PROCEEDINGS OF SCIENCE

Optimization pipeline for in-ice radio neutrino detectors

Martin Ravn,[a,*] **Philipp Pilar,**[b] **Christian Glaser,**[a,c] **Niklas Wahlström**[b] **and Thorsten Glüsenkamp**[c]

[a] *Dept. of Physics and Astronomy, Uppsala University, Box 516, 751 37 Uppsala, Sweden*
[b] *Dept. of Information Technology, Uppsala University, Box 337, 751 05 Uppsala, Sweden*
[c] *Dept. of Physics, TU Dortmund University, Dortmund, Germany*
[d] *Oskar Klein Centre and Dept. of Physics, Stockholm University, SE-10691 Stockholm, Sweden*

E-mail: martin.ravn@physics.uu.se

In-ice radio detection of neutrinos is a rapidly growing field and a promising technique for discovering the predicted but yet unobserved ultra-high-energy astrophysical neutrino flux. With the ongoing construction of the Radio Neutrino Observatory in Greenland (RNO-G) and the planned radio extension of IceCube-Gen2, we have a unique opportunity to improve the detector design now and accelerate the experimental outcome in the field for the coming decades. In this contribution, we present the first end-to-end in-ice radio neutrino simulation, detection, and reconstruction pipeline using generative machine learning models and differentiable programming. We demonstrate how this framework can be used to optimize the antenna layout of detectors to achieve the best possible reconstruction resolution of neutrino parameters.

Fifth MODE Workshop on Differentiable Programming for Experiment Design (MODE2025)
8-13 June 2025
Kolymbari, Crete, Greece

*Speaker

 ISSN 1824-8039 . Published by SISSA Medialab.

PoS(MODE2025)031

1. Introduction

Over the past decades, neutrino physics has transitioned from studying terrestrial and solar sources to exploring the astrophysical realm. Large-scale neutrino experiments, like the IceCube Neutrino Observatory, have successfully proven that a high-energy astrophysical neutrino flux exists and that neutrinos serve as excellent astrophysical messengers [1]. Since neutrinos travel undisturbed through the interstellar and intergalactic medium, they point directly back to their origin and can be used to probe the most extreme events in the Universe. They are therefore a key component of multi-messenger astrophysics and promise great insights in the coming years [2].

At the ultra-high-energy scale (10 PeV and above), the neutrino flux is too faint for traditional optical-based telescopes, whose limited observation volumes prevent consistent detection. However, a flux at these energies is predicted by several models, and the science cases for detecting it are plentiful for both particle physics (neutrino-nucleon cross-section, beyond-standard-model effects, etc.) and astrophysics (particle acceleration in sources, cosmogenic neutrinos, etc.) [3]. One of the most promising techniques for measuring the ultra-high-energy neutrino flux in a cost-effective and scalable way is in-ice radio detector arrays [4], such as the prototype arrays ARIANNA [5] and ARA [6], the under-construction RNO-G [7], and the planned IceCube-Gen2 radio array [8].

In-ice radio detectors aim to exploit the fact that particle showers from ultra-high-energy neutrino interactions emit radio pulses via the Askaryan effect and that ice is mostly transparent to radio signals with attenuation lengths of hundreds of meters to a kilometer. This means that with only a few antennas deployed near the surface of an ice sheet, e.g., in Greenland or at the South Pole, several cubic kilometers of ice can be monitored for neutrino interactions. The IceCube-Gen2 radio array is planned to consist of around 200 shallow stations with 8 antennas buried a few meters below the surface of the ice, and around 150 hybrid stations with a shallow section and a deep section of 16 antennas up to 150 m below the surface. The IceCube-Gen2 shallow radio station layout is shown in Figure 1, which will be the focus of this contribution.

