Bridging Particle Physics and AI with the ATHANOR Pipeline

Ziping Ye[1], **Shaokai Yang**[2], **Aksel Hallin**[2], **Joshua Klein**[1]

[1]University of Pennsylvania
zipingye@sas.upenn.edu, jrk@hep.upenn.edu
[2]University of Alberta
aksel.hallin@ualberta.ca, shaokai1@ualberta.ca

Abstract

Neutrinos, characterized by their small mass and weak interactions, are among the most elusive particles within the Standard Model (SM) of particle physics. Nevertheless, gaining a deep understanding of neutrinos is indispensable for unveiling the fundamental physical laws of the Universe, comprehending its evolutionary mechanisms, and potentially providing critical clues for new physics beyond the SM. The SNO+ experiment aims to investigate neutrino properties through high precision detection. SNO+ produces huge amount of data with complex interaction properties and random fluctuation that constitute challenges for precise interaction classification, reconstruction, and simulation. To address these challenges, we have developed ATHANOR, a deep learning analysis pipeline with a modular design that automates data preprocessing, model construction, training, and result visualization. By integrating high-quality simulated data, which follows all known physical laws, and calibration data, ATHANOR can also provide a unique platform for understanding deep learning model uncertainties and enhancing interpretability. In preliminary applications to event position reconstruction in SNO+, ATHANOR outperformed traditional methods in a more efficient way.

Introduction

1.1 The Importance and Challenges of Neutrino Physics

Neutrinos are among the most mysterious fundamental particles in the observable Universe. They possess an extremely small mass and interact very weakly with other matter. This unique property makes neutrinos indispensable in understanding cosmic evolution, supernova explosions, solar energy production, and the processes of nucleosynthesis. In-depth studies of neutrinos not only help reveal the fundamental rules governing the Universe and its evolutionary history but may also guide us toward discovering new physