

Likelihood reconstruction for radio detectors of neutrinos and cosmic rays

Martin Ravn¹, Christian Glaser^{1,2}, Thorsten Glüsenskamp^{1,3}, Ayca Özcelikkale⁴, and Alan Coleman¹

¹*Department of Physics and Astronomy, Uppsala University, Box 516, S-75120 Uppsala, Sweden*

²*Department of Physics, TU Dortmund University, D-44221 Dortmund, Germany*

³*Oskar Klein Centre and Department of Physics, Stockholm University, SE-10691 Stockholm, Sweden*

⁴*Department of Electrical Engineering, Uppsala University, Box 65, 75103 Uppsala, Sweden*



(Received 4 November 2025; accepted 11 June 2026; published 11 August 2026)

Ultra-high-energy neutrinos and cosmic rays are excellent probes of astroparticle physics phenomena. For astroparticle physics analyses, robust and accurate reconstruction of signal parameters such as arrival direction and energy is essential. Radio detection is an established detector concept explored by many observatories; however, current reconstruction methods ignore bin-to-bin noise correlations, which limits reconstruction resolution and, so far, has prevented calculations of event-by-event uncertainties. In this work, we present a likelihood description of neutrino or cosmic-ray signals in radio detectors with correlated noise, as present in all neutrino and cosmic-ray radio detectors. We demonstrate, with simulation studies of both neutrinos and cosmic-ray radio signals, that signal parameters such as energy and direction, including event-by-event uncertainties with correct coverage, can be obtained. This method reduces reconstruction uncertainties and biases compared to previous approaches. Additionally, the Likelihood can be used for event selection and enables differentiable end-to-end detector optimization. The reconstruction code is available through the open-source software NuRadioReco.

DOI: [10.1103/m22t-cbgw](https://doi.org/10.1103/m22t-cbgw)

I. INTRODUCTION

Radio detection of ultra-high-energy (UHE) cosmic rays and neutrinos is a novel technique with the prospect of substantially advancing astroparticle physics in the coming years. For cosmic rays, the radio technique has been successfully used for many years at different air-shower observatories such as the Pierre Auger Observatory and LOFAR [1,2]. Its strengths include accurate energy estimation with small systematic uncertainties, as the emission can be calculated from first principles, insensitivity of the detector system to aging and environmental/atmospheric changes [3–8], as well as precise measurement of the position of the shower maximum if the shower footprint is sampled with sufficient fidelity [9–12]. With the Auger-Prime upgrade, all surface detector stations of the Pierre Auger Observatory have been upgraded with radio antennas with the aim of measuring the chemical composition of cosmic rays at the highest energies [13,14]. It is planned to enhance the IceTop air shower array of the IceCube detector at the South Pole with radio antennas [15,16] and to equip the Square Kilometer Array (SKA) with cosmic-ray capabilities for precision air-shower physics [17].

For neutrinos, the radio detection technique enables cost-effective instrumentation of large enough volumes to measure the faint flux of UHE neutrinos [18] through ground-based air-shower arrays to detect Earth-skimming neutrinos [19–21], balloon missions [22,23], and through arrays of in-ice radio detector stations [16,24–27]. The latter, which will be the focus of this paper, was explored in the pilot arrays RICE [24,28], ARA [25,29–31], and ARIANNA [26,32–38] at the South Pole and Moore’s Bay. With RNO-G [27], a substantially larger detector, with the potential of measuring the first UHE neutrino, is currently being constructed in Greenland, and an order-of-magnitude more sensitive in-ice radio array is planned for IceCube-Gen2 [16,39].

All of these detectors have the same key challenge: Extracting the signal parameters from time series with low signal-to-noise ratios (SNRs). This key challenge will be addressed in this paper. We present, for the first time, a full likelihood reconstruction which promises to improve reconstruction performance. Furthermore, due to our correct probabilistic treatment, the likelihood not only estimates the true signal parameters but also provides the probability of any signal hypothesis, i.e., event-by-event uncertainties or the posterior probability density functions with arbitrary correlations between signal parameters. In the following paragraphs, we provide a brief summary of the current status of existing reconstruction methods before we present our work.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Funded by Bibsam.

A. Reconstructing radio signals from air showers

Air shower arrays such as the Pierre Auger Observatory, LOFAR, or GRAND typically measure the radio signal on the ground with several dual-polarized antennas separated by $\mathcal{O}(100\text{ m})$ (an exception is LOFAR, which has a much denser sampling). Some antennas will have SNRs significantly above the noise floor, but due to the rapidly falling lateral signal distribution [40–42], antennas further away from the shower axis will have small SNR. However, these antennas are crucial to determine the electromagnetic energy content and mass-sensitive parameters such as X_{max} , and L/R (see e.g. [3,9,43]). So far, these analyses do not use the full information of the recorded time series but calculate a scalar variable, e.g., the energy fluence for two different polarization components of the signal.

To improve on this, the *forward folding* technique [44–46] was developed, which uses a signal model whose predicted electric field is forward folded through the antenna and detector response and compared to the measured time series by subtracting the predicted from the measured signal in each time bin and adding the squares, i.e., calculating the chi-square (χ^2). By minimizing this objective function, the best-fitting model parameters can be determined. It was found that the *forward folding* method improved the reconstruction performance significantly at low SNRs [46]. This method was successfully used to reconstruct air showers with the ARIANNA detector [36], and a simulation study predicted sensitivity to the cosmic-ray arrival direction and energy with a single detector station [47].

We note that the signal model is ideally an *end-to-end* signal model that directly predicts the time series at each antenna for the parameters of interest such as cosmic-ray energy, direction, and X_{max} . However, such a model is often difficult to formulate, has theoretical uncertainties, or is computationally very expensive. Alternatively, a more low-level reconstruction can be performed, where the electric field is reconstructed from each dual-polarized antenna, i.e., the signal arriving at the antennas. It turns out that the electric field can be parametrized analytically with high accuracy with just a few parameters, allowing a largely model-independent extraction of the signal parameters at each antenna location [46].

B. Reconstructing radio signals in dense media

In-ice radio detectors for neutrinos are built as compact detector stations, each consisting of several (single polarization) antennas of different types with different orientations and positions with small separations of $\mathcal{O}(10\text{ m})$. By drilling holes in the ice, antennas can be lowered up to $\mathcal{O}(100\text{ m})$ beneath the surface, allowing for three-dimensional detector layouts. The stations are separated by $\mathcal{O}(1\text{ km})$ to observe mostly independent ice volumes and thus to maximize the total effective volume of the detector. Hence, each detector station must be able to *independently* determine the relevant

neutrino parameters. As a consequence, most observed neutrinos will have SNRs close to the trigger threshold. Even more, there are many ongoing efforts to further reduce trigger thresholds as much as possible, as an increased neutrino detection rate substantially improves the scientific outcome [29,37,48,49]. Hence, extracting the signal parameters at low SNRs is of the utmost importance.

The station design of in-ice detectors includes several single polarization antennas with $\mathcal{O}(1\text{m}–10\text{m})$ spacing and thus each antenna observes a different electric field. Thus, unlike for the dual-polarized antennas that are used in air shower arrays, a reconstruction of the electric field per antenna is not uniquely defined. Hence, to use the *forward folding* technique, a model is needed that simultaneously predicts the signal in all antennas. Fortunately, this is readily available via parametrizations or semianalytic calculations of the *Askaryan emission* (see [50] for an overview of the available models) that predict the electric field for any observer position (relative to the in-ice particle shower), shower energy, and interaction type. Effects from propagating the signal through ice are added, after which the signal is folded with the antenna and detector response to obtain the time series that every antenna is expected to observe. Simulation studies show that the *forward folding* technique can be used to estimate the neutrino direction [35,51–53] and it achieves the best resolution of all traditional methods, i.e., not considering deep learning.

The forward folding method has been shown to reliably obtain best-fit parameters of radio neutrino and cosmic-ray signals; however, published implementations of the method have not developed a complete likelihood description and have therefore not been able to provide reliable event-by-event uncertainties. The main reason for this is that the noise observed in radio detectors has strong correlations between neighboring bins, but these correlations are ignored in current reconstruction methods. For instance, in the case of the forward folding method, when calculating the difference between data and model prediction bin-by-bin, the sum of the differences squared divided by the variance of the noise can not be interpreted as a χ^2 . The main missing ingredient in properly accounting for these correlations and obtaining event-by-event uncertainties is a probabilistic description of the noise observed in radio detectors, which we will develop in this paper. In general, a measured trace consists of a signal plus one realization of probabilistic noise. A probabilistic description of the noise will then allow us to calculate the likelihood of a signal, given the measured trace. Using the likelihood in a reconstruction is the most robust and powerful way of extracting the signal parameters from a measurement if the signal model can be fully specified.

We note that in practice, it can be challenging to fully specify the signal model and its dependence on the parameters of interest, especially for nondeterministic cases such as high-energy electron neutrino interactions where the

LPM effect introduces significant shower-to-shower fluctuations [54–56]. In such cases, better performance might be achievable with deep learning. Either only the signal model is parametrized via neural networks and used in a likelihood fit (as successfully done by IceCube to discover neutrino emission from the Galactic plane [57]), or a deep neural network is directly trained to predict the parameters of interest [58]. By incorporating conditional normalizing flows into neural networks, the network can predict the probability of any model parameter including correlations, just as in the likelihood method [59] (which was recently demonstrated for IceCube [60] and in-ice radio detectors [61]). In this paper, we do not consider deep learning.

In this paper, we present a probabilistic description of the noise observed in radio neutrino or cosmic ray detectors. The model can be used to express a likelihood for a neutrino or cosmic-ray radio signal, given a set of measured traces in an arbitrary number of antennas, which can be used in a reconstruction of signal parameters. In Sec. II, we describe the probabilistic noise model and demonstrate its compatibility with simulated data. In Sec. III, we demonstrate the various applications of the noise model: neutrino signal reconstruction for an in-ice radio detector (Sec. III A), fast uncertainty estimation for differentiable end-to-end detector optimization (Sec. III B), reconstruction of electric fields from cosmic ray air showers with a dual-polarized antenna (Sec. III C), and identification of cosmic-ray signals in an in-ice radio detector (Sec. III D).

II. PROBABILISTIC NOISE MODEL

In this section, we describe the noise present in radio detectors and propose how to describe it probabilistically. We discuss various theoretical aspects of the proposed model and verify it against simulated noise data.

A. Noise in radio detectors of cosmic rays and neutrinos

The noise present in all radio detectors falls into two main categories: anthropogenic noise and noise natural to the environment or the detector. Anthropogenic noise consists primarily of spurious pulsed signals. They are, in general, considered to be an identifiable background and can be differentiated from the target physics signals in event selections. Additionally, many modern radio detectors are built in remote places that are sufficiently radio-quiet that noise pulses are rare and unlikely to overlap with the neutrino and cosmic-ray signals that only last a few nanoseconds. Therefore, we do not need to model pulsed noise for event reconstruction. We note that pulsed noise is still a challenge for self-triggering and identifying the few cosmic rays and neutrinos from all triggered events. Only the latter is addressed in this paper. Anthropogenic noise from radio broadcasting and communication can form continuous signals in detectors; however, these often span very narrow frequency ranges or single frequencies and can

be filtered out in the frequency domain, e.g., by notch filters.

In contrast, natural noise sources form a diffuse stochastic background that is continuously present in the detector across all frequencies. The main contributions are thermal radio noise from the environment, radio noise from the Milky Way galaxy [62], and noise from the detector electronics themselves. The environmental thermal noise and galactic noise are observed through the antennas and are hence folded with the antenna response, whereas the electronic noise is added in various steps in the signal chain and is heavily dependent on the hardware, e.g., amplifiers (see e.g. [63]). In total, these sources form a continuous background that can not be eliminated but is well understood and can be characterized.

Noise in radio detectors of cosmic rays and neutrinos has been studied extensively in the past. After basic data cleaning (e.g., high- and low-pass filtering, removal of frequency bins contaminated with continuous wave signals), the noise can be described as bandwidth-limited Gaussian noise. A procedure was developed to simulate this noise in the Fourier domain (see, e.g., the implementation in NuRadioReco [46]), which can be transformed back into the time domain via an inverse Fourier transform. Studies of noise from several radio neutrino experiments show Rayleigh-distributed frequency amplitudes (equivalent to our assumption of bandwidth-limited Gaussian noise, see below) [64–67]. In addition, multiple analyses have shown excellent agreement between measurements and simulations, e.g., [37,38], where a deep neural network trained on simulated data performed as expected on experimental data from the ARIANNA detector, providing indirect evidence that the noise has been simulated correctly. We describe the noise generation method and our notation in the following.

B. Notation and formalism

A radio detector measures the voltages in a number of antennas as a function of time, which we denote $V(t_n)$. The data are sampled at discrete times t_n , called bins or samples, and recorded in traces of n_t consecutive measurements. We have $t_n = t_0 + n \times \delta t$, for $n = 0, \dots, n_t - 1$ where δt is the time between the samples, and $T = n_t \times \delta t$ is the duration of the trace. In the derivation below, it is assumed that n_t is even and t_0 is set as $t_0 = 0$ for convenience. The discrete Fourier transform of a trace $V(t_0), \dots, V(t_{n_t-1})$ is given by

$$\mathcal{V}(\omega_k) = \sqrt{2\delta t} \left(\sum_{n=0}^{n_t-1} V(t_n) e^{-i\omega_k t_n} \right), \quad (1)$$

where $\mathcal{V}(\omega_k)$ is the complex amplitude of the Fourier mode corresponding to the frequency $\omega_k = k \frac{2\pi}{T}$, where $-(\frac{n_t}{2} - 1) \leq k \leq \frac{n_t}{2}$. Since $V(t)$ is real-valued, there is Hermitian symmetry in the Fourier spectrum, i.e., $\mathcal{V}(\omega_k) = \mathcal{V}^*(-\omega_k)$ where $*$ denotes complex conjugate. Hence,

we only need to calculate $\mathcal{V}(\omega_k)$ from $k = 0$ to $k = n_t/2$, i.e., the Nyquist frequency. The trace can be recovered from the spectrum using the following inverse discrete Fourier transform:

$$\begin{aligned} V(t_n) &= \frac{1}{\sqrt{2\delta t}} \cdot \frac{1}{n_t} \sum_{k=-(n_t/2-1)}^{n_t/2} \mathcal{V}(\omega_k) e^{i\omega_k t_n} \\ &= \frac{1}{\sqrt{2\delta t}} \cdot \frac{1}{n_t} \left(\mathcal{V}(\omega_0) e^{i\omega_0 t_n} + \mathcal{V}(\omega_{n_t/2}) e^{i\omega_{n_t/2} t_n} \right. \\ &\quad \left. + 2 \sum_{k=1}^{n_t/2-1} \Re\{\mathcal{V}(\omega_k) e^{i\omega_k t_n}\} \right), \end{aligned} \quad (2)$$

where we used Hermitian symmetry to represent $V(t_n)$ using only the non-negative frequencies.

