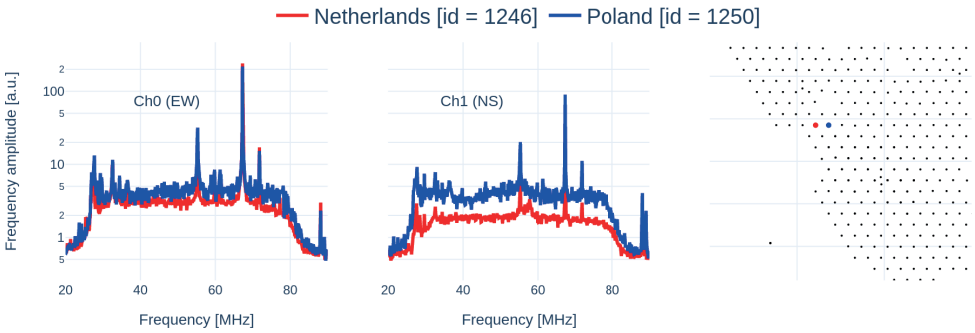


Figure 3.23: Comparison of frequency spectra between two stations, Netherlands (id = 1246) and Poland (id = 1250). The first two panels show the EW and NS polarization, while the last panel indicates their location on the array (South-West). Poland exhibits stronger frequency amplitudes, especially within the 67-MHz range.

Another deployment issue where MoRD data were essential occurred when the antennas were mounted on top of WCD tanks without proper support.
As a result, under strong winds, the antennas rotated, causing them to no longer face north, leading to data quality problems. This issue was covered [104] as the antenna rotation is determined by comparing the ratio of signals received in the two antenna polarizations at a certain frequency (TV line of 67-MHz) to the expected ratio from the antenna model. This comparison helps assess alignment, though there is some ambiguity in quadrant determination. However, the method is sufficient to determine if an antenna is well-aligned or misaligned beyond 20 degrees. To ensure accuracy, background radio noise is subtracted from the measured signal. Hence, misalignment can be calculated for any station with a reasonable 67-MHz strength, namely, closer to south-west region of the array. These data can be found in map of misalignment in the RD Bad period page of MoRD. The commissioning team identified this issue in November 2023 and responded by installing rivets to support the mounted antennas. This requirement was then incorporated into future deployments, saving significant effort before mass deployment began in 2024 and eliminating the need to revisit all stations. Additional data quality checks are discussed further in chapter 5.

Chapter 4

Potential photon sensitivity of WCD and RD

The Pierre Auger Observatory has extensively searched for UHE photons, setting strong flux limits (section 1.5) and no event has been unambiguously confirmed as a photon. Prior searches focused on vertical air showers ($\theta \leq 60°$), where a FD observable can support identification in conjunction with the WCD. This leaves a gap in the detection of very inclined photons ($\theta \gtrsim 65°$). For inclined photon-induced air showers, the electromagnetic cascade produces few muons, and the long atmospheric slant depth causes further muon attenuation, so the electromagnetic component never reaches the WCD. This suppresses the WCD signal, making a trigger unlikely. This gap presents a unique opportunity for the RD array. The atmosphere is largely transparent to the radio emission from the electromagnetic cascade of a photon-induced air shower. Consequently, the RD can potentially trigger on signals with more stations than the muon-signal-dominated WCD. This capability can be enabled by a radio self-trigger, where antennas initiate detection independently, unlocking sensitivity to the population of very inclined UHE photons. Currently, the radio system operates in a slave-mode, requiring a WCD trigger. In this chapter, we present a comprehensive investigation into the potential of an RD self-trigger to enhance the observatory's photon sensitivity. We move beyond a conceptual discussion to a quantitative analysis, using simulations to assess trigger efficiencies, reconstruction performance, and the final effective aperture. We demonstrate that a self-triggering RD array can recover nearly all photon-induced air showers missed by the WCD at high zenith angles and achieves a lower energy threshold. This analysis provides support for the development of a new, promising channel in the search for UHE photons.

4.1 Threshold-Based Assessment of Hybrid Trigger Efficiency

The potential of the radio self-trigger can be illustrated with a typical scenario. A photon-induced air shower of 10 EeV arriving at $\theta = 75°$ will, on average, trigger only two WCD stations falling short of the detection requirement (see section 2.4). However, that same air

shower can potentially trigger more than ten radio antennas, based on conservative estimates in section 4.2. This contrast demonstrates that a strong radio signal, especially when coupled with a weak or absent WCD trigger, provides a distinctive signature for photon candidates. This radio self-triggering allows the RD to detect air showers independently when the WCD fails or augment the WCD trigger events with more stations. To test and quantify how much gain in photon detection is maximally achievable from a pure physics point of view, before the full reconstruction chain implementing the detector response, as well as accounting for measured noise using the simulation set described in appendix A, we investigate the hadronic and electromagnetic components at the observation level based on two parameters:

- The muonic component, inferred from the number of muons per m^2 reaching the ground, labeled ρ.
- The electromagnetic component, represented by the absolute maximum electric field strength recorded by the radio antennas, labeled $\vec{E}_{\text{max}}$ measured in μV/m, which scales linearly with the electromagnetic air-shower energy.

A simple threshold is then used as a trigger condition for both quantities. This approach leaves the possibility open of applying more complex trigger conditions in future analyses. Since the simulations use fixed antenna positions and record all particles that reach the ground; air-shower thinning is accounted for in the muon–density maps. For each simulated event, we construct a map of the muon density ρ in the air-shower plane to determine its value at the location of each WCD tank. The radio signal strength $\vec{E}_{\text{max}}$ is obtained from the maximum of the electric field in the $-\vec{v}\times\vec{B}$ direction after applying a bandpass filter from 30−80 MHz to match the frequency range of the RD.

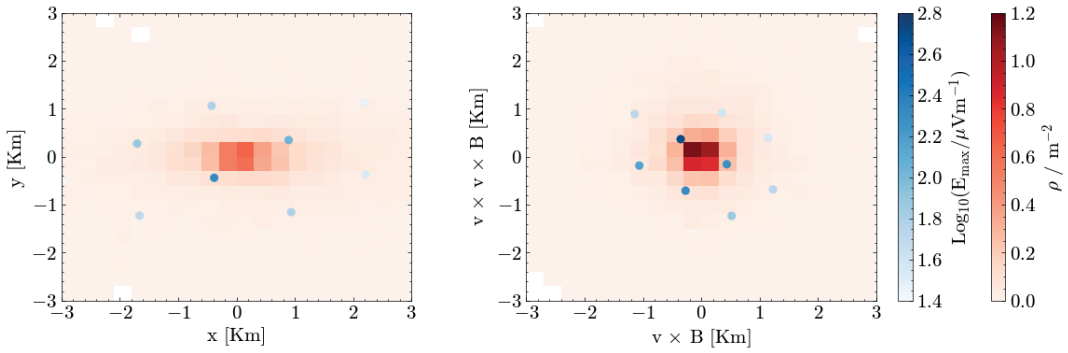


Figure 4.1: A typical photon air shower muon map ($E_{MC} = 4$ EeV, $\theta_{MC} = 65°$) in the observation plane (Left) and air-shower plane (Right). The dots mark the WCD tank positions while the color scale in both plots represents $\vec{E}_{\text{max}}$ and ρ recorded per station, respectively.

An example of a muon map and electric field map is shown in Figure 4.1. As air showers develop, both muons and the electromagnetic component become more dispersed in the observation plane. This increased dispersion reduces the local muon density and radio signal strength due to the diluted charge separation and at the same time spreads the signal over a larger area,

covering more detector stations.
Each parameter was evaluated using two thresholds: optimistic and conservative estimates. For ρ, we adopt an optimistic threshold of 0.15 m^{-2} and a more conservative threshold of 0.3 m^{-2}. Representing estimates for 10 EeV air showers at $\theta = 70°$ (including atmospheric attenuation and path-length corrections), these values are motivated by the properties of the surface detectors. Given an effective tank area of $A_{\text{tank}} \approx 10\,m^2$, these densities yield an average of 1.5 and 3 muons per tank, respectively:

$$\langle N_\mu^{\text{tank}} \rangle = \rho A_{\text{tank}}$$

Since an inclined muon produces between 0 and 3 VEM depending on its track length [105], these thresholds probe the regime where one to a few inclined muons can generate a trigger-level signal in an individual tank."
For hadron-induced air showers, the number of muons approximately scales as

$$N_\mu \propto E_0^{0.85}$$

[22, 106]. This scaling is not directly applicable to photon-induced air showers, which are muon-poor because their development is dominated by electromagnetic interactions, with muons produced mainly through photonuclear interactions.
In vertical cosmic-ray air showers, many muons reach the ground and produce characteristic signals in water-Cherenkov detectors. For very inclined air showers, the electromagnetic component is largely attenuated before reaching the detector, so the ground signal is dominated by muons. Inclined muons can produce larger signals mainly because they cross longer paths through the water-Cherenkov tanks. Thus, the lower threshold used for inclined showers should be interpreted as an approximate proxy for station triggering rather than a full detector-response simulation [59, 107].
For $\vec{E}_{\text{max}}$ we used 50 and 300 μV/m thresholds, a preliminary choice made at the beginning of the study based on expected radio detection signal values. To validate these thresholds, we analyzed commissioning data from events with azimuthal directions between 250° and 290°, a range where the geomagnetic signal is predominantly deposited in the EW direction.
As shown in Figure 4.2, a stable linear correlation is observed between the maximum of the EW channel in ADC counts and the absolute maximum of the reconstructed air-shower-plane signal in the $-\vec{v}\times\vec{B}$ direction. This linear behavior is most pronounced in the range of approximately 10–300 ADC counts, corresponding to reconstructed air-shower-plane signal amplitudes of about 50–1000 μV/m. The analysis confirms that the selected thresholds are well positioned within the stable correlation region, starting from the minimum stable signal and spanning a substantial portion of the linear range, where 50 μV/m and 300 μV/m correspond on average to 12 and 67 ADC counts, respectively.
The first step in evaluating a potential trigger is assessing the station multiplicity. Figure 4.3 shows the cumulative number of photon-induced air showers in the simulation set described

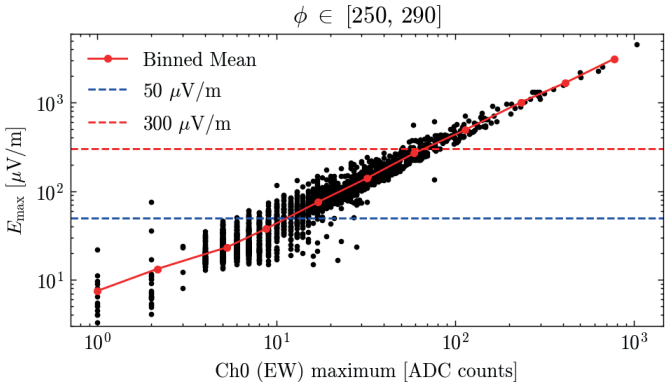


Figure 4.2: Correlation between the maximum signal in the EW polarization channel (ADC) and the reconstructed air-shower-plane signal (μV/m) for events with azimuthal directions between 250° and 290°.

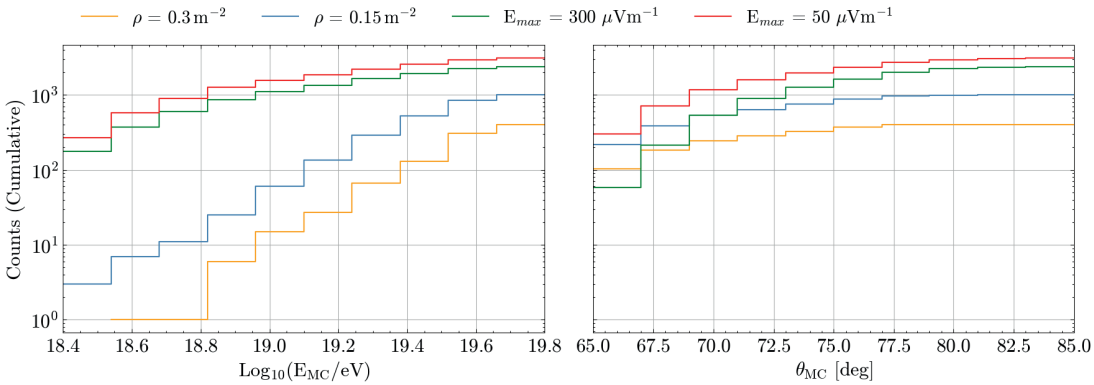


Figure 4.3: Cumulative number of simulated photon-induced air showers per threshold meeting trigger condition, versus energy (left) and zenith (right).

in appendix A that satisfy a $\geq$3-station condition versus energy and zenith. The selection is evaluated separately, for ρ, which acts as a proxy for the WCD trigger, and $E_{\max}$, which serves as a proxy for the RD response. The relative performance of the WCD and radio detectors depends strongly on the analysis scenario. At small to moderate inclinations ($\theta \lesssim 68°$), WCD efficiency dominates in the pessimistic case, while radio response dominates in the optimistic case. With increasing zenith, the WCD rapidly loses efficiency and radio acceptance rises with the expanding footprint. Beyond $\sim$ 82°, the WCD no longer meets the 3-station condition, whereas the radio detector continues to trigger. The reconstruction quality is driven by the number and layout of triggered radio (N_{RD}) and WCD (N_{WCD}) stations. The Figure 4.4 summarizes its dependence on energy and zenith. The dominant trend is geometric: N_{RD} grows much more strongly with zenith than with energy. Above $\sim$75° the footprint broadens rapidly; in the 80°–85° bin N_{RD} is typically a factor of 3–4 larger than in 75°–80°. A lower radio threshold increases station multiplicity in all bins, especially at very inclined angles.

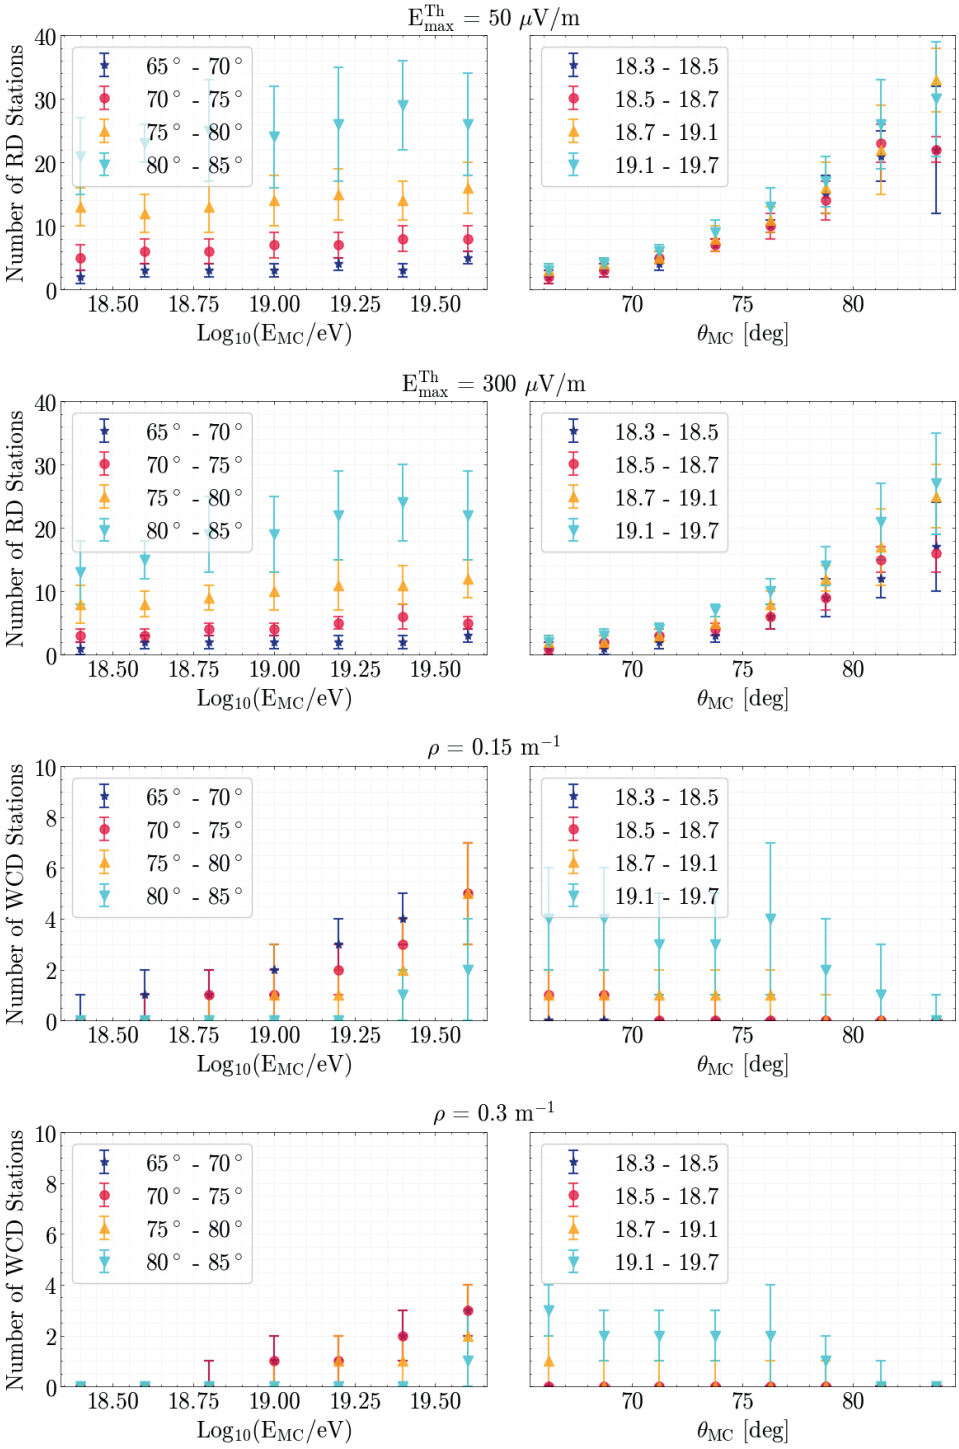


Figure 4.4: Number of stations as a function of energy (left, binned in zenith angle) and zenith angle (right, binned in energy), shown for two radio-based thresholds, $E^{\mathrm{Th}}_{\mathrm{max}} = 50\ \mu\mathrm{V\,m^{-1}}$ and $300\ \mu\mathrm{V\,m^{-1}}$, and two particle-based thresholds, $\rho = 0.15\ \mathrm{m^{-1}}$ and $0.3\ \mathrm{m^{-1}}$. Error bars represent the mean $\pm 1\sigma$ (68% confidence interval for a normal distribution).

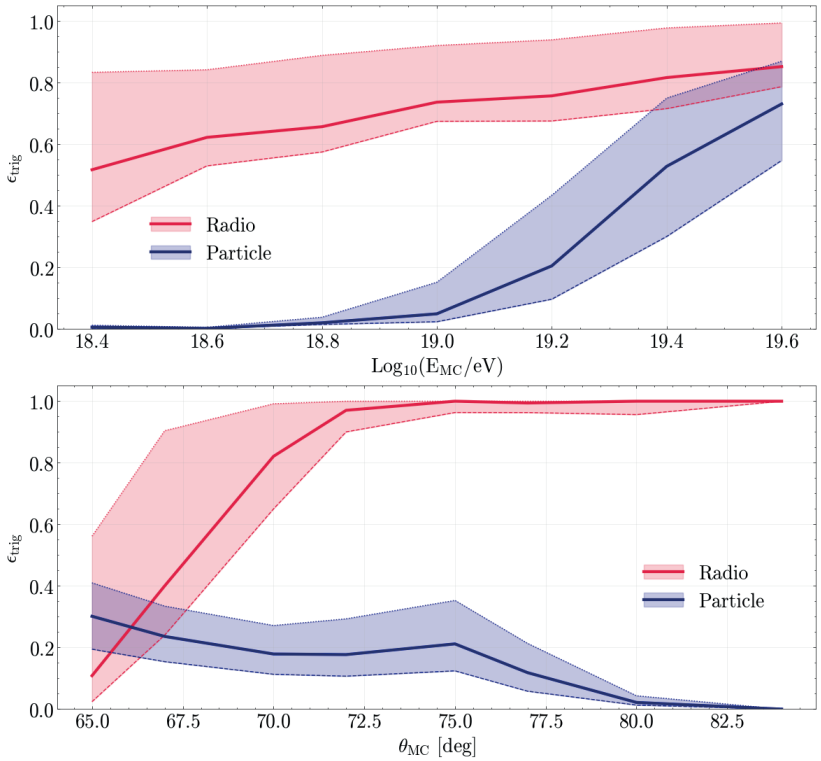


Figure 4.5: Trigger efficiency comparison between Radio (Red) and Particles (Blue). The dashed, solid, and dotted lines for Particles represent thresholds ρ_μ of 0.15, 0.22, and 0.3 m^{-2}, respectively, when using particle information alone. For Radio, the dashed, solid, and dotted lines correspond to E_{max}^{Th} values of 50, 300, and 500 μV/m, respectively, when using radio information alone.

The trigger efficiency (ϵ_{trig}) for photons is defined as the fraction of events that meet the trigger condition in each energy bin.

$$\epsilon_{trig} = \frac{N(n \geq 3)}{N} \tag{4.1}$$

where n is the station multiplicity per detector and N is the entries per bin.

In Figure 4.5, the trigger efficiency is compared for both radio and particle detection thresholds as a function of primary energy and zenith angle.

Radio detection consistently exhibits a higher ϵ_{trig} across all energies compared to particle detection. The radio efficiency remains above 0.5 even at lower energies ($\sim 10^{18.4}$ eV) and steadily increases. Unfortunately, for cosmogenic photons this energy is relatively high. This gradual increase suggests that radio-triggered events are not as strongly dependent on energy as particle-triggered events. In contrast, the particle detection efficiency remains very low until approximately 10^{19} eV, after which it exhibits a sharp rise. This sharp increase reflects the energy-dependent nature of air-shower development, where higher-energy air showers produce larger muon densities at ground level.

Despite this rise, particle efficiency still lags behind radio detection. The dependence of ϵ_{trig} on the zenith angle is where the differences between the two detection methods are most pronounced. For radio detection, the efficiency rises rapidly with increasing zenith angle, reaching a full efficiency at angles above 75°. This is because the radio footprint expands significantly at larger inclinations, ensuring more stations are triggered. Particle trigger efficiency is at its highest at lower zenith angles and then declines as the angle increases, dropping to near zero above 80°. The particle efficiency shows a notable inflection point around 75°, corresponding to a transition where atmospheric attenuation begins to dominate over the production of secondary muons.

The significant gap in trigger efficiency between the two methods, particularly in the intermediate zenith range of 65° to 75°, suggests that a hybrid detection system would be highly beneficial. Radio detection can effectively complement particle detection by providing robust event triggers at higher inclinations and across a wider energy range. At higher energies and intermediate zenith angles, a hybrid system could offer increased detection redundancy and improved event reconstruction quality.

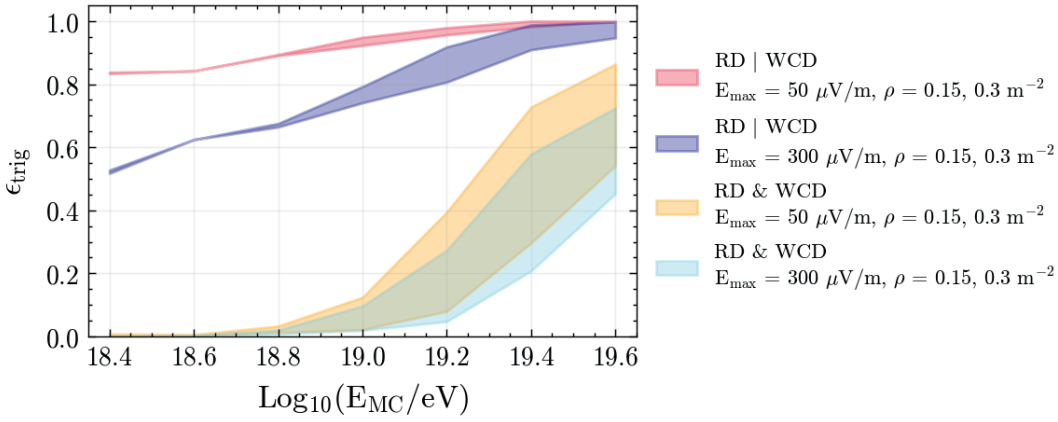


Figure 4.6: Trigger efficiency comparison for different configurations using WCD and RD. Triggered events are defined as having at least 3 antennas AND 3 tanks above their respective thresholds (AND configuration) or at least 3 antennas OR 3 tanks above their thresholds (OR configuration).

We further explore hybrid trigger possibilities for photon-induced air showers. For combined WCD and RD triggers, the second trigger level of TOT (T2-ToT) at 2.2 VEM ($\rho \sim 0.2\,\text{m}^{-2}$) becomes relevant only when the radio footprint is small. As shown in Figure 4.6, we assess the trigger efficiency for different combined detection configurations, comparing two main scenarios:

- **OR Configuration (RD — WCD):** This configuration requires a trigger from at least three stations of either RD or WCD at fixed ρ. This OR logic achieves the highest trigger efficiency, reaching nearly 100% at the highest energies. Lower radio thresholds significantly increase efficiency, especially at lower energies where efficiency remains above 50% throughout the energy range.

- **AND Configuration (RD+WCD):** This configuration requires simultaneous triggers from at least three stations of each detector type. The AND efficiency is significantly lower than the OR configuration, remaining near zero at lower energies and rising sharply only above $10^{19.2}$ eV. This indicates that simultaneous triggers from both detectors are rare at lower energies.

This initial assessment, while powerful for establishing a theoretical upper bound on performance, employs several simplifications. The trigger condition used here—a simple threshold on station-level observables—is not a full simulation of the Auger Observatory's complex multi-level triggering logic. Furthermore, the WCD response is approximated using only muon density, neglecting the contributions of other particles and the detailed tank response to an electromagnetic component. On the radio side, this analysis does not yet incorporate the effects of ambient noise. Finally, we evaluate only the trigger stage, without applying the full reconstruction and selection cuts necessary to identify a photon candidate. Consequently, the efficiencies derived represent a maximum achievable limit from a pure physics perspective.
They provide a crucial baseline against which the performance of a more realistic, full-detector simulation can be compared. Consequently, the efficiencies and apertures derived here represent a maximum achievable limit from a pure physics perspective. They provide an optimistic baseline against which the performance of a more realistic, full-detector simulation—including background rejection—must be compared.

4.2 Photon Reconstruction with WCD, Radio, and Hybrid Triggers

Having established the theoretical maximum detector efficiency using idealized thresholds, we now turn to a more realistic performance assessment. In this section, we reconstruct the photon simulation set by incorporating the full detector response and measured noise. The goal is to evaluate and compare the actual reconstruction efficiency and quality for photons using three distinct operational modes: standard WCD-only, radio-only, and a hybrid approach.
The reconstruction was carried out using **Offline** [108, 109]. In this study, the reconstruction quality and consistency were not evaluated, as the primary motivation was to determine if events can be triggered and establish radio detection potential. However, more robust station and event quality cuts will be applied in chapter 5 and chapter 7.
Accordingly, an event was deemed successfully reconstructed if the pipeline returned any physical (finite, non-null) values for the energy and arrival direction To comprehensively assess the potential for photon detection, we employ three distinct reconstruction modes to isolate the contribution of a specific detection strategy (the reconstruction sequence and configuration are described in the appendix C):

1. Standard WCD-only reconstruction (WCD-mode)
   This mode also follows the standard WCD reconstruction of horizontal air showers (details

in section 5.1). This mode establishes the baseline performance of the Auger Observatory, revealing the inherent limitations of the WCD array for identifying inclined, muon-poor photon-induced air showers.

2. Standard Radio-only reconstruction (RD-mode)
   We employ this mode to quantify the standalone potential of a radio self-trigger. In this mode, stations trigger independently using a fixed SNR threshold, emulating realistic self-trigger conditions (see section 5.2). The reconstruction uses the true Monte Carlo arrival direction to evaluate antenna gain patterns, providing an optimistic performance baseline by eliminating directional reconstruction uncertainties. While this approach yields an upper bound on achievable performance, it directly answers a key question: how effectively can radio detect photons without WCD assistance?.

3. Hybrid reconstruction (Hybrid-mode)
   An application [110] allows stations to trigger on both WCD and RD signals, whichever arrives first and satisfies the T1 condition where radio T1 automatically is promoted to T2. The T3 events are then formed by requiring a coincidence of at least three stations of a single detector type (three WCD or three RD), independent of their initial trigger source. This mode tests the efficiency of a combined RD–WCD trigger and evaluates its impact on event selection and reconstruction.
   For radio self-triggering, a fixed threshold (T_{thr}) is applied at the ADC level. Each channel (EW, NS) is offset-corrected, and the peak value and timing are extracted. A station is marked as triggered if any sample in the signal window exceeds T_{thr} in either channel.
   For this analysis, three threshold values are considered: T_{thr} of 15,20,30 ADC counts.

By evaluating the reconstruction efficiency, defined as:

$$\eta(E,\theta) = \frac{N_{\mathrm{Reconstructed}}(E,\theta)}{N_{\mathrm{Simulations}}(E,\theta)}, \tag{4.2}$$

where $N_{\mathrm{Reconstructed}}$ and $N_{\mathrm{Simulations}}$ represent the number of successfully reconstructed and total simulated events, respectively, at a given energy and zenith angle. This reconstruction efficiency is conceptually different from Equation 4.1, which evaluates the potential trigger efficiency using pure simulation output without realistic reconstruction.

Figure 4.7 shows the reconstruction efficiency η obtained with the full reconstruction chain for photon-induced showers. The left panel shows the efficiency as a function of primary energy, while the right panel shows the efficiency as a function of zenith angle. The three colored curves correspond to Hybrid-mode reconstructions using fixed radio thresholds of 15, 20, and 30 ADC counts. Based on the ADC-to-field-amplitude relation shown in Figure 4.2, these thresholds correspond approximately to 65, 90, and 130 μV/m, respectively.
Since the WCD trigger and reconstruction settings are kept fixed between the three Hybrid-mode configurations, the differences between the colored curves mainly reflect the effect of

changing the radio threshold. Lowering the radio threshold generally increases the Hybrid reconstruction efficiency, especially at lower energies and in regions where events are close to the radio trigger threshold.

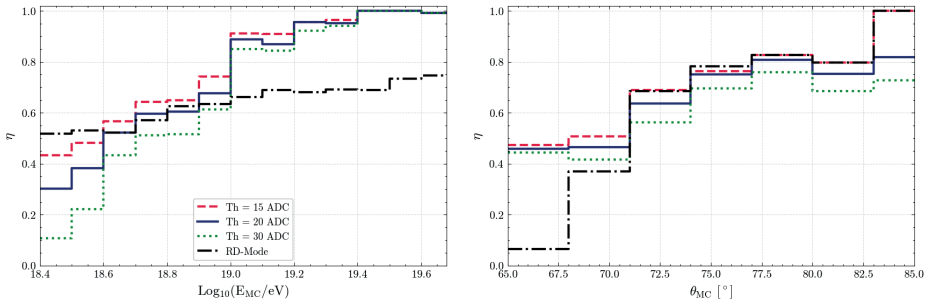


Figure 4.7: Reconstruction efficiency η for photon-induced air showers as a function of primary energy (left) and zenith angle (right). The colored curves show the Hybrid-mode reconstruction for fixed radio thresholds of 15, 20, and 30 ADC counts, corresponding approximately to 65, 90, and 130 μV/m based on Figure 4.2. The black dash-dotted curve shows the RD-mode reconstruction using the SNR-based self-trigger. The WCD trigger and reconstruction settings are kept fixed for the Hybrid-mode curves; therefore, their differences mainly reflect the impact of the radio threshold.

This trend should not be interpreted as a strict inverse proportionality between efficiency and threshold. Rather, increasing the threshold reduces the number of usable radio stations and therefore increases the probability that an event fails the radio reconstruction or quality requirements. At higher energies, the radio signals become stronger and more stations exceed the threshold. The Hybrid-mode efficiency therefore approaches saturation and becomes less sensitive to the precise threshold choice. A similar behaviour is observed at large zenith angles, where the radio footprint is larger and the station multiplicity is higher. This explains why lowering the threshold is most important in the lower-energy and lower-multiplicity regime. The RD-mode curve behaves differently because it uses an SNR-based self-trigger rather than a fixed ADC threshold. The SNR criterion normalizes the peak signal by the station-specific RMS noise level. Therefore, the RD-mode efficiency depends not only on the signal amplitude, but also on the local noise conditions. At low energies or in less favourable geometries, this can reduce the probability of obtaining enough usable radio stations for a successful reconstruction. At higher energies and larger zenith angles, the stronger radio signal and larger footprint improve the RD-mode efficiency, although it does not directly follow the fixed-threshold Hybrid-mode curves. The remaining inefficiency in the Hybrid-mode reconstruction can be attributed to events that pass the trigger-level requirement but fail the full reconstruction or quality selection. This can occur when the radio signal is too weak, when too few stations pass the quality criteria, or when the spatial distribution of triggered stations is nearly linear, making the direction reconstruction unreliable. Such cases could be improved by lowering the effective radio threshold or by developing methods that recover subthreshold signals.

Impact of Noise on Reconstruction

Since the radio threshold is highly sensitive to simulation noise, real noise traces from the early deployed RD stations were recorded during March 15, 2022. Data from both channels were collected manually by visiting each operational at that time station with a USB data storage device, recording noise traces without external triggering. Each station provided data from a different 24-hour period. The noise level in the array, defined here as the absolute maximum ADC count observed in the absence of a signal, fluctuates over time and across stations. To safely superimpose background noise onto Monte Carlo (MC) simulations without contaminating air shower signals with transient radio frequency interference (RFI) or high-amplitude anthropogenic spikes, the traces were filtered to isolate stationary ambient noise [111]. For each station in the event, a random noise trace was selected for both channels. Note that this selection was applied during the air shower reconstruction.

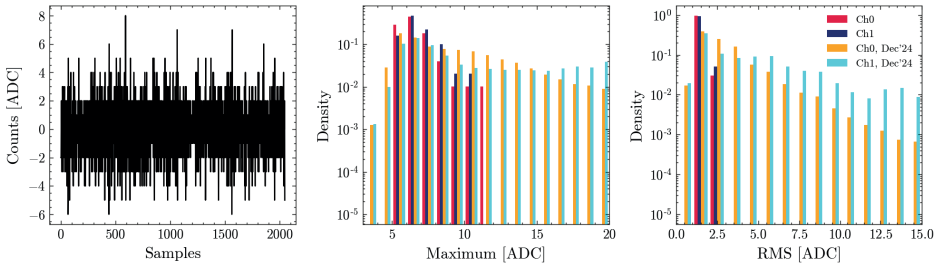


Figure 4.8: Radio-noise characterization for simulation studies. Left: example noise waveform from the noise library injected into the reconstruction. Middle and right: distributions of the absolute maximum amplitude and RMS for the injected-noise library (Ch0, Ch1), compared to a representative sample of measured noise data from December 2024.

Figure 4.8 characterizes the noise properties. The example trace (left panel) demonstrates typical noise behavior, while the maximum amplitude distribution (center panel) shows that peak noise levels predominantly fall between 5–6 ADC counts across both channels, as the NS channel more noisier. This range is well below the threshold used. The RMS distribution (right panel) confirms stable noise characteristics, with both channels centered at approximately 1.4–2 ADC counts[1], indicating consistent low-level background noise. The December 2024 sample is included as an external cross-check against current operating conditions; a time-dependent noise treatment using continuous monitoring (and its impact on the high-noise tail). The noise level is peaking well with real data; however, it does not follow the long tail, which can significantly affect the reconstruction performance, especially for weak signals. The Weekly RMS distributions[2] show median values for EW and NS channels within 1–2 ADC.

A comprehensive noise library is currently under development using continuous measurements from all operational stations, which will provide more representative noise conditions across

[1]In data recorded from October 2024 to April 2025, 90% of EW channels and 82% of NS channels have absolute maxima below 15 ADC.

[2]Weekly RMS distribution: https://hef.ru.nl/mord/data/Pages/[3]WeeklyRMSPage.html

diverse environmental and seasonal conditions; it was unavailable for the analyses presented here.

The Reconstruction Performance

The geometries of the reconstructed simulation set are presented in Figure 4.9. WCD reconstruction shows strong energy and zenith dependence, reconstructing only photon-induced air showers below 10^{19} eV due to their hadron-like muon signatures that mimic proton-induced air showers. Coverage extends to larger zenith angles above $10^{19.2}$, eV, but the detection efficiency remains significantly higher for lower zenith angle air showers. RD-mode provides broader coverage but exhibits a deficit below $10^{18.8}$ eV at 65° - 70° zenith angles. This gap weakens at higher energies but persists. The hybrid mode achieves better coverage by combining both detectors' strengths. It maintains reasonable efficiency down to $10^{18.5}$ eV, significantly reduces RD's low-energy deficit, and eliminates WCD's near vertical bias.
Figure 4.10 quantifies the zenith angle dependence across energy bins. WCD efficiency scales steeply with energy, reaching nearly full efficiency for less inclined air showers at the highest energies. RD-mode maintains consistently high efficiency (> 95%) at large zenith angles but remains limited at 65° - 70°. Hybrid mode achieves intermediate performance, but in the 70°–75° bin below 10^{19} eV it is slightly less efficient than RD-mode because it uses an SNR ≥ 7.22 as radio trigger (and the true MC direction for antenna-response corrections), whereas the Hybrid chain applies a fixed ADC threshold and relies on full geometry reconstruction. Near threshold this reduces antenna multiplicity and increases the fraction of events that fail reconstruction or quality requirements, while the WCD contribution for photons in this bin remains small.

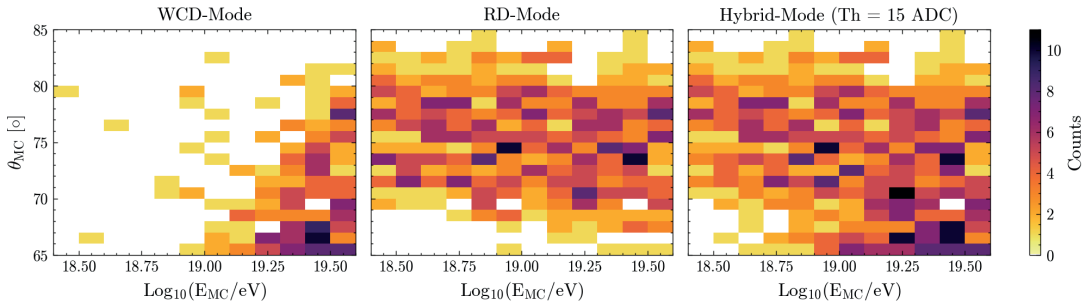


Figure 4.9: Reconstructed event geometries for the detection modes: WCD-mode, RD-mode, and Hybrid-mode with a threshold of 15 ADC.

The reconstruction efficiency values presented in Table 4.1 are derived from Monte Carlo simulations of photon-induced air showers across multiple energy and zenith angle bins.
The trends are consistent with Figs. 4.5 and 4.6. At low energies ($10^{18.4}$–$10^{18.8}$ eV), the RD triggers efficiently while the WCD rarely does, so the hybrid efficiency closely follows RD. With increasing energy, the WCD efficiency rises sharply, and the hybrid efficiency correspondingly improves. In zenith, the RD efficiency increases from ∼65° and saturates at 100% near 70°–75°, whereas the WCD efficiency decreases at the largest angles; consequently, the RD dominates at

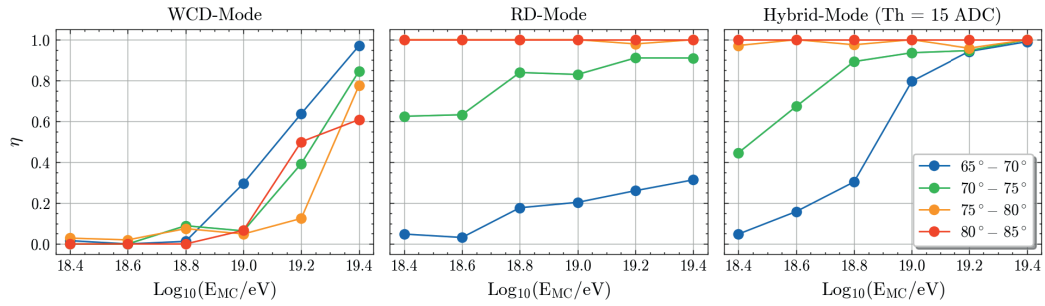


Figure 4.10: The angular dependence of trigger efficiency for three different modes: WCD-mode, RD-mode, and Hybrid-mode (with a threshold of 15 ADC).

high zenith and the hybrid configuration inherits this advantage. For each bin combination, the efficiency η is given in Equation 4.2. Statistical uncertainties are estimated assuming binomial statistics. For bins with $0 < k < N_{\text{sim}}$ we use $\delta\eta = \sqrt{\eta(1-\eta)/N_{\text{sim}}}$, treating each simulated event as an independent trial. For boundary cases ($k = 0$ or $k = N_{\text{sim}}$), where this expression yields zero, we quote a one-sided uncertainty of $1/N_{\text{sim}}$ (upper for $k = 0$, lower for $k = N_{\text{sim}}$), reflecting the finite Monte Carlo statistics.

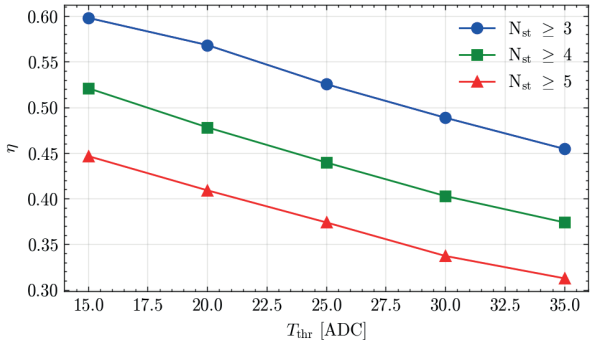


Figure 4.11: RD-mode trigger efficiency versus radio trigger threshold for different station multiplicity requirements: $N_{\text{st}} \geq 3$ (blue circles), $N_{\text{st}} \geq 4$ (green squares), and $N_{\text{st}} \geq 5$ (red triangles). The efficiency is evaluated for reconstructed photon simulations in the range $18.4 \leq \log_{10}(E_{\text{RD}}/\text{eV}) \leq 20.7$ and $65° \leq \theta_{\text{RD}} \leq 85°$.

Figure 4.11 shows the impact of T_{thr} on RD-mode trigger efficiency for different station multiplicity requirements across all energies and zenith angles. Efficiency decreases monotonically with increasing threshold for all cases, in this threshold range.

Over the explored range, the efficiency decreases monotonically with T_{thr}: raising the threshold from 15 to 35 ADC lowers η by ∼0.14–0.15 in absolute terms for $N_{\text{st}} \geq 3, 4, 5$, indicating a broadly uniform threshold penalty across multiplicity selections.

To highlight the combined potential of running a radio trigger and applying a missing-signal analysis for photon searches: events reconstructed with a weak or absent WCD signal but a strong radio signal are characteristic of neutral primaries—i.e., photons. Hence, building on the

Energy bin	Zenith bin	N_{sim}	η_{WCD}	η_{RD}	η_{Hyb}
$10^{18.4}-10^{18.6}$	$65-70$	60	0.017 ± 0.017	0.050 ± 0.028	0.050 ± 0.028
	$70-75$	56	$0.000^{+0.018}_{-0.000}$	0.625 ± 0.065	0.446 ± 0.066
	$75-80$	35	0.029 ± 0.028	$1.000^{+0.000}_{-0.077}$	0.971 ± 0.028
	$80-85$	13	$0.000^{+0.018}_{-0.000}$	$1.000^{+0.000}_{-0.077}$	$1.000^{+0.000}_{-0.077}$
$10^{18.6}-10^{18.8}$	$65-70$	63	$0.000^{+0.018}_{-0.000}$	0.032 ± 0.022	0.159 ± 0.046
	$70-75$	49	$0.000^{+0.018}_{-0.000}$	0.633 ± 0.069	0.673 ± 0.067
	$75-80$	50	0.020 ± 0.020	$1.000^{+0.000}_{-0.077}$	$1.000^{+0.000}_{-0.077}$
	$80-85$	12	$0.000^{+0.018}_{-0.000}$	$1.000^{+0.000}_{-0.077}$	$1.000^{+0.000}_{-0.077}$
$10^{18.8}-10^{19.0}$	$65-70$	73	0.014 ± 0.014	0.192 ± 0.046	0.329 ± 0.055
	$70-75$	56	0.089 ± 0.038	0.839 ± 0.049	0.893 ± 0.041
	$75-80$	40	0.075 ± 0.042	$1.000^{+0.000}_{-0.077}$	0.975 ± 0.025
	$80-85$	15	$0.000^{+0.018}_{-0.000}$	$1.000^{+0.000}_{-0.077}$	$1.000^{+0.000}_{-0.077}$
$10^{19.0}-10^{19.2}$	$65-70$	54	0.296 ± 0.062	0.204 ± 0.055	0.796 ± 0.055
	$70-75$	47	0.064 ± 0.036	0.830 ± 0.055	0.936 ± 0.036
	$75-80$	41	0.049 ± 0.034	$1.000^{+0.000}_{-0.077}$	$1.000^{+0.000}_{-0.077}$
	$80-85$	15	0.067 ± 0.064	$1.000^{+0.000}_{-0.077}$	$1.000^{+0.000}_{-0.077}$
$10^{19.2}-10^{19.4}$	$65-70$	68	0.647 ± 0.058	0.265 ± 0.054	0.956 ± 0.025
	$70-75$	56	0.393 ± 0.065	0.911 ± 0.038	0.946 ± 0.030
	$75-80$	47	0.128 ± 0.049	$1.000^{+0.000}_{-0.077}$	0.979 ± 0.021

Table 4.1: Reconstruction efficiency for photon-induced air showers across energy and zenith angle bins.

reconstruction efficiency analysis, Figure 4.12 quantifies the practical benefit of radio addition by showing the fraction of events missed by WCD-mode reconstruction.

This directly addresses the limitations identified earlier: WCD's sharp energy threshold at 10^{19} eV and strong zenith bias. For photons, radio detection recovers nearly 95% of events below the WCD threshold at $10^{18.5}$ eV, directly compensating for WCD's poor low-energy performance. Even above 10^{19} eV where WCD efficiency improves, radio still recovers 40-70% additional photon events, particularly at large zenith angles where WCD remains inefficient. For protons, the main background in this context, the enhancement is minimal above 10^{19} eV (consistent with WCD's full efficiency) and remains below 30% across all conditions.

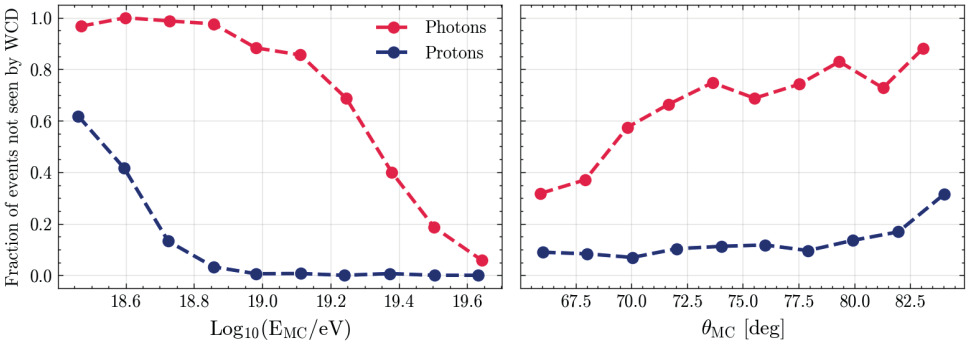


Figure 4.12: Fraction of events not detected by WCD-only mode, showing the relative gain from radio detection for photons (red) and protons (blue) as functions of energy (left) and zenith angle (right).

