

Enhancing events in neutrino telescopes through deep-learning-driven superresolution

Felix J. Yu^{1,2,*}, Nicholas Kamp^{2,†} and Carlos A. Argüelles^{1,2,‡}

¹*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions*

²*Department of Physics and Laboratory for Particle Physics and Cosmology, Harvard University, Cambridge, Massachusetts 02138, USA*



(Received 3 September 2024; accepted 6 January 2025; published 5 February 2025)

Recent discoveries by neutrino telescopes, such as the IceCube Neutrino Observatory, relied extensively on machine learning (ML) tools to infer physical quantities from the raw photon hits detected. Neutrino telescope reconstruction algorithms are limited by the sparse sampling of photons by the optical modules due to the relatively large spacing (10–100 m) between them. In this Letter, we propose a novel technique that learns photon transport through the detector medium through the use of deep-learning-driven superresolution of data events. These “improved” events can then be reconstructed using traditional or ML techniques, resulting in improved resolution. Our strategy arranges additional “virtual” optical modules within an existing detector geometry and trains a convolutional neural network to predict the hits on these virtual optical modules. We show that this technique improves the angular reconstruction of muons in a generic ice-based neutrino telescope. Our results readily extend to water-based neutrino telescopes and other event morphologies.

DOI: [10.1103/PhysRevD.111.L041301](https://doi.org/10.1103/PhysRevD.111.L041301)

Introduction. Neutrino telescopes have opened a new window into the Universe beyond the energies and distances available to radio and γ -ray astronomy, respectively [1]. This is facilitated by the largest particle detectors ever built, typically an array of light detectors distributed throughout a gigaton of water or ice. Neutrino astronomy is currently led by the IceCube Neutrino Observatory, which has instrumented a cubic kilometer of South Pole ice with 5160 digital optical modules (DOMs) [2] arranged in 86 vertical strings, such that DOMs are spaced by 17 m (125 m) in the vertical (horizontal) direction. IceCube has made a series of breakthrough discoveries over the past decade, including the first observation of high-energy astrophysical neutrinos [3,4], as well as the first neutrinos from a blazar, TXS 0506 + 056 [5], and an active galaxy, NGC 1068 [6]. More recently, IceCube made the first detection of neutrinos from our own Galactic plane [7], relying extensively on machine learning (ML) tools to determine the energy and direction of the original neutrino [8,9]. This is only one example of the emerging importance

of ML techniques in neutrino telescopes [9–19] and neutrino physics as a whole [20–27].

Building upon the success of IceCube, several new ice- and water-based neutrino telescopes are being built or planned. In the former category are KM3NeT in the Mediterranean Sea [28], Baikal-GVD in Lake Baikal [29], and P-ONE in the Pacific Ocean [30]. In the latter category are IceCube-Gen2 [31] and a series of water-based detectors in China, including TRIDENT [32] and HUNT [33], all of which would increase the instrumented volume by an order of magnitude. Because of the larger size, the light sensors in these next-generation detectors will be spaced further apart. This will reduce the information available for each detected neutrino interaction, especially at lower energies.

To address this, we introduce the concept of “super-resolution” for neutrino telescope data. This technique is used in the image-processing community to increase the definition of the pictures at the subpixel level [34]. For neutrino telescope data, superresolution translates to sampling photon depositions between optical modules (OMs). In this Letter, this is achieved by training a deep-learning algorithm to predict detected photons on virtual OMs arranged between the physical OMs, as depicted in Fig. 1. We show that angular reconstruction algorithms perform better on superresolved events with information from the virtual DOMs. This will be essential for improving the sensitivity of existing and future neutrino telescopes to astrophysical neutrino sources within and beyond our Galaxy. Further, the techniques proposed in this Letter can

*Contact author: felixyu@g.harvard.edu

†Contact author: nkamp@g.harvard.edu

‡Contact author: carguelles@g.harvard.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Funded by SCOAP³.