The normalization factor $\sqrt{2\delta t}$ in Eq. (1) ensures that the units in the frequency domain is voltage per frequency, and that the energy is conserved under reasonable assumptions that approximately hold in practice in a typical scenario. In particular, we consider the following conditions: (A1) there is no bias in the noise, i.e., the average value of the noise trace (when there is no signal present) is zero so that $\frac{1}{n_t} \sum_{n=0}^{n_t-1} V(t_n) = 0$, hence $\mathcal{V}(0) = 0$; (A2) $\mathcal{V}(\omega_{n_t/2}) = 0$, i.e., the sampling rate is sufficiently high so that the nonzero part of the noise spectrum is captured by sampling, and hence the contribution from the Fourier mode at the Nyquist frequency is zero. Then, the energy is conserved as follows:

$$\sum_{n=0}^{n_t-1} |V(t_n)|^2 \delta t = \sum_{k=0}^{n_t/2} |\mathcal{V}(\omega_k)|^2 \delta f, \quad (3)$$

where $\delta f = \frac{1}{n_t \delta t}$ is the length of the frequency bins. Under (A1) and (A2), Eq. (2) can be conveniently rewritten as

$$\begin{aligned} V(t_n) &= \frac{\sqrt{2}}{n_t \delta t} \sum_{k=1}^{n_t/2-1} [\Re\{\mathcal{V}(\omega_k)\} \cos(\omega_k t_n) \\ &\quad - \Im\{\mathcal{V}(\omega_k)\} \sin(\omega_k t_n)]. \end{aligned} \quad (4)$$

The noise generation procedure is then as follows: first, a real-valued spectrum $A(\omega_k)$ is chosen for the noise generation process. This spectrum can either be a theoretical spectrum or be measured directly from data. Second, the noise realization is generated in the frequency domain by pulling the real and imaginary parts of $\mathcal{V}(\omega_k)$ from Gaussian distributions with zero mean and equal variance of $\frac{1}{2}(A(\omega_k))^2$, statistically independent over the frequency index k , and over real and imaginary parts for a given k . We note that for $k = 0$ and $k = n_t/2$, purely real $\mathcal{V}(\omega_k)$ is formed using a Gaussian distribution with variance $(A(\omega_k))^2$ to satisfy the condition $V(t_n)$ is real-valued. This procedure results in Rayleigh-distributed amplitudes and uniformly distributed phases for $\mathcal{V}(\omega_k)$ for

$1 \leq k \leq \frac{n_t}{2} - 1$. The relationship between $A(\omega_k)$ and the expectations of the Fourier mode amplitudes is given by $E[|\mathcal{V}(\omega_k)|^2] = A(\omega_k)^2$. Finally, the noise is inverse Fourier transformed to the time domain according to Eq. (2), or equivalently according to Eq. (4) if $A(\omega_0) = 0$ and $A(\omega_{n_t/2}) = 0$. This formalism of generating noise is the established standard in the *in-ice radio* community, the same as that used in NuRadioMC [50], and it was verified numerous times to describe experimentally measured noise (see e.g. [37,38]).

If the spectrum is chosen to be a box function extending from 0 to the Nyquist frequency, the generated noise will be uncorrelated Gaussian noise, also called white noise. If, however, a band-limited spectrum with a cutoff at low and/or high frequencies is chosen, the generated noise will have correlations between neighboring samples. The latter is the case for the noise observed in radio detectors. We show a few examples of different noise scenarios in Fig. 1 with an injected simulated neutrino signal. The first row shows white noise, the second row shows unrealistically strongly band-limited noise, and the last row shows a realistic noise spectrum. The realistic spectrum is based on the averaged noise spectrum for a single RNO-G antenna presented in Fig. 3 in Ref. [68]. Additionally, we added two notch filters to demonstrate the robustness of our method in cases where anthropogenic continuous-wave (CW) signals need to be filtered out. We note that the method presented in this section works for any noise spectrum with Rayleigh-distributed amplitudes, as long as the spectrum is known or can be measured.

In any problem where successive measurements are Gaussian and uncorrelated, the standard objective function to minimize in order to fit a curve, e.g., $\mu(t_n; \theta)$, parametrized by $\theta = (\theta_0, \dots, \theta_{n_\theta})$, to a measured dataset, e.g., $V(t_n)$, is the chi-square,

$$\chi^2 \equiv \sum_{n=0}^{n_t-1} \frac{[V(t_n) - \mu(t_n; \theta)]^2}{\sigma^2}, \quad (5)$$

where σ is the standard deviation of the noise. However, if this objective function is used for bandwidth-limited noise, all information about correlations is ignored, and it will not be chi-square distributed for repeated measurements. It can therefore not be interpreted statistically like a chi-square, and it can not be used to estimate uncertainties on reconstructed parameters. This is the reason why previous work on reconstruction using the forward folding technique, i.e., a chi-square reconstruction, did not predict any event-by-event uncertainties [35,36,46,51–53]. To obtain an objective function that can be interpreted statistically, the correct probabilistic description of band-limited noise is needed. We present such a model in this paper, which allows us to calculate the likelihood of a signal in a radio detector given a measured trace.




FIG. 1. Different noise spectra and resulting time domain noise realizations with an injected simulated neutrino signal. The realistic spectrum is based on the averaged noise spectrum from a single RNO-G antenna presented in Ref. [68], to which we added two notch filters.

C. Probabilistic noise model

We model the noise plus a deterministic signal in a radio detector as a multivariate normal distribution in n_t dimensions. The mean of each bin is the signal, and the deviation is a random number from a normal distribution correlated with the neighboring bins. By representing the trace and signal as vectors, $\mathbf{x} = (V(t_0), \dots, V(t_{n_t-1}))^T$ and $\boldsymbol{\mu}(\boldsymbol{\theta}) = (\mu(t_0; \boldsymbol{\theta}), \dots, \mu(t_{n_t-1}; \boldsymbol{\theta}))^T$, the probability density function for measuring a trace given a signal is then

$$p(\mathbf{x}; \boldsymbol{\mu}(\boldsymbol{\theta}), \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^{n_t} |\boldsymbol{\Sigma}|}} \times \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))\right), \quad (6)$$

where $\boldsymbol{\Sigma}$ is the covariance matrix of the noise and $|\boldsymbol{\Sigma}|$ is its determinant. Here, it is assumed that $\boldsymbol{\Sigma}$ is invertible, whereas the noninvertible case is discussed in Sec. II E. The covariance matrix parameterizes the correlations between the noise in different time bins, and we discuss how to determine it in the following sections. The multivariate normal distribution hence describes correlated noise and

can be used to calculate the probability that a trace belongs to the expected distribution, potentially including a neutrino or cosmic ray signal.

The probability density function in Eq. (6) can then be seen as the likelihood for a signal given a measured trace. For statistical analyses, minus 2 times the natural logarithm of the likelihood is of special significance,

$$-2 \ln \mathcal{L}(\boldsymbol{\mu}(\boldsymbol{\theta}); \mathbf{x}, \boldsymbol{\Sigma}) = n_t \ln(2\pi) + n_t \ln(|\boldsymbol{\Sigma}|) + (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta})). \quad (7)$$

This is the ideal objective function to minimize in a reconstruction, and we denote the maximum likelihood estimate of the signal parameters as $\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} (-2 \ln \mathcal{L}(\boldsymbol{\mu}(\boldsymbol{\theta}); \mathbf{x}, \boldsymbol{\Sigma}))$. For this purpose, the first two constant terms can be ignored, and the log-likelihood can easily be summed over the different channels of a radio detector if the signal is present in several antennas. Additionally, this objective function, and differences in the objective function between two hypotheses,

$$-2\Delta \ln \mathcal{L}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) = -2[\ln \mathcal{L}(\boldsymbol{\mu}(\boldsymbol{\theta}_2); \mathbf{x}, \boldsymbol{\Sigma}) - \ln \mathcal{L}(\boldsymbol{\mu}(\boldsymbol{\theta}_1); \mathbf{x}, \boldsymbol{\Sigma})], \quad (8)$$

can be interpreted statistically since they are chi-square distributed according to Wilks' theorem [69]. It can hence be used to derive event-by-event uncertainties on the reconstructed parameters and estimate a goodness-of-fit for the reconstructed signal. The main goal of this paper is to demonstrate that the likelihood can be used for reconstructing signals across the field of radio detection with correct event-by-event uncertainties and that it improves upon the reconstruction uncertainties compared to previous methods.

In the following, we will discuss how to calculate the covariance matrix and its inverse and show the model's compatibility with simulated noise. All aspects discussed here are implemented in a noise model class available in NuRadioReco, which includes methods for calculating the likelihood for any radio signal with Gaussian noise.

D. The covariance matrix

The covariance of measurements at two different times, t_i and t_j , is defined as

$$\text{Cov}(t_i, t_j) = \mathbb{E}[(V(t_i) - \mu_i)(V(t_j) - \mu_j)], \quad (9)$$

where $\mathbb{E}[\cdot]$ denotes the expectation, i.e., statistical average, operator and $\mu_i = \mathbb{E}[V(t_i)]$ and $\mu_j = \mathbb{E}[V(t_j)]$ are the averages of the measurements at t_i and t_j . In practice, the expectations, $\hat{\mu}_i$ and $\hat{\mu}_j$, can be calculated as the average over N independent realizations of traces containing only noise, either measured or simulated. Hence, the empirical covariance values can be estimated numerically from the traces as follows:

$$\widehat{\text{Cov}}(t_i, t_j) = \frac{1}{N} \sum_{m=1}^N (x_{m,i} - \hat{\mu}_i)(x_{m,j} - \hat{\mu}_j), \quad (10)$$

where $x_{m,i}$ and $x_{m,j}$ are the m^{th} realizations or measurement of $V(t_i)$ and $V(t_j)$, and $\hat{\mu}_i$ and $\hat{\mu}_j$ are expected to be zero for traces containing purely noise.

When $i = j$ the covariance reduces to the variance, $\text{Cov}(t_i, t_i) = \text{Var}(t_i)$. The covariance matrix, Σ , then has entries $\Sigma_{ij} = \text{Cov}(t_i, t_j)$ and has the structure,

$$\Sigma = \begin{pmatrix} \text{Var}(t_0) & \text{Cov}(t_1, t_0) & \cdots & \text{Cov}(t_{n_t-1}, t_0) \\ \text{Cov}(t_0, t_1) & \text{Var}(t_1) & \cdots & \text{Cov}(t_{n_t-1}, t_1) \\ \vdots & \vdots & \ddots & \vdots \\ \text{Cov}(t_0, t_{n_t-1}) & \text{Cov}(t_1, t_{n_t-1}) & \cdots & \text{Var}(t_{n_t-1}) \end{pmatrix}. \quad (11)$$

Note that $\text{Cov}(t_i, t_j) = \text{Cov}(t_j, t_i)$, $\forall t_i, t_j$; hence the covariance matrices are always symmetric.

For the noise we are considering here, the variance is assumed to be constant within the time period under consideration. Hence, the diagonal of the covariance matrix

is constant. For band-limited noise, measurements close to each other (small $\Delta t = t_j - t_i$) are expected to be strongly correlated, and the correlations should decrease for longer time intervals. The entries of the covariance matrix close to the diagonal should thus be large and decrease going away from the diagonal. In Fig. 2 (top left and top right), we show an example of a typical covariance matrix. We calculated the covariance matrix according to Eq. (10) from 2,000 traces of simulated band-limited thermal noise using the realistic spectrum shown in Fig. 1. The left panel shows the full covariance matrix, and the right panel shows a zoom-in of the relevant features close to the diagonal.

We assume that the correlation between the values of noise at two time instants solely depends on how close these two instants are; i.e., $\text{Cov}(t_i, t_j)$ only depends on $|t_i - t_j|$. As a result of this translation invariance in time and the fact that the traces are obtained using equidistant samples, the covariance matrix has a symmetric Toeplitz structure. In particular, the values on the diagonals are constant, and the matrix is symmetric around the main diagonal. A symmetric Toeplitz matrix has the following structure:

$$\mathbf{T}_s = \begin{pmatrix} c_0 & c_1 & \cdots & c_{n_t-2} & c_{n_t-1} \\ c_1 & c_0 & c_1 & \cdots & c_{n_t-2} \\ \vdots & c_1 & c_0 & \ddots & \vdots \\ c_{n_t-2} & \vdots & \ddots & \ddots & c_1 \\ c_{n_t-1} & c_{n_t-2} & \cdots & c_1 & c_0 \end{pmatrix}. \quad (12)$$

When constructing an estimate of the covariance matrix from traces, we need to make assumptions about the structure of the covariance matrix. The symmetric Toeplitz matrix assumption is attractive due to both the intuitively appealing translational invariance and also due to the fact that the statistical uncertainty on the estimate of the covariance matrix can be reduced by averaging over the diagonals.

We can make another simplification by observing that the formalism to generate noise (as described above) will result in periodic signals; i.e., the last time bins are correlated with the first time bins of the trace. This is because we represent a real-valued finite-duration sampled time series using a finite number of modes in the Fourier domain. Although this periodicity is unphysical and not present in data, it is not of any practical concern because of the following reasons. The bin-to-bin correlations of noise are local; i.e., they only span over a few time bins as we will show below. Hence, ignoring the first and last few time bins will remove the unphysical correlations. Furthermore, the expected neutrino and cosmic-ray signals are short and well contained within a trace. Therefore, it is common practice in a typical reconstruction pipeline to switch between representing the measurement in the time domain and the Fourier domain, depending on what is more appropriate for the analysis step [46]. Any potential problem can be further mitigated by filtering the edges of

the trace in the time domain to zero, e.g., using a Hanning window, and prepending/appending the traces with zeros. Hence, representing and modifying signals in the Fourier domain using a finite number of Fourier modes, which may implicitly make the trace cyclic, is not a practical problem, but simplifies many processing steps, and we will adopt it in the following.

Hence, the structure of the covariance matrix can be further restricted as a symmetrical circulant matrix, which is a special case of the Toeplitz matrix. It has the structure,

$$\mathbf{C}_s = \begin{pmatrix} c_0 & c_1 & \cdots & c_2 & c_1 \\ c_1 & c_0 & c_1 & \cdots & c_2 \\ \vdots & c_1 & c_0 & \ddots & \vdots \\ c_2 & \vdots & \ddots & \ddots & c_1 \\ c_1 & c_2 & \cdots & c_1 & c_0 \end{pmatrix}. \quad (13)$$

The symmetric circulant matrix has half the number of parameters as the Toeplitz matrix, and the statistical uncertainty of each entry is thus decreased.

