



PAPER • OPEN ACCESS

Deep learning framework for enhanced neutrino reconstruction of single-line events in the ANTARES telescope

To cite this article: A Albert *et al* 2026 *Mach. Learn.: Sci. Technol.* **7** 035004

View the [article online](#) for updates and enhancements.

You may also like

- [Deep learning reconstruction in ANTARES](#)
J. García-Méndez, N. Geißelbrecht, T. Eberl et al.
- [Event identification for KM3NeT/ARCA](#)
Thomas Heid and on behalf of the KM3NeT Collaboration
- [Event reconstruction for KM3NeT/ORCA using convolutional neural networks](#)
S. Aiello, A. Albert, S. Alves Garre et al.



PAPER

OPEN ACCESS

RECEIVED
12 January 2026

REVISED
30 March 2026

ACCEPTED FOR PUBLICATION
9 April 2026

PUBLISHED
29 April 2026

Original content from
this work may be used
under the terms of the
Creative Commons
Attribution 4.0 licence.

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Deep learning framework for enhanced neutrino reconstruction of single-line events in the ANTARES telescope

A Albert^{1,2}, S Alves³, M André⁴, M Ardid^{5,*} , S Ardid^{5,*} , J-J Aubert⁶, J Aublin⁷, B Baret⁷, S Basa⁸, Y Becherini⁷ , B Belhorma⁹, F Benfenati^{10,11}, V Bertin⁶, S Biagi¹², J Boumaaza¹³, M Bouta¹⁴, M C Bouwhuis¹⁵, H Brânzaş¹⁶, R Bruijn^{15,17}, J Brunner⁶ , J Busto⁶, B Caiffi¹⁸, D Calvo³, S Campion^{19,20}, A Capone^{19,20}, F Carenini^{10,11}, J Carr⁶, V Carretero³, T Cartraud⁷, S Celli^{19,20}, L Cerisy⁶, M Chabab²¹, R Cherkaoui El Moursli¹³, T Chiarusi¹⁰, M Circella²², J A B Coelho⁷, A Coleiro⁷, R Coniglione¹², P Coyle⁶, A Creusot⁷, A F Díaz²³, B De Martino⁶, C Distefano¹², I Di Palma^{19,20}, C Donzaud^{7,24}, D Dornic⁶, D Drouhin¹, T Eberl²⁵ , A Eddymaoui¹³, T van Eeden¹⁵, D van Eijk¹⁵, S El Hedri⁷, N El Khayati¹³, A Enzenhöfer⁶, P Fermani^{19,20}, G Ferrara¹², F Filippini^{10,11}, L Fusco²⁶, S Gagliardini^{19,20}, J García-Méndez^{5,*} , C Gatus Oliver¹⁵, P Gay^{7,27}, N Geißelbrecht²⁵, H Glotin²⁸, R Gozzini³, R Gracia Ruiz²⁵, K Graf²⁵, C Guidi^{18,29}, L Haegel⁷, H van Haren³⁰, A J Heijboer¹⁵, Y Hello³¹, L Hennig²⁵, J J Hernández-Rey³, J Höbl²⁵, F Huang⁶, G Illuminati^{10,11}, B Jisse-Jung¹⁵, M de Jong^{15,32}, P de Jong^{15,17}, M Kadler³³, O Kalekin²⁵, U Katz²⁵, A Kouchner⁷ , I Kreykenbohm³⁴, V Kulikovskiy¹⁸, R Lahmann²⁵, M Lamoureux⁷, A Lazo³, D Lefèvre^{35,36}, E Leonora⁷, G Levi^{10,11}, S Le Stum⁶, S Loucatos³⁷, J Manczak³, M Marcelin⁸, A Margiotta^{10,11}, A Marinelli^{39,40}, J A Martínez-Mora⁵, P Migliozi³⁹, A Moussa¹⁴, R Muller¹⁵, S Navas⁴¹ , E Nezri⁸, B Ó Fearraigh¹⁵ , E Oukacha⁷, A M Păun¹⁶, G E Păvălaş¹⁶, S Peña-Martínez⁷, M Perrin-Terrin⁶, P Piattelli¹², C Poiré²⁶, V Popa¹⁶, T Pradier¹, N Randazzo³⁷, D Real³, G Riccobene¹², A Romanov^{18,29}, A Sánchez Losa³, A Saina³, F Salesa Greus³, D F E Samtleben^{15,32}, M Sanguineti^{18,29}, P Sapienza¹², F Schüssler³⁸, J Seneca¹⁵, M Spurio^{10,11} , Th Stolarczyk³⁸, M Taiuti^{18,29}, Y Tayalati¹³, B Vallage^{7,38}, G Vannoye⁷, V Van Elewyck^{7,42}, S Viola¹², D Vivolo^{39,43}, J Wilms³⁴, S Zavatarelli¹⁸, A Zegarelli^{19,20}, J D Zornoza³ and J Zúñiga³

¹ Université de Strasbourg, CNRS, IPHC UMR 7178, F-67000 Strasbourg, France

² Université de Haute Alsace, F-68100 Mulhouse, France

³ IFIC—Instituto de Física Corpuscular (CSIC—Universitat de València) c/ Catedrático José Beltrán, 2 E-46980 Paterna, Valencia, Spain

⁴ Technical University of Catalonia, Laboratory of Applied Bioacoustics, Rambla Exposició, 08800 Vilanova i la Geltrú, Barcelona, Spain

⁵ Institut d'Investigació per a la Gestió Integrada de les Zones Costaneres (IGIC) - Universitat Politècnica de València, C/Paranimf 1, 46730 Gandia, Spain

⁶ Aix Marseille Univ, CNRS/IN2P3, CPPM, Marseille, France

⁷ Université Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France

⁸ Aix Marseille Univ, CNRS, CNES, LAM, Marseille, France

⁹ National Center for Energy Sciences and Nuclear Techniques, B.P.1382, R. P.10001 Rabat, Morocco

¹⁰ INFN—Sezione di Bologna, Viale Berti-Pichat 6/2, 40127 Bologna, Italy

¹¹ Dipartimento di Fisica e Astronomia dell'Università di Bologna, Viale Berti-Pichat 6/2, 40127 Bologna, Italy

¹² INFN—Laboratori Nazionali del Sud (LNS), Via S. Sofia 62, 95123 Catania, Italy

¹³ University Mohammed V in Rabat, Faculty of Sciences, 4 av. Ibn Battouta, B.P. 1014, R.P. 10000 Rabat, Morocco

¹⁴ University Mohammed I, Laboratory of Physics of Matter and Radiations, B.P.717, Oujda 6000, Morocco

¹⁵ Nikhef, Science Park, Amsterdam, The Netherlands

¹⁶ Institute of Space Science—INFLPR subsidiary, 409 Atomistilor Street, Măgurele, Ilfov 077125 Romania

¹⁷ Universiteit van Amsterdam, Instituut voor Hoge-Energie Fysica, Science Park 105, 1098 XG Amsterdam, The Netherlands

¹⁸ INFN—Sezione di Genova, Via Dodecaneso 33, 16146 Genova, Italy

¹⁹ INFN—Sezione di Roma, P.le Aldo Moro 2, 00185 Roma, Italy

²⁰ Dipartimento di Fisica dell'Università La Sapienza, P.le Aldo Moro 2, 00185 Roma, Italy

²¹ LPHEA, Faculty of Science—Semplali, Cadi Ayyad University, P.O.B. 2390 Marrakech, Morocco

²² INFN—Sezione di Bari, Via E. Orabona 4, 70126 Bari, Italy

²³ Department of Computer Architecture and Technology/CITIC, University of Granada, 18071 Granada, Spain

²⁴ Université Paris-Sud, 91405 Orsay Cedex, France

²⁵ Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen Centre for Astroparticle Physics, Erwin-Rommel-Str. 1, 91058 Erlangen, Germany

²⁶ Dipartimento di Fisica, Università di Salerno e INFN Gruppo Collegato di Salerno, Via Giovanni Paolo II 132, Fisciano 84084 Italy

²⁷ Laboratoire de Physique Corpusculaire, Clermont Université, Université Blaise Pascal, CNRS/IN2P3, BP 10448, F-63000 Clermont-Ferrand, France

²⁸ LIS, UMR Université de Toulon, Aix Marseille Université, CNRS, 83041 Toulon, France

²⁹ Dipartimento di Fisica dell'Università, Via Dodecaneso 33, 16146 Genova, Italy

³⁰ Royal Netherlands Institute for Sea Research (NIOZ), Landsdiep 4, 1797 SZ 't Horntje (Texel), The Netherlands

³¹ Géoazur, UCA, CNRS, IRD, Observatoire de la Côte d'Azur, Sophia Antipolis, France

³² Huygens-Kamerlingh Onnes Laboratorium, Universiteit Leiden, Leiden, The Netherlands

³³ Institut für Theoretische Physik und Astrophysik, Universität Würzburg, Emil-Fischer Str. 31, 97074 Würzburg, Germany

³⁴ Dr Remeis-Sternwarte and ECAP, Friedrich-Alexander-Universität Erlangen-Nürnberg, Sternwartstr. 7, 96049 Bamberg, Germany

³⁵ Mediterranean Institute of Oceanography (MIO), Aix-Marseille University, 13288 Marseille, Cedex 9, France

³⁶ Université du Sud Toulon-Var, CNRS-INSU/IRD UM 110, 83957 La Garde Cedex, France

³⁷ INFN—Sezione di Catania, Via S. Sofia 64, 95123 Catania, Italy

³⁸ IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France

³⁹ INFN—Sezione di Napoli, Via Cintia 80126 Napoli, Italy

⁴⁰ Dipartimento di Fisica dell'Università Federico II di Napoli, Via Cintia 80126, Napoli, Italy

⁴¹ Dpto. de Física Teórica y del Cosmos & C.A.F.P.E., University of Granada, 18071 Granada, Spain

⁴² Institut Universitaire de France, 75005 Paris, France

⁴³ Dipartimento di Matematica e Fisica dell'Università della Campania L. Vanvitelli, Via A. Lincoln, 81100 Caserta, Italy

* Authors to whom any correspondence should be addressed.

E-mail: mardid@fis.upv.es, sardid@upv.es and jgarmen@upvnet.upv.es

Keywords: deep neural network, neutrino event reconstruction, submarine telescope, transfer learning, dimensionality reduction, event classification, multimessenger astrophysics

Supplementary material for this article is available [online](#)

Abstract

We present the N -fit algorithm designed to improve the reconstruction of neutrino events detected by a single line of the ANTARES underwater telescope, usually associated with low energy neutrino events (~ 100 GeV). N -Fit is a neural network model that relies on deep learning and combines several advanced techniques in machine learning—deep convolutional layers, mixture density output layers, and transfer learning (TL). This framework divides the reconstruction process into two dedicated branches for each neutrino event topology—tracks and showers—composed of sub-models for spatial estimation—direction and position—and energy inference, which later on are combined for event classification. Regarding the direction of single-line (SL) events, the N -Fit algorithm significantly refines the estimation of the zenithal angle, and delivers reliable azimuthal angle predictions that were previously unattainable with traditional χ^2 -fit methods. Improving on energy estimation of SL events is a tall order; N -Fit benefits from TL to efficiently integrate key characteristics, such as the estimation of the closest distance from the event to the detector. N -Fit also takes advantage from TL in event topology classification by freezing convolutional layers of the pretrained branches. Tests on Monte Carlo simulations and data demonstrate a significant reduction in mean and median absolute errors across all reconstructed parameters. The improvements achieved by N -Fit highlight its potential for advancing multimessenger astrophysics and enhancing our ability to probe fundamental physics beyond the Standard Model using SL events from ANTARES data.

1. Introduction

Neutrino telescopes, such as the ANTARES deep-sea detector [1], are pivotal in advancing our understanding of high-energy astrophysical neutrinos, offering insights into cosmic sources and fundamental physics beyond the Standard Model. Accurate reconstruction of neutrino-induced events is essential for extracting meaningful physical information from the collected data. Hence, the inherent complexity of neutrino interactions, coupled with challenges posed by high-dimensional, limited information and correlated detector signals, can benefit from the development of advanced reconstruction techniques that surpass traditional methods. In recent years, machine learning (ML), particularly deep learning (DL), has emerged as a transformative tool in particle physics and astrophysics. The capacity of these techniques to process large volumes of data and extract intricate patterns has proven invaluable in neutrino experiments. For instance, the NOvA experiment implemented convolutional neural networks (CNNs) for event classification, achieving substantial improvements in neutrino identification accuracy [2]. Similarly, the IceCube and KM3NeT neutrino telescopes applied DL methods to enhance the reconstruction of neutrino direction and energy, resulting in improved angular resolution and energy precision [3–5].

The integration of ML into neutrino physics is part of a broader trend where Artificial Intelligence (AI) accelerates scientific discovery across various domains. Notably, AI models based on artificial neural networks leveraging DL have been applied to physical systems reaching remarkable accuracy, enabling

advancements in fields such as weather forecasting and climate modeling [6]. In fundamental physics, ML has been instrumental in detecting rare events, such as observations of the Higgs boson using boosted decision trees [7], and assessing potential cosmic phenomena, thereby expanding our understanding of the Universe [8]. Conventional reconstruction algorithms used by the ANTARES experiment primarily rely on traditional techniques, such as χ^2 regression [9, 10] and maximum likelihood approaches [11, 12]. While effective for certain event types, these methods exhibit limitations in capturing the full complexity of event signatures, particularly in single-line (SL) events, i.e. events reconstructed with data from only one detection line (see section 2). In such cases, reconstructing the azimuthal angle has remained a significant challenge, limiting the precision of directional and energy estimations. To address these issues, we introduce a novel DL-based reconstruction model, *N-Fit*, designed to enhance the accuracy and robustness of neutrino event reconstruction for this type of events.

N-Fit integrates multiple advanced neural network architectures, including deep convolutional networks (DCN) [13], mixture density networks [14], and transfer learning (TL) strategies [15], to effectively process detector data. This modular framework decomposes the reconstruction task into dedicated branches for each neutrino event topology—tracks and showers—with their own networks for directional estimation and energy inference. Both branches are ultimately combined for event classification through TL, ensuring a comprehensive characterization of neutrino interactions. The results of our study, using supervised training through Monte Carlo (MC) simulations and validated by both simulations and real data, demonstrate a significant improvement in reconstruction accuracy compared to conventional methods. These advancements underscore the potential of DL methodologies in multimessenger astrophysics, offering new opportunities to explore the high-energy neutrino sky and investigate phenomena beyond the standard model.

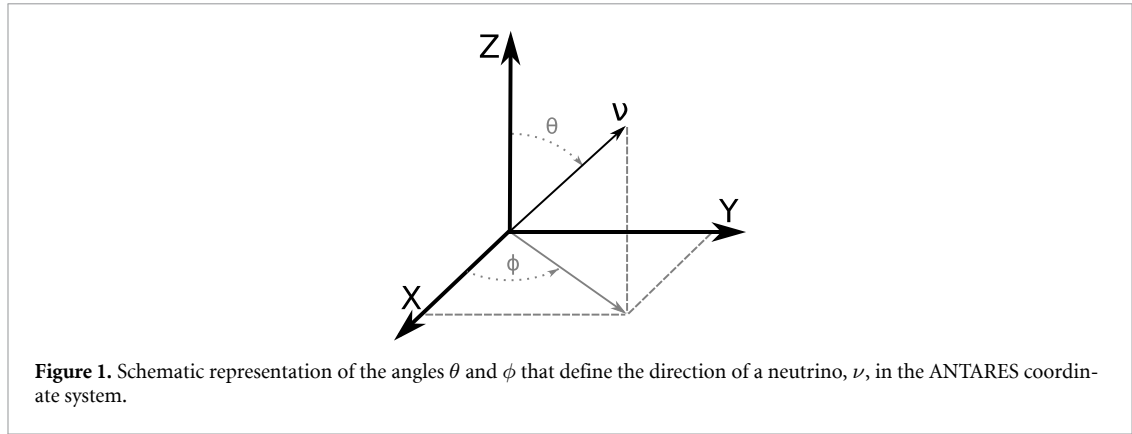
Remarkably, *N-fit* surpasses the SL direction reconstruction over χ^2 -fit in two main aspects; first, by significantly improving zenithal angle estimates; and second, by reporting azimuthal angle estimates, which were impossible to disentangle in χ^2 -fit when all the hits are coplanar. In addition, *N-fit* benefited from knowledge in pre-trained direction reconstruction through *direct* and *indirect* TL, leveraging improvements beyond direction reconstruction. In the next sections, we first describe the ANTARES detector and the characteristics of the dataset used in the study (section 2), then fully describe *N-fit* components in detail (section 3) and its application in ANTARES (section 4). After that, we compare its performance with respect to the standard reference for low-energy SL events in ANTARES (section 5), and illustrate the use of *N-fit* in physics analyses (section 6).