4.3 Exploring lower energies

The strong performance of RD in photon detection (Figure 4.10) motivates exploring its potential at lower energies. Since no photon simulation library was available below 4 EeV at the moment, we approximate this regime by rescaling the radio signal amplitude with primary energy. The underlying assumption is that the radio emission strength is proportional to the number of e^+/e^- in the air shower, which itself scales with the primary energy. This relation is illustrated in Figure 4.13, where the slopes of the electron number and the maximum recorded electric field are nearly identical. Thus, we rescale the electric field amplitude accordingly to estimate trigger efficiency below the simulated energy range.

This rescaling modifies only the overall radio amplitude according to the assumed energy dependence. It does not re-simulate the footprint shape, air-shower geometry, $X_{\max}$ distribution, antenna-gain variations with azimuth or zenith, or the noise environment. The results should therefore be regarded as an indication of performance rather than a full prediction under realistic fluctuations and systematics. This approach provides a rough assessment of RD array

limitations in the absence of dedicated low-energy simulations, but it bypasses the full reconstruction process since only the original simulations were rescaled. Unlike section 4.2, where station-level reconstruction is performed with ADC thresholds, the rescaling study uses a simpler conservative criterion: at least three stations with $E_{\max} \geq 300\,\mu\mathrm{V/m}$. This isolates the effect of amplitude scaling and enables clean energy comparisons, but the results are not directly comparable to the ADC-based performance shown in Figure 4.2.

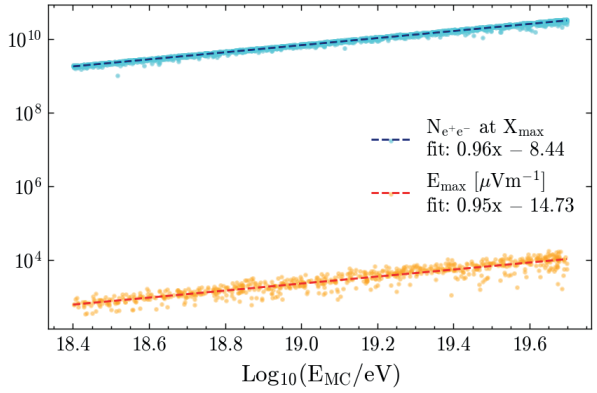


Figure 4.13: The correlation of energy versus the maximum electric field per event and the number of e^+/e^- at air-shower maximum. The x-axis shows $\log_{10}(E_{\mathrm{MC}}/\mathrm{eV})$, while the y-axis is logarithmic and shared by both datasets ($E_{\max}$ and $N_{e^+e^-}$). Dashed lines indicate power-law trends obtained from linear fits performed in $\log_{10}$ space.

For the efficiency estimate in Figure 4.14, an event is considered triggered if at least three stations record $E_{\max} \geq 300\,\mu\mathrm{V/m}$. The results indicate that photon-induced air showers are scarcely detectable below 1 EeV across all zeniths, while inclined geometries benefit from their larger footprint and achieve higher efficiencies than the 65°–70° range.

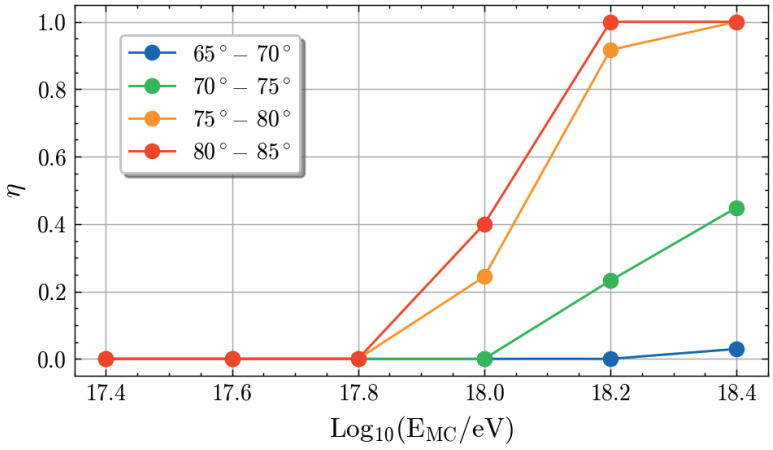


Figure 4.14: The trigger efficiency from rescaled events and its zenith dependence.

Given the RD sensitivity, antenna gain, and RFI conditions, the estimation likely represents an optimistic projection of full efficiency. Nonetheless, data recorded by the RD have demonstrated the capability to capture such low-energy events, making directional and transient searches feasible.

An example of an event falling in the SD750 area is shown in Figure 4.15, recorded in October 2024 with eight RD stations. The event yielded consistent energy reconstructions between RD and WCD, with $\log_{10}(E_{\mathrm{WCD}}/\mathrm{eV}) = 17.91 \pm 0.19$ ($\log_{10}(E_{\mathrm{RD}}/\mathrm{eV}) = 17.89 \pm 0.18$) and $\theta_{\mathrm{WCD}} = 71.6° \pm 1.4$ ($\theta_{\mathrm{RD}} = 73.1° \pm 0.12$).

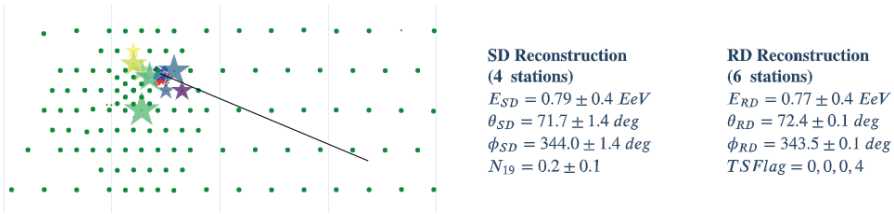


Figure 4.15: EventID: 76911083. An example of a low-energy event reconstructed by RD and triggered by WCD.

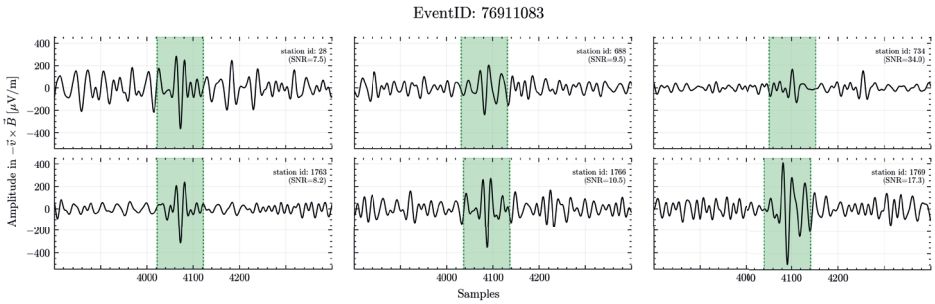


Figure 4.16: The reconstructed waveforms in $-\vec{v}\times\vec{B}$ direction with SNR ≥ 7.22 (see section 5.2) for the event in Figure 4.15.

The reconstructed waveforms shown in Figure 4.16 show a detectable signal in the signal window indicated by the green band.

4.4 Aperture for Ultra-High-Energy Photons

In the following, we define the effective aperture $A_{\mathrm{eff}}(E)$ of the WCD and RD arrays for measuring UHE photons in the zenith range 65°−85°. Throughout, an event is considered detected when it passes the trigger criteria described in section 4.2. The effective aperture is determined in energy bins, reflecting the detector's capability to register events at a given energy integrated over the specified zenith range. It condenses into a single quantity both the geometric acceptance of the array and the probability that an extensive air shower of energy E arriving from direction (θ, φ) at time t satisfies the trigger (quality criteria not applied in this context).

We quote A_{eff} in km^{2}sr. Both the WCD and the RD arrays operate with a duty cycle close to 100%, and their active area remains stable; under these conditions $A_{\text{eff}}(E;t)$ is effectively time independent.

Assuming an azimuth-independent efficiency (or, equivalently, using its azimuth average), the effective aperture (derived in Appendix B) can be expressed in terms of the zenith-averaged detection efficiency, $\bar{\eta}(E)$, as:

$$A_{\text{eff}}(E) \;=\; \pi\, A\, \bar{\eta}(E)\left(\cos^2\theta_1 - \cos^2\theta_2\right), \tag{4.3}$$

where $A = 3000\,\text{km}^2$ is the geometric area of the surface array (flat-array approximation). Here, $\bar{\eta}(E)$ is the detection efficiency from the full simulation–reconstruction chain, appropriately weighted and integrated over the zenith range $\theta_1 \le \theta \le \theta_2$ (with $\theta_1 = 65°$ and $\theta_2 = 85°$ in this analysis), rather than assuming the efficiency is independent of the zenith angle (see section 4.2).

If the full angular dependence (e.g., antenna gain/polarization or specific zenith effects) must be retained, the general form is:

$$A_{\text{eff}}(E) \;=\; A\int_{\theta_1}^{\theta_2}\int_0^{2\pi} \eta(E,\theta,\varphi)\,\cos\theta\,\sin\theta\,\mathrm{d}\varphi\,\mathrm{d}\theta,$$

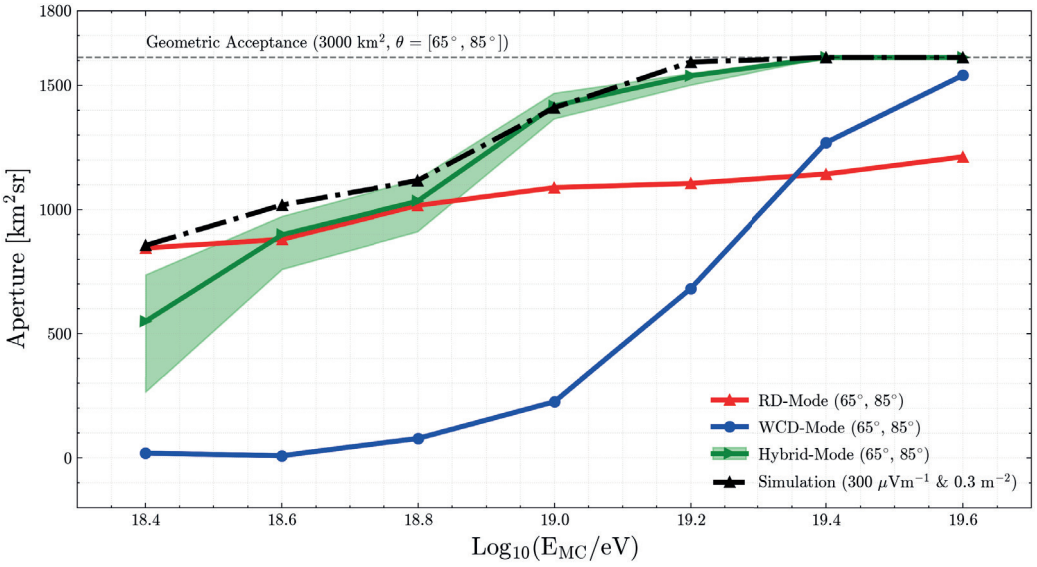


Figure 4.17: Photon aperture of different Auger trigger systems in the zenith angle range from 65° to 85°. The red line shows the RD self-trigger mode, the blue line the standard WCD trigger, and the green line the WCD+RD Hybrid trigger at 20 ADC. The shaded green band illustrates the threshold variation between 15 and 30 ADC. The black line gives a detector-agnostic baseline from simulation using a hybrid approach without reconstruction. The grey dashed line marks the geometric acceptance for $\theta = 65°-85°$ and $A = 3000$ km^2.

The equation (4.3) neglects three secondary effects—minor duty–cycle losses, residual azimuthal asymmetries in the radio signal, and time–dependent vetoes used to suppress radio–frequency interference. Because the radio emission and detector response depend on arrival direction, while our limited simulation set samples only discrete azimuths, the present analysis cannot yet provide a fully azimuth–integrated efficiency and aperture; where relevant, we therefore quote azimuth–averaged values.

Figure 4.17 summarizes the photon aperture in the investigated trigger modes and compares them with recent Auger photon searches. The WCD-only aperture remains below $500\,\mathrm{km^2\,sr}$ up to $10^{19}\,\mathrm{eV}$ and nearly reaches the geometric aperture at $\sim$1500 $\mathrm{km^2\,sr}$ at $10^{19.6}\,\mathrm{eV}$. The RD-mode rises constantly from $10^{18.4}\,\mathrm{eV}$ onwards and exceeds $1000\,\mathrm{km^2\,sr}$ already at $10^{19}\,\mathrm{eV}$. The RD-mode starts off at low energy much larger than WCD it saturates below the Hybrid case, underlining the benefit of combining WCD and RD triggers. Note that the RD-mode curve is an optimistic reference (SNR-based trigger and true MC direction for antenna response), whereas Hybrid-mode uses fixed ADC thresholds and the full reconstruction chain; at low energies, where WCD contributes little, this can lead to $\eta_{\mathrm{Hybrid}}(E) < \eta_{\mathrm{RD}}(E)$ and hence a smaller aperture.

Overlaying the detector-agnostic estimation from section 4.1 helps gauge the effects of trigger and reconstruction. Below $\sim 10^{19.2}\,\mathrm{eV}$, the Hybrid aperture falls short of the geometric acceptance by $\sim 7\%$, consistent with thresholding and noise. At higher energies, the Hybrid converges to the geometric acceptance and the black simulation curve, indicating near-optimal efficiency for photon-induced air showers. The robustness of the Hybrid-mode is further supported by the green shaded band: varying the trigger threshold between 15 and 30 ADC changes the aperture only modestly.

Overall, the radio-enhanced detection substantially extends photon sensitivity at large zenith angles. The gain in accessible sky coverage is estimated at least $\sim$21% [112], and at the highest energies the apertures approach the geometric limit, showing that these configurations operate close to their maximum collection efficiency.

This chapter has established, through simulation, the strong potential of a radio self-trigger and hybrid detection to significantly enhance the sensitivity of the Auger Observatory to inclined UHE photons. However, realizing this potential requires addressing two key challenges: first, the implementation of an actual working self-trigger in the field, and second, the development of analysis techniques to maintain purity against background. As with current Auger analyses, we expect typical efficiency losses of $\sim 50\%$ for photon identification after applying necessary background rejection cuts. The development of these crucial quality criteria and identification techniques will be addressed in the next chapters.

Chapter 5

Data Reconstruction and Validation

This chapter details the end-to-end process of reconstructing inclined cosmic-ray air showers using both WCDs and RD subsystems at the Pierre Auger Observatory. After a successful T3 trigger from the WCD array, radio data is retrieved from the local stations, inherently linking all radio measurements to established particle detector events. As a result, reconstruction of radio events always follows and builds upon WCD event identification. While the WCD and RD subsystems are conceptually independent, practical integration of WCD-derived geometry parameters (such as air-shower direction and core position) serves to enhance the robustness and convergence speed of radio reconstruction algorithms.
This part presents the methodologies, algorithms, and dependencies involved in reconstructing air showers in both detector systems, laying the groundwork for subsequent data validation and selection.

5.1 Water Cherenkov Detector Event Reconstruction

Since the particle footprint of inclined air showers ($\theta \geq 60°$) is dominated by muons, measured signals in the WCD stations are fitted to reference muon–density maps to estimate the air-shower size and core position; the arrival direction is obtained independently from a timing fit to the air-shower front [89].

WCD Arrival Direction Reconstruction

The arrival direction is determined from the relative start times of the signals. The start times recorded at triggered WCD stations are fitted to a spherical (parabolic) air-shower front [113]. For a station at position $\mathbf{r}_i$, the arrival time is parameterized as

$$t_i \;=\; t_0 \;+\; \frac{1}{c}\left[\,\hat{\mathbf{n}}\cdot(\mathbf{r}_i-\mathbf{r}_c) \;+\; \frac{|\mathbf{r}_i-\mathbf{r}_c|^2}{2R}\,\right], \tag{5.1}$$

where t_0 is a reference time, $\mathbf{r}_c$ is the air-shower-core position, $\hat{\mathbf{n}}$ is the air-shower-axis unit vector (defined by θ, ϕ), R is the radius of curvature, and c is the speed of light. Station timestamps are weighted using the Auger time-variance model, which parametrizes each station's timing

uncertainty as a function of the measured signal and the signal rise-time (the interval for the integrated signal to go from 10% to 50% of its final value). This weighting improves the geometry reconstruction [114].
This procedure yields an angular resolution better than 0.5° (improving to ∼0.35° at the highest energies). Typically, the resolution improves with zenith angle up to ∼ 80° primarily due to the longer lever arm provided by the larger shower footprint on the ground, in addition to the larger number of triggered stations.

WCD Energy

With the arrival direction determined, the expected muon density at ground level is modeled as a two-dimensional function of zenith and azimuth angles. The muon density $\rho_\mu(x, y; \theta, \phi)$ at ground position (x, y) is expressed as:

$$\rho_\mu(x, y; \theta, \phi) = N_{19} \cdot \rho_{\text{ref}}(x, y; \theta, \phi), \tag{5.2}$$

where N_{19} is the air-shower size parameter relative to a 10 EeV (10^{19} eV) proton-induced air shower, and ρ_{ref} is a reference muon density derived from atmospheric cascade simulations using hadronic interaction models such as QGSJET or EPOS.
A likelihood fit of the observed station signals to these expected muon densities $\rho_\mu(\mathbf{r})$ is then performed, defined by:

$$\mathcal{L} = \prod_i P(S_i|\rho_\mu(x_i, y_i; \theta, \phi, N_{19}, \vec{r}_c)), \tag{5.3}$$

where S_i is the signal at station i, and P is the probability of observing S_i given the modeled density. The fit accounts for the residual electromagnetic fraction separated using a parameterized ratio $f_{\text{em}} = S_{\text{em}}/S_\mu$ from Monte Carlo simulations, where S_{em} and S_μ are the electromagnetic and muonic signal components, respectively, ensuring accurate isolation of the muonic signal and geomagnetic/attenuation effects encoded in the reference maps. Including both triggered and non-triggered stations enhances the fit's robustness by constraining low-signal regions. This procedure yields both the N_{19} and the best-fit core position $\vec{r}_c$.
Then, N_{19} is converted to the primary cosmic-ray energy E_{WCD} using a calibration derived from hybrid events detected by both WCD and FD. The latter provides a direct, calorimetric energy measurement, and the relationship is modeled as a power law:

$$E_{\text{WCD}} = 10^{19}\,\text{eV} \left(\frac{N_{19}}{A}\right)^{1/B}. \tag{5.4}$$

where $A = 1.723 \pm 0.023$ and $B = 0.984 \pm 0.020$.
The WCD energy resolution is $(19.3 \pm 1.4)\%$, inferred from the distribution of the ratio $E_{\text{WCD}}/E_{\text{FD}}$ for the hybrid set with the FD energy resolution fixed to 7.6%; this resolution is attributed to the combined effect of air-shower-to-air-shower fluctuations and the reconstruction measurement uncertainty. In addition, the calibration contributes an energy-dependent uncertainty from the propagation of the N_{19}–E_{FD} fit parameters, ranging from 1.5% at 10^{19} eV

to 5% at 10^{20} eV. A further $\sim 2\%$ accounts for the different zenith-angle distributions between the hybrid calibration set and the full inclined data set. The FD energy scale carries an overall 14% systematic fluorescence yield (3.6%), absolute detector calibration (9.9%), and air-shower reconstruction (6.5–5.6%) which propagates directly to E_{WCD}. Adding these in quadrature gives a total systematic uncertainty on E_{WCD} of $\sim 14\%$ at 10^{19} eV (nearly energy-independent below 10^{19} eV) rising to $\sim 17\%$ at 10^{20} eV.

5.2 Radio Detector Event Reconstruction

Each WCD station is equipped with a radio antenna. If the antenna was operating and produced a readout for the event, the station is included in the radio-processing pipeline [111]. The data-reconstruction pipeline below applies to real measurements, accounting for actual detector conditions and environmental noise. The simulation-reconstruction pipeline is constructed to closely mirror the data methodology. For simulation, at key stages (such as the antenna-response correction), Gaussian-distributed deviations are injected to realistically imitate hardware and environmental effects.

1. **ADC-to-voltage conversion:** The reconstructed radio processing starts from ADC traces. These traces are converted into voltage traces using the detector calibration. At this stage, the signal is still a channel-level voltage trace and not yet a reconstructed electric-field quantity.

2. **Antenna-response correction:** Apply the forward and inverse antenna response functions to correct the detector's frequency response: a forward response (20–90 MHz by default) followed by an inverse response (30–80 MHz by default).

3. **RFI cleaning:** Analyze each channel's spectrum using a discrete-time Fourier transform. Iteratively identify and subtract dominant narrow-band interference in the time domain until the target noise-reduction improvement is reached or until a maximum of five iterations (per channel) is performed.

4. **Windowing and resampling:** Reduce edge effects on time-series traces by applying a Hann window covering 1% of the relative width at both ends of each trace. Increase time resolution either by a fixed upsampling factor (default 4×) or by interpolating to a target sampling frequency. Remove DC offsets.

5. **Station-level quantities:** Convert raw channel data to station-level vectors using the air-shower geometry and a model for the antenna. For inclined air showers, corrections are not performed to the event-by-event value of the air-shower maximum X_{max}, which is in that case not directly known. Rather, all corrections are made with respect to a reference depth, determined from a parameterisation built using the hadronic models.

6. **Pulse reconstruction:** For each station, identify a distinct pulse and determine the noise level using predefined windows. Estimate the SNR and extract key observables (peak

amplitude and timing). Perform signal integration (fluence) and propagate uncertainties from signal amplitude, antenna calibration, and noise.

Stations meeting the SNR criterion of ≥ 7.22 in the standard application [95] are retained for the subsequent air-shower reconstructions. The quadratic SNR is defined as:

$$\mathrm{SNR} = \left(\frac{A_{\mathrm{peak}}}{\mathrm{RMS}_{\mathrm{noise}}}\right)^2 \tag{5.5}$$

where A_{peak} is the absolute peak amplitude in the signal window and $\mathrm{RMS}_{\mathrm{noise}}$ is the root-mean-square of the noise measured in a dedicated noise window. The SNR can be computed at different processing levels [115], using either the magnitude of the vector field $|\mathbf{E}|$, individual components (E_x, E_y, E_z), or projections onto specific plane e.g., the horizontal (x, y)/ground plane, or the physically motivated ($\vec{v}\times\vec{B}$, $\vec{v}\times(\vec{v}\times\vec{B})$) planes. By default, the SNR is computed from the Hilbert-envelope amplitude after baseline removal on the absolute magnitude of the field vector.

RD Arrival Direction

In reconstructing the arrival direction from RD data, the direction previously determined by WCD serves as an initial guess (more details in 6.3). Since WCDs detect secondary particles at ground level while RDs measure coherent radio emission from the air shower's electromagnetic component, comparing their reconstructed directions offers valuable cross-validation to detect anomalies or complex air-shower geometries.

The implementation proceeds by initializing the RD timing fit with the WCD axis–core pair $(\mathbf{a}_{\mathrm{WCD}}, \mathbf{r}_{\mathrm{WCD}})$. In this timing-only stage, the RD core is fixed to the WCD core, $\mathbf{r}_{\mathrm{RD}} \equiv \mathbf{r}_{\mathrm{WCD}}$, while the RD axis is determined by the fit. A local right-handed Cartesian system $\mathcal{C}_{\mathrm{loc}}$ is centered at $\mathbf{r}_{\mathrm{RD}}$;$\mathbf{a}_{\mathrm{WCD}}$ provides the initial angles (θ_0, ϕ_0). To seed the source position, a simplified slant-depth model converts a typical depth of $750\,\mathrm{g\,cm^{-2}}$ into a geometric distance R_0 along $\mathbf{a}_{\mathrm{WCD}}$, yielding

$$\mathbf{s}_0 = R_0 \begin{pmatrix} \sin\theta_0 \cos\phi_0 \\ \sin\theta_0 \sin\phi_0 \\ \cos\theta_0 \end{pmatrix}_{\mathcal{C}_{\mathrm{loc}}}$$

For each radio station k, the inputs are the displacement vector $\mathbf{x}_k$ from $\mathbf{r}_{\mathrm{RD}}$, the measured arrival time t_k^{meas}, and the timing uncertainty σ_k. The timing fit follows the spherical-wave formulation [116]: assuming propagation at the speed of light [30], the expected arrival time is

$$t_k^{\mathrm{exp}}(R, \theta, \phi) = \frac{\left\|\mathbf{x}_k - \mathbf{s}(R, \theta, \phi)\right\|}{c}, \qquad \mathbf{s}(R, \theta, \phi) = R \begin{pmatrix} \sin\theta \cos\phi \\ \sin\theta \sin\phi \\ \cos\theta \end{pmatrix}_{\mathcal{C}_{\mathrm{loc}}}$$

Since absolute trigger times are arbitrary, the fit uses time differences by subtracting the mean measured time $\bar{t}^{\text{meas}}$ and mean expected time $\bar{t}^{\text{exp}}$. The fit goodness is

$$\chi^2(R,\theta,\phi) = \sum_k \left[\frac{(t_k^{\text{meas}} - \bar{t}^{\text{meas}}) - (t_k^{\text{exp}} - \bar{t}^{\text{exp}})}{\sigma_k} \right]^2$$

A numerical minimizer adjusts R, θ, and ϕ to locate a (local) minimum of χ^2. The best-fit (θ, ϕ) define the unit RD direction $\mathbf{a}_{\text{RD}}$; uncertainties are propagated to Cartesian components.

RD Core and Energy

The air shower electromagnetic energy is reconstructed directly from the LDF of the dominant geomagnetic radio emission [117]. After correcting for early-late projection effects and subtracting a small, parameterized charge-excess contribution, the remaining $\vec{v} \times \vec{B}$-polarized energy fluence becomes rotationally symmetric. This fluence is fitted with the analytic form:

$$f_{\text{GS}}(r) = f_0 \left[\exp\left(- \left[\frac{r - r_0^{\text{fit}}}{\sigma} \right]^{p(r)} \right) + \frac{a_{\text{rel}}}{1 + \exp\left(s \cdot \left[\frac{r}{r_0^{\text{fit}}} - r_{02} \right] \right)} \right] \tag{5.6}$$

A Gaussian[1] plus sigmoid function with six shape parameters $(r_0^{\text{fit}}, \sigma, p(r), a_{\text{rel}}, s, r_{02})$. Simulations show that each shape parameter depends solely on the geometrical distance $d_{\max}$ between the air-shower core and the air-shower maximum. Thus, the radio-emission footprint is described by two observables—the geomagnetic radiation energy E_{geo} and the geometrical distance $d_{\max}$, together with two core coordinates specifying the impact position at ground. The air-shower core is obtained by fitting the radio footprint: the radio–symmetry centre is determined with a geomagnetic LDF fit in which the core coordinates are free.

To account for environmental and geometric effects, E_{geo} is corrected for the geomagnetic angle α and the air density $\rho_{\max}$ at the depth of air-shower maximum $X_{\max}$, yielding the scale variable:

$$S_{\text{geo}} = \frac{E_{\text{geo}}}{\sin^2 \alpha} \left[1 - p_0 + p_0 \exp\left(p_1 (\rho_{\max} - \langle \rho \rangle) \right) \right]^{-2},$$

where α denotes the geomagnetic angle between the air-shower axis and the geomagnetic field, $\sin \alpha = \| \hat{\vec{a}} \times \hat{\vec{B}} \|$. The constants are from a combined fit [117]:

$$p_0 = 0.5045, \quad p_1 = -2.7083\,\text{kg}^{-1}\text{m}^3, \quad \langle \rho \rangle = 0.300\,\text{kg m}^{-3}.$$

The corrected scale variable S_{geo} follows a power-law relationship with the electromagnetic energy E_{em}:

$$E_{\text{em}} = 10\,\text{EeV} \left(\frac{S_{\text{geo}}}{S_{19}} \right)^{1/\gamma}, \quad S_{19} = 3.1461\,\text{GeV}, \quad \gamma = 1.9997.$$

For simulated air showers with energies between $10^{18.4}$ and 10^{20} eV and zenith angles between 65° and 85°, this method achieves an intrinsic energy resolution better than 5% (using at least

[1] when $p(r) = 2$

5 stations) with negligible bias.
Geometrically, a minimum number of stations is required for a constrained reconstruction, with additional stations improving robustness and reducing uncertainties. In the signal-model reconstruction of the paper, the footprint (LDF) is fitted with four free parameters ($E_{\text{geo}}, d_{\max}, x_{\text{core}}, y_{\text{core}}$), i.e. two physical LDF parameters plus two (core/symmetry-center) position coordinates.
Consequently, the effective number of fit parameters depends on whether the core is fitted from radio data or fixed from external information:

- **Three-station configurations (reduced fit, only if the core is fixed externally)**: With three stations, the full four-parameter fit is underconstrained. However, if an independent core estimate ($x_{\text{core}}, y_{\text{core}}$) is available (e.g. from another detector), one can fit only the two LDF parameters ($E_{\text{geo}}, d_{\max}$). With three data points and two parameters, this yields one degree of freedom[2], allowing an assessment of fit quality (e.g. via χ^2).

- **Four-station configurations**: The reconstruction fits two LDF parameters and two core-position parameters ($E_{\text{geo}}, d_{\max}, x_{\text{core}}, y_{\text{core}}$) using four data points. This results in zero degrees of freedom, making the fit exactly determined but providing no internal statistical goodness-of-fit test.

- **Five or more station configurations**: The same four parameters are fitted. With more than four stations, the degrees of freedom become positive (e.g. one for five stations, two for six), enabling statistical evaluation of the fit quality.

5.2.1 $\overline{\text{Off}}\underline{\text{line}}$ Antenna Configuration

We validate the antenna model implementation in the reconstruction pipeline by verifying that the reconstructed polarisation is dominated by the geomagnetic component and follows the $\hat{\vec{v}} \times \hat{\vec{B}}$ expectation. We adopt a right–handed geographic (ENU) frame with unit axes $\hat{\boldsymbol{e}}_x$ (East), $\hat{\boldsymbol{e}}_y$ (North), $\hat{\boldsymbol{e}}_z$ (Up); the azimuth ϕ is measured in the horizontal plane from $+x$ (East) toward $+y$ (North).
We are concerned here with the polarisation direction only. Let $\hat{\vec{v}}$ and $\hat{\vec{B}}$ be unit vectors along the air-shower propagation and the geomagnetic field, respectively. The geomagnetic polarisation is

$$\hat{\vec{e}}_{\text{geo}} \;\propto\; \hat{\vec{v}} \times \hat{\vec{B}}$$

with magnitude $\|\hat{\vec{v}} \times \hat{\vec{B}}\| = \sin\alpha$ (where α is the angle between $\boldsymbol{v}$ and $\boldsymbol{B}$). For a down-going air shower of zenith θ and azimuth ϕ,

$$\hat{\vec{v}} = (-\sin\theta\cos\phi,\; -\sin\theta\sin\phi,\; -\cos\theta)$$

[2]The degrees of freedom (d.o.f.) in the LDF fit are defined as the number of radio stations contributing one independent measurement minus the number of free parameters in the fit.

and for a geomagnetic field with inclination I (taken negative when directed downward) and declination δ (east of geographic north),

$$\hat{\vec{B}} = (\cos I \, \sin\delta, \; \cos I \, \cos\delta, \; \sin I)$$

(At the Pierre Auger Observatory, approximate values are $I \approx -35°$, $\delta \approx -5°$.) Writing $\hat{\vec{e}}_{\text{geo}} = (e_x, e_y, e_z)$, the fractional contribution seen by the Cartesian channels is

$$\eta_i = \frac{e_i^2}{e_x^2 + e_y^2 + e_z^2}, \qquad i \in \{x, y, z\},$$

which depends only on the direction $\hat{\vec{v}} \times \hat{\vec{B}}$ (not on the absolute field amplitude).

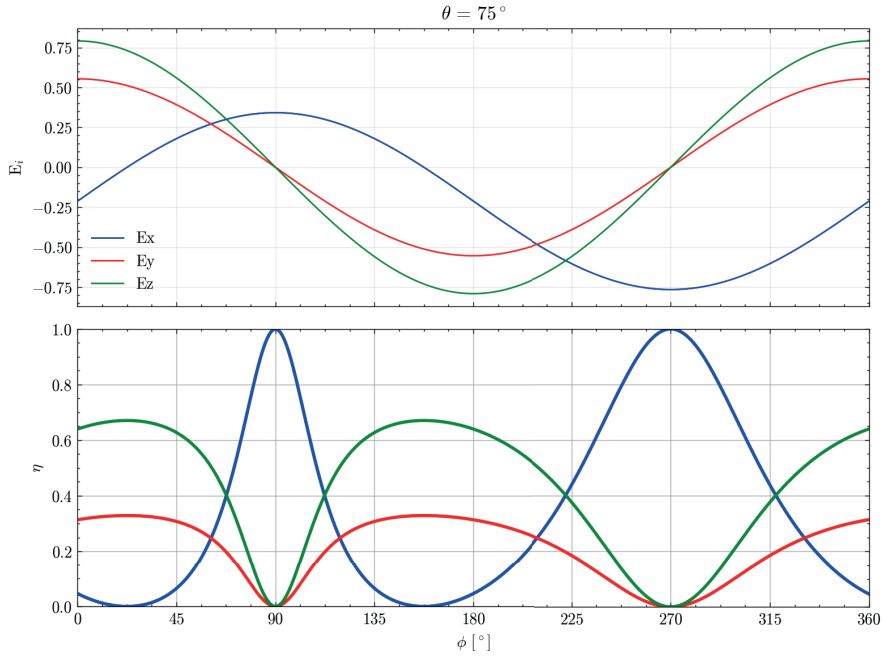


Figure 5.1: Theoretical electric field components for a cosmic ray air shower with zenith angle $\theta = 75°$, versus azimuthal angle ϕ. The top panel plots E_x, E_y, and E_z magnitudes (arbitrary units) from $-\vec{v}\times\vec{B}$. The bottom panel shows each component's fractional contribution η to the total squared electric field. The Earth's magnetic field has $I = -35°$ and $\delta = -5°$, typical of the Auger Observatory's location.

Figure 5.1 illustrates these contributions for $\theta = 75°$, showing the electric field components' magnitudes and their fractional contributions as functions of ϕ. The minima of E_x and E_y are displaced due to the magnetic field's declination.

We then assess a reconstructed simulation set with added real noise, serving as a consistency check. Starting from raw ADC traces, the process converts these to voltage, applies the antenna pattern to derive the 3D electric field, and transforms it to the air-shower plane, where geomagnetic emission should dominate.

Relative amplitudes are calculated at key stages voltage, 3D electric field, and air-shower plane by summing absolute traces:

$$A_{\text{comb}}(t) = \sum_{i=1}^{N} |X_i(t)|$$

where $X_i(t)$ represents the signal amplitude (voltage, electric field component, or air-shower-plane component) of the i-th channel or component at time t. The peak time is identified as:

$$t_{\text{peak}} = \arg\max_t A_{\text{comb}}(t)$$

The relative amplitude of each component at t_{peak} is:

$$\text{Rel}(X_i) = \frac{|X_i(t_{\text{peak}})|}{A_{\text{comb}}(t_{\text{peak}})}$$

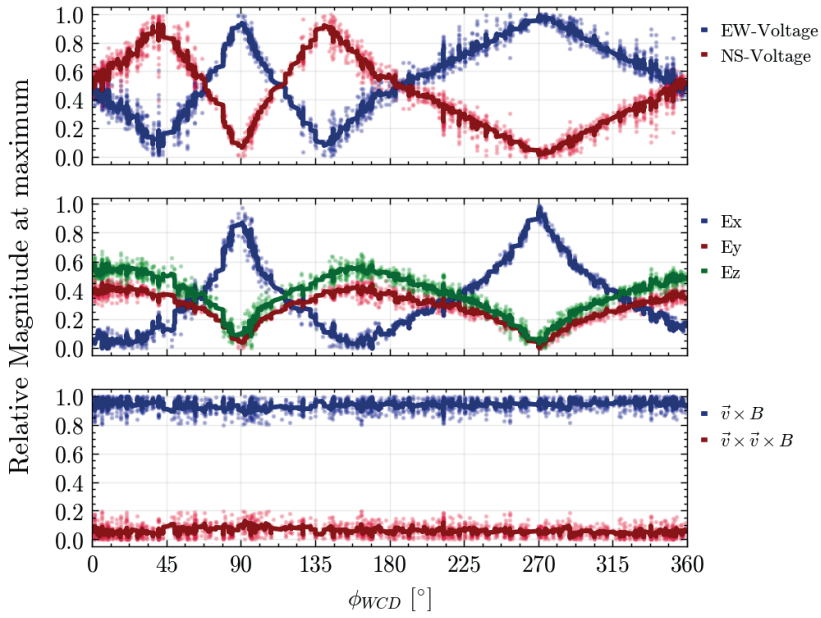


Figure 5.2: Relative magnitudes of radio emission components from reconstructed simulations versus air-shower azimuthal direction ϕ_{WCD}. Top: EW and NS voltage components with sinusoidal azimuthal dependence. Middle: Electric field components E_x, E_y, and E_z, showing dominant horizontal polarization. Bottom: Geomagnetic $(-\vec{v}\times\vec{B})$ versus $(\vec{v}\times(\vec{v}\times\vec{B}))$ contributions. Solid lines are running averages.

Figure 5.2 displays these results versus the air shower's azimuthal arrival direction ϕ_{WCD}. The top panel shows sinusoidal variations in EW and NS voltage components, the middle panel shows that E_x and E_y dominate for most azimuthal directions, though E_z becomes significant near $\phi \approx 0°$ and 180°. The bottom panel shows the geomagnetic polarization fraction. Its level is set by polarization leakage driven by the noise: more noise causes more cross-polarized

contamination, which lowers the geomagnetic fraction and increases its scatter. In the absence of noise, the reconstructed geomagnetic fraction is expected to be close to unity, but not exactly equal to one. This is because the measured radio signal is not purely geomagnetic: the charge-excess emission provides an additional, radially polarized component whose projection onto the analysed polarization basis slightly reduces the fraction attributed to the geomagnetic component. Early data analysis, however, revealed discrepancies (Figure 5.3).
While voltage patterns aligned with expectations, the 3D electric field components were incorrect, distorting the air-shower-plane transformation. This was traced to a software error in the **Offline** antenna configuration.

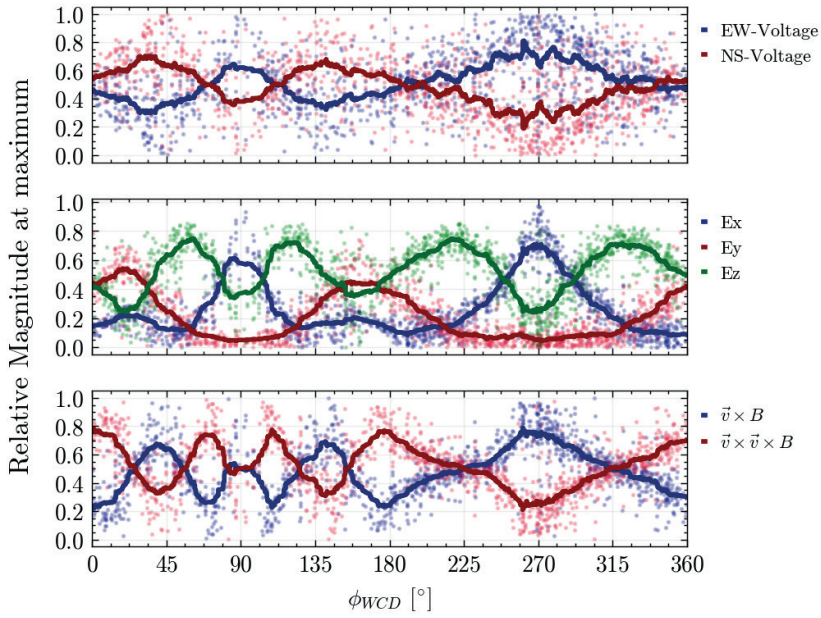


Figure 5.3: Relative magnitudes of radio emission components from reconstructed radio data (January–September 2024) with original loop azimuths $\phi_{\mathrm{EW}} = 0°$, $\phi_{\mathrm{NS}} = 90°$ versus air-shower azimuthal direction ϕ_{WCD}. Top: EW and NS voltage components with sinusoidal azimuthal dependence. Middle: Electric field components E_x, E_y, and E_z, showing dominant horizontal polarization. Bottom: Geomagnetic ($-\vec{v}\times\vec{B}$) versus ($\vec{v}\times(\vec{v}\times\vec{B})$) contributions, with geomagnetic emission dominating. Solid lines are running averages.

The RD SALLA antenna's Vector Effective Length (VEL) was modeled using `4NEC2` [118], with separate EW and NS loop responses stored in the **Offline** database [119]. In the Auger coordinate system (x, y, z = East, North, Up), with θ from $+z$ and ϕ from $+x$, the loop normals were initially:

$$(\theta, \phi)_{\mathrm{EW}} = (90°, 0°), \quad (\theta, \phi)_{\mathrm{NS}} = (90°, 90°).$$

However, the reconstructed E_x polarity was reversed, indicating the EW loop's effective orientation was 180° opposite the nominal setting, possibly due to a polarity inversion (e.g., wiring or model mismatch). Adjusting $\phi_{\mathrm{EW}} = 180°$ while retaining $\phi_{\mathrm{NS}} = 90°$ corrected the VEL

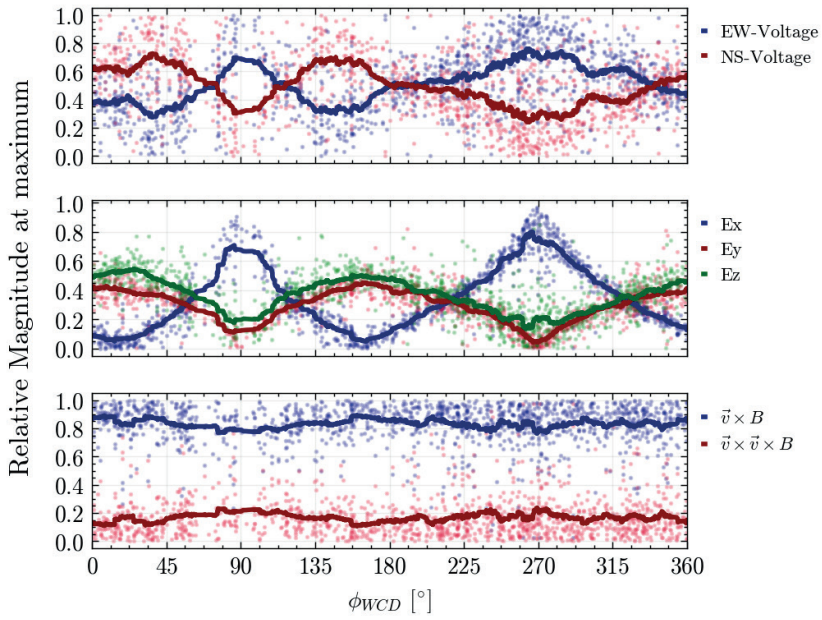


Figure 5.4: Relative magnitudes of radio emission components from reconstructed radio data (January–September 2024) after setting $\phi_{\mathrm{EW}} = 180°$, with $\phi_{\mathrm{NS}} = 90°$ versus air-shower azimuthal direction ϕ_{WCD}. Top: EW and NS voltage components with sinusoidal azimuthal dependence. Middle: Electric field components E_x, E_y, and E_z, showing dominant horizontal polarization. Bottom: Geomagnetic ($-\vec{v}\times\vec{B}$) versus ($\vec{v}\times(\vec{v}\times\vec{B})$) contributions, with geomagnetic emission dominating. Solid lines are running averages.

vector's direction, restoring the expected geomagnetic pattern (Figure 5.4) without altering gain or bandwidth.

Post-correction, Figure 5.4 aligns with expectations, though noise weakens peaks and dips. Figure 5.5 exemplifies the fix's impact: the perpendicular component ($\vec{v} \times (\vec{v} \times \vec{B})$) is suppressed from a dominant 1500 μV/m to 200 μV/m, while the geomagnetic component ($-\vec{v}\times\vec{B}$) reverses polarity (from -1500 to +1500 μV/m) with similar amplitude.

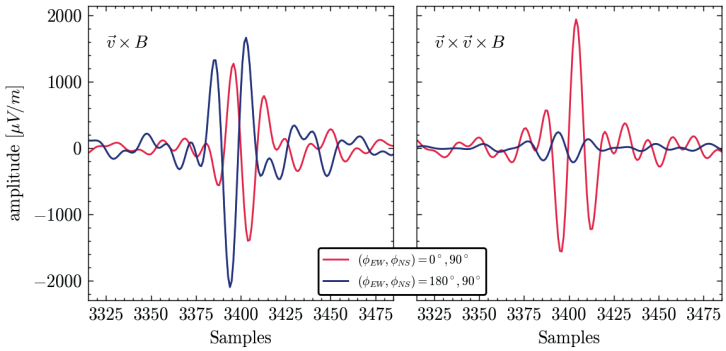


Figure 5.5: Waveforms for $\vec{v}\times\vec{B}$ (left) and $\vec{v}\times(\vec{v}\times\vec{B})$ (right) at station ID 922, event ID 75376213, showing pre- and post-correction effects.

The accurate physical and processing antenna orientation is thus critical for reliable radio reconstruction, minimizing systematic errors in cosmic ray measurements. However, data availability constraints, such as deadtime during acquisition, can further impact reconstruction quality, as detailed next.

5.3 Radio Detector Performance and Efficiency

In this section, we characterize the performance and efficiency of the radio detector using the full available data set. Typical physics analyses employ more stringent quality cuts (section 6.2), so they use a narrower, higher-purity event sample than shown here. The RD stations were installed incrementally; the deployment of the RD array was completed in late November 2024, after which the array entered a notably stable operating state, as shown in Figure 5.6. Three distinct phases are evident: from June (EA) through October 2023 (Super Hexagon), the baseline event rate remained relatively constant ($\sim$ 20–30 events per month). Beginning in November 2023, coinciding with the deployment of the digitizers, the monthly event rate increased steeply, scaling approximately 40-fold to reach $\sim$ 1100 events per month by November 2024. Following full deployment completion, the event rate plateaued at approximately 1100–1200 reconstructed events per month. A temporary 13% reduction was observed in December 2024 due to an operational disturbance in the WCD array, which impacted RD performance. Full recovery was achieved by January 2025. Throughout this period, events satisfying the stricter $N_{\mathrm{RD}} \geq 5$ criterion comprised roughly 40–45% of the total reconstructed events, demonstrating consistent trigger efficiency across the stable operating phase.

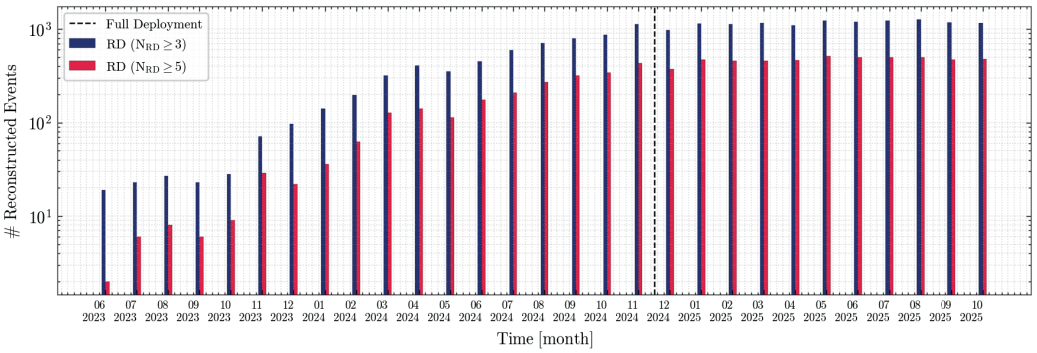


Figure 5.6: Monthly RD reconstruction event rate from June 2023 to October 2025. Blue: $N_{\mathrm{RD}} \geq 3$; Red: $N_{\mathrm{RD}} \geq 5$. The dashed line marks full deployment completion in November 2024.

Figure 5.7 shows all reconstructed[3] events by WCD and RD within the period when the array was fully deployed up to October 31, 2025.