Assuming that the covariance matrix is circulant mainly affects its upper right and lower left corners, where the correlations between the first and last few bins are described. We experience that for long traces, the decreased statistical uncertainty outweighs any negative effects from assuming periodicity. In this paper, we thus use the symmetrical circulant matrix to describe the covariance matrix. A zoomed in part of the covariance matrix averaged over the diagonals according to the circulant structure is shown in Fig. 2 (bottom left). Due to the symmetry of the matrix, the covariance matrix can be represented in 1D as the value of the covariance as a function of the distance to the diagonal, i.e., the time difference between the two bins. We show this representation in Fig. 2 (bottom right). We observe what also corresponds to our intuition that the variance ($\Delta t = 0$) is the largest value and that the magnitude of the correlation quickly decays with increasing time separation. We also see that neighboring bins have a positive correlation, but for slightly larger time separation, the bins are negatively correlated.

E. Inversion of the covariance matrix and the pseudoinverse

In order to evaluate the multivariate normal probability of a trace given a signal [Eq. (6)], the inverse of the covariance matrix is required. In cases where the covariance matrix is of full rank, the inverse can simply be calculated with standard inversion techniques. However, this is often not the case, and its inverse, hence, does not exist. Here, we describe the origin of the problem and how it can be solved.

In essentially all realistic applications, some frequency amplitudes are zero (or numerically so close to zero that they can be treated as being exactly zero). This comes in

part from the antenna response and the filters built into the signal chain of the detector, as well as additional digital processing where frequency filters are often applied to data to remove unwanted frequencies. This could, for instance, be high and low-pass filters to remove regions of low SNR, and notch filters to remove anthropogenic CW signals. Every complex-valued frequency bin filtered out removes 2 degrees of freedom in the time domain. Hence, every removed frequency subtracts two degrees of freedom from a trace in the time domain, and the degrees of freedom are $n_{\text{dof}} = n_t - 2n_{f,\text{removed}}$. Since the rank of the covariance matrix is equal to the degrees of freedom, the covariance matrix will not be of full rank if frequency filtering is applied, which describes a degenerate distribution. The covariance matrix is therefore not invertible, and the probability density function in Eq. (6) cannot be evaluated.

The correct way to describe a degenerate multivariate normal distribution is to replace the inverse covariance matrix with the Moore-Penrose pseudoinverse $\Sigma^{-1} \rightarrow \Sigma^+$, and the determinant with the pseudodeterminant $|\Sigma| \rightarrow |\Sigma|_+$ [70]. The Moore-Penrose pseudoinverse of a covariance matrix is computed by first performing the eigendecomposition $\Sigma = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}$, where $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_{n_t})$ is a diagonal matrix comprised of the eigenvalues of Σ . By inverting all nonzero entries of the diagonal matrix, $\mathbf{\Lambda}^+ = \text{diag}(\lambda_1^+, \dots, \lambda_{n_t}^+)$, where $\lambda_k^+ = 1/\lambda_k$ if $\lambda_k > 0$ and 0 if $\lambda_k = 0$, the Moore-Penrose pseudoinverse is obtained, $\Sigma^+ = \mathbf{Q}\mathbf{\Lambda}^+\mathbf{Q}^{-1}$. The pseudodeterminant is then the product of all nonzero eigenvalues, $|\Sigma|_+ = \prod_{\lambda_k > 0} \lambda_k$. In numerical implementation, the Moore-Penrose pseudoinverse and its determinant are typically computed using a threshold ϵ , where eigenvalues below this value are regarded as zero to ensure numerical stability. The degenerate multivariate normal distribution can then be expressed as

$$p(\mathbf{x}; \boldsymbol{\mu}(\boldsymbol{\theta}), \Sigma) = \frac{1}{\sqrt{(2\pi)^{n_{\text{dof}}} |\Sigma|_+}} \times \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \Sigma^+ (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))\right), \quad (14)$$

where $n_{\text{dof}} = n_t - 2n_{f,\text{removed}}$. The Moore-Penrose pseudoinverse and determinant exist and are unique for all matrices, and they coincide with the regular inverse and determinant if the covariance matrix is of full rank. Throughout this work, we use the notation of the regular and pseudoinverse interchangeably and apply the Moore-Penrose pseudoinverse where needed.

An additional benefit of using the pseudoinverse of the covariance is that it improves the numerical stability of evaluating the probability density function. If numerical errors are added to a point from the distribution, it may no longer lie exactly in the hyperplane where the distribution exists, and the likelihood can thus not be evaluated.




FIG. 2. *Top left:* The covariance matrix calculated empirically from 10,000 traces of generated noise using the realistic spectrum shown in Fig. 1. *Top right:* Zoom in on the first 30 ns by 30 ns of the covariance matrix. *Bottom left:* Covariance matrix calculated empirically from 10,000 traces of generated noise averaged over the diagonals according to a symmetric circulant structure (zoom in on the first 30 ns by 30 ns). *Bottom right:* The values of the diagonals of the averaged covariance matrix as a function of $\Delta t = t_j - t_i$.

However, when using the pseudoinverse, the width of the distribution is infinite in the filtered-out dimensions, and it is thus not sensitive to numerical errors.

F. Analytical expression of covariance matrix and its inverse

Under the noise generation process described after Eq. (4) and assuming $A(\omega_0) = 0$ and $A(\omega_{n_t/2}) = 0$, the covariance of the two measurements $V(t_i)$ and $V(t_j)$ can be shown to have the following form:

$$\text{Cov}(t_i, t_j) = \frac{2}{n_t^2 \delta t^2} \frac{1}{2} \sum_{k=1}^{n_t/2-1} A(\omega_k)^2 \cos(\omega_k \Delta t_{i,j}), \quad (15)$$

where $\Delta t_{i,j} = t_j - t_i$. We see here that the covariance matrix depends only on Δt , i.e., the distance from the diagonal and not t_i or t_j alone, and is symmetrical around both $\Delta t = 0$

(the diagonal) and $\Delta t = T/2$. It is thus a symmetrical circulant matrix as discussed in Sec. II D.

To find the inverse of the covariance matrix, Eq. (15) can be expressed as

$$\Sigma = \mathbf{F}^H \text{diag} \left(\frac{1}{2n_t \delta t^2} A'(\omega_k)^2 \right) \mathbf{F}, \quad (16)$$

where $\mathbf{F}$ is the discrete Fourier transform matrix with elements $F_{nk'} = \frac{1}{\sqrt{n_t}} e^{i\omega_{k'} \Delta t_{0,n}}$, $\mathbf{F}^H = \mathbf{F}^{-1}$ is its Hermitian transpose and inverse, and we let k' go from 0 to $n_t - 1$. The different scaling of the Fourier transform here compared to Eq. (1) ensures that the column vectors of $\mathbf{F}^H$ form a normalized basis. The matrix $\text{diag}(\frac{1}{2n_t \delta t^2} A'(\omega_k)^2)$ is a $n_t \times n_t$ diagonal matrix and $A'(\omega_k)$ are the amplitudes defined for $0 \leq k' \leq n_t - 1$ with $A'(\omega_{k'}) = A'(\omega_{n_t-k'}) = A(\omega_k)$. Since $\mathbf{F}^H$ is the matrix with normalized column vectors that diagonalizes Σ , the columns of $\mathbf{F}^H$ are the eigenvectors and




FIG. 3. One row of the covariance matrix (left) and its inverse (right) for simulated noise calculated empirically from 10,000 traces of noise using Eq. (9) and analytically from the spectrum using Eq. (15). The difference between the two calculations is less than 4% of the element on the diagonal ($\Delta t = 0$), which we attribute to the statistical uncertainties of calculating the covariance matrix from a limited number of noise realizations.

the entries of the diagonal matrix are the eigenvalues of Σ . We observe that the covariance matrix's rank (number of nonzero eigenvalues) is equal to 2 times the number of nonzero frequency amplitudes, $\text{rank}(\Sigma) = 2n_{A(\omega_k) > 0}$. Additionally, the inverse covariance matrix is

$$\Sigma^{-1} = \mathbf{F}^H \text{diag} \left(\frac{2n_t \delta t^2}{A'(\omega_k)^2} \right) \mathbf{F}, \quad (17)$$

or

$$\Sigma_{ij}^{-1} = \frac{2}{n_t} \sum_{k=1}^{n_t/2-1} \frac{2n_t \delta t^2}{A(\omega_k)^2} \cos(\omega_k \Delta t_{i,j}). \quad (18)$$

From this expression, we see that if any of the amplitudes are equal to zero, the inverse of the covariance matrix is not defined, which is consistent with the discussion in Sec. II E. When this is the case, the sum is calculated only for $A(\omega_k) > 0$, which gives the Moore-Penrose pseudoinverse of the covariance matrix. In numerical implementations, a nonzero but small threshold amplitude is set to avoid numerical uncertainties, and the sum is calculated for $A(\omega_k) > A_{\text{threshold}}$. Additionally, one can choose to ignore frequency components of signals by omitting them from the sum above.

Equations (15) and (17) provide us with an analytical way of calculating the covariance matrix of noise and its inverse directly from the spectrum. One row of the covariance matrix and one row of its inverse calculated analytically from the spectrum is shown in Fig. 3 alongside the empirical calculation from many noise realizations using Eq. (9). The two results are visually indistinct from each other and agree well. As the average noise spectrum is easily obtainable from data or can be calculated from measurements of the signal chain (S21 parameter), using the analytical expressions

presented in this section is a robust and convenient way to calculate the covariance matrix and its inverse, and will be used in the following.

The development discussed here naturally leads to a frequency domain formulation of the multivariate normal distribution, which is presented in Appendix A. This expression is numerically faster to evaluate and more stable. It will hence be employed in the following. Furthermore, the relation between the likelihood and *matched filtering* is discussed in Appendix B. The matched filtering approach will be used as part of a reconstruction algorithm (see Sec. III C 3) to improve stability and for efficient and optimal signal identification (see Sec. III D).

G. Verification of noise model

To verify that the multivariate normal distribution describes noise in radio detectors well, we evaluate its consistency with noise generated by the established and well-tested procedure described in Sec. II B. This is done by calculating the minus two log probability for many realizations of noise,

$$\begin{aligned} & -2 \ln p(\mathbf{x}; \boldsymbol{\mu}(\boldsymbol{\theta}), \Sigma) \\ & = (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta})) + \text{norm}. \end{aligned} \quad (19)$$

If the normalization is ignored, Eq. (19) is equivalent to a chi-square with correlations taken into account. It should hence be chi-square distributed with degrees of freedom equal to the number of degrees of freedom of the multivariate normal distribution, if the distribution describes the noise correctly. If all frequency amplitudes are nonzero and the covariance matrix is full rank, we have $n_{\text{dof}} = n_t$; otherwise, the degrees of freedom are equal to 2 times the number of nonzero frequency amplitudes or, equivalently, the rank of the (inverse) covariance matrix.




FIG. 4. $-2\Delta \ln p$ distributions for 10,000 realizations of simulated noise using the realistic spectrum in Fig. 1. The distributions have been calculated with the inverse covariance matrix [Eq. (18)] using the full spectrum except the first, the last, and the two filtered out frequency bins (left) and ignoring all frequency amplitudes below 1% of the maximal amplitude of the spectrum (right). Alongside the distributions are chi-square distributions with the expected degrees of freedom for the two cases shown.

In Fig. 4, the $-2\Delta \ln p$ (normalization subtracted) distribution is shown for 10,000 realizations of noise using the realistic spectrum shown in Fig. 1. The resulting distributions are shown for two different cases: first (left), using the full spectrum except for the first, the last, and the two filtered out frequency bins, which gives $n_{\text{dof}} = n_{\text{samples}} - 6 = 1018$ degrees of freedom. Second (right), ignoring all frequency amplitudes below 1% of the maximal amplitude which results in $n_{\text{dof}} = \text{rank}(\Sigma^+) = 762$. The inverse covariance matrices were calculated analytically from the spectrum using Eq. (18), and the calculations were repeated with the empirically determined covariance matrices, giving the same result within numerical uncertainties. The figure shows perfect agreement between the $-2\Delta \ln p$ distributions of the noise realization and the chi-square distribution with the corresponding degrees of freedom. This experiment again illustrates that Eq. (15) provides the covariance values for noise generated according to Sec. II B, and the resulting time-domain variables are Gaussian distributed as expected.

III. APPLICATIONS

The probabilistic noise model and the corresponding likelihood have vast applicability and could substantially improve the analysis of radio-based detection of cosmic rays and neutrinos. This is because all detectors record bandwidth-limited signals; i.e., the noise is correlated, and all detectors need to extract the signal properties from low signal-to-noise ratio data, as discussed in the Introduction. In the following, we quantify the impact of our new model on several typical reconstruction and event selection tasks.

A. Reconstruction of neutrinos for in-ice radio detectors

With a probabilistic description of noise in radio detectors, it is possible to perform a maximum-likelihood

reconstruction of neutrino signals in in-ice radio detectors. In this section, we demonstrate with a Monte Carlo (MC) simulation study of a radio detector station foreseen for the IceCube-Gen2 radio array [16], that using the likelihood description presented in this paper, signal parameters can be accurately reconstructed with correct uncertainties. Additionally, we study the improvements in reconstruction performance compared to previous work, which used an uncorrelated χ^2 [Eq. (5), e.g., used in [46,51,52]].

A preliminary study has been presented in Ref. [71]. In this paper, we improve the realism of the study by using the detector layout proposed for IceCube-Gen2 radio and taking into account the bending of radio signals in ice due to a depth-dependent refractive index.

1. Neutrino signal simulation

All simulations of neutrino signals described here are performed with NuRadioMC [50]. The Askaryan radio emission [72,73] resulting from the time-varying charge excess of a particle shower initiated by a neutrino can be calculated analytically in the frequency domain. We employ one such parametrization, which is referred to as *Alvarez2009* [74–76]. The parametrization provides an efficient way to calculate a realistic neutrino signal and is well-suited for a simulation-based reconstruction study. In this study, we consider only the so-called *hadronic shower*; i.e., in-ice cascades initiated by the breakup of the nucleus through Z or W exchange with the incoming neutrino. Due to the many initial hadronic interactions, these cascades have low shower-to-shower fluctuations, and the emitted signal can hence be considered deterministic, unlike electron-induced showers from electron-neutrino charged-current interactions, which have large stochasticities due to the LPM effect (see, e.g., [50,77]).