Since RNO-G is under construction and the IceCube-Gen2 radio array is planned, it is now

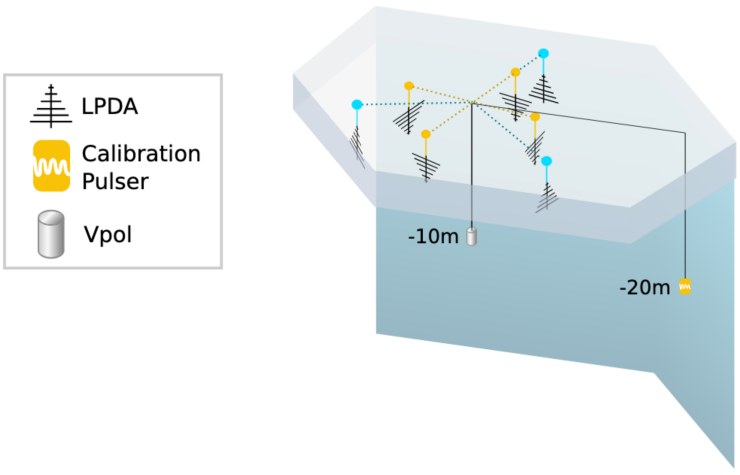

Figure 1: Illustration of the planned radio station layouts for the IceCube-Gen2 radio array. The figure is adapted from [8].

crucial to investigate how the detectors can be improved to maximize their scientific outcome. This is the focus of the NuRadioOpt project [9], which aims to improve radio detectors in two ways: with deep learning-based triggers and end-to-end optimization using differentiable programming. The latter is the focus of this contribution.

In recent years, innovations in differentiable programming, one of the core concepts behind deep learning, have proven useful for the optimization of particle physics detectors [10]. Traditional detector optimization is slow and often infeasible in cases where Monte Carlo generation is expensive and the detector parameter space is large. If, however, a full simulation, reconstruction, and analysis pipeline (see illustration in Figure 2) can be implemented in a differentiable framework, e.g., PyTorch or JAX, the gradient of a loss with respect to the parameters of a detector can automatically be calculated. If the loss is chosen to be a combination of the sensitivity to a specific science output and the cost, this method allows for efficient gradient-based detector optimization of detectors with large numbers of parameters.

In this contribution, we present the first differentiable end-to-end optimization pipeline for in-ice radio neutrino detectors. For the simulation, we ported a state-of-the-art radio neutrino simulation code to a differentiable framework. As a proxy for the reconstruction, we use the Fisher information to estimate the uncertainty of the reconstructed parameters, and the loss is the neutrino arrival direction reconstruction uncertainty, which is equivalent to optimizing for point source discovery. Several simplifications have been made to the pipeline, which we will outline in the following and discuss how they can be addressed in future work. As a proof-of-concept, we optimize a shallow IceCube-Gen2 radio station with four antennas, demonstrating the viability and potential of the approach.

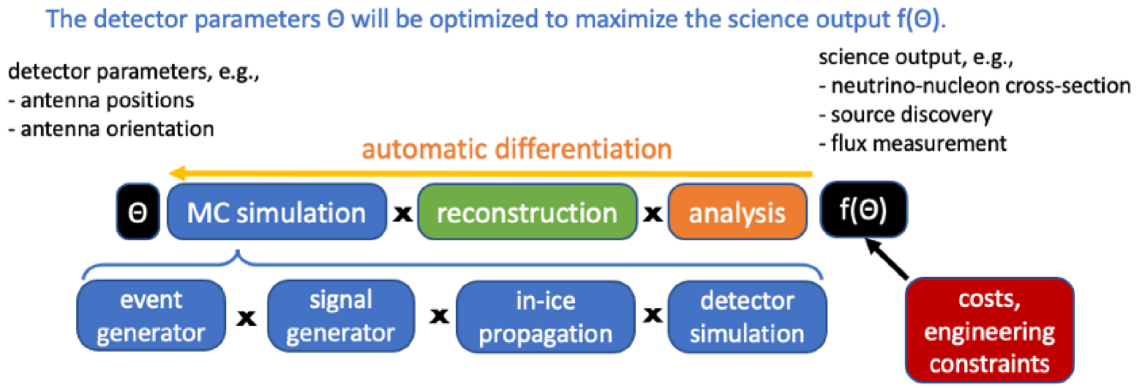

Figure 2: Illustration of a differentiable simulation, reconstruction, and analysis pipeline needed for end-to-end detector optimization. The figure is adapted from [9].