2. ANTARES dataset

ANTARES, located 40 km offshore from Toulon at 2475 m depth, was the first undersea neutrino telescope [1]. It was completed on 29 May, 2008, becoming the largest neutrino telescope in the sea until February 2022, when it was decommissioned and KM3NeT took over. Results obtained by the ANTARES detector are summarized in [16]. ANTARES was made of 12 flexible vertical lines anchored to the seabed. Each line had 25 storeys with a vertical distance of 14.5 m. In each storey there were 3 optical modules (OMs) –consisting of a glass sphere containing a large photomultiplier (PMT)– separated 120° in the horizontal plane and looking downwards at 45° . The storey also contained a titanium cylinder for the electronics and other devices needed for the operation of the telescope, such as positioning and timing calibration systems and connections.

The data acquisition in ANTARES followed an *all-data-to-shore* philosophy. This means that every PMT signal surpassing a predefined voltage threshold is digitized by an analogue ring sampler (ARS) chip offshore and sent to the shore station where they are filtered, recording an event if certain triggers were fulfilled. A *run* was defined as a time period in which the telescope took data with a given configuration, typically covering a few hours. The information for each event included the run number, the position of OMs, their direction, the environmental conditions and the information of the PMTs in terms of hits, i.e. discrete signals that were triggered when the amplitude of the PMT reached the predefined threshold of 0.3 photoelectrons. The hit information was basically the time, voltage amplitude, and the OM identifier.

The information recorded in the runs is used to reconstruct the event parameters, such as the neutrino direction or energy. An event can be classified as a *track* or as a *shower* event, depending on the signature left in the detector. Track-like events originate mainly from charged current (CC) interactions of muon-(anti)neutrinos ($(\bar{\nu}_\mu)^{CC}$) where a muon is produced, which crosses the detector and typically leaves a straight line signature. Shower-like events originate from neutral current (NC) interactions ($(\bar{\nu}_e)^{NC}$ and $(\bar{\nu}_\mu)^{NC}$) and CCs of electron-(anti)neutrinos ($(\bar{\nu}_e)^{CC}$). They can be seen as localized point sources of



light in the detector. Every event is reconstructed as a track and as a shower, and then selected as one or the other for the physics analyses according to reconstruction quality criteria.

Tracks are the most used events in current physics analyses. There are two main methods developed by the ANTARES Collaboration to reconstruct track events. The λ -fit method [11, 12] is based on maximum likelihood estimates and it is particularly suitable for mid- and high-energy events ($\gtrsim 150$ GeV), whereas the χ^2 -fit method is most efficient for lower energies ($\lesssim 150$ GeV). The χ^2 -fit method relies on χ^2 analyses [9], and it allows discrimination between neutrino events that are detected by a SL vs. detection in multiple lines (multi-line). SL events can occur due to several factors, such as the energy of the neutrino, the distance of its interaction from the detector, or its high verticality. SL events for the current study were selected according to χ^2 -fit, i.e. those events reconstructed by χ^2 -fit with only hits from a single line. Shower-like events were reconstructed by a different χ^2 method presented in [10].

The ANTARES Collaboration developed MC simulations of the detector performed *run-by-run*, which are particularly useful to develop and test reconstruction algorithms [17]. These simulations take into account the telescope configuration, and the set of parameters ruling the run condition. Thus, they are well suited for a supervised learning approach. For the development and training of the N -fit framework, a representative subset of approximately 10% of the available MC runs was selected, as detailed in section 4.2.1. We used charged and NC interactions of electron and muon neutrinos for the development of the method. This data was subdivided in training, validation and test datasets in a 3 : 1 : 1 ratio. After the development phase, we also used other events types, such as simulations of atmospheric muons and random noise background events, for control analyses (see section 5.3 and appendix B).

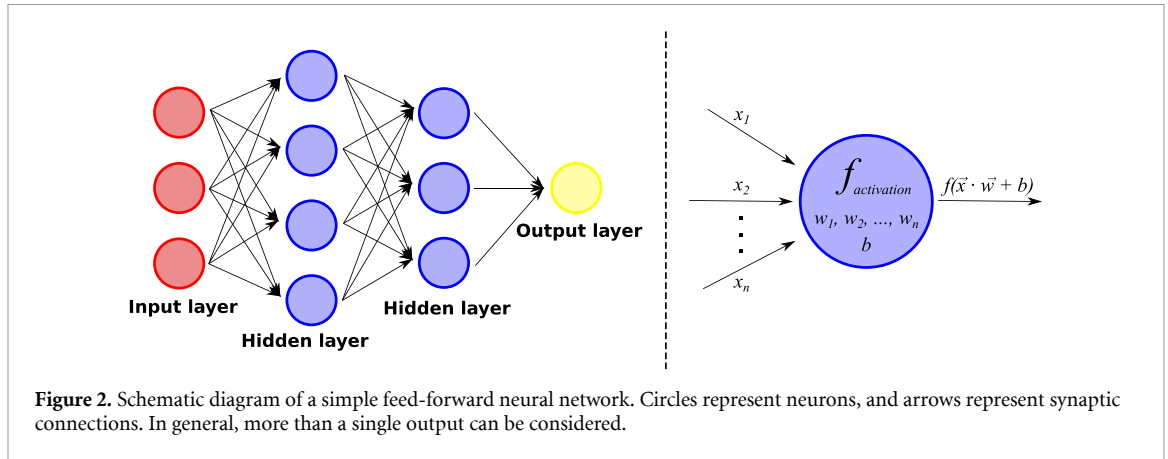
Before proceeding with the N -fit development, we must first clarify the coordinates we use in ANTARES. We describe the spherical coordinates (θ, ϕ) as the zenithal and azimuthal angles, which do not strictly represent the same magnitudes as (Zenith, Azimuth), even though they are strongly correlated. The coordinates (Zenith, Azimuth) represent the direction where the neutrino comes from in the sky, while (θ, ϕ) represent the direction to where neutrino moves to (see figure 1). The transformation between both coordinates is the following:

$$\text{Zenith} = 180^\circ - \theta \quad \text{Azimuth} = \phi + 180^\circ. \quad (1)$$

Note that the range of the Azimuth is $(0^\circ, 360^\circ)$, while the range $(-180^\circ, 180^\circ)$ is used for ϕ . Both of these coordinates are in the universal Transverse Mercator (UTM) system. In this system, the X coordinate points to the UTM East, while the Y coordinate points to the UTM North, which means that the Azimuth runs anti-clockwise.

3. Components of N -fit

The key feature for the successful performance of the N -fit algorithm is the integration of several DL tools. In this section, we unravel the algorithm: first, we introduce the very basics of artificial neural networks, on which the algorithm is based; then, we describe DCNs and MDNs, which play an essential role in N -fit's capabilities. Lastly, we describe how we utilized different aspects of TL to enhance challenging reconstruction analyses, such as energy estimation and event classification.



3.1. Basics of artificial neural networks

Neurons are the core units of neural network models. The basic neuron model in artificial neural networks is the perceptron [18], which is inspired by the non-linear transduction of synaptic input summation in biological neurons towards action potential firing. Mathematically, the output (y) is described as the result of a nonlinear *activation* function of a weighted linear combination of synaptic inputs:

$$y = f\left(\sum_i w_i \cdot x_i + b\right) = f(\vec{w} \cdot \vec{x} + b), \quad (2)$$

where \vec{x} , \vec{w} , and b represent inputs, weights, and the neuron bias, respectively. Even though this represents a strong oversimplification of the nonlinear dynamics of biological neurons in the brain, such a computation is capable of mapping arbitrary input-output functions efficiently if multiple perceptrons are present in layers, creating a *feed-forward* network (figure 2). In N -fit, the rectified linear unit [19]—defined as $\text{ReLU}(x) = \max\{0, x\}$ —is used as the activation function of neurons in the hidden layers.

Each layer in the network model processes a level of internal representation that transmits information from one layer to the next. One of the main characteristics of deep vs. shallow learning is precisely its power of abstraction, which is achieved by significantly increasing the number of layers. Data fixes the number of neurons in the input, whereas the number of outputs is concomitant to the question that the model is designed to address. In contrast, the number of hidden layers and the number of neurons in them can vary. These are examples of hyper-parameters that need to be explored while testing the network performance. Similarly, the learning process is also sensitive to the network initialization [20]. We initialized weights and biases randomly following [21] for all the feed-forward layers in the N -fit algorithm.

During training in supervised ML, the loss function computes the error committed between the network output and the desired output in data batches (of 64 events in N -fit). Weights and biases are updated following error backpropagation, which uses gradient descent to propagate the output loss backwards through the network from the output layer, through the hidden layers, to the input layer. While there are several optimization algorithms based on error backpropagation [22], N -fit uses the Adam algorithm [23] because it dynamically self-regulates the learning rate, which controls the pace at which learnable parameters in the network adapt. The learning rate was initialized at 0.001 in every N -fit training process.

Typically, the learning procedure iterates over multiple epochs until a predetermined number of cycles is completed, where an epoch constitutes one complete pass of the entire training dataset through the neural network. However, an optimization strategy called Early Stopping enables premature termination of training before reaching this preset epoch threshold in order to diminish the risk of over-fitting [24]. This method monitors the loss of the validation dataset after each epoch and halts the process when two conditions are met: the validation loss reaches a minimum value, and a specified number of subsequent epochs (termed the *patience* parameter) elapse without further reduction in validation loss. Upon stopping, the algorithm automatically restores the model parameters corresponding to the optimal validation performance observed during training. We set the maximum number of epochs to 150 and the patience to 10 epochs.

Another technique employed in order to avoid overfitting is called DropOut [25]. It consists of randomly switching off a preset percentage of random neurons within a layer so that the active neurons

vary in each batch. In this way, DropOut prevents the neurons from becoming too specialized. In the N -fit algorithm, 20% of the neurons from the convolutional block outputs were randomly dropped (i.e. their activation was set to 0) before feeding the first feed-forward layer.

3.2. Deep convolutional networks

CNNs are a class of neural network models designed for processing structured grid data, such as images and image-like tensors [13]. Inspired by the human visual system, they are built to recognize patterns and spatial hierarchies efficiently. Unlike traditional fully connected neural networks such as feed-forward ones, which treat every input as independent, CNNs leverage local connectivity, weight sharing, and hierarchical feature extraction to reduce the computational complexity while enhancing learning efficiency. At the heart of a CNN are convolutional layers, which are composed of small, trainable filters that slide over the input image to extract meaningful patterns. These filters evolve over time through training, adapting to detect increasingly sophisticated features in deeper layers. Initially, filters may identify simple edges or gradients, but as they propagate through deeper layers, they start recognizing textures and shapes. The weight updates are governed by backpropagation and gradient descent, ensuring that the filters become more attuned to distinguishing relevant features from noise. After convolutional operations, the activation function introduces non-linearity to the network. In N -fit convolutional layers, ReLU is used, as well as in the fully connected layers. Without the activation function, CNNs would merely perform linear transformations, limiting their ability to capture complex representations.

Pooling layers follow the convolutional layers to refine the feature maps by reducing their spatial dimensions. These layers do not have trainable weights but instead apply fixed operations such as max pooling or average pooling. We used max pooling, the most common method, that retains the most prominent feature in a given region, preserving the strongest activations while discarding less relevant information. This downsampling process helps mitigate overfitting by reducing the number of parameters, thereby improving generalization. By compressing the feature maps, pooling layers also allow subsequent convolutional layers to focus on more abstract features without excessive computational overhead.

DCNs, meaning Deep CNNs, are the main tools used in the development of N -fit. DCNs extend traditional CNNs by stacking multiple convolutional, activation, and pooling layers to extract increasingly abstract and hierarchical features from input data. We processed ANTARES data generating image-like tensors that serve as inputs (see section 4.1). Weights initialization followed [26] for every convolutional layer in N -fit.

While basic CNNs can capture low-level patterns such as edges and textures, DCNs leverage deeper architectures to detect complex structures, object parts and entire objects with high accuracy. As the depth increases, feature maps transition from simple to highly abstract representations, enabling superior generalization and performance in tasks like image recognition and object detection. Training DCNs introduces challenges such as vanishing gradients and computational inefficiency. They are addressed through techniques like batch normalization, which we considered in N -fit. By deepening the network, DCNs significantly enhance the ability to learn intricate patterns, making them the backbone of state-of-the-art DL applications. Lastly, the extracted features from the convolutional layers are fed into fully connected layers, such as those in the feed-forward network explained earlier. At this stage, the network shifts from learning spatial hierarchies to making final predictions.

3.3. MDNs

MDNs extend traditional regression neural networks by predicting complex, multimodal probability distributions instead of deterministic single-point estimates [14], making them ideal for estimating uncertainty in data. A measure of uncertainty is essential for selecting the most reliable estimates in posterior physics analyses.

The outputs of a MDN model represent the necessary parameters to create the multimodal probability density function (PDF) from Gaussian kernels. Thus, instead of the typical output \vec{y}_p for a single multidimensional prediction, the outputs of the network consist of means $\vec{\mu}_i$, variances σ_i and mixture weights α_i (assuming that the outputs are statistically independent):

$$\text{PDF}(\vec{y}_p) = \sum_i^n \alpha_i \cdot N(\vec{\mu}_i, \sigma_i; \vec{y}_p), \quad \text{such that} \quad \sum_i^n \alpha_i = 1, \quad (3)$$

where $N(\vec{\mu}, \sigma; \vec{y}_p)$ is a multi-dimensional isotropic Gaussian function. The number of Gaussian distributions (n) must be hand fixed and predicted variances must always be positive, using the appropriate activation function N -fit uses $\text{ELU} + 1$ (ELU being the Exponential Linear Unit function). Similarly, the

mixture weights use a softmax function to ensure they add to one. The network model learns to predict the parameters of these Gaussians—means, variances, and mixture weights—using a specialized loss function based on maximum likelihood estimation, where the network is optimized to maximize the probability of the observed data under the predicted mixture model (4):

$$\mathcal{L} = -\ln \left\{ \sum_i^n \frac{\alpha_i}{(2\pi)^{c/2} \sigma_i^c} \exp \left(-\frac{\|\vec{y}_t - \vec{\mu}_i\|^2}{2\sigma_i^2} \right) \right\}, \quad (4)$$

where c is the dimension of the output vector \vec{y}_t . In the development of N -fit, each reconstructed parameter follows a single Gaussian mode ($n=1$). As demonstrated by the signed error distributions in figure 10, these errors are symmetric and centered at zero, confirming that a single-mode MDN output effectively characterizes the reconstruction uncertainty. With $c=1$ and $n=1$, the loss function is reduced to:

$$\mathcal{L} = \ln(\sqrt{2\pi}\sigma) + \left(\frac{(y_t - \mu)^2}{2\sigma^2} \right). \quad (5)$$

This is different from traditional loss functions like mean-squared error, as it allows the model to learn distributions that can accommodate multiple correct answers, such as in inverse problems, rather than forcing a single prediction. However, training MDNs can be challenging due to numerical instability, and difficulties in optimizing multiple parameters simultaneously. These effects can be mitigated using regularization techniques, careful initialization strategies, and robust optimization methods like Adam's. In addition to careful initialization and Adam's optimization, N -fit avoided numerical instability issues by adding a small contribution to σ 's activation function: $\text{ELU} + 1 + \epsilon$. The value $\epsilon = 10^{-5}$ was selected after a quick exploration, in which a balance was achieved between being small enough not to alter the results, but large enough to avoid instability.

By integrating neural networks with probabilistic mixture models, MDNs bridge the gap between DL and statistical modeling.

3.4. TL

TL refers to the technique of leveraging knowledge gained from previously trained neural network models to improve performance on related tasks [15]. In the context of N -fit, we adopt two distinct paradigms of TL: *direct TL*, based on model reuse through layer freezing, and *indirect TL*, relying on knowledge distillation via dimensionality reduction.