We approximate the index-of-refraction profile of the ice with an exponential function, which was used in previous studies and provides a good description of measured data [78]. The likelihood reconstruction will work with arbitrary ice models, but we choose this simple model as it provides a good trade-off between speed and accuracy for this proof-of-concept study.

The resulting radio neutrino signal depends on seven parameters: the energy of the particle shower (E_{shower}), the neutrino arrival direction (zenith angle θ_ν and azimuth angle ϕ_ν), the neutrino vertex position (expressed in spherical coordinates through zenith angle θ_{vertex} , azimuth angle ϕ_{vertex} , and distance r_{vertex}), and an overall time offset (t_0). The neutrino vertex position in spherical coordinates is defined in accordance with NuRadioMC, i.e., relative to an observer located at $x = y = z = 0$. These seven parameters are the ones we aim to estimate in a reconstruction, where neutrino energy and arrival direction are of special interest for astrophysical analyses.

The detector simulation is performed using NuRadioReco [46]. It provides a useful data structure for events and a wide range of antenna and signal chain descriptions. With NuRadioReco, the detector response to an electrical field,

TABLE I. Monte Carlo true parameter values of the three representative example events shown in this work.

Event	E_{shower}	θ_ν	ϕ_ν	r_{vertex}	θ_{vertex}	ϕ_{vertex}	t_0
1	15 PeV	62°	322°	220 m	149°	-70°	0 ns
2	640 PeV	64°	312°	2576 m	120°	-57°	0 ns
3	18 PeV	91°	191°	321 m	158°	-168°	0 ns

as calculated by NuRadioMC, can be simulated, and a trace for a neutrino signal in a given antenna and position is obtained.

This reconstruction study is performed using the station layout foreseen for the IceCube-Gen2 radio array, which is shown in Fig. 5. Specifically, we study events in the deep part of the detector, which consists of 16 antennas located on strings in three boreholes. At the bottom of the main string, called the *power string*, at 150m depth, four vertically polarized (VPol) antennas are located with a vertical distance of 1m. These antennas act as a phased array and trigger the station [79]. Directly above the phased array, two horizontally polarized (HPol) antennas are positioned, and four additional vertically polarized antennas are distributed along the power string, extending to a depth of 45 meters. The two additional helper strings sit ~ 34.6 m horizontal distance from the power string, forming a unilateral triangle. Each helper string has two VPol antennas and an HPol antenna at 144m depth with 1m spacing. Noise is added to the traces using the realistic spectrum shown in Fig. 1 and scaled to a noise temperature of $T_{\text{noise}} = 300$ K. Additionally, the neutrino signals are filtered using the same filter shape as the noise spectrum, representing the signal chain response. This choice ensures that the signal and noise are filtered consistently.

We performed an end-to-end simulation of neutrinos with various energies and incoming directions. In the following, we show the reconstruction results for three typical examples with varying event quality, but all fulfill the trigger criteria. The MC true parameter values of the events are shown in Table I. Figure 6 shows the simulated neutrino signal in all 16 antennas for event 3 with one realization of noise.

2. Likelihood reconstruction

We estimate the shower parameters by minimizing the minus two log-likelihood. For a signal in many antennas, the minus two log-likelihood for each antenna can be summed to give a likelihood for the entire event:

$$\begin{aligned}
 & -2 \ln \mathcal{L}(\boldsymbol{\mu}(\boldsymbol{\theta}); \mathbf{x}, \boldsymbol{\Sigma}) \\
 & = \sum_{\text{antennas}} (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta})) + \text{const.} \quad (20)
 \end{aligned}$$

In a reconstruction, the constants can be ignored since they do not depend on the neutrino signal parameters. In our reconstruction, $\boldsymbol{\mu}(\boldsymbol{\theta})$ is the neutrino signal parametrized




FIG. 5. Sketch of the planned IceCube-Gen2 hybrid station radio detector layout. Figure adapted from Ref. [16].




FIG. 6. Simulated neutrino signal in 16 antennas with one realization of noise for event 3 in Table I. The panels show the types and (x, y, z) coordinates of the antennas, and the signals fall outside the readout windows in the last three panels.

by the seven parameters, $\theta = (E_{\text{shower}}, \theta_\nu, \phi_\nu, r_{\text{vertex}}, \theta_{\text{vertex}}, \phi_{\text{vertex}}, t_0)$, described in Sec. III A 1 and $\mathbf{x}$ the simulated signal in the antennas including noise.

We use the `iminuit` [80] Python library, specifically the MIGRAD [81] algorithm, to minimize Eq. (20). The minimizer is initialized at the MC true parameters, assuming that the global optimum of the likelihood can be found. This guarantees fast and stable convergence and allows us to study the impact of the likelihood itself, independent of the rest of the reconstruction chain. Strategies exist for finding the global minimum when the true parameters are not known, e.g., by first only reconstructing the signal arrival direction, which was successfully done in previous work [52,82] or using neural networks to get an initial estimate of the parameters [58]. We leave this to later work as it does not affect the conclusions of this paper.

The likelihood not only allows us to reconstruct the neutrino parameters, but also provides a goodness-of-fit measure. If the signal model correctly describes the signal present in the data, the reconstructed minus two log-likelihood, $-2 \ln \mathcal{L}(\mu(\hat{\theta}); \mathbf{x}, \Sigma)$, should follow a chi-square distribution with degrees of freedom equal to the rank of the inverse covariance matrix (or 2 times the number of nonzero amplitudes in the noise spectrum) minus the number of reconstructed parameters. From this, a p-value for the signal can be calculated, which is crucial for verifying that a reconstructed signal is in agreement with the data. In this simulation study, we verified that all p-values are within the expected distribution.

3. Profile likelihood scans

For astrophysical analyses, the reconstructed parameters of interest are often the neutrino arrival direction or the energy. In this case, one can focus on the parameters of interest and treat the rest as nuisance parameters.

To obtain uncertainties on the reconstructed parameters of interest, a profile likelihood scan is performed. The procedure is as follows: First, the reconstruction is performed with all parameters free. This gives the best-fit parameters and a best-fit likelihood value. Second, a scan is performed on a grid in the parameters of interest, either arrival direction or energy. For every point in the grid, a fit is performed with the arrival direction or energy fixed to the grid value and the nuisance parameters free. This gives a likelihood for an alternative hypothesis. The minus two delta log-likelihood can then be used to determine confidence limits (uncertainty contours) for, e.g., 1σ , 2σ , and 3σ confidence levels (68.3%, 95.5%, and 99.7% containment, respectively), using Wilks' theorem. This is shown for the neutrino arrival direction in Fig. 7 (top row) for the three example events. For two degrees of freedom (θ_ν and ϕ_ν), the 68.3% uncertainty contour corresponds to a $-2\Delta \ln \mathcal{L}$ of 2.30.

To investigate whether the uncertainties estimated with this method are correct, we quantify their coverage, i.e., the expected number of times the MC true values of the parameters fall within the uncertainties. This is done by repeating the reconstruction of the same event with different realizations of noise 500 times. The best-fit likelihood values are then compared to the likelihoods of the MC true




FIG. 7. *Top row:* Profile likelihood scans for the three example events in Table I with one noise realization each. *Middle row:* Distribution of $-2\Delta \ln \mathcal{L}$ values between the reconstructed neutrino parameters and the true arrival direction profiled over the nuisance parameters for the three events with 500 different noise realizations each. A χ^2 -distribution with 2 degrees of freedom is shown for reference. The dotted histograms show the $\Delta \chi^2$ values between the reconstructed and true (profiled over the nuisance parameters) parameter values using the χ^2 as the objective function (see Sec. III A 4). *Bottom row:* Coverage for any confidence limit for the three events using the $-2 \ln \mathcal{L}$ or the χ^2 as the objective function in the reconstruction and to estimate the uncertainties.

neutrino arrival direction profiled over the nuisance parameters. According to Wilks' theorem, the $-2\Delta \ln \mathcal{L}$ should then follow a χ^2 distribution with 2 degrees of freedom, which we used to convert the likelihood to uncertainty contours. This is shown in Fig. 7 (Middle row) and shows good agreement between the expected and real distributions. By comparing the integrated χ^2 distribution to the integrated distribution of $-2\Delta \ln \mathcal{L}$ values, we obtain the coverage for any confidence level, which is shown in Fig. 7 (bottom row). The figure shows that the coverage is correct for all confidence levels. Hence, using the

likelihood proposed in this paper, we obtain uncertainties on the reconstructed parameters with correct coverage on an event-by-event basis.

Additionally, the likelihood values of the best-fit parameters were compared to the true parameter likelihood values, without profiling over the nuisance parameters. This again gives $-2\Delta \ln \mathcal{L}$ distributions that follow the expected χ^2 distributions (not shown here) with 7 degrees of freedom, and the coverages are correct for all confidence levels. Hence, if uncertainty contours were drawn in the seven-dimensional parameter space, they would be correct.

TABLE II. Standard deviations, i.e., uncertainties, of best-fit signal parameters for the three neutrino events in Table I with 500 different realizations of noise each using the $-2 \ln \mathcal{L}$ or the χ^2 as the objective function in the reconstruction. The uncertainties on the shower energy is given as a percentage of the true shower energy.

Event	Method	σ_{E_ν}	σ_{θ_ν}	σ_{ϕ_ν}	$\sigma_{r_{\text{vertex}}}$	$\sigma_{\theta_{\text{vertex}}}$	$\sigma_{\phi_{\text{vertex}}}$	σ_{t_0}
1	$\mathcal{L}$	5.7%	0.8°	1.0°	4.0 m	0.769°	0.120°	0.04 ns
	χ^2	6.8%	1.1°	1.4°	4.5 m	0.861°	0.138°	0.05 ns
2	$\mathcal{L}$	11.8%	1.6°	8.8°	79.6 m	0.066°	0.024°	0.04 ns
	χ^2	18.3%	1.7°	9.0°	131.0 m	0.107°	0.026°	0.05 ns
3	$\mathcal{L}$	11.6%	0.6°	3.7°	20.5 m	0.988°	0.129°	0.11 ns
	χ^2	13.1%	0.8°	4.7°	23.8 m	1.141°	0.150°	0.12 ns

4. Importance of considering bin-to-bin correlations

To verify the importance of using the correct probabilistic description of noise in radio detectors, we perform the same reconstruction ignoring noise correlations; i.e., we use a chi-square as expressed in Eq. (5) as the objective function in the reconstruction. This has been used as the objective function in all previous work employing the forward folding technique, e.g., [36,46,51,52]. The resulting $\Delta\chi^2$ distribution (true arrival direction profiled over nuisance parameters minus best fit) for 500 trials is shown in the middle row of Fig. 7. The distribution does not follow a χ^2 -distribution with 2 degrees of freedom since the objective function does not take into account the correlations in the noise. It can hence not be used to estimate the uncertainty on the reconstructed parameters. In the bottom row of Fig. 7, it is shown that if uncertainties are estimated assuming the $\Delta\chi^2$ is χ^2 -distributed, it would result in undercoverage for the three events.

In addition to not being able to estimate uncertainties, the uncertainties themselves are also larger when minimizing a χ^2 compared to a $-2 \ln \mathcal{L}$. The standard deviations of the seven reconstructed parameters for the 500 reconstructions of the same neutrino signals with different realizations of noise using both methods are shown in Table II. It shows that the spread of best-fit parameters is consistently lower when minimizing the $-2 \ln \mathcal{L}$ compared to the χ^2 . The reason for this is that by taking the information about the correlations in the noise into account, the fitted parameters are better constrained close to the true values. A similar study for the biases of the two methods is shown in Appendix C, which shows a similar improvement for most parameters using the likelihood. Hence, the likelihood method proposed in this paper substantially improves the reconstruction uncertainties, with the magnitude of the improvement depending on the frequency spectra of the noise and signals.

B. Fast uncertainty estimation

In the following, we discuss how the noise model and likelihood description enable us to estimate the reconstruction resolution of radio experiments in a computationally

efficient way, which can be used for the optimization of future detectors.

Estimating the reconstruction capabilities of an experiment can be done in several ways. The most robust, though computationally expensive, method is to generate a large Monte Carlo dataset and reconstruct the events to determine the uncertainties on the reconstructed parameters. This method has successfully been used in the past, for example, to estimate the reconstruction resolution of in-ice neutrino detectors [52,58,82]. It is, however, often extremely computationally expensive, especially when averaging over a large sample of events, and hence too inconvenient for, e.g., detector optimization, where many different detector layouts are investigated.

The likelihood description presented in this work enables two additional methods. The first of which is a so-called *Asimov* likelihood scan [83]. The Asimov dataset is the “most representative” dataset, i.e., the average of all traces, for a given set of true parameter values. It is, hence, simply equal to the signal prediction in an antenna without any noise added. We can then perform a likelihood scan, or a profile likelihood scan as described in Sec. III A 3, treating the Monte Carlo truth signal as the observed trace. The resulting contour is the median reconstruction uncertainty for the given set of true parameter values and should match the distribution of reconstructed values if we were to do many reconstructions of the same event with different realizations of noise. The Asimov profile likelihood contour for event 3 in Table I is shown in Fig. 8 (right) along with 500 reconstructions with different realizations of noise. The figure shows that the Asimov contour captures the shape of the distribution well.

The most computationally efficient way to estimate reconstruction uncertainty, enabled by the likelihood description, is to calculate the Fisher information. The Fisher information matrix tells us how much information a sample from a distribution, e.g., a deterministic signal plus random noise, carries about the parameters that model the distribution, e.g., the signal parameters. For data described by a multivariate normal distribution, where the covariance matrix does not depend on the signal parameters, i.e., signals in radio detectors, the Fisher information matrix is




FIG. 8. *Left*: Inverse Fisher information matrix for event 3 in Table I. The units of the entries are the product of the units of the two components. Here the units have been scaled for visibility and are 1 PeV, 1° , 10° , 1° , 0.1° , 10m, and 0.1 ns for the seven parameters, respectively. *Right*: Asimov profile likelihood scan for the same event with 68% containment contour and 500 reconstructions with different realizations of noise. The marginalized Fisher information estimate of the uncertainty (68% containment) on the neutrino arrival direction is shown centered around the MC true parameter values.

given by

$$\mathcal{I}_{nm} = \sum_{\text{antennas}} \frac{\partial \boldsymbol{\mu}^T}{\partial \theta_n} \boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\mu}}{\partial \theta_m}, \quad (21)$$

which depends only on the derivatives of the signal with respect to the parameters and the covariance matrix of the noise, and it can easily be summed over the antennas measuring the signal. The Crámer-Rao bound [84,85] states that for an unbiased estimator, the inverse Fisher information matrix is a lower bound on the covariance matrix of the parameters,

$$\text{Cov}(\boldsymbol{\theta}) \geq \mathcal{I}^{-1}. \quad (22)$$

This bound is saturated for large signal amplitudes. A maximum likelihood estimator is, however, not guaranteed to be unbiased, which can be seen from the asymmetric distribution of reconstructed parameters in Fig. 8 (right) and quantized in Appendix C. In this case, the Crámer-Rao bound can be slightly broken, and the inequality turns into an approximate estimate of the covariance matrix. We note that this is of little concern for optimization studies because by optimizing the estimated uncertainties, the real uncertainties should also improve, which can be verified with a full sensitivity study of an initial and optimized detector design.