2. Neutrino signal pipeline

The central component of our differentiable detector optimization pipeline is a neutrino and detector simulation framework, adapted from the open-source Monte Carlo software NuRadioMC [11]. We first ported the relevant modules to a stand-alone NumPy implementation, which is used for validation and debugging. The code was then ported to PyTorch, with deep learning-based surrogate models introduced for non-differentiable components.

The neutrino signal pipeline is a function that receives the parameters of a neutrino event and an antenna as input and returns the signal as observed by the antenna as the output. The event parameters are the energy of the particle shower ($E_{\text{show.}}$), the neutrino arrival direction spherical angles (θ_ν and ϕ_ν), the vertex position in spherical coordinates ($\theta_{\text{vert.}}$, $\phi_{\text{vert.}}$, and $r_{\text{vert.}}$), and interaction time (t_0). The antenna is characterized by its position ($x_{\text{ant.}}$, $y_{\text{ant.}}$, and $z_{\text{ant.}}$) and three angles ($\theta_{\text{ant.}}$, $\phi_{\text{ant.}}$, $\psi_{\text{ant.}}$), parameterizing its orientation and rotation. The output is the measured voltage trace $V(t)$, sampled at n_t regularly spaced discrete time steps. In this work, we use a sampling rate of 2.0 GHz and $n_t = 1024$, which is typical for in-ice radio neutrino detectors.

The simulation proceeds in four steps: ray-tracing, Askaryan emission, propagation effects, and detector response. For computational efficiency, the simulation is performed in a different order than the underlying physical processes, and thus differs from the sequence shown in Figure 2. The first step of the simulation is the ray-tracing, which finds the paths a radio signal can travel between the position of the neutrino interaction vertex and the antenna. Since ice sheets have a non-constant index of refraction, which results in bending of rays, the paths that radio waves travel are non-trivial, and various algorithms exist to find them. The algorithms are, in general, difficult to implement in a differentiable framework since they rely on internal minimization algorithms. For this work, we chose to simplify the propagation by using straight-line ray-tracing, i.e., corresponding to a constant refractive index. This can be improved upon in future work by implementing a deep learning-based surrogate model to approximate ray-tracing for a realistic index of refraction profile.

The second part of the pipeline is to calculate the Askaryan emission for the viewing angle returned by the ray-tracing. This step is also not trivial to implement in a differentiable framework since it includes stochastic processes. We employ a deep learning surrogate model described in Ref. [12], which allows us to differentiate through the Askaryan emission process.

The surrogate model employs a modular model architecture that splits the signal generation into three stages: a generator that produces reference signals at a fixed viewing angle (65.57°) and distance (1 km), a neural network to adjust the signal to the desired viewing angle, and a neural network that predicts the signal amplitude. This separation has several advantages. First, by isolating the viewing-angle transformation, derivatives do not need to propagate through the generator itself, allowing the use of computationally expensive or even non-differentiable generators. Second, it ensures that the radio signals from the same neutrino event can be generated consistently for different antennas. Finally, introducing the amplitude only at the end of the signal generation enables the preceding modules to operate on normalized data, which is often required for effective training of machine learning models.

The distance dependence of the Askaryan signal is approximated by scaling it with the inverse of the travel distance. In future work, the full distance dependence can be treated with a neural network in a similar way to the viewing-angle dependence. After the emitted electric field at the antenna is estimated, the effects on the signal from propagating through the ice are applied. This includes frequency-dependent attenuation, which ideally should be averaged over the path. For simplicity, we approximate it with the attenuation at the center of the path.