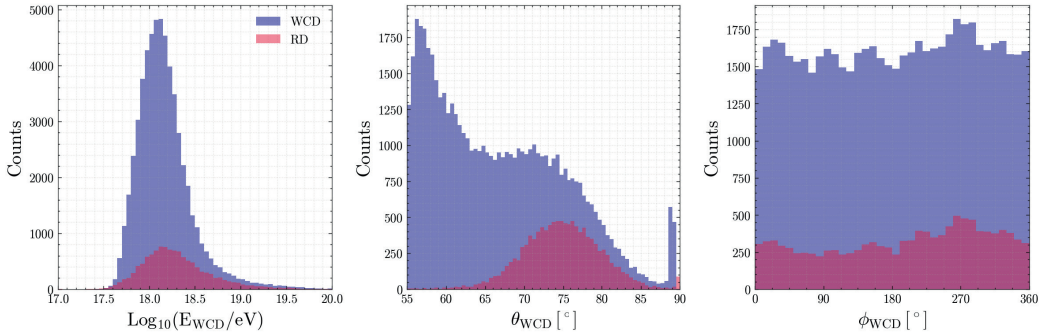


Figure 5.7: WCD (blue) vs. RD (red) event distributions in energy, zenith, and azimuth during the period from December 1, 2024, to October 31, 2025.

The RD is skewed to higher energies and shows a longer high-energy tail, reflecting the radio detection threshold, the growth of the radio footprint with energy, and WCD trigger dependency. Focusing on zenith, the WCD distribution decreases smoothly from 55° to 90°, RD events cluster strongly at very inclined geometries between 70° to 80° where radio signals are stronger and the footprint is largest; the small WCD spike near 90° are miss-reconstructed events

The WCD azimuth distribution is nearly uniform, as expected for an isotropic exposure. In contrast, the RD shows a pronounced east–west modulation—enhanced around ($\phi \approx 240° - 300°$), consistent with geomagnetic polarization.

[3]For the RD, we require a successful radio reconstruction based on the geomagnetic LDF fit with at least three radio stations providing usable waveforms above the analysis SNR threshold. For the WCD, we require a standard SD inclined air air-shower reconstruction with at least three triggered stations where T4/T5 triggers are applied.

We define the RD reconstruction efficiency as the fraction of WCD-reconstructed events that is also successfully reconstructed by RD ($N_{\mathrm{RD}} \geq 3$ with SNR≥ 7.22). Because the RD readout is gated by the WCD T3, the quoted efficiencies are conditional on WCD triggering. As shown in Figure 5.8(left), the efficiency rises steeply with both energy and zenith, exceeding $\sim 90\%$ for $\log_{10}(E_{\mathrm{WCD}}/\mathrm{eV}) \gtrsim 18.4$ at $\theta \gtrsim 70^{\circ}$ and approaching unity once the WCD is fully efficient ($\log_{10} E \gtrsim 18.6$).

At smaller zenith or lower energies, many events fall below the SNR/multiplicity threshold, and the ratio decreases. Figure 5.8(right) shows the corresponding increase in the mean station multiplicity, $\langle N_{\mathrm{RD}} \rangle$, which tracks the efficiency plateau: higher energy and more inclined geometries yield larger footprints and thus larger N_{RD} (color scale capped at 9).

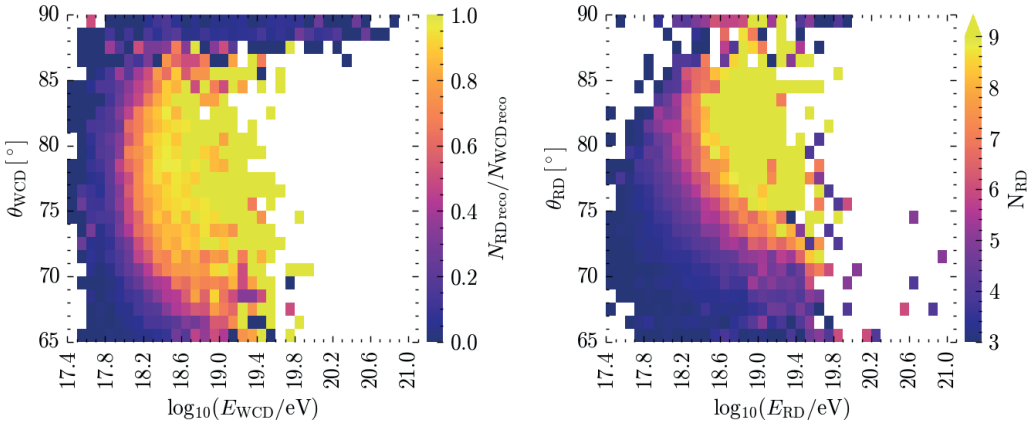


Figure 5.8: Left: RD-to-WCD reconstruction ratio, $N_{\mathrm{RD\,reco}}/N_{\mathrm{WCD\,reco}}$, as a function of $\log_{10}(E_{\mathrm{WCD}}/\mathrm{eV})$ and θ_{WCD}. Right: Mean number of RD stations, $\langle N_{\mathrm{RD}} \rangle$, versus $\log_{10}(E_{\mathrm{RD}}/\mathrm{eV})$ and θ_{RD} (color scale capped at 9).

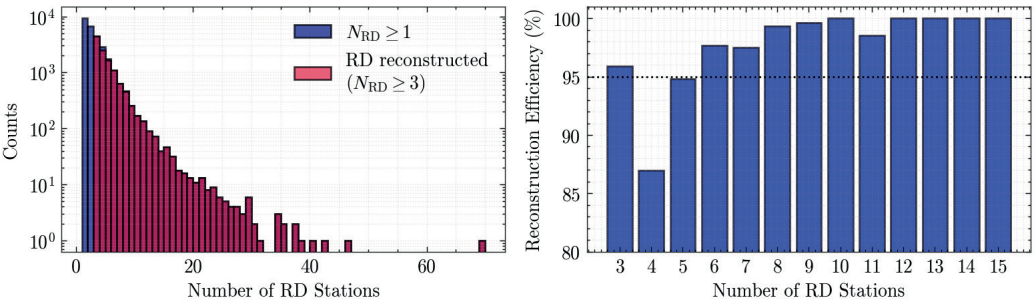


Figure 5.9: Multiplicity and reconstruction efficiency of the RD. Left: distribution of the number of RD stations with signal per WCD event (blue) and for RD-reconstructed events requiring $N_{\mathrm{RD}} \geq 3$ (red). Right: RD reconstruction efficiency versus the minimum station multiplicity, where $\geq 95\%$ is indicated by the horizontal black dotted line.

In the following, we are going to estimate the efficiency of the RD reconstruction, which gives an idea when the timing and signal fits (section 5.2) fail. We define the RD full-Yield re-

gion as the phase space where the RD reconstruction efficiency—conditional on the WCD T3 and our SNR/multiplicity cuts is $\geq 95\%$. In this regime, all air showers are measured and fully reconstructed after quality cuts. The left panel of Figure 5.9 shows a steeply falling RD multiplicity: most events have only a few stations above threshold (with the reconstruction requirement $N_{\mathrm{RD}} \geq 3$), with $\sim 60\%$ of radio air showers having $N_{\mathrm{RD}} \leq 4$; high-multiplicity footprints ($N_{\mathrm{RD}} \gtrsim 10$) are rare ($\sim 5\%$). A small subset of events satisfies the multiplicity minimum, yet the fit does not converge at $N_{\mathrm{RD}} = 4$.

The right panel quantifies the missed events as the efficiency increases monotonically from $\sim 95\%$ at $N_{\mathrm{RD}} = 3$ and reaches (near-)unity for $N_{\mathrm{RD}} \geq 7$, with a local dip at $N_{\mathrm{RD}} = 4$.

To assess the consistency of the core reconstruction between detectors, we compare WCD and RD core positions. As shown in Figure 5.10, the spatial distributions overlap closely, indicating that both techniques sample the same exposure region.

The histogram of core-distance differences between the WCD and RD shows a distribution peaked at $\mathcal{O}(10^2\,\mathrm{m})$, with most events exhibiting differences at the hundred-meter scale. The distribution is asymmetric with a long tail extending to kilometre scales, attributable to limited station multiplicity, edge-of-array geometries, or reduced timing/SNR situations where model choices and quality cuts influence the fit. The bulk concentration near $10^2\,\mathrm{m}$ indicates typical agreement between the two reconstruction methods, while the sparse tail represents a small fraction (1%) of events with larger discrepancies consistent with statistical fluctuations and reconstruction-model differences rather than a systematic shift.

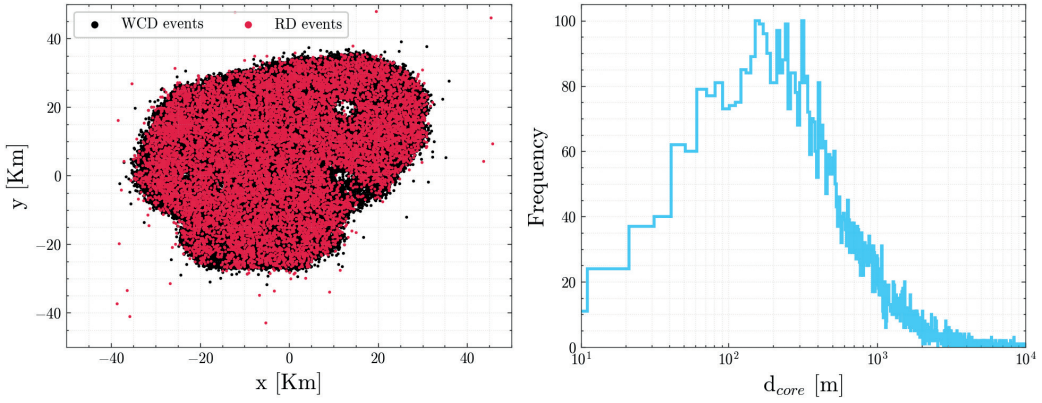


Figure 5.10: Comparison of air-shower-core reconstruction between WCD and RD. Left: Core locations over the array for WCD (black) and RD (red) events. The range has been adjusted to focus on the cores on the array; however, there is a small set of outliers (1%) Right: Histogram of the core–distance (d_{core}) difference (WCD−RD).

Figure 5.11 shows event-by-event differences between the WCD and RD reconstructions. The angular residuals are tightly peaked around zero with negligible bias: azimuth shows mean $\approx 0.00^\circ$ and spread $\sigma \approx 0.71^\circ$, while zenith has mean $\approx -0.03^\circ$ and $\sigma \approx 0.63^\circ$. In contrast, the energy difference distribution is asymmetric and shifted to negative values, with $\langle E_{\mathrm{WCD}} - E_{\mathrm{RD}} \rangle \approx -0.05$ EeV (median 0.17 EeV, $\sigma \approx 1.5$ EeV), indicating a systematic off-

set where WCD energies tend to be higher than RD. The long positive tail suggests a subset of events with larger discrepancies, plausibly linked to limited station multiplicity or extreme zeniths. Outliers are identified by masking events whose energy, azimuth, or zenith differences fall outside strict bounds: only values within $[-10, 10]$ EeV for energy and $[-5, 5]^\circ$ for angles are considered valid. Events outside these ranges are excluded from final statistics and visualizations, ensuring the analysis focuses on the central bulk of the distribution. These outliers account for approximately 2% for energy, 4% for azimuth, and 3% for zenith, while the majority of events demonstrate sub-degree directional agreement between the two detectors.

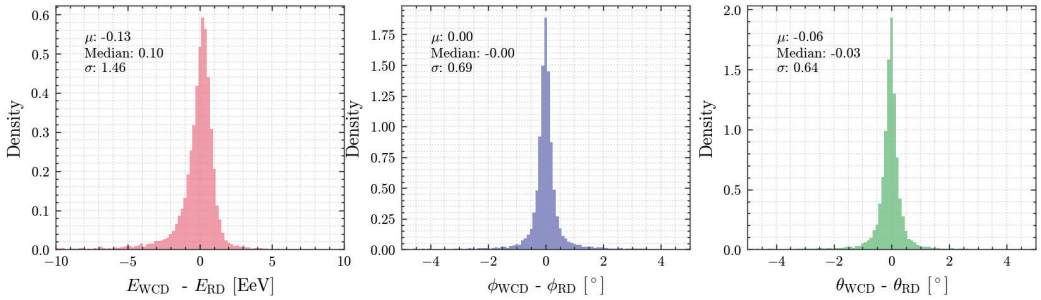


Figure 5.11: Event-by-event differences between WCD and RD reconstructions. Left: energy difference $E_{\mathrm{WCD}} - E_{\mathrm{RD}}$ (EeV). Middle/Right: azimuth and zenith differences ($\phi_{\mathrm{WCD}} - \phi_{\mathrm{RD}}$, $\theta_{\mathrm{WCD}} - \theta_{\mathrm{RD}}$).

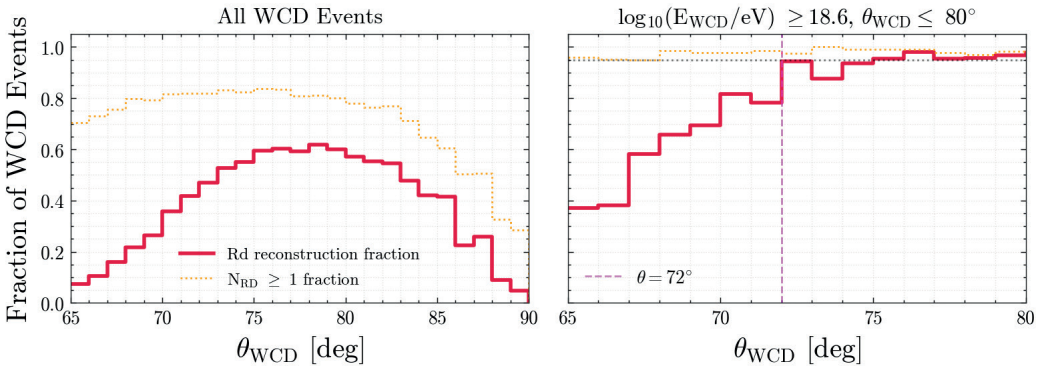


Figure 5.12: fraction of WCD events as a function of zenith angle. Left: Fraction of all WCD events with at least one RD station providing a proper signal (orange dotted line), and the fraction of events successfully reconstructed by the RD (red solid line). Right: Same as the left panel, but for events with energy $\log_{10}(E_{\mathrm{WCD}}/\mathrm{eV}) \geq 18.6$ and zenith angle $\theta_{\mathrm{WCD}} \leq 80°$. The magenta vertical line marks the zenith angle ($\theta \approx 72°$) where the full efficiency regime begins, and the black horizontal line indicates the 95% efficiency threshold.

Now, we will explore how the RD efficiency evolves as a function of zenith angle and energy. Figure 5.12 illustrates the evolution of the RD reconstruction fraction across zenith angles for all WCD events (left) and for events with higher energy and lower zenith angles (right).

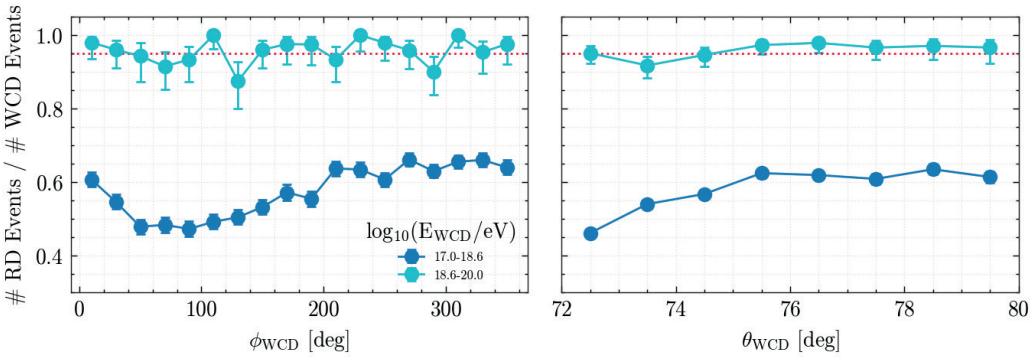


Figure 5.13: Fraction of reconstructed RD to reconstructed WCD events as a function of azimuth (ϕ_{WCD}) and zenith (θ_{WCD}) angle, for two energy ranges. The dotted red line marks the threshold at which full RD efficiency is reached. Error bars represent 1σ binomial confidence intervals.

It shows that 75% of all WCD air showers trigger at least one RD station with a proper signal. In the left panel, as the zenith angle increases, the RD reconstruction fraction initially rises, reaching a plateau at higher zeniths. A noticeable drop occurs beyond 80°, where the WCD detection efficiency starts to decrease due to geometric limitations, such as muon absorption. This WCD multiplicity bottleneck leads to a decrease in RD efficiency as it also limits RD station multiplicity, as the majority of the missed events have only 4 WCD stations. The right panel zooms into the high-energy and low-zenith region with a threshold of $\log_{10}(E_{WCD}/eV) \geq 18.6$ and $\theta_{WCD} \leq 80°$.

In this regime, the efficiency saturates, and we observe an increase in the fraction of reconstructed events, as indicated by the red line. The magenta vertical line indicates the zenith angle ($\theta \approx 72°$) where the full-efficiency geometry begins.

This high-performance region, which accounts for approximately 12% of the total number of events, represents a phase space where the RD reconstruction achieves at least 95% efficiency. The important thing is that above these cuts, we measure all air showers and therefore are unbiased. The trade-off is that this efficiency comes at the cost of a significant reduction in the number of events, as the stricter energy and zenith angle cuts naturally exclude a large fraction of the data. Figure 5.13 shows the fraction of reconstructed RD events to WCD events for two different energy ranges. The left panel indicates how the RD reconstruction fraction varies with azimuth. For $\log_{10}(E_{WCD}/eV) \geq 18.6$ and $\theta_{WCD} \geq 72°$, it demonstrates a largely uniform distribution fluctuating near the 95% threshold, whereas the low-energy range retains the strong east–west modulation also visible in Figure 5.7. The right panel shows the zenith angle dependence, where the reconstruction fraction increases steadily and consistently maintains full efficiency ($\geq 95\%$) for $\theta_{WCD} \gtrsim 75°$.

As we move to the next chapter, we focus on station and event quality, examining how station performance, noise reduction, and event validation impact the reliability of the reconstruction process, aiming to optimize data quality and minimize systematic errors.

Chapter 6

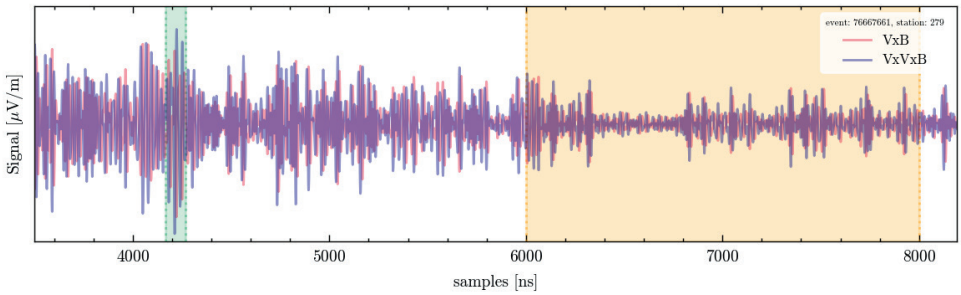
Data Quality Selection and Bias Assessment

Following event reconstruction, rigorous attention must be given to the quality of the analyzed dataset. In this chapter, we develop and apply a comprehensive quality selection strategy for the radio detector and assess its impact on the reconstruction of inclined air showers. We begin by revisiting the station-level selection, going beyond a simple SNR threshold to include additional, physics-motivated checks on signal presence, temporal morphology, and polarization in the air-shower plane. Using a pilot dataset, these criteria are tuned and benchmarked in terms of signal efficiency and background rejection, and their combined effect on the purity of the selected station sample is quantified. We then introduce a set of event-level requirements, based on the number and geometry of contributing radio stations, the quality of the lateral-distribution fit, the consistency between radio and WCD reconstructions, and the identification of periods affected by thunderstorms. Finally, we use the resulting high-quality hybrid data set to compare radio and water-Cherenkov reconstructions, to study the dependence of the radio performance on station multiplicity and geometry, and to derive the angular resolution achieved by the radio detector for very inclined air showers.

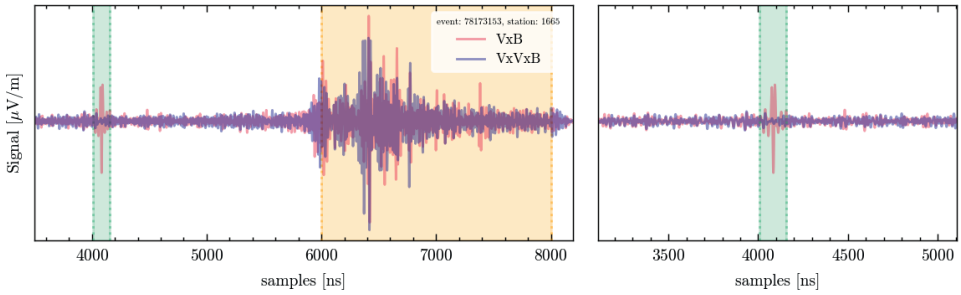
6.1 Signal-Quality Selection

Although selecting stations for radio event reconstruction based purely on a quadratic SNR threshold is fairly efficient (section 5.2), it introduces an excessive number of false positives, which consequently degrades the overall reconstruction quality. Aside from the threshold itself, a critical limitation of this method is its lack of contextual awareness within the signal window, rendering it incapable of distinguishing a genuine radio pulse from RFI spikes, transient bursts that may potentially originate from, for instance, a faulty LNA, or short, man-made transient bursts. Furthermore, the practice of defining the noise window as the final 2 microseconds of the trace is demonstrably suboptimal; any localized interference or elevated noise within this narrow segment can artificially depress the computed SNR, leading to the rejection of legitimate signals.

We adopt a new set of station selection criteria that prioritize quality over quantity. The guiding idea is to work with a high-purity subset of stations so that no single spurious trace can bias the lateral-distribution (fluence) fit or the spherical wavefront timing fit described in section 5.2. Quantitative benchmarks are derived from a pilot dataset recorded between October 2024 and February 2025 and processed with the collaboration standard reconstruction described in this chapter (see also appendix C). The sample comprises 6075 events and 62187 station traces (unique EvID–StID pairs). The radio-station multiplicity with pulse (SNR $\geq$ 7.22) spans $N_{\mathrm{RD}} \in [3, 62]$ with a median of 4 and a mean of 5; 2401 events (39.5%) have $N_{\mathrm{RD}} \geq 5$. Since the geomagnetic component is expected to dominate the radio emission, all station-level quality quantities are computed in the air-shower plane, i.e. on the $-\vec{v}\times\vec{B}$ and $\vec{v}\times(\vec{v}\times\vec{B})$ components obtained from the antenna traces. In air shower signals, the pulse is expected to concentrate in $-\vec{v}\times\vec{B}$, while man-made RFI generally shows no such geomagnetic polarization pattern, so this choice directly supports the polarization and waveform-based rejection of RFI.



(a) Trace that passes the standard SNR criterion.



(b) Trace that fails the standard SNR criterion.

Figure 6.1: Shower-plane signal components for two stations, evaluated with an SNR threshold of 7.22. Green shading indicates the signal window; orange shading indicates the noise window. (a) A trace (event: 76667661, station: 279) accepted (false positive). (b) A trace (event: 78173153, station: 1665) rejected (true negative), with a 1000-sample zoom on the right.

To illustrate typical failure modes of the standard SNR definition, we discuss two representative traces. In this context, we use the term "false positive" for stations that pass the standard SNR

cut even though the pulse in the signal window is clearly not an air-shower signal (e.g. because its shape matches noise bursts elsewhere in the trace), and "false negative" for stations that fail the SNR cut despite showing a clean, air-shower-like pulse. The trace shown in Figure 6.1(a) met the standard SNR condition, as the signal window contains a prominent pulse, while the noise window is relatively quiet. However, this signal cannot be a legitimate cosmic-ray signal because the shape of the large pulse in the signal window is similar to other significant noise pulses found elsewhere in the trace, including within the noise window. In contrast, Figure 6.1(b) did not meet the SNR condition due to the prominent noise in the noise window that boosted RMS to suppress the SNR; however, the signal window and its vicinity suggest a clean signal in the $\vec{v}\times\vec{B}$ direction with a fairly low $\vec{v}\times(\vec{v}\times\vec{B})$ component, which support a legitimate cosmic ray signal before the high simultaneous noise hits both trace components.
To minimize false positives while retaining high-quality stations, we apply three complementary, physics-motivated checks in the air-shower plane: verifying signal presence, enforcing the expected waveform shape, and confirming the expected polarization.

6.1.1 Dual-Window Linear SNR

We introduce a refined Signal-to-Noise ratio based on the linear peak amplitude (SNR′). In contrast to the standard, power-based SNR definition (section 5.2), We work directly with the non-enveloped time series of the two air-shower-plane components $\vec{v}\times\vec{B}$ and $\vec{v}\times(\vec{v}\times\vec{B})$. For each component, the signal amplitude $A_{\text{peak}}^{\vec{v}\times\vec{B}}$ and $A_{\text{peak}}^{\vec{v}\times(\vec{v}\times\vec{B})}$ is defined as the maximum absolute signal value within the signal window.
For each station, a fixed signal window is defined around the expected potential pulse time; in the present analysis, this corresponds to a width of 100 ns[1]. The noise RMS, on the other hand, is computed from the noise window; hence, taking absolute values would not change σ. The morphology of this new SNR definition and the placement of the signal and noise windows are illustrated in Figure 6.2.
A dual-window approach is employed to accurately estimate background noise, with measurements taken from a pre-signal window (900 ns) and a post-signal window (700 ns), each separated from the signal window by a 100 ns buffer. This buffer accounts for the finite duration of the cosmic-ray air shower's broadband radio pulse (lasting tens of nanoseconds), preventing residual signal components from inflating the RMS noise estimate and underestimating the SNR. Using two noise windows immediately before and after the signal increases the number of samples in the noise estimate while keeping it local in time to the pulse. This is particularly useful to study how fluctuations in the vicinity of the signal region affect the SNR. Transient interference is also diluted: for example, if a short noise burst contaminates only the pre-signal window, its effect is reduced when combined with the post-signal window. However, this approach assumes approximately stationary noise across both windows, which may not always be the case.

[1]The signal search window depends on the operated firmware; in the current version, it lives within the range from 3900 to 4300 ns

Strong interference in one window could skew the combined RMS, affecting the SNR, as shown in cases of inter-window noise or concentrated noise spans (see Figure 6.1).

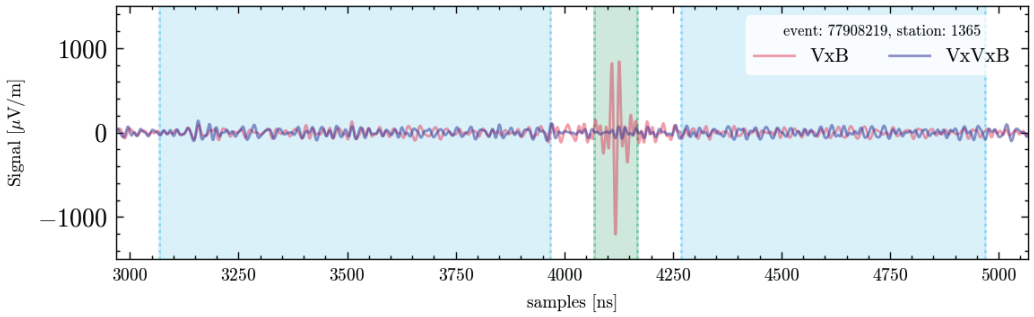


Figure 6.2: Signal and Noise Window Configuration for SNR estimation from station 1365, event 77008219. Green-shaded area indicates the signal window, blue-shaded areas denote noise windows (pre- and post-signal) with $\vec{v}\times\vec{B}$ trace in red and $\vec{v}\times(\vec{v}\times\vec{B})$ in blue.

The combined dual window RMS of each component is calculated as:

$$\sigma = \sqrt{\frac{\sum_{i=1}^{n} x_i^2 + \sum_{i=1}^{m} y_i^2}{n+m}}$$

where n is the number of samples in the pre-signal noise window, m is the number of samples in the post-signal noise window, x_i represents the samples in the pre-signal window, and y_i represents the samples in the post-signal window.
Hence, the SNR$'$ for each component is computed as:

$$\mathrm{SNR}'_{\mathrm{vxB}} = \frac{A_{\mathrm{vxB}}}{\sigma_{\mathrm{vxB}}}, \qquad \mathrm{SNR}'_{\mathrm{vxvxB}} = \frac{A_{\mathrm{vxvxB}}}{\sigma_{\mathrm{vxvxB}}},$$

In practice, the final station selection is determined by the combination of SNR definition and the chosen cut value: an appropriately tuned threshold in a power-based SNR can reach similar efficiency and background rejection as the linear case. In this work, we adopt the linear definition mainly for conceptual simplicity and consistency with the waveform-based quality cuts introduced below, rather than because it is intrinsically less or more sensitive to noise spikes.
Passing the SNR$'$ criteria require meeting two conditions; when the SNR$'_{\mathrm{vxB}}$ component exceeds a minimum threshold ($\mathrm{snr}_{\mathrm{cut}}$) and dominates the SNR$'_{\mathrm{vxvB}}$ as:

$$\mathrm{snr}_{\mathrm{pass}} = (\mathrm{SNR}'_{\mathrm{vxB}} \geq \mathrm{snr}_{\mathrm{cut}}) \quad \wedge \quad (\mathrm{SNR}'_{\mathrm{vxB}} > \mathrm{SNR}'_{\mathrm{vxvxB}}) \tag{6.1}$$

The extra polarization check requires the SNR of the $\vec{v}\times(\vec{v}\times\vec{B})$ component to be strictly less than that of the $\vec{v}\times\vec{B}$ component, directly probing the expected dominance of geomagnetic emission. To quantitatively evaluate the performance, a Figure of Merit (FoM) analysis was conducted to assess the trade-off between signal detection efficiency and noise rejection capability across a range of thresholds. The methodology consisted of three key stages.

First, SNR′ values were computed for each trace in two distinct pulse-search windows.
The "cosmic-ray window" corresponds to the expected air-shower-pulse region from the SD trigger (3900–4300 ns), with pre- (2900–3800 ns) and post- (4400–5100 ns) noise windows providing the RMS estimate. The "control window" is a noise-dominated region where no cosmic-ray pulse is expected (1000–1400 ns), with neighboring quiet segments (0–900 ns and 1500–2200 ns) used as its noise windows; stations passing the cut in this window are interpreted as false positives.
Second, we define two performance metrics based on these windows. For each station, we evaluate the condition in Eq. (6.1) once in the cosmic-ray window and once in the control window. The signal detection efficiency η_s is the fraction of stations in the cosmic-ray window that pass the cut, whereas the noise detection rate η_n is the fraction of stations in the control window that pass the cut (i.e. are falsely tagged as signal). These quantities are obtained from a dense scan of $\mathrm{snr}_{\mathrm{cut}}$.
Third, the resulting η_s and η_n values were used to construct the FoM, as shown in Figure 6.3. Three complementary FoMs were defined to capture distinct performance aspects:

1. The Statistical Significance FoM, $\eta_s/\sqrt{\eta_n}$, which maximizes the detection of true signals relative to Poisson fluctuations in the noise background;

2. The Multiplicative FoM, $\eta_s(1-\eta_n)$, which prioritizes signal purity by weighting detection efficiency by the noise rejection rate $(1-\eta_n)$

3. Youden's Statistic, $\eta_s - \eta_n$, which balances sensitivity and specificity through their arithmetic difference.

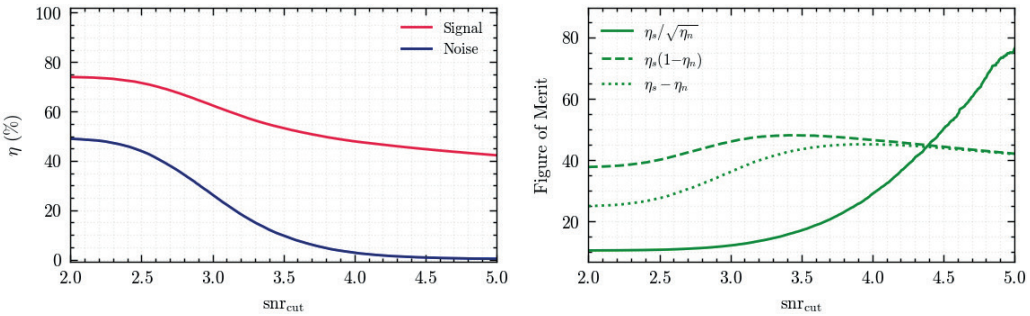


Figure 6.3: Left: Retention efficiency $\eta = N_{\mathrm{kept}}/N_{\mathrm{total}}$ for potential signal window (red) and pure-noise windows (blue) as a function of the $\mathrm{snr}_{\mathrm{cut}}$. Right: Three common figures-of-merit (FoM) built from the same efficiencies: solid line $\eta_s/\sqrt{\eta_n}$; dashed $\eta_s\,(1-\eta_n)$; dotted $\eta_s - \eta_n$.

The left panel shows that tightening the $\mathrm{snr}_{\mathrm{cut}}$ from 2.5 to 4 significantly reduces noise efficiency (η_n) while gradually affecting signal efficiency (η_s). This provides a robust balance between signal and background. The right panel shows that this range, specifically $\mathrm{snr}_{\mathrm{cut}} \approx 3.4-4.0$, is optimal for metrics like the Youden statistic and the multiplicative FoM. In this range, about half the signal is kept while the background is cut by two orders of magnitude to roughly 10^{-2}.

Therefore, a threshold between 2.5 and 4 is a promising starting point, offering a compromise between event purity and yield.

6.1.2 Template Waveform Matching

A SNR′ cut selects stations with signals strong enough to stand out against background noise. However, significant signals are not always genuine air-shower emissions, as RFI can mimic them, producing high-amplitude, high-SNR pulses whose time-domain shape differs from that of genuine air-shower signals.
To distinguish true air-shower signals from RFI and other spurious events, we augment the SNR′ criterion with a cosine-similarity test. This test quantifies the overlap between the observed trace and a normalized template of a typical air-shower signal's pulse in the time domain, shaped by the air-shower geometry and Earth's magnetic field.
The template was developed using air-shower traces in the $\vec{v}\times\vec{B}$ direction that passed the SNR′ criteria as follows:

1. Temporal alignment via peak synchronisation: the signal pulse should lie inside the signal window between sample indices t_{start} and t_{stop}. For each selected trace $v(t)$ we

   (a) locate its maximum bin
   $$t^{\max} \;=\; \underset{t\in[t_{\text{start}},\,t_{\text{stop}}]}{\arg\max} \Big|v(t)\Big|,$$

   (b) compute the shift needed to place that peak at a common reference index t_{ref},
   $$\Delta t = t_{\text{ref}} - t^{\max},$$

   (c) and circularly shift and clip by extracting a fixed window of width $2w$ around the reference:
   $$\tilde{v}(t) \;=\; \Big[\mathcal{R}_{\Delta t}\, v\Big](t), \quad t \in [\,t_{\text{ref}} - w,\; t_{\text{ref}} + w\,],$$
   where $\mathcal{R}_\tau$ denotes the roll (circular–shift) operator. We apply a circular shift so that the pulse maximum of each trace is aligned at a common reference sample before clipping a fixed window around it. The wrapped part of the trace is discarded and does not affect the template.

2. Unit normalisation of waveform shape: to remove amplitude information and focus on the shape, each clipped, time–aligned snippet $\tilde{v}(t)$ is demeaned and scaled to unit Euclidean norm,
   $$\mu = \frac{1}{2w}\sum_{t=t_{\text{ref}}-w}^{t_{\text{ref}}+w} \tilde{v}(t), \qquad \hat{v}(t) = \frac{\tilde{v}(t)-\mu}{\|\tilde{v}-\mu\|_2},$$
   where $\|\tilde{v}-\mu\|_2$ denotes the Euclidean (2-)norm, $\|\tilde{v}-\mu\|_2 = \sqrt{\sum_t(\tilde{v}(t)-\mu)^2}$.

3. Template generation via statistical aggregation: applying the above procedure to all N selected traces yields a set of normalised snippets $\hat{v}_i(t)$, $i = 1, \ldots, N$.

From these we construct the mean and median templates,

$$T_{\text{mean}}(t) = \frac{1}{N}\sum_{i=1}^{N} \hat{v}_i(t), \qquad T_{\text{median}}(t) = \text{median}\Big\{\hat{v}_1(t), \ldots, \hat{v}_N(t)\Big\}.$$

The resultant templates (using 32040 traces) are shown in Figure 6.4. Both templates capture the characteristic shape of an EAS radio pulse, with only minor differences in the bipolar peaks, indicating that outliers have minimal influence on the final averaged waveform. However, in the cosine matching test, the shape only matters, not the amplitudes.

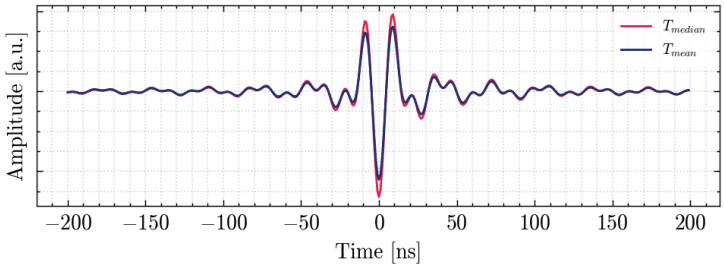


Figure 6.4: Median (red) and mean (blue) pulse templates, constructed from high-SNR′ geomagnetically polarized traces. Each waveform has been time-aligned so that the air-shower-pulse peak falls at $t_{\text{ref}} = 200\,\text{ns}$, clipped to a 100 ns window, and normalized to unity norm.

For the cosine-similarity test, a candidate trace $v(t)$ is compared to the template $T(t)$ via

$$\cos\theta = \frac{\sum_{t=t_{\text{ref}}-w}^{t_{\text{ref}}+w} v(t)\,T(t)}{\|v\|_2\,\|T\|_2}, \qquad \|v\|_2 = \sqrt{\sum_t v(t)^2},$$

and similarly for $\|T\|_2$. A trace passes the waveform cut if

$$\text{corr}_{\text{pass}} = (\cos\theta \geq \text{corr}_{\text{score}})\,.$$

By requiring $\cos\theta_i$ to exceed a certain threshold ($\text{corr}_{\text{score}}$), we ensure that only traces whose waveform morphology closely matches that of an average extensive-air-shower pulse pass the cut.

Constructing the template requires settling three design choices: (i) whether to correlate against the raw bipolar waveform or its envelope; (ii) the width of the correlation window, w, which trades bias (too narrow) against noise leakage (too wide); and (iii) whether to encode geometry's station position and arrival direction via multiple, bin-dependent templates or adopt a single universal template for all stations.

Figure 6.5 stacks high-SNR′ traces. For each aligned snippet, we define the pulse energy as the sum of $v(t)^2$ over the $\pm 200\,\text{ns}$ window; the median fraction contained in the central $|t| \leq 100\,\text{ns}$ interval (between the green dashed lines) is about 95%. Beyond this band, the colour scale is flat, and the ensemble mean (orange) sits at the noise floor, so a 100 ns window captures the full

air-shower waveform while the side bands (out to 200 ns, between the blue lines) remain clean. Because the bipolar pulse retains its sign information, random noise rarely reproduces both lobes with the correct phase, whereas an envelope collapses the trace into a unipolar shape and loses this discriminating power. We therefore adopt a bipolar template and set the correlation window to $w = 100$ or $w = 200$ ns.

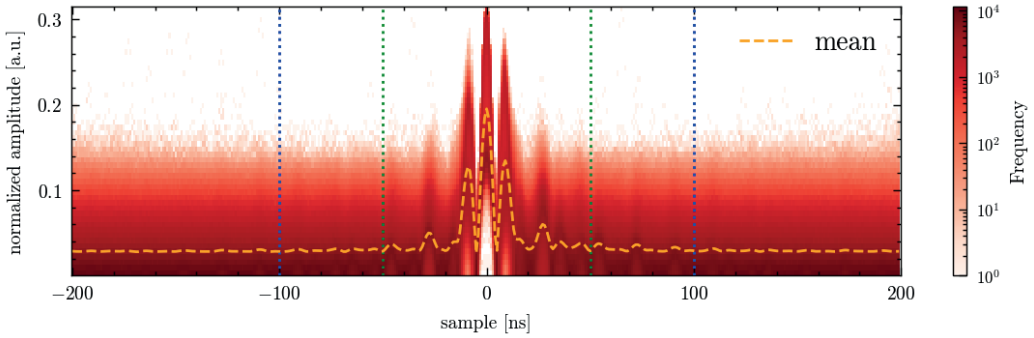


Figure 6.5: Log-scaled 2-D histogram of the absolute, aligned, normalized $\vec{v} \times \vec{B}$ traces after clipping and temporal alignment. The colour map shows how often a given power appears at each sample; the dashed orange curve is the ensemble mean. Vertical green lines mark the 100 ns signal window, while blue lines mark a 200 ns window.

Figure 6.6 overlays the universal bipolar template (black dashed) on the mean waveform in each of the 30 $(d, \theta_{\mathrm{RD}})$ bins. The observed broadening of the radio pulses toward larger lateral distances from the air-shower core reflects simple geometry: emission from different parts of the air-shower front around $X_{\max}$ travels along different slant paths to a distant antenna, so the arrival times are more spread out, and the pulse becomes wider and slightly smoother. This behaviour is expected from geometrical time dispersion and has been discussed in dedicated studies of the pulse-shape dependence on air-shower geometry [120, 121]. Close to the core, these path–length differences are small, and the waveforms are narrow and nearly symmetric, whereas at larger d the time spread increases and the side lobes gain relative weight. The trend with zenith angle follows the same logic: for more inclined air showers, the projected path lengths are longer, enhancing the geometrical time dispersion at a given core distance.
Overall agreement is strong: 24/30 bins have a cosine similarity ≥ 0.95 and 22/30 already exceed 0.98. For all core distances $d \leq 600$ m, the similarity is ≥ 0.988 across every zenith slice. Noticeable deviations appear only at the largest distances: in the shallowest zenith slice $(65^\circ - 68^\circ)$ the similarity falls to 0.86 in the 800–1000 m ring and to 0.30 at 1–2 km; in the $69^\circ - 72^\circ$ slice it remains between 0.91 and 0.88 for $d > 800$ m. These stronger variations in the least-inclined slices are consistent with limited statistics and lower SNR at large d: for 65°–72° the footprint is more compact, so the far-distance rings sample the tail of the LDF and are closer to threshold, making the mean, unit-normalised waveform more sensitive to residual noise and timing jitter. When pooling all zeniths $(65^\circ - 82^\circ)$, the far rings still show high consistency (0.985 at 800–1000 m, 0.960 at 1–2 km). These distant, low-SNR stations carry little statistical weight after the standard quality cuts (section 6.2), so the marginal gain from

distance- or angle-dependent templates is small; we therefore retain a single universal template for the cosine-similarity cut. This also suggests that a single, fixed template could in principle be used in the online waveform trigger-level.

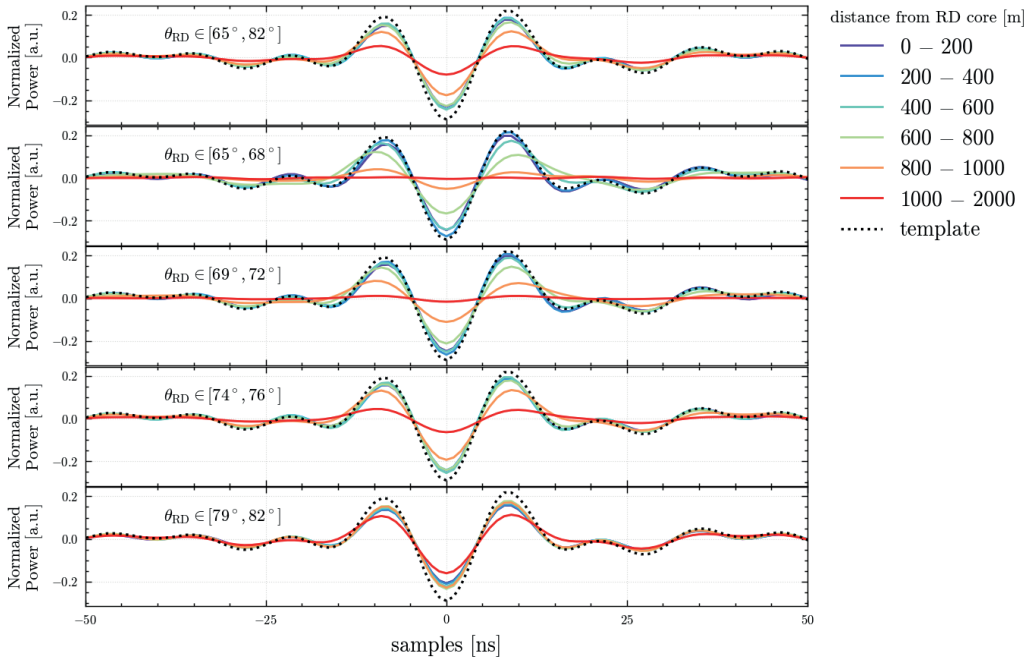


Figure 6.6: Geometry–dependence of the air-shower pulse shape. Each panel corresponds to a zenith range θ_{RD}; within a panel, the colored curves are the mean, time-aligned, unit-normalised $\mathbf{v}\times\mathbf{B}$ waveforms for different core-distance ranges. All traces satisfy $SNR_{cut} \geq 3.8$. The universal bipolar template used in the cosine-similarity test is overlaid as the black dashed curve.

To quantify the effect of the waveform template cut, we scan the correlation threshold $corr_{score}$ and measure the survival efficiencies in the cosmic-ray window and in the control window.
For the cosine cut the main goal is to reject non–air-shower–like pulses, which should be rare in the control window; there η_n directly measures how often pure noise is misidentified as an EAS-like waveform. Results for two correlation windows, $w = 100$ ns and $w = 200$ ns, are shown in Figure 6.7.
For the broader $w = 200$ ns window, tightening $corr_{score}$ from about 0.6 to 0.7 reduces the noise survival in the control window from $\mathcal{O}(40-50)\%$ to a few percent (reduction by roughly an order of magnitude), while the efficiency in the cosmic-ray window decreases from $\sim 90\%$ to $\sim 55-60\%$. The narrower $w = 100$ ns window suppresses noise more gradually, its dashed curves fall at higher thresholds and retains a larger fraction of signal at a given threshold up to ~ 0.8, after which both windows become similarly selective.
The right panel summarizes this behavior in terms of the same three composite FoMs used for the SNR scan. Here $\eta_s/\sqrt{\eta_n}$ should not be interpreted as a formal discovery significance,

but rather as a convenient way to visualize how strongly the cut suppresses the control window background for a given loss in cosmic-ray-window efficiency. In practice, we choose the operating point by requiring very low noise survival in the control window with still acceptable retention in the cosmic-ray window.

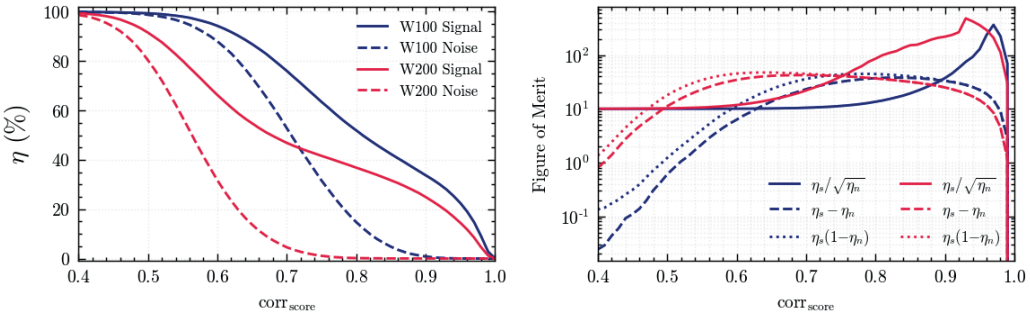


Figure 6.7: Left: Survival efficiency η for signal (solid) and pure–noise windows (dashed) versus correlation threshold $\mathrm{corr}_{\mathrm{score}}$ for $w = 100$ ns (blue) and $w = 200$ ns (red). Right: Corresponding figures of merit: solid $\eta_s/\sqrt{\eta_n}$, dashed $\eta_s - \eta_n$, dotted $\eta_s(1-\eta_n)$. The broader 200 ns window delivers stronger noise suppression at intermediate thresholds, while the significance metric peaks only at very high thresholds.