FIG. 1. Event displays, showing the masked, superresolved, and unmasked event. The unmasked and masked events are obtained from simulation, representing ideal and realistic detector configurations, while the superresolution network attempts to enhance the masked event into the unmasked. The top plot shows the photon arrival time series from the superresolution network and the pretrained VAE on two particular virtual OM in the superresolved event.

improve particle reconstruction in other light-based neutrino experiments, such as Super-Kamiokande [35], Hyper-Kamiokande [36], JUNO [37], and THEIA [38], among others.

The rest of the Letter is organized as follows. Section II describes the architecture of the superresolution network. Next, Sec. III discusses the training procedure and dataset generation. Section IV demonstrates the improved angular resolution for superresolved muon events in an IceCube-like detector. Finally, Sec. V summarizes the results and potential future directions of this work. The code implemented for this Letter is made available at [39].

Architecture. The basic premise of the superresolution method is to start by simulating events within a denser geometry and then mask a portion of these strings during training. The network is trained to reconstruct the sensors on the masked strings, which we will refer to as virtual strings. Our architecture is independent of the specific detector geometry, while simultaneously preserving as much timing information as possible. This versatility enables the network to adapt to various geometries and detectors, as well as to different event morphologies.

Our implementation of superresolution is organized by a high-level algorithm that compares predicted superresolved events with their true, unmasked events, as depicted in Fig. 2. This is achieved by encoding each event into a 2D image and passing these images through a UNet-type network. The

output is also a 2D image representing the superresolved event. One of the main challenges with this pipeline is adequately encoding and decoding the timing information, which is needed for the prediction of the time series on virtual OM. This difficulty arises from the geometrically sparse distribution of photon hits for any given neutrino interaction and the highly variable length of the photon hit time series in a given OM. Recently, Ref. [9] demonstrated that an asymmetric Gaussian mixture model can effectively represent the time series of showerlike events in IceCube. Building upon this concept, we train a variational autoencoder (VAE) to encode the time series into a compressed latent space [17]. The core idea of this process is to take a finely binned distribution of photon hit times (we chose to use 5000 bins, each 1 ns wide) and compress it into a more compact yet effective representation—a 64-parameter latent vector. The UNet then operates fully in the latent space, taking in a point in the latent space for each physical OM and predicting one for each virtual OM. Examples of the VAE representation and superresolution network predicted time series for two specific virtual OM are shown in the top plot of Fig. 1. While the shapes of most predicted timing probability density functions closely match that of the true hits, those farther from the interaction point, with fewer statistics, are harder to reconstruct accurately. Given the limited data on these OM, it is often sufficient to capture key features like the first hit time and the position of the peak, rather than the exact shape of the distribution.




FIG. 2. Pipeline of the superresolution framework. The OM timing information is encoded into a 64-parameter latent vector. Neutrino telescope events are arranged into 2D images by string and sensor on a string. For the network inputs, each OM contains 68 features: 3 (3D sensor position) + 1 (number of photon hits) + 64 (timing latent vector).

Event Generation and Training.

Virtual strings: To produce our dataset, we use the open-source neutrino telescope simulation toolkit Prometheus [40]. This allows us to define arbitrary detector geometries. As a benchmark, we use an orthogonal grid of OMs in ice comprising a 15×15 array of strings, each with 61 OMs, for a total of 13,705 OMs. The interstring spacing is 60 m, and the OMs vertical spacing is 15 m. As input to the network, we mask out every other string, resulting in an 8×8 grid of strings with 120 m spacing. We note that the virtual OM shadow effect, caused by photons being absorbed when hitting virtual OMs in the denser geometry, is negligible. Thus, in this particular geometry, characterized by a dense arrangement of 225 strings with 61 OMs

per string, we generate 225×61 images, as illustrated in Fig. 2.