The final step of the simulation is to apply the response of the antenna to the electric field. We implemented the response of a log-periodic dipole antenna (LPDA), planned to be used for the IceCube-Gen2 shallow stations, which has an analytical approximation in NuRadioMC. Additional antenna models, such as the vertically or horizontally polarized dipole antennas planned for the deep

stations, can be implemented with similar approximations in future work. Additionally, we apply a band-pass filter (2nd order high pass Butterworth filter at 80 MHz and 10th order Butterworth low pass filter at 400 MHz) to approximate the rest of the signal chain.

To validate the differentiable pipeline, we compare identical neutrino events simulated with the PyTorch implementation, the NumPy implementation, and NuRadioMC under comparable settings and simplifications. Figure 3 shows perfect agreement between the three pipelines. The differentiable pipeline is hence suitable for the proof-of-concept optimization study presented here.

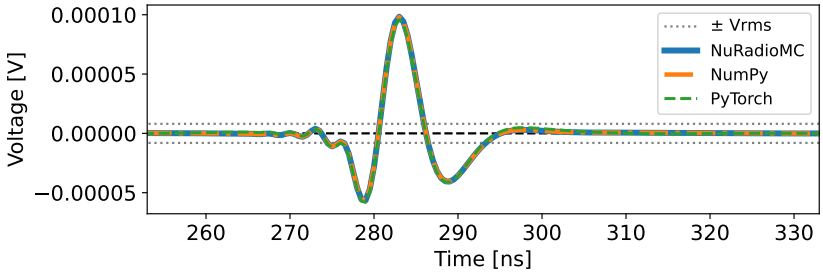

Figure 3: Comparison of neutrino signal calculated by the three simulation pipelines: the differentiable PyTorch implementation, the standalone NumPy implementation, and NuRadioMC under comparable settings and simplifications. The amplitude of the assumed noise (Vrms) is shown for reference.

3. Reconstruction proxy: Fisher information

A maximum likelihood method for reconstructing in-ice radio neutrino events was recently presented in Ref. [13]. This approach has the advantage of correctly estimating uncertainties on reconstructed parameters. It is, however, difficult to directly implement it in a differentiable framework because it relies on minimization algorithms. As an alternative, we use the Fisher information matrix to estimate reconstruction uncertainties. This can be viewed as an approximation to the maximum-likelihood uncertainty contours, averaged over noise realizations. For Gaussian noise, as present in radio detectors, the Fisher information matrix takes the form:

$$\mathcal{I}_{nm} = \sum_{\text{ant.}} \frac{\partial \boldsymbol{\mu}^{\mathrm{T}}}{\partial \theta_n} \boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\mu}}{\partial \theta_m} + \frac{1}{2} \mathrm{tr}\left(\boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \theta_n} \boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \theta_m} \right), \tag{1}$$

where $\boldsymbol{\mu} = \big(V(t_0),\, V(t_1),\, \ldots,\, V(t_{n_t-1})\big)$ is the neutrino signal in an antenna, i.e., the output of the simulation pipeline, $\boldsymbol{\theta} = \big(E_{\text{show.}},\, \theta_\nu,\, \phi_\nu,\, \theta_{\text{vert.}},\, \phi_{\text{vert.}},\, r_{\text{vert.}},\, t_0\big)$ are the neutrino parameters, and $\boldsymbol{\Sigma}$ is the covariance matrix of the noise, derived from the frequency spectrum of the noise. The second term vanishes because the noise in a radio detector is independent of the neutrino parameters. The inverse Fisher matrix serves, via the Cramér–Rao bound, as an estimate of the covariance of the reconstructed parameters, σ_{nm}. The inverse Fisher information matrix for an example neutrino event is shown in Figure 4 (left).