The inverse Fisher information matrix for event 3 in Table I has been calculated using numerical differentiation and is shown in Fig. 8 (left). By selecting the 2×2 sub-matrix associated with the neutrino arrival direction, we obtain the marginalized covariance matrix which parametrizes an ellipse corresponding to the uncertainties of these parameters. The Fisher estimate of the uncertainties is

shown in Fig. 8 (right). We see that the Fisher estimate is an elliptical (second order) approximation that does not capture the full asymmetries of the distribution, but its area sufficiently matches the Asimov contour. Since the Fisher estimate can be calculated many orders of magnitude faster than the full reconstruction and profile likelihood scan, it is useful in analyses where the exact shapes of the uncertainties are superfluous and only their approximate magnitudes are of interest.

The Fisher information matrix is an important tool in the prospects of end-to-end detector optimization of radio detectors [86]. With the use of differentiable programming, e.g., PyTorch or JAX, and deep learning-based surrogate models, a fully differentiable radio neutrino signal pipeline can be implemented. The signal will then be automatically differentiable with respect to both the signal parameters and the antenna positions and orientations. Additionally, the Fisher uncertainty calculation can be implemented as a differential proxy for the reconstruction uncertainty, which means that the gradients of the reconstruction uncertainty with respect to the detector layout can be calculated using backpropagation. It hence allows us to efficiently optimize the layout of radio detectors by minimizing the reconstruction uncertainty with respect to the detector parameters. This concept holds great potential for improving the scientific outcome of future radio detectors and will be the subject of future studies.

C. Electric-field reconstruction for cosmic rays

The likelihood can improve another key reconstruction task. Air-shower detectors are typically built with dual-polarized antennas, i.e., two antennas in one mechanical structure that measure two orthogonal polarization

components of the electric field simultaneously [1]. This measurement, together with the knowledge of the arrival direction, allows for the reconstruction of the incident electric field. The electric-field reconstruction constitutes the low-level component of a full cosmic-ray air shower reconstruction, where, e.g., the energy fluences calculated from the electric field at each antenna are combined to estimate the air-shower energy (e.g. [3]) and $X_{\max}$ (e.g. [9]) of the shower. Hence, an improvement in this first step will result in an improvement in the reconstruction of the key observables, the shower energy and $X_{\max}$.

In this section, we compare the likelihood method to the currently most widely used method, which obtains the electric field via unfolding (see Sec. III C 2 for details). This method does not utilize the full time-domain information of the waveform, which limits reconstruction resolution, and it is known to be biased at low SNR. We also compare our method to the *forward folding* method proposed in Ref. [46], which uses an analytic electric field model and minimizes a χ^2 objective function without considering correlations in the noise. Our method adopts the same analytic electric field model, but minimizes the $-2\ln\mathcal{L}$ instead, hence taking noise correlations into account. In the following, we show that this improves the polarization and energy fluence reconstruction uncertainties over both previous methods in a simulation study.

The results of this section have been previously shown at a conference [87].

1. Simulation dataset

To evaluate the performance of the reconstruction methods discussed here, we generate a simulation dataset of cosmic-ray air showers using the CoREAS extension [88] of the CORSIKA framework. The dataset consists of a large set of proton-initiated showers with energies uniformly distributed in $\log_{10}(E)$ between 10^{15} and 10^{19} eV, and arrival directions uniformly distributed over a zenith angle range of 60° to 85° . We simulate the showers for the magnetic field and atmospheric conditions at the South Pole. For each air shower, these simulations produce the time-dependent electric field at 160 observer positions placed on the ground in the standard star-shaped grid around the shower axis. We use these electric fields directly in the subsequent calculations and do not apply any additional randomization of positions and interpolation. As we present all our findings as a function of signal-to-noise ratio, this is the most accurate approach.

Each simulated electric field is observed by a dual-polarized antenna. The antenna is a logarithmic periodic dipole antenna (LPDA) that measures with a sampling rate of 500 MHz in the frequency band from 30 MHz to 80 MHz with two readout channels, one sensitive to north-south polarized electric fields and one sensitive to east-west polarizations. We employ a simple analytical




FIG. 9. Example air-shower signal from CoREAS simulation as measured by a dual-polarized antenna. The black curve shows the noiseless signal, whereas the gray curve shows the signal with measurement noise.

approximation of the LPDA response available in NuRadioReco [46], which we use to simulate the response of the antenna to the electric field.

Finally, we add noise to the traces and apply several filters to simulate hardware- and analysis-level filters, which remove environmental noise at low and high frequencies. The generated noise follows a smooth spectrum in the 30 MHz to 80 MHz range, implemented using a third-order Butterworth high-pass filter and an eighth-order Butterworth low-pass filter. This frequency band is chosen as it represents a typical range for air-shower radio arrays [89]. Additionally, as very low SNR observers are overrepresented, we randomly skip some of them to make the distribution of events over SNR more uniform. The SNR is calculated using the maximum absolute amplitude of the noiseless traces divided by the standard deviation of the noise. The resulting dataset contains 73,459 simulated cosmic-ray electric-field waveforms with SNR values below 20, which we will reconstruct in the following sections. We put our emphasis on low SNRs, as this is the region that is difficult to reconstruct. An example of the resulting traces as measured by the dual-polarized antenna is shown in Fig. 9.

For the purpose of this study, we assume the cosmic-ray arrival direction is known from external reconstruction (e.g., from triangulation using the signal times in several observers), and the signal arrival time is known to within 30 ns, which we call the *search window*. This is a reasonable assumption, as the arrival direction can often be well constrained by fitting a wavefront to the arrival times of the signal in antennas with high SNR. The wavefront can then be extrapolated to estimate the arrival time in antennas with lower SNR.

2. Previous method: Unfolding and noise subtraction

The *unfolding* and *noise subtraction* method is the most commonly used method to estimate the fluence and polarization of an electric field in a dual-polarized antenna [44]. The method comprises three simple steps. First, the measured trace is unfolded with the antenna response for the known arrival direction of the shower. This obtains an estimate of the three-dimensional (polarizations), time-dependent electric field, $\mathbf{E}$ of shape $n_{\text{pol}} \times n_t$, that includes unfolded noise from the trace. It is worth noting that the unfolded noise is not a correct physical representation of the radio noise that the antenna sees, since all noise is unfolded as if it is coming from the direction of the cosmic ray and includes noise from the electronics. This can lead to biases in reconstructed quantities.

To estimate the time of the peak of the electric field, a Hilbert envelope is applied in a *search window*, the norm of the three-dimensional envelope is calculated, and the peak is identified as t_{peak} . In this work, the search window is defined as $t_{\text{true}} \pm 30$ ns where t_{true} is the peak of the norm of the Hilbert envelope of the noiseless electric field. A *signal window* is then defined as $[t_1, t_2]$ at ± 30 ns around t_{peak} , and a *noise window* is defined as $[t_3, t_4]$, which is a 200 ns period outside the signal window. The electric field energy fluence of the three components is then estimated as

$$\hat{f}_{\text{pol}} = \epsilon_0 c \left(\delta t \sum_{t_1 \leq t_n \leq t_2} |\mathbf{E}_{\text{pol},n}|^2 - \delta t \frac{t_2 - t_1}{t_4 - t_3} \sum_{t_3 \leq t_n \leq t_4} |\mathbf{E}_{\text{pol},n}|^2 \right), \quad (23)$$

where ϵ_0 is the dielectric constant and c is the speed of light. The method hence estimates the total fluence in the signal window and subtracts the contribution from the noise estimated from the noise window. We set f_{pol} to zero in cases where the estimated noise fluence is larger than the fluence in the signal window, to avoid nonphysical negative fluence estimates.

The basis vectors of the coordinate system chosen for three polarizations of the electric field are along the shower axis, $\hat{r}$ (with $\mathbf{E}_r$ assumed to be 0), perpendicular to the shower axis in the vertical plane, $\hat{\theta}$, and perpendicular to the shower axis in the horizontal plane, $\hat{\phi}$. The estimate of the total fluence is then $\hat{f}_{\text{tot}} = \hat{f}_{\theta} + \hat{f}_{\phi}$ and the polarization of the electric field is

$$\hat{P} = \arctan \left(\sqrt{\hat{f}_{\phi}} / \sqrt{\hat{f}_{\theta}} \right). \quad (24)$$

We note that using this convention, all measured polarizations will be between 0° and 90° .

The uncertainties of the reconstructed fluences are estimated by assuming the noise in the noise window is Gaussian and uncorrelated, and the uncertainty on the polarization is derived using error propagation.

3. Forward folding reconstructions: Chi-square and likelihood

The benefits of forward folding methods are twofold. First, it avoids unfolding system and ambient noise as if it was coming from one specific direction, and second, it constrains the signal to a parametrized model representing the physical process being observed. The parametrized model can either be a full simulation (as in Sec. III A) or a simpler approximation. Since cosmic-ray air-shower simulations are computationally expensive, we employ a simple frequency domain parametrization of the electric field, which has been shown to describe radio pulses from cosmic-ray air showers well [46],

$$\begin{pmatrix} \mathcal{E}_{\theta,k} \\ \mathcal{E}_{\phi,k} \end{pmatrix} = \begin{pmatrix} A_{\theta} \\ A_{\phi} \end{pmatrix} 10^{\alpha \cdot f_k + \beta (f_k - f_{\text{offset}})^2} e^{-i(2\pi f_k \cdot t_{\text{offset}} - \psi)}, \quad (25)$$

where $\mathcal{E}$ is the electric field in the frequency domain, f_k are the frequencies, and A_{pol} are the electric field amplitudes. The frequency dependence of $\mathcal{E}$ is controlled by α , which sets the slope of the electric field in the frequency domain and is often constrained to negative values, and β , which controls a second-order correction where f_{offset} is fixed to 30 MHz. The last exponential controls the complex phases, where t_{offset} is the time of the pulse relative to the start of the trace, and ψ is a phase offset. Additionally, we numerically renormalize the electric field, such that the parameters we reconstruct are directly the fluences of the two polarizations, f_{θ} and f_{ϕ} . In our reconstruction, we allow f_{θ} and f_{ϕ} to take on negative values, corresponding to a sign flip of A_{θ} and A_{ϕ} , respectively. This makes the method sensitive to polarizations across the 0° to 180° range (polarization flips of 180° are degenerate with changing ψ by π). However, for a fair comparison, we restrict it here to the same range as the noise subtraction method in a consistent way.

For a given set of parameter values, the electric field is multiplied by the vector effective length (VEL) of the antenna for the known arrival direction and filtered accordingly. By applying the inverse Fourier transform, the estimate of the noiseless signal, μ , is obtained, which can then be compared to the measured trace in an objective function. The resulting signal depends on six parameters to be reconstructed, $\theta = (f_{\theta}, f_{\phi}, \alpha, \beta, t_{\text{offset}}, \psi)$.

Traditionally, the objective function to be minimized is a χ^2 , as expressed in Eq. (5), which does not consider the correlation in the noise. We propose minimizing the $-2 \ln \mathcal{L}$ expressed in Eq. (20), which correctly describes the band-limited nature of the noise. In this work, the minimization algorithms for both objective functions were improved to achieve better stability. For the $-2 \ln \mathcal{L}$ minimization, the algorithm consists of two minimization processes. In the first, a matched filter (see Appendix B) is applied in every step of the minimization to analytically profile over amplitudes and efficiently profile over a grid of t_{offset} in the search window. This drastically reduces the number of local optima

of the likelihood, and allows us to fit the parameters that control the shape of the signal, α , β , ψ , and the ratio of the fluences $r_{\theta/\phi} = f_{\theta}/f_{\phi}$. In the second minimization, the $-2\ln\mathcal{L}$ is minimized with all parameters free, including the time and overall amplitude, to fine-tune the result. Since the $-2\ln\mathcal{L}$ has local minima for opposite polarizations, we run the algorithm four times, initialized in each quadrant of (f_{ϕ}, f_{θ}) to guarantee convergence.

The χ^2 minimization algorithm was improved in a corresponding way by using a normalized cross-correlation [Eq. (26)] to profile over time and amplitude in the first minimization process. The normalized cross-correlation is equivalent to a matched filter with a diagonal covariance matrix, i.e., no correlations taken into account. This allows us to evaluate the impact of the objective function directly, i.e., $-2\ln\mathcal{L}$ or χ^2 , with the same improvements to the minimization algorithm.

For both objective functions, the uncertainties on the reconstructed parameters are extracted from the Hessian

matrix at the minimum of the $-2\ln\mathcal{L}$ or χ^2 . The uncertainty on the polarization can then be estimated using error propagation.

4. Reconstruction performances

All events in the simulation dataset were reconstructed using the three methods discussed in the previous sections. The relative errors (reconstructed minus true divided by true) of the reconstructed fluences and the errors (reconstructed minus true) of the reconstructed polarizations are shown in Fig. 10. The results are binned in SNR, and the medians and central 68%-quantiles, $\pm\sigma_{68\%}$, of each bin are shown in the figure. The noise subtraction method has two competing biases in the fluence at low SNR. At $\text{SNR} < 3$, there is a positive bias originating from the peak-finding, which will always find the noisiest part of the unfolded electric field in the search window when almost no signal is present. The negative bias at $\text{SNR} > 3$ is caused by the fluence estimator not taking into account that the signal and




FIG. 10. Reconstruction results for the three methods. The relative errors (reconstructed minus true divided by true) of the total fluences and errors (reconstructed minus true) of the polarizations are shown for individual events. Many of the individual reconstructions (transparent gray dots) fall outside the plotted range. The events were binned in SNR to calculate the median, i.e., the bias, and asymmetric central 68%-quantiles, i.e., the uncertainties ($\pm\sigma_{68\%}$), of each bin.

noise can be destructively aligned. Additionally, since f_θ and f_ϕ are not allowed to be negative, the reconstructed fluences pile up at 0, and the polarizations pile up at 0° and 90° , which causes very large errors at low SNR.