Guided by the FoM scans, the practically optimal region for the cosine–similarity threshold lies between $\mathrm{corr}_{\mathrm{score}} = 0.70$ and 0.80 for the $w = 200$ ns window. Here we quote efficiencies conditional on the SNR-based preselection (SNR ≥ 7.22): tightening the cut from 0.70 to 0.80 reduces the signal survival from 87.9% to 74.6%, while the control-window noise survival collapses from 4.5% to 0.2%. For $w = 100$ ns, a comparably control regime is not observed. Moreover, thresholds near 0.80 still leave 14.9% noise (with 86.9% signal), whereas a looser cut at 0.70 retains 95.3% signal but still 53.0% noise.

Figure 6.8 maps the Youden statistic over the $(\mathrm{snr}_{\mathrm{cut}}, \mathrm{corr}_{\mathrm{score}})$ plane for $w = 100$ ns and $w = 200$ ns. A broad ridge of optimal performance runs through $\mathrm{corr}_{\mathrm{score}} \simeq 0.6$–0.7 for both windows; tightening the correlation cut beyond ~ 0.85 or loosening it below ~ 0.5 reduces this FoM. The $w = 200$ ns map is uniformly higher, indicating better noise rejection at comparable signal efficiency. Consistent with the SNR′ study, useful $\mathrm{snr}_{\mathrm{cut}}$ values lie around 2.5–4. If we remain in the high-purity SNR′ regime (e.g., $\mathrm{snr}_{\mathrm{cut}} \approx 3.0$–3.8), the practically optimal correlation threshold falls in the $\mathrm{corr}_{\mathrm{score}} \approx 0.65$–0.75 band; moving toward $\gtrsim 0.8$ trades additional purity for a noticeable loss in efficiency, while $\lesssim 0.6$ buys efficiency at the expense of background leakage.

Hence, the global maximum of FoM = 43.91% at $(\mathrm{snr}_{\mathrm{cut}}, \mathrm{corr}_{\mathrm{score}}) = (3.8, 0.6)$, however, we adopted $(\mathrm{snr}_{\mathrm{cut}}, \mathrm{corr}_{\mathrm{score}}) = (3.8, 0.7)$ scoring FoM = 41.52% as we prioritize quality.

6.1.3 Polarization

The significant contribution of the $-\vec{v}\times\vec{B}$ component to the signal, contrary to expectations, prompted an investigation into antenna configuration misalignment (subsection 5.2.1) and the

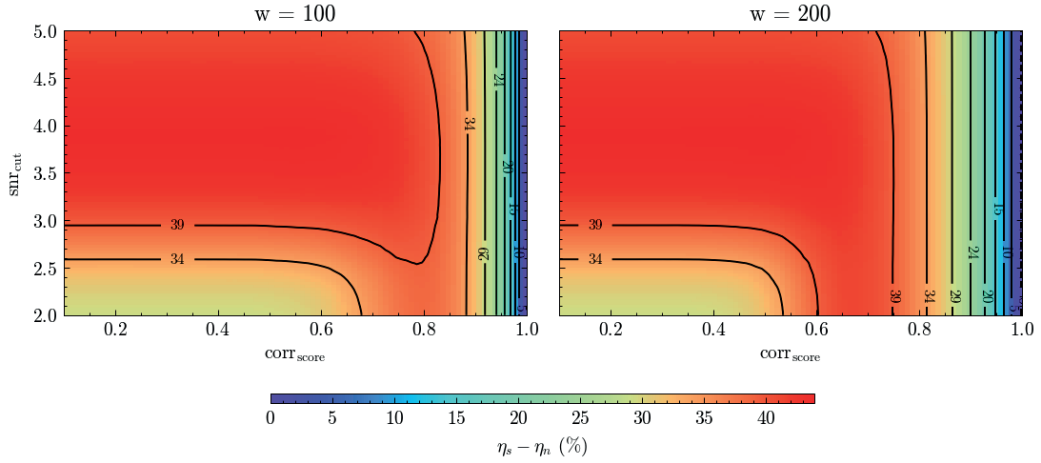


Figure 6.8: Youden statistic FoM for combined cuts. Contours of FoM in percentage as a function of the $\mathrm{snr}_{\mathrm{cut}}$ and cosine-similarity threshold $\mathrm{corr}_{\mathrm{score}}$ for correlation windows $w = 100$ ns (left) and $w = 200$ ns (right).

development of new station quality cuts. The two-stage station-quality selection, based on $\mathrm{snr}_{\mathrm{pass}}$ and $\mathrm{corr}_{\mathrm{pass}}$, enhances background rejection while retaining genuine air shower signals. However, these criteria do not fully ensure clean polarization, a characteristic of air shower events, as they rely on statistical measures of signal power or shape over a time window rather than instantaneous signal purity.

Given the expected dominance of the $-\vec{v}\times\vec{B}$ component in geomagnetic emission, we require its contribution to the pulse strength to exceed 60%. This still allows a substantial fraction of power in the $\vec{v}\times(\vec{v}\times\vec{B})$ component, accommodating non-geomagnetic contributions and residual noise, especially at lower SNR. Signals that do not meet this 60% threshold are treated as polarization outliers, typically cases where $\vec{v}\times(\vec{v}\times\vec{B})$ carries power comparable to or larger than $\vec{v}\times\vec{B}$.

To address this, a dedicated polarization purity check is introduced, focusing on the instantaneous polarization ratio at the point of maximum signal amplitude (subsection 5.2.1).

This is defined as:

$$\xi_{\mathrm{pass}} = \mathrm{Rel}(\vec{v}\times\vec{B}) \geq 0.6$$

Figure 6.9 on the left, the standard RD reconstruction using simple SNR has a significant number of outlier traces (some stations in some events) where the $\vec{v}\times(\vec{v}\times\vec{B})$ contribution exceeds the $\vec{v}\times\vec{B}$ component, especially in air showers coming from the south to south-west. In the tested dataset, the standard SNR cut leaves 2641 polarization outlier traces identified. Adding $\mathrm{snr}_{\mathrm{pass}}$ removes 30% of them, and $\mathrm{corr}_{\mathrm{pass}}$ removes another 77% of the remaining ones. The middle panel shows that applying only $\mathrm{snr}_{\mathrm{pass}} + \mathrm{corr}_{\mathrm{pass}}$ already suppresses most of the clear outliers and yields a much cleaner clustering at large $\vec{v}\times\vec{B}$ contribution, while a small residual population with $\mathrm{Rel}(\vec{v}\times\vec{B}) < 0.6$ remains. The final $\xi_{\mathrm{pass}} \geq 0.6$ requirement eliminates the remainder.

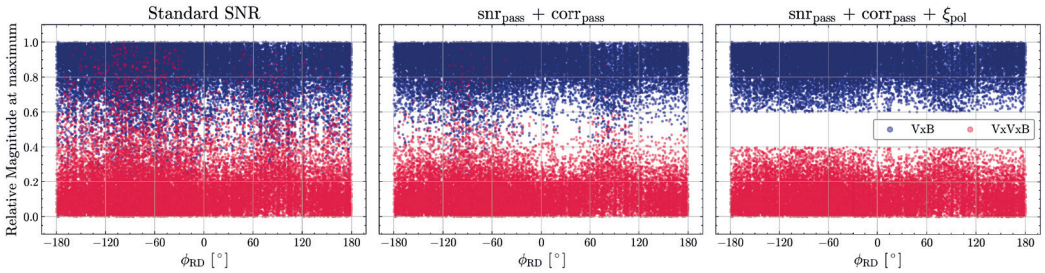


Figure 6.9: Standard reconstructed radio data from the pilot dataset (October 2024 to February 2025). Relative magnitude at the pulse maximum in the air-shower-plane components $\vec{v}\times\vec{B}$ (blue) and $\vec{v}\times(\vec{v}\times\vec{B})$ (red) as a function of reconstructed azimuth ϕ_{RD}. Left: standard station selection with SNR ≥ 7.22. Middle: selection with $\mathrm{snr}_{\mathrm{pass}} + \mathrm{corr}_{\mathrm{pass}}$ using $(\mathrm{snr}_{\mathrm{cut}}, \mathrm{corr}_{\mathrm{cut}}, w) = (3.5, 0.70, 200\ \mathrm{ns})$. Right: full selection with $\mathrm{snr}_{\mathrm{pass}} + \mathrm{corr}_{\mathrm{pass}} + \xi_{\mathrm{pol}}$ and $\xi_{\mathrm{pol}} \geq 0.6$. See the text for how the figure was made.

Finally, applying the new selection criteria (QC) defined as:

$$\mathrm{QC}_{\mathrm{pass}} = \mathrm{snr}_{\mathrm{pass}} \quad \wedge \quad \mathrm{corr}_{\mathrm{pass}} \quad \wedge \quad \xi_{\mathrm{pass}}$$

has eliminated most of the clear outliers, making the data more in line with the reconstructed simulation with noise expectation.

Discussion

The refined station-selection criteria introduced combine physical principles with empirical optimization to suppress noise while retaining genuine air-shower signals.

Selection Criteria	Threshold	Primary Justification	Impact
SNR′	3.5	Maximizes $\eta_s(1-\eta_n)$; 15× more selective than SNR$\geq$ 7.22 (Figure 6.3)	Residual noise <0.1%; ∼3% signal loss
Cosine similarity	0.70 (200 ns)	Retains 96.7% signals, rejects >99.9% noise (Figure 6.7)	Primary rejection cause (45.6% of failures)
Polarization (ξ_{pass})	0.6	Ensures geomagnetic dominance (subsection 5.2.1)	Eliminates 15.9% residual outliers

Table 6.1: Summary of custom station-level quality cuts and their optimization criteria

Table 6.1 consolidates the thresholds, justifications, and performance impacts of these cuts, addressing the trade-offs between signal efficiency (η_s) and noise rejection ($1-\eta_n$) quantified before. The unified QC_pass achieves a >15× reduction in false positives compared to the standard selection, at the cost of discarding about 20% of valid signals.

Practically, by reconstructing the pilot dataset using the new quality criteria resulted in the successful recovery of 80% of all events, and 72% of those that originally had at least five stations. This recovery rate increases to 79% within the full efficiency range defined in section 5.3. In addition, the QC method successfully reconstructed approximately 9% more events overall (or 2% more among events with at least five stations) that were previously missed by the standard method. However, the QC approach still fails to reconstruct about 28% of events handled by the standard method, though this drops to just 3.5% when considering only events with five or more stations.
To quantify the effect of the QC method on RD reconstruction accuracy, comparisons were made for a common set of events with at least five stations.

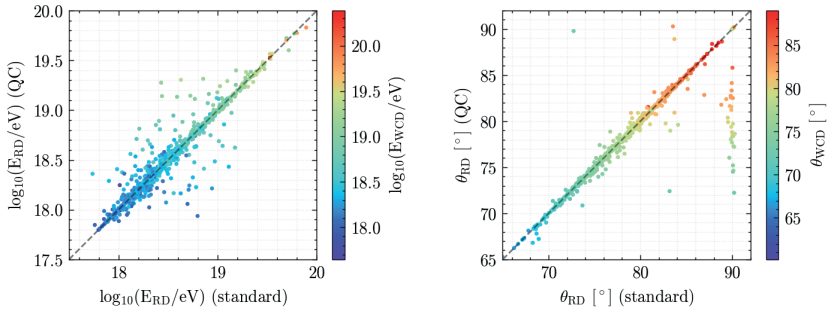


Figure 6.10: Comparison of the reconstructed radio energy and zenith angle between the standard and QC methods. The color bars represent the reconstructed energy and zenith angle as measured by the WCD.

Figure 6.10 focuses on two key parameters: reconstructed energy and zenith angle. The left plot shows the comparison between the logarithm of the reconstructed energy. The majority of events clustered along the diagonal, indicating a high degree of agreement. The color bar represents the reconstructed energy values in terms of $\log_{10}(E_{\mathrm{WCD}})$, which shows that both models generally reconstruct energy with very small deviations. The statistical analysis confirms this with negligible bias and very low variance, a mean difference of –0.01, and a standard deviation of 0.12. On the right, the plot compares the zenith angle differences. Here, the data points also align along the diagonal, but with noticeable outliers, especially at higher zenith angles. A significant set (38 events represent 2.3% shown in the right panel of Figure 6.10) of events have been reconstructed to $\theta_{\mathrm{RD}} \approx 90^\circ$ by the standard reconstruction while the QC method fixed them to agree with the WCD angular reconstruction. The statistical metrics show a mean difference of -0.47° with a large standard deviation of 7.52°; however, this spread is dominated by a small outlier population. The median difference is 0.0° and the central 68% of events lies within $[-0.07^\circ, 0.04^\circ]$, while only 3.04% of events have $|\Delta\theta| > 5^\circ$. Restricting to $|\Delta\theta| \leq 5^\circ$ yields a core standard deviation of 0.37° (mean -0.025°), consistent with the tight clustering around the diagonal in Figure 6.10. The largest deviations correspond to the 38 events (2.36%) reconstructed by the standard method close to $\theta_{\mathrm{RD}} \simeq 90^\circ$, which are corrected by the QC reconstruction.

In these common events, around 6.5% of stations included in the standard reconstruction have not passed the QC. The $\text{corr}_{\text{pass}}$ was the most frequent reason for rejection, being involved in 80.40% of all rejected cases. Polarization was involved in 46.87% of rejections, while snr_{pass} was involved in 26.10%. When considering the sole reason for rejection, cosine matching again stood out, being the only reason for rejection in 45.59% of rejected traces. Polarization ξ_{pass} was the sole reason for rejection in 15.85% of cases, and SNR′ was the sole reason in only 1.18%.

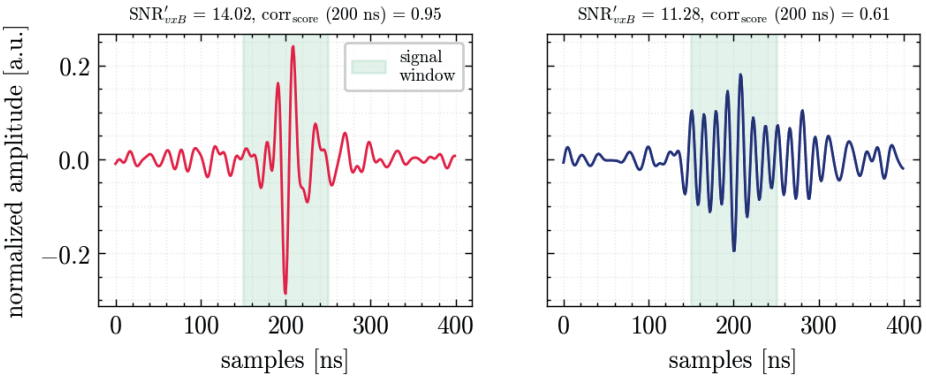


Figure 6.11: Time-aligned, clipped, unit-normalized $\vec{v} \times \vec{B}$ waveforms used in the cosine-similarity test. Both stations pass the SNR′ cut. Left: clean bipolar pulse with high correlation to the template ($\text{SNR}'_{\vec{v}\times\vec{B}} = 14.0$, $\text{corr}_{\text{score}} = 0.95$), accepted by the QC. Right: high-amplitude but irregular waveform ($\text{SNR}'_{\vec{v}\times\vec{B}} = 11.3$, $\text{corr}_{\text{score}} = 0.61$), rejected solely by the cosine-similarity cut.

Looking at the impact across all traces, cosine matching issues led to 59.43% of total traces being rejected, with 33.70% rejected solely due to cosine matching. This confirms that the waveform cut is intentionally the dominant discriminator: it removes many marginal stations, while the overall event statistics and RD–WCD agreement remain largely preserved and the latter even improved (zenith).

To illustrate the typical morphologies accepted and rejected by the waveform cut, Figure 6.11 shows two stations that both satisfy the SNR′ threshold. The left trace exhibits a clean, bipolar air-shower pulse and has a high correlation with the template, whereas the right trace has a distorted, multi-cycle waveform and is rejected solely by the cosine-similarity criterion.

6.2 Events Quality Cuts

The station–level criteria of the previous section are designed to select antenna traces that are clean and dominated by a genuine air-shower signal. However, this alone does not guarantee a high-quality reconstruction. In this section, we introduce additional event-level quality cuts aimed at improving the reliability of the reconstructed energy and arrival directions. These cuts were derived from the pilot dataset and will be used in subsequent chapters.

6.2.1 Number of Radio Stations

For RD reconstruction, the number of contributing stations is one of the key criteria for event selection. You cannot perform any meaningful reconstruction without a sufficient number of stations (section 5.2). Figure 6.12 shows the reconstructed energy and direction for different numbers of stations.

The first panel contrasts the energies reconstructed by RD, E_{RD}, and those obtained with WCD, E_{WCD}. High-multiplicity events cluster tightly around the diagonal, while low-multiplicity events show a broader scatter. The corresponding residual distributions for the same sample are shown in Figure 6.13. For all multiplicities the histogram of $E_{\mathrm{WCD}} - E_{\mathrm{RD}}$ peaks close to zero; even for $N_{\mathrm{RD}} = 3$–5 the mean offset remains at the level of a few EeV (about -1.6 to -3.0 EeV), and it reduces to $\simeq -0.5$ EeV for $N_{\mathrm{RD}} = 8$. This confirms that the radio energy scale is broadly consistent with the WCD reference, with increasing precision as more stations contribute to the fit.

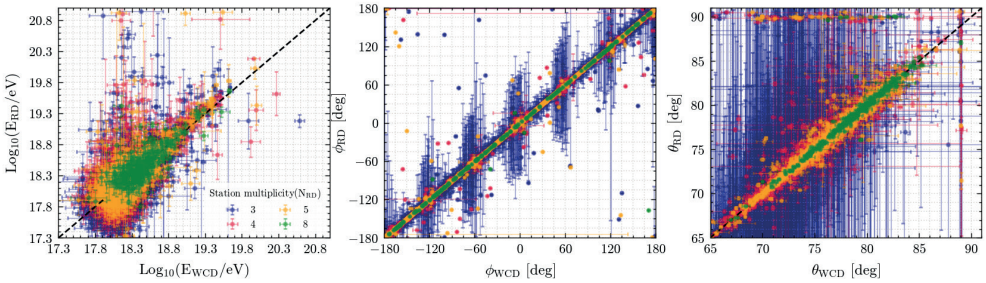


Figure 6.12: Comparison of air-shower properties measured by WCD and RD. The three panels show the correlation for energy (left), azimuth (middle), and zenith (right). Points are colour-coded by the number of radio detector stations (N_{RD}) used in the event reconstruction.

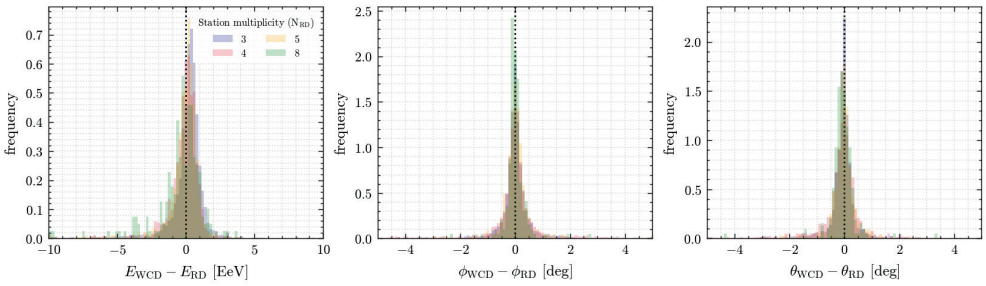


Figure 6.13: Residuals between WCD and RD reconstructions for the events in Figure 6.12, shown for different radio-station multiplicities N_{RD}. Left: $E_{\mathrm{WCD}} - E_{\mathrm{RD}}$ (EeV); middle: $\phi_{\mathrm{WCD}} - \phi_{\mathrm{RD}}$; right: $\theta_{\mathrm{WCD}} - \theta_{\mathrm{RD}}$. The vertical dashed line indicates zero residual.

The other two panels compare the arrival direction (θ, ϕ) from the two detectors. The data in the middle panel follow a strong linear trend, but with a clear $\sim 60°$ periodic pattern in azimuth. This modulation reflects the hexagonal layout of the surface array: for certain azimuths, the stations are nearly collinear, which degrades the triangulation and leads to repeated structures

when the angle is wrapped. In these azimuth bands, the statistical uncertainties also increase, as visible from the larger error bars for low-multiplicity events.

The residual histograms in Figure 6.13 show that the angular differences $\phi_{\mathrm{WCD}} - \phi_{\mathrm{RD}}$ and $\theta_{\mathrm{WCD}} - \theta_{\mathrm{RD}}$ are centred close to zero for all multiplicities, with mean residuals at the level of $\sim 1^\circ$ in azimuth and $\sim 0.5^\circ$ in zenith. The zenith residual width decreases with N_{RD} (from $\sigma_{d\theta} \simeq 3.0^\circ$ at $N_{\mathrm{RD}} = 3$ to $\simeq 1.9^\circ$ at $N_{\mathrm{RD}} = 8$), consistent with improved geometry constraints when more radio stations contribute to the fit. To assess the agreement in units of the quoted reconstruction uncertainties, we also study the corresponding discrepancy (pull) distributions. For the energy, the pull widths are of order unity ($\sigma_{z_E} \simeq 0.9$–1.3), indicating broadly consistent uncertainty estimates. For the angular pulls, the core widths are below unity and become tighter for higher multiplicities (68% widths $\sigma_{68} \simeq 0.58$–0.76 for ϕ and $\sigma_{68} \simeq 0.61$–0.71 for θ at $N_{\mathrm{RD}} \geq 4$), while the rms is dominated by a small tail of events with poorly constrained directional uncertainties, particularly in ϕ (as reflected by the very wide distribution of the combined azimuth uncertainty, reaching tens of degrees in the upper percentiles). Overall, both representations support the same trend: increasing radio-station multiplicity improves the precision and stability of the RD reconstruction with respect to the WCD reference.

Based on the visual evidence in the figure, we adopt $N_{\mathrm{RD}} \geq 5$ as a pragmatic quality cut to obtain a more stable radio reconstruction. This selection reduces the total number of events by 65%, but it suppresses low-multiplicity events in which the directional reconstruction is often poorly constrained, as also reflected by the long tail of large quoted angular uncertainties. This multiplicity yields better agreement between RD and WCD and selects a dataset with higher statistical precision than lower-multiplicity events.

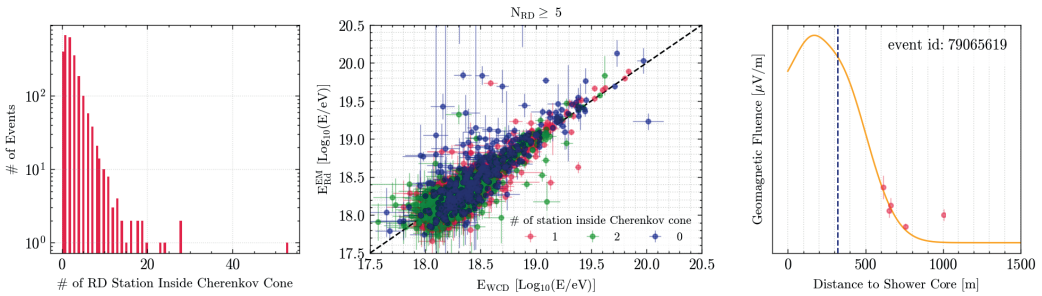


Figure 6.14: Impact of the Cherenkov cone on radio energy reconstruction for at least 5 stations in the event. (left) The distribution of the number of radio stations with pulse located within the Cherenkov cone. (middle) WCD and RD energy comparison for events with 0 (blue), 1 (red), and 2 or more (green) stations inside the cone. (right) Example event (id: 79065619) where no stations are within the Cherenkov cone (indicated by the dashed blue line), where the LDF fit is shown as an orange curve and measurements in red.

The radio emission from air showers is not uniform but concentrated in a Cherenkov-like ring where the coherent signal is at its maximum (section 1.2). Hence, to determine the air shower's energy from the total radio fluence, the measurements must effectively sample this high-signal region. Without at least one station in the Cherenkov cone, a fit to the emission profile is

potentially poorly constrained. This leads to large uncertainties in the reconstructed energy and a likely overestimation of the air shower's true energy. Figure 6.14 shows the importance of RD station placement in three panels. The first panel shows the distribution of the number of RD stations that lie inside the Cherenkov cone for events with at least five stations. Here, "inside the Cherenkov cone" means that the viewing angle between the station and the air-shower axis is close to the local Cherenkov angle, computed from the effective refractive index n_{eff} at the slant depth around $X_{\max}$ via Equation 1.4. While most events have at least one station in the cone, there is a small population ($\sim 10\%$) of events with no stations inside it.
The second panel compares reconstructed energy from WCD and RD for events with zero, one, or two stations in the cone. A diagonal line shows a 1:1 relation agreement. If there are zero stations in the Cherenkov cone, the LDF is constrained only by off-cone measurements and the correlation with E_{WCD} is poorest: the residuals are broad ($\sigma_{\text{resid}} = 0.36$, RMSE $= 0.37$) and essentially uncorrelated with the reference ($R^2 \approx -0.03$). Introducing exactly one in-cone station already provides a precise amplitude anchor at the peak of the radio footprint and markedly improves the fit quality ($\sigma_{\text{resid}} = 0.16$, RMSE $= 0.17$, $R^2 = 0.73$), with only a small mean bias of order 3×10^{-2}. Events with two or more in-cone stations show a similar scatter ($\sigma_{\text{resid}} = 0.18$, RMSE $= 0.18$, $R^2 = 0.65$), indicating that the dominant gain in accuracy comes from having at least one well-placed station in the Cherenkov ring.
The third panel provides a clear example of an LDF fit where no station is located within the Cherenkov ring. All data points are located at a distance from the air-shower core that corresponds to the steep, off-cone tail of the radio footprint. With the peak of the emission left unconstrained. The LDF fit overshoots the measured data points by a large margin, leading to a significant overestimate of the air shower's RD energy of 12.9 EeV compared to the WCD energy of 2.6 EeV. This highlights the importance of sampling the Cherenkov cone to obtain an unbiased radio energy estimate. In particular, photon searches rely on identifying events with an excess of energy in radio compared to WCD, as will be discussed in chapter 7. An overestimate of E_{RD} caused by poor cone sampling would mimic such a signature and could therefore lead to false photon candidates.

6.2.2 LDF Fit Quality

The reduced chi-squared $\chi^2_\nu \equiv \chi^2/\text{d.o.f.}$, is the picked goodness-of-fit metric for the LDF reconstruction. Figure 6.15 illustrates its impact on the event sample that already fulfills the multiplicity and geometry conditions of subsection 6.2.1. This analysis is performed on events that have already passed the prior cuts, including a minimum of 5 radio stations and at least one station within the Cherenkov cone, as detailed in subsection 6.2.1. The left panel presents the distribution of χ^2_ν values for events that passed the previous selections. The histogram shows different station multiplicity groups, with higher multiplicity events (green) exhibiting a narrower χ^2_ν distribution, highlighting better fitting precision. For $N_{\text{RD}} \geq 5$, the distribution has a median around 1.25, with the peak around 0.1. For higher multiplicities, the peak shifts closer to 1, and the median decreases, indicating a better central tendency for the χ^2_ν values.

Events with $\chi^2_\nu \geq 3$ (red dashed line) are rejected, removing approximately 13% of the data. This demonstrates that higher station multiplicities lead to better fits, with a more centered distribution around 1.

The central panel shows the reconstructed energies by fit-quality bands: $\chi^2_\nu \leq 3$, $3 < \chi^2_\nu \geq 4$, and $\chi^2_\nu > 4$. High-quality fits (red) cluster tightly around the 1:1 line, with a bias of 0.0018 and a scatter of 0.1782. In contrast, poorer fits (blue) show progressively larger scatter (0.2808) and bias (-0.0358), with intermediate-quality fits (green) exhibiting a bias of -0.0121 and scatter of 0.2040. This quantification demonstrates that as the χ^2_ν value increases, both the bias and scatter increase, leading to larger deviations from the ideal 1:1 line.

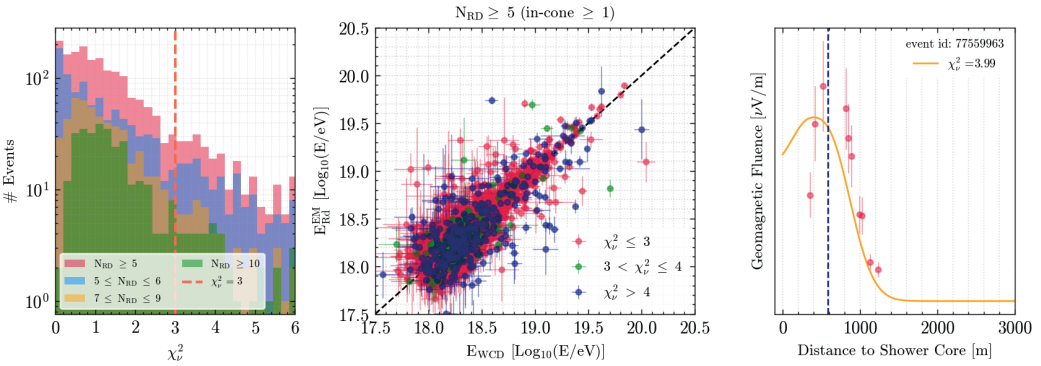


Figure 6.15: Impact of the reduced LDF fit-quality parameter χ^2_ν. (a) Distribution of χ^2_ν of RD for events with different station multiplicities. The vertical dashed line indicates the threshold which is the cutoff used for further analysis. (b) Correlation of E^{EM}_{RD} with E_{WCD}, color-coded by fit-quality band. The R^2 for $\chi^2_\nu \leq 3$, $3 < \chi^2_\nu \leq 4$, and $\chi^2_\nu > 4$ are 0.7311, 0.4192, and 0.3537, respectively. (c) Example event ($\chi^2_\nu = 3.99$) where the dashed line indicates the Cherenkov radius.

Occasionally, the fit yields high χ^2_ν even with multiple clear stations as in the right panel displays an example where the event has 10 stations (3 are inside the Cherenkov cone), and the χ^2_ν is 3.99. The third panel also illustrates a typical failure mode: even for moderate χ^2_ν, the LDF fit can be dominated by a single small-distance, high-weight station, so χ^2_ν is a necessary but not sufficient indicator of a robust energy reconstruction.

Hence, accordingly, we retain only air showers with $\chi^2_\nu \leq 3$. This threshold is strict enough to remove the worst-fitting tails while still preserving most of the dataset.

This phenomenon highlights a specific limitation within the current reconstruction software's implementation rather than an inherent deficiency of the LDF model itself. The transition in the number of fitted parameters when moving from three to four stations, leading to a perfectly determined system, creates an artificial barrier for quality control based on standard statistical metrics. Since the present analysis requires $N_{RD} \geq 5$, the $\nu \simeq 0$ regime of the reduced-χ^2 statistic is avoided. This detailed understanding is crucial for optimizing reconstruction pipelines and interpreting efficiency plots in cosmic ray air shower analysis, as well as benefiting from all recorded data.

6.2.3 Opening Angle

For events which are reconstructed by both the radio and WCD, the opening angle ($\Delta\Psi$) serves as a key empirical observable. For each event j, if $\Psi_{\mathrm{WCD},j}$ and $\Psi_{\mathrm{RD},j}$ represent the reconstructed arrival directions from WCD and RD, respectively, then the opening angle is defined as:

$$\Delta\Psi_j = \arccos(\hat{\mathbf{k}}_{\mathrm{WCD},j} \cdot \hat{\mathbf{k}}_{\mathrm{RD},j})$$

where $\hat{\mathbf{k}}$ are the unit direction vectors corresponding to the reconstructed zenith and azimuth angles. The distribution of these $\Delta\Psi$ values across the dataset of hybrid events provides a powerful insight into the consistency and relative precision of the two reconstruction techniques.

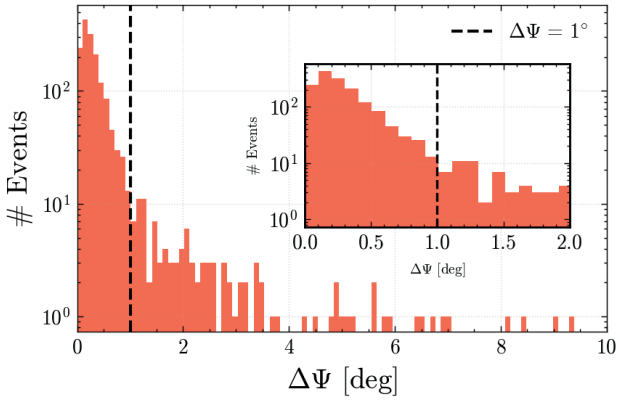


Figure 6.16: Histogram of the inter-detector opening angle ($\Delta\Psi$) between RD and WCD reconstructed air-shower directions for the pilot dataset after applying the multiplicity ($N_{\mathrm{RD}} \geq 5$), Cherenkov-cone, and LDF-quality cuts, but before the $\Delta\Psi$ cut. The main panel shows the full $\Delta\Psi$ range on a logarithmic y-axis to emphasize the tail, while the inset shows a zoom of the core region (0°–2°). The vertical dashed black line indicates the cut value used ($\Delta\Psi = 1°$).

The distribution of the opening angle, $\Delta\Psi$ shown in Figure 6.16 highlights this deviation in the distribution's tail. The distribution is peaked with a median $\Delta\Psi = 0.24°$ and 95^{th} percentile of 0.92°, indicating that for the vast majority of events, the WCD and radio reconstructions yield very similar arrival directions.

Approximately 8% of events (139 events) show $\Delta\Psi > 1.0°$, with the tail extending up to $\Delta\Psi \approx 10°$. The presence of events with $\Delta\Psi$ values exceeding a few degrees signifies instances where the radio reconstruction, despite its WCD-seeded initialization, finds a notably different optimal direction based on its precise timing information. These events are crucial for understanding the specific conditions and limitations of directional reconstruction under which the two reconstruction methods may diverge or where one method might offer better precision.

To ensure the quality of the selected hybrid events for further analysis, a quality cut is applied to the opening angle distribution. A value of $\Delta\Psi < 1.0°$ is chosen for this cut. This decision is motivated by the statistical properties of the distribution.

6.2.4 Thunderstorm Flagging

Thunderstorms generate strong electric fields that can either mimic radio signals from air showers or significantly alter the radio emission process itself. This necessitates a systematic approach to flag potentially affected events [122]. While direct measurements from Electric Field Mills (EFM) provide local ground-level information at six stations across the array [123], they cannot capture the full three-dimensional distribution of atmospheric electrical conditions relevant to radio detection, especially for very inclined events.

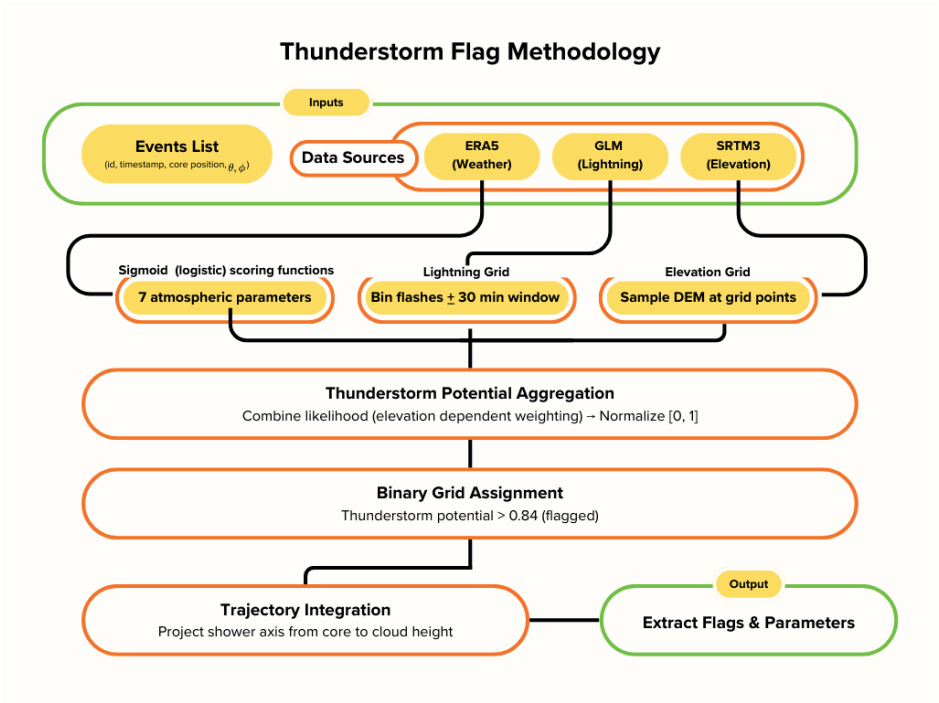


Figure 6.17: Thunderstorm-flagging pipeline. Event metadata are combined with ERA5 meteorological fields, GOES-16 GLM lightning observations (binned in ±30 min windows), and SRTM3 topography. Seven atmospheric variables are mapped to continuous scores in $[0, 1]$ using sigmoid functions and aggregated into a gridded thunderstorm potential. GLM flashes override this score, forcing the potential to 1 (always flagged) in cells with detected lightning. Finally, the potential is sampled along the reconstructed air-shower axis up to 16 km and weighted by the air-shower profile (peaked at $X_{\max}$) to obtain the event-level flag ($\mathrm{TS}_{\mathrm{flag}}$).

This limitation motivated the development of a weather-based thunderstorm flag that exploits meteorological data to assess thunderstorm potential across the entire air-shower development path. The code was originally developed by B. Pont [124]. For this analysis, the implementation was extended in three respects: (i) the I/O was adapted to the RD event dataset, reading event metadata (core position, zenith, and azimuth) to store per-event flags; (ii) the elevation lookup was replaced by a static local Digital Elevation Model (DEM) grid to ensure fast and reproducible elevation estimates over the full domain; and (iii) event-level diagnostic metrics were added by sampling the gridded thunderstorm potential along the reconstructed air-shower

axis. This allows for weighting the potential by the air-shower development profile around the expected $X_{\max}$.

As shown in Figure 6.17, the flag combines meteorology-based predictions using ERA5[2] with validation from satellite-based lightning observations by the GOES-16 Geostationary Lightning Mapper (GLM). It incorporates local topography from the Shuttle Radar Topography Mission (SRTM3, 3-arcsec resolution) to enable elevation-dependent weighting. ERA5 provides hourly meteorological fields on a 0.25° grid (approx. 28 km spacing). While this resolution is coarse compared to the timescales of an air shower, ERA5 represents the best available global product for retrospective analysis. The spatial domain extends ±2.5° around the observatory to encompass the atmospheric volume where inclined air showers interact with thunderstorm conditions up to a cloud height of 16 km.

Thunderstorm potential is assessed using seven near-surface and thermodynamic variables: 2 m air temperature ($t_{2\mathrm{m}}$), 2 m dewpoint ($d_{2\mathrm{m}}$), their difference $\Delta(T-T_d)$, mean sea-level pressure (msl), total cloud cover (tcc), surface wind speed, and Convective Available Potential Energy (CAPE). These variables represent the canonical ingredients for convective storms: surface heating, moisture availability, buoyant instability, and low-level dynamic organization. The parameter settings, summarized in Table 6.2, utilize reference values x_0 rather than hard cuts. A key methodological choice is mapping each variable to a continuous score in $[0, 1]$ using a smooth sigmoid (logistic) function:

$$L(x) = [1 + \exp(\mp k(x - x_0))]^{-1}$$

where x_0 is the midpoint and k controls the transition sharpness. The sign of k is chosen such that $L(x) \to 1$ indicates storm conditions (e.g., high temperature or low pressure). This avoids artificial discontinuities and captures marginal conditions that binary cuts might miss. The individual scores are aggregated into a single gridded thunderstorm potential $L \in [0, 1]$. This aggregation uses the SRTM3 elevation to distinguish between plains ($h < 1.5$ km) and foothills ($h \geq 1.5$ km), applying different weights to the predictors to account for the observatory's orography. Figure 6.18 illustrates the resulting fields: the continuous potential L, the binary flag (where $L > 0.84$), and the validation layers.

Validation is provided by GLM lightning data. For each event, GLM flashes within ±30 min are binned onto the grid. Any cell with detected lightning is assigned maximum potential ($L = 1$), ensuring that confirmed electrical activity is always flagged regardless of the meteorological model.

The final event-level flag integrates this 2D potential along the reconstructed 3D air-shower trajectory. The air-shower axis is projected up to 16 km altitude (or $3 \times X_{\max}$). The grid values are interpolated onto this path and weighted by a Gaussian profile centered at $X_{\max}$ (width 5 km) to represent the longitudinal air-shower development. While metrics such as the weighted average potential $\langle L \rangle_w$ and flagged fraction f_{flag} provide continuous diagnostics, we apply a strict binary flag for the energy reconstruction analysis.

[2] A global reanalysis dataset produced by the European Centre for Medium-Range Weather Forecasts.

Variable	Symbol	Condition (x_0)	Δx	Rationale
2-m temperature	t_{2m}	$> 18\,°C$	$2\,°C$	Surface heating
2-m dewpoint	d_{2m}	$> 4\,°C$	$2\,°C$	Moisture presence
Temp–dewpoint diff.	$\Delta(T-T_d)$	$> 15\,K$	$5\,K$	Moisture-deficit proxy
Mean sea-level pressure	msl	$< 1012\,hPa$	$1\,hPa$	Synoptic ascent proxy
Total cloud cover	tcc	> 0.03	0.01	Cloud development
Surface wind magnitude	$\|u\|$	$> 2\,m\,s^{-1}$	$2\,m\,s^{-1}$	Near-surface flow proxy
CAPE	CAPE	$> 12\,J\,kg^{-1}$	$40\,J\,kg^{-1}$	Buoyant instability

Table 6.2: Parameter settings for the sigmoid scoring functions used to construct the thunderstorm potential. Each variable is mapped to a score in $[0, 1]$ with midpoint x_0 and decay window Δx (transition width).

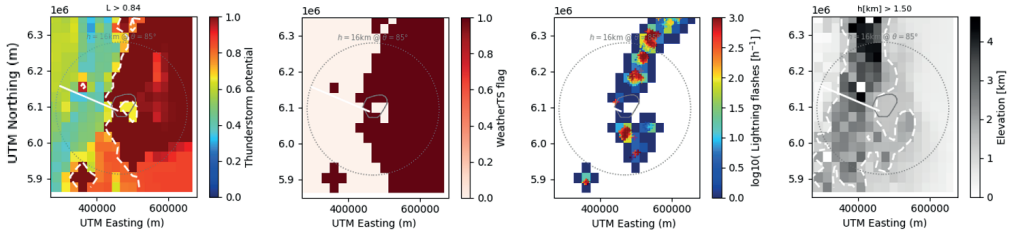


Figure 6.18: Thunderstorm diagnostics for a single event. Left to right: (i) Gridded thunderstorm potential $L \in [0, 1]$ from ERA5 (white contour denotes $L > 0.84$); (ii) Binary weather flag; (iii) GLM lightning density ($\log_{10}$ flashes), which overrides the weather score by setting $L = 1$; (iv) SRTM3 terrain elevation, showing the 1.5 km contour used for aggregation weighting. The reconstructed air-shower axis is overlaid in white.

The event is flagged (TS_{flag}) if any grid cell within $\pm 0.5\,X_{max}$ of the air-shower maximum exceeds the thunderstorm potential threshold of 0.84. This threshold was empirically established by optimizing agreement with GLM ground truth and local electric field mill data during the 2024 validation campaign. The flag ensures that no event is used for reconstruction if atmospheric conditions at the signal generation height are likely to distort the radio emission. The impact of this criterion on the dataset statistics is summarized in Table 6.3.

Summary of Quality Cuts

By applying the three station-level quality cuts discussed to the pilot dataset, with standard reconstruction followed by event-level quality selection, the compounded efficiency is approximately 16%. The most aggressive cut is the number of stations required, as shown in Table 6.3.

However, despite that the Figure 6.19, which shows data after quality cuts, has a significant improvement in the correlation plots compared to Figure 6.12, a few outliers are still present and require further investigation to understand the underlying causes, rather than adjusting the cut thresholds. There are a few events with θ beyond 80° where WCD angular reconstruction is not optimized; however, due to the accurate RD zenith reconstruction, the resultant opening

Criteria	N events	efficiency [%]
Reconstructed	6383	100%
≥ 5 stations	2517	39%
≥ 1 in-cone	2116	84%
$\chi^2_\nu \leq 3$	1656	78%
$\Psi \leq 1$	1517	92%
TS_{flag} = False	1005	66%
Overall	**1005**	**16%**

Table 6.3: Summary of selection criteria and their impact on event counts for the pilot dataset.

angle is still below 1°. The azimuth reconstruction shows a similarly tight clustering around the diagonal as the zenith. Azimuth error bars are included, but they are typically at the level of $\mathcal{O}(1°)$ and therefore appear visually negligible on the full $[-180°, 180°]$ scale. A single event exhibits an anomalously large ϕ_{WCD} uncertainty ($\sigma_{\phi,\mathrm{WCD}} \sim -61.6°$), which reflects a poorly constrained SD directional covariance estimate rather than the RD reconstruction quality (the event still satisfies all RD-based selections and has a small opening-angle difference).

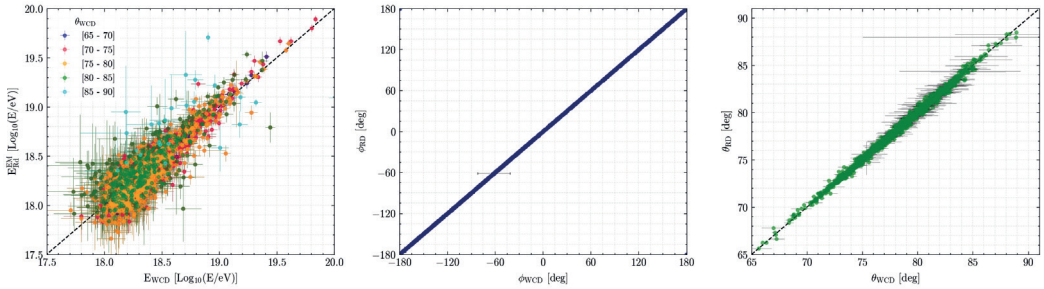


Figure 6.19: Performance comparison of RD and WCD cosmic-ray air-shower reconstructions. The panels illustrate the agreement in: (left) reconstructed energy ($\log_{10}(E/\mathrm{eV})$) for different zenith angle (θ_{WCD}) ranges; (middle) azimuthal angle (ϕ_{WCD}); and (right) zenith angle (θ_{WCD}).

The general behavior of the event quality cuts does not exhibit a dramatic difference when using the signal quality checks (section 6.1), but rather with a significantly reduced number of outliers, as it is also reduced to 16%.

6.3 Angular Resolution

Since the RD operates in slave-mode triggered by the WCD, an absolute truth for the arrival direction is not available for RD-only events. To characterize the performance, we examined the opening angle $\Delta\Psi$ between the RD and WCD angular reconstructions. This comparison represents the convolution of the resolutions of both detectors.