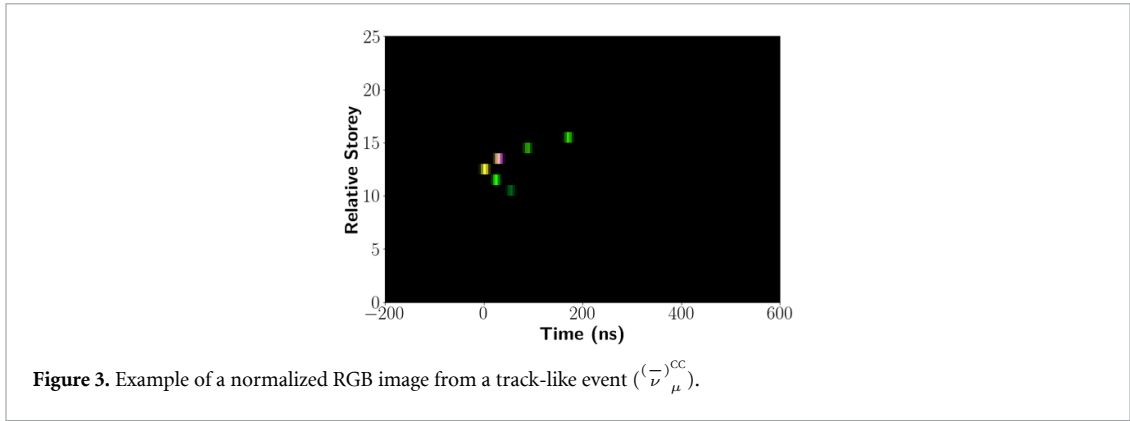
In direct TL, convolutional blocks from pretrained models are reused as fixed feature extractors in subsequent models. Specifically, networks trained for spatial reconstruction are repurposed for classification in N -fit between track and shower neutrino events by freezing their convolutional layers. These frozen components are connected in parallel to a feed-forward neural network (FFN), effectively enriching the classifier input space with spatially encoded representations without requiring retraining of the feature extractors. This strategy benefits from the spatial specialization of each pretrained model and mitigates overfitting by reducing the number of trainable parameters for classification.

Indirect TL in N -fit is achieved through knowledge distillation [27] from pretrained networks by extracting internal activations and transforming them into compact, informative representations using dimensionality reduction techniques. Specifically, this is accomplished using principal component analysis (PCA), a linear dimensionality reduction technique that identifies orthogonal directions (principal components) along which data variability is maximized [28].

PCA is applied to neuron activations from the hidden layers of the networks trained for spatial reconstruction in N -fit, which are assumed to encode latent features relevant for energy estimation. These activations form a high-dimensional feature space, potentially redundant or noisy. With the PCA, we can reduce this space by projecting the original activations onto a lower-dimensional subspace defined by the principal components. Components are ordered by their explained variance, and only those contributing significantly to the total variance are retained. This selection ensures that the most informative and least correlated features are preserved.

The resulting low-dimensional vectors serve as inputs to a downstream FFN tasked with energy regression. This setup effectively transfers high-level abstract knowledge learned during direction and distance reconstruction, enabling the energy model to exploit interdependencies among event features not explicitly provided by the original inputs. More generally, this PCA-based approach enables a form of model-agnostic knowledge transfer, where the distilled representations serve as a bridge between pretrained networks and the target model.

Specific details on how TL is applied to every reconstruction task are left to section 4.2.



4. Application to the ANTARES dataset

4.1. Dataset preprocessing

The first step to apply N -fit to SL neutrino events was to preprocess the available data as image-like tensors. In this view, the ANTARES telescope is akin to a camera collecting 3D images at each time step, having as many pixels as OMs. For SL events, we created colored 2D images, with time in one dimension (X -axis) and the storey of the line in the other (Y -axis). The time stamp of each event was relative to the reference hit of the χ^2 -fit strategy, which is the one that occurred first in time. Next, all hits recorded by the OMs on the line covering a time window of $[-200, 600]$ ns relative to the reference hit are selected. Cleaning noisy or background hits is left to the networks themselves. Colors result from the superposition of the signals provided by the three OMs in each storey, which are evenly distributed in the horizontal plane. In each pixel of the 2D image, we inserted RGB channels to be informative of the event reconstruction in the horizontal (XY) plane, helping to estimate the azimuthal angle (ϕ). Specifically, for each OM i in a storey, we weighted its orientation angle α_i in the horizontal plane by its recorded voltage amplitude A_i . The final color is obtained as the sum of these three individual contributions. Mathematically, the transformation is represented as:

$$(R; G; B) = \sum_{i=1}^3 \begin{cases} A_i \cdot \left(1 - \frac{\alpha_i}{2\pi/3}; \frac{\alpha_i}{2\pi/3}; 0\right) & \text{if } \alpha_i \in \left[0, \frac{2\pi}{3}\right] \\ A_i \cdot \left(0; 2 - \frac{\alpha_i}{2\pi/3}; \frac{\alpha_i}{2\pi/3} - 1\right) & \text{if } \alpha_i \in \left[\frac{2\pi}{3}, \frac{4\pi}{3}\right] \\ A_i \cdot \left(\alpha_i \cdot \frac{3}{2\pi} - 2; 0; 3 - \alpha_i \cdot \frac{3}{2\pi}\right) & \text{if } \alpha_i \in \left[\frac{4\pi}{3}, \frac{6\pi}{3}\right] \end{cases} \quad (6)$$

From the discrete hits, we built a continuous signal representing the hit amplitudes over time by applying a Gaussian kernel smoothing with $\sigma = 5$ ns, and then discretized this signal at regular intervals of 5 ns for each individual color: $\{R(t); G(t); B(t)\}$. Finally, we re-centred the image based on the floor of the reference hit (a representative example is shown in figure 3).

The collected images of SL events were separated into track-like $(\bar{\nu}_{\mu}^{\text{CC}})$ and shower-like $(\bar{\nu}_{\mu}^{\text{NC}}, \bar{\nu}_{e}^{\text{CC}}, \bar{\nu}_{e}^{\text{NC}})$ MC simulations datasets, in order to train specialized models for each type of event. Before feeding images to the network, a Z -score normalization was separately applied to each set. For this, we computed the mean and standard deviation of the respective train sub-sets and applied the transformation to all sub-sets (including real data) to control for distribution shifts.

4.2. N -fit modular organization

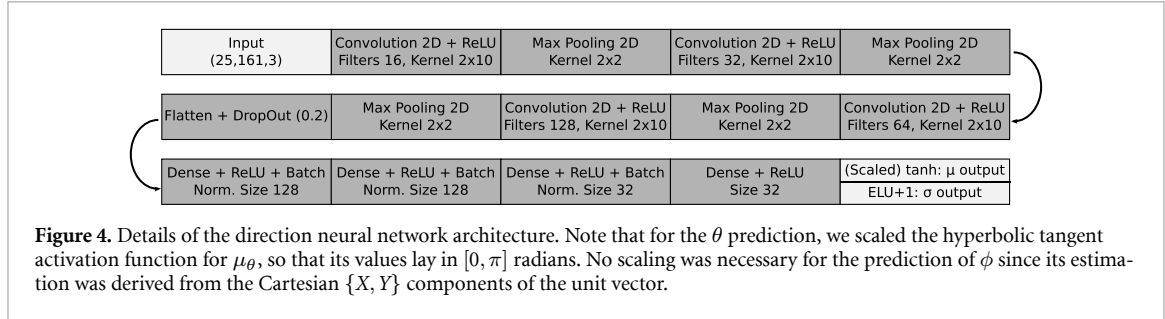
The N -fit reconstruction strategy is subdivided into two branches specialized in each type of event: tracks and showers. Each branch covers the different aspects that fully describe neutrino events. The track-branch reconstructs the neutrino direction, the closest point of the secondary muon track to the detector line and the muon energy. The shower-branch also reconstructs the direction of the neutrino, the vertex point of the neutrino interaction and its energy. The two branches converge through TL in the classifier that divides the events in tracks or showers. The next subsections explain the details of these network models.

4.2.1. Direction reconstruction

Several architectures of increasing complexity were analyzed during the optimization of direction reconstruction models for track events. In all convolutional layers, a stride of 1 and ‘same’ padding were used

Table 1. Evolution of the MAE of the test data set along the key steps which were followed in the development and optimization of the N -fit algorithm for the direction reconstruction of tracks. The estimation of σ was introduced in step (iii).

MAE	<i>i</i>	<i>ii</i>	<i>iii</i>	<i>iv</i>	<i>v</i>	<i>vi</i>
θ	11.2°	10.5°	10.5°	9.6°	8.4°	7.4°
ϕ	50.1°	49.5°	49.5°	46.2°	44.1°	41.4°



to maintain the dimensionality of the feature maps before the downsampling layers. The key steps which were followed to measure and improve the model performance are summarized below:

- i*) Baseline control model: Feed-forward network of four hidden layers implemented to fully reconstruct the Cartesian components of the direction unit vector: $\{X; Y; Z\}$.
- ii*) Components separation: Next, the model was divided to reconstruct Z and $\{X, Y\}$ separately. The reconstructions were sequential: first Z , then $\{X, Y\}$. We regularized the value of the components $\{X, Y\}$ to penalize deviations of $\{X; Y; Z\}$ from unit vectors.
- iii*) MDNs and θ representation: Point predictions were replaced by inferring probability distributions (MDN). The change did not affect the performance but provided event uncertainty estimation, which is of major relevance as a quality metric in posterior physics analyses. Also, the network output Z was replaced by the angle θ , gaining a best estimation of the angle uncertainty, since no error propagation was needed to infer it.
- iv*) Convolutional layers: A significant improvement was observed by adding two convolutional layers before the dense network.
- v*) Image centering: Aligning events to a centred reference frame further enhanced performance, despite the architecture remained unaltered.
- vi*) Final adjustments: The final architecture was selected through a heuristic-guided manual tuning process. During this optimization, we explored a range of hyperparameters, including the number of convolutional blocks (0 to 4), filter depths (16 to 128), and dense layer sizes (32 to 256 neurons). We also experimented with various kernel shapes, ultimately selecting the 2×10 rectangular kernel to better capture the spatio-temporal correlations in the SL images. The final configuration, as detailed in figure 4, was chosen based on the saturation of performance gains (see table 1). From there on, additional layers provided only marginal changes in performance.

The progression of the Mean Absolute Error (MAE) through these steps is presented in table 1.

The final architecture is illustrated in figure 4, which relies on rectangular 2×10 kernels. This specific geometry is designed to capture the unique spatio-temporal signature of SL events: the 10-step temporal width covers a 50 ns window (given the 5 ns discretization), which is optimal for detecting the timing correlations of the hits, while the 2-storey vertical height extracts local spatial patterns between adjacent storeys. This allows the network to effectively learn the characteristic patterns of Cherenkov light propagation along the detector line.

The loss function applied to the angle θ becomes:

$$\mathcal{L} = \ln\left(\sqrt{2\pi}\sigma_\theta\right) + \frac{1}{2} \frac{(\theta_t - \mu_\theta)^2}{\sigma_\theta^2}, \quad (7)$$

where θ_t represents the true value of the angle θ for each neutrino event in the simulation.

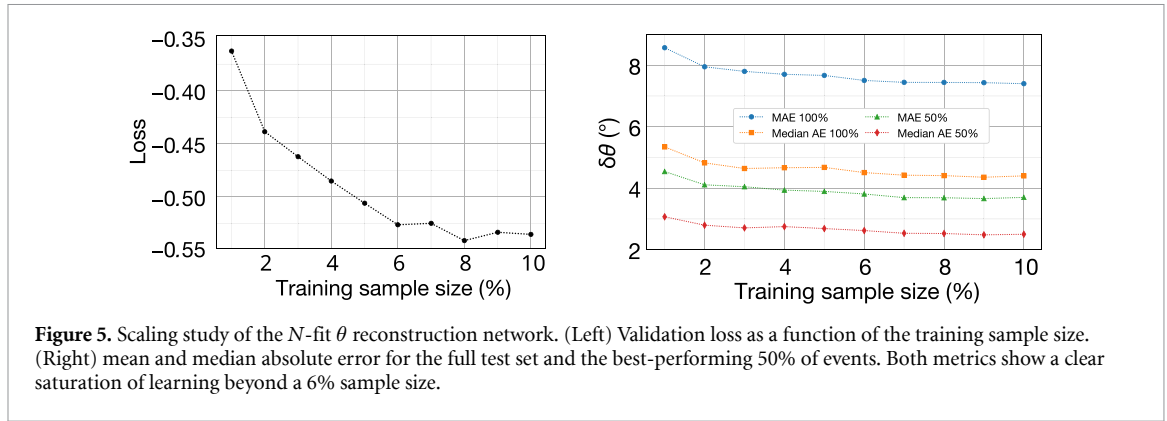


Figure 5. Scaling study of the N -fit θ reconstruction network. (Left) Validation loss as a function of the training sample size. (Right) mean and median absolute error for the full test set and the best-performing 50% of events. Both metrics show a clear saturation of learning beyond a 6% sample size.

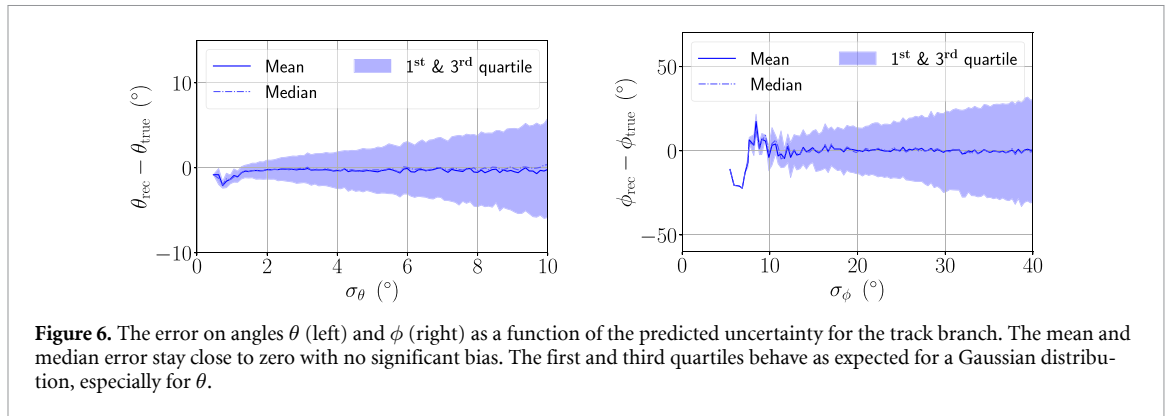


Figure 6. The error on angles θ (left) and ϕ (right) as a function of the predicted uncertainty for the track branch. The mean and median error stay close to zero with no significant bias. The first and third quartiles behave as expected for a Gaussian distribution, especially for θ .

To ensure that the training dataset size is sufficient for the model to reach its full potential, a scaling study was performed using the θ reconstruction network. The model was trained with increasing sample sizes ranging from 1% to 10% of the available data (where 10% corresponds to approximately 1250 000 events). As illustrated in figure 5, both the validation loss and the reconstruction error reach a clear plateau beyond 6% of the sample. The improvement in precision beyond 6% is marginal, indicating that the model's learning has effectively saturated and that the selected subset provides an optimal balance between performance and computational efficiency.

A second network predicts the Cartesian components (μ_X, μ_Y) and their uncertainties (σ_X, σ_Y) to ensure a robust Gaussian error distribution and avoid periodicity issues during training. The loss function was transformed consequently as:

$$\mathcal{L} = \ln(\sqrt{2\pi}\sigma_X) + \frac{1}{2} \frac{(X_t - \mu_X)^2}{\sigma_X^2} + \ln(\sqrt{2\pi}\sigma_Y) + \frac{1}{2} \frac{(Y_t - \mu_Y)^2}{\sigma_Y^2} + (1 - [\mu_X^2 + \mu_Y^2 + \cos^2(\mu_\theta)])^2. \quad (8)$$

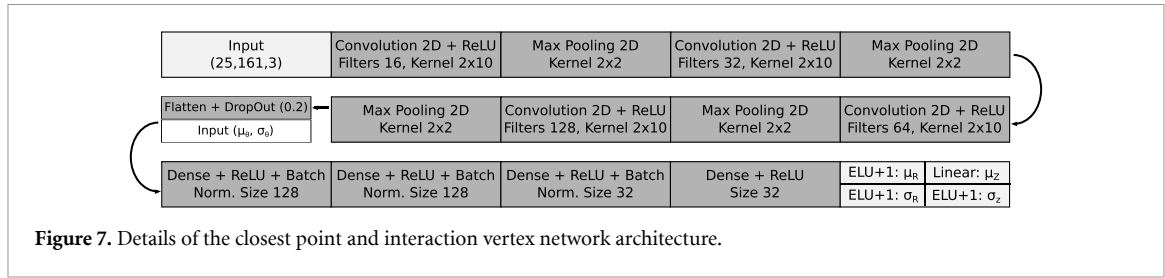
The last term in equation (8) regularizes the output of the network, penalizing deviations of $(X; Y; Z)$ from unit vectors. To infer the uncertainty of ϕ , we performed a quadratic error propagation:

$$\sigma_\phi^2 = \left(\frac{\partial\phi}{\partial\mu_X} \cdot \sigma_X \right)^2 + \left(\frac{\partial\phi}{\partial\mu_Y} \cdot \sigma_Y \right)^2 \Rightarrow \sigma_\phi = \frac{\sqrt{\sigma_X^2 \cdot \mu_Y^2 + \sigma_Y^2 \cdot \mu_X^2}}{\mu_X^2 + \mu_Y^2}. \quad (9)$$

Note that the uncertainty estimation assumes a Gaussian distribution. This means that the distribution of errors, although centred, should become wider with growing σ values. For most of the error ranges this is precisely the case, especially for θ (see figure 6 for track events, and figure 16 in the Supplementary Material for shower events). As for the uncertainty of the direction reconstruction (σ_Ω), we used the following expression derived from the solid angle definition:

$$\sigma_\Omega = \sqrt{\sin^2(\theta) \sigma_\phi^2 + \sigma_\theta^2}. \quad (10)$$

This optimized architecture for the track branch was then applied in the shower branch without further optimization. Further analyses and results can be found in section 5.