Figure 3 shows the top-down view of the virtual and physical string layout. An example event is shown to demonstrate the effect of masking. For visualization purposes, only the top OM in each string is shown, and its color indicates the time of the first hit. Recall that the VAE-based treatment of the hit time distributions described in Sec. II allows this method to take advantage of the full timing information of the data, as shown in Fig. 1.

Dataset and training details: In this Letter, we specifically investigate tracklike events, corresponding to muons produced in ν_μ charged-current interactions. The energies of the incident neutrinos are sampled from a power-law distribution with a spectral index of -1 between 100 GeV and 1 PeV. We select only events that hit at least eight distinct OMs in the masked geometry. This procedure results in approximately 500,000 selected events for the training dataset, and 75,000 selected events for the test dataset. We implement all network components using the PyTorch deep-learning framework and use the UNet++ implementation provided by Ref. [41]. More details can be found in Appendix A.

Results. The main result of this study is an observed improvement in the angular reconstruction of superresolved events. This is presented in Sec. IV A. We have also investigated the robustness of our algorithm to uncertainties in the detector medium optical properties, which is discussed in Sec. IV B.

Directional reconstruction improvements: To assess the performance of the superresolution network, we run a




FIG. 3. Top-down view of a simulated track event, showing both real and virtual OM strings. Color indicates a hit and its timing.




FIG. 4. Log-scale angular resolution. The median lines are drawn in solid color as a function of the true neutrino energy, produced by a baseline SSCNN method. The 20 and 80 percentiles are denoted by the dashed lines and shaded regions.

baseline directional reconstruction method on the masked input events, the superresolved events produced by the network, and the true unmasked events from the simulation. We used a sparse submanifold convolutional neural network (SSCNN), as described in Ref. [42]; however, we note that the angular resolution improvements should extend to other reconstruction methods. We train a separate SSCNN model for each of the three datasets.

Figure 4 shows the angular resolution as predicted by the SSCNN algorithm. As expected, the superresolved events do not outperform the unmasked events, as the super-resolution network uses the unmasked events as ground truth. There is noticeable improvement across the whole energy range for the superresolved dataset, corresponding to about 0.8° on average. The gain is higher in the lower energy bins below 1 TeV, around 1.3° on average.

The runtime performance is also efficient due to the compression of events into the latent space. During inference on an NVIDIA A100 GPU, with a batch size of 128, an event takes an average of 0.311 ms on the forward pass. This is negligible compared to standard angular reconstruction algorithm computational times, which can be tens of milliseconds, and comparable to the SSCNN forward pass time of ~ 0.1 ms [42].

Sensitivity to detector medium optical properties: Similar to many ML-based reconstruction methods, the superresolution process is highly dependent on the simulation and detector properties. This dependency is particularly significant for superresolution, as we treat the “unmasked” events in the denser geometry from simulations as ground truth. This naturally poses a challenge to the robustness against varying optical detector properties. As a case study, we examine the network’s performance on simulation sets

with different scattering lengths λ_s . The superresolution network and subsequent SSCNN reconstruction are both trained on the nominal training sets. We then generate new test datasets with varied scattering lengths and analyze the reconstruction performance on each set.

We find that the network is robust to optical medium uncertainties up to approximately 5%. This tolerance is sufficient to use this method in current experimental settings. For example, IceCube reports to have optical property uncertainties at the level of a few percent [43] and are expected to improve with the IceCube Upgrade. For more details, see Appendix B.

Conclusions and Future Directions. In this Letter, we have introduced a novel technique, superresolution, for enhancing neutrino telescope event reconstruction. Our results have demonstrated significant improvements in the angular reconstruction of ν_μ events in a generic ice-based neutrino telescope. The approach is also generalizable to different event morphologies and detector types, including water-based neutrino telescopes such as the KM3NeT detector currently being deployed in the Mediterranean Sea. Since all neutrino telescope analyses rely on the directional reconstruction of events to pinpoint neutrino sources, the superresolution technique can enhance the discovery potential of these analyses by improving the performance of existing datasets.