We first assess the reliability of the reported uncertainties by examining pull distributions. Figure 6.20 shows the difference between the WCD and RD reconstructed angles normalized by their combined uncertainties. The zenith angle pull distribution is centered at $\mu = 0.129 \pm 0.016$,

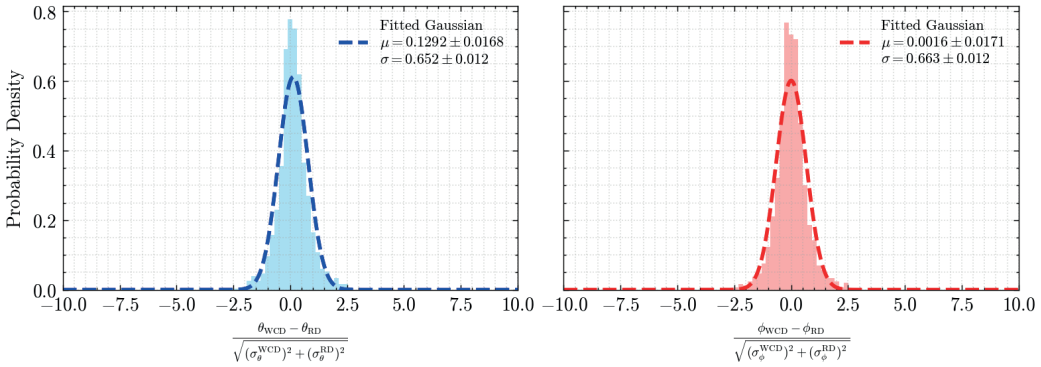


Figure 6.20: Histograms of the normalized differences (pull distributions) for zenith (left, light blue) and azimuth (right, light red) angles between WCD and radio reconstructions, after a 3-sigma cut. A Gaussian fit is overlaid on each distribution (dashed blue for zenith, dashed red for azimuth), with the fitted mean (μ) and standard deviation (σ) indicated in the legend.

indicating a minor systematic offset between the detectors. The azimuth distribution is consistent with zero ($\mu = 0.0016 \pm 0.017$). Crucially, both distributions exhibit standard deviations significantly smaller than unity ($\sigma \approx 0.65 - 0.66$). This implies that the combined uncertainties are overestimated by a factor of $1/\sigma \approx 1.5$. Given the precision of radio wavefront timing, this overestimation is likely dominated by conservative error estimation in the WCD reconstruction. It should be noted that no standard Surface Detector (SD) quality cuts or energy thresholds are applied at this stage. Consequently, the sample remains biased toward low-energy events near the detection threshold; because WCD reconstruction performance and error estimation improve significantly at higher energies, the presence of these marginal events accounts for the apparent overestimation of the combined uncertainties

To disentangle the contributions of each detector, we examine the individual uncertainties in Figure 6.21. The left panel reveals two key features. First, WCD uncertainties systematically exceed RD uncertainties by approximately one order of magnitude. Second, there is a strong correlation with station multiplicity. Events with high $N_{\rm RD}$ cluster in the bottom-left corner, exhibiting the smallest uncertainties for both detectors. This confirms that high-multiplicity radio events correspond to higher-energy air showers that also trigger more WCD stations, leading to a simultaneous improvement in both WCD and RD resolutions.

The middle and right panels contrast the zenith dependence. RD uncertainties remain consistently low ($< 0.2^\circ$) and relatively flat across the zenith range. In contrast, reported WCD uncertainties show a clear dependence on inclination, increasing significantly as the air shower becomes more horizontal ($> 80^\circ$).

The intrinsic improvement of the radio reconstruction is quantified in Figure 6.22. The uncertainty drops rapidly from $\sim 0.12^\circ$ at $N_{\rm RD} = 5$ to a plateau of $\sim 0.04^\circ$ for $N_{\rm RD} \geq 15$.

Combining these observations, the median opening angle of $\Delta\Psi \approx 0.16^\circ$ observed for high-quality events serves as a conservative upper bound on the RD resolution.

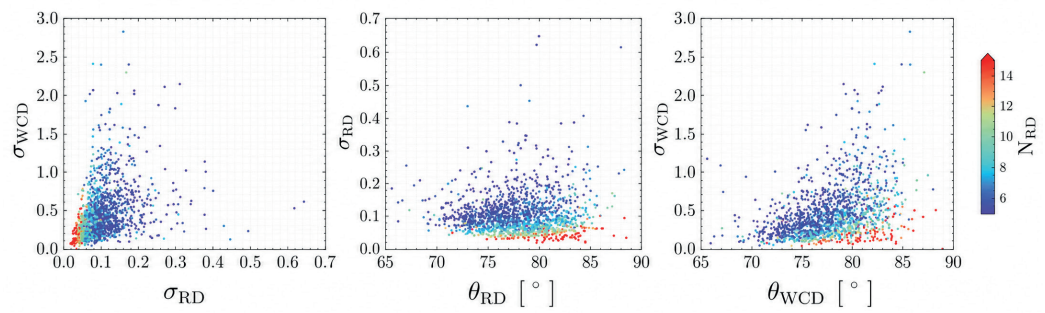


Figure 6.21: Opening angle $\Delta\Psi$ as a function of air-shower energy (left) and zenith angle (right), color-coded by the number of RD stations N_{RD}.

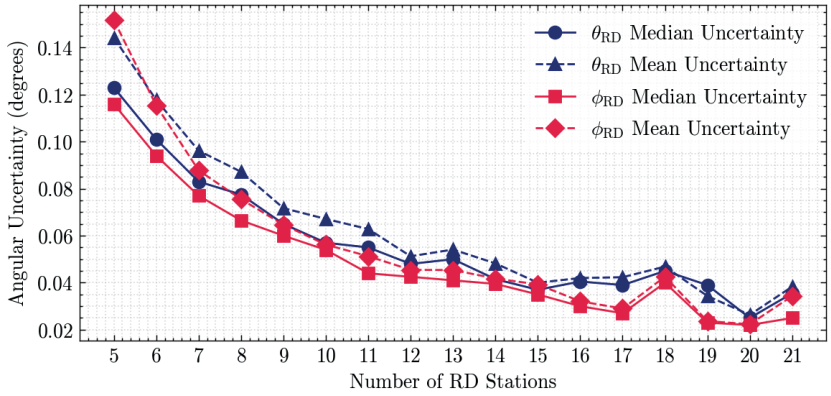


Figure 6.22: Median and mean angular uncertainty for reconstructed zenith (θ_{RD}) and azimuth (ϕ_{RD}) angles for RD as a function of the number of RD stations.

Since $\Delta\Psi$, and the scatter plots confirm $\sigma_{WCD} \gg \sigma_{RD}$, the contribution of the radio detector to the opening angle is sub-dominant. This demonstrates that the RD angular reconstruction achieves a precision significantly better than the standard WCD limits, making it especially relevant for improving point-source photon searches.

Chapter 7

Photon-hadron discrimination using WCD and RD

Detecting UHE photons requires identifying the extensive air showers they produce. The primary background in these studies comes from hadron-induced air showers.
The discrimination of UHE primaries is limited by large air-shower-to-air-shower fluctuations and by degeneracies in single-detector energy reconstruction. In WCD-only reconstruction, the ground signal depends on both the electromagnetic and muonic shower components. Fluctuations in these components can cause hadronic and photon-induced showers to overlap within detector and reconstruction uncertainties. Hybrid events reduce this ambiguity by providing two complementary energy estimates: one derived from the WCD array, which reflects the particle content at ground, and one derived from the radio array, which is closely linked to the electromagnetic cascade.
In this chapter, we study photon–hadron separation using Monte Carlo simulations reconstructed with the standard WCD and RD pipelines. The labeled signal (photons) and background (hadrons) samples enable controlled tests of the chosen observables and of bias-mitigation procedures. Our objectives are (i) to identify a simple, physically motivated separation criterion for inclined hybrid events and (ii) to validate the AugerPrime concept by quantifying the improvement achieved with combined WCD+RD measurements. A multivariate machine-learning approach was also investigated during the development phase. For completeness, the corresponding feature engineering, polynomial detrending, classifier optimisation, and performance study are documented in Appendix D.

7.1 Simulation Reconstruction and Selection

For this task, all simulations were reconstructed using the standard procedures for WCD and RD in slave mode (section 5.2) to ensure consistency with the data and to reproduce the expected synergy between detectors. The simulated signals were reconstructed with added measured noise, which is nevertheless relatively clean compared to real data. To arrive at similar quality as the real data, we applied both the previously developed station-level quality criteria and the

event-level selection cuts. This resulted in an event loss of approximately 50%.
Because these cuts substantially reduce the available statistics-limited simulated data sets, it is essential to include all proton simulations generated with different hadronic interaction models, such as SIBYLL 2.3d and QGSJETII-04. These models are needed to account for uncertainties in modeling high-energy cosmic ray interactions with the atmosphere. Using multiple models improves the robustness of the photon identification by capturing a wider range of possible air-shower characteristics, because real cosmic-ray data may not correspond closely to any single model currently available. Similarly, iron-initiated air showers were added as a background sample to better approximate the mixed composition of UHECRs, as these represent the heavy-mass end of the spectrum. Iron primaries produce air showers with higher muon content and different development profiles than protons, and incorporating them provides a broader representation of potential background characteristics than with protons alone, even though the true composition includes intermediate masses not included in the available simulation sets. Hence, the *background* in this analysis is a composite of both protons and iron.

Cut	proton (QGSJETII-04)	proton (SIBYLL 2.3d)	iron (SIBYLL 2.3d)	photon (SIBYLL 2.3d)
Reconstructed (WCD+RD)	1370	3344	1620	495
$65^\circ \leq \theta_{\mathrm{RD}} \leq 82^\circ$	1277	2681	1514	491
$N_{\mathrm{RD}} \geq 5$	909	2267	1081	321
$N_{\mathrm{cone}}^{\mathrm{RD}} \geq 1$	747	1950	899	252
Final selection	725	1923	879	237

Table 7.1: Cutflow counts for each simulated dataset after successive selection cuts. Hadronic interactions are simulated using SIBYLL 2.3d (protons and iron) and QGSJETII-04 (protons).

Table 7.1 summarizes the progressive reduction of event counts after each selection cut for the different primary particle types. The cuts on reconstructed zenith angle ($65^\circ \leq \theta_{\mathrm{RD}} \leq 82^\circ$) and basic radio-detector quality ($N_{\mathrm{RD}} \geq 5$, $(\chi^2_\nu)_{\mathrm{RD}} \leq 3$) reduce all samples. The strongest model-dependent difference for protons appears after the requirement $N_{\mathrm{RD}} \geq 5$: this cut retains $909/1277 \simeq 71\%$ of the QGSJETII-04 protons after the zenith-angle selection, compared with $2267/2681 \simeq 85\%$ of the SIBYLL 2.3d protons. The hybrid directional consistency requirement, $\psi_{\mathrm{RD-WCD}} \leq 1^\circ$, is then applied to require agreement between the air-shower axes reconstructed independently by the RD and WCD arrays. This cut removes only ~ 0.5–$1\,\%$ of the events, indicating that well-reconstructed hybrid air showers yield highly compatible directions in both subsystems. Finally, the cone multiplicity requirement ($N_{\mathrm{cone}}^{\mathrm{RD}} \geq 1$) produces the sharpest drop, especially for protons, since it demands at least one radio station in the expected Cherenkov-cone region and therefore suppresses events with insufficient or poorly reconstructed radio information.
The two proton models also differ in their total survival fractions. After all cuts, $725/1370 \simeq 53\%$ of the reconstructed QGSJETII-04 proton sample remains, compared with $1923/3344 \simeq 58\%$ of the SIBYLL 2.3d proton sample. These differences may reflect hadronic-model-

dependent shower properties, such as fluctuations and the relative electromagnetic and muonic content, but also residual differences in the generated phase space, for example in the (E, θ) distributions. For this reason, both proton models are retained in the background definition to improve robustness against hadronic-model and phase-space systematics.
Despite these losses, the photon sample retains sufficient statistics after the final selection, with 237 events remaining, enabling a meaningful proof-of-concept photon–hadron separation while preserving the intrinsic differences in air-shower development across primaries.

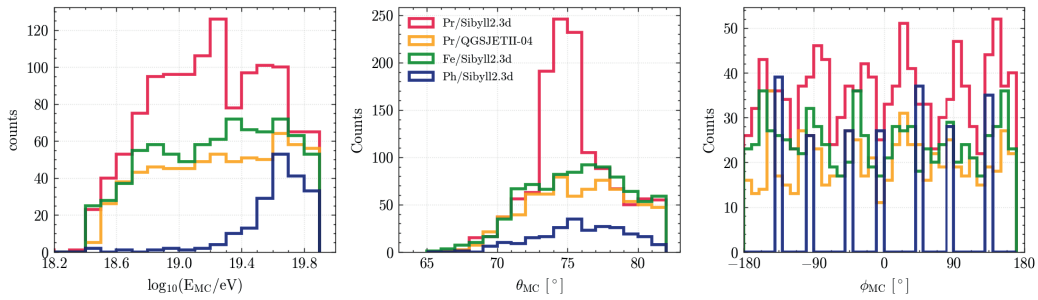


Figure 7.1: Distributions of the true Monte Carlo primary energy E_{MC} (left), zenith angle θ_{MC} (center), and azimuth angle ϕ_{MC} (right) for proton (red, orange), photon (blue), and iron (green) events after event-quality selection.

Figure 7.1 shows the distributions of the true Monte Carlo geometrical parameters for the selected hybrid dataset. To increase statistics in the region of highest radio-detector performance, SIBYLL 2.3d proton simulations were oversampled in the zenith range 72°–77° (section 5.3), while an additional SIBYLL 2.3d proton set was also included. The resulting peak in the proton θ_{MC} distribution reflects this intentional oversampling. Consequently, the different primaries (and hadronic models) are not drawn from identical zenith-angle mixtures. Since reconstruction quality and selection efficiencies depend on θ, such differences can propagate into the cutflow and feature distributions. Moreover, due to the fixed simulation geometry, the photon azimuth distribution appears discrete, whereas real data exhibit an almost flat ϕ acceptance.
In the current hybrid configuration, the RD is operated in slave mode to the WCD (chapter 5). As a result, the sample is strongly shaped by the WCD trigger threshold, which is visible in the left panel of Figure 7.1 as a significant shift in photon energies compared to the hadronic background. Below $\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) \simeq 19$, the probability for photon-induced air showers to trigger the WCD is very low. For the discrimination study, however, the analysis threshold is imposed on the reconstructed radio energy, $\log_{10}(E_{\mathrm{RD}}/\mathrm{eV}) \geq 19$, which is the reconstructed-energy variable used for the photon-candidate selection in data.
This choice reduces the impact of trigger-driven selection bias and ensures that the comparison between photons and hadrons is performed in a regime where both populations are meaningfully represented.
Without accounting for the WCD trigger threshold, the apparent separation between photons and hadrons would be dominated by a trigger-driven energy bias rather than by genuine air-shower-physics differences. Energy detrending (Section 7.3) mitigates the residual energy

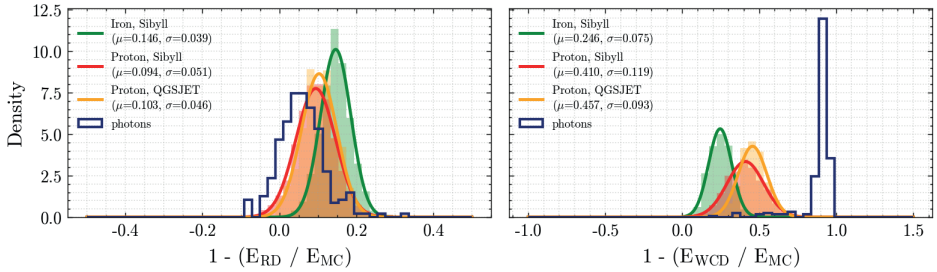
dependence and ensures that the separation is driven primarily by the relative WCD–RD energy mismatch, rather than by the absolute energy scale.
Having established the dataset and its limitations, we now examine the detector observables that will form the basis of the separation strategy.

7.2 WCD and RD Observables

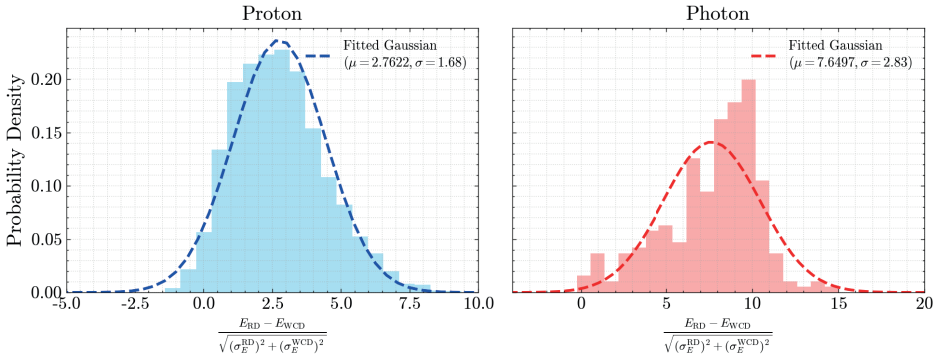
The UHE photon searches at the Pierre Auger Observatory and the Telescope Array have consistently been restricted to less inclined air showers section 1.5. This study focuses on highly inclined air showers with $\theta \geq 65°$. While Auger is capable of reconstructing cosmic-ray events at such large zenith angles, with dedicated algorithms, no inclined photon searches have been conducted so far, and observables have not yet been studied in this range.
Specifically, for $\theta \geq 65°$, reconstructing $X_{\max}$ becomes significantly more challenging due to unfavorable geometries and atmospheric light-propagation effects. Several radio-based techniques have been developed but not yet used in a dedicated photon search.
Observables such as risetime, curvature, S_{800}, S_b, and multivariate classifiers trained for $\theta \leq 60°$ become at high inclination ineffective or heavily biased [125], because inclined air showers are overwhelmingly muon-dominated and carry little electromagnetic air-shower signature.
As shown previously, the RD adds a powerful independent measurement of the electromagnetic component, which for inclined air showers completes the picture. The simplest way to quantify the electromagnetic and hadronic components is the reconstructed energy: for highly inclined events, the WCD energy estimator is largely driven by the muonic signal at ground, while in the RD it is derived from the fitted radio-fluence pattern, which is generated by the electromagnetic cascade throughout the air-shower development.

The Reconstructed Energies

In Figure 7.2, we compare two key performance metrics of the WCD and RD energy reconstructions for inclined background- and photon-initiated air showers. All separation results shown here are obtained after applying the reconstructed-radio-energy threshold $\log_{10}(E_{\mathrm{RD}}/\mathrm{eV}) \geq 19$. We use two complementary representations. Figure 7.2a shows the fractional deviation of each reconstructed energy from the true Monte Carlo energy E_{MC}. Figure 7.2b shows the inconsistency between the two reconstructions in units of their combined statistical uncertainty. The RD reconstruction exhibits a weak and approximately Gaussian bias. For protons, the mean fractional deviation is at the level of $\sim$ 9–10% (depending on the hadronic interaction model), while for iron it is $\sim$ 15%, as indicated by the fitted means in the left panel of Figure 7.2a. In contrast, the WCD reconstruction shows a substantially larger positive shift: for protons the mean bias is $\sim$ 41–46%, while for iron it is $\sim$ 25% (right panel of Figure 7.2a). This behavior is consistent with the large positive shifts reported in [89] and reflects the known muon-related discrepancies of simulations in the inclined regime.



(a) Fractional reconstruction error $1 - E_{\mathrm{rec}}/E_{\mathrm{MC}}$ for RD (left) and WCD (right).



(b) Energy-pull distributions $(E_{\mathrm{RD}} - E_{\mathrm{WCD}})/\sqrt{\sigma_E^2(\mathrm{RD}) + \sigma_E^2(\mathrm{WCD})}$ for proton (left) and photon (right).

Figure 7.2: (a) Reconstruction deviation from true air-shower energy. (b) Energy-pull using Gaussian fits quantifying the deviation of WCD energy from the RD energy for each air shower.

The WCD–RD pull distribution for protons is well described by a Gaussian fit with mean $\mu = 2.7622$ and width $\sigma = 1.68$ (see Figure 7.2b). This indicates a systematic offset between the two energy estimators in units of their quoted statistical uncertainties, and points to a combination of WCD energy bias and/or misestimated uncertainties in the WCD and RD reconstructions.

Photon pulls exhibit an even larger mean shift, with a Gaussian fit yielding $\mu = 7.6497$ and $\sigma = 2.83$. However, as noted below, the photon distribution is non-Gaussian; these fit parameters serve only as summary statistics. Nevertheless, the large positive mean confirms that the standard WCD calibration severely underestimates the photon energy, while the RD reconstruction remains comparatively stable.

While hadronic reconstruction deviations remain approximately Gaussian, photon-induced air showers exhibit non-Gaussian sub-structures in Figure 7.2a and Figure 7.2b. Specifically, the photon pull distribution contains a secondary, lower-bias component where events are reconstructed in a more hadron-like manner—a subset characterized by increased relative RD uncertainty $\epsilon = \sigma_E^{\mathrm{RD}}/E_{\mathrm{RD}}$. Consequently, standard Gaussian fits serve as useful summary parameters but fail to fully capture the tails and the inherent complexities of the photon distributions.

A reliable characterization of these systematic shifts and pull widths is crucial, because any residual bias or misestimated uncertainty directly impacts the interpretation of WCD–RD energy inconsistency as a separation observable. Large radio-energy uncertainties can broaden the distributions and increase overlap, and therefore must be controlled before defining a robust cut-based photon-candidate selection.

7.3 Feature Engineering and Detrending

The separation strategy relies on a fundamental physical principle: the energy deposition difference between electromagnetic (measured by RD; E_{RD}) and muonic components (measured by WCD; E_{WCD}) substantially exceeds systematic biases and statistical uncertainties in the measurements. To exploit this principle, we construct the Relative Energy Difference (δ) as our primary observable:

$$\delta \;=\; \frac{E_{\mathrm{WCD}} \;-\; E_{\mathrm{RD}}}{E_{\mathrm{RD}}}\,.$$

This observable captures the relative energy mismatch while providing an approximate energy-scale independence through normalization by E_{RD}, enabling direct comparison across events of different energies. Figure 7.3 presents the δ distribution and its residual energy dependence for several primary particle types after applying the analysis threshold $\log_{10}(E_{\mathrm{RD}}/\mathrm{eV}) \geq 19$. This cut is performed on the measured energy (E_{RD}) to ensure that hadronic backgrounds with reconstructed energies similar to the signal are correctly retained for analysis. A negative correlation with energy is observed for all species, although its strength depends on the primary and on the hadronic interaction model. Here, r denotes the Pearson correlation coefficient, and the quoted p-values correspond to a two-sided test of the null hypothesis $\rho = 0$ (no linear correlation), i.e. the probability to obtain $|r| \geq |r_{\mathrm{obs}}|$ for a sample of size N if the true correlation is zero. For proton (SIBYLL 2.3d) the correlation is weak, $r = -0.086$ with $N = 1998$ ($p = 1.18 \times 10^{-4}$), while for proton (QGSJETII-04) and iron (SIBYLL 2.3d) the correlation is stronger, $r = -0.279$ ($N = 522$, $p = 8.09 \times 10^{-11}$) and $r = -0.258$ ($N = 623$, $p = 5.97 \times 10^{-11}$), respectively. Photons also show a negative correlation, $r = -0.260$ with $N = 231$ ($p = 6.16 \times 10^{-5}$).

Detrending of δ

To reduce residual energy-dependent trends, we detrend the relative energy difference δ as a function of $x = \log_{10}(E_{\mathrm{RD}}/\mathrm{eV})$. As shown in Figure 7.3, δ still exhibits an energy dependence after the event-quality selection, mainly induced by trigger/selection effects and reconstruction systematics. To avoid bias from the photon sample, the detrending function is determined using background only and then applied unchanged to both background and signal. We define the detrended variable by subtracting the fitted background trend,

$$\delta' = \delta - \delta_{\mathrm{fit}}(x)\,.$$

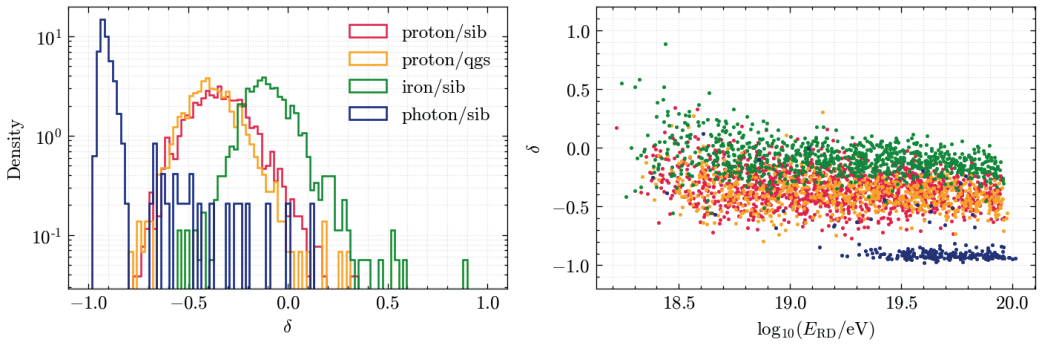


Figure 7.3: For background and signal primaries before the final reconstructed-energy threshold. (Left) distributions and energy dependence of the raw relative energy difference δ. (Right) δ as a function of the reconstructed radio energy $\log_{10}(E_{\mathrm{RD}}/\mathrm{eV})$.

The detrending procedure is:

- We bin x in steps of 0.2 over the selected range $19 \leq x < 20$, where $x = \log_{10}(E_{\mathrm{RD}}/\mathrm{eV})$, and compute, per bin, the background mean $\langle\delta\rangle$, the standard deviation σ_δ, and the uncertainty on the binned mean
$$\sigma_{\langle\delta\rangle} = \frac{\sigma_\delta}{\sqrt{N}},$$
where N is the number of background events in that bin.
- We retain only bins with at least $N \geq 10$ background events and finite $\langle\delta\rangle$ and $\sigma_{\langle\delta\rangle}$.
- We fit a straight line to the retained binned background means,
$$\delta_{\mathrm{fit}}(x) = a\,x + b,$$
using weights $w = 1/\sigma^2_{\langle\delta\rangle}$ so that bins with a better-determined mean contribute more.
- To prevent a small number of outlying bins from pulling the fit, we apply an iterative 2σ rejection on the bin residuals. After each fit, bins farther than 2σ from the median residual are removed, where σ is estimated robustly from the median absolute deviation (MAD). We refit and repeat until the set of retained bins no longer changes, or after at most 20 iterations.
- The final fitted trend is subtracted from all events (background and signal) to obtain δ'.

This procedure removes energy-dependent systematics captured by the background while keeping the signal sample independent of the detrending.
The fit uses 6 eligible background bins, of which 5 are retained after the iterative 2σ MAD-based rejection. The best-fit parameters are $a = -0.04998 \pm 0.018$ and $b = 0.669 \pm 0.35$, with a reduced chi-square of $\chi^2_\nu = 1.92$ (dof= 3). The residual scatter around the trend is RMSE $= 8.42 \times 10^{-3}$, while the median-absolute-deviation estimate gives a comparable value of $\sigma_{\mathrm{MAD}} = 7.67 \times 10^{-3}$. The fitted slope is only marginally significant ($p_{\mathrm{slope}} = 7.09 \times 10^{-2}$).

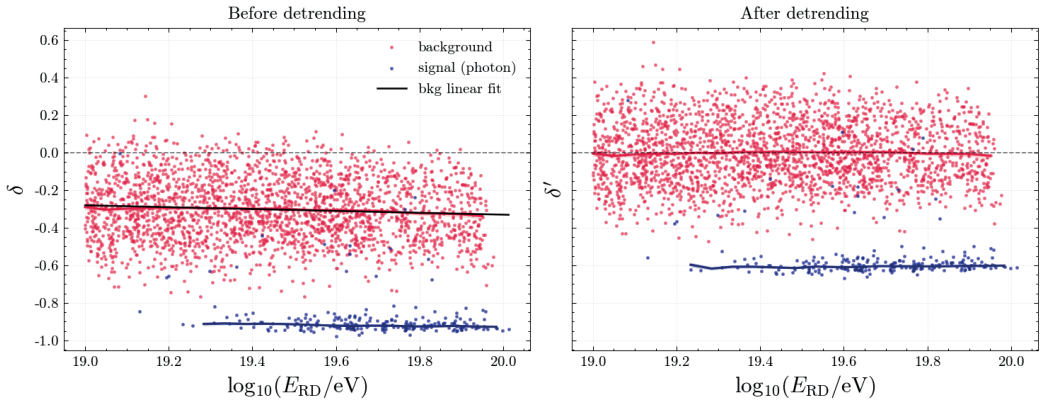


Figure 7.4: Effect of detrending on the energy dependence of δ for background (red) and signal (blue). The left panel shows δ as a function of $x = \log_{10}(E_{\mathrm{RD}}/\mathrm{eV})$, together with the fitted background trend (black). The right panel shows the detrended residuals $\delta' = \delta - \delta_{\mathrm{fit}}(x)$ obtained by subtracting the same trend from all events.

After subtracting δ_{fit}, the residual δ' shows no significant energy dependence in the background sample. The Pearson coefficient, computed on event-level δ', becomes $r = +0.005$ with a 16–84% bootstrap interval $[-0.015, +0.025]$ ($p = 8.09 \times 10^{-1}$, $N = 2303$), consistent with zero within uncertainties. For photons, the residual correlation is reduced but remains nonzero, $r = -0.156$ with interval $[-0.243, -0.053]$ ($p = 1.76 \times 10^{-2}$, $N = 230$), indicating that a mild energy dependence persists in the signal population and is not fully captured by a background-only detrending.

Having removed the background energy trend, we use δ' as the primary separation observable in the following sections.

7.4 Photon–hadron separation with δ'

In rare-signal searches, high purity is often prioritized over total efficiency. Consequently, any identified photon candidate must be supported by a strong and transparent separation criterion. The detrended energy mismatch δ' provides such a criterion: it is physically motivated, one-dimensional, and directly applicable to data. Figure 7.5 shows the normalised δ' distributions for background and photons after the full quality selection and the background-only detrending described in Section 7.3. Photons populate a narrow peak around $\delta' \simeq -0.6$, while the hadronic background is broadly centred around $\delta' \simeq 0$. The logarithmic-scale panel highlights the distribution tails: the available background Monte Carlo does not visibly populate the extreme left tail close to the photon peak. The finite size of the simulated background sample means that the absence of events in the photon-like region cannot, by itself, exclude leakage at the rejection levels required for discovery searches ($\sim 10^{-4}$). To probe beyond this statistical limit, we parameterise the background δ' distribution and extrapolate it into the signal region. Here we use a Gaussian model as a simple approximation to obtain a rough estimate of the residual

leakage below the statistical reach of the finite Monte Carlo sample. Since this choice is not unique and does not specifically model the far-left tail, the extrapolated leakage should be interpreted as a model-dependent estimate rather than a directly validated prediction. Because such far-tail estimates are sensitive to the fitted parameters, especially the width, we propagate the fit uncertainty with a bootstrap procedure using resamplings of the trimmed background sample (central 1–99% range). The resulting extrapolation provides a rough, model-dependent estimate of the residual contamination. For simplicity, we treat the full background as a single pooled population; a more complete treatment would instead separate the different simulated cosmic-ray composition components, fit each component individually, fit each with an independent Gaussian component, and combine them as a mixture energy depedent model.

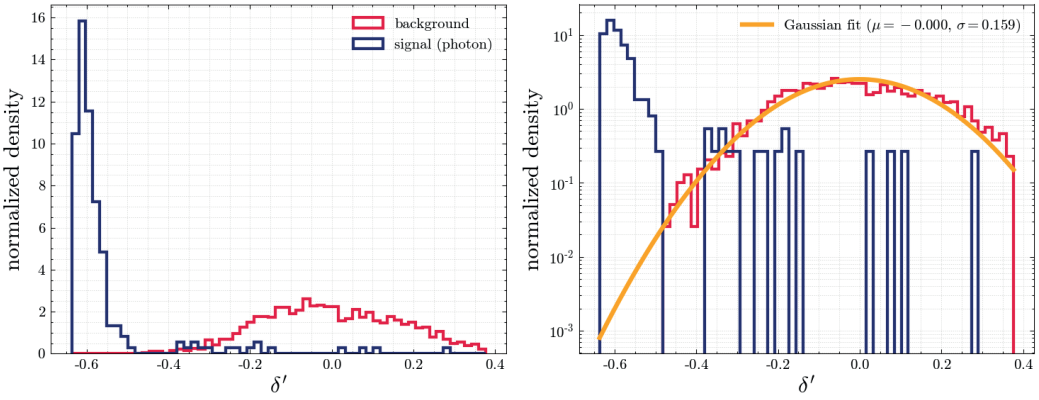


Figure 7.5: Normalised δ' distributions for background (red) and photons (blue). Left: linear y scale. Right: logarithmic y scale highlighting the tails. The orange curve shows a Gaussian fit to the background distribution, used to extrapolate the far-left tail.

Quantitatively, the photon distribution has a robust core width $\sigma_{68,\gamma} = (P_{84} - P_{16})/2 = 3.38 \times 10^{-2}$ around a median of median$(\delta') = -0.6017$. The background distribution is broader, with $\sigma_{68,\mathrm{bkg}} = 1.77 \times 10^{-1}$ and median$(\delta') = -0.0103$. The median separation between the two populations is therefore $\Delta = 0.5914$, corresponding to $3.28\,\sigma$ in units of the combined core widths. We define photon-like events by a simple threshold,

$$\delta' < \delta'_{\mathrm{cut}}\,,$$

and quantify the performance by the photon efficiency and background leakage,

$$\epsilon_\gamma(\delta'_{\mathrm{cut}}) = \frac{N_\gamma(\delta' < \delta'_{\mathrm{cut}})}{N_\gamma^{\mathrm{tot}}}, \qquad \epsilon_{\mathrm{bkg}}(\delta'_{\mathrm{cut}}) = \frac{N_{\mathrm{bkg}}(\delta' < \delta'_{\mathrm{cut}})}{N_{\mathrm{bkg}}^{\mathrm{tot}}}\,.$$

To achieve a target photon selection efficiency of approximately 80%, we define a reference working point at $\delta'_{\mathrm{cut}} = -0.5653$. For the reconstructed-radio-energy-selected sample used in this discrimination study, this selection retains 185/230 photons, corresponding to $\epsilon_\gamma = 0.80$, and yields zero observed background leakage in the available simulated sample (0/2303).

Since the finite background sample cannot directly constrain leakage at the level relevant for discovery searches, we estimate the residual contamination by extrapolating a Gaussian fit to the background tail. At the reference cut, this gives

$$P(\delta' < \delta'_{\text{cut}}) = 1.83 \times 10^{-4},$$

with a 68% confidence interval of $[1.53,\ 2.14]\times10^{-4}$ and a 95% confidence interval of $[1.29,\ 2.52]\times 10^{-4}$. This corresponds to an expected leaked background of $N_{\text{leak}} = 0.421$ events in the present simulated sample.

The dependence of ϵ_γ and ϵ_{bkg} on δ'_{cut} is shown explicitly in Figure 7.6, and alternative operating points are discussed below. For comparison, a discovery-level working point defined by the Gaussian-tail requirement $P(\delta' < \delta'_{\text{cut}}) = 10^{-4}$ corresponds to $\delta'_{\text{cut}} = -0.58996$ and yields a photon efficiency of $\epsilon_\gamma = 143/230 = 0.622$ in the present photon sample. Alternatively, requiring an expected leakage of 0.5 background events in the present simulated sample ($N_{\text{bkg}} = 2303$) corresponds, under the same Gaussian-tail model, to $\delta'_{\text{cut}} = -0.55814$, for which the photon efficiency is $\epsilon_\gamma = 192/230 = 0.835$. These examples illustrate the trade-off between photon efficiency and the rejection level implied by the tail model.

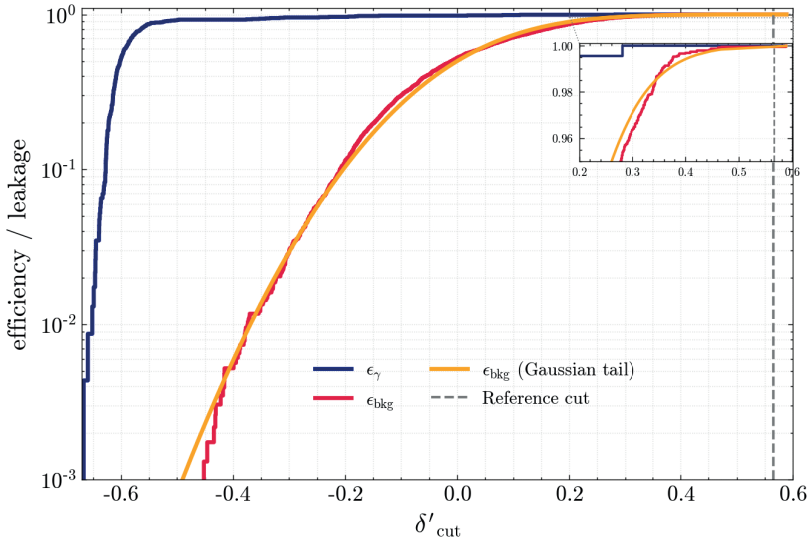


Figure 7.6: Photon efficiency ϵ_γ (blue) and background leakage ϵ_{bkg} (red) as a function of the threshold δ'_{cut} for the selection rule $\delta' < \delta'_{\text{cut}}$, shown on a logarithmic scale. The orange curve represents the Gaussian-tail extrapolation of the background leakage ($\mu \approx -1.83 \times 10^{-4}$, $\sigma \approx 0.15858$). The vertical dashed line indicates the reference cut at $\delta'_{\text{cut}} = -0.5653$. The inset provides a linear-scale magnification of the high-efficiency region between 0.95 and 1.0.

For the most restrictive cuts, the leakage inferred directly from the finite simulated sample is limited by statistics: in the present background set ($N^{\text{tot}}_{\text{bkg}} = 2303$), the first nonzero leakage occurs at $\epsilon_\gamma \simeq 0.926$, where $\epsilon_{\text{bkg}} = 1/2303 = 4.34 \times 10^{-4}$, while tighter cuts yield zero observed background events and therefore only upper limits of order $\mathcal{O}(1/N_{\text{bkg}})$.

To illustrate the behaviour beyond this statistical reach, we overlay the Gaussian-tail extrapolation

$$\epsilon_{\mathrm{bkg}}^{\mathrm{Gauss}}(\delta'_{\mathrm{cut}}) = \Phi\left(\frac{\delta'_{\mathrm{cut}} - \mu}{\sigma}\right),$$

with $(\mu, \sigma) = (-1.83 \times 10^{-4}, 0.15858)$ obtained from a fit to the central 1–99% range of the background δ' values.

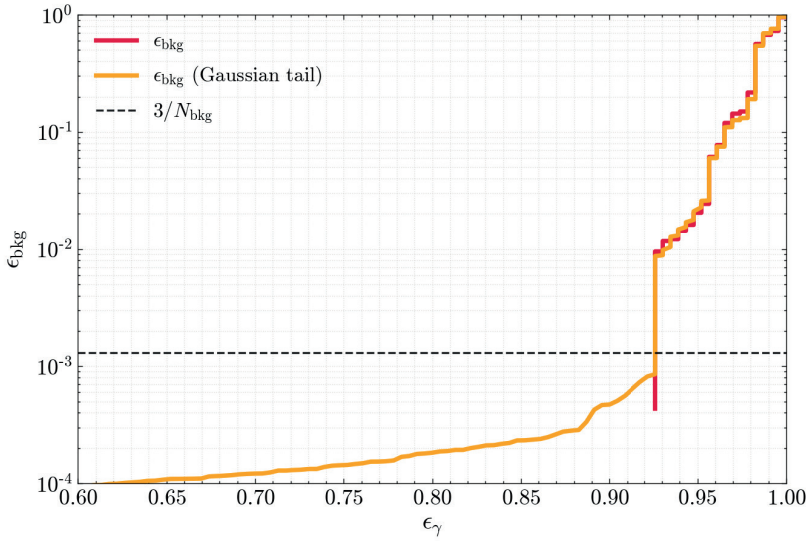


Figure 7.7: Background leakage ϵ_{bkg} as a function of photon efficiency ϵ_γ obtained by scanning the cut δ'_{cut} for the selection rule $\delta' < \delta'_{\mathrm{cut}}$, after applying the energy cut $\log_{10}(E_{\mathrm{RD}}/\mathrm{eV}) \geq 19$. The red curve shows the empirical leakage in the finite simulated background sample ($N_{\mathrm{bkg}} = 2303$). The orange curve shows a Gaussian-tail extrapolation of the background δ' distribution (single-Gaussian fit to the central 1–99% range). The dashed line marks the finite-sample reference level $\epsilon_{\mathrm{bkg}} \simeq 3/N_{\mathrm{bkg}} = 1.30 \times 10^{-3}$ for $N_{\mathrm{bkg}} = 2303$, corresponding approximately to the 95% upper-limit scale for zero observed background events.

Figure 7.7 visualises the trade-off between photon efficiency and background leakage for a simple global cut on δ'. The empirical leakage curve (red) is limited by the finite simulated background statistics: for sufficiently restrictive cuts, zero background events are observed, so the leakage can only be constrained at the level $\mathcal{O}(1/N_{\mathrm{bkg}})$ rather than at the 10^{-4} level relevant for discovery searches. In the present background sample ($N_{\mathrm{bkg}} = 2303$), the first nonzero leakage appears at $\epsilon_\gamma \simeq 0.926$ with $\epsilon_{\mathrm{bkg}} = 1/2303 = 4.34 \times 10^{-4}$, demonstrating that the simulation alone does not probe the far tail below this scale.

The quoted bootstrap intervals quantify the statistical uncertainty of this extrapolation within the Gaussian assumption, but they do not capture possible non-Gaussian structure in the true far-left tail. The inferred leakage should therefore be interpreted as a model-dependent estimate rather than a directly validated background level.

In summary, the detrended observable δ' provides a simple and physically motivated discriminator that yields strong photon–hadron separation in the inclined hybrid simulations. Because

the selection is transparent and relies on a single cut, it is straightforward to validate and to apply to data. While the finite background Monte Carlo sample limits direct validation of the most extreme rejection levels, the overall separation performance remains strong in the simulated sample. We therefore adopt δ' as the primary separation observable in the subsequent analysis.

Chapter 8

Towards the search for ultra-high energy photons in data

This chapter establishes a workflow for conducting a photon search using the hybrid WCD+RD detector configuration at the Pierre Auger Observatory. We begin by quantifying the detector exposure through accounting of operational periods and data quality, then compute expected sensitivity limits based on the photon–hadron separation criterion developed in previous chapters. Finally, we propose a burn-sample strategy to validate the separation performance on real data while preserving the integrity of an unbiased photon search.

8.1 The Hybrid Exposure

The exposure calculation is an important aspect of any rare-event search, translating detector geometry and operational history into the effective observation time and area required for flux determination. For the hybrid of WCD and RD configuration, this calculation must account for the time-varying detector status and quality.

Data Processing

This section describes the data processing steps used to compute the hybrid WCD+RD exposure. The analysis covers the data-taking period from 1 December 2024 to 31 October 2025. During this time, the readout of the RD was operated in slave-mode trigger, recording data only when the corresponding WCD station registered an event.
The exposure calculation requires accurate knowledge of the operational status of each detector station at all times. Our processing pipeline integrates three primary data sources:

1. Metadata file (see Chapter 3.3) listing WCD T3 events together with the presence/absence of the corresponding RD readout, from which we infer the time-dependent RD station availability.

2. Detailed bad period records identifying time intervals when specific stations were offline or their data quality was compromised.

3. The WCD stations list specifying the physical positions in the array.

These inputs are combined to produce timestamped records of the active detector configuration—that is, the subset of stations suitable for physics analysis at each moment. The processing pipeline begins by extracting the operational RD instances for all main-grid (SD-1500) detector stations, including trigger timestamp, station ID, northing, and easting.
The trigger timestamps in the WCD T3 metadata table (timestamp, station ID, northing, easting) are converted to UTC and then rounded to the nearest hour to enable temporal aggregation. This is sufficient for the exposure calculation, since sub-hour variations in detector status have a negligible impact on the final result. Subsequently, a complete hourly time series is constructed by aggregating the station data for each hour. This results in a continuous record where each hour is associated with a list of IDs of operational stations.

Bad Period Assessment

An operational WCD tank doesn't guarantee the quality of RD data; hence, additional data quality information is needed to exclude periods when individual stations were known to be malfunctioning or producing unreliable data. This information is encoded in XML data described in the bad period page described in Chapter 3.4. These files specify time intervals during which particular RD stations should be excluded from analysis, along with the reasons for exclusion (e.g., hardware failures, calibration issues, or data acquisition problems).
We parse the XML file for the bad period, extracting the station ID and the start and end times of the bad period. Then retain only those bad period records that overlap with our analysis time window between $t_{\text{start}} = 1$ Dec 2024 and $t_{\text{end}} = 31$ Oct 2025.
A critical aspect of the bad-period handling is how the start and end timestamps of each station-level interval are mapped onto the hourly grid used in the exposure calculation. Because station activity is evaluated hour by hour, interval boundaries that fall within an hour must be assigned to hourly timestamps. To motivate the boundary convention, it is useful to first examine the distribution of bad-period durations.

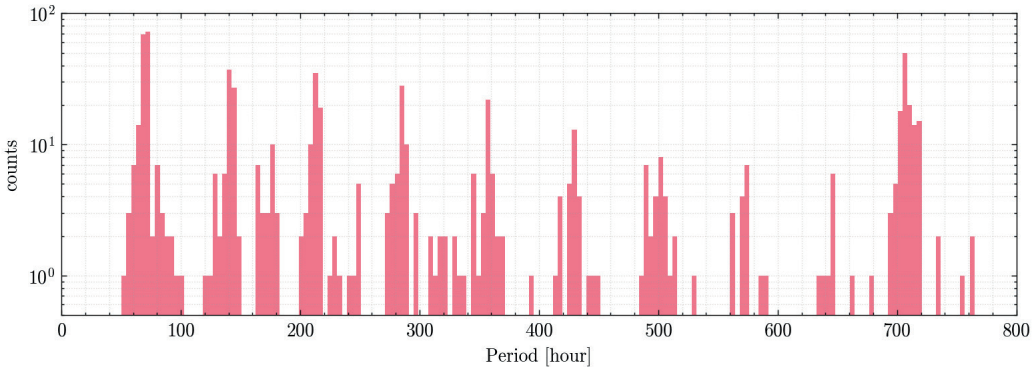


Figure 8.1: Distribution of station bad-period durations before mapping the interval boundaries onto the hourly grid.

As shown in Figure 8.1, most bad periods are much longer than one hour and cluster at preferred multi-hour or multi-day values. This indicates that the precise sub-hour boundary convention is not expected to be a leading effect in the exposure calculation. The clustered structure also reflects the way the bad periods are constructed from coarse predefined segments and later merged using fixed gap parameters.

We apply a 15-minute tolerance rule to the start and end timestamps. A comparison with a 30-minute tolerance showed that the resulting hourly active-station counts differed at only 23 of the 646 timestamps considered ($\sim 3.6\%$). The relative difference had mean -0.05%, median 0.00%, and standard deviation 0.49%, with minimum and maximum values of -6.97% and $+1.05\%$, respectively. This confirms that the precise sub-hour tolerance is not a leading effect in the exposure calculation.

If a bad period begins within the first 15 minutes of an hour, the start time is rounded down to the hour boundary (floor); otherwise it is rounded up to the next hour (ceiling). The same rule is applied to the end time[1].

After this hourly mapping, station status is determined for each hour, and the resulting hourly active-station lists are later aggregated into coarser 12-hour exposure bins, as described in section 8.1.

The bad-period durations cluster at a set of discrete values. This structure is expected from the bad-period production procedure rather than from a detector-specific physical timescale. The station checks are evaluated on predefined time segments ("N-day" segments and monthly segments), and when a segment is flagged, the full segment interval is written as a candidate bad period. These candidate intervals are then merged using fixed gap parameters defined in hours, which correspond to day-scale values (e.g., 4 days for DuplicateTraces, 5 days for LowGain, 7–14 days for misalignment handling, and 14 days for BrokenLNA). In the final merging step, short gaps of less than 10-hours between neighboring bad-period candidates are bridged, so that they are treated as one continuous interval. As a result, the production imposes preferred durations, which appear as the pronounced peaks in the distribution, even before any additional rounding is applied in the exposure pipeline. With the nominal operational station lists and the bad-period records, we build the final active-station list by excluding stations that are flagged as bad at a given time. For each hourly timestamp t_i, we start from the set of station IDs reported as operational at t_i and remove any station whose bad-period interval (after applying the threshold rounding) covers t_i. The remaining station IDs define the active stations for that hour.