4.2.2. Closest point and interaction vertex

In addition to the reconstruction of direction, estimating the energy of neutrino events is fundamental for physics analyses. Energy and distance to the detector are, however, intermingled in SL events: distant high-energy events may appear as SL as well as near low-energy events. Then, to improve the energy estimation, we included in N -fit the reconstruction of the *closest point* of track events to the ANTARES detector line, and the *interaction vertex* position of shower events (figure 7), since these two magnitudes are characterized by the horizontal distance in meters (R_c for closest point of tracks, R_v for interaction vertex of showers) and their vertical position (Z_c for tracks, Z_v for showers) defined in the ANTARES frame of reference.

These spatial magnitudes (R and Z) constitute entirely novel outputs for SL events in the ANTARES framework, as traditional χ^2 -fit methods did not provide reliable estimates for a single detector line. To evaluate the performance of N -fit, we have implemented a benchmark comparison using the χ^2 -fit. As illustrated in the Supplementary Material (figures 18 and 20 for track and shower events, respectively), N -fit significantly outperforms this baseline, achieving improved spatial characterization with median errors well below 5 m, whereas the χ^2 -fit results show broad distributions and high uncertainties. Differences lay in input and output of the networks to fulfil the characteristics of these reconstructions. Thus, events were not centred across images in these network models to ease the reconstruction of Z . In addition, these networks also incorporated the reconstruction of θ as input, which was introduced in parallel to the convolutional layers. Lastly, the loss function was adapted in consonance with the outputs of these networks:

$$\mathcal{L} = \ln\left(\sqrt{2\pi}\sigma_R\right) + \frac{1}{2} \frac{(R_t - \mu_R)^2}{\sigma_R^2} + \ln\left(\sqrt{2\pi}\sigma_Z\right) + \frac{1}{2} \frac{(Z_t - \mu_Z)^2}{\sigma_Z^2}. \quad (11)$$

4.2.3. Energy

As described earlier, estimating the energy is particularly challenging for SL events, especially for track events, given their physical characteristics and the limited information available by a single line of the detector. Shower events are better suited for energy reconstruction because of their physical topology. Moreover, the neutrino energy is directly inferred by N -fit in shower events, whereas for track events, the reconstructed energy by N -fit is limited to that of the secondary muon. This is due to the stochastic energy loss in neutrino interactions producing track events. The neutrino energy of track events for physics analyses can, however, be inferred indirectly, by combining N -fit energy reconstruction and the statistical properties of the interactions.


As a first approximation, we applied the same model architecture of the direction reconstruction to infer energy without further tuning. The energy reconstructions by such model presented very low accuracy. Reasons underlying the poor performance include that the secondary muon can escape the detector in track events, or that events could be very close or far away from the detector line in shower events. To better guide training in the N -fit energy reconstruction, we preselected events that had good direction reconstruction (i.e. 50% of the events with the lowest predicted σ_θ), and that were close to the line (according to the $\{R, Z\}$ reconstruction). The specific cuts appear in expression (12) for tracks, and in expression (13) for showers.

$$R_c \leq 50 \text{ m} \quad Z_c \in (-150, 150) \text{ m} \quad \sigma_\theta \leq 8.6^\circ \quad (12)$$

$$\sigma_{R_c} \leq 4.3 \text{ m} \quad \sigma_{Z_c} \leq 6.2 \text{ m},$$

$$R_v \leq 50 \text{ m} \quad Z_v \in (-150, 150) \text{ m} \quad \sigma_\theta \leq 14.9^\circ \quad (13)$$

$$\sigma_{R_v} \leq 5.6 \text{ m} \quad \sigma_{Z_v} \leq 3.3 \text{ m}.$$



Although each spatial constraint applies a 50% selection, the final cumulative acceptance from the energy network is approximately 40%. This indicates a correlation between the spatial networks; events that are well-reconstructed in direction also tend to yield reliable distance and vertex estimates. We emphasize that this architecture establishes a sequential reconstruction pipeline, where the spatial outputs and their uncertainties are first independently determined and subsequently used as selection filters for the energy module. While this entails a significant reduction in the event sample, this filtering process drastically enhances the quality and reliability of the energy reconstruction, which is relevant for subsequent physics analyses.

After training with these cuts, energy inference slightly improved. We considered this model as the *benchmark* to evaluate further improvements, particularly that from indirect TL through PCA-based knowledge distillation (see section 3.4). In such an approach, we aimed to exploit the relationship between the energy of the neutrino and all other physical parameters of the event processed by N -fit direction and distance models. We assumed that internal representations from those models could then benefit energy predictions. To test the hypothesis, neuron activations in hidden layers from the θ and $\{R, Z\}$ networks were taken as feature dimensions from which to infer the energy of the events. Specifically, we extracted the internal activations from all intermediate layers of these models, including the four convolutional blocks (with integrated pooling) and the four dense hidden layers. This process yielded 28 480 features per spatial network, resulting in a high-dimensional input space of 56 960 dimensions. We disregarded activations from the ϕ network for two reasons. First, the azimuthal information is physically irrelevant for energy estimation due to the approximately cylindrical symmetry of the detector line. Second, tests conducted during development showed that including ϕ activations introduced additional noise into the PCA feature space; specifically, the validation loss was found to be nearly an order of magnitude better (lower) when ϕ activations were excluded (-2.4×10^{-3} vs -3.5×10^{-4}).

We linearly transformed all the selected feature dimensions, using the PCA, to rank components according to their variability. This approach allowed us to compress the initial space of 59 960 dimensions by three orders of magnitude, isolating the most informative latent features while eliminating the vast redundancy inherent in the raw neural activations. We then selected those most relevant features as inputs to a FFN trained to infer the energy as point predictions (figure 8). As illustrated in the architecture diagram, this network features a high-capacity hidden layer of 1280 neurons. This choice is a specific design to re-expand the compressed information from the PCA into a higher-dimensional latent space. By increasing the layer size ten-fold compared to the standard 128-neuron layers used in our spatial networks, we provide the energy model with the representational power needed to model the complex, non-linear relationships of energy loss from a limited set of distilled features. The final number of components, 63 features for the track branch and 68 for the shower branch, was determined applying the elbow rule, i.e. stop incorporating more features once the cumulated variance reached a *plateau*. Concurrently, we realized that inferring $\log(E)$ performed better than estimating E directly, due to the large energy range of neutrino events (from 5 GeV up to 20 TeV).

In a subsequent phase aimed to further improve energy predictions for physics analyses, we designed a specialized FFN by feeding the model during training with only the 50% best reconstructed events from the original model shown in figure 8. In addition, this specialized FFN incorporated a MDN output layer to provide uncertainty estimation of energy predictions, as well as loss weighting factors to balance the impact of the non-uniform energy distribution.

4.2.4. Classification

N -fit includes a neural network for track vs. shower classification, which leverages TL from the specialized track and shower branches. The output is a track classifier score (S), with shower-likeness defined as $1 - S$. To train, validate and test the classifier, we used a combined dataset made of 200 000 MC events, half of them representing track events ($(\bar{\nu}_\mu)^{\text{CC}}$) and the other half representing shower events ($(\bar{\nu}_\mu)^{\text{NC}}$, $(\bar{\nu}_e)^{\text{NC}}$, $(\bar{\nu}_e)^{\text{CC}}$). This new dataset, as well as the previous used in the training phase, are selected according to the

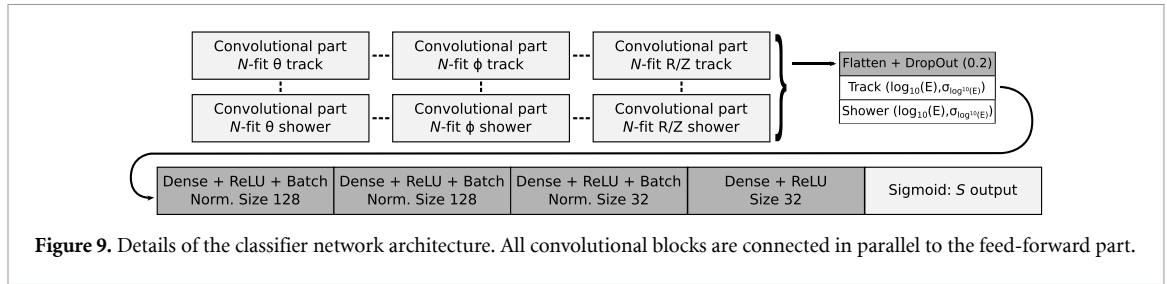


Figure 9. Details of the classifier network architecture. All convolutional blocks are connected in parallel to the feed-forward part.

χ^2 -fit SL criteria with no further refinement. All events were randomly sampled in a uniform manner covering the full period in which ANTARES collected data.

The *control* network model for classification consisted of the same model architecture of the direction reconstruction without further tuning. Only the output was adjusted: a single neuron represented the track classifier score, using the sigmoid as its activation function. Given that track and shower events typically present different trace characteristics and that these traits were presumably present as internal representations in the convolutional layers of *N*-fit direction and distance reconstruction models, we decided to exploit this knowledge in a second, more sophisticated classifier endowed with TL. Specifically, we incorporated the convolutional blocks of the θ , ϕ and $\{R, Z\}$ network models from both *N*-fit branches as frozen components that were plugged in parallel into a FFN. This network model also incorporated the reconstruction of the energy and its uncertainty from the networks specialized in track and shower events. The complete architecture is shown in figure 9.

4.2.5. Code implementation and computational efficiency

All neural network models were implemented using the TensorFlow 2.4.1 framework [29], with model evaluation and analysis conducted in Python 3.9.15. The implementation relied on standard scientific computing libraries, including NumPy 1.23.4, Scikit-learn 1.1.3, and Pandas 1.5.1, along with several auxiliary packages.

Training times varied depending on the network architecture. Specifically, direction and position networks required approximately 12 hours each to train, leveraging GPU acceleration via TensorFlow. The PCA, applied separately to each branch (track and shower), required roughly 15 minutes per branch. The subsequent training of the energy reconstruction networks took between 5 and 15 minutes, being much faster than the spatial models primarily due to their lack of convolutional layers, which are typically the most computationally demanding components. Training the track vs. shower classifier, which reuses frozen convolutional layers transferred from the specialized branches, was comparatively lightweight, requiring only about 20 minutes. GPU-accelerated training was conducted on a workstation running KDE Neon 24.04, equipped with two NVIDIA Quadro RTX 8000 GPUs (TU102GL architecture, 64-bit interface), each of them having a 48 GB GDDR6 memory and supporting 672 GB s⁻¹ memory bandwidth.

The final application of the *N*-fit algorithm to ANTARES data was organized into three main phases: (i) reading the ANTARES data files, (ii) processing the raw data into the standardized *N*-fit input format, and (iii) performing the actual reconstruction. The first phase is common to any reconstruction pipeline and does not affect the evaluation of *N*-fit's computational performance. The image generation stage, where detector hits are transformed into 2D representations for network input, took an average of 0.2 seconds per event. The complete reconstruction process required approximately 0.1 seconds per event. These estimates are based on the averaged processing times over 10 000 events, using conservative rounding to ensure upper-bound accuracy. In both cases, model parameters and dependencies were pre-loaded into memory cache, eliminating the overhead of repeated I/O operations during batch processing.

The whole application of *N*-fit to the ANTARES data was executed on the high-throughput computing (HTC) partition of the IN2P3 Computing Center, managed via the SLURM workload manager. The HTC partition consists of heterogeneous CPU-only nodes running Red Hat Enterprise Linux release 9.6 (64-bit). Most nodes are equipped with AMD EPYC 7302 16-core, EPYC 7453 28-core, or EPYC 9334 32-core processors, with memory per node ranging from 192 GB to over 1.2 TB. For our runs, a single CPU core and less than 2 GB of RAM were sufficient for efficient reconstruction inference. GPU resources were not required for deployment, although training leveraged TensorFlow's GPU acceleration when available.

Table 2. Mean and median absolute error of reconstructed neutrino angles θ , ϕ , and total deviation, Ω , for the track branch. The 50% best events are selected according to the lowest values of the parameters σ for N -fit and χ^2 for the χ^2 -fit.

Mean/Median	χ^2 -fit		N -fit		χ^2 -fit (50%)		N -fit (50%)	
θ	15.5°	8.5°	7.4°	4.4°	9.7°	5.5°	3.7°	2.5°
ϕ	—	—	41.4°	31.5°	—	—	29.2°	23.6°
Ω	—	—	28.3°	22.7°	—	—	18.7°	13.7°

5. Results

This section presents the main performance results of the N -fit algorithm. All evaluations are performed on the corresponding test subsets associated with each reconstruction branch, as defined in section 4.1, except for the classifier performance, which is tested with its own dataset. Finally, additional validation studies—addressing the algorithm robustness with respect to training data variations, resilience to noise, and applicability to real ANTARES data—are briefly presented.

5.1. Direction reconstruction improvement

For the track branch of N -fit, statistics of the error distribution for the reconstructed angles can be seen in table 2. Errors are computed as the absolute values of the difference between the reconstructed (μ_α) and the simulated (α_t) angles for the $(\bar{\nu}_\mu)^{\text{CC}}$ MC simulations. The error distributions are shown in figure 10(a). The N -fit algorithm clearly outperforms the χ^2 -fit SL reconstruction method: results show a significantly larger proportion of low errors for the θ angle and a first estimation of the ϕ angle, which was previously missing for SL events. Remarkably, our approach allows to focus on best predictions by selecting the events with the lowest values of σ provided by N -fit. Applying a selection criterion on σ improves results (see table 2 and figure 10(b)). While a mean ϕ error of $\sim 29^\circ$ for high-quality events may seem large, it represents a significant advance over previous SL reconstruction methods where such information was unavailable. In practical physics analyses, this resolution is sufficient to define asymmetrical Regions of Interest that can reduce background contamination by a factor of 6 while retaining approximately 2/3 of the signal acceptance, as demonstrated in the follow-up study in section 6. These plots for the shower branch—where $(\bar{\nu}_\mu)^{\text{NC}}$, $(\bar{\nu}_e)^{\text{NC}}$ and $(\bar{\nu}_e)^{\text{CC}}$ MC simulations are employed—are shown in the supplementary material, figure 17, as well as the statistics of the error for the reconstructed angles, table 4.

By using expression (10), we can also apply the selection criterion to the total angular deviation for both track and shower branches (figure 11(a)). The direction reconstruction for showers is worse than for tracks due to the physical geometry of the shower signature. It is easier to determine the direction of a source emitting light following a straight path than from an almost point-like source, typical of showers. This issue, however, does not prevent the use of N -fit and the estimated uncertainty can still be used to choose the best reconstructed events.

Finally, we checked that N -fit direction reconstructions show no bias for the whole MC simulated energy range of neutrinos, in contrast to the χ^2 -fit (figure 11(b)). This is particularly relevant for using this method in physics analyses that involve low energy events.

5.2. TL applications

5.2.1. Energy reconstruction

As mentioned in section 4.2.3, the energy was difficult to reconstruct due to its dependency on the distance of the neutrino interaction and its physical properties. Thus, we first reconstructed the closest point of tracks to the detector line and the interaction vertex position of showers. These reconstructions were very satisfactory, as can be seen in figure 12. Supplementary plots regarding these reconstructions can be seen in figures 18–21. After this, the reconstruction of the energy was addressed using TL with the PCA, as explained in section 4.2.3.

Results for the pre-selected test dataset according to expressions (12) and (13) can be seen in figure 13, showing the utility of the $\sigma_{\log_{10}(E)}$ in selecting the best predictions. Compared to the benchmark reconstruction, the results from the track fit show an improvement when using inheritance from previous networks, even if the results are still moderately good due to the physical complications. The results for showers are slightly better because they are better suited for energy reconstructions due to its topology. Compared to the benchmark, we can see that the uncertainty parameter improves considerably in order to select the best reconstructed events.