Additionally, the use of a VAE to encode finely binned timing distributions into a compressed latent space has proven to be an effective strategy. As a future direction, we aim to explore how to use the VAE to compress neutrino telescope data, enhance and accelerate our reconstructions, and also integrate it into an ML-based event generation model.

As we advance toward constructing larger detectors with larger string spacing, the superresolution method could play a crucial role in maximizing the scientific output of these experiments. Furthermore, as we develop more diverse types of detectors, it becomes increasingly important to implement reconstruction methods that work seamlessly across different experiments.

Acknowledgments. F.J.Y. is supported by the National Science Foundation (NSF) through the NSF AI Institute for Artificial Intelligence and Fundamental Interactions. C.A.A. is supported by the Faculty of Arts and Sciences of Harvard University, the National Science Foundation (NSF), and the NSF AI Institute for Artificial Intelligence and Fundamental Interactions, the Research Corporation for Science Advancement. N.K. is supported by the National Science Foundation (NSF) CAREER Award No. 2239795 and the David and Lucile Packard Foundation.

Data availability. The code used in this Letter is openly available [39].

Appendix A: More Details On The UNet. This appendix is intended to provide additional details on the architecture and training methods of the UNet++ used in this Letter.

One of the main strengths of using a VAE to encode timing information on OM is that the output becomes a fixed-length vector, as defined by the user (the size of the latent dimension). This allows us to effectively compress complex neutrino telescope events into a simple 2D image, where the latent vector is included as part of the features at each pixel. This also enables us to use UNet-type architectures, which has seen much success on 2D images across many scientific disciplines. As mentioned in the main text, we use a latent dimension size of 64.

The features at each OM contains its 3D coordinates, the number of photon hits, and the 64-parameter latent vector containing the timing information. The inputs to UNet++ are masked, while the unmasked events are retained only for loss computation. As is standard for UNet-type architectures, the outputs are also images of the same spatial size. The output has 65 features at each pixel, corresponding to the number of photon hits and the timing latent vector. Mean-squared error (MSE) loss is used for both the prediction of the number of photon hits and the timing latent vector. The timing latent vector prediction is only trained on OMs that saw light in the event. We use the AdamW optimizer with an initial learning rate of 0.001, with periodic drops when the validation loss plateaus. The network is trained for 30 epochs with a batch size of 128 on an NVIDIA A100 GPU.




FIG. 5. Comparison of angular reconstruction performance when varying scattering length. The masked and superresolved resolutions are obtained from the nominal test dataset, while the dash-dotted median lines are inferred from test datasets with altered scattering lengths.

Appendix B: More details on sensitivity to detector properties. As demonstrated in Fig. 5, variations in optical properties at 5% exhibited negligible differences in performance from the nominal test dataset. Increasing the scattering length change to 10% reveals a noticeable decline in reconstruction performance. Interestingly, at lower energies, where reconstruction is more constrained by the lack of light and information, there is minimal impact on performance even with a 10% scattering length change. The performance loss due to the scattering length change is primarily observed at higher energies, where accurate light modeling is more critical.