The χ^2 minimization method reduces the bias at low SNR for both the fluence and polarization reconstruction. The $-2 \ln \mathcal{L}$ minimization further improves the bias and remains unbiased in the fluence down to $\text{SNR} > 2$, and is effectively unbiased in the polarization. The remaining positive bias in the fluence at very low SNR is caused by the minimization fitting the pulse to the largest fluctuation in the noise when no signal is present.

In order to compare the uncertainties of the three methods, we plot half the width of the central 68%-quantiles of the errors, which we here call the uncertainties, of each bin in Fig. 11 (left). Using the likelihood method, the fluence relative uncertainties are decreased by a factor ranging from 1.6 to 2.3 in the $5 < \text{SNR} < 15$ range compared to the noise subtraction method, and the polarization uncertainties are decreased by a factor ranging from 1.7 to 13.1 in the same range. Compared to the χ^2 method, the likelihood method decreases fluence relative uncertainties by a factor ranging between 1.2 to 1.3 and the polarization uncertainties by a factor ranging between

1.2 to 1.5 in the same SNR range. The latter improvement arises purely from changing the objective function, which underlines the importance of taking bin-to-bin correlations into account. The improvement in reconstruction resolution is hence substantial, especially at low SNR.

An estimate of the uncertainty on the reconstructed parameters can be obtained event-by-event from the three methods. To verify if the estimated 68% uncertainties are correct, we calculate the coverage and show it in Fig. 11 (right). The coverage is calculated by counting how often the true value is within the estimate plus/minus the estimated 68% uncertainty, which should ideally be 68% of the time. Both the noise subtraction and χ^2 methods severely under-predict the uncertainties, whereas the likelihood method predicts the uncertainties well across most SNRs. The remaining slight under-coverage is likely caused by the simplicity of the electric field model, which can not perfectly model the cosmic-ray signals at high SNR.

The unfolding method can be improved by estimating the energy fluence from the noisy electric fields based on Rice distributions as recently proposed in Ref. [90]. This method takes into account the band-limited nature of the noise but does not constrain the electric field to a physical model, and keeps the problem of the unfolding of amplifying the noise




FIG. 11. *Left:* Half the width of the asymmetric central 68%-quantiles, i.e., uncertainties, of the fluence relative uncertainties (reconstructed minus true divided by true) and polarization uncertainties (reconstructed minus true) of the three methods. *Right:* Coverage of the event-by-event estimated 68% uncertainties of the fluence and polarization for the three methods, i.e., how often the true value is within the estimate plus/minus the estimated uncertainty, which should ideally be 68%.

at small SNR. When applied to our simulation dataset, we observed an equivalent bias and similar uncertainties to the noise subtraction method at low SNR. However, the method correctly estimated the uncertainties across all SNRs, and it may hence be beneficial to use it at high SNRs.

The likelihood method for reconstructing electric fields from several antennas with close proximity to each other is implemented in NuRadioReco. Our method has already been applied to RNO-G cosmic-ray data with high signal-to-noise ratios and gives consistent results to the simpler χ^2 method [91], thus showing the applicability of our method to experimental data.

D. Signal identification from background

The final application of the likelihood we present is identification of signals from background. The ratio of the likelihoods for two different hypotheses is the most powerful way of comparing two hypotheses [92]. Thus, using the likelihood is the most effective method for separating a signal from a background when searching for known signals in radio antenna traces. Separating the signal from the background is a common challenge in all radio detection experiments, as the signal purity is extremely low. The first primary analysis step is to identify the cosmic-ray or neutrino-induced radio signals from background events that are several orders of magnitude more numerous. In the following, we demonstrate the method by conducting a toy study of searching for cosmic-ray signals in in-ice radio neutrino detectors, as one of many possible applications.

Cosmic-ray signals in radio neutrino detectors make up a substantial background. Since cosmic-ray air showers are several orders of magnitude more abundant than UHE neutrinos, and radio emission from air showers is a

well-understood phenomenon, they can effectively be used as a calibration source and end-to-end detector and analysis verification and demonstration. Searches for cosmic rays have been carried out with ARIANNA [36] using shallowly deployed LPDA antennas, demonstrating the reconstruction capabilities of the experiment. A similar analysis was recently published with RNO-G [91,93,94].

When identifying cosmic ray signals, the main background is thermal noise that triggers the detector. The measured event rate of cosmic rays by ARIANNA is on the order of 1 per station per day; however, the stations are often configured to trigger and save traces at a 1 Hz rate. The vast majority of triggered events are thus thermal noise fluctuations, and a background rejection rate of $\mathcal{O}(\sim 10^5)$ per signal event is needed. To separate signal from background, previous analyses have relied on correlation scores of known signal templates with the measured trace in one antenna,

$$\rho = \max(\rho(\Delta n)) = \max\left(\frac{\sum_i^{n_t} \mu_i \cdot x_{i+\Delta n}}{\sqrt{\sum_i^{n_t} \mu_i^2} \cdot \sqrt{\sum_{j=\Delta n}^{n_t+\Delta n} x_j^2}}\right), \quad (26)$$

where x is the measured trace, μ is the template, and Δn represents a time offset. This correlation acts as a test statistic for which a cut is set to determine which triggered events are likely cosmic-ray signals. In the following, we demonstrate, using a toy-model study, that utilizing the maximized two delta log-likelihood between no signal (zeros) and the signal template achieves a better separation between the signal and background.

The toy-model study is carried out as follows: first, a set of 1,000,000 background traces for one antenna is simulated, containing thermal noise that fulfills the trigger criteria. We use a sampling rate of 0.8 GHz and a trace length of $n_t = 512$ samples, and noise with a temperature of 300 K that follows a smooth spectrum using Butterworth filters at 80 MHz (second-order high-pass) and 220 MHz (tenth-order low-pass). The narrower frequency band is chosen to optimize for cosmic-ray signals, which have more power at lower frequencies than neutrino signals, and targets the frequency range where the LPDAs are most sensitive. We set the trigger criteria such that the station has a 1 Hz trigger rate and corresponds to a high-low trigger threshold of 3.3σ . Second, we generate a set of 100,000 signal traces, by generating unbiased thermal noise (that typically does not satisfy the trigger criteria) and injecting a cosmic-ray signal with an SNR of 3.3, i.e., at the trigger threshold. An example of a trace with the injected cosmic-ray signal is shown in Fig. 12 (top), and Fig. 12 (bottom) shows a trace of triggered thermal noise with high signalness.

To quantify the performance of identifying signals using the correlation score, we calculate the correlation score for the background and signal traces using the injected signal




FIG. 12. *Top*: Trace with injected cosmic ray signal template and one realization of thermal noise. *Bottom*: Simulated thermal noise that would trigger the detector (1 Hz trigger), which is the main background for identifying cosmic-ray signals in single antennas.




FIG. 13. *Left and Middle:* Correlation score and two delta log-likelihood distributions for 1,000,000 background traces containing triggered thermal noise and 100,000 signal traces containing thermal noise and an injected cosmic ray signal. *Right:* Signal efficiency versus background passing fraction (ROC curves) of the toy-model cosmic ray identification analysis using the correlation score and the likelihood as the test statistic. The gray dotted line indicates the target background passing fraction of 10^{-5} , which gives a 55% signal efficiency for the correlation score and 77% for the likelihood.

as the template. Similarly, we calculate the two delta log-likelihood between no signal and the template maximized over time and amplitude. The maximization is performed using the matched filter search approach described in Appendix B, due to its efficiency for a pre-calculated template. The resulting test statistic distributions are shown in Fig. 13 (left and middle). The histograms show a clear separation between the two classes using both methods, with a smaller overlap for the likelihood method.

To fully assess the gain of the likelihood approach, we calculate the signal efficiency (true positive rate) and a background passing fraction (false positive rate) for any cut on the test statistic. The resulting ROC curves are shown in Fig. 13 (right). Especially at low background passing fractions, the likelihood achieves a better signal efficiency. At the target background passing fraction of 10^{-5} , the remaining fraction of cosmic ray events after the cut is 55% using the correlation score, whereas it is 77% using the likelihood. For smaller passing fractions, as required for the less abundant neutrinos, the benefit of using the likelihood further increases.

This simple toy-model study demonstrates that the identification of signals in radio antennas can be significantly improved by exploiting the information about the correlations in the noise. The method can be applied in any scenario where known signals are searched for in one or more radio antennas with band-limited noise.

IV. CONCLUSION

In this paper, we have proposed a probabilistic description of noise in radio detectors that captures its band-limited/correlated nature. This framework enables the formulation of a likelihood for deterministic signals in such detectors, providing an ideal objective function for reconstruction in

cases where a full signal model, or a reasonable proxy, can be formulated. The central code for calculating and optimizing this likelihood with respect to a signal is available in NuRadioMC. Through simulation studies, we have demonstrated that the method accurately reconstructs neutrino signals in in-ice radio detectors as well as electric fields from cosmic-ray air showers using dual-polarized antennas. In both cases, the method provides reliable uncertainty estimates on the reconstructed parameters and substantially improves the reconstruction precision compared to previous approaches. We further highlighted the potential of the likelihood to improve signal efficiency in event selections. Finally, the likelihood formulation enables a fast estimation of reconstruction uncertainties using the Fisher information, enabling differentiable end-to-end detector optimization. The approach presented in this paper has the potential to advance analyses across the field of radio detection of particle physics signals and may play a central role in optimizing future detectors, future discoveries, or precision measurements.

ACKNOWLEDGMENTS

We thank Nils Heyer for generating the neutrino simulation dataset used in Sec. III A and for useful discussions. We thank the RNO-G Collaboration for useful feedback and discussions. The work is supported by the European Union (ERC, NuRadioOpt, Grant No. 101116890) and by the Swedish Research Council (VR) via the Project No. 2021-05449.

DATA AVAILABILITY

NuRadioMC [50] and NuRadioReco [46] were used to analyze the data in this article. The code implementing the methods described in this article are publicly available in

NuRadioReco. Analysis scripts and data files are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: FREQUENCY DOMAIN FORMULATION OF MULTIVARIATE NORMAL DISTRIBUTION

It is often convenient to perform calculations in the frequency domain instead of the time domain. We here present a formulation of the noise model [multivariate normal distribution, Eq. (6)] using Fourier transforms.

We first define the noise power spectral density as $S_n(\omega_k) \equiv 2A(\omega_k)^2 \delta f$. From Eq. (17), the following relation can then be derived:

$$\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \mathbf{b} = 4\Re \left\{ \sum_{k=1}^{n_t/2-1} \frac{\tilde{b}(\omega_k) \tilde{a}^*(\omega_k)}{S_n(\omega_k)} \delta f \right\}, \quad (\text{A1})$$

where $\tilde{a}(\omega_k)$ and $\tilde{b}(\omega_k)$ are the discrete Fourier transforms [as defined in Eq. (1)] of real-valued vectors $\mathbf{a}$ and $\mathbf{b}$. This relation states that the multiplication of any two real-valued vectors with the inverse of the covariance matrix can be expressed using the Fourier transforms of the vectors and the noise power spectral density. The matrix multiplication in Eq. (6) can hence be expressed as

$$\begin{aligned} (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta})) \\ = 4 \sum_{k=1}^{n_t/2-1} \frac{|\tilde{x}(\omega_k) - \tilde{\mu}(\omega_k, \boldsymbol{\theta})|^2}{S_n(\omega_k)} \delta f. \end{aligned} \quad (\text{A2})$$

By comparing the definition of $S_n(\omega_k)$ to the eigenvalues of the covariance matrix [e.g., Eq. (16)], we see that they can be expressed as $\lambda_k = \frac{1}{4\delta t} S_n(\omega_k)$ with two degenerate eigenvalues per ω_k . The determinant is then

$$|\boldsymbol{\Sigma}| = \left(\prod_{k=1}^{n_t/2-1} \frac{1}{4\delta t} S_n(\omega_k) \right)^2. \quad (\text{A3})$$

The multivariate normal distribution can hence be expressed as

$$\begin{aligned} p(\mathbf{x}; \boldsymbol{\mu}(\boldsymbol{\theta}), S_n) &= \frac{1}{\sqrt{(2\pi)^{n_t}} \left(\prod_{k=1}^{n_t/2-1} \frac{1}{4\delta t} S_n(\omega_k) \right)} \\ &\times \exp \left(-\frac{1}{2} 4 \sum_{k=1}^{n_t/2-1} \frac{|\tilde{x}(\omega_k) - \tilde{\mu}(\omega_k, \boldsymbol{\theta})|^2}{S_n(\omega_k)} \delta f \right), \end{aligned} \quad (\text{A4})$$

from which the likelihood can be derived as in Eq. (7). This expression is convenient since it circumvents the calculation

of the covariance matrix and its inverse. It is also numerically more robust compared to the inversion of the covariance matrix, which has a combination of very large and small numbers. The combination of Fourier transforms and sums is also computationally faster by about a factor of 4 compared to the time-domain calculation.

Similar to the discussion in Sec. II E, if any frequency amplitudes are equal to zero, Eq. (A4) is undefined. In this case, the multivariate normal probability can be calculated by letting the sum and product run over all nonzero values of the noise power spectral density, i.e., $S_n(\omega_k) > 0$. This is equivalent to using the pseudoinverse of the covariance matrix as in Eq. (14).

APPENDIX B: RELATION TO MATCHED FILTERING

A common tool used to search for signals in noisy data is called a *matched filter*. It is the filter that maximizes the signal-to-noise ratio for a known signal template in a data stream with an arbitrary known noise spectrum. This method is for instance used in gravitational wave detection, where a large signal template bank is correlated using the *matched filter* with the continuous data stream of the gravitational wave detector, e.g., LIGO, VIRGO, and KAGRA [95–99]. In many ways, gravitational wave detection is similar to radio detection of cosmic rays or neutrinos, i.e., searching for rare known pulsed signals in a continuous stream of data with a known noise spectrum.