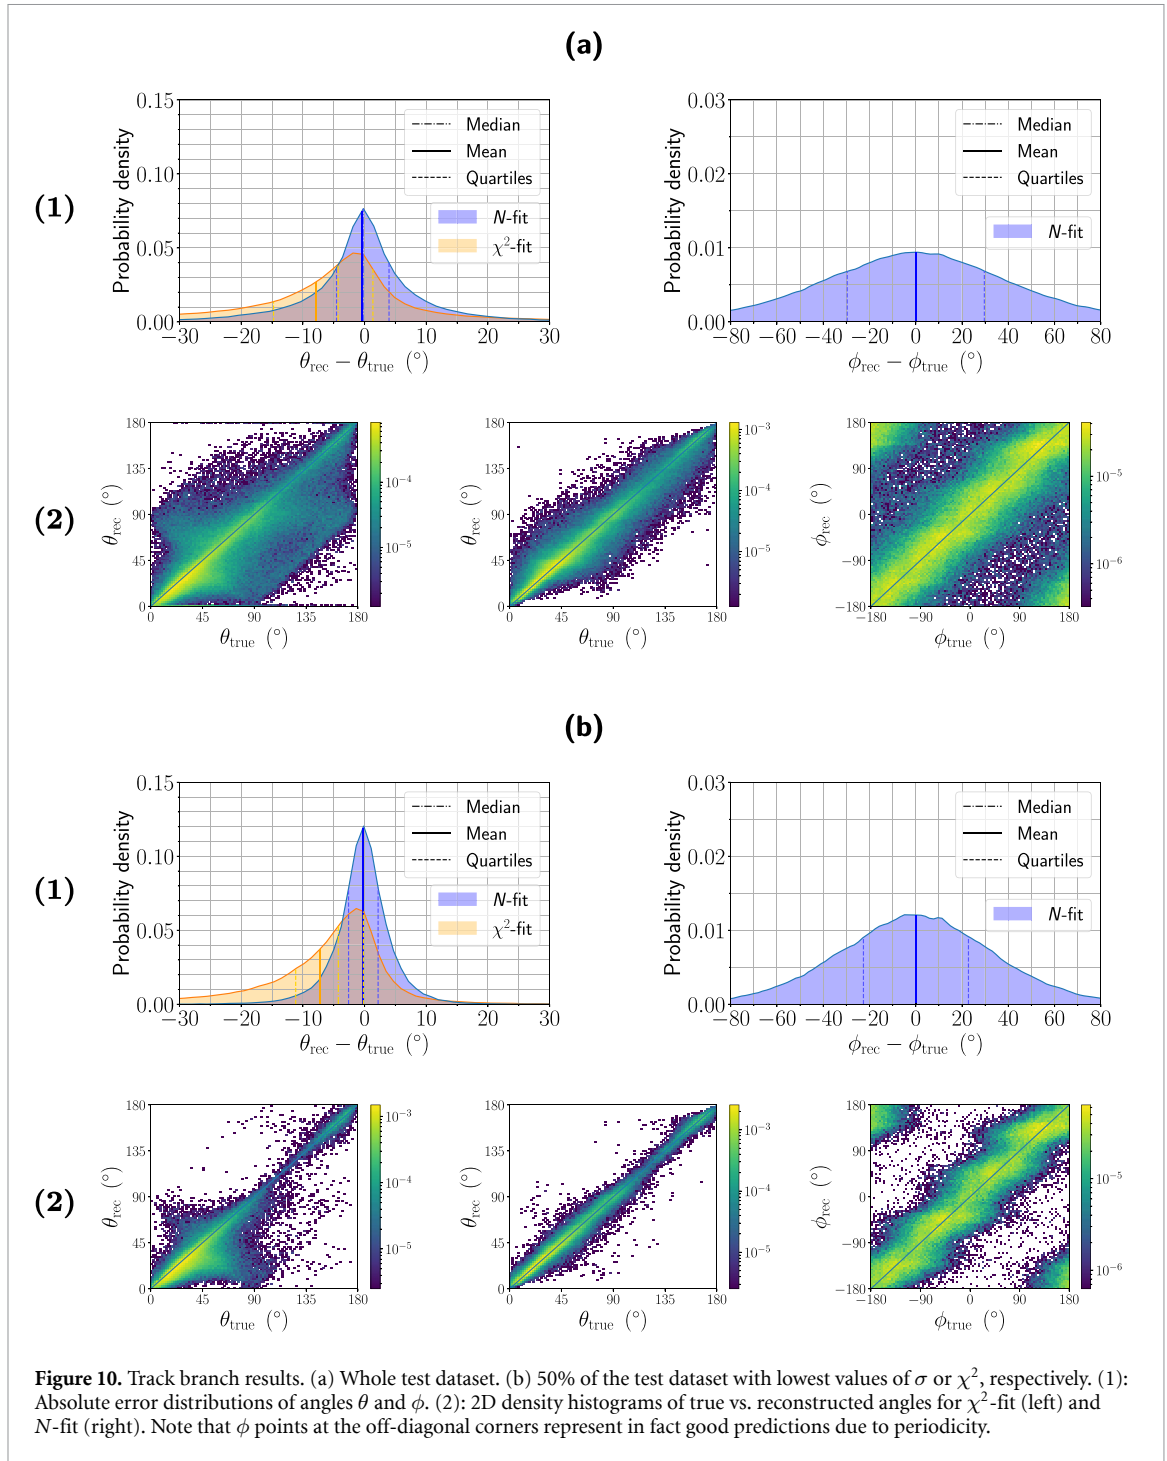


Figure 10. Track branch results. (a) Whole test dataset. (b) 50% of the test dataset with lowest values of σ or χ^2 , respectively. (1): Absolute error distributions of angles θ and ϕ . (2): 2D density histograms of true vs. reconstructed angles for χ^2 -fit (left) and N -fit (right). Note that ϕ points at the off-diagonal corners represent in fact good predictions due to periodicity.

In order to quantify the reconstruction resolution, we have computed the median relative error as a function of the true energy:

$$\left| \frac{E_{\text{rec}} - E_{\text{true}}}{E_{\text{true}}} \right|, \quad (14)$$

for 25% of events with the best energy reconstruction according to the uncertainty prediction. For tracks, this quantity has a stable value of 0.55 between a true muon energy of 20 and 200 GeV, worsening considerably outside this range. For showers, the error is stable for true neutrino energies higher than 100 GeV, with a value of 0.50 and worsening significantly for lower energies.

5.2.2. Classifier performance

To measure the performance of the classifier, we computed the overall accuracy as well as the bias, recall and precision for both types of events on the independent 20% test dataset. These metrics, summarized

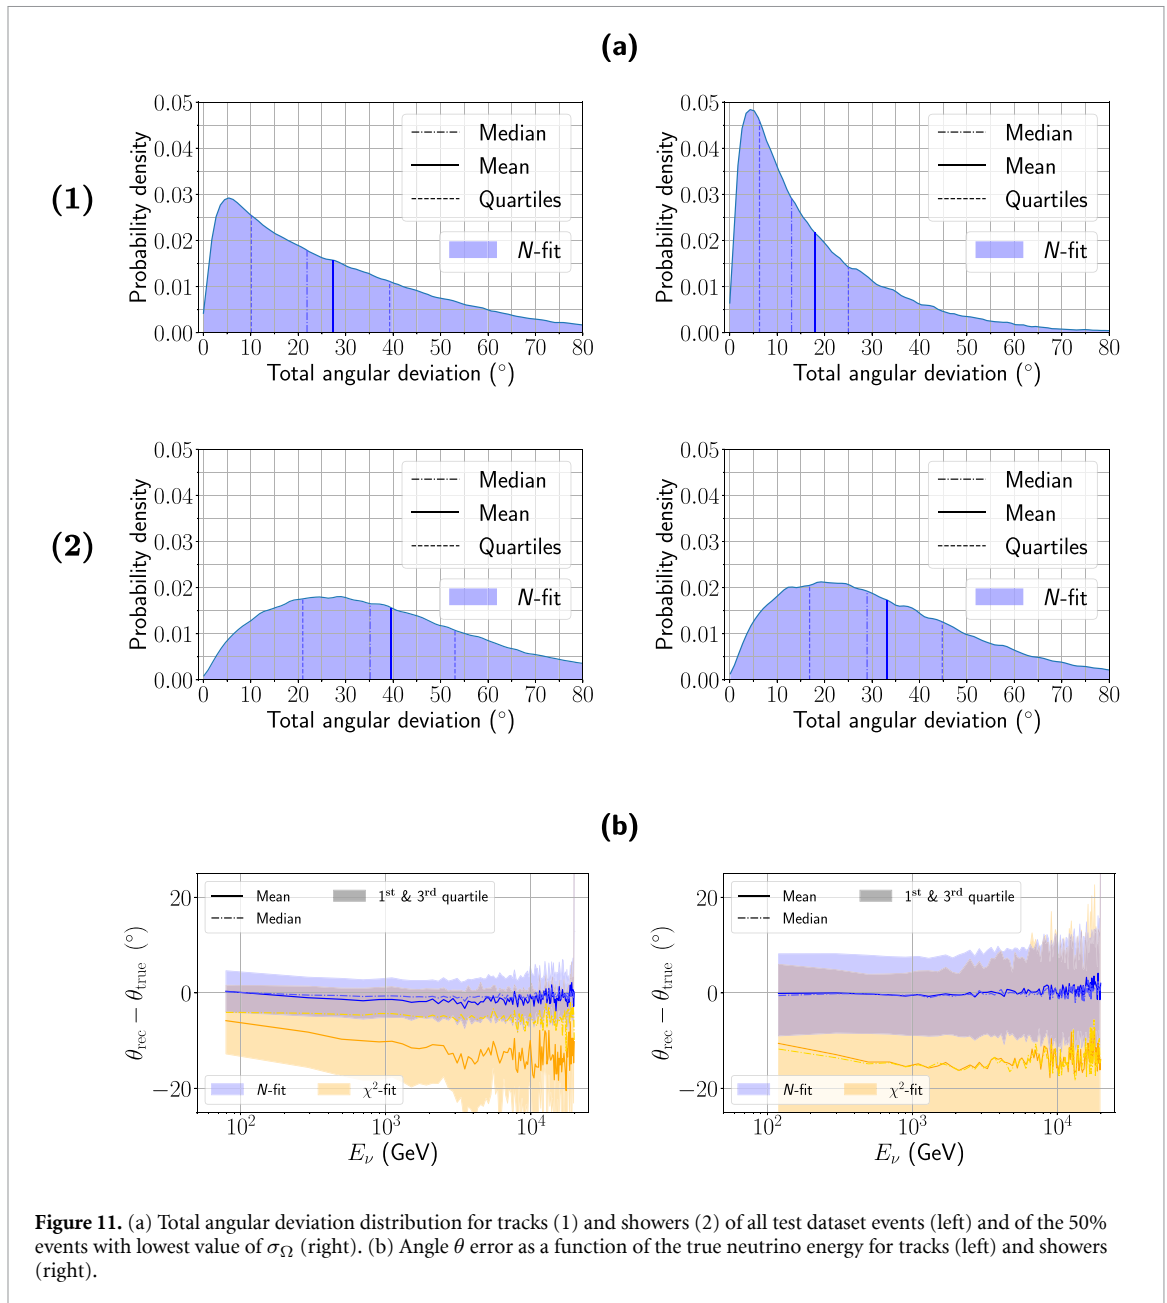


Figure 11. (a) Total angular deviation distribution for tracks (1) and showers (2) of all test dataset events (left) and of the 50% events with lowest value of σ_Ω (right). (b) Angle θ error as a function of the true neutrino energy for tracks (left) and showers (right).

in table 3, provide a comprehensive view of the model's ability to distinguish between track and shower topologies.

It is important to note that the current N -fit classifier is presented as a proof-of-concept. In practical physics analyses, event selection is typically not restricted to a simple 0.5 threshold. Instead, researchers often define high-purity regions (e.g. > 0.75 for tracks and < 0.25 for showers), discarding the intermediate 'inconclusive' band to minimize cross-contamination. This strategy effectively enhances the sample purity beyond the nominal metrics reported here.

To provide a formal benchmark, we performed a classification study using the standard χ^2 -fit values as a selection criterion for SL events, selecting the reconstruction (track or shower) with the lowest χ^2 . As shown in table 3, N -fit significantly outperforms this established baseline, particularly in track recall (75% vs 53%), while maintaining a more balanced track/shower prediction distribution (45/55%) compared to the 40/60% bias observed in the χ^2 -fit.

The final model achieved an accuracy of 80%, with a recall of 75% for tracks and 85% for showers, and a precision of 81% for tracks and 77% for showers. This implies that while 25% of tracks are misclassified as showers, there is an 81% confidence that an event classified as a track is indeed a track interaction.

Figure 14 illustrates the distributions of the track classifier score for track-like and shower-like events, both from the final transfer learning approach and from the benchmark without it ('no TL' tag). There is

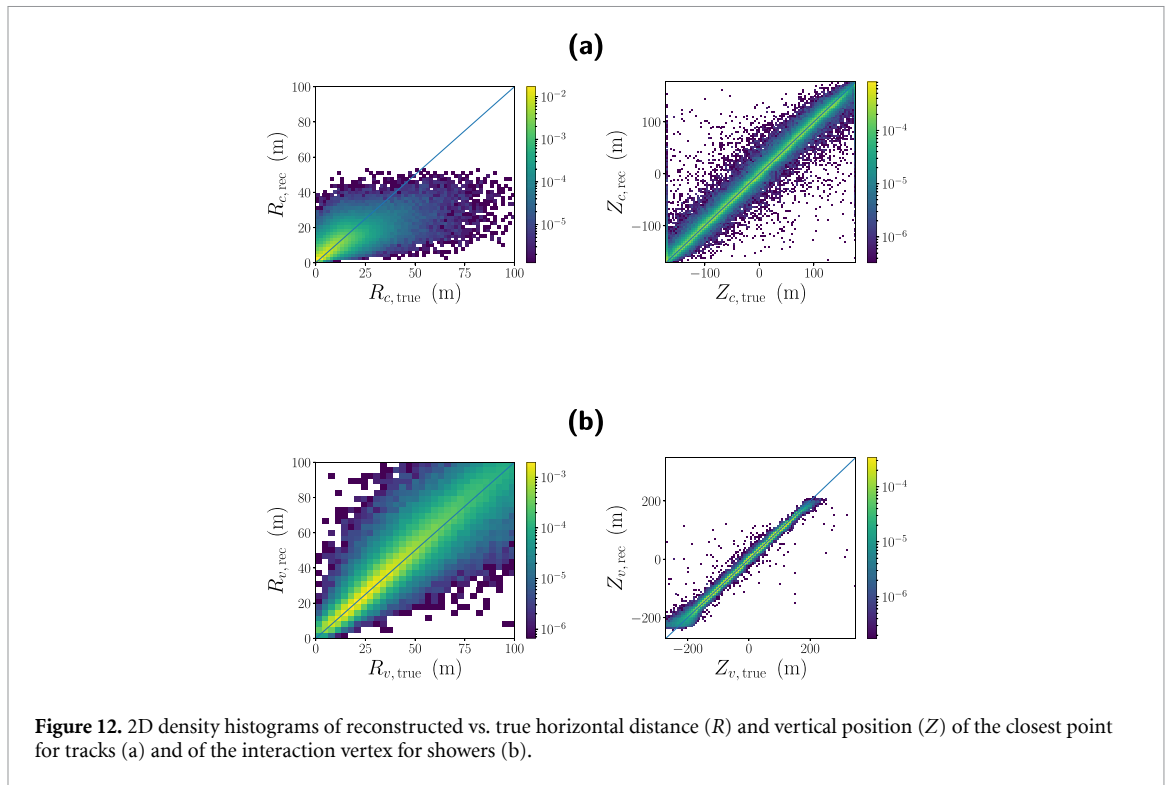


Figure 12. 2D density histograms of reconstructed vs. true horizontal distance (R) and vertical position (Z) of the closest point for tracks (a) and of the interaction vertex for showers (b).

a higher peak for tracks near to the score equal to one for the TL approach. Also, more showers are classified as tracks without the TL. As shown in table 3, the overall accuracy reduces to 77% when Transfer Learning is not employed, with a corresponding decrease in recall and precision for both classes.

5.3. Robustness tests

To ensure that the N -fit algorithm is stable, unbiased, and suitable for real physics analyses, a series of robustness checks were conducted. These tests aimed to assess the model dependence on the training dataset, its behaviour under extreme or unphysical inputs, and its capability to reproduce realistic results when applied to experimental data. The analyses were carried out using the track branch of N -fit and are described in detail in appendix B.

The first test consisted of a K-fold cross-validation, a standard method to evaluate the stability of ML models against data partitioning. The dataset used in the track branch development was divided into five equally sized folds, cyclically used for training, validation, and testing. The resulting mean absolute errors remained consistent across folds, indicating that the network's performance does not depend on a particular training subset. Additionally, a date-sorted K-fold test was performed to verify robustness over the operational history of the ANTARES detector. Despite small fluctuations related to detector conditions at different periods, the results confirmed that N -fit can handle data from various stages of the experiment without degradation in accuracy. See appendix B.1 for details.