- [1] C. A. Argüelles, F. Halzen, and N. Kurahashi, From the dawn of neutrino astronomy to a new view of the extreme universe, [arXiv:2405.17623](#).
- [2] M. G. Aartsen *et al.* (IceCube Collaboration), The IceCube neutrino observatory: Instrumentation and online systems, *J. Instrum.* **12**, P03012 (2017); **19**, E05001 (2024).
- [3] M. G. Aartsen *et al.* (IceCube Collaboration), First observation of PeV-energy neutrinos with IceCube, *Phys. Rev. Lett.* **111**, 021103 (2013).
- [4] IceCube Collaboration, Evidence for high-energy extraterrestrial neutrinos at the IceCube detector, *Science* **342**, 1242856 (2013).
- [5] M. G. Aartsen *et al.* (IceCube Collaboration), Neutrino emission from the direction of the blazar TXS 0506 + 056 prior to the IceCube-170922A alert, *Science* **361**, 147 (2018).
- [6] R. Abbasi *et al.* (IceCube Collaboration), Evidence for neutrino emission from the nearby active galaxy NGC 1068, *Science* **378**, 538 (2022).
- [7] R. Abbasi *et al.* (IceCube Collaboration), Observation of high-energy neutrinos from the Galactic plane, *Science* **380**, adc9818 (2023).
- [8] R. Abbasi *et al.*, A convolutional neural network based cascade reconstruction for the IceCube neutrino observatory, *J. Instrum.* **16**, P07041 (2021).
- [9] M. Huennefeld *et al.* (IceCube Collaboration), Combining maximum-likelihood with deep learning for event reconstruction in IceCube, *Proc. Sci. ICRC2021* (2021) 1065 [[arXiv:2107.12110](#)].
- [10] R. Abbasi *et al.* (IceCube Collaboration), Graph neural networks for low-energy event classification & reconstruction in IceCube, *J. Instrum.* **17**, P11003 (2022).
- [11] I. Kharuk, G. Safronov, A. Matseiko, and A. Leonov, Machine learning in Baikal-GVD experiment, *Proc. Sci. ICRC2023* (2023) 1077.
- [12] P. Eller (IceCube Collaboration), Public kaggle competition “IceCube—neutrinos in deep ice”, in *Proceedings of the 38th*