By extracting the 2x2 submatrix corresponding to the neutrino arrival direction (θ_ν, ϕ_ν), we obtain the marginalized uncertainty of those parameters. This submatrix defines the uncertainty ellipse with major and minor axes $l_{\text{maj./min.}} = \left(\frac{1}{2}(\sigma_{\theta_\nu} + \sigma_{\phi_\nu}) \pm \sqrt{\frac{1}{4}(\sigma_{\theta_\nu} - \sigma_{\phi_\nu})^2 + \sigma^2_{\theta_\nu \phi_\nu}} \right)^{1/2}$ and

area $A = \pi\, l_{\text{min.}}\, l_{\text{maj.}}$. An example of a reconstructed simulated event with one noise realization and its uncertainty contour is shown in Figure 4 (right) along with the Fisher uncertainty estimate centered at the true parameter values. This approach offers an extremely efficient way to estimate reconstruction uncertainties for neutrino events, as Eq. 1 is straightforward to evaluate and the required signal derivatives are obtained directly from the differentiable pipeline. We hence adopt the area of the Fisher uncertainty ellipse as the loss function to minimize in the detector optimization.

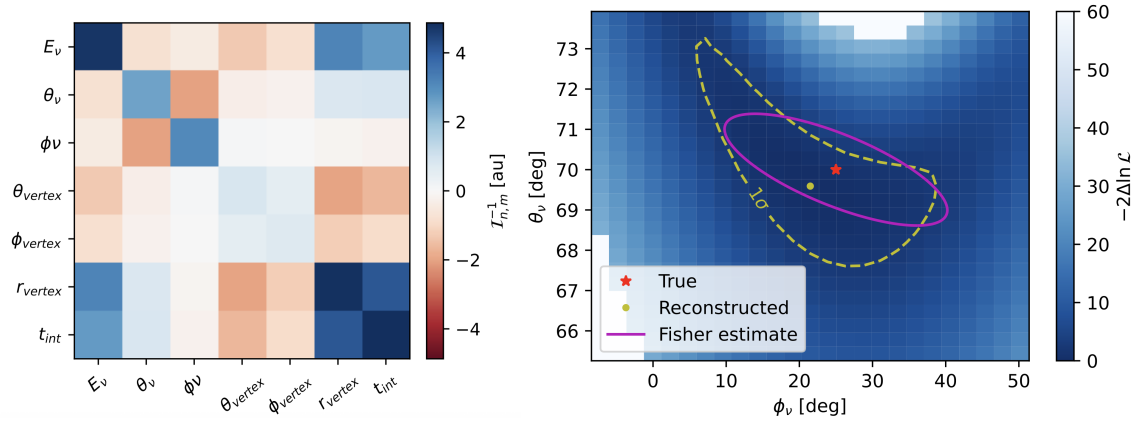

Figure 4: Left: Example inverse Fisher information matrix for a neutrino event. Right: An example of a reconstructed simulated neutrino event and its uncertainty contour. The Fisher information estimate of the uncertainty is shown as an ellipse centered at the true neutrino parameters, which should on average have the same area as the uncertainty contour.

4. Detector optimization experiment

To demonstrate the viability and potential of the method, we perform a proof-of-concept detector optimization. The detector is a simple shallow station with four downward-facing LPDAs. The initial positions and orientations match those planned for the shallow IceCube-Gen2 radio stations (Figure 1), i.e., a cross-like pattern with antennas pointing straight down. The depths of the antennas were fixed at 3 m below the surface during the optimization. The optimization objective is to minimize the natural logarithm of the neutrino direction uncertainty area by adjusting the antenna positions and orientations.

For event generation, we select neutrino parameters that are likely to give a detectable signal. The arrival direction, θ_ν, is drawn uniformly between horizontal and 45° above the horizon, while ϕ_ν is uniform between 0° and 360°. The interaction vertex is then chosen such that the top of the Askaryan cone intersects the detector center, with random offsets of ±10° in θ_{vert} and ϕ_{vert}, and $r_{\text{vert.}}$ between 1 km and 3 km. Shower energies are drawn uniformly between 0.1 EeV and 1 EeV, and interaction times are set so that signals fall near the center of the trace. While this distribution is not fully realistic, it efficiently produces plausible signals suitable for a proof-of-concept study.