The resulting record contains the list of active stations per hour. Figure 8.2 shows the total number of active hours within the used period after ingesting the bad periods per station. Most stations exhibit high availability, with more than 7,000 active hours, corresponding to uptime fractions exceeding 85% over the considered 11 months period. However, scattered

[1] A more conservative choice would be to declare an entire hour as bad if it is partially covered by a bad period. We do not adopt this rule because the bad-period intervals are produced on coarse predefined segments and merged using fixed gap parameters; an "any-overlap" criterion on an hourly grid would therefore tend to over-mask stations and bias the exposure low.

stations show significantly reduced uptime. Those indicating 1000–4000 active hours are predominantly located at the array boundaries or in specific clusters in the East[2] or Northeastern regions, where the strength of the 67-MHz TV line affects the performance of the antenna-alignment bad-period check (sections 3.4 and 3.6). A possible future improvement would be to make the alignment bad-period check more robust against persistent local 67-MHz interference. This is motivated by the fact that the current bad-period production already includes a region-dependent 50-MHz veto to mitigate known TV-induced biases in the alignment observable, suggesting that a more interference-aware treatment could reduce unnecessary masking in these regions.

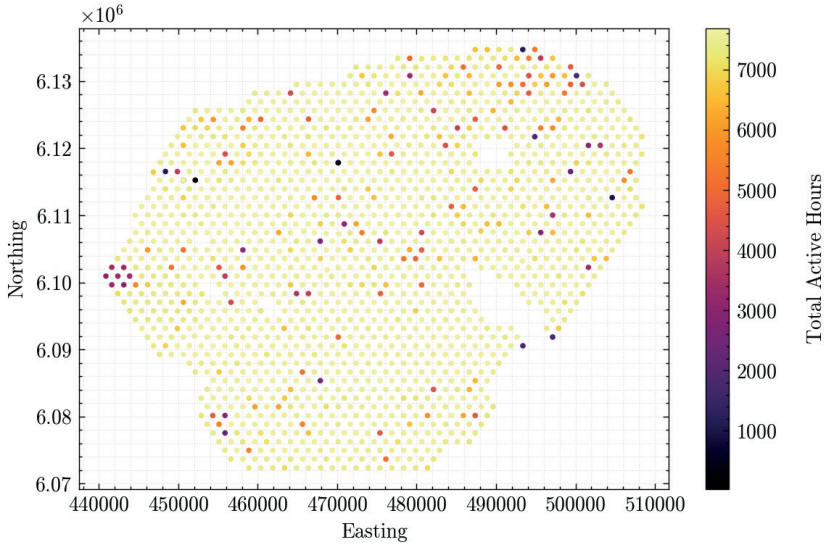


Figure 8.2: Spatial distribution of RD stations. Station positions are shown in UTM coordinates (Easting, Northing in meters). The color scale indicates the cumulative number of active hours from December 2024 to October 2025.

Temporal Aggregation to 12-Hour Bins

For the hybrid WCD+RD exposure calculations and the associated systematic uncertainty studies, we require coarser time binning than the hourly resolution because the exposure depends on the time-integrated array configuration, and finer bins would mainly increase sensitivity to short-lived station glitches and hour-boundary artefacts without materially improving the integrated exposure estimate. We therefore aggregate the hourly active station lists into 12-hour bins aligned to midnight UTC. For every 12 hours starting at time T_j (where $T_j = 0, 12, 24, \ldots$ hours since the analysis start), we compute the union of all active station IDs that appeared in any hour within that period. This union operation ensures that a station is included in the 12-hour bin if it was operational for any portion of that period. Consequently, short outages

[2] A 7-station configuration forms what is called a *test hexagon*, where the radio self-trigger is tested, leading to reduced readout into real event participation.

within a 12-hour interval can be smoothed over at the aggregated level, so the 12-hour representation is not intended to resolve fine hour-scale detector downtime but rather to capture the time-integrated array configuration relevant for the exposure calculation.

Number of Active Hexagons

The 12-hour granularity serves as the input for the active hexagon counting, which relies on the surface detector array of the Pierre Auger Observatory. The array is configured in a triangular grid, with detectors positioned at the vertices of equilateral triangles. This triangular layout can be described as a hexagonal configuration for event reconstruction, where each effective hexagon consists of a central station surrounded by six neighboring stations [126].
Active hexagons represent the fundamental building blocks of the detector's aperture, as the hexagonal geometry defines the minimum unit that can provide full angular coverage for event reconstruction [127]. The count of these active hexagons determines the array's time-dependent effective area.

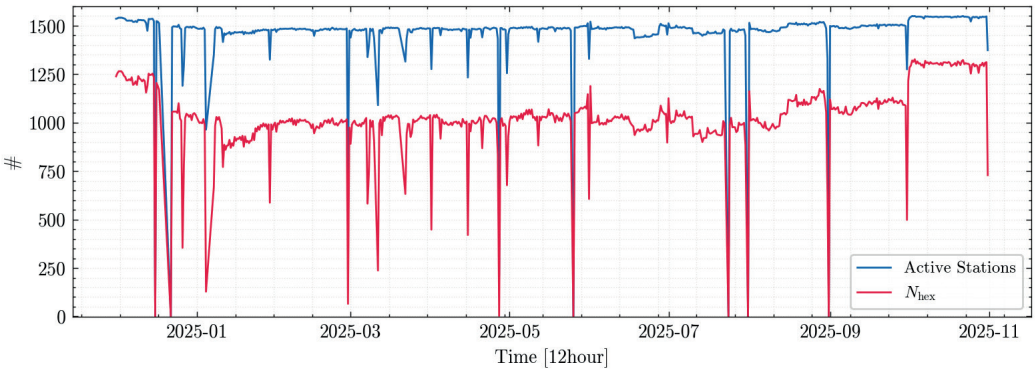


Figure 8.3: Time evolution of the number of active stations (blue) and active hexagons (red) over the analysis period.

We construct a neighbor map that identifies the six nearest neighbors of each station. In an ideal hexagonal lattice, each interior station has exactly six neighbors at a distance of 1500 meters. In practice, some deviations from the ideal geometry occur due to non-deployment areas and array edges.
Hence, to find the neighboring stations:

1. Compute the pairwise Euclidean distance matrix for all stations:

$$d_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$$

2. For each station i, identify all stations j within a distance tolerance of the nominal SD-1500 grid spacing of 1500 m:

$$1500\,\mathrm{m} - \delta < d_{ij} < 1500\,\mathrm{m} + \delta$$

where we use $\delta = 300\,\mathrm{m}$ to account for geometric irregularities while excluding false matches.

3. Select the six closest stations satisfying this criterion as the neighbors of station i.

The choice of $\delta = 300\,\mathrm{m}$ represents a balance between accommodating real geometric variations and maintaining strict hexagonal topology. We define a station as an active hexagon center if and only if:

- The station itself is in the active station list: $s \in \text{active_ids}(t)$
- All six of its neighbors are also in the active station list: $\{n_1, n_2, \ldots, n_6\} \subseteq \text{active_ids}(t)$

We compute the total number of hexagons (N_{hex}) for each 12-hour time bin, as shown in Figure 8.3. The time plateau from 2025-10-01 to 2025-11-01 may be taken as a representative benchmark for future detector performance. In this period, the median numbers of active stations and active hexagons were 1515 and 1151, respectively, and 98.4% of the 12-hour bins remained within ±5% of these median values, indicating a highly stable operational state. The plot displays the temporal variation in active stations, highlighting significant drops in both station and hexagon counts due to WCD operation. These fluctuations represent the intermittent availability of active stations and their corresponding hexagons, with maintenance or downtime periods leading to reduced operational units. The average number of active surface detector hexagons is around 1250 [128]. Due to regions where no RD antennas are installed, such as the eastern region (see Figure 2.1), along with periods affected by radio bad periods as shown in Figure 8.2, the number of combined WCD+RD active hexagons fluctuates between 1000 and 1200.

The number of active hexagons is directly related to the effective area of the detector. We define A_{cell} as the area of the elementary surface-detector unit of the SD-1500 array. Since one active hexagon is composed of six such cells, the geometric hexagon area of approximately 11.7 km^2 corresponds to

$$A_{\text{cell}} = \frac{A_{\text{hex}}}{6} \approx 1.95\ \text{km}^2. \tag{8.1}$$

The factor of six reflects the overlap of neighboring active hexagons in the array geometry [100]. This relation is approximate, particularly near the array boundaries where edge effects become relevant. The total number of active hexagons at any given time determines the detector's effective area, which is used in the exposure calculation for the photon search.

Exposure

Following the aperture derivation in Appendix B, the effective aperture in a zenith band $[\theta_1, \theta_2]$ can be written as

$$A_{\text{eff}}(E_i, t) = \pi\, A(t)\, \bar{\eta}_{\text{hyb}}(E_i) \left(\cos^2\theta_1 - \cos^2\theta_2\right), \tag{8.2}$$

where $\bar{\eta}_{\text{hyb}}(E_i)$ denotes the zenith-averaged hybrid detection efficiency in reconstructed WCD energy bin $E_i \equiv E_{\text{WCD},i}$ over the selected zenith interval. The efficiencies are therefore evaluated

in the same energy variable used for the real-data selection, namely E_{WCD}. In our case, the time-dependent area is obtained from the number of active hexagons,

$$A(t) = N_{\mathrm{hex}}(t)\, A_{\mathrm{cell}}, \tag{8.3}$$

where $A_{\mathrm{cell}} \approx 1.95$ km^2 is the effective area associated with one counted active hexagon in the hexagon-counting method, corresponding to one elementary SD-1500 cell, i.e. one sixth of the geometric hexagon area because of the sixfold overlap between neighboring hexagons. The zenith-averaged hybrid efficiency factorizes as

$$\bar{\eta}_{\mathrm{hyb}}(E_i) = \eta_{\mathrm{WCD}}(E_i)\, \eta_{\mathrm{RD}}(E_i)\, \eta_{\mathrm{quality}}(E_i), \tag{8.4}$$

where the three terms correspond to WCD reconstruction, conditional RD reconstruction, and the combined quality selections (Sections 5.3 and 6.2). All component efficiencies entering this factorization are evaluated for the selected zenith interval and are therefore not differential in θ. The exposure is then computed as the time integral of the effective aperture,

$$\mathcal{E}(E_i) = \int_{t_0}^{t_1} A_{\mathrm{eff}}(E_i, t)\, \mathrm{d}t. \tag{8.5}$$

The detection efficiency is implicitly accounted for within $A_{\mathrm{eff}}(E_i, t)$, which includes the effects of both WCD and RD efficiencies as well as data quality cuts. The analysis is restricted to the full WCD efficiency domain at energies $E_i \geq 10^{18.6}$ eV, where $\eta_{\mathrm{WCD}}(E_i) = 1.0$ over seven $\log_{10}(E_{\mathrm{WCD}}/\mathrm{eV})$ bins. This reconstructed-WCD-energy range corresponds to the full-efficiency domain of the WCD array for the selected zenith interval, so that the remaining efficiency losses are dominated by RD reconstruction and subsequent hybrid quality cuts rather than by WCD-limited detection.
Two distinct zenith angle ranges are analyzed:

1. $65° \leq \theta_{\mathrm{WCD}} \leq 80°$: This is the subset regime where the WCD full efficiency space intersects with the RD sensitivity.
   In this case, $\eta_{\mathrm{RD}}(E) = [\,0.710 \pm 0.034,\ 0.789 \pm 0.054,\ 0.776 \pm 0.076,\ 0.859 \pm 0.132,\ 0.872 \pm 0.186,\ 0.778 \pm 0.392,\ 1.000 \pm 0.707]$ derived from Figure 5.8 and Figure 5.12. After applying the quality cuts in section 6.2, the resulting per-bin quality-cut efficiencies are given in Table 8.1.
   Individual bins range from 31% to 67%, but the highest-energy bins contain only a few events with large statistical uncertainties. In the exposure calculation we use the measured values $\eta_{\mathrm{quality}}(E_i)$ from Table 8.1 on a bin-by-bin basis. The quality-cut efficiency in the range 65°–80° exhibits bin-to-bin variations that are dominated by a strong zenith-angle dependence of the cuts. Averaged over all energies, the survival fraction rises from ∼1% near 65°–68° to ∼76% near 77°–80°. Figure 8.4 summarizes this behaviour and clarifies how limited statistics at the highest energies enhance the observed scatter.
   The top-right panel breaks down the survival fraction after each quality cut. The first

Energy bin	N_{before}	N_{after}	$\eta_{\text{quality}}(E) \pm \sigma_{\text{binom}}$
18.60–18.80	555	221	0.40 ± 0.02
18.80–19.00	268	123	0.46 ± 0.03
19.00–19.20	125	57	0.46 ± 0.04
19.20–19.40	48	15	0.31 ± 0.07
19.40–19.60	27	8	0.30 ± 0.09
19.60–19.80	5	2	0.40 ± 0.22
19.80–20.00	3	2	0.67 ± 0.27

Table 8.1: Quality cuts efficiency per energy bin for $65° \leq \theta_{\text{WCD}} \leq 80°$.

step, $N_{\text{RD}} \geq 5$, removes a sizeable fraction of events in all bins, consistent with requiring a minimum number of radio stations for a reliable reconstruction. The next requirement, $\ln(\text{Cone}) \geq 1$, produces the second major reduction and dominates the additional loss in several bins. The final two selections, $\chi^2 < 3$ and $\Delta\Psi \leq 1$, act mainly as consistency checks and typically remove a smaller additional fraction.

The bottom-left panel shows η_{quality} in fine bins of (θ_{WCD}, E), making the zenith dependence explicit: efficiencies are very low near 65° and increase towards 77°–80°, where values of 60–90% are typical. The bottom-right panel shows the corresponding event counts; in the highest-energy bins several cells contain only a few events, which explains the appearance of extreme values (0% or 100%) and the larger bin-to-bin scatter in $\eta_{\text{quality}}(E)$. Within the available statistics, the dominant variation is with zenith angle, while no clear additional systematic dependence on energy is resolved.

2. $73° \leq \theta_{\text{WCD}} \leq 80°$: The regime where RD is almost fully efficient (reconstructs essentially all WCD events).

   In this case, $\eta_{\text{RD}}(E) \approx 1$ for all energy bins, while $\eta_{\text{quality}}(E)$ is obtained from Table 8.2.

Energy bin	N_{before}	N_{after}	$\eta_{\text{quality}}(E) \pm \sigma_{\text{binom}}$
18.60–18.80	261	183	0.70 ± 0.03
18.80–19.00	129	95	0.74 ± 0.04
19.00–19.20	61	42	0.69 ± 0.06
19.20–19.40	18	13	0.72 ± 0.11
19.40–19.60	10	4	0.40 ± 0.15
19.60–19.80	2	1	0.50 ± 0.35
19.80–20.00	1	1	1.00 ± 0.00

Table 8.2: Quality cuts efficiency per energy bin for $73° \leq \theta_{\text{WCD}} \leq 80°$.

Individual bins range from 40% to 100%, but the highest-energy bins contain only a few events with large statistical uncertainties. As for the broader zenith range, the exposure calculation uses the per-bin values $\eta_{\text{quality}}(E_i)$ from Table 8.2 rather than enforcing a single global efficiency.

Figure 8.5 shows the cumulative geometric exposure for different zenith angle ranges. This quantity is calculated by first computing the incremental geometric exposure in each time bin

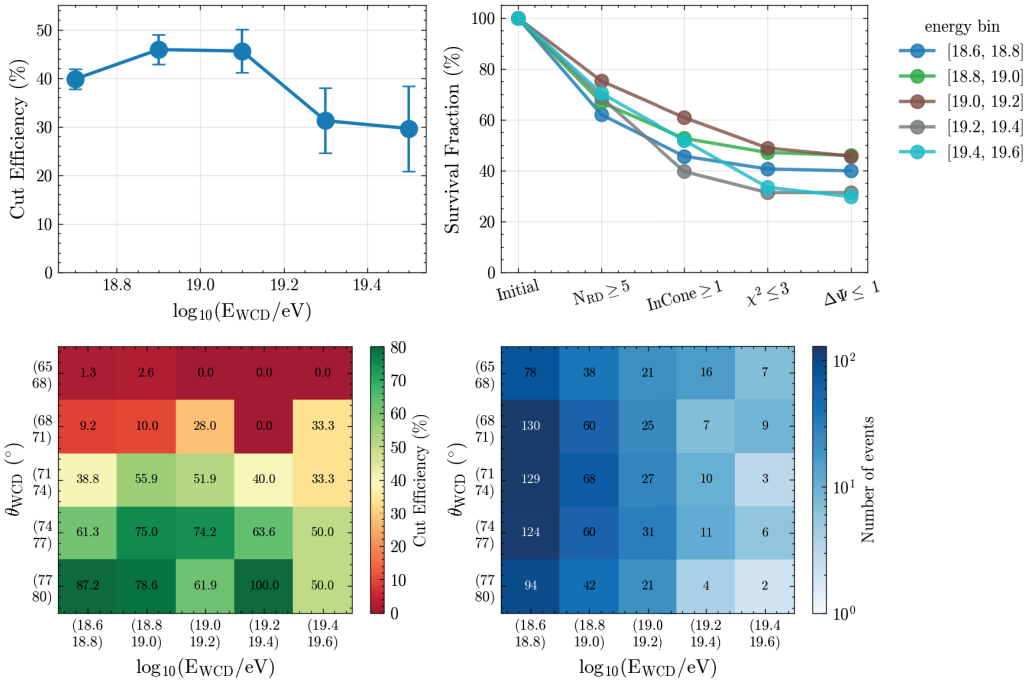


Figure 8.4: Dependence of the hybrid WCD+RD quality cuts on energy and zenith angle for $65° \leq \theta_{\mathrm{WCD}} \leq 80°$. Top-left: $\eta_{\mathrm{quality}}(E)$ per energy bin (as in Table 8.1). Top-right: survival fraction after successive quality requirements for selected energy bins. Bottom-left: η_{quality} in $(\theta_{\mathrm{WCD}}, E)$ sub-bins, showing a strong zenith dependence. Bottom-right: corresponding event counts per $(\theta_{\mathrm{WCD}}, E)$ sub-bin.

(in km^2 sr yr) as the product of the effective area and the time interval in years. The cumulative exposure is then obtained by summing these incremental contributions over time, while the total exposure for each range is the sum over all time bins. The cumulative exposure represents the total effective area accumulated over the entire analysis period. The broad zenith range (65°–80°) accumulates 819.7 km^2 sr yr, while the fully efficient range (73°–80°) reaches 305.5 km^2 sr yr. Exposure increases nearly linearly throughout the analysis period, reflecting the detector's stable operation.

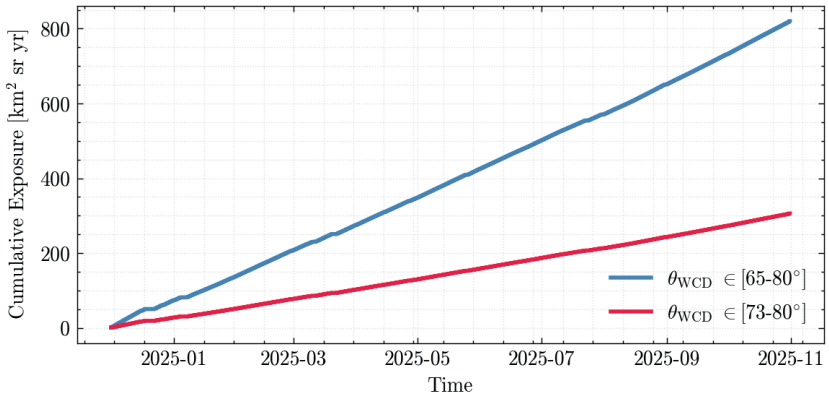


Figure 8.5: Cumulative geometric exposure (km^2 sr yr) as a function of time for two zenith angle (θ_{WCD}) ranges. The blue line corresponds to $\theta_{\mathrm{WCD}} \in [65°, 80°]$, while the red line corresponds to $\theta_{\mathrm{WCD}} \in [73°, 80°]$.

The differential exposure per reconstructed WCD-energy bin, shown in Figure 8.6, is the effective exposure in each reconstructed energy bin after applying the corresponding efficiency factors. The use of E_{WCD} at this stage is motivated by the fact that the real-data selection, the radio-detection efficiency η_{RD}, and the quality-cut efficiency η_{quality} are all evaluated in the same reconstructed energy variable. Thus, this figure characterises the operational exposure of the selected data sample.
Since the analysis is restricted to the full WCD-efficiency domain, $\eta_{\mathrm{WCD}} = 1$, the energy dependence of the differential exposure follows the zenith-averaged hybrid efficiency

$$\bar{\eta}_{\mathrm{hyb}}(E_i) \simeq \eta_{\mathrm{RD}}(E_i)\, \eta_{\mathrm{quality}}(E_i).$$

The horizontal steps in Figure 8.6 represent the value in each energy bin, while the markers are placed at the bin centres to guide the eye. The exposure is therefore an operational exposure expressed as a function of the reconstructed WCD energy, E_{WCD}, and is used to characterise the data selection and the behaviour of the hybrid efficiency. It is not used directly as the photon exposure entering the final flux-limit calculation. For the photon-flux upper limits, the exposure is re-evaluated using photon simulations in bins of true primary photon energy, E_{MC}. In addition, the final sensitivity estimate is restricted to the stable high-energy region, $\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) \geq 19.2$, to avoid the threshold-dominated low-energy part where the efficiency

turn-on is not accurately constrained.

For 65°–80°, the exposure is of order 200–300 $\mathrm{km^2\,sr\,yr}$ over most reconstructed-energy bins, with bin-to-bin variations driven mainly by the measured quality-cut efficiency. The last bin, 19.80–20.00, reaches 541 $\mathrm{km^2\,sr\,yr}$; however, this value is driven by the high measured η_{quality} in that bin and should be interpreted with care because the underlying statistics are very limited (Table 8.1). The mean effective exposure over the seven bins is 290 $\mathrm{km^2\,sr\,yr}$.

For 73°–80°, the exposure is approximately 200 $\mathrm{km^2\,sr\,yr}$ in the first four bins, followed by larger bin-to-bin variations at higher energies. These variations are dominated by the limited event counts entering the per-bin quality-efficiency estimate (Table 8.2). The mean effective exposure over the seven bins is 200 $\mathrm{km^2\,sr\,yr}$.

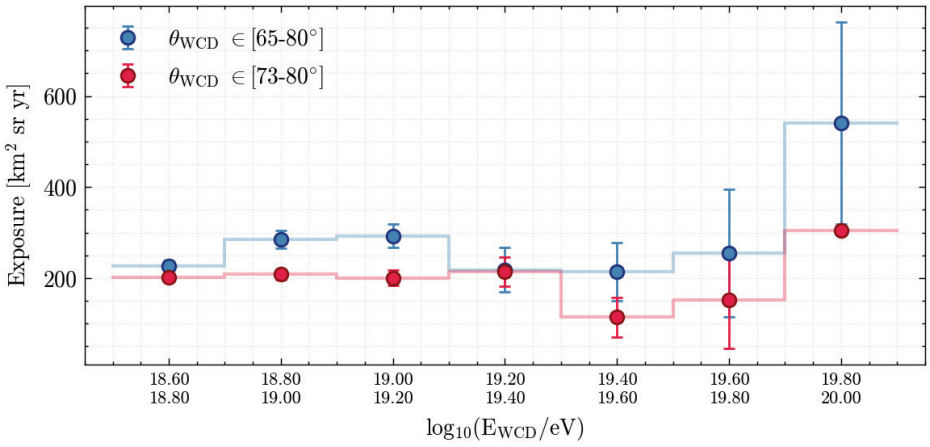


Figure 8.6: Differential effective exposure per reconstructed WCD-energy bin for $\theta_{\mathrm{WCD}} \in [65°, 80°]$ (blue) and $[73°, 80°]$ (red), using the per-bin quality-cut efficiencies from Table 8.1 and Table 8.2. Error bars indicate the statistical uncertainty obtained from the binomial errors on η_{quality} and propagated to the effective exposure. The low-energy bins are close to the analysis threshold and are therefore shown only to illustrate the behaviour of the operational exposure; the final photon-flux upper limits are derived later using true primary photon energy, E_{MC}, and only in the stable high-energy region.

The mean effective exposures are 290 $\mathrm{km^2\,sr\,yr}$ and 200 $\mathrm{km^2\,sr\,yr}$ for the 65°–80° and 73°–80° ranges, respectively. They differ by a factor of ~ 1.46 despite the 2.7× difference in geometric exposure, because the overall quality-cut efficiency is higher in the 73°–80° range, with total $\eta_{\mathrm{quality}} = 0.703$, than in the 65°–80° range, with total $\eta_{\mathrm{quality}} = 0.415$, partially compensating for the smaller geometric aperture.

8.2 Expected photon flux upper limits

The photon–hadron separation cut developed in section 7.4 is evaluated for simulated photons in bins of true primary energy, $E_i \equiv E_{\mathrm{MC},i}$. The exposure and expected upper-limit curves are

shown over the full analysis range,

$$18.6 \leq \log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) \leq 19.8,$$

to illustrate the low-energy turn-on of the analysis. However, the first three energy bins are close to the photon reconstruction and selection threshold. In this region, the effective exposure is extremely small and changes rapidly with energy, so the corresponding flux limits are dominated by the threshold behaviour rather than by a stable sensitivity estimate. For this reason, the low-energy bins are shown only for completeness and to indicate the analysis turn-on, while the quoted photon sensitivity is based only on

$$\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) \geq 19.2.$$

For the benchmark sensitivity estimate, Equation 8.5 is evaluated using a constant operated geometric area of approximately $2465\,\mathrm{km}^2$, derived from the stable late-time reference period. The calculation accounts for events within a zenith-angle range of 65° to 80°, corresponding to a solid angle of $0.466\,\mathrm{sr}$.

While the WCD array is effectively fully efficient for hadronic air showers in this energy range, the response to photon-induced air showers remains energy dependent. The photon exposure is therefore obtained by folding the operational exposure with the photon-specific WCD trigger efficiency, RD reconstruction efficiency, quality-selection efficiency, and photon-selection efficiency. The zenith-averaged total efficiency relevant for the photon search is defined as

$$\bar{\eta}_{\gamma}(E_i) = \eta_{\mathrm{WCD}}(E_i)\,\eta_{\mathrm{RD}}(E_i)\,\eta_{\mathrm{quality}}(E_i)\,\eta_{\mathrm{photon}}(E_i)\,, \tag{8.6}$$

where $E_i \equiv E_{\mathrm{MC},i}$ denotes a true primary photon-energy bin.

The resulting energy-dependent efficiencies are summarized in Table 8.3. The first three bins have very small but nonzero total efficiency and are retained in the exposure and upper-limit figures only to illustrate the low-energy turn-on of the analysis; they are not used for the quoted photon sensitivity limits.

$\log_{10}(E/\mathrm{eV})$	η_{WCD}	η_{RD}	η_{quality}	η_{photon}	$\bar{\eta}_{\gamma}$
18.60	0.006	1.000	0.400	0.001	2.4×10^{-6}
18.80	0.049	0.777	0.400	0.001	1.5×10^{-5}
19.00	0.148	0.363	0.400	0.001	2.1×10^{-5}
19.20	0.423	0.351	0.400	0.643	0.0382
19.40	0.788	0.538	0.400	0.905	0.1535
19.60	0.956	0.666	0.400	0.809	0.2060
19.80	0.997	0.800	0.400	0.782	0.2495

Table 8.3: Energy-dependent efficiencies entering the photon-exposure calculation. The first three bins are threshold dominated and are shown only to illustrate the low-energy turn-on; they are not used for the quoted photon sensitivity limits.

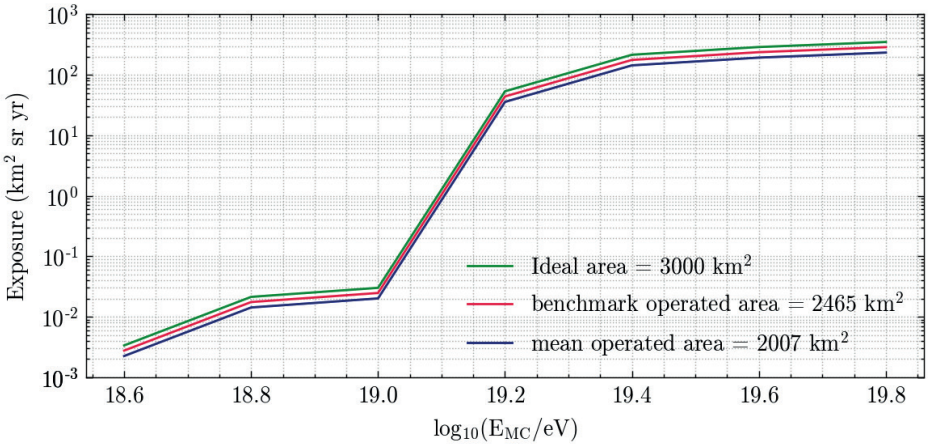


Figure 8.7: Expected photon exposure as a function of true primary photon energy, E_{MC}, in the zenith-angle range $65^\circ \leq \theta \leq 80^\circ$. The green curve corresponds to the ideal case with a constant geometrical area of $3000\,\mathrm{km}^2$, the red curve to the benchmark operated area of $2465\,\mathrm{km}^2$, and the blue curve to the mean operated area over the full analysis period of $2007\,\mathrm{km}^2$.

The resulting exposure per energy bin is shown in Figure 8.7. The comparison between the ideal case, the benchmark operated case, and the mean operated case illustrates the impact of detector downtime and partial operation on the effective exposure. At low energies, the benchmark operated exposure is extremely small, with values of only 2.8×10^{-3}, 1.8×10^{-2}, and 2.5×10^{-2} $\mathrm{km}^2\,\mathrm{sr}\,\mathrm{yr}$ at $\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) = 18.6$, 18.8, and 19.0, respectively. These bins therefore mainly illustrate the analysis threshold. Above $\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) = 19.2$, the exposure becomes large enough to provide a meaningful sensitivity estimate, rising from approximately 44 $\mathrm{km}^2\,\mathrm{sr}\,\mathrm{yr}$ at $\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) = 19.2$ to approximately 287 $\mathrm{km}^2\,\mathrm{sr}\,\mathrm{yr}$ at $\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) = 19.8$ for the benchmark operated configuration. The expected photon-flux upper limits are evaluated using the same energy bins shown in the exposure curve. However, the bins below

$$\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) = 19.2$$

are interpreted only as a visualization of the low-energy turn-on of the analysis. In this region, the photon exposure is extremely small and changes rapidly with energy. As a result, the corresponding limits would be driven mainly by the uncertain low-energy efficiency turn-on rather than by a stable exposure estimate. The quoted sensitivity is restricted to the stable high-energy region,

$$\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) \geq 19.2.$$

The analysis employs the Feldman–Cousins construction [129] to derive a 95% confidence-level upper limit on the signal yield for a Poisson counting experiment with an expected background.

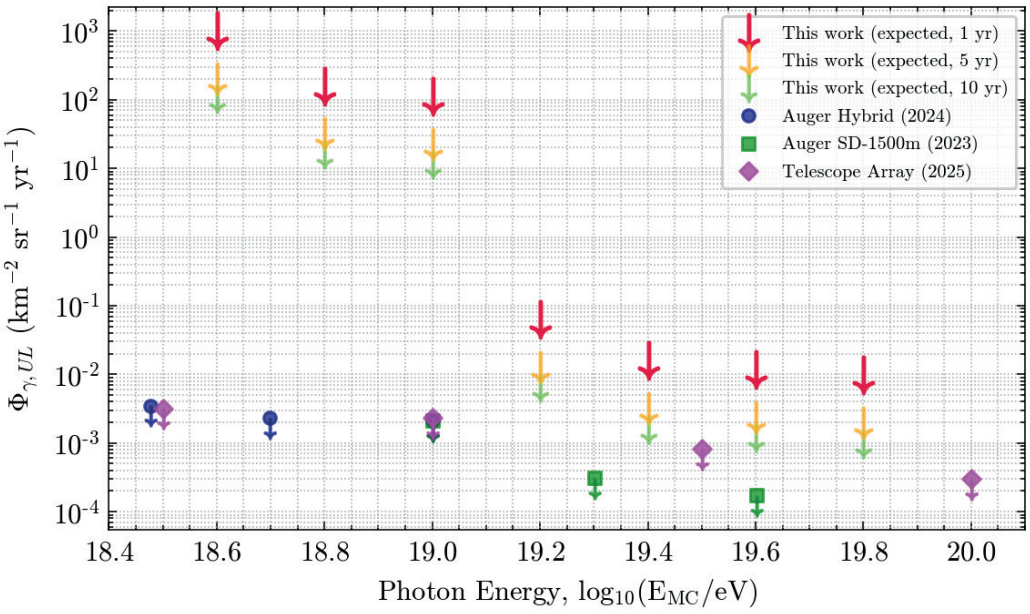


Figure 8.8: Expected 95% C.L. upper limits on the photon flux as a function of true primary photon energy. The limits are computed using the benchmark operated exposure $\mathcal{E}(E_i)$ and the Feldman–Cousins upper limit on the signal counts for $n_{\text{obs}} = 0$ and an expected background $b_{\text{exp}} = 0.308$, giving $\mu_{\text{UL}} = 2.784$ events. The 1-year expectation is projected to 5 and 10 years by scaling the exposure under the assumption of stable detector performance. Published limits from the Pierre Auger Observatory and the Telescope Array are overlaid for comparison; these are reported as integral limits above a threshold energy and are therefore not directly equivalent to the bin-wise limits shown here.

The background expectation is estimated from the simulated hadronic leakage after the final δ'-based photon-selection cut. Since no simulated background event survives this cut, 0/2303, we assign a finite-statistics 95% upper bound to the leakage probability,

$$p_{\mathrm{bkg}}^{95} = 1 - 0.05^{1/2303} = 1.30 \times 10^{-3}.$$

For the fixed-background benchmark used in this projection, based on $N_{\mathrm{data}} = 237$ quality-selected real-data events entering the final photon-selection stage, this gives

$$b_{\mathrm{exp}} = N_{\mathrm{data}}\, p_{\mathrm{bkg}}^{95} = 0.308.$$

The resulting expected limits are displayed over the full plotted energy range, while the quoted photon sensitivity is restricted to the stable high-energy region rather than to the low-energy threshold regime. The Feldman–Cousins upper limit is evaluated with $n_{\mathrm{obs}} = 0$ and $b_{\mathrm{exp}} = 0.308$, giving

$$\mu_{\mathrm{UL}} = 2.784.$$

The corresponding photon-flux upper limits are computed independently in each true-energy bin shown in the figure as

$$\Phi_{\gamma,\mathrm{UL}}(E_i) = \frac{\mu_{\mathrm{UL}}}{\mathcal{E}(E_i)}\,, \tag{8.7}$$

where $E_i \equiv E_{\mathrm{MC},i}$ denotes a true primary photon-energy bin, and $\mathcal{E}(E_i)$ is the benchmark operated exposure in that bin, derived from the stable late-time reference period and including all photon reconstruction, quality, and selection efficiency factors. For the plotted energy-bin centers,

$$\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) = [18.60,\ 18.80,\ 19.00,\ 19.20,\ 19.40,\ 19.60,\ 19.80],$$

the benchmark operated exposures are

$$\mathcal{E}(E_i) = [2.759 \times 10^{-3},\ 1.751 \times 10^{-2},\ 2.470 \times 10^{-2},\ 43.90,\ 176.42,\ 236.85,\ 286.81]\ \mathrm{km^2\,sr\,yr}.$$

Using $\mu_{\mathrm{UL}} = 2.784427$, the corresponding 1-year expected photon-flux upper limits are

$$\begin{aligned}\Phi_{\gamma,\mathrm{UL}}(E_i) \simeq [&1.009 \times 10^{3},\ 1.590 \times 10^{2},\ 1.127 \times 10^{2},\\ &6.343 \times 10^{-2},\ 1.578 \times 10^{-2},\ 1.176 \times 10^{-2},\\ &9.709 \times 10^{-3}]\ \mathrm{km^{-2}\,sr^{-1}\,yr^{-1}}.\end{aligned}$$

Thus, over the plotted range, the expected 1-year limits span from approximately $1.009 \times 10^3\ \mathrm{km^{-2}\,sr^{-1}\,yr^{-1}}$ at $\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) = 18.60$ down to $9.709 \times 10^{-3}\ \mathrm{km^{-2}\,sr^{-1}\,yr^{-1}}$ at $\log_{10}(E_{\mathrm{MC}}/\mathrm{eV}) = 19.80$. The 5-year and 10-year curves shown in Figure 8.8 are obtained by scaling the 1-year flux limits by factors of 1/5 and 1/10, respectively, assuming stable detector performance and proportional exposure growth.[3]

[3] If the expected background also scales with exposure, the Feldman–Cousins limit should be recomputed

The comparison limits from Auger and Telescope Array are integral constraints of the form $\Phi_\gamma(E > E_0)$. They represent the cumulative photon flux above a threshold energy E_0 (and depend on the spectral assumptions adopted in those analyses), while the limits derived here constrain the flux normalisation within each energy bin. The overlay therefore serves as an indicative comparison of sensitivity scales, rather than a one-to-one comparison between equivalent observables.

8.3 Validation: Burn Sample Proposal

This section proposes a statistically robust and practically feasible burn-sample (BS) strategy specifically designed for photon search analyses within the Pierre Auger Collaboration, ensuring transparent and reproducible use of real data during development and validation of the photon–hadron separation criterion.
Currently, several official burn samples exist for previous searches, covering the following detector configurations:

- Main surface detector array (SD-1500)[4]
- Infill array (SD-750)
- Hybrid configurations combining the fluorescence detector (FD) with the main array (SD-1500)
- Hybrid superdense array (SD-433) with the underground muon detector (UMD)

With the implementation of the WCD+RD upgrade, a new burn sample is required to accommodate the updated configuration and ensure consistency in future photon search analyses. In each case, a fixed fraction of the available data (typically 2–10%) is unblinded using selection rules based on event identifiers or GPS timestamps, with only the corresponding detector components burned. These samples are centrally maintained, can be extended upon request, and are designed for reuse across different photon analyses to ensure consistent blinding practices. Since the RD operates in slave-trigger mode, a deterministic burn selection based on the SD event identifier (or GPS time) can in principle also select a representative subset of hybrid events, because the hybrid sample is a subset of SD events. However, the existing SD-1500 burn sample is limited to the phase-I period (2004–2021) and thus does not provide a corresponding unblinded subset of RD readout, RD reconstruction, and hybrid observables for the 2024–2025 production considered in this work. We therefore propose defining an explicit WCD+RD burn sample for the 2024–2025 hybrid dataset, using the same type of deterministic rule but applied consistently to both detector subsystems and all derived hybrid quantities used in the photon search.

with the corresponding time-dependent background expectation. In the background-dominated regime, the long-term sensitivity scaling would become weaker and approach $1/\sqrt{T}$, where T denotes the observation time.

[4]The current SD-1500 burn sample was defined for phase-I SD data (2004–2021), i.e. prior to RD deployment, and therefore does not cover the 2024–2025 hybrid WCD+RD dataset used here.

However, all existing burn samples were designed for analyses involving the surface detector arrays, fluorescence telescopes, and underground muon detectors. With the radio detector now in stable operation, it is useful to define a dedicated burn subsample within the SD-triggered data for events with usable WCD+RD information, in order to develop, tune, and validate the photon–hadron separation strategy for this configuration.
The development and validation of a hybrid WCD+RD photon–hadron separation criterion present unique challenges that necessitate access to representative subsets of real data. Photon searches are inherently sensitive to background contamination, detector response uncertainties, and systematic effects that cannot be fully captured by Monte Carlo simulations alone. The burn sample serves three critical functions in this context:

1. It enables feature engineering and selection tuning on real detector data. While simulations provide the expected response to photon-induced air showers, real data are essential for characterizing the background, assessing the stability of the observables under realistic detector conditions, and identifying possible mismatches between simulated and measured distributions.

2. It allows validation of any classifier or separation-cut performance beyond Monte Carlo assumptions. Photon searches depend critically on accurate modeling of hadronic backgrounds, which constitute the overwhelming majority of cosmic-ray events. Subtle mismatches between simulated and real-data distributions, especially in the tails of discriminating observables, can lead to biased background estimates and reduced sensitivity.

3. It provides a means to quantify and mitigate detector-specific systematic uncertainties. The radio detector introduces additional sources of variability, including atmospheric effects on radio propagation, antenna-response variations, transient radio noise, and reconstruction biases specific to the combined WCD+RD geometry. Testing the analysis chain on a real burn subsample helps expose these effects before unblinding the full dataset.

By establishing a dedicated BS, we ensure that the tuning and validation of the separation strategy are data-driven yet statistically independent of the final signal region, thereby maintaining the rigor required for credible upper limits or potential discovery claims in photon searches.
We propose establishing an event-ID-based BS comprising 10% of hybrid WCD+RD data collected during January–December 2025. Events are selected when the last digit of the event identifier `sd_event_id` equals 3, ensuring statistical independence from physics-based criteria while maintaining consistency with established collaboration blinding practices. This event-ID-based selection yields an approximately uniform subsample over the full January–December 2025 period and across all detector subsystems and reconstruction algorithms. It therefore naturally includes the range of operational conditions present during the analysis period. This also encompasses vertical air showers ($\theta_{\mathrm{WCD}} < 65^{\circ}$), which are not included in the current photon search zenith range but may be relevant for future analyses, as well as coincident detections from FD, HEAT, SSD, and AMIGA.

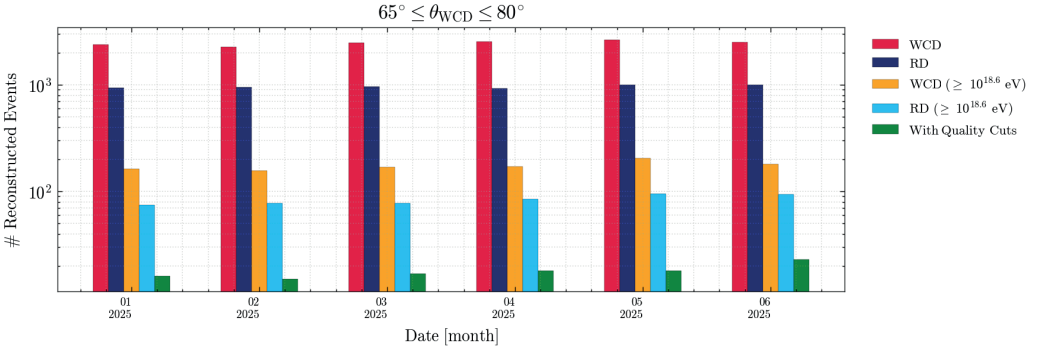


Figure 8.9: Monthly reconstructed event statistics for hybrid WCD+RD data (January–June 2025, $65^\circ \leq \theta_{\mathrm{WCD}} \leq 80^\circ$). Only the first six months are shown, following the original burn-sample proposal; annual totals are obtained by extrapolation. Red: all WCD events; blue: all RD events; orange/cyan: events with $E_{\mathrm{WCD}} \geq 10^{18.6}$ eV for WCD/RD; green: events passing quality cuts.

Advantages of this method are:

- **Detector-agnostic unblinding**: Constant burn fraction regardless of reconstruction method, trigger configuration, or detector uptime variations.
- **Uniform temporal coverage**: Proportional representation of seasonal variations in atmospheric conditions, detector performance, and background composition.
- **Reproducibility**: Deterministic selection independently verifiable by any collaboration member.

Figure 8.9 shows monthly event statistics for January–June 2025 within the zenith range $65^\circ \leq \theta_{\mathrm{WCD}} \leq 80^\circ$, demonstrating stable hybrid WCD+RD operation with 2400–2700 WCD events per month, of which 900–1000 satisfy hybrid reconstruction requirements.
Table 8.4 quantifies event yields through successive selection stages. The six-month dataset contains 14,945 WCD events and 5,782 hybrid RD events, reducing to 1,046 and 505 events, respectively, above the photon search threshold ($E_{\mathrm{WCD}} \geq 10^{18.6}$ eV). After WCD and RD quality cuts are applied (Table 6.3), 237 events remain.
Extrapolating to the full year yields approximately 450 quality-selected events, making a 10% burn sample of roughly 45 events sufficient to validate and tune the separation criterion while preserving 90% ($\sim$400 events/year) for final analysis with near-optimal statistical sensitivity. This balance avoids both overfitting to Monte Carlo predictions (too few burned events) and unnecessarily sacrificing discovery power (too many burned events).
This burn-sample proposal was subsequently adopted by the Pierre Auger Collaboration for the WCD+RD dataset [130]. However, the dedicated burn sample became available only at a late stage of this thesis work, and a full analysis of the unblinded burn subsample could not be completed within the available time.

Month	WCD	RD	WCD ($\geq 10^{18.6}$ eV)	RD ($\geq 10^{18.6}$ eV)	Quality cuts
Jan	2406	943	162	75	37
Feb	2287	953	157	78	33
Mar	2499	961	169	78	33
Apr	2564	923	172	85	39
May	2660	998	205	95	47
Jun	2529	1004	181	94	48
6-month total	**14,945**	**5,782**	**1,046**	**505**	**237**
Annual projection	**29,890**	**11,564**	**2,092**	**1,010**	**474**
10% burn sample	**2,989**	**1,156**	**209**	**101**	**47**

Table 8.4: Monthly (Jan–Jun 2025) and projected annual event counts for hybrid WCD+RD data at successive photon search selection stages.

The present section therefore documents the motivation, definition, and expected statistical size of the burn sample, while its detailed use for data-driven validation is left for future work.

Summary and Outlook

AugerPrime will play a key role in understanding the origin of Ultra-High Energy Cosmic Rays by precisely measuring the mass composition. Neutral particles will be decisive in constraining models. The identification of photon primaries depends on the ability to distinguish their air showers from a background dominated by protons and heavier nuclei primaries. More statistics are needed, as well as an accurate measurement. This thesis establishes a methodology for photon identification using the Radio Detector array, spanning commissioning, reconstruction, and quality selection, and culminating in a transparent cut-based photon–hadron separation strategy. The work quantifies detection capabilities and develops analysis tools that substantially improve event quality and enable high-purity photon-candidate selections.

Chapter 3 establishes the Radio Detector's operational status and data quality through the development and deployment of the Monitoring Radio Detector (MoRD) system, which was essential for validating the array as a reliable instrument. Deployment progressed from the 8-station Engineering Array in 2021 through the 40-station Super Hexagon in 2023 to full array completion in November 2024. The MoRD platform, developed as part of this thesis, processes data every 12 hours, generating interactive visualizations that enable real-time identification of hardware failures, antenna misalignments, and radio frequency interference. Critical commissioning issues were identified and resolved using MoRD diagnostics: swapped antenna arms were detected via anomalous polarization ratios in the 67-MHz TV transmitter line; malfunctioning low-noise amplifiers were identified through systematically reduced gain in specific polarization channels; and antenna rotation under wind loading was quantified through time-dependent alignment monitoring.

Chapter 4 investigates the Radio Detector's potential for photon detection through comprehensive simulation studies. Using photon-induced air shower simulations spanning $10^{18.4}$ to $10^{19.4}$ eV and zenith angles from 65° to 85°, the chapter establishes that radio detection can recover many photon-induced air showers missed by Water Cherenkov Detectors at extreme zenith angles. Three reconstruction modes were evaluated: WCD-only, RD-only, and hybrid operation. The hybrid configuration achieves reconstruction efficiencies exceeding 95% above 10^{19} eV for zenith angles greater than 70°, where WCD efficiency drops below 20%. The effective aperture for photons could reach 1500–2000 km sr in the hybrid mode, approaching the geometric limit of 2400 km sr and representing more than a factor of three improvement over WCD-only operation. These results demonstrate that a radio self-trigger could extend photon sensitivity to energies as low as $10^{18.5}$ eV at large zenith angles, addressing the observational

gap identified in Chapter 1.