The *matched filter* can be derived by maximizing the likelihood with respect to the amplitude of a signal template. Let us assume that a signal can be expressed as $\boldsymbol{\mu}(s_0, t_s, \boldsymbol{\theta}) = s_0 \boldsymbol{\mu}_0(t_s, \boldsymbol{\theta})$, where s_0 governs the amplitude of the signal and $\boldsymbol{\mu}_0(t_s, \boldsymbol{\theta})$ is a signal template with normalized amplitude which depends on parameters $\boldsymbol{\theta}$, and t_s represents a unitary translation/shift of the signal in time. The Fourier transform of the signal template can then be expressed as $\tilde{\mu}(\omega_k; s_0, t_s, \boldsymbol{\theta}) = s_0 \tilde{\mu}_{00}(\omega_k; \boldsymbol{\theta}) e^{i\omega_k t_s}$, where $\tilde{\mu}_{00}(\omega_k; \boldsymbol{\theta})$ is the Fourier transform of the normalized template evaluated at $t_s = 0$. We first define the matched filtering quantities,

$$\begin{aligned} y_{\text{mf}}(t_s, \boldsymbol{\theta}) &\equiv \boldsymbol{\mu}_0(t_s, \boldsymbol{\theta})^T \boldsymbol{\Sigma}^{-1} \mathbf{x} \\ &= 4\Re \left\{ \sum_{k=1}^{n_t/2-1} \frac{\tilde{x}(\omega_k) \tilde{\mu}_{00}(\omega_k; \boldsymbol{\theta})^*}{S_n(\omega_k)} e^{i\omega_k t_s} \delta f \right\}, \\ y_{\boldsymbol{\mu}}(\boldsymbol{\theta}) &\equiv \boldsymbol{\mu}_0(t_s, \boldsymbol{\theta})^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_0(t_s, \boldsymbol{\theta}) \\ &= 4 \sum_{k=1}^{n_t/2-1} \frac{|\tilde{\mu}_{00}(\omega_k; \boldsymbol{\theta})|^2}{S_n(\omega_k)} \delta f, \\ y_{\mathbf{x}} &\equiv \mathbf{x}^T \boldsymbol{\Sigma}^{-1} \mathbf{x} \\ &= 4 \sum_{k=1}^{n_t/2-1} \frac{|\tilde{x}(\omega_k)|^2}{S_n(\omega_k)} \delta f, \end{aligned} \quad (\text{B1})$$

which are here expressed in both the time domain and frequency domain. The frequency domain expressions are computationally efficient to calculate; however, the time domain expressions are also shown here for clarity. By maximizing $\ln \mathcal{L}(\mu(s_0, t_s, \theta); \mathbf{x}, \Sigma)$ with respect to s_0 ,

$$\frac{\partial}{\partial s_0} \left(-\frac{1}{2} (\mathbf{x} - s_0 \mu_0(t_s, \theta))^T \Sigma^{-1} (\mathbf{x} - s_0 \mu_0(t_s, \theta)) \right) = 0, \quad (\text{B2})$$

we get

$$\hat{s}_0(t_s, \theta) = \frac{y_{\text{mf}}(t_s, \theta)}{y_{\mu}(\theta)}, \quad (\text{B3})$$

which is the well-known matched filter estimate of s_0 . By normalizing the matched filter output with the template normalization factor, $\sigma(\theta) = y_{\mu}(\theta)^{-1/2}$, the SNR of the template can be calculated as follows:

$$\text{SNR}_{\text{mf}}(t_s, \theta) = \frac{|y_{\text{mf}}(t_s, \theta)|}{\sqrt{y_{\mu}(\theta)}}. \quad (\text{B4})$$

A *matched filter search* is performed by calculating $\text{SNR}_{\text{mf}}(t_s, \theta)$ for one or more templates for all possible t_s . By selecting the t_s and θ that maximizes $\text{SNR}_{\text{mf}}(t_s, \theta)$, the best matching template and time is found, and $\text{SNR}_{\text{mf}}(t_s, \theta)$ serves as a significance for the signal, i.e., if $\max(\text{SNR}_{\text{mf}}(t_s, \theta)) \gg 1$, the data likely contain a signal consistent with the template. The matched filter quantities can be related to the likelihood by [100]

$$\ln \mathcal{L}(t_s, \theta; \mathbf{x}, S_n) |_{\hat{s}_0(t_s)} = -\frac{1}{2} y_{\mathbf{x}} + \frac{1}{2} \frac{|y_{\text{mf}}(t_s, \theta)|^2}{y_{\mu}(\theta)}. \quad (\text{B5})$$

Since the matched filter likelihood has been analytically maximized over the signal amplitude in Eq. (B2) and is easy to scan for all possible times, it is an extremely computationally efficient way to calculate the maximum likelihood (over time and amplitude) for an already simulated signal template.

As we demonstrate in Sec. III C 3, the matched filter can be used in a likelihood-based reconstruction where the parameters θ are treated as free parameters. This would eliminate two of the parameters from the reconstruction, namely the time and amplitude, and can improve the stability and speed of a reconstruction algorithm. In the case of a radio neutrino signal, t_0 represents the interaction time and s_0 corresponds to the energy of the particle shower.

The matched filter can be generalized to more antennas by adding a sum over the antennas, e.g., in Eq. (B4), and having the trace, signal, and noise power spectral density per antenna. The relative time delay of the signal should be included in the signal templates, and the time t_0 is then the global time offset of the signal in all antennas. In the frequency domain, the most efficient implementation can be achieved by flattening the traces, signals, and noise power spectral densities to a signal vector, and letting k run from 1 to $n_{\text{antennas}} \cdot n_t/2$ to avoid the sum over antennas. This optimized matched filter method is implemented in NuRadioReco.

APPENDIX C: BIASES OF RECONSTRUCTED IN-ICE RADIO NEUTRINO PARAMETERS

To further assess the importance of considering noise correlations in reconstruction of radio neutrino signals (as discussed in Sec. III A 4), we quantify the biases of the reconstruction methods using the $-2 \ln \mathcal{L}$ or χ^2 as objective functions. We define the bias as $\psi_{\theta} = E[\theta_{\text{reco}} - \theta_{\text{true}}]$, which we calculate using 500 reconstructions of the same three neutrino events with different realizations of noise. The resulting biases for each neutrino parameter using the two methods is shown in Table III, including their statistical uncertainties (error of mean). In general, both methods have small biases, often on the order of 0.2σ compared to the uncertainties in Table II. For the parameters of most importance, the neutrino energy and arrival direction, the likelihood method has smaller biases than the χ^2 method. For the parameters related to the vertex position, the results are less conclusive due to statistical uncertainties.

TABLE III. Biases, $\psi_{\theta} = E[\theta_{\text{reco}} - \theta_{\text{true}}]$, of reconstructed signal parameters for the three neutrino events in Table I with 500 different realizations of noise each using the $-2 \ln \mathcal{L}$ or the χ^2 as the objective function in the reconstruction. The bias on the shower energy is given as a percentage of the true shower energy. The statistical uncertainties on the biases are the errors of the means, i.e., the spreads (Table II) divided by the square root of the number of events.

Event	Method	ψ_{E_s}	ψ_{θ_s}	ψ_{ϕ_s}	$\psi_{r_{\text{vertex}}}$	$\psi_{\theta_{\text{vertex}}}$	$\psi_{\phi_{\text{vertex}}}$	ψ_{t_0}
1	$\mathcal{L}$	$(1.5 \pm 0.3)\%$	$(0.13 \pm 0.03)^\circ$	$(0.18 \pm 0.05)^\circ$	$(-0.1 \pm 0.2) \text{ m}$	$(0.034 \pm 0.034)^\circ$	$(-0.009 \pm 0.005)^\circ$	$(-0.001 \pm 0.002) \text{ ns}$
	χ^2	$(2.1 \pm 0.3)\%$	$(0.17 \pm 0.05)^\circ$	$(0.22 \pm 0.06)^\circ$	$(0.1 \pm 0.2) \text{ m}$	$(0.001 \pm 0.039)^\circ$	$(-0.005 \pm 0.006)^\circ$	$(-0.001 \pm 0.002) \text{ ns}$
2	$\mathcal{L}$	$(2.2 \pm 0.5)\%$	$(0.64 \pm 0.07)^\circ$	$(-0.33 \pm 0.39)^\circ$	$(-3.3 \pm 3.6) \text{ m}$	$(0.006 \pm 0.003)^\circ$	$(0.001 \pm 0.001)^\circ$	$(-0.004 \pm 0.002) \text{ ns}$
	χ^2	$(4.2 \pm 0.8)\%$	$(0.68 \pm 0.08)^\circ$	$(-0.33 \pm 0.40)^\circ$	$(3.6 \pm 5.9) \text{ m}$	$(0.004 \pm 0.005)^\circ$	$(0.001 \pm 0.001)^\circ$	$(-0.005 \pm 0.002) \text{ ns}$
3	$\mathcal{L}$	$(3.4 \pm 0.5)\%$	$(0.21 \pm 0.03)^\circ$	$(-0.90 \pm 0.17)^\circ$	$(1.5 \pm 0.9) \text{ m}$	$(-0.003 \pm 0.044)^\circ$	$(0.004 \pm 0.006)^\circ$	$(-0.004 \pm 0.005) \text{ ns}$
	χ^2	$(3.3 \pm 0.6)\%$	$(0.31 \pm 0.03)^\circ$	$(-1.21 \pm 0.21)^\circ$	$(-0.6 \pm 1.1) \text{ m}$	$(0.124 \pm 0.051)^\circ$	$(-0.010 \pm 0.007)^\circ$	$(-0.007 \pm 0.005) \text{ ns}$

- [1] T. Huege, Radio detection of cosmic ray air showers in the digital era, *Phys. Rep.* **620**, 1 (2016).
- [2] F. G. Schröder, Radio detection of cosmic-ray air showers and high-energy neutrinos, *Prog. Part. Nucl. Phys.* **93**, 1 (2017).
- [3] A. Aab *et al.* (Pierre Auger Collaboration), Energy estimation of cosmic rays with the engineering radio array of the Pierre Auger Observatory, *Phys. Rev. D* **93**, 122005 (2016).
- [4] A. Aab *et al.* (Pierre Auger Collaboration), Measurement of the radiation energy in the radio signal of extensive air showers as a universal estimator of cosmic-ray energy, *Phys. Rev. Lett.* **116**, 241101 (2016).
- [5] C. Glaser, M. Erdmann, J. R. Hörandel, T. Huege, and J. Schulz, Simulation of radiation energy release in air showers, *J. Cosmol. Astropart. Phys.* **09** (2016) 024.
- [6] M. Gottowik, C. Glaser, T. Huege, and J. Rautenberg, Determination of the absolute energy scale of extensive air showers via radio emission: Systematic uncertainty of underlying first-principle calculations, *Astropart. Phys.* **103**, 87 (2018).
- [7] K. Mulrey *et al.*, On the cosmic-ray energy scale of the LOFAR radio telescope, *J. Cosmol. Astropart. Phys.* **11** (2020) 017.
- [8] R. de Almeida *et al.* (Pierre Auger Collaboration), Long-term calibration and stability of the Auger Engineering Radio array using the diffuse galactic radio emission, *Proc. Sci. ICRC2023* (2023) 279.
- [9] S. Buitink *et al.*, Method for high precision reconstruction of air shower $X_{\max}$ using two-dimensional radio intensity profiles, *Phys. Rev. D* **90**, 082003 (2014).
- [10] S. Buitink *et al.*, A large light-mass component of cosmic rays at 1017–1017.5 eV from radio observations, *Nature (London)* **531**, 70 (2016).
- [11] A. Abdul Halim *et al.* (Pierre Auger Collaboration), Demonstrating agreement between radio and fluorescence measurements of the depth of maximum of extensive air showers at the Pierre Auger Observatory, *Phys. Rev. Lett.* **132**, 021001 (2024).
- [12] A. Abdul Halim *et al.* (Pierre Auger Collaboration), Radio measurements of the depth of air-shower maximum at the Pierre Auger Observatory, *Phys. Rev. D* **109**, 022002 (2024).
- [13] A. Aab *et al.* (Pierre Auger Collaboration), The Pierre Auger Observatory upgrade—preliminary design report, [arXiv:1604.03637](https://arxiv.org/abs/1604.03637).
- [14] A. Abdul Halim *et al.* (Pierre Auger Collaboration), Status and expected performance of the AugerPrime radio detector, *Proc. Sci. ICRC2023* (2023) 344.
- [15] A. Haungs (IceCube Collaboration), A scintillator and radio enhancement of the IceCube surface detector array, *EPJ Web Conf.* **210**, 06009 (2019).
- [16] IceCube-Gen2 Collaboration, IceCube-Gen2 Technical Design Report, <https://icecube-gen2.wisc.edu/science/publications/TDR>.
- [17] S. Buitink *et al.*, High-resolution air shower observations with the square kilometer array, *Proc. Sci. ICRC2023* (2023) 503.
- [18] S. W. Barwick and C. Glaser, Chapter 6: Radio detection of high energy neutrinos in ice, [arXiv:2208.04971](https://arxiv.org/abs/2208.04971).
- [19] J. Alvarez-Muniz, R. A. Batista, A. B. V., J. Bolmont, M. Bustamante *et al.* (GRAND Collaboration), The Giant Radio array for neutrino detection (GRAND): Science and design, *Sci. China Phys. Mech. Astron.* **63**, 219501 (2020).
- [20] S. Wissel, A. Romero-Wolf, H. Schoorlemmer, W. R. C. Jr., J. Alvarez-Muñiz, E. Zas *et al.*, Prospects for high-elevation radio detection of > 100 PeV tau neutrinos, *J. Cosmol. Astropart. Phys.* **11** (2020) 065.
- [21] A. Anker *et al.* (TAROG, ARIANNA Collaborations), TAROG-M: Radio antenna array on antarctic high mountain for detecting near-horizontal ultra-high energy air showers, *J. Cosmol. Astropart. Phys.* **11** (2022) 022.
- [22] P. W. Gorham *et al.* (ANITA Collaboration), The antarctic impulsive transient antenna ultra-high energy neutrino detector design, performance, and sensitivity for 2006–2007 balloon flight, *Astropart. Phys.* **32**, 10 (2009).
- [23] Q. Abarr *et al.* (PUEO Collaboration), The Payload for Ultrahigh Energy Observations (PUEO): A white paper, *J. Instrum.* **16**, P08035 (2021).
- [24] I. Kravchenko *et al.* (RICE Collaboration), Performance and simulation of the RICE detector, *Astropart. Phys.* **19**, 15 (2003).
- [25] P. Allison *et al.* (ARA Collaboration), Design and initial performance of the Askaryan radio array prototype EeV neutrino detector at the south pole, *Astropart. Phys.* **35**, 457 (2012).
- [26] A. Anker *et al.* (ARIANNA Collaboration), Targeting ultra-high energy neutrinos with the ARIANNA experiment, *Adv. Space Res.* **64**, 2595 (2019).
- [27] J. A. Aguilar *et al.* (RNO-G Collaboration), Design and sensitivity of the radio neutrino observatory in Greenland (RNO-G), *J. Instrum.* **16**, P03025 (2021).
- [28] I. Kravchenko *et al.*, Rice limits on the diffuse ultrahigh energy neutrino flux, *Phys. Rev. D* **73**, 082002 (2006).
- [29] P. Allison, S. Archambault, R. Bard, J. J. Beatty, M. Beheler-Amass, D. Z. Besson *et al.* (ARA Collaboration), Design and performance of an interferometric trigger array for radio detection of high-energy neutrinos, *Nucl. Instrum. Methods Phys. Res., Sect. A* **930**, 112 (2019).
- [30] P. Allison *et al.* (ARA Collaboration), Constraints on the diffuse flux of ultrahigh energy neutrinos from four years of askaryan radio array data in two stations, *Phys. Rev. D* **102**, 043021 (2020).
- [31] P. Allison *et al.* (ARA Collaboration), Low-threshold ultrahigh-energy neutrino search with the Askaryan Radio Array, *Phys. Rev. D* **105**, 122006 (2022).
- [32] A. Anker *et al.* (ARIANNA Collaboration), A search for cosmogenic neutrinos with the ARIANNA test bed using 4.5 years of data, *J. Cosmol. Astropart. Phys.* **03** (2020) 053.
- [33] A. Anker *et al.* (ARIANNA Collaboration), Neutrino vertex reconstruction with in-ice radio detectors using surface reflections and implications for the neutrino energy resolution, *J. Cosmol. Astropart. Phys.* **11** (2019) 030.
- [34] A. Anker *et al.* (ARIANNA Collaboration), Probing the angular and polarization reconstruction of the ARIANNA detector at the south pole, *J. Instrum.* **15**, P09039 (2020).
- [35] S. Barwick *et al.* (ARIANNA Collaboration), Capabilities of ARIANNA: Neutrino pointing resolution and