The optimization was performed using the Adam optimizer with a learning rate of 0.003 and a batch size of 98 events, averaging the loss across each batch. An early stopping condition was implemented to terminate the training if the loss failed to improve over 1000 iterations. Figure 5 (left) shows $\log_{10}$ of the neutrino direction uncertainty area for each event over the training, with the

running mean indicated in red. We show the running mean (calculated in $\log_{10}$ space) on a linear scale in Figure 5 (right). The optimized layout achieves an improvement of roughly a factor of 3 in the neutrino direction reconstruction uncertainty.

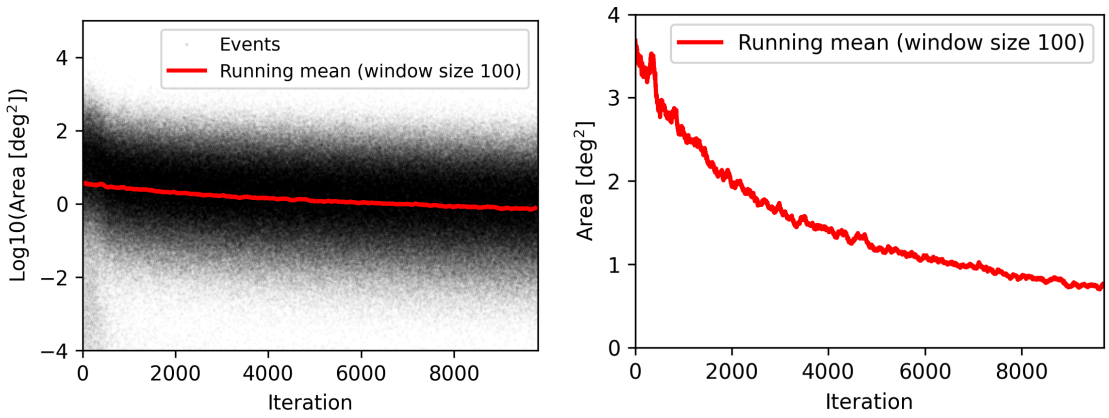

Figure 5: Left: The logarithm of the neutrino arrival direction reconstruction uncertainty in square degrees for each event over the training with 98 events per iteration. The running mean with a window size of 100 iterations is shown in red. Right: A zoom in on the running mean (calculated in $\log_{10}$ space) on a linear scale, which shows that the uncertainties improve during the optimization.

The evolution of antenna positions and orientations during optimization is shown in Figure 6. The azimuth angles, $\theta_{\text{ant.}}$, are initialized at 180°, corresponding to antennas pointing straight downward, and rapidly converge to about 35° from vertical. This adjustment is physically reasonable, as neutrino signals typically arrive from the sides, and antennas are more sensitive when oriented toward the incoming electric field. In addition, antennas 0 and 2 move apart along the x-axis, while antennas 1 and 3 do the same along the y-axis, effectively enlarging the detector. A wider layout is expected to be beneficial for reconstruction, as it provides a larger lever arm and samples a larger portion of the Askaryan radiation cone, and thereby improves the constraint on the neutrino arrival direction. Overall, the end-to-end optimization pipeline successfully enhances the reconstruction resolution of this simple detector, and the adjustments it makes are physically interpretable.

5. Conclusion

In this work, we presented the first differentiable end-to-end optimization pipeline for in-ice radio neutrino detectors. To demonstrate its viability, we applied the method to a proof-of-concept optimization of a simple IceCube-Gen2 station with four antennas, showing that the pipeline can produce physically interpretable improvements in detector performance. The current simulation is simplified in a few key areas; these can, however, be addressed in future work, e.g., through deep learning surrogate models. In particular, differentiable ray-tracing and more realistic event generation are needed. With this framework now established, the next step is to enhance the realism of the simulation and extend it to include additional antenna types and more complex detector designs. Ultimately, the goal is to optimize the full IceCube-Gen2 radio array to maximize its scientific potential.