The second test addressed the model's response to *background noise*. By feeding the networks with background-only images, we evaluated whether N -fit could mistakenly reconstruct a direction in the absence of a signal. As expected, the reconstructed directions followed a random uniform distribution, confirming that the model does not infer spurious correlations from pure noise. Moreover, the predicted uncertainties were appropriately large, reflecting the low confidence of the algorithm when the input lacks physical features of neutrino events. See appendix B.2 for details.

The last test compared the model's predictions on ANTARES data with those obtained from MC simulations. Since N -fit was trained exclusively on MC samples, this validation step was essential in verifying its applicability to experimental data. Using runs not employed during training, and ensuring the availability of the corresponding simulation files for a *run-by-run* comparison, we obtained distributions of reconstructed angles where simulations were in close agreement with the data. After applying quality cuts based on the estimated angular uncertainty (σ_{Ω}), most background and atmospheric muon events were rejected for up-going events, and the agreement further improved. This demonstrates that the N -fit method is reliable for use in real physics analyses. See appendix B.3 for details.

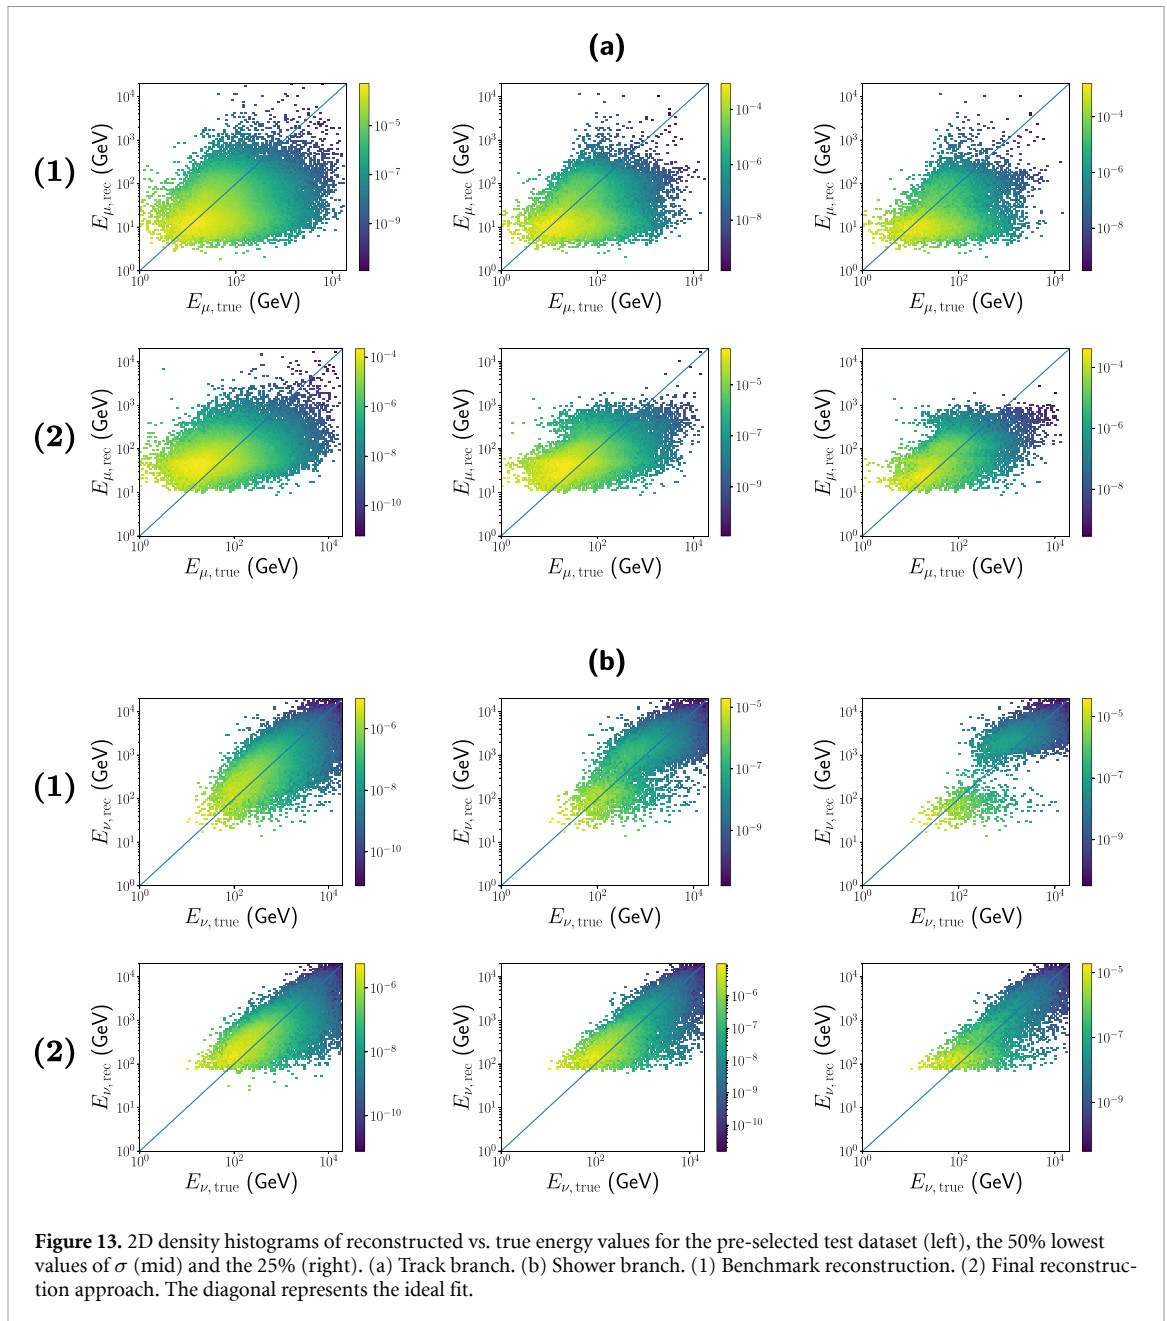


Figure 13. 2D density histograms of reconstructed vs. true energy values for the pre-selected test dataset (left), the 50% lowest values of σ (mid) and the 25% (right). (a) Track branch. (b) Shower branch. (1) Benchmark reconstruction. (2) Final reconstruction approach. The diagonal represents the ideal fit.

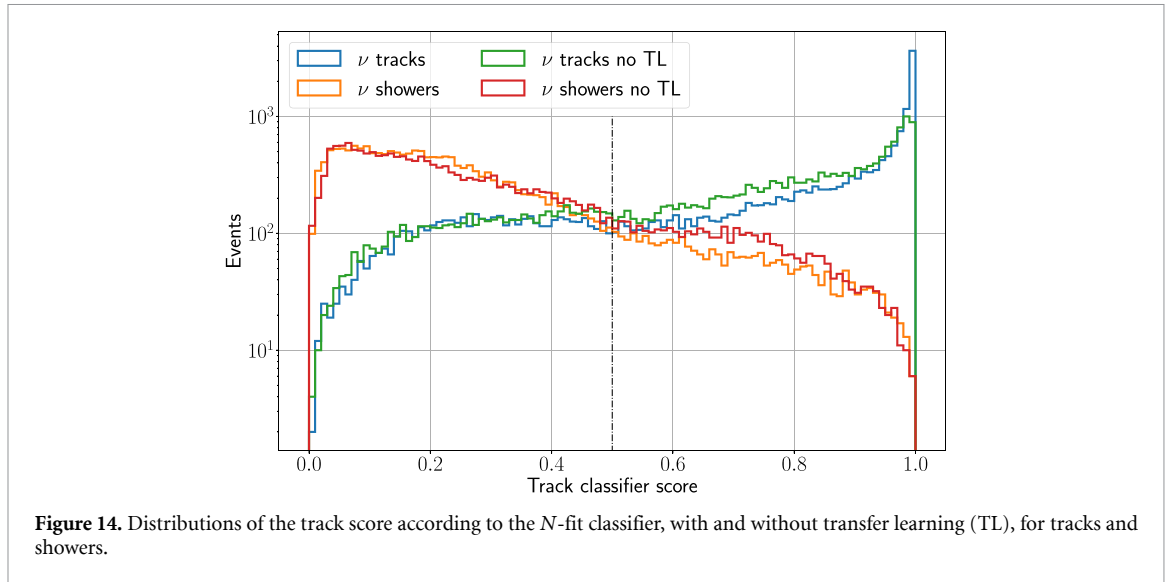
Table 3. Accuracy, bias, precision, and recall metrics comparison between the N -fit classifier of tracks (T) and showers (S) single-line events (with and without transfer learning) and the χ^2 -fit selection.

Method	Accuracy	Bias T/S	Precision T/S	Recall T/S
χ^2 -fit	63%	40/60%	67/61%	53/74%
N -fit (no TL)	77%	46/54%	80/75%	73/81%
N -fit (TL)	80%	45/55%	81/77%	75/85%

6. Physics application example

The N -fit track branch has been already used for different physics analyses. A binned methodology has been applied for an indirect Dark Matter search towards the Sun [30] and a final unbinned analysis is under development using the whole ANTARES dataset [31]. Also, the potential of the method has been tested with a point source search following an IceCube alert, as shown in this section.

The IceCube Collaboration reported a neutrino track-like event with a moderate probability of being of astrophysical origin that occurred on December 8, 2021 at 20:02:51.1 UT [32]. The reported position,



in equatorial coordinates referred to J2000 and with 90% point-spread function (PSF) containment, was:

$$RA = 114.52^\circ \begin{matrix} +2.82^\circ \\ -2.50^\circ \end{matrix} \quad \delta = 15.56^\circ \begin{matrix} +1.81^\circ \\ -1.39^\circ \end{matrix}.$$

The iceCube collaboration pointed to two gamma-ray sources listed in the 4FGL-DR2 Fermi-LAT catalog and located within the error region presented for the event: 4FGL J0738.4 + 1539 and 4FGL J0743.1 + 1713. Hence, the Collaboration encouraged additional follow-up by ground and space-based instruments. The Fermi-LAT Collaboration inspected the vicinity of the neutrino event with all-sky survey data from the Large Area Telescope (LAT) [33]. They did not find any significant photon detection from the aforementioned sources nor did they find a new transient source compatible with the best-fit direction of the event. However, outside the 90% error region, at a distance of 2.1° , the cataloged gamma-ray source 4FGL J0738.1 + 1742 was significantly detected. This source is associated with the BL Lac object PKS 0735 + 17, that is a blazar located at $RA = 114.5^\circ$ and $\delta = 17.7^\circ$.

The ANTARES Collaboration performed a follow-up of the blazar PKS 0735 + 17 in a sensible time window around the time of the IceCube alert. Two types of analyses were carried out: one binned analysis and another one unbinned [34]. Both of them covered a wide neutrino energy range above 300 GeV. The N -fit track branch is designed to be sensitive to even lower energies, so we performed a binned analysis for the follow-up of the blazar PKS 0735 + 17 using N -fit reconstructed parameters as a complementary study to that already performed. This study served as the first real physics analysis for the new reconstruction method.

The methodology followed here is based on [35]. It is based on comparing the measured data with the expected background inside a determined solid angle region—or region of interest (RoI)—and then computing the significance of the signal excess [35]. Classically, a circular RoI over the sphere is used, defined by its radius. However, in contrast to traditional methods analyzing multi-line events, our technique is asymmetric for the direction angles, to take into account the better reconstruction of the angle θ . Thus, we decided to modify the methodology by using an asymmetrical RoI, i.e. a rectangle in (θ, ϕ) coordinates defined by the half-length of the sides: R_θ and R_ϕ . In addition, to optimize the set of cuts, typically the equatorial coordinates are used, for which the right ascension is scrambled (blinded). Due to the highly asymmetric resolution of the N -fit reconstruction in θ and ϕ , and to ensure consistency with the predicted uncertainties, we performed the background estimation by scrambling the azimuthal coordinate ϕ in the local detector frame. Given the uniform distribution of event arrival times, this method is statistically equivalent to the standard right-ascension scrambling. Furthermore, optimizing the asymmetric (rectangular) RoI directly in the (θ, ϕ) system avoids potential biases that could arise from coordinate transformations between reference frames. This approach provides an unbiased background expectation and does not affect the derived sensitivity.

After completing the optimization, the set of parameters found for the study are:

$$\begin{aligned} R_\theta &= 5^\circ & R_\phi &= 70^\circ \\ \sigma_{\Omega, \max} &= 21^\circ, \end{aligned} \tag{15}$$

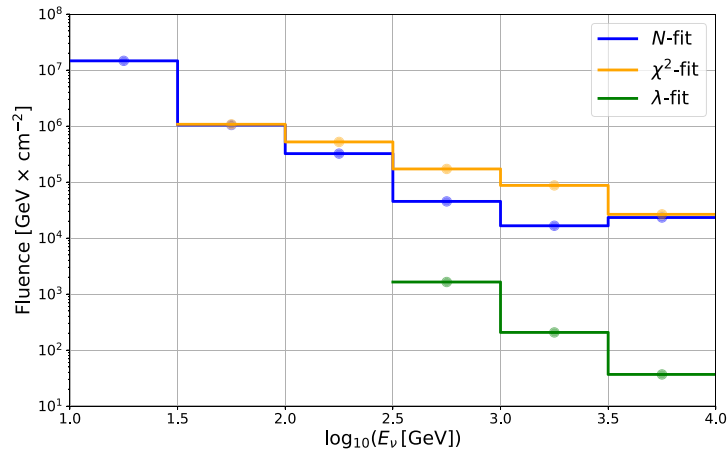


Figure 15. 90% confidence upper limits on the neutrino fluence as a function of the energy for N -fit, χ^2 -fit and λ -fit from PKS 0735 + 17 over ± 1 day around 20:02:51.1 UT on December 8, 2021. The values of the study performed with the λ -fit reconstruction are taken from [34], while the values for χ^2 -fit—using only SL events—were computed just for comparison purposes.

where $\sigma_{\Omega, \max}$ is the maximum angular error estimation allowed for the event selection. We did not find any signal in this region for the given date of the study. Therefore, we present the 90% confidence level (CL) upper limits on the neutrino fluence (\mathcal{F}).

Using the fact that the 90% CL upper limit compatible with no event detection is $n_s^{90\%} = 2.44$, according to the Feldman and Cousins method [36] and assuming an E^{-2} energy spectrum, the upper limit for the fluence is:

$$\mathcal{F}^{90\%} = \frac{2.44}{\text{Acc}} \cdot \log(E_{\max}/E_{\min}) \cdot T \quad [\text{GeV cm}^{-2}], \quad (16)$$

where Acc is the acceptance of the detector for our reconstruction technique, and T is the search time (~ 2 days). This computation is performed for a discrete histogram, in such a way that for each bin we have that $\log(E_{\max}/E_{\min}) = 0.5 \cdot \log(10)$. The same procedure was applied for the χ^2 -fit reconstructions, just skipping the steps that include a computation with the azimuthal angle, since its reconstruction is missing. The RoI for the χ^2 -fit method was simply a full zenithal band characterized by its width. The results are shown in figure 15, along with the ones of the previous study carried out with the usual λ -fit reconstruction method.

The N -fit study is sensitive to low energies, being complementary to the previous λ -fit analysis in this sense. Only N -fit is able to retain enough statistics to compute the limits below 30 GeV. The limits of the N -fit method are slightly better to those of χ^2 -fit using SL events in the energy range between 30 and 300 GeV. Above 300 GeV, where all methods can be compared, the λ -fit provides more stringent limits than N -fit, as expected.

7. Discussion and conclusions

In this work, we have developed and presented a new reconstruction and classification approach based on DL for SL events in the ANTARES telescope. The exploration to achieve the best possible reconstruction has led us to a novel combination of ML techniques, such as the use of deep convolutional layers, a mixture density output layer, and TL. The reconstruction of the spatio-temporal parameters of the SL event, such as the incoming neutrino direction and location, can directly be obtained using deep CNNs. Combining a DCN with a mixture density output layer allowed us to estimate the quality of the prediction. Other more complex tasks, however, such as energy estimation or event classification, benefitted from using TL. For the energy reconstruction, this study presents an innovative *indirect* TL approach based on PCA knowledge distillation. Using this dimensionality reduction technique, we selected the most informative neuron activations from networks performing spatial tasks, and used them as input features for training a feed-forward network in energy reconstruction. This method provided results that can be exploited in physics analyses. *Direct* TL was applied to classification tasks, by freezing the convolutional layers of pretrained network models specialized in event-type-specific spatial information.

Thus, we developed N -fit as a repertoire of specialized modules that together fully characterized SL neutrino events in terms of energy, spatial properties and type of event, track or shower (appendix B.3).