Chapter 5 establishes the reconstruction performance of the Radio Detector and validates its reliability through systematic comparison with Water Cherenkov measurements. The Water Cherenkov reconstruction relies on fitting the muon lateral distribution and spherical air-shower front, achieving angular resolution better than 0.5° and energy resolution of approximately 19%. Radio reconstruction employs a two-stage process: timing fits using WCD-seeded geometry followed by lateral distribution fitting to extract electromagnetic energy with better than 5% intrinsic resolution. A critical detector calibration issue was identified and corrected: an antenna orientation error in the Offline software had reversed the east-west polarization by 180°, which was distorting the geomagnetic emission pattern and corrupting energy reconstruction. After correction, data from January-September 2024 showed the expected geomagnetic dominance. For the full deployment period (December 2024 to October 2025), approximately 12,000 hybrid events were successfully reconstructed. The radio detector achieves reconstruction efficiency exceeding 95% at zenith angles above 73° and energies above $10^{18.6}$ eV, with angular precision improving from $\sim 0.3°$ at 5 stations to $<0.15°$ for events with more than 15 stations. The median opening angle between radio and WCD reconstructed directions is 0.25°, with 68% of events showing agreement better than 0.36°. Measured data validating the reliable and unbiased reconstruction performance of two independent energy measurements for inclined air showers: E_{WCD} (muon-dominated energy) and E_{RD} (calorimetric electromagnetic energy), establishing the necessary foundations for the hybrid photon-hadron discrimination parameter used in Chapter 7.

Chapter 6 develops a comprehensive quality selection framework that ensures high-purity data essential for photon searches by operating at both station and event levels to suppress background contamination while preserving genuine air-shower signals. At the station level, three complementary physics-motivated checks replace the standard SNR threshold to achieve substantial background rejection: (1) a dual-window linear SNR criterion requiring $\mathrm{SNR}' \geq 3.5$ in the $\mathbf{v} \times \mathbf{B}$ component, which achieves 15× better noise rejection than the standard SNR ≥ 7.22 cut by adapting to local noise conditions rather than using fixed amplitude thresholds; (2) a cosine-similarity test against a bipolar template waveform constructed from 32,000 high-quality traces, requiring correlation ≥ 0.70 over a 200 ns window, which rejects 45.6% of false positives (primarily RFI and hardware transients) while retaining 96.7% of genuine air-shower signals; and (3) a polarization purity check requiring the geomagnetic $\mathbf{v} \times \mathbf{B}$ component to contribute at least 60% of the signal power at peak amplitude, eliminating residual outliers where the perpendicular component dominates contrary to air-shower physics. These station-level cuts, when combined, achieve a signal efficiency minus noise leakage fraction of 41.5% and reduce false positives by more than an order of magnitude compared to the standard method. At the event level, strict geometric and reconstruction quality requirements ensure reliable photon-hadron discrimination: minimum station multiplicity $N_{\mathrm{RD}} \geq 5$ (39% of events), at least one station sampling the Cherenkov cone where radio emission peaks (84% of 5+ station events), lateral distribution fit quality $\chi^2_\nu \leq 3$ (78% retention), directional consistency between radio and WCD

with opening angle $\leq 1^\circ$ (92% retention), and thunderstorm flagging using meteorological data and satellite lightning observations to exclude periods of atmospheric electric field contamination (66% retention). Applied to the pilot dataset spanning October 2024 to February 2025, the complete quality selection retains 16% of initially reconstructed events (1005 of 6383).

Chapter 7 establishes photon–hadron separation using the hybrid energy mismatch between radio and Water Cherenkov reconstructions. The relative energy difference $\delta = (E_{\mathrm{WCD}} - E_{\mathrm{RD}})/E_{\mathrm{RD}}$ serves as the primary discriminator, as photon-induced air showers yield systematically higher radio energies relative to WCD measurements compared to hadronic air showers. A background-only robust linear detrending removes residual instrumental and reconstruction trends, yielding an energy-stabilized observable δ'. A single transparent cut, $\delta' < \delta'_{\mathrm{cut}}$, provides strong separation in the inclined hybrid simulations: for a representative working point retaining $\sim 80\%$ of photons, no background event is observed to pass in the available hadronic sample, and the first nonzero leakage appears only when loosening the cut to achieve $\epsilon_\gamma \simeq 92\%$. These results validate the hybrid detection strategy for UHE photon searches at large zenith angles and demonstrate that a simple, physically motivated selection can deliver the high purity required for a rare-photon search. Chapter 8 builds the analysis chain needed to turn the selected hybrid data set into a quantitative photon search, defining how the available observing time and detector coverage translate into an exposure that can be used for flux constraints.

Chapter 8 develops the exposure calculation and establishes a framework for applying the photon–hadron separation selection to real data. Using the period from December 2024 to October 2025, station operational status was tracked hourly by integrating trace-based activity, bad-period records, and array geometry. Active hexagons were counted in 12-hour bins to construct a time-dependent geometric exposure. For the zenith range 65°–80°, the cumulative geometric exposure reaches 850.4 $\mathrm{km}^2\,\mathrm{sr}\,\mathrm{yr}$, while the fully efficient domain 73°–80° accumulates 316.9 $\mathrm{km}^2\,\mathrm{sr}\,\mathrm{yr}$. While this full-efficiency statement applies to hadronic air showers in the chosen domain, the expected photon sensitivity is obtained by folding the operated exposure with photon-specific reconstruction and selection efficiencies, which rise with energy. The chapter concludes by proposing a burn-sample strategy: 10% of the 2025 data (selected via event ID modulo 10 = 3) would provide ~ 45 quality-selected events to validate the separation cut on real data while preserving 90% (corresponding to ~ 400 events/year) for the final photon search, ensuring statistical rigor through blinded analysis protocols.

In conclusion, rigorous detector commissioning, data quality selection, and a transparent hybrid separation strategy provide the necessary tools and a statistical framework to conduct the next-generation search for ultra-high-energy photons at the Pierre Auger Observatory. Crucially, the Radio Detector array is now fully deployed and operating smoothly, supported by the validated monitoring and quality frameworks developed herein, allowing for trusted data collection and enabling this new photon search to contribute significantly to the existing methods.

Outlook

It is essential to validate the photon–hadron separation cut on real data through the proposed burn-sample strategy. The 10% sample (∼45 quality-selected events from 2025 data) will test whether the simulation-derived selection efficiency and background leakage remain valid under realistic detector conditions, including time-dependent calibration drifts, seasonal atmospheric variations, and the full diversity of hadronic backgrounds not captured by SIBYLL and QGSJET models. Discrepancies between the burn sample and Monte Carlo performance will inform systematic uncertainty budgets before unblinding the remaining 90% of data. Following successful validation, the first photon flux upper limits from the hybrid detector at large zenith angles can be established.

Operating the Radio Detector in slave mode (dependent on Water Cherenkov triggers) fundamentally limits photon sensitivity below $\sim 10^{18.6}$ eV and at the most inclined geometries ($\theta > 80°$), where WCD efficiency drops precipitously. Chapter 4 demonstrated via simulation that a radio self-trigger could recover essentially all photons missed by WCD, extending energy reach toward 10^{18} eV and improving aperture by factors of 2–3 at large zenith angles. This upgrade transforms the Radio Detector from a passive slave into an independent, photon-sensitive instrument. More on how we can implement this can be found in [131].

The current separation strategy relies primarily on the energy mismatch $(E_{\mathrm{WCD}} - E_{\mathrm{RD}})/E_{\mathrm{RD}}$ (after background-only detrending), which proves remarkably effective for producing high-purity photon-candidate selections, but leaves substantial information unused. Future iterations could incorporate additional hybrid observables, such as the number of triggered radio stations relative to WCD stations (especially in a self-trigger configuration where the two multiplicities are independent), the steepness of the radio lateral distribution function, and the ratio of signal rise times between detectors. Preliminary analyses suggest these variables are only weakly correlated with δ', indicating potential for improved separation power. If a multivariate approach is pursued, residual energy and zenith dependencies should be controlled explicitly (e.g. by debiasing or by stability checks) to ensure the selection remains driven by air-shower physics rather than triggering thresholds or geometric acceptance.

Samenvatting

Kosmische straling bestaat uit hoogenergetische deeltjes, voornamelijk protonen en zwaardere atoomkernen, die de aarde voortdurend vanuit de ruimte bereiken. De meest energetische deeltjes, met energieën boven 10^{18} eV, worden ultrahoge-energie kosmische straling genoemd. Hoewel duidelijk is dat deze deeltjes grotendeels van buiten ons Melkwegstelsel afkomstig zijn, blijven hun bronnen en versnellingsmechanismen grotendeels onbekend. Omdat geladen deeltjes worden afgebogen door galactische en extragalactische magneetvelden, is hun aankomstrichting op aarde moeilijk rechtstreeks te herleiden naar hun bron.

Ultrahoge-energie fotonen vormen daarom een bijzonder waardevol alternatief. Als elektrisch neutrale deeltjes worden zij niet door magneetvelden afgebogen en kunnen zij, in principe, direct terugwijzen naar hun productieplaats. Daarnaast bieden zij een gevoelige toets voor modellen van kosmische-stralingsbronnen, propagatie-effecten en mogelijke nieuwe natuurkunde. Tot nu toe zijn zulke fotonen echter niet overtuigend waargenomen, waardoor een hogere detectiegevoeligheid essentieel is.

Wanneer een kosmisch deeltje de atmosfeer binnendringt, veroorzaakt het een uitgebreide luchtdouche. Hadronische douches bevatten zowel een elektromagnetische component als veel muonen, terwijl foton-geïnduceerde douches vooral elektromagnetisch zijn en veel minder muonen bevatten. Dit verschil vormt de basis voor foton-hadron-scheiding. Bij grote zenithoeken wordt de elektromagnetische component echter grotendeels geabsorbeerd voordat zij de grond bereikt. Water-Cherenkovdetectoren meten dan vooral muonen, waardoor fotondouches zwakke signalen geven. Radiostraling, geproduceerd door de elektromagnetische component, plant zich daarentegen grotendeels ongehinderd door de atmosfeer voort. Radiodetectie is daarom bijzonder geschikt voor sterk hellende luchtdouches.

Dit proefschrift onderzoekt hoe de combinatie van de Water-Cherenkovdetectoren en de Radiodetector van AugerPrime kan worden gebruikt in de zoektocht naar ultrahoge-energie fotonen. De Radiodetector bestaat uit ongeveer 1660 antennes die gevoelig zijn voor radiostraling tussen 30 en 80 MHz. Door de muongevoelige Water-Cherenkovdetectoren te combineren met de elektromagnetisch gevoelige Radiodetector ontstaat een hybride meetmethode waarmee foton-geïnduceerde douches beter kunnen worden onderscheiden van de veel talrijkere hadronische achtergrond.

Een belangrijk onderdeel van dit werk is de ingebruikname en kwaliteitsbewaking van de Radiodetector. Hiervoor is het Monitoring Radio Detector-systeem, MoRD, ontwikkeld. Dit systeem volgt de toestand van de radiostations en helpt hardwareproblemen, foutieve aansluitin-

gen en veranderingen in antenneoriëntatie te identificeren. Daarnaast zijn reconstructie- en kwaliteitsselecties ontwikkeld om betrouwbare radiogebeurtenissen te selecteren, onder meer op basis van signaal-ruisverhouding, pulsvorm, polarisatie, laterale signaalverdeling en overeenstemming tussen radio- en WCD-richting.

Met simulaties is aangetoond dat de hybride combinatie van WCD en RD de gevoeligheid voor fotondouches bij grote zenithoeken aanzienlijk vergroot. De centrale discriminatievariabele is het verschil tussen de energie gereconstrueerd met de Water-Cherenkovdetectoren en die met de Radiodetector. Omdat fotondouches weinig muonen bevatten, wijkt de verhouding tussen E_{WCD} en E_{RD} karakteristiek af van die van hadronische douches. Na correctie voor trends in energie en zenithoek levert deze grootheid een eenvoudige en transparante foton-hadron-scheiding.

Ten slotte is een raamwerk opgesteld voor toepassing op echte data, inclusief blootstellingsberekeningen en een geblindeerde analyse met een burn sample ter validatie. Samenvattend laat dit proefschrift zien dat de hybride combinatie van de Water-Cherenkovdetectoren en de Radiodetector van AugerPrime een krachtig nieuw instrument vormt voor de zoektocht naar ultrahoge-energie fotonen bij grote zenithoeken.

Appendix A: Simulation Sets

In this study, the simulation of inclined air showers was conducted using a local dataset generated by [110], as there is currently no official simulation framework for photons in AugerPrime that incorporates the Radio Detector. The dataset utilized CORSIKA in conjunction with the CoREAS radio extension, incorporating the SIBYLL2.3d hadronic model and the ideal Auger sparse grid. The simulated air showers were inclined at angles greater than 65°, with energies falling within the range of $10^{18.4}$ eV to $10^{20.2}$ eV. The focus of the analysis was on primary energies below $10^{19.7}$ eV to account for preshower conversion in the Earth's magnetosphere. The zenith angle of the air showers ranged from 65° to 85°. It is important to note that, as depicted in Figure A.1, highly inclined air showers exhibit lower statistics across all energy bins. However, the dataset utilized in this analysis has certain limitations, particularly in terms of azimuthal coverage. The simulation provides discrete values at azimuthal angles of 0°, 45°, 90°, 135°, 180°, 225°, 270°, and 315°. This discrete azimuthal coverage imposes constraints on the ability to perform detailed directional analyses beyond these specific angles. This will limit the generalization of the study's findings to a broader range of air-shower geometries. Furthermore,

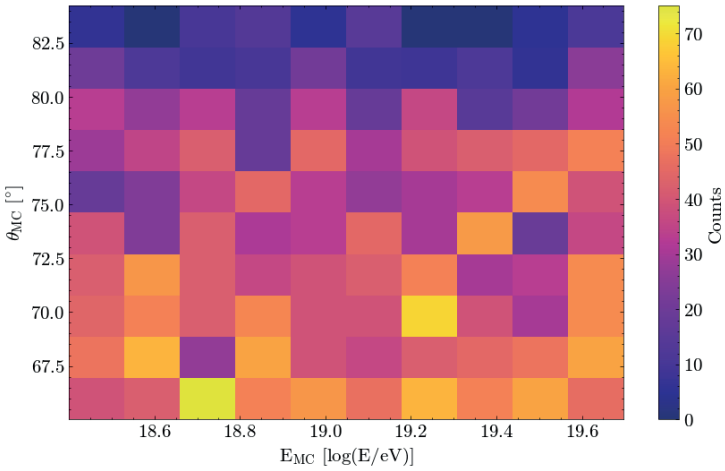


Figure A.1: Energy and zenith distribution of the selected simulation set.

the determination of X_{max} for each simulation, obtained through the Gaisser-Hillas (GH) fit of the longitudinal profile of electrons and positrons. Few simulations have $X_{max} = 0$ due to the vertical depth arrangement of the gauge planes; by the time the air shower has developed sufficiently for a successful GH fit, it has already terminated. Moreover, dozens of GH fits fail

and return $X_{max} = -1$ as illustrated in Figure A.2. This usually indicates that the simulation did not produce a valid result for the air-shower maximum. There are several possible reasons, such as:

- The simulation may not have run successfully or encountered issues during the execution.
- The simulation time might be insufficient to allow the air shower to develop and reach its maximum depth.

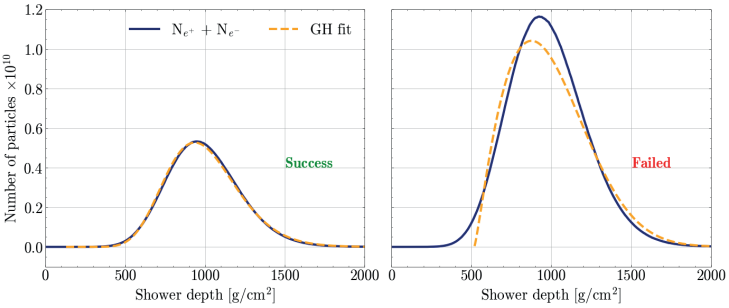


Figure A.2: An example illustrating a successful fit (Left) and a failed fit (Right) of the Gaisser-Hillas (GH) function to the number of particles versus the measurement depth. The blue line shows the total number of electrons and positrons at each depth measurement during the simulation.

Consequently, the majority of such events occur at very high inclinations ($\theta > 80°$) and relatively low energies ($E < 10\,EeV$) as shown in Figure A.3. This scenario necessitates the simulation of a substantial number of stations for each event. The increased number of stations, combined with the computationally intensive nature of the radio simulations, contributes to a considerable simulation time for each event in this high-inclination range.

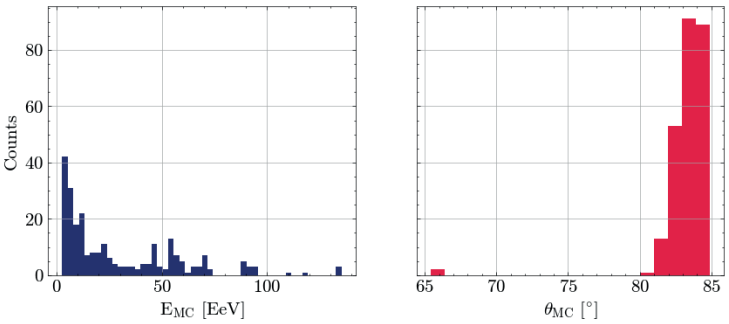


Figure A.3: The distribution of energy and zenith for omitted event with $X_{max} \in \{0, -1\}$.

In conclusion, 3421 simulations were selected to proceed into the analyses.

Appendix B: Aperture Derivation

For a flat detector with physical area A, the effective aperture is derived by considering the detector's projected area and the solid angle over which particles arrive. We define the detection efficiency, $\eta(E, \theta, \phi)$, as the probability that an extensive air shower of energy E arriving from direction (θ, ϕ) will be detected and satisfy all quality criteria.
The differential effective area, $dA_{\text{eff}}(E, \theta, \phi)$, for primary particles coming from direction (θ, ϕ) with energy E is given by:

$$dA_{\text{eff}}(E, \theta, \phi) = A_{\text{proj}} \cdot \eta(E, \theta, \phi) \cdot d\Omega,$$

where $A_{\text{proj}} = A \cos\theta$ is the projected area of the detector perpendicular to the incoming particle direction, and $d\Omega = \sin\theta \, d\theta \, d\phi$ is the solid angle element in spherical coordinates. Substituting these terms, we get:

$$dA_{\text{eff}}(E, \theta, \phi) = A \cos\theta \cdot \eta(E, \theta, \phi) \cdot \sin\theta \, d\theta \, d\phi,$$

For detectors with azimuthal symmetry, or in cases where the azimuthal dependence is averaged out (as in our simulation set), the efficiency can be approximated as independent of ϕ, i.e. $\eta(E, \theta)$. We can then integrate over all azimuthal angles ϕ from 0 to 2π:

$$\begin{aligned} dA_{\text{eff}}(E, \theta) &= A \cos\theta \cdot \eta(E, \theta) \cdot \sin\theta \, d\theta \cdot \int_0^{2\pi} d\phi \\ &= 2\pi \cdot A \cdot \eta(E, \theta) \cdot \sin\theta \cos\theta \, d\theta. \end{aligned}$$

To obtain the effective aperture $A_{\text{eff}}(E_i, t)$ for a given energy bin E_i and time bin t within the zenith angle range $[\theta_1, \theta_2]$, the differential aperture must in general be integrated over zenith angle while retaining the zenith dependence of the detection efficiency:

$$A_{\text{eff}}(E_i, t) = 2\pi \, A(t) \int_{\theta_1}^{\theta_2} \eta(E_i, \theta) \, \sin\theta \cos\theta \, d\theta. \tag{B.1}$$

In the present analysis, however, the radio detection efficiency is not evaluated differentially in zenith, but is obtained as a single value for each energy bin over the full sensitive zenith range. The resulting efficiency therefore represents a zenith-averaged quantity, which we denote by

$\bar{\eta}(E_i)$, defined as

$$\bar{\eta}(E_i) = \frac{\int_{\theta_1}^{\theta_2} \eta(E_i, \theta)\, \sin\theta \cos\theta \, d\theta}{\int_{\theta_1}^{\theta_2} \sin\theta \cos\theta \, d\theta}. \tag{B.2}$$

Using this definition, Eq. (B.1) can be written as

$$A_{\text{eff}}(E_i, t) = \pi \, A(t) \, \bar{\eta}(E_i) \, \left(\cos^2\theta_1 - \cos^2\theta_2\right). \tag{B.3}$$

Thus, the aperture used in this work is expressed in terms of the zenith-averaged detection efficiency over the selected angular interval, rather than assuming that the efficiency is independent of zenith angle. The term $(\cos^2\theta_1 - \cos^2\theta_2)$ accounts for the solid-angle acceptance of the detector within the defined zenith range, modulated by the projected-area effect.

Appendix C: OfflineB Configurations

Rd-mode Version: `cd289f9` (Mar 10, 2025)

```
<moduleControl>
  <loop numTimes="unbounded" pushEventToStack="yes">
    <module> EventFileReaderOG                    </module>
    <loop numTimes="1" pushEventToStack="yes">
      <module> EventGeneratorOG                   </module>
      <module> RdStationAssociator                </module>
      <module> RdEventInitializer                 </module>
      <module> RdAntennaStationToChannelConverter </module>
      <module> RdChannelResponseIncorporator      </module>
      <module> RdChannelResampler                 </module>
      <module> RdChannelTimeSeriesClipper         </module>
      <module> RdChannelVoltageToADCConverter     </module>
      <module> RdTimeJitterAdder                  </module>
      <module> RdChannelASCIINoiseImporterRD      </module>
      <module> RdChannelADCClipper                </module>
      <module> RdChannelADCToVoltageConverter     </module>
      <module> RdChannelUpsampler                 </module>
      <module> RdChannelPedestalRemover           </module>
      <module> RdChannelResponseIncorporator      </module>
      <module> RdChannelTimeSeriesTaperer         </module>
      <module> RdAntennaChannelToStationConverter </module>
      <module> RdStationSignalReconstructor       </module>
      <module> RdSphericalFit                     </module>
      <module> RdHASLDFFitter                     </module>
      <module> RecDataWriterNG                    </module>
    </loop>
  </loop>
</moduleControl>
```

WCD-mode Version: `cd289f9` (Mar 10, 2025)

```
<moduleControl>
  <loop numTimes="unbounded" pushEventToStack="yes">
    <module> EventFileReaderOG                          </module>
    <loop numTimes="1" pushEventToStack="yes">
      <module> EventGeneratorOG                         </module>
      <module> SdAccidentalInjectorKG                   </module>
      <module> G4StationSimulatorOG                     </module>
      <loop numTimes="unbounded" pushEventToStack="no">
        <module> CachedShowerRegeneratorOG              </module>
        <module> G4StationSimulatorOG                   </module>
      </loop>
      <module> SdSimulationCalibrationFillerOG          </module>
      <module> SdPMTSimulatorOG                         </module>
      <module> SdFilterFADCSimulatorMTU                 </module>
      <module> SdBaselineSimulatorOG                    </module>
      <module> TankTriggerSimulatorOG                   </module>
```

```
      <module> TankGPSSimulatorOG                          </module>
      <module> ClearParticleLists                          </module>
      <module> CentralTriggerSimulatorXb                   </module>
      <module> CentralTriggerEventBuilderOG                </module>
      <module> EventBuilderOG                              </module>
      <module> EventCheckerOG                              </module>
      <module> SdGainRatioCorrectorKG                      </module>
      <module> SdStationCheckerOG                          </module>
      <module> SdHistogramFitterKG                         </module>
      <module> SdBaselineFinderKG                          </module>
      <module> SdTraceCalibratorOG                         </module>
      <module> SdSignalRecoveryKLT                         </module>
      <module> SdMonteCarloEventSelectorOG                 </module>
      <module> SdEventSelectorOG                           </module>
      <module> TopDownSelector                             </module>
      <module> SdPlaneFitOG                                </module>
      <module> SdHorizontalReconstruction                  </module>
      <module> SdHorizontalEventSelectorHOG                </module>
      <module> RecDataWriterNG                             </module>
    </loop>
  </loop>
</moduleControl>
```

Hybrid-mode Version: cd289f9 (Mar 10, 2025)

```
<moduleControl>
  <loop numTimes="unbounded" pushEventToStack="yes">
    <module> EventFileReaderOG                    </module>
    <loop numTimes="1" pushEventToStack="yes">
      <module> EventGeneratorOG                   </module>
      <module> RdStationAssociator                </module>
      <module> RdEventInitializer                 </module>
      <module> RdAntennaStationToChannelConverter </module>
      <module> RdChannelResponseIncorporator      </module>
      <module> RdChannelResampler                 </module>
      <module> RdChannelTimeSeriesClipper         </module>
      <module> RdChannelVoltageToADCConverter     </module>
      <module> RdTimeJitterAdder                  </module>
      <module> SdAccidentalInjectorKG             </module>
      <module> G4StationSimulatorOG               </module>
      <loop numTimes="unbounded" pushEventToStack="no">
        <module> CachedShowerRegeneratorOG        </module>
        <module> G4StationSimulatorOG             </module>
      </loop>
      <module> SdSimulationCalibrationFillerOG    </module>
      <module> SdPMTSimulatorOG                   </module>
      <module> SdFilterFADCSimulatorMTU           </module>
      <module> SdBaselineSimulatorOG              </module>
      <module> ClearParticleLists                 </module>
      <module> RdSimulationRadioTrigger           </module>
      <module> TankTriggerSimulatorOG             </module>
      <module> TankGPSSimulatorOG                 </module>
      <module> CentralTriggerSimulatorXb          </module>
      <module> CentralTriggerEventBuilderOG       </module>
      <module> RdStationSimulationRejector        </module>
      <module> EventBuilderOG                     </module>
      <module> EventCheckerOG                     </module>
      <try>
        <module> EventCheckerOG                   </module>
        <module> SdGainRatioCorrectorKG           </module>
        <module> SdStationCheckerOG               </module>
```

```
      <module> SdHistogramFitterKG              </module>
      <module> SdBaselineFinderKG               </module>
      <module> SdTraceCalibratorOG              </module>
      <module> SdSignalRecoveryKLT              </module>
      <module> SdMonteCarloEventSelectorOG      </module>
      <module> SdEventSelectorOG                </module>
      <module> TopDownSelector                  </module>
      <module> SdPlaneFitOG                     </module>
      <module> SdHorizontalReconstruction       </module>
      <module> SdHorizontalEventSelectorHOG     </module>
    </try>
    <module> RdPreWaveFitter                    </module>
    <loop numTimes="unbounded">
      <module> RdDirectionConvergenceChecker    </module>
      <module> RdAntennaChannelToStationConverter </module>
      <module> RdStationSignalReconstructor     </module>
      <module> RdSphericalFit                   </module>
    </loop>
    <module> RdHASLDFFitter                     </module>
    <module> RecDataWriterNG                    </module>
  </loop>
 </loop>
</moduleControl>
```

Data Reconstruction Version: 0d2ece64 (Aug 21, 2025)

```
<moduleControl>
  <loop numTimes = "unbounded">
    <module> EventFileReaderOG                  </module>
    <module> RdEventPreSelector                 </module>
    <module> SdGainRatioCorrectorKG </module>
    <module> TriggerTimeCorrection </module>
    <module> SdStationCheckerOG </module>
    <module> SdHistogramFitterKG </module>
    <module> SdBaselineFinderKG </module>
    <module> SdTraceCalibratorOG </module>
    <module> SdStationPositionCorrection </module>
    <module> SdBadStationRejectorKG </module>
    <module> SdSignalRecoveryKLT </module>
    <module> SdEventSelectorOG </module>
    <module> TopDownSelector </module>
    <module> SdPlaneFitOG </module>
    <module> SdHorizontalReconstruction </module>
    <module> SdHorizontalEventSelectorHOG </module>
    <!-- RdHASObserver -->
    <module> RdEventInitializer                 </module>
    <module> RdChannelADCToVoltageConverter     </module>
    <module> RdChannelResponseIncorporator      </module>
    <module> RdChannelPedestalRemover           </module>
    <module> RdChannelDTFTFilter                </module>
    <module> RdChannelTimeSeriesTaperer         </module>
    <module> RdChannelUpsampler                 </module>
    <module> RdAntennaChannelToStationConverter </module>
    <module> RdStationTimingCalibrator          </module>
    <module> RdStationSignalReconstructor       </module>
      <try>
    <module> RdSphericalFit                     </module>
    <module> RdHASLDFFitter                     </module>
      </try>
    <module> RecDataWriterNG                    </module>
  </loop>
```

```
  </moduleControl>
```

Bootstrap override:

```
<configLink id="RdStationSignalReconstructor"><VectorialComponent> 4 </VectorialComponent></configLink>
```

Appendix D: Machine learning for photon–hadron discrimination

During the method-development phase of this work, we investigated whether a multivariate machine-learning (ML) classifier could improve the separation of photon-induced air showers from the hadronic background using hybrid WCD+RD information. The feature construction and detrending procedure follow exactly the description in Section 7.3 of the main text. The goal was to combine several reconstructed, energy-related observables into a single discriminant, while restricting the input to quantities that are minimally affected by acceptance and selection effects.

Although the ML models achieve competitive performance in simulation, the gain in photon efficiency at the extremely low background-leakage levels relevant for the search is modest compared to a direct cut on the detrended energy-mismatch observable. More importantly, the assessment of such low leakage is limited by the available background statistics in the simulations, making zero-leakage operating points difficult to interpret as stringent bounds. In contrast, a cut-based selection enables a transparent leakage bound using the full available background sample, and provides a selection that is straightforward to validate on real data. For these reasons, the final analysis presented in Chapter 7 adopts a linear/direct separation strategy based on the detrended energy-mismatch observable, prioritising robustness and interpretability.

For completeness, this appendix documents the ML-based study. It reproduces the development starting from the feature construction and the polynomial detrending of the relative energy difference, and then continues with the classifier setup, optimisation strategy, and performance evaluation on simulations.

The Classification Model Design and Performance

This section presents a proof-of-concept classification study of photon-induced versus hadron-induced air showers using hybrid WCD+RD energy information. The model design follows the four main considerations introduced earlier. We restrict the input to reconstructed energy-related features only. This minimizes selection effects that would otherwise require careful treatment. For example, the WCD station multiplicity depends on arrival direction and is not identically distributed between background and signal. Similarly, the RD station multiplicity depends on the WCD trigger and on the imposed quality criteria. By focusing on the energy

mismatch observable, we isolate its discriminating power while reducing sensitivity to acceptance effects.
The classifier uses three event-level input features:

- The detrended relative energy difference δ': the mismatch between the muon-dominated energy estimate from the WCD and the electromagnetic energy estimate from the RD.
- The per-event uncertainty $\sigma_{\delta'}$ obtained by error propagation: a measure of how precisely δ' is determined for each event. Events with large $|\delta'|$ and small $\sigma_{\delta'}$ correspond to a high-significance disagreement between the two energy estimators.
- The ratio of energy uncertainties, $\sigma_{E_{\mathrm{WCD}}}/\sigma_{E_{\mathrm{RD}}}$: complementary information on the relative reconstruction quality of the two detectors.

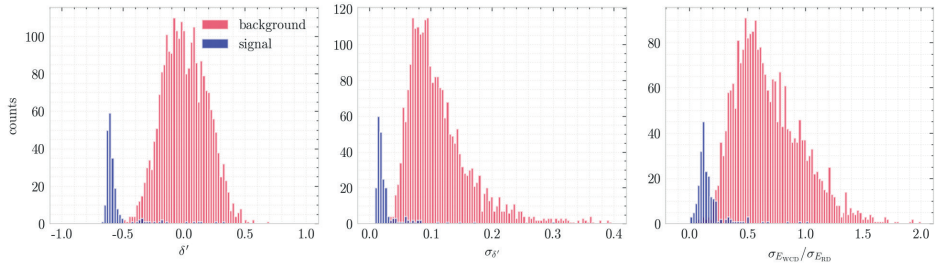


Figure D.1: Distributions of the three input features. Signal (photons, blue) and background (hadrons, red).

Figure D.1 shows the distributions of the three input features. Signal (blue, photons) and background (red, hadrons) are separated in δ', and the other two variables provide complementary information through their different widths and shifts. Correlation matrices in Figure D.2 show non-zero inter-feature dependencies. We therefore use both a linear discriminant analysis (LDA) model and a non-linear Random Forest (RF) classifier.

Model Design

Class labels are mapped to $y \in \{0, 1\}$ with $y = 1$ for signal and $y = 0$ for background. The Monte Carlo sample is split into a training set and an independent test set; all preprocessing, hyperparameter tuning, and score calibration are performed using the training set only, while the final quoted performance is evaluated on the held-out test set. The classification models are implemented using the SCIKIT-LEARN library [132].
For the LDA model, the three inputs are standardized to have zero mean and unit variance using statistics computed on the training set only, thereby avoiding data leakage. The Random Forest uses the raw features directly, as tree-based methods are invariant to feature scaling.
Given class means $\boldsymbol{\mu}_0, \boldsymbol{\mu}_1$ and within-class scatter

$$S_w = \sum_{k\in\{0,1\}} \sum_{i\in\mathcal{C}_k} (\mathbf{x}_i - \boldsymbol{\mu}_k)(\mathbf{x}_i - \boldsymbol{\mu}_k)^\top,$$

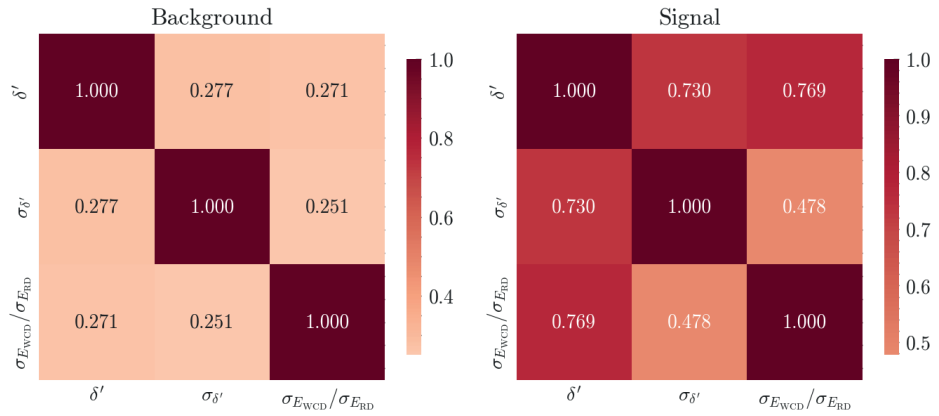


Figure D.2: Feature correlations for background (left) and signal (right). Background shows weaker relationships between variables; signal shows stronger relationships.

where $\mathcal{C}_k$ denotes the set of events belonging to class k, the LDA [133] projection vector is

$$\mathbf{w} = S_w^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0),$$

and the discriminant score is $s(\mathbf{x}) = \mathbf{w}^\top \mathbf{x} + b$, with b an offset (intercept) that can equivalently be absorbed into the decision threshold.

To allow for non-linear decision boundaries and feature interactions, we use a RF classifier [134] with class balancing. For each training split, we conduct a small hyperparameter search and then calibrate the scores so that probability thresholds can be applied consistently. Hyperparameters are tuned via a randomized search with stratified 10-fold cross-validation, maximizing the ROC AUC (area under the receiver-operating-characteristic curve, i.e. the integral of true-positive rate versus false-positive rate over all thresholds) over:

$$\begin{aligned} \texttt{max_depth} &\in \{\texttt{None}, 12, 18, 24\}, \\ \texttt{min_samples_leaf} &\in \{1, 2, 4\}, \\ \texttt{min_samples_split} &\in \{2, 5, 10\}, \\ \texttt{bootstrap} &\in \{\texttt{True}, \texttt{False}\}. \end{aligned}$$

Here, `max_depth` controls tree complexity (with `None` allowing full growth). `min_samples_split` and `min_samples_leaf` set the minimum number of events required to create a split or a leaf, respectively; larger values regularize the trees and reduce variance. `bootstrap` selects whether each tree is trained on a bootstrap resample of the training set. Other settings are fixed to robust defaults. The tuned RF is calibrated using isotonic regression with 5-fold cross-validation fitted on the training set. Isotonic regression learns a monotonic mapping from raw classifier scores to calibrated probabilities, improving the interpretability of $s \in [0, 1]$ as $P(\text{signal})$. Calibrated scores allow fixed decision thresholds to be applied consistently across splits and energies.

Performance

As discussed in section 7.1, the available simulation sample is limited, both in photon statistics and in its ability to sample the extreme tails of the hadronic background distribution. After all selections, the dataset contains approximately 223 photon events and 2746 background events. With a 30% test fraction, each evaluation uses $N_{\text{test}} \sim 8 \times 10^2$ background events in the test set, varying slightly with the random split. This finite test-set size sets a hard limit on what can be claimed about false-positive rates (FPR). In a single split, the FPR takes discrete values

$$\text{FPR} = \frac{k}{N_{\text{test}}}, \qquad k = 0, 1, 2, \ldots$$

so the smallest non-zero value that can be resolved in a single split is $1/N_{\text{test}} \sim 1.4 \times 10^{-3}$. If zero false positives are observed, one can only place an upper limit on the FPR. Using the standard approximation for zero observed events at 95% confidence level (rule of three), this corresponds to

$$\text{FPR}_{95} \lesssim \frac{3}{N_{\text{test}}} \sim 4 \times 10^{-3}.$$

The goal of this section is not to provide a final classifier for data, but to demonstrate, on the available simulation set, that the hybrid WCD+RD energy information yields photon–hadron separation in the core of the distributions. All quoted performance numbers are constrained by the size of the simulation sample, in particular by the number of background events in the test sets.

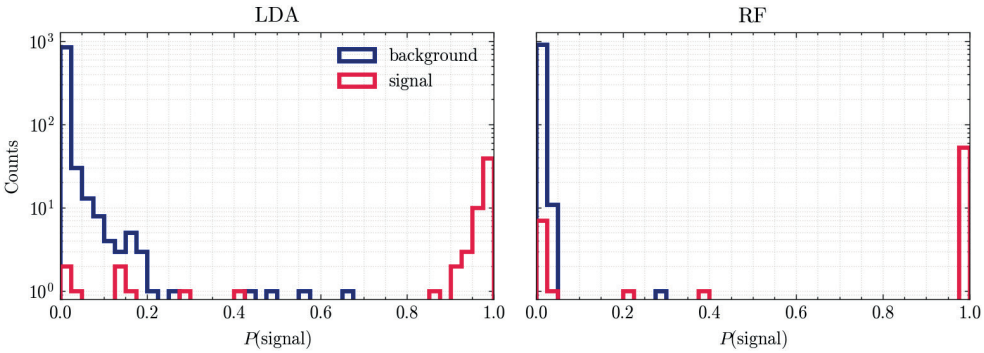


Figure D.3: Distributions of the calibrated classifier score $P(\text{signal})$ in the test set for a representative train/test split (seed 0). Background (navy) and signal (crimson) are shown for LDA (left) and RF (right). The vertical axis is logarithmic to visualize the low-probability tail of the background alongside the high-probability signal peak.

We first inspect the distributions of the calibrated classifier score $P(\text{signal})$ in the test set for signal (photons) and background (hadrons) (see Figure D.3). Both models assign most background events $P(\text{signal}) \approx 0$, while signal events cluster near $P(\text{signal}) \approx 1$; the logarithmic scale highlights the small background tail at intermediate scores, which is more pronounced for LDA than for RF. We then evaluate the classifier at three decision thresholds, $t = 0.90$,

0.99, and 0.999, and compute the signal efficiency (true-positive rate, i.e. the fraction of photons classified as signal) and the false-positive rate (fraction of background events classified as signal). We repeat the stratified train/test split for 100 independent random seeds (30% test fraction). Repeating the split does not increase the effective statistics of the simulation set, but quantifies the sensitivity of the results to statistical fluctuations in the limited sample.
Table D.1 presents the performance of LDA and RF at the three thresholds, averaged over 100 seeds.

Metric/Model	LDA 0.9	LDA 0.99	LDA 0.999	RF 0.9	RF 0.99	RF 0.999
Recall_mean	0.883	0.153	0.000	0.871	0.850	0.842
Recall_std	0.029	0.085	0.000	0.048	0.049	0.050
FPR_mean	0.000067	0.0	0.0	0.000135	0.0	0.0
FPR_std	0.000301	0.0	0.0	0.000415	0.0	0.0
BkgRej_mean	0.999933	1.0	1.0	0.999865	1.0	1.0

Table D.1: Mean and standard deviation of performance metrics across 100 seeds for LDA and RF models at thresholds $t = 0.90$, 0.99, and 0.999. For FPR, values near zero reflect that many seeds yield $k = 0$ false positives in the test set. With $N_{\text{test}} \sim 7 \times 10^2$ background events per split, the smallest non-zero per-split FPR is $1/N_{\text{test}} \sim 1.4 \times 10^{-3}$.

At the strictest threshold $t = 0.999$, the RF retains 0.842 ± 0.050 of photon events on average across seeds. In the available test sets, no hadronic events are observed to pass this cut for RF. This indicates that, in the present simulation library, the selected region does not overlap with the sampled background. However, zero observed false positives does not justify quoting FPR $= 0$. With $N_{\text{test}} \sim 7 \times 10^2$ background events per split, the corresponding 95% confidence level upper limit is FPR $\lesssim 3/N_{\text{test}} \sim 4 \times 10^{-3}$ (Poisson zero-count approximation) [135]. Background rejection below the $\sim 10^{-3}$ level therefore cannot be validated with the current simulation statistics.

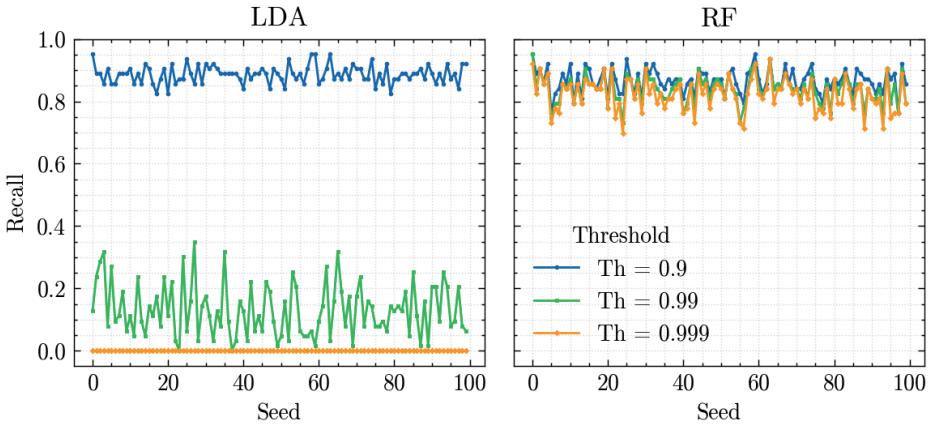


Figure D.4: Recall versus seed for LDA (left) and RF (right) at thresholds $t = 0.90$, 0.99, and 0.999.

Increasing the threshold from $t = 0.90$ to $t = 0.999$ reduces the RF recall from 0.871 to 0.842. LDA yields recall consistent with zero at $t = 0.999$ on the available test sets; with the present statistics, this cannot be interpreted as evidence against linear separability in the extreme tail. The stability of classifier performance across repeated stratified train/test splits (30% held-out test; all tuning and calibration performed within the training set via cross-validation) is shown in Figure D.4 and Figure D.5, which plot signal efficiency (TPR) and FPR versus seed for both models at all three thresholds.

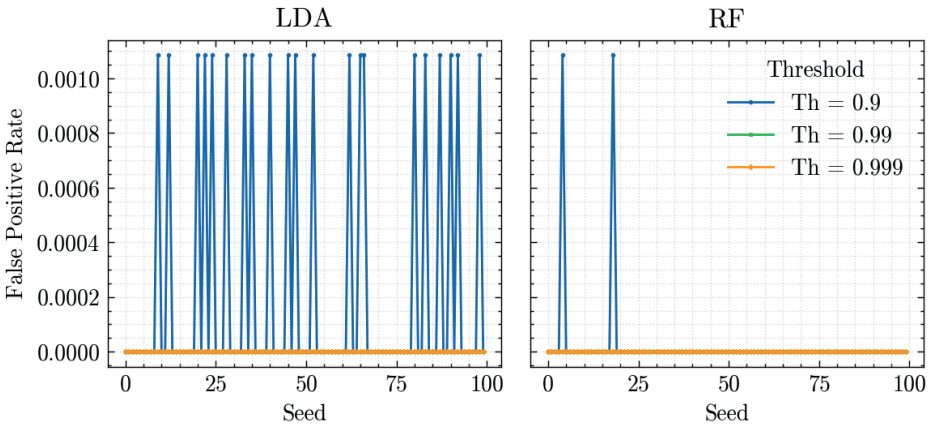


Figure D.5: False positive rate versus seed for LDA (left) and RF (right) at thresholds $t = 0.90$, 0.99, and 0.999, evaluated on the test sets of the repeated train/test splits. With $N_{\text{test}} \sim 7 \times 10^2$ background events per split, the per-split FPR is quantized in steps of $1/N_{\text{test}} \sim 1.4 \times 10^{-3}$. Non-zero points therefore correspond to one or a few leaked background events and are statistically volatile under changes of the random train/test split.

RF recall varies weakly across seeds (standard deviation $\sim$ 0.05). The FPR curves show occasional spikes at the $\sim 10^{-3}$ level, corresponding to one or a few misclassified background events in the finite held-out test set. Because the absolute number of leaked background events is very small, the exact seeds for which a non-zero FPR is observed are statistically volatile: rerunning the repeated train/test splits with different random seeds can change the total number of false positives by a few events. Such changes, for example from several leaked events to only a couple of leaked events across all splits, should not be interpreted as a meaningful change in classifier performance. They reflect the discreteness of the background counts and the limited size of the simulated background sample. LDA shows larger variability at $t = 0.99$.

Figure D.6 summarizes the efficiency–purity trade-off by plotting the mean signal efficiency (TPR) versus the mean false-positive rate (FPR) for the three thresholds, using the same 100 repeated stratified train/test splits reported in Table D.1. Because each split contains only $N_{\text{test}} \sim 7 \times 10^2$ background events, the per-split FPR is quantized in steps of $1/N_{\text{test}} \sim 10^{-3}$. Therefore, the near-zero mean FPR values in Table D.1 (and the corresponding points in Figure D.6) occur when most splits yield $k = 0$ false positives and should be interpreted as being consistent with an upper limit, rather than as a direct measurement of rejection at the 10^{-4}–10^{-5} level. In Figure D.6, the vertical error bars represent the seed-to-seed variation of the

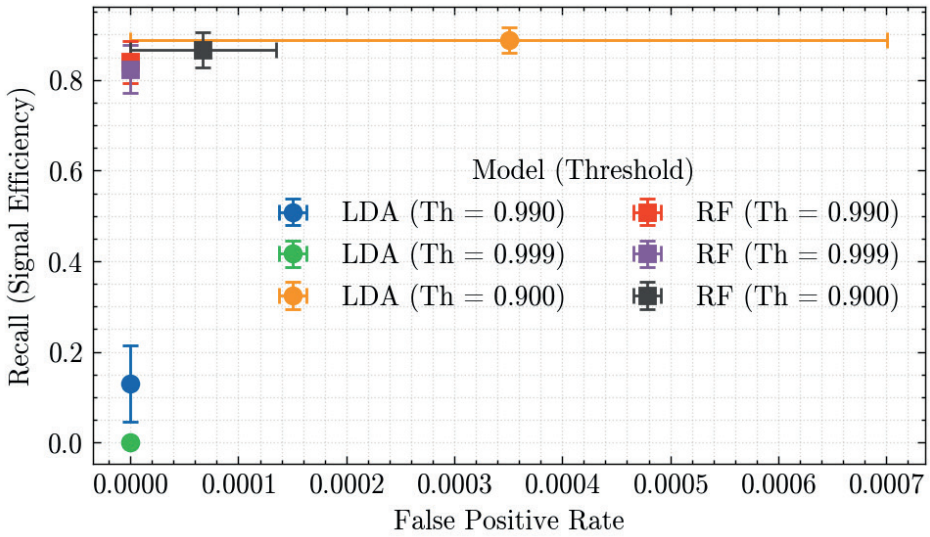


Figure D.6: Signal efficiency (recall) versus false positive rate for LDA and RF at various thresholds, evaluated on the test sets of the repeated train/test splits.

signal efficiency, while the horizontal error bars reflect the finite-background counting statistics (quantization and one-sided upper intervals when $k = 0$), not a Gaussian $\pm\sigma$ spread.