- implications for future ultra-high energy neutrino astronomy, *Proc. Sci. ICRC2021* (**2021**) 1151.
- [36] A. Anker *et al.* (Arianna Collaboration), Measuring the polarization reconstruction resolution of the ARIANNA neutrino detector with cosmic rays, *J. Cosmol. Astropart. Phys.* **04** (2022) 022.
- [37] A. Anker *et al.* (Arianna Collaboration), Improving sensitivity of the ARIANNA detector by rejecting thermal noise with deep learning, *J. Instrum.* **17**, P03007 (2022).
- [38] A. Anker *et al.* (Arianna Collaboration), Developing new analysis tools for near surface radio-based neutrino detectors, *J. Cosmol. Astropart. Phys.* **10** (2023) 060.
- [39] S. Hallmann, B. Clark, C. Glaser, D. Smith *et al.* (IceCube-Gen2 Collaboration), Sensitivity studies for the IceCube-Gen2 radio array, *Proc. Sci. ICRC2021* (**2021**) 1183 [arXiv:2107.08910].
- [40] A. Nelles *et al.*, The radio emission pattern of air showers as measured with LOFAR—a tool for the reconstruction of the energy and the shower maximum, *J. Cosmol. Astropart. Phys.* **05** (2015) 018.
- [41] C. Glaser, S. de Jong, M. Erdmann, and J. R. Hörandel, An analytic description of the radio emission of air showers based on its emission mechanisms, *Astropart. Phys.* **104**, 64 (2019).
- [42] A. Coleman, Model of the lateral distribution of the radio emission from air showers in the 70–350 MHz frequency band, *Proc. Sci. ARENA2022* (**2023**) 038.
- [43] A. Corstanje *et al.*, Prospects for measuring the longitudinal particle distribution of cosmic-ray air showers with SKA, *Proc. Sci. ARENA2022* (**2023**) 024 [arXiv:2303.09249].
- [44] C. Glaser, Absolute energy calibration of the Pierre Auger Observatory using radio emission of extensive air showers, PhD thesis, RWTH Aachen University, 2017, 10.18154/RWTH-2017-02960.
- [45] C. Glaser (ARIANNA Collaboration), ARIANNA: Measurement of cosmic rays with a radio neutrino detector in antarctica, *EPJ Web Conf.* **216**, 02008 (2019).
- [46] C. Glaser, A. Nelles, I. Plaisier, C. Welling, S. W. Barwick, D. García-Fernández, G. Gaswint, R. Lahmann, and C. Persichilli, NuRadioReco: A reconstruction framework for radio neutrino detectors, *Eur. Phys. J. C* **79**, 464 (2019).
- [47] C. Welling, C. Glaser, and A. Nelles, Reconstructing the cosmic-ray energy from the radio signal measured in one single station, *J. Cosmol. Astropart. Phys.* **10** (2019) 075.
- [48] C. Glaser and S. W. Barwick, An improved trigger for Askaryan radio detectors, *J. Instrum.* **16**, T05001 (2021).
- [49] C. Glaser, A. Coleman, and T. Glusenkamp, NuRadioOpt: Optimization of radio detectors of ultra-high energy neutrinos through deep learning and differential programming, *Proc. Sci. ICRC2023* (**2023**) 1114.
- [50] C. Glaser *et al.*, NuRadioMC: Simulating the radio emission of neutrinos from interaction to detector, *Eur. Phys. J. C* **80**, 77 (2020).
- [51] G. G. Gaswint, Quantifying the neutrino energy and pointing resolution of the ARIANNA detector, PhD thesis, UC, Irvine, 2021.
- [52] I. Plaisier, S. Bouma, and A. Nelles, Reconstructing the arrival direction of neutrinos in deep in-ice radio detectors, *Eur. Phys. J. C* **83**, 443 (2023).
- [53] S. Bouma *et al.* (IceCube-Gen2 Collaboration), Direction reconstruction performance for IceCube-Gen2 radio, *Proc. Sci. ICRC2023* (**2023**) 1045.
- [54] L. D. Landau and I. Pomeranchuk, Limits of applicability of the theory of bremsstrahlung electrons and pair production at high-energies, *Dokl. Akad. Nauk Ser. Fiz.* **92**, 535 (1953).
- [55] L. D. Landau and I. Pomeranchuk, Electron-Cascade processes at ultra-high energies, *Dokl. Akad. Nauk SSSR* **92**, 589 (1965).
- [56] A. B. Migdal, Bremsstrahlung and pair production in condensed media at high-energies, *Phys. Rev.* **103**, 1811 (1956).
- [57] R. Abbasi *et al.* (IceCube Collaboration), Observation of high-energy neutrinos from the galactic plane, *Science* **380**, adc9818 (2023).
- [58] C. Glaser, S. McAleer, S. Stjärnholm, P. Baldi, and S. W. Barwick, Deep-learning-based reconstruction of the neutrino direction and energy for in-ice radio detectors, *Astropart. Phys.* **145**, 102781 (2023).
- [59] T. Glusenkamp, Unifying supervised learning and VAEs: Coverage, systematics and goodness-of-fit in normalizing-flow based neural network models for astro-particle reconstructions, *Eur. Phys. J. C* **84**, 163 (2024).
- [60] T. Glusenkamp *et al.* (IceCube Collaboration), Conditional normalizing flows for IceCube event reconstruction, *Proc. Sci. ICRC2023* (**2023**) 1003 [arXiv:2309.16380].
- [61] N. Heyer *et al.* (IceCube-Gen2 Collaboration), Deep learning based event reconstruction for the IceCube-Gen2 radio detector, *Proc. Sci. ICRC2023* (**2023**) 1102 [arXiv:2308.00164].
- [62] I. R. C. Committee, Recommendations and Reports of the CCIR, 1982: Spectrum utilization and monetary. Recommendations and Reports of the CCIR, 1982: . International Telecommunication Union, 1982.
- [63] K. Mulrey *et al.*, Calibration of the LOFAR low-band antennas using the galaxy and a model of the signal chain, *Astropart. Phys.* **111**, 1 (2019).
- [64] E. S. Hong, Searching for ultra-high energy neutrinos with data from a prototype station of the askaryan radio array, PhD thesis, The Ohio State University, 2014.
- [65] T. Meures, Development of a sub-glacial array of radio antennas for the detection of the flux of GZK neutrinos, PhD thesis, U. Brussels, Brussels U., IIHE, 2014, 10.1007/978-3-319-18756-3.
- [66] L. Cremonesi *et al.* (ANITA Collaboration), The simulation of the sensitivity of the Antarctic Impulsive Transient Antenna (ANITA) to Askaryan radiation from cosmogenic neutrinos interacting in the Antarctic Ice, *J. Instrum.* **14**, P08011 (2019).
- [67] S. Ali *et al.* (ARA Collaboration), Progress towards an array-wide diffuse UHE neutrino search with the Askaryan radio array, arXiv:2409.03854.
- [68] S. Agarwal, J. A. Aguilar, N. Alden, S. Ali, P. Allison, M. Betts *et al.*, Calibration of the radio neutrino observatory in Greenland using thermal noise, *Proc. Sci. ICRC2025* (**2025**) 1003.
- [69] S. S. Wilks, The large-sample distribution of the likelihood ratio for testing composite hypotheses, *Ann. Math. Stat.* **9**, 60 (1938).

- [70] C. R. Rao, *Linear Statistical Inference and its Applications* (Wiley, New York, 1973).
- [71] M. Ravn, C. Glaser, T. Glüsenkamp, and A. Coleman, Likelihood reconstruction of radio signals of neutrinos and cosmic rays, in *10th International Workshop on Acoustic and Radio EeV Neutrino Detection Activities* (2024), 9, [arXiv:2409.11888](#).
- [72] G. A. Askar'yan, Excess negative charge of an electron-photon shower and its coherent radio emission, *Zh. Eksp. Teor. Fiz.* **41**, 616 (1961).
- [73] G. A. Askar'yan, Coherent radio emission from cosmic showers in air and in dense media, *Sov. J. Exp. Theor. Phys.* **21**, 658 (1965).
- [74] J. Alvarez-Muniz, E. Marques, R. A. Vazquez, and E. Zas, Coherent radio pulses from showers in different media: A unified parameterization, *Phys. Rev. D* **74**, 023007 (2006).
- [75] J. Alvarez-Muniz, C. W. James, R. J. Protheroe, and E. Zas, Thinned simulations of extremely energetic showers in dense media for radio applications, *Astropart. Phys.* **32**, 100 (2009).
- [76] J. Alvarez-Muniz, W. R. Carvalho, Jr., M. Tueros, and E. Zas, Coherent Cherenkov radio pulses from hadronic showers up to EeV energies, *Astropart. Phys.* **35**, 287 (2012).
- [77] A. Coleman, O. Ericsson, C. Glaser, and M. Bustamante, Flavor composition of ultrahigh-energy cosmic neutrinos: Measurement forecasts for in-ice radio-based EeV neutrino telescopes, *Phys. Rev. D* **110**, 023044 (2024).
- [78] S. W. Barwick *et al.* (ARIANNA Collaboration), Observation of classically ‘forbidden’ electromagnetic wave propagation and implications for neutrino detection, *J. Cosmol. Astropart. Phys.* **07** (2018) 055.
- [79] P. Allison *et al.* (ARA Collaboration), Design and performance of an interferometric trigger array for radio detection of high-energy neutrinos, *Nucl. Instrum. Methods Phys. Res., Sect. A* **930**, 112 (2019).
- [80] H. Dembinski, P. Ongmongkolkul *et al.*, scikit-hep/iminuit.
- [81] F. James and M. Roos, Minuit: A system for function minimization and analysis of the parameter errors and correlations, *Comput. Phys. Commun.* **10**, 343 (1975).
- [82] J. A. Aguilar *et al.*, Reconstructing the neutrino energy for in-ice radio detectors: A study for the Radio Neutrino Observatory Greenland (RNO-G), *Eur. Phys. J. C* **82**, 147 (2022).
- [83] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, Asymptotic formulae for likelihood-based tests of new physics, *Eur. Phys. J. C* **71**, 1554 (2011).
- [84] H. Cramér, *Mathematical Methods of Statistics (PMS-9)* (Princeton University Press, Princeton, NJ, 1999).
- [85] C. Radhakrishna Rao, Information and the accuracy attainable in the estimation of statistical parameters, *Bull. Calcutta Math. Soc.* **37**, 81 (1945).
- [86] C. Glaser, A. Coleman, and T. Glüsenkamp, NuRadioOpt: Optimization of radio detectors of ultra-high energy neutrinos through deep learning and differential programming, *Proc. Sci. ICRC2023* (2023) 1114.
- [87] C. Glaser, M. L. Ravn, T. Glüsenkamp, A. Özcelikkale, and A. Coleman, Likelihood-reconstruction for radio detectors of cosmic rays and neutrinos, *Proc. Sci. ICRC2025* (2025) 271.
- [88] T. Huege, M. Ludwig, and C. W. James, Simulating radio emission from air showers with CoREAS, *AIP Conf. Proc.* **1535**, 128 (2013).
- [89] T. Huege (Pierre Auger Collaboration), The radio detector of the Pierre Auger Observatory—status and expected performance, *EPJ Web Conf.* **283**, 06002 (2023).
- [90] S. Martinelli, T. Huege, D. Ravignani, and H. Schoorlemmer, Quantifying energy fluence and its uncertainty for radio emission from particle cascades in the presence of noise, *Astropart. Phys.* **168**, 103091 (2025).
- [91] S. Agarwal *et al.*, Validating the performance of the radio neutrino observatory in greenland using cosmic-ray air showers, [arXiv:2512.17664](#).
- [92] J. Neyman and E. S. Pearson, Ix. on the problem of the most efficient tests of statistical hypotheses, *Phil. Trans. R. Soc. A* **231**, 289 (1933).
- [93] J. Henrichs, A. Nelles, J. A. A. Aguilar, P. Allison, D. Z. Besson, A. Bishop *et al.*, Searching for cosmic-ray air showers with RNO-G, *Proc. Sci. ICRC2023* (2023) 259.
- [94] A. Nelles, Cosmic ray detection with RNO-G, *Proc. Sci. ARENA2024* (2024) 006.
- [95] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries, *Phys. Rev. D* **85**, 122006 (2012).
- [96] C. M. Biwer, C. D. Capano, S. De, M. Cabero, D. A. Brown, A. H. Nitz, and V. Raymond, PyCBC inference: A Python-based parameter estimation toolkit for compact binary coalescence signals, *Publ. Astron. Soc. Pac.* **131**, 024503 (2019).
- [97] K. Cannon, S. Caudill, C. Chan, B. Cousins, J. D. E. Creighton, B. Ewing *et al.*, GStLAL: A software framework for gravitational wave discovery, *SoftwareX* **14**, 100680 (2021).
- [98] F. Aubin *et al.*, The MBTA pipeline for detecting compact binary coalescences in the third LIGO–Virgo observing run, *Classical Quantum Gravity* **38**, 095004 (2021).
- [99] Q. Chu *et al.*, SPIIR online coherent pipeline to search for gravitational waves from compact binary coalescences, *Phys. Rev. D* **105**, 024023 (2022).
- [100] I. n. Zubeldia, A. Rotti, J. Chluba, and R. Battye, Understanding matched filters for precision cosmology, *Mon. Not. R. Astron. Soc.* **507**, 4852 (2021).