Compared to previous techniques applied to SL events in ANTARES, N -fit outperforms the θ angle reconstruction of SL events provided by the χ^2 -fit. For track-like events, the mean θ error decreased from $\sim 9.7^\circ$ (standard χ^2 -fit) to $\sim 3.7^\circ$ for the best 50% reconstructed events. Moreover, N -fit provides a sufficiently accurate estimation for the ϕ angle—with a mean error of $\sim 29^\circ$ for the best 50% track events—that is useful in physics analyses. Such information was not attempted in χ^2 -fit due to the approximations made in the method. Specifically for track events, N -fit provides a very precise reconstruction of the closest point of the track event to the ANTARES detector line and a first approximation of the muon energy reconstruction, taking into account the physical difficulties in this endeavour, particularly in tracks. To address these challenges, the energy module employs a sequential pipeline where spatial reconstruction and its associated uncertainties act as quality filters. Although this selection criteria is based on the best 50% of spatial reconstructions, the high correlation between the spatial sub-models results in a final cumulative acceptance of approximately 40%. This selective process ensures that energy estimations are derived from a high-confidence sample, enhancing the reliability of downstream physics results. Regarding shower events, their interaction vertex reconstruction was very precisely estimated in N -fit. The neutrino energy in shower events was reconstructed with a 0.50 relative error, compared to a 0.55 relative error for muon energy in track events. Lastly, discrimination between the two types of neutrino events by the N -fit classifier reached a high accuracy level, of around 80%. When comparing data and MC simulations, N -fit demonstrates a good agreement in the track branch, heavily reducing the background noise and atmospheric muons of the up-going sample when quality cutoffs are applied.

N -fit modules characterizing track events have already been used in physics analyses as shown in section 6. We have a promising improvement in the low energy range for dark matter WIMP searches towards the Sun [30, 31]. Also, we have demonstrated that the new methodology is able to complement classical reconstruction methods in point sources neutrino searches. In addition, we are actively applying N -fit reconstructions in different analyses considering the complete ANTARES dataset.

More generally, the repertoire and combination of ML techniques used in N -fit may be adapted to reconstructions in other particle detectors, such as KM3NeT, as well as inspiring new applications in other areas of computational physics.

Acknowledgment

The authors acknowledge the financial support of the funding agencies: Centre National de la Recherche Scientifique (CNRS), Commissariat à l'énergie atomique et aux énergies alternatives (CEA), Commission Européenne (FEDER fund and Marie Curie Program), LabEx UnivEarthS (ANR-10-LABX-0023 and ANR-18-IDEX-0001), Région Alsace (contrat CPER), Région Provence-Alpes-Côte d'Azur, Département du Var and Ville de La Seyne-sur-Mer, France; Bundesministerium für Bildung und Forschung (BMBF), Germany; Istituto Nazionale di Fisica Nucleare (INFN), Italy; Nederlandse organisatie voor Wetenschappelijk Onderzoek (NWO), the Netherlands; Ministry of Education and Scientific Research, Romania; MCIN for PID2024-156285NB-C41, -C42, -C43, PID2021-124591NB-C41, -C42, -C43, PID2020-120037GA-I00, and PDC2023-145913-I00 funded by MICIU/AEI/10.13039/501100011033 and by 'ERDF A way of making Europe', for ASFAE/2022/014 and ASFAE/2022 /023 with funding from the EU NextGenerationEU (PRTR-C17.I01) and Generalitat Valenciana, for Grant AST22_6.2 with funding from Consejería de Universidad, Investigación e Innovación and Gobierno de España and European Union—NextGenerationEU, for CSIC-INFRA23013 and for CNS2023-144099, Generalitat Valenciana for CIDEAGENT/2019/043, CIDEAGENT/2020/049, CIDEAGENT/2021/23, CIDEIG/2023/20, ESGENT/2024/24, CIPROM/2023/51, GRISOLIAP/2021/192 and INNVA1/2024/110 (IVACE+i), Spain; Ministry of Higher Education, Scientific Research and Innovation, Morocco, and the Arab Fund for Economic and Social Development, Kuwait. We also acknowledge the technical support of Ifremer, AIM and Foselev Marine for the sea operation and the CC-IN2P3 for the computing facilities.

Data availability statement

Due to the huge volume of raw data collected throughout the 16 years of continuous operation of the ANTARES telescope, no suitable public repository currently exists for hosting the complete dataset. Furthermore, the data are restricted by the ANTARES collaboration policy and require specific authorization for their use.

Therefore, the data cannot be made publicly available upon publication because no suitable repository exists for hosting data in this field of study. The data that support the findings of this study are available upon reasonable request from the authors.

Supplementary data 1 available at: <https://doi.org/10.1088/2632-2153/ae5d84/data1>.

Author contributions

M Ardid  0000-0002-3199-594X

Conceptualization (equal), Investigation (equal), Supervision (equal), Writing – original draft (equal), Writing – review & editing (equal)

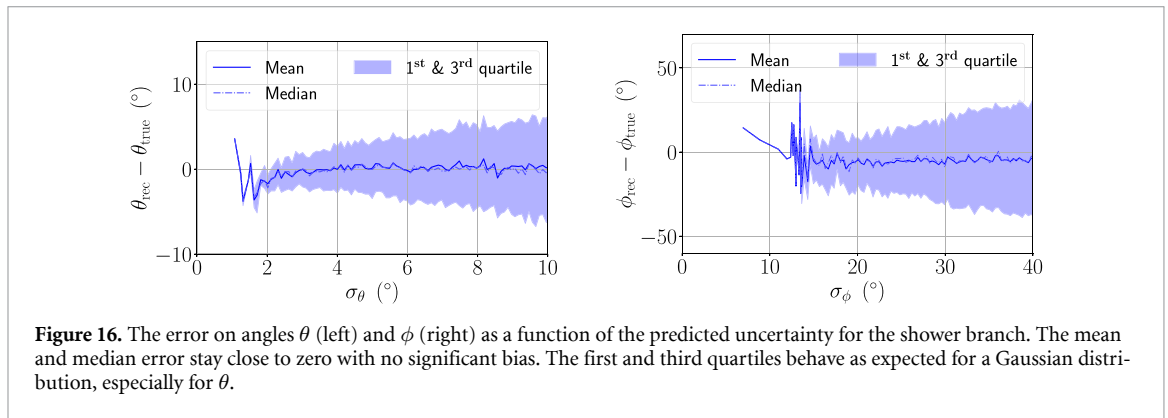
S Ardid  0000-0003-4821-6655

Conceptualization (equal), Investigation (equal), Supervision (equal), Writing – original draft (equal), Writing – review & editing (equal)

Appendix A. Supplementary material

Table 4. Mean and median absolute error of reconstructed neutrino angles θ , ϕ , and total deviation, Ω , for the shower branch. The 50% best events are selected according to the lowest values of the parameters σ for N -fit and χ^2 for the χ^2 -fit.

Mean/Median	χ^2 -fit		N -fit		χ^2 -fit (50%)		N -fit (50%)	
θ	27.4°	22.8°	12.3°	8.9°	24.9	20.5	7.9°	5.8°
ϕ	—	—	47.0°	37.7°	—	—	34.2°	28.8°
Ω	—	—	39.4°	34.9°	—	—	32.2°	27.8°



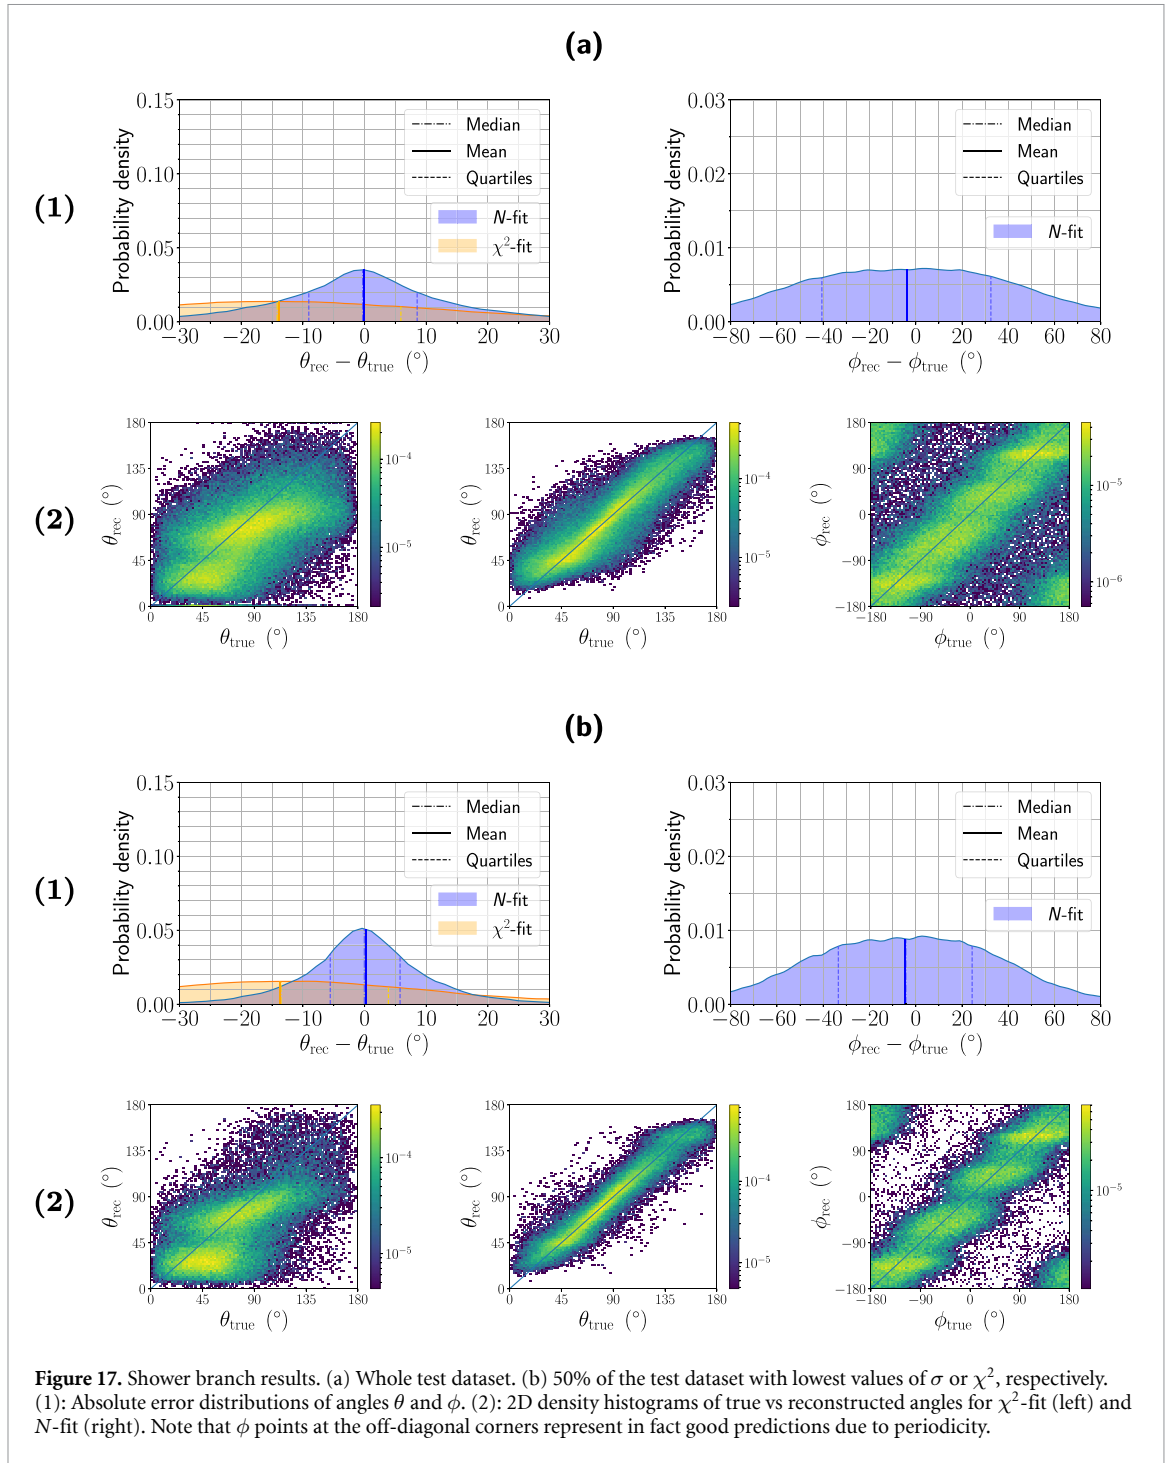


Figure 17. Shower branch results. (a) Whole test dataset. (b) 50% of the test dataset with lowest values of σ or χ^2 , respectively. (1): Absolute error distributions of angles θ and ϕ . (2): 2D density histograms of true vs reconstructed angles for χ^2 -fit (left) and N-fit (right). Note that ϕ points at the off-diagonal corners represent in fact good predictions due to periodicity.

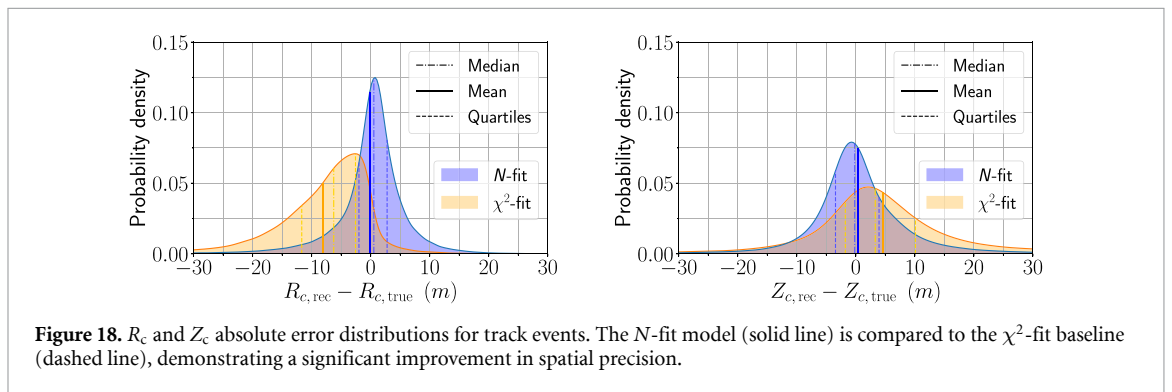
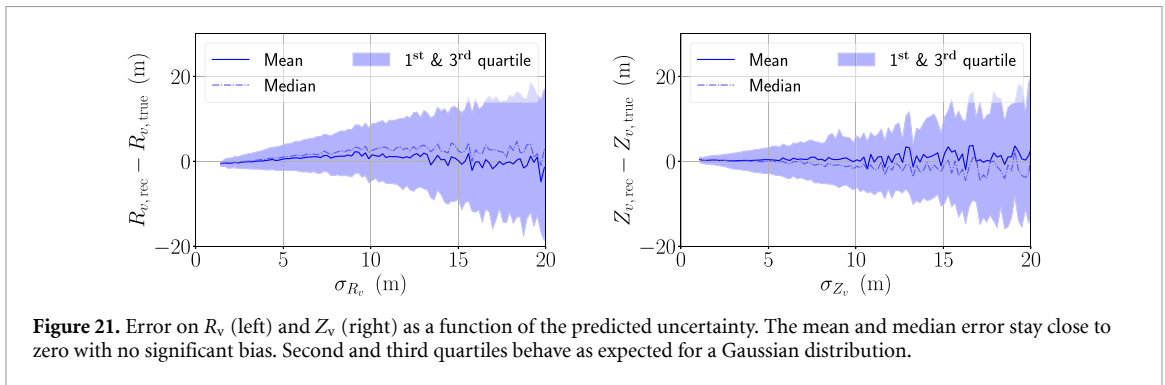
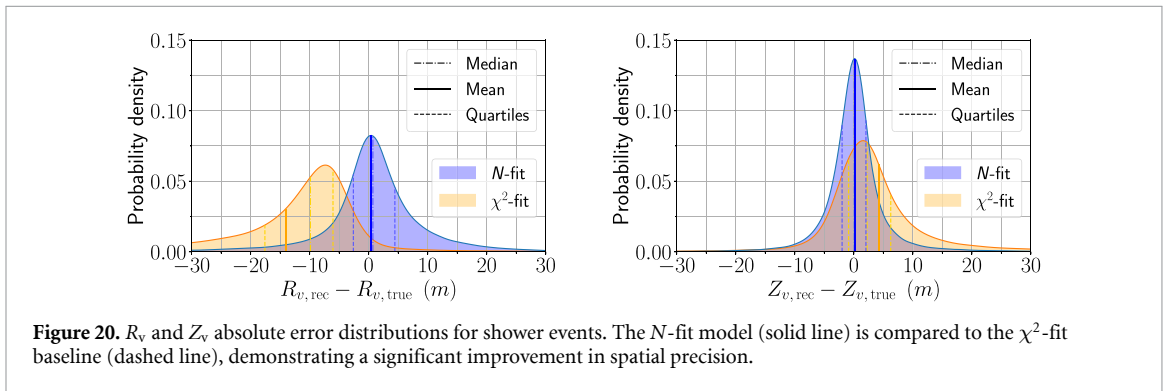
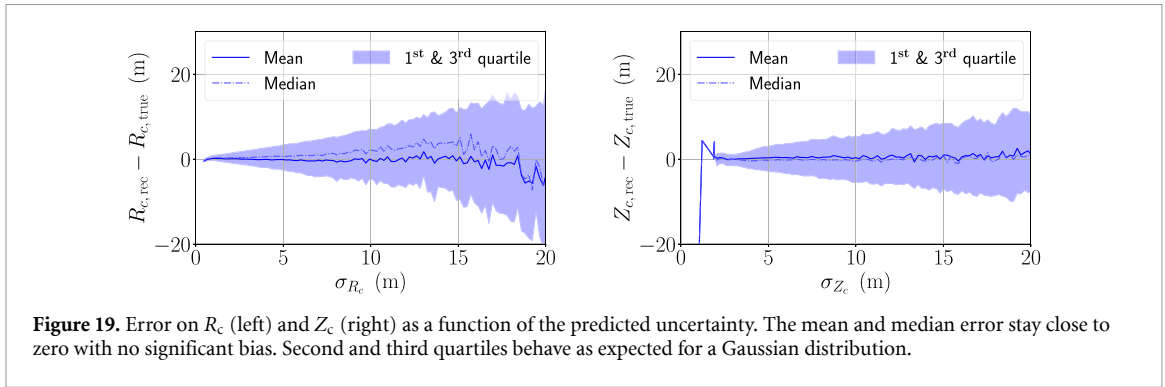


Figure 18. R_c and Z_c absolute error distributions for track events. The N-fit model (solid line) is compared to the χ^2 -fit baseline (dashed line), demonstrating a significant improvement in spatial precision.