Per-event stability is quantified in Figure D.7, which shows, for each event, the standard deviation of the calibrated score $P(\text{signal})$ across the 100 random seeds. For both models, most events have small variability ($\sigma \lesssim 0.05$), indicating that the classification of the bulk of the sample is robust against the choice of train/test split. Larger fluctuations occur mainly at intermediate mean scores, consistent with events close to the decision boundary.

The RF exhibits a broader tail toward higher σ than LDA, reflecting the higher variance of a more flexible, non-linear model under limited training statistics.

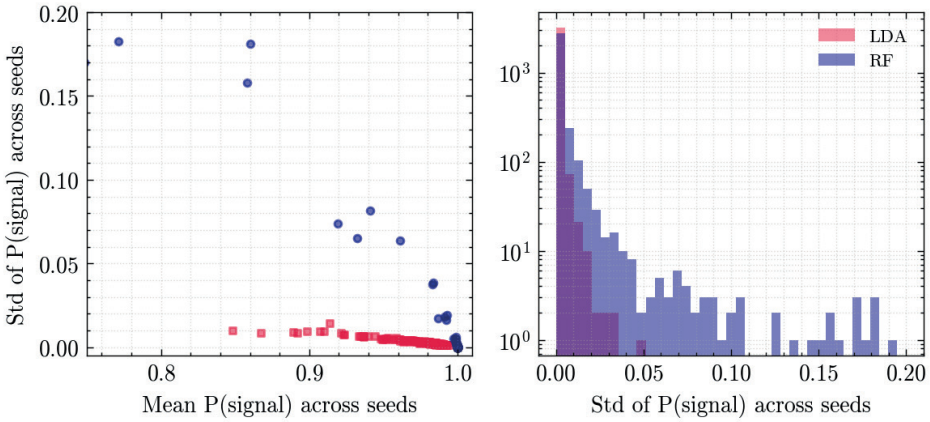


Figure D.7: Standard deviation of $P(\text{signal})$ across seeds versus mean $P(\text{signal})$ (left) and histogram of standard deviations (right) for LDA (red) and RF (blue).

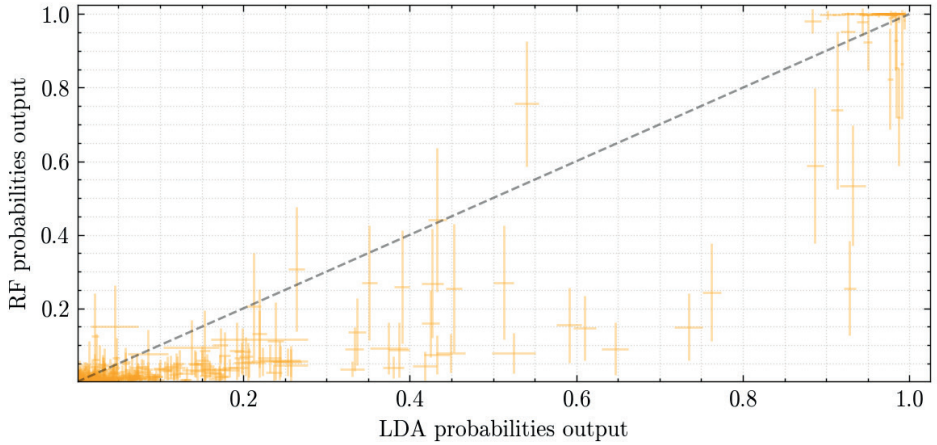


Figure D.8: LDA versus RF predicted signal probabilities per event (probability that this event is a photon). Error bars show variability across seeds.

Figure D.8 compares the predicted probabilities from both models for each event, with error bars indicating variability across seeds. The LDA assigns higher probabilities than the RF for many events. At probabilities close to 1.0, both models give similar predictions, while the RF shows a larger spread. The diagonal is a reference for equal probabilities; deviations are expected because LDA and RF are different models calibrated independently, with differences largest for borderline events.

This is consistent with isotonic calibration producing less overconfident probabilities in regions with limited training support [136].

Motivated by the energy-dependent selection effects identified in Section 7.1, we examine how performance varies with reconstructed radio energy. Figure D.9 shows mean recall as a function of $\log_{10}(E_{\mathrm{RD}}/\mathrm{eV})$ for both models at all three thresholds.

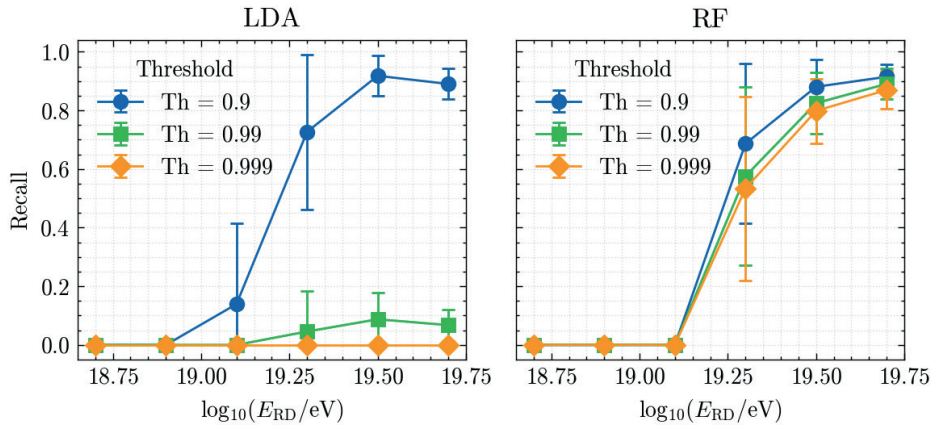


Figure D.9: Mean recall as a function of $\log_{10}(E_{\mathrm{RD}}/\mathrm{eV})$ for LDA (left) and RF (right) at thresholds $t = 0.90$, 0.99, and 0.999. Error bars are standard deviations across seeds.

At higher energies, the RF recall increases in the available simulation sample. Efficiency corrections as a function of energy are required when applying the classifier in a search.
To connect classifier performance with rare-signal conditions, we compute the precision assuming a prior photon fraction π:

$$\text{Precision} = \frac{\pi\,\text{TPR}}{\pi\,\text{TPR} + (1-\pi)\,\text{FPR}}.$$

Figure D.10 shows precision–recall points for assumed photon fractions $\pi = 10^{-3}$ and $\pi = 10^{-4}$. These points inherit the statistical limitations of the available background test sets: with $N_{\text{test}} \sim 10^3$ background events per split, false-positive rates below $\mathcal{O}(10^{-3})$ cannot be resolved. In particular, at $\pi = 10^{-4}$ the inferred precision is dominated by the background upper limit: when zero false positives are observed, the FPR entering the precision estimate is constrained only by an upper limit (rule-of-three), so the resulting high precision values should be interpreted as illustrative rather than as a validated background rejection at the 10^{-4} level.

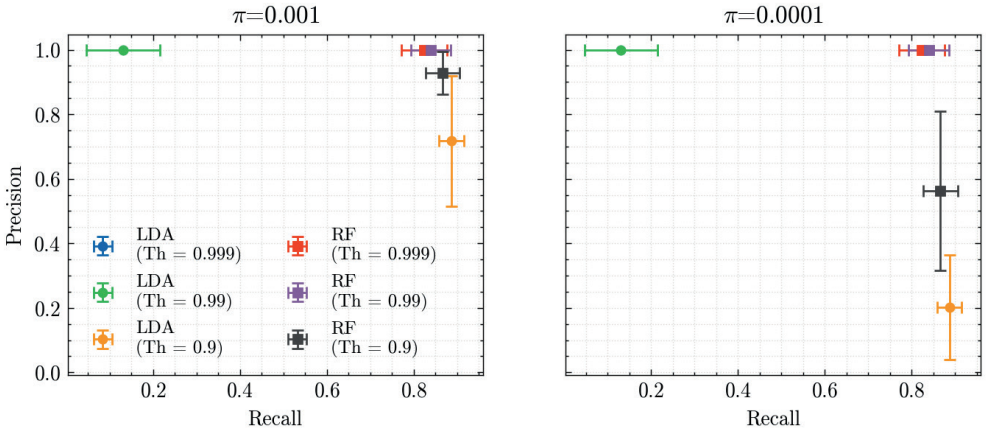


Figure D.10: Precision–recall points at photon fractions $\pi = 10^{-3}$ (left) and $\pi = 10^{-4}$ (right). Error bars reflect seed-to-seed variation of the TPR and the statistics-limited constraint on the FPR (upper limits when zero false positives are observed). LDA at $t = 0.999$ yields TPR$= 0$ and is not shown.

Overall, these results show that the RF classifier selects a photon-candidate region with $\sim 84\%$ signal efficiency on the available sample. The present background statistics are sufficient to demonstrate separation of the core photon and hadron distributions, but they do not constrain background leakage in the extreme tails. For a realistic photon fraction π, observing only a few photon candidates ($S \sim$ 1–3) requires the expected number of leaked background events B to remain $\mathcal{O}(1)$, i.e.

$$B \simeq \frac{S}{\pi\,\text{TPR}}\,\text{FPR} \lesssim 1 \quad \Rightarrow \quad \text{FPR} \lesssim \frac{\pi\,\text{TPR}}{S}.$$

With TPR $\simeq 0.84$ and $\pi \sim 10^{-4}$ this implies FPR $\lesssim$ (3–8) $\times\, 10^{-5}$ for $S =$ 3–1, i.e. background rejection at the 10^{-5}–10^{-4} level. Larger background simulation sets are therefore required to

validate leakage rates in this regime and to support statements on rejection below 10^{-4}.

Outlook: Toward Application on Real Data

The classification framework developed in this appendix demonstrates that hybrid WCD+RD energy information can be used in a multivariate photon–hadron discriminator at large zenith angles. On simulated data, the Random Forest classifier achieves 0.842 ± 0.050 signal efficiency at threshold $t = 0.999$, with no background events observed in the test samples. Since each test split contains only $\mathcal{O}(10^3)$ background events, this zero-count result corresponds to an upper limit on the false-positive rate of order $3/N_{\text{test}} \sim \text{few} \times 10^{-3}$, rather than a direct validation of rejection at the 10^{-4} level or below.

This performance is comparable to the cut-based analysis adopted in Chapter 7. There, the reference cut on the detrended energy-mismatch observable, $\delta' < -0.5653$, retains $185/230 \simeq 0.80$ photons and gives zero observed background events in the post-threshold simulated background sample, $0/2303$. The finite-sample constraint is therefore $3/2303 \simeq 1.3 \times 10^{-3}$, while the smaller quoted leakage estimate of 1.83×10^{-4} comes from the model-dependent Gaussian-tail extrapolation of the background δ' distribution. At the stricter nominal 10^{-4} tail-leakage working point, the cut-based photon efficiency is $143/230 \simeq 0.62$.

Thus, within the statistical reach of the present simulations, the multivariate analysis does not provide a decisive improvement over the one-dimensional δ' cut. The RF classifier gives a photon efficiency comparable to the 80% cut-based working point, but its false-positive rate is limited by the smaller held-out test samples and is not validated in the extreme background tail. Moreover, the ML classifier uses δ' as its dominant physics input, with event-by-event uncertainty information providing supplementary separation. Any gain over the direct cut should therefore be interpreted as incremental rather than as an independent separation mechanism. For the final analysis, the direct cut on δ' is preferred because it is physically transparent, uses the full post-threshold background sample for the empirical leakage check, and is simpler to validate on real data. The multivariate classifier is kept as a complementary proof-of-concept study, with possible future value if larger simulated or data-driven background samples become available.

All performance metrics presented here rely on Monte Carlo simulations. Before applying such a classifier to the real Auger dataset, validation on a limited real-data sample is essential to:

- verify that simulation-derived efficiencies hold for real detector responses, including time-dependent calibration effects and station-level systematics not fully captured in MC;
- confirm that background rejection remains stable when confronted with the full diversity of hadronic air showers in nature, which may differ from the SIBYLL and QGSJET models used in training;
- establish credibility for any photon candidates by demonstrating that classification thresholds were not tuned post-hoc based on observed results.

To this end, a real test-sample analysis protocol has been proposed to the Pierre Auger Collaboration. This protocol defines a statistically representative subset of real data that remains blinded during classifier development and is unblinded only once for validation, following established practices in particle-physics searches. If the unblinded sample shows significant disagreement with the simulation expectations, the classifier is not tuned on this sample. Instead, the discrepancy is treated as evidence of mismodelling and motivates a dedicated investigation of reconstruction effects and systematic uncertainties.

Appendix E: Data Management

This thesis research has been carried out under the institute research data management policy of the Institute for Mathematics, Astrophysics and Particle Physics as documented in https://www.ru.nl/sites/default/files/2023-05/imapp_rdm_policy_1.pdf.

- The data, CORSIKA simulations and software used in this thesis are available within the Pierre Auger Collaboration.
- Raw data used in chapter 3, 5, 6 and 8 are stored on the Pierre Auger/CNCZ computing infrastructure at:
  /vol/augerprime/data/pauger/Malargue/Raid/data/
  while the extracted radio data are stored at:
  /vol/augerprime/data/RdDataTraces/pauger/Malargue/Raid/data/.
- The raw simulation sets used in chapter 4 and 7 are stored on the Pierre Auger/CNCZ computing infrastructure at:
  /vol/augerprime/simulations/pauger/Radio/Simulations/library.
- The reconstructed data and simulation products used in this thesis are too large to be archived in the GitLab repositories, with several datasets exceeding 100 GB. They are therefore stored on the Pierre Auger/CNCZ computing infrastructure and are available internally to collaboration members. Since user-specific cluster paths are not guaranteed to be persistent, the analysis code, configuration files, file lists, and scripts required to reproduce the derived data products are archived in the GitLab repository listed below.
- The reconstructed simulation sets used in chapter 4 and 7 were produced from the raw simulation library listed above. The production scripts, configuration files, and file lists required to reproduce these samples are archived in the analysis-code repository listed below.
- The reconstructed data used in chapter 5 and 6 were produced from the raw data listed above. The reconstruction and selection procedures are documented through the scripts, configuration files, and file lists archived in the analysis-code repository listed below.
- All analysis codes, together with configuration files and dataset file lists needed to reproduce the thesis results, can be found at:
  https://gitlab.iap.kit.edu/auger-observatory/sandboxes/memam94

- Some of the codes use H. Schoorlemmer's packages, which can be found at: https://gitlab.iap.kit.edu/auger-observatory/sandboxes/hschoorlemmer

Bibliography

[1] Pedro Abreu et al. Searches for Ultra-High-Energy Photons at the Pierre Auger Observatory. *Universe*, 8(11):579, 2022.

[2] K. Greisen. End to the Cosmic-Ray Spectrum? *Phys. Rev. Lett.*, 16(17):748–750, 1966.

[3] P. Abreu et al. Cosmological implications of photon-flux upper limits at ultra-high energies in scenarios of Planckian-interacting massive particles for dark matter. *Phys. Rev. D*, 107(4):042002, 2023.

[4] Pedro Abreu et al. Testing effects of Lorentz invariance violation in the propagation of astroparticles with the Pierre Auger Observatory. *JCAP*, 01(01):023, 2022.

[5] O. Adriani et al. PAMELA Measurements of Cosmic-ray Proton and Helium Spectra. *Science*, 332:69–72, 2011.

[6] M. Aguilar et al. Precision Measurement of the Proton Flux in Primary Cosmic Rays from Rigidity 1 GV to 1.8 TV with the Alpha Magnetic Spectrometer on the International Space Station. *Phys. Rev. Lett.*, 114:171103, 2015.

[7] S. Mollerach and E. Roulet. Progress in high-energy cosmic ray physics. *Prog. Part. Nucl. Phys.*, 98:85–118, 2018.

[8] Alexander Aab et al. Features of the Energy Spectrum of Cosmic Rays above 2.5×10^{18} eV Using the Pierre Auger Observatory. *Phys. Rev. Lett.*, 125(12):121106, 2020.

[9] Johannes Blümer, Ralph Engel, and Jörg R. Hörandel. Cosmic rays from the knee to the highest energies. *Prog. Part. Nucl. Phys.*, 63:293–338, 2009.

[10] B. Peters. Primary cosmic radiation and extensive air showers. *Il Nuovo Cimento*, 22(4):800–819, nov 1961.

[11] Andreas Haungs et al. Investigating the 2nd knee: The KASCADE-Grande experiment. *J. Phys. Conf. Ser.*, 47:238, 2006.

[12] J. Abraham et al. Measurement of the energy spectrum of cosmic rays above 10^{18} eV using the Pierre Auger Observatory. *Phys. Lett. B*, 685:239–246, 2010.

[13] A. Abdul Halim et al. Inference of the Mass Composition of Cosmic Rays with Energies from $10^{18.5}$ to 10^{20} eV Using the Pierre Auger Observatory and Deep Learning. *Phys. Rev. Lett.*, 134(2):021001, 2025.

[14] V. Berezinsky, A. Gazizov, and S. Grigorieva. On astrophysical solution to ultrahigh-energy cosmic rays. *Phys. Rev. D*, 74(4):043005, 2006.

[15] G. T. Zatsepin and V. A. Kuzmin. Upper Limit of the Spectrum of Cosmic Rays. *Journal of Experimental and Theoretical Physics Letters*, 4:78–80, 1966.

[16] J. L. Puget, F. W. Stecker, and J. H. Bredekamp. Photonuclear interactions of ultrahigh energy cosmic rays and their astrophysical consequences. *Astrophys. J.*, 205:638–654, apr 1976.

[17] A. M. Hillas. The Origin of Ultra-High-Energy Cosmic Rays. *Annual Review of Astronomy and Astrophysics*, 22:425–444, 1984.

[18] Karl-Heinz Kampert and Alan A. Watson. Extensive Air Showers and Ultra High-Energy Cosmic Rays: A Historical Review. *The European Physical Journal H*, 37:359–412, 2012.

[19] Ralph Engel, Dieter Heck, and Tanguy Pierog. Extensive Air Showers and Hadronic Interactions at High Energy. *Annual Review of Nuclear and Particle Science*, 61:467–489, 2011.

[20] Pedro Abreu et al. Measurement of the proton-air cross-section at $\sqrt{s} = 57$ TeV with the Pierre Auger Observatory. *Phys. Rev. Lett.*, 109:062002, 2012.

[21] Kumiko Kotera and Angela V. Olinto. The Astrophysics of Ultrahigh-Energy Cosmic Rays. *Annual Review of Astronomy and Astrophysics*, 49:119–153, 2011.

[22] James Matthews. A Heitler model of extensive air showers. *Astroparticle Physics*, 22:387–397, 2005.

[23] Tanguy Pierog. The Impact of Accelerator Data on our Understanding of Extensive Air Showers. In *Proceedings of the 39th International Cosmic Ray Conference — PoS(ICRC2025)*, 2025.

[24] T. K. Gaisser and A. M. Hillas. Reliability of the method of constant intensity cuts for reconstructing the average development of vertical showers. In *Proceedings of the 15th International Cosmic Ray Conference*, volume 8, pages 353–357, Plovdiv, Bulgaria, 1977.

[25] P. K. F. Grieder. Extensive Air Showers: High Energy Phenomena and Astrophysical Aspects. *Springer-Verlag Berlin Heidelberg*, 1-2, 2010.

[26] M. Unger, B. R. Dawson, R. Engel, F. Schüssler, and R. Ulrich. Reconstruction of longitudinal profiles of ultra-high energy cosmic ray showers from fluorescence and Cherenkov light measurements. *Nucl. Instrum. Meth. A*, 588(3):433–441, 2008.

[27] H. M. J. Barbosa, F. Catalani, J. A. Chinellato, and C. Dobrigkeit. Determination of the calorimetric energy in extensive air showers. *Astroparticle Physics*, 22(2):159–166, 2004.

[28] Xavier Bertou, Pierre Billoir, and Sylvie Dagoret-Campagne. LPM effect and pair production in the geomagnetic field: a signature of ultra-high energy photons in the Pierre Auger Observatory. *Astroparticle Physics*, 14(2):121–130, 2000.

[29] P. Homola, Markus Risse, R. Engel, D. Gora, J. Pekala, B. Wilczynska, and H. Wilczynski. Characteristics of geomagnetic cascading of ultrahigh energy photons at the southern and northern sites of the Pierre Auger Observatory. *Astroparticle Physics*, 27:174–184, 2007.

[30] T. Huege. Radio detection of cosmic ray air showers in the digital era. *Phys. Rept.*, 620:1–52, 2016.

[31] J. Schulz. *Cosmic Radiation. Reconstruction of Cosmic-Ray Properties from Radio Emission of Extensive Air Showers*. Phd thesis, Radboud University Nijmegen, 2016.

[32] Tomáš Fodran. *Commissioning the Auger Prime Radio Detector: Paving the Way for Mass Separation*. Phd thesis, Radboud University Nijmegen, September 2024.

[33] Felix Andreas Schlüter. *Expected sensitivity of the AugerPrime Radio Detector to the masses of ultra-high-energy cosmic rays using inclined air showers*. PhD thesis, KIT, Karlsruhe, IAP, 2022.

[34] D. Ardouin et al. Geomagnetic origin of the radio emission from cosmic ray induced air showers observed by CODALEMA. *Astroparticle Physics*, 31(3):192–200, 2009.

[35] G. A. Askaryan. Excess negative charge of an electron-photon shower and its coherent radio emission. *Soviet Physics JETP*, 14:441–443, 1962.

[36] Alexander Aab et al. Probing the Radio Emission from Air Showers with Polarization Measurements. *Phys. Rev. D*, 89(5):052002, 2014.

[37] A. Nelles et al. Measuring a Cherenkov ring in the radio emission from air showers at 110–190 MHz with LOFAR. *Astroparticle Physics*, 65:11–21, 2015.

[38] T. Huege, R. Ulrich, and R. Engel. Energy and composition sensitivity of geosynchrotron radio emission from cosmic ray air showers. *Astroparticle Physics*, 30:96, 2008.

[39] A. Abdul Halim et al. Radio Measurements of the Depth of Air-Shower Maximum at the Pierre Auger Observatory. *Phys. Rev. D*, 109(2):022002, 2024.

[40] Pierre Auger Collaboration. The scintillator surface detector of the pierre auger observatory. *Journal of Instrumentation*, 20:P08002, 2025.

[41] Pierre Auger Collaboration. The Pierre Auger Cosmic Ray Observatory. *Nucl. Instrum. Meth. A*, 798:172–213, 2015.

[42] A. M. Hillas. Evolution of ground-based gamma-ray astronomy from the early days to the Cherenkov Telescope Arrays. *Astroparticle Physics*, 43:19–43, 2013.

[43] Pedro Abreu et al. Antennas for the Detection of Radio Emission Pulses from Cosmic-Ray. *JINST*, 7:P10011, 2012.

[44] Frank G. Schröder. Radio detection of cosmic-ray air showers and high-energy neutrinos. *Prog. Part. Nucl. Phys.*, 93:1–68, 2017.

[45] H. Tokuno et al. New air fluorescence detectors employed in the Telescope Array experiment. *Nucl. Instrum. Meth. A*, 676:54–65, 2012.

[46] Pierre Auger Collaboration. Deep-learning based reconstruction of the shower maximum using the water-cherenkov detectors of the pierre auger observatory. *Journal of Instrumentation*, 16:P07019, 2021.

[47] Karl-Heinz Kampert and Michael Unger. Measurements of the cosmic ray composition with air shower experiments. *Astroparticle Physics*, 35(10):660–678, 2012.

[48] Benjamin Fuchs. The Auger Engineering Radio Array. *Nucl. Instrum. Meth. A*, 692:93–97, 2012. 3rd Roma International Conference on Astroparticle Physics.

[49] Alexander Aab et al. Measurement of the cosmic-ray mass composition with the Pierre Auger Observatory. *Phys. Rev. D*, 100(8):082002, 2019.

[50] A. Aab et al. Depth of maximum of air-shower profiles at the Pierre Auger Observatory. I. Measurements at energies above $10^{17.8}$ eV. *Phys. Rev. D*, 90:122005, 2014.

[51] A. Abdul Halim et al. Measurement of the depth of maximum of air-shower profiles with energies between $10^{18.5}$ and 10^{20} eV using the surface detector of the Pierre Auger Observatory and deep learning. *Phys. Rev. D*, 111(2):022003, 2025.

[52] R. U. Abbasi et al. Mass composition of ultra-high-energy cosmic rays with the Telescope Array Surface Detector data. *Phys. Rev. D*, 97:102005, 2018.

[53] Dan Hooper, Andrew Taylor, and Subir Sarkar. The Impact of heavy nuclei on the cosmogenic neutrino flux. *Astroparticle Physics*, 23:11–17, 2005.

[54] Lyndon Evans and Philip Bryant. LHC Machine. *JINST*, 3:S08001, 2008.

[55] Pierre Auger Collaboration. Measurement of the proton-air cross section at $\sqrt{s} = 57$ TeV with the Pierre Auger Observatory. *Phys. Rev. D*, 102:062005, 2019.

[56] Andrew M. Taylor. The need for a local source of UHECR nuclei. *Astroparticle Physics*, 54:48–53, 2015.

[57] J. Abraham et al. An upper limit to the photon fraction in cosmic rays above 10^{19} eV from the Pierre Auger Observatory. *Astroparticle Physics*, 27:155–168, 2007.

[58] Luis A. Anchordoqui. Ultra-high energy cosmic rays. *Phys. Rept.*, 801:1–93, 2019.

[59] S. Navas et al. Review of particle physics. *Phys. Rev. D*, 110(3):030001, 2024.

[60] Markus Ahlers et al. GZK neutrinos after the Fermi-LAT diffuse photon flux measurement. *Astroparticle Physics*, 34(2):106–115, 2010.

[61] P. Mészáros. Gamma-ray Bursts. *Reports on Progress in Physics*, 69(8):2259–2322, 2006.

[62] R. D. Blandford and R. L. Znajek. Electromagnetic Extraction of Energy from Kerr Black Holes. *Monthly Notices of the Royal Astronomical Society*, 179:433–456, 1977.

[63] Marco Chianese, Damiano F. G. Fiorillo, Rasmi Hajjar, Gennaro Miele, and Ninetta Saviano. Constraints on heavy decaying dark matter with current gamma-ray measurements. *JCAP*, 11:035, 2021.

[64] R. Aloisio, S. Matarrese, and A. V. Olinto. Super Heavy Dark Matter in Light of BICEP2, Planck and Ultra High Energy Cosmic Rays Observations. *JCAP*, 08:024, 2015.

[65] Rafael Alves Batista, Andrej Dundovic, Martin Erdmann, Karl-Heinz Kampert, Daniel Kuempel, Gero Müller, Guenter Sigl, Arjen van Vliet, David Walz, and Tobias Winchen. CRPropa 3 - a Public Astrophysical Simulation Framework for Propagating Extraterrestrial Ultra-High Energy Particles. *JCAP*, 05:038, 2016.

[66] Markus Risse and Piotr Homola. Search for ultra-high-energy photons using air showers. *Modern Physics Letters A*, 22(11):749–766, 2007.

[67] A. Aab et al. Search for photons with energies above 10^{18} eV using the Pierre Auger Observatory. *JCAP*, 04:009, 2017.

[68] P. Homola, M. Risse, R. Engel, D. Góra, J. Pe˛kala, B. Wilczyńska, and H. Wilczyński. Characteristics of geomagnetic cascading of ultra-high energy photons at the southern and northern sites of the Pierre Auger Observatory. *Astroparticle Physics*, 27(2):174–184, 2007.

[69] A. Abdul Halim et al. Search for photons above 10^{18} eV by simultaneously measuring the atmospheric depth and the muon content of air showers at the Pierre Auger Observatory. *Phys. Rev. D*, 110(6):062005, 2024.

[70] P. Abreu et al. Search for Photons above 10^{19} eV with the Surface Detector of the Pierre Auger Observatory. *JCAP*, 05:021, 2023.

[71] Biswajit Sarkar, Karl-Heinz Kampert, Joerg Kulbartz, Luca Maccione, Nils Niersten-hoefer, and Guenter Sigl. Ultra-High Energy Photon and Neutrino Fluxes in Realistic Astrophysical Scenarios. In *32nd International Cosmic Ray Conference, Beijing 2011*. ICRC2011, 2011.

[72] Graciela B. Gelmini, Oleg Kalashev, and Dmitri Semikoz. Upper Limit on the Diffuse Radio Background from GZK Photon Observation. *Universe*, 8(8):402, 2022.

[73] Anna Bobrikova, Marcus Niechciol, Markus Risse, and Philip Ruehl. Predicting the UHE photon flux from GZK-interactions of hadronic cosmic rays using CRPropa 3. *PoS*, ICRC2021:449, 2021.

[74] Alexander Aab et al. Erratum: Search for photons with energies above 10^{18} eV using the hybrid detector of the Pierre Auger Observatory. *JCAP*, 2020(9):E02, Sep 2020.

[75] L. A. Anchordoqui et al. Hunting super-heavy dark matter with ultra-high energy photons. *Astroparticle Physics*, 132:102614, 2021.

[76] P. Abreu et al. Cosmological implications of photon-flux upper limits at ultra-high energies in scenarios of Planckian-interacting massive particles for dark matter. *Phys. Rev. D*, 107(042002), 2023.

[77] P. Abreu et al. Limits to gauge coupling in the dark sector set by the non-observation of instanton-induced decay of Super-Heavy Dark Matter in the Pierre Auger Observatory data. *Phys. Rev. Lett.*, 130(061001), 2023.

[78] O. E. Kalashev and M. Yu. Kuznetsov. Constraining heavy decaying dark matter with the high energy gamma-ray limits. *Phys. Rev. D*, 94:063535, Sep 2016.

[79] Alexander Aab et al. A targeted search for point sources of EeV photons with the Pierre Auger Observatory. *Astrophys. J. Lett.*, 837(2):L25, 2017.

[80] Alexander Aab et al. A Search for Point Sources of EeV Photons. *Astrophys. J.*, 789(2):160, 2014.

[81] Adila Abdul Halim et al. Search for Ultra-high-energy Photons from Gravitational Wave Sources with the Pierre Auger Observatory. *Astrophys. J.*, 952(1):91, 2023.

[82] T. Hermann-Josef Mathes. The HEAT telescopes of the Pierre Auger Observatory: Status and first data. In *32nd International Cosmic Ray Conference*, volume 3, page 153, 8 2011.

[83] P. Abreu et al. The energy spectrum of cosmic rays beyond the turn-down around 10^{17} eV as measured with the surface detector of the Pierre Auger Observatory. *Eur. Phys. J. C*, 81(11):966, 2021.

[84] M. Unger, B. R. Dawson, R. Engel, F. Schüssler, and R. Ulrich. Reconstruction of longitudinal profiles of ultra-high energy cosmic ray showers from fluorescence and cherenkov light measurements. *Nuclear Instruments and Methods in Physics Research A*, 588:433–441, 2008.

[85] Pierre Auger Collaboration. Data-driven estimation of the invisible energy of cosmic ray showers with the pierre auger observatory. *Physical Review D*, 100:082003, 2019.

[86] J. Abraham et al. The Fluorescence Detector of the Pierre Auger Observatory. *Nucl. Instrum. Meth. A*, 620:227–251, 07 2010.

[87] I. Allekotte et al. The Surface Detector System of the Pierre Auger Observatory. *Nucl. Instrum. Meth. A*, 586:409–420, 2008.

[88] Pierre Auger Collaboration. The Pierre Auger Cosmic Ray Observatory. *Nucl. Instrum. Meth. A*, 798:172–213, 2015.

[89] Pierre Auger Collaboration. Reconstruction of inclined air showers detected with the Pierre Auger Observatory. *JCAP*, 2014(8):019, aug 2014.

[90] Valerio Verzi. The Energy Scale of the Pierre Auger Observatory. In *Proceedings of the 33rd International Cosmic Ray Conference (ICRC 2013)*, page 0928, Rio de Janeiro, Brazil, 2013.

[91] Alexander Aab et al. The Pierre Auger Observatory upgrade-preliminary design report. *arXiv preprint arXiv:1604.03637*, 2016.

[92] Allan Machado Payeras and Anderson Campos Fauth. Use of a small photomultiplier tube to extend the dynamic range of the surface detectors of the pierre auger observatory. *Proceedings of Science*, BHCB2018:002, 2019.

[93] Mart Pothast. Ultra-high-energy cosmic particle identification. Master's thesis, Radboud University, 2022.

[94] Harm Schoorlemmer. *Tuning in on Cosmic Rays: Polarization of Radio Signals from Air Showers as a Probe of Emission Mechanisms*. PhD thesis, Radboud University Nijmegen, Nijmegen, The Netherlands, 2012.

[95] Tim Huege. The Radio Detector of the Pierre Auger Observatory – status and expected performance. *EPJ Web Conf.*, 283:06002, 2023.

[96] Alexander Aab et al. Observation of inclined EeV air showers with the radio detector of the Pierre Auger Observatory. *JCAP*, 10:026, 2018.

[97] K. Louedec. Radio-detection of extensive air showers at the Pierre Auger Observatory - Results and next enhancements. In *Proceedings of The European Physical Society Conference on High Energy Physics (EPS-HEP 2013). 18-24 July*, page 27, jan 2013.

[98] Jörg R. Hörandel. Precision measurements of cosmic rays up to the highest energies with a large radio array at the Pierre Auger Observatory. In *European Physical Journal Web of Conferences*, volume 210 of *European Physical Journal Web of Conferences*, page 06005, oct 2019.

[99] The Radio Group. The science case for the Radio Upgrade of the Pierre Auger Observatory. *Pierre Auger Observatory internal GAP notes*, (04), 2020.

[100] Pierre Auger Collaboration. Trigger and aperture of the surface detector array of the Pierre Auger Observatory. *Nucl. Instrum. Meth. A*, 613(1):29–39, 2010.

[101] Jim Pivarski, Pratyush Das, Chris Burr, Dmitri Smirnov, Matthew Feickert, Tamas Gal, Luke Kreczko, Nicholas Smith, Noah Biederbeck, Oksana Shadura, Mason Proffitt, benkrikler, Hans Dembinski, Henry Schreiner, Jonas Rembser, Marcel R., Chao Gu, Jonas Rübenach, Michele Peresano, and Ruggero Turra. https://github.com/scikit-hep/uproot: 3.12.0, jul 2020.

[102] Alex Johnson, Étienne Tétreault-Pinard, and Mojtaba Samimi. Open source Plotly charting library, 2026.

[103] Bokeh Development Team. *Bokeh: Python library for interactive visualization*, 2018.

[104] Bjarni Pont. RD Bad Periods. *Pierre Auger Observatory internal GAP notes*, (42), 2024.

[105] G. Rodriguez, L. Cazon, V. M. Olmos-Gilbaja, R. A. Vazquez, and E. Zas. Simulation of the Surface Detector response for inclined shower reconstruction. *Pierre Auger Observatory internal GAP notes*, (49), 2009.

[106] H. P. Dembinski et al. Report on Tests and Measurements of Hadronic Interaction Properties with Air Showers. *EPJ Web Conf.*, 210:02004, 2019.

[107] A. Abdul Halim et al. Search for photons above 10^{18} eV by simultaneously measuring the atmospheric depth and the muon content of air showers at the Pierre Auger Observatory. *Phys. Rev. D*, 110(6):062005, 2024.

[108] S. Argirò, S.L.C. Barroso, J. Gonzalez, L. Nellen, T. Paul, T.A. Porter, L. Prado Jr., M. Roth, R. Ulrich, and D. Veberič. The offline software framework of the Pierre Auger Observatory. *Nucl. Instrum. Meth. A*, 580(3):1485–1496, 2007.

[109] Pedro Abreu, M Aglietta, EJ Ahn, IFM Albuquerque, D Allard, I Allekotte, J Allen, P Allison, J Alvarez Castillo, J Alvarez-Muñiz, et al. Advanced functionality for radio analysis in the Offline software framework of the Pierre Auger Observatory. *Nucl. Instrum. Meth. A*, 635(1):92–102, 2011.

[110] Jannis Pawlowsky. Simulation Study of the Detection of Inclined Photon Air Showers with the AugerPrime Radio Detector. *PoS*, ARENA2022:054, 2023.

[111] Tomáš Fodran. *Commissioning the Auger Prime Radio Detector*. PhD thesis, Radboud University Nijmegen, Nijmegen, The Netherlands, 2025.

[112] P. Abreu et al. Expected performance of the Auger Radio Detector. *PoS*, ARENA2022:028, 2023.

[113] Carla Bonifazi, A. Letessier-Selvon, E. M. Santos, et al. A model for the time uncertainty measurements in the Auger surface detector array. *Astroparticle Physics*, 28(6):523–528, 2008.

[114] Giuseppe De Mauro. *Tasting the ultra-high-energy universe: Weighing in on mass composition measurement of ultra-high-energy cosmic rays*. PhD thesis, Radboud University Nijmegen, Nijmegen, The Netherlands, oct 2020.

[115] P. Abreu et al. Advanced functionality for radio analysis in the Offline software framework of the Pierre Auger Observatory. *Nucl. Instrum. Meth. A*, 635(1):92–102, 2011.

[116] Lars Mohrmann. Measurement of Radio Emission from Cosmic Ray Induced Air Showers at the Pierre Auger Observatory with a Spherical Wave Reconstruction. Masterarbeit in physik, Fakultät für Mathematik, Informatik und Naturwissenschaften, Rheinisch-Westfälische Technische Hochschule Aachen, Aachen, Germany, sep 2011.

[117] Felix Schlüter and Tim Huege. Signal model and event reconstruction for the radio detection of inclined air showers. *JCAP*, 2023(01):008, 2023.

[118] Gerald J. Burke and Andrew J. Poggio. Numerical Electromagnetics Code (NEC-2): Program Description and User's Guide (Parts I–III). Technical report ucid-18834, Lawrence Livermore National Laboratory, 1981. Contains Part I (Theory, nec2prt1.pdf), Part II (Code, nec2prt2.pdf) and Part III (User's Guide, nec2prt3.pdf).

[119] Ugo Gregorio Giaccari et al. Simulations of the antenna response for the Auger Radio Detector. *PoS*, ARENA2022:042, 2023.

[120] Stefan Grebe. *Finger on the pulse of cosmic rays: dependence of the radio pulse shape on the air shower geometry*. PhD thesis, Radboud University Nijmegen, sep 2013.

[121] Stefan Jansen. *Radio for the masses: Cosmic ray mass composition measurements in the radio frequency domain*. PhD thesis, Radboud University Nijmegen, apr 2016.

[122] S. Buitink et al. Amplified radio emission from cosmic ray air showers in thunderstorms. *Astron. Astrophys.*, 467:385–394, 2007.

[123] Max Büsken. A new network of electric field mills at the Pierre Auger Observatory. *J. Phys. Conf. Ser.*, 2398(1):012004, 2022.

[124] B. Pont. Radio Detector Thunderstorm Flag. `https://gitlab.iap.kit.edu/auger-observatory/sandboxes/bpont/RDWeatherTSFlag`, 2026.

[125] Ewa M. Holt, Frank G. Schröder, and Andreas Haungs. Enhancing the cosmic-ray mass sensitivity of air-shower arrays by combining radio and muon detectors. *Eur. Phys. J. C*, 79(5):371, 2019.

[126] Maria Clementina Medina, M. Gomez Berisso, I. Allekotte, A. Etchegoyen, G. Medina Tanco, and A. D. Supanitsky. Enhancing the Pierre Auger Observatory to the 10^{17}–$10^{18.5}$ eV Range: Capabilities of an Infill Surface Array. *Nucl. Instrum. Meth. A*, 566:302–311, 2006.

[127] Alan A. Watson. Highlights from the Pierre Auger Observatory, the Birth of the Hybrid Era. In *30th International Cosmic Ray Conference*, volume 6, pages 67–78, 7 2007.

[128] Alexander Schulz. *Measurement of the Energy Spectrum and Mass Composition of Ultra-high Energy Cosmic Rays*. PhD thesis, Karlsruhe Institute of Technology (KIT), 2016. Auger internal note GAP 2016-021.

[129] Gary J. Feldman and Robert D. Cousins. A Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998.

[130] Jaime Alvarez-Muñiz and Marcus Niechciol. A Burnt Sample for Neutral-Particle Searches with the SD in Phase II. *Pierre Auger Observatory internal GAP notes*, (003), 2026.

[131] Jannis Pawlowsky. *Search for inclined photon induced air showers with the AugerPrime Radio Detector*. Dissertation, Bergische Universität Wuppertal, Wuppertal, Germany, oct 2024.

[132] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[133] R. A. Fisher. The Use of Multiple Measurements in Taxonomic Problems. *Annals of Eugenics*, 7:179–188, 1936.

[134] Leo Breiman. Random Forests. *Machine Learning*, 45(1):5–32, 2001.

[135] Pierpaolo Savina. A search for ultra-high-energy photons at the Pierre Auger Observatory exploiting air-shower Universality. In *Proceedings of Science, 37th International Cosmic Ray Conference (ICRC2021)*, page 373, 2021.

[136] Bianca Zadrozny and Charles Elkan. Transforming Classifier Scores into Accurate Multiclass Probability Estimates. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2002)*, pages 694–699, 2002.

Use of Generative AI

During the preparation and revision of this thesis, generative AI tools were used in a limited supporting role. The author engaged with these tools with awareness of their limitations and risks, and used them critically and selectively throughout the preparation of the thesis.
They were used to help reorganize and refine draft text written by the author, improve grammar, style, readability, paragraph structure, and overall coherence, identify redundancy, and detect unclear or orphan definitions in the manuscript. They also helped the author interpret supervisors' comments and explore possible ways of revising the text in line with that feedback. Generative AI tools were also used to improve the readability and structure of code written by the author for generating figures.
However, these tools were not used as a source of scientific evidence, nor did they replace the author's own scientific analysis, judgement, interpretation, or methodological decision-making. All AI-assisted output was critically reviewed, edited, and verified by the author before inclusion in the thesis. The underlying calculations, figures, code, methodological choices, results, and conclusions were produced and validated independently by the author. No personal data, sensitive information, or confidential research material were entered into these systems.
The generative AI applications used for these support tasks were ChatGPT (including GPT-4.5 and GPT-5.2) and Gemini (including Gemini 2.5).

Acknowledgements

Of everything in this thesis, these acknowledgements were the hardest to write. I can't really capture in words what these people and moments have meant to me, so I'll keep it short. They belong here because these years would not have been the same without them.

I want to thank my supervisors, Sijbrand and Harm. Beyond their scientific guidance and the many hours of discussion, they supported me as a person. They understood my situation, gave me space when I needed it, and stood by me when things were difficult. Their trust in me is what brought me to this point.

Anthony, I'm especially grateful for the friendship we built over these years. We both came from Africa and were building a life far from home, and that gave us a connection from the very beginning. I valued our conversations, which often went well beyond research. Thank you for being part of my family life too. I appreciate the time you spent with us, the kindness you showed my daughter, and the thoughtful gifts you brought from Zambia, some of which have a special place in our home.

Mohit, thank you for your company and kindness. I appreciated our conversations at work, as well as the time you spent with my family. My daughter still loves watching the video of you playing with her, and it always brings a smile to our faces.

Mart, although our time at Radboud overlapped only briefly, I always appreciated how helpful and welcoming you were.

I have only a few close friends, and they mean a lot to me. Even when they were not part of my everyday PhD life, they stayed in the background, asking how things were going and reminding me that there was a world outside the office. Sara, Jackline, and Mohamed Ibrahim: thank you for staying close, for checking in, and for being there through the different stages of my life.

To my brother Mahmoud, my sister Marwa, and my whole family: thank you for the support that is difficult to summarize in a few sentences. Distance, work, and life have taken us in different directions, but knowing that you were there has always mattered.

Lilas, you have been part of this journey from almost the very beginning. We met around the time I was finishing my Master's, only a few months after I lost my mother, and you were there, standing and applauding, when I defended my thesis. Not long after that, we moved together to the Netherlands. What I thought was going to be another academic step became four years of building a home, a family, and a life together. Throughout these years, you have been the person beside me, supporting me, taking care of the things around me when I needed the space to work, and making many difficult periods easier than they otherwise would have been. Now,

as I approach another defense, after losing my father only a few months ago, I find myself in a strangely familiar place. And once again, you are beside me. I could write much more, but I think I will leave it at that: thank you for being there through all of it.

For me, these years contained science, frustration, learning, family, friendship, moving between countries, becoming a husband and a father, and saying goodbye. So I finish this thesis with mixed feelings, however, I am proud to have reached the end.

Curriculum Vitae

Mohamed Ismaiel was born on January 21, 1994, in Jeddah, Saudi Arabia. He studied Space Science at Helwan University in Egypt, receiving his B.Sc. degree in 2015. During his undergraduate studies, he developed interests in experimental physics, space weather, and cosmic rays. His bachelor's thesis investigated the temperature dependence of the electrical properties of semiconductor materials, while other projects introduced him to scintillator detectors and Resistive Plate Chambers.

After graduation, Mohamed combined research with teaching. Between 2015 and 2019, he taught physics and mathematics within the International Baccalaureate system and also contributed to university-level teaching in astrophysics and fundamental physics, including electromagnetism and classical mechanics. During the same period, his research interests increasingly shifted toward cosmic rays and space weather. In 2016, under the mentorship of Prof. Rolf Bütikofer of the University of Bern, he studied the Ground Level Enhancement event GLE-69 using neutron-monitor data. He later contributed to research on Forbush decreases associated with coronal mass ejections. He also worked as a freelance graphic designer, maintaining an interest in visual communication alongside his scientific work.

In 2019, Mohamed moved to Italy to pursue a Master's degree in Astrophysics at the University of Padova. After working on exoplanet characterization, he joined research connected to the GERDA experiment for his Master's thesis, supervised by Prof. Riccardo Brugnera and Dr. Mariia Redchuk. His work focused on pulse-shape discrimination in germanium detectors used in searches for neutrinoless double-beta decay. He completed his Master's degree in 2021.

In December 2021, Mohamed joined the High Energy Physics department at Radboud University in Nijmegen as a doctoral researcher under the supervision of Prof. Sijbrand de Jong and Dr. Harm Schoorlemmer. As a member of the Pierre Auger Collaboration, he worked on searches for ultra-high-energy photons using hybrid measurements of extensive air showers with Water-Cherenkov Detectors and radio antennas. Alongside his research, he contributed to teaching several physics courses. He completed his doctoral research in December 2025.

Following his Ph.D., Mohamed moved into industry, applying his experience in programming, data analysis, and scientific problem solving to large and complex datasets.

Outside work, he enjoys football and chess and is a long-time follower of *One Piece*, whose visual world inspired elements of the artistic concept of this dissertation. His scientific journey has taken him from space weather and cosmic rays, through underground neutrino experiments, to the search for photons at the highest energies observed in nature.

The Photon That Wasn't There

Yesterday, in the data stream,
I sought a photon, so it would seem.
It wasn't there again today,
I wish, I wish it'd come my way...

When I scanned the tanks at three,
Just muons were crowding round for me,
But when I checked the radio trace,
I couldn't find it in its place!
I searched, I searched through every shower ...
Oh, photon, show yourself this hour.
Go away, go away, protons, don't you mimic any more!
Go away, go away, and let a photon through the door ...

Through inclined events, slant and rare,
I looked for signs that should be there:
A quiet muon trace below,
A radio footprint wide in glow.
But when the final cuts were through,
Not one good photon made it through.

And somewhere far beyond our sight,
A cosmic source may send its light.
Perhaps tomorrow, in the data stream,
We'll catch the particle of the dream.
Until that day, we search the air
For the photon that wasn't there ...