- International Cosmic Ray Conference* (2023), [arXiv:2307.15289](#).
- [13] C. De Sio (KM3NeT Collaboration), Machine learning in KM3NeT, *EPJ Web Conf.* **207**, 05004 (2019).
 - [14] C. Mo, F. Hu, L. Li, and D. Xu (TRIDENT Collaboration), Reconstruction of track-like event in TRIDENT based on a submanifold sparse convolutional network, *Proc. Sci. ICRC2023 (2023)* 1205.
 - [15] S. Reck, D. Guderian, G. Vermariën, and A. Domi (KM3NeT Collaboration), Graph neural networks for reconstruction and classification in KM3NeT, *J. Instrum.* **16**, C10011 (2021).
 - [16] N. Choma, F. Monti, L. Gerhardt, T. Palczewski, Z. Ronaghi, Prabhat, W. Bhimji, M. M. Bronstein, S. R. Klein, and J. Bruna (IceCube Collaboration), Graph neural networks for IceCube signal classification, [arXiv:1809.06166](#).
 - [17] F. J. Yu, N. Kamp, and C. A. Argüelles, Learning efficient representations of neutrino telescope events, [arXiv:2410.13148](#).
 - [18] M. Jin, Y. Hu, and C. A. Argüelles, Two Watts is all you need: Enabling in-detector real-time machine learning for neutrino telescopes via edge computing, *J. Cosmol. Astropart. Phys.* **06** (2024) 026.
 - [19] T. Zhu, M. Jin, and C. A. Argüelles, Comparison of geometrical layouts for next-generation large-volume Cherenkov neutrino telescopes, [arXiv:2407.19010](#).
 - [20] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, Machine learning at the energy and intensity frontiers of particle physics, *Nature (London)* **560**, 41 (2018).
 - [21] A. Li, Z. Fu, L. A. Winslow, C. Grant, H. Song, H. Ozaki, I. Shimizu, and A. Takeuchi, KamNet: An integrated spatio-temporal deep neural network for rare event searches in KamLAND-Zen, *Phys. Rev. C* **107**, 014323 (2023).
 - [22] Z. Imani, T. Wongjirad, and S. Aeron, Score-based diffusion models for generating liquid argon time projection chamber images, *Phys. Rev. D* **109**, 072011 (2024).
 - [23] C. Adams *et al.* (MicroBooNE Collaboration), Deep neural network for pixel-level electromagnetic particle identification in the MicroBooNE liquid argon time projection chamber, *Phys. Rev. D* **99**, 092001 (2019).
 - [24] P. Abratenko *et al.* (MicroBooNE Collaboration), Semantic segmentation with sparse convolutional neural network for event reconstruction in MicroBooNE, *Phys. Rev. D* **103**, 052012 (2021).
 - [25] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle, A convolutional neural network neutrino event classifier, *J. Instrum.* **11**, P09001 (2016).
 - [26] L. Dominé and K. Terao (DeepLearnPhysics Collaboration), Scalable deep convolutional neural networks for sparse, locally dense liquid argon time projection chamber data, *Phys. Rev. D* **102**, 012005 (2020).
 - [27] F. Drielsma, K. Terao, L. Dominé, and D. H. Koh, Scalable, end-to-end, deep-learning-based data reconstruction chain for particle imaging detectors, *Proceedings of the 34th Conference on Neural Information Processing Systems* (2021), [arXiv:2102.01033](#).
 - [28] S. Aiello *et al.* (KM3NeT Collaboration), Sensitivity of the KM3NeT/ARCA neutrino telescope to point-like neutrino sources, *Astropart. Phys.* **111**, 100 (2019).
 - [29] A. D. Avrorin *et al.* (Baikal-GVD Collaboration), Baikal-GVD: Status and prospects, *EPJ Web Conf.* **191**, 01006 (2018).
 - [30] M. Agostini *et al.* (P-ONE Collaboration), The pacific ocean neutrino experiment, *Nat. Astron.* **4**, 913 (2020).
 - [31] M. G. Aartsen *et al.* (IceCube-Gen2 Collaboration), IceCube-Gen2: The window to the extreme Universe, *J. Phys. G* **48**, 060501 (2021).
 - [32] Z. P. Ye *et al.*, A multi-cubic-kilometre neutrino telescope in the western Pacific Ocean, [arXiv:2207.04519](#).
 - [33] T.-Q. Huang, Z. Cao, M. Chen, J. Liu, Z. Wang, X. You, and Y. Qi, Proposal for the high energy neutrino telescope, *Proc. Sci. ICRC2023 (2023)* 1080.
 - [34] Z. Wang, J. Chen, and S. C. Hoi, Deep learning for image superresolution: A survey, *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3365 (2020).
 - [35] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), The Super-Kamiokande detector, *Nucl. Instrum. Methods Phys. Res., Sect. A* **501**, 418 (2003).
 - [36] K. Abe *et al.* (Hyper-Kamiokande Collaboration), Hyper-Kamiokande design report, [arXiv:1805.04163](#).
 - [37] A. Abusleme *et al.* (JUNO Collaboration), JUNO physics and detector, *Prog. Part. Nucl. Phys.* **123**, 103927 (2022).
 - [38] M. Askins *et al.* (Theia Collaboration), THEIA: An advanced optical neutrino detector, *Eur. Phys. J. C* **80**, 416 (2020).
 - [39] F. J. Yu, https://github.com/felixyu7/nt_mlreco (2024).
 - [40] J. Lazar, S. Meighen-Berger, C. Haack, D. Kim, S. Giner, and C. A. Argüelles, Prometheus: An open-source neutrino telescope simulation, *Comput. Phys. Commun.* **304**, 109298 (2024).
 - [41] P. Iakubovskii, Segmentation models PyTorch, https://github.com/qubvel/segmentation_models.pytorch (2019).
 - [42] F. J. Yu, J. Lazar, and C. A. Argüelles, Trigger-level event reconstruction for neutrino telescopes using sparse submanifold convolutional neural networks, *Phys. Rev. D* **108**, 063017 (2023).
 - [43] R. Abbasi *et al.* (IceCube Collaboration), An improved mapping of ice layer undulations for the IceCube neutrino observatory, *Proc. Sci. ICRC2023 (2023)* 975 [[arXiv:2307.13951](#)].