Appendix B. Robustness tests

B.1. K-fold cross validation (KCV)

The KCV technique [37] is largely used in ML to check the robustness of the networks and their possible dependence with the dataset chosen for the training process. The technique consists in splitting the dataset in K folds and use some of them for training while the others are used for validation and test. Then, the data is shifted circularly 1 fold and a second training is performed. The process is repeated until reaching K circular shifts. This ensures that all the folds take all the different positions in training, validation and test.

We have randomly split the dataset in 5 folds with 20% of the data in each one. Then, we have used 3 folds for training, 1 for validation and 1 for test. Consequently, we have performed the training process 5 times. Statistics of the results can be seen in table 5. The errors are stable and similar to those in table 2, showing that the method has no dependence on a particular data split.

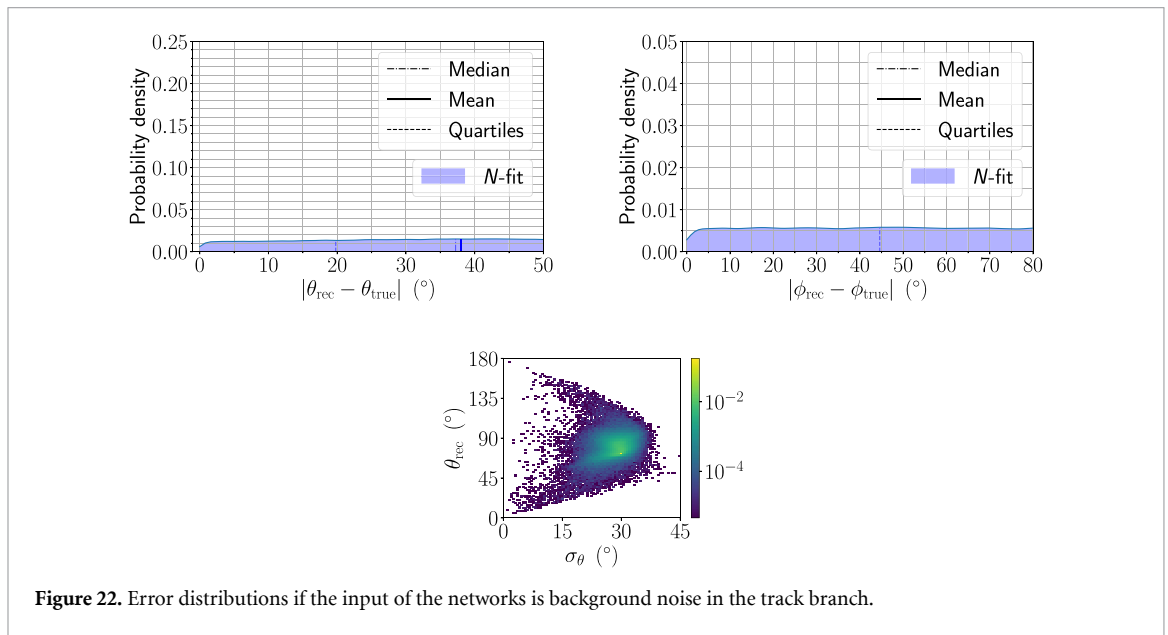
Since the conditions of the ANTARES telescope have changed through the years, we need to ensure that the network model is able to fit different stages of the telescope properly. To that end, we performed another KCV, but this time with folds sorted by date. Thus, we split the data in 5 folds, each one containing events within the following dates:

Table 5. Mean absolute error of predicted neutrino angles θ and ϕ for the 5-fold cross-validation. The fold number denotes the fold used as test.

MAE	fold 5	fold 4	fold 3	fold 2	fold 1
θ	7.4°	7.3°	7.4°	7.4°	7.4°
ϕ	41.5°	41.3°	41.6°	41.4°	41.3°
Ω	28.2°	28.1°	28.4°	28.1°	28.1°

Table 6. Mean absolute error of predicted neutrino angles θ and ϕ for the 5-fold dates-sorted cross-validation. The fold number is the one used as test.

MAE	fold 5	fold 4	fold 3	fold 2	fold 1
θ	8.1°	7.6°	7.4°	7.5°	6.9°
ϕ	44.4°	42.2°	41.4°	41.7°	39.7°
Ω	30.0°	28.7°	28.4°	29.1°	26.3°



- Fold 1: from 06/02/2007 to 13/01/2010
- Fold 2: from 13/01/2010 to 17/05/2011
- Fold 3: from 22/05/2011 to 02/09/2012
- Fold 4: from 03/09/2012 to 24/04/2015
- Fold 5: from 24/04/2015 to 27/12/2017

Statistics of the results can be seen in table 6. We acknowledge modest fluctuations in the results that emerge from the fact that not all lines and OMs were always operative in ANTARES at all times. The performance was slightly better if the first stages of the telescope are used as the test dataset. This is possibly due to the fact that these data had better quality, compared to the others, since the telescope was not affected by ageing effects. Thus, having better data in the test set gives raise to better performance.

B.2. Background noise

As a control of the predictions generated by our model, we processed images of the background and fed them to the N -fit track direction networks. The expectation of using the background as the signal was to obtain direction reconstructions following a random uniform distribution, and that was indeed what we obtained as reported in figure 22.

Moreover, the predicted uncertainty was typically large, as seen in the bottom panel of figure figure 22 for the θ angle.




Figure 23. Comparison between data and MC simulations.




Figure 24. Comparison between data and MC simulations using a cut of $\sigma_{\Omega} < 21^{\circ}$.

B.3. Comparison with data and systematic uncertainties

The primary objective of this work is the development and validation of the N -fit reconstruction framework. Consequently, this study utilizes established MC simulation protocols and standards that have been extensively documented and validated throughout the operational history of the ANTARES Collaboration. This validation is performed using run-by-run simulations, which incorporate the specific telescope configuration and environmental conditions for each period, as detailed in [17]. Furthermore, every simulated event is provided with weights to account for the interaction, propagation, and detection physics, as well as the expected flux of every type of event [38].

For the comparison, we selected new runs that were not used for training. We also made sure that, for every selected run, all the MC simulation files were available. Figures 23 and 24 illustrate the results from this comparison.

As seen in figure 23, the dataset is dominated by downgoing atmospheric muons, showing good agreement for $\theta > 90^{\circ}$. The initial visual discrepancies observed at $\theta < 80^{\circ}$ and in the ϕ distribution are primarily due to uncorrelated optical noise from ^{40}K radioactivity and bioluminescence. These noise-driven events typically exhibit large reconstruction uncertainties, as they lack the structured spatio-temporal features of neutrino interactions.

We observe that the agreement in the upgoing region improves significantly when a cutoff based on the quality parameter $\sigma_{\Omega} < 21^{\circ}$ is applied (figure 24), which effectively filters the background noise and atmospheric muons. A residual discrepancy may remain for nearly vertical tracks ($\theta \lesssim 10^{\circ}$), but this corresponds to a very small region of the sky and is handled through dedicated systematic evaluations in specific physics analyses.

Following established ANTARES procedures (e.g. [39]), we evaluated the sensitivity of the reconstruction to systematic uncertainties by varying key parameters. The overall OM sensitivity was varied by $\pm 10\%$, accounting for uncertainties in photon-to-photoelectron conversion efficiency, while modifications to the angular acceptance were found to affect measurements by less than 10%. Additionally, the light absorption length in water was scaled by $\pm 10\%$. Investigations specific to SL events, such as detector line verticality, showed negligible impact (see [40]). The overall systematic uncertainty is estimated at 15% or less. Notably, since SL events have larger intrinsic reconstruction uncertainties than multi-line events, the relative contribution of these systematic effects is reduced, further supporting the robustness of the N -fit results.

References

- [1] ANTARES Collaboration 2011 ANTARES: the first undersea neutrino telescope *Nucl. Instrum. Methods Phys. Res. A* **656** 11–38
- [2] Psihas F, Groh M, Tunnell C and Warburton K 2020 A review on machine learning for neutrino experiments *Int. J. Mod. Phys. A* **35** 2043005
- [3] Bukhari H, Chakraborty D, Eller P, Ito T, Shugaev M V and Ørsøe R 2024 Icecube – neutrinos in deep ice *Eur. Phys. J. C* **84** 646
- [4] KM3NeT Collaboration 2020 Event reconstruction for KM3NeT/ORCA using convolutional neural networks *J. Instrum.* **15** 10005
- [5] Reck S, Guderian D, Vermariën G and Domi A KM3NeT collaboration 2021 Graph neural networks for reconstruction and classification in KM3NeT *J. Instrum.* **16** C10011
- [6] Jones A, Kuehnert J, Fraccaro P, Meuriot O, Ishikawa T, Edwards B, Stoyanov N, Remy S L, Weldemariam K and Assefa S 2023 AI for climate impacts: applications in flood risk *npj Clim. Atmos. Sci.* **6** 63
- [7] Glaysher P 2018 BDTs in the search for $t\bar{t}$ production with higgs decays to $b\bar{b}$ at ATLAS PoS **EPS-HEP2017** 698
- [8] Duev D A et al 2019 Real-bogus classification for the Zwicky Transient Facility using deep learning *Mon. Not. R. Astron. Soc.* **489** 3582–90
- [9] ANTARES Collaboration 2011 A fast algorithm for muon track reconstruction and its application to the ANTARES neutrino telescope *Astropart. Phys.* **34** 652–62
- [10] ANTARES Collaboration 2017 An algorithm for the reconstruction of neutrino-induced showers in the ANTARES neutrino telescope *Astron. J.* **154** 275
- [11] Heijboer A 2004 Track reconstruction and point source searches with ANTARES *PhD Thesis* Amsterdam U
- [12] ANTARES Collaboration 2012 Search for cosmic neutrino point sources with four years of data from the ANTARES telescope *Astrophys. J.* **760** 53
- [13] Gu J et al 2018 Recent advances in convolutional neural networks *Pattern Recognit.* **77** 354–77
- [14] Bishop C M 1994 Mixture density networks *NCRG/94/004* (available at: <https://research.aston.ac.uk/en/publications/mixture-density-networks>)
- [15] Zhuang F, Qi Z, Duan K, Xi D, Zhu Y, Zhu H, Xiong H and He Q 2021 A comprehensive survey on transfer learning *Proc. IEEE* **109** 43–76
- [16] ANTARES Collaboration 2025 The ANTARES detector: two decades of neutrino searches in the Mediterranean Sea *Phys. Rep.* **1121-1124** 1–46
- [17] Albert A ANTARES Collaboration 2021 Monte Carlo simulations for the ANTARES underwater neutrino telescope *J. Cosmol. Astropart. Phys.* **JCAP01(2021)064**
- [18] Rosenblatt F 1958 The perceptron: a probabilistic model for information storage and organization in the brain *Psychol. Rev.* **65** 386–408
- [19] Daubechies I, DeVore R, Foucart S, Hanin B and Petrova G 2022 Nonlinear approximation and (deep) ReLU networks *Construct. Approx.* **55** 127–72
- [20] Kumar S K 2017 On weight initialization in deep neural networks *Technical Report* (arXiv:1704.08863 [cs])
- [21] He K, Zhang X, Ren S and Sun J 2015 Delving deep into rectifiers: surpassing human-level performance on imagenet classification *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*
- [22] Ruder S 2017 An overview of gradient descent optimization algorithms *Technical Report* (arXiv:1609.04747 [cs])
- [23] Kingma D P and Ba J 2017 Adam: a method for stochastic optimization, *Technical Report* (arXiv:1412.6980 [cs])
- [24] Prechelt L 1998 *Early Stopping - But When?* (*Lecture Notes in Computer Science*) (Springer) pp 55–69
- [25] Srivastava N, Hinton G, Krizhevsky A, Sutskever I and Salakhutdinov R 2014 Dropout: a simple way to prevent neural networks from overfitting *J. Mach.: Learn. Res.* **15** 1929–58
- [26] Glorot X and Bengio Y 2010 Understanding the difficulty of training deep feedforward neural networks *JMLR Workshop and Conf. Proc.* pp 249–56
- [27] Moslemi A, Briskina A, Dang Z and Li J 2024 A survey on knowledge distillation: Recent advancements *Mach. Learn. Appl.* **18** 100605
- [28] Kherif F and Latyпова A 2020 Chapter 12 - Principal component analysis *Machine Learning* ed A Mechelli and S Vieira (Academic) pp 209–25
- [29] Abadi M et al 2015 TensorFlow: large-scale machine learning on heterogeneous systems (available at: www.tensorflow.org/)
- [30] García-Méndez J, Ardid S and Ardid M 2023 Dark matter search towards the Sun using Machine Learning reconstructions of single-line events in ANTARES PoS **ICRC2023** 1443
- [31] Poiré C 2024 Indirect dark matter searches towards the Sun using the full ANTARES data set (available at: <https://doi.org/10.5281/zenodo.13350925>)
- [32] IceCube Collaboration 2021 GCN circular #31191 (available at: <https://gcn.gsfc.nasa.gov/gcn/gcn3/31191.gcn3>)
- [33] Fermi-LAT Collaboration 2021 GCN circular #31194 (available at: <https://gcn.gsfc.nasa.gov/gcn3/31194.gcn3>)
- [34] Alves S ANTARES collaboration 2022 ANTARES follow-up of IceCube event IC211208A coincident with flaring blazar PKS 0735+17 (available at: <https://zenodo.org/record/6785295>)
- [35] Alves Garre S, Versari F, Avrorin A D, Dzhlkibaev Z-A M, Shelepov M D and Suvorova O V the ANTARES collaboration 2021 ANTARES offline study of three alerts after Baikal-GVD follow-up found coincident cascade neutrino events *37th Int. Cosmic Ray Conf.* vol ICRC2021 p 1121
- [36] Feldman G J and Cousins R D 1998 Unified approach to the classical statistical analysis of small signals *Phys. Rev. D* **57** 3873–89
- [37] Anguita D, Ghelardoni L, Ghio A, Oneto L and Ridella S 2012 The ‘K’in K-fold cross validation *20th European Symp. on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)* (i6doc.com publ) pp 441–6
- [38] Colnard C M M 2009 Ultra-high energy neutrino simulations *PhD Thesis* (available at: <https://inspirehep.net/literature/823582>)
- [39] Adrián-Martínez S et al 2013 Measurement of the atmospheric ν_μ energy spectrum from 100 GeV to 200 TeV with the ANTARES telescope *Eur. Phys. J. C* **73** 2606
- [40] Méndez J G 2025 Light at the end of the line: ML-enhanced reconstruction of single-line events in ANTARES and an indirect search for dark matter in the Sun *PhD Thesis* Universitat Politècnica de València