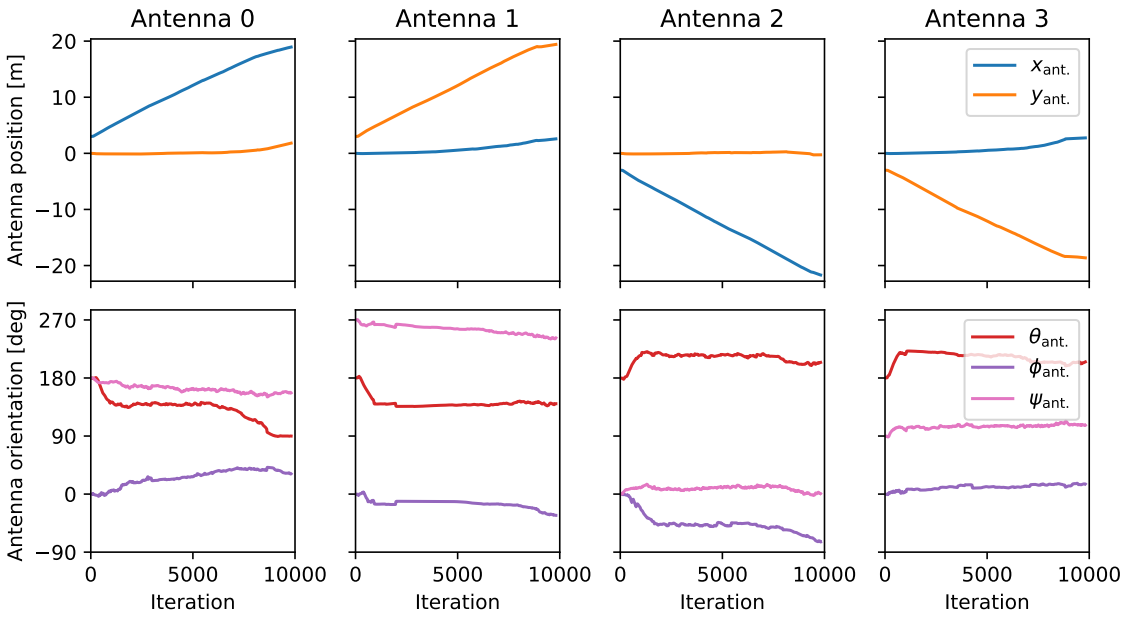

Figure 6: The positions and orientations of the four antennas during the optimization process.

Acknowledgements

The work is supported by the European Union (ERC, NuRadioOpt, 101116890) and by the Swedish Research Council (VR) via the projects 2021-04321 and 2021-05449. This work used resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

References

[1] **IceCube** Collaboration, M. G. Aartsen *et al.*, *Phys. Rev. Lett.* **113** (2014) 101101.
[2] P. Mészáros, D. B. Fox, C. Hanna, and K. Murase, *Nature Rev. Phys.* **1** (2019) 585–599.
[3] M. Ackermann *et al.*, *JHEAp* **36** (2022) 55–110.
[4] S. Barwick and C. Glaser, *Neutrino Physics and Astrophysics, World Scientific* (2023) .
[5] **ARIANNA** Collaboration, A. Anker *et al.*, *Adv. Space Res.* **64** (2019) 2595–2609.
[6] **ARA** Collaboration, P. Allison *et al.*, *Astropart. Phys.* **35** (2012) 457–477.
[7] **RNO-G** Collaboration, J. A. Aguilar *et al.*, *JINST* **16** (2021) P03025.
[8] **IceCube-Gen2** Collaboration, "IceCube-Gen2 Technical Design Report." https://icecube-gen2.wisc.edu/science/publications/TDR.
[9] C. Glaser, A. Coleman, and T. Glusenkamp, *PoS* **ICRC2023** (2023) 1114.
[10] **MODE** Collaboration, T. Dorigo *et al.*, *Rev. Phys.* **10** (2023) 100085.
[11] C. Glaser *et al.*, *Eur. Phys. J. C* **80** no. 2, (2020) 77.
[12] P. Pilar, M. Ravn, C. Glaser, and N. Wahlström, *arXiv:2509.10274* (9, 2025) .
[13] M. Ravn, C. Glaser, T. Glüsenkamp, and A. Coleman, *PoS(ARENA2024)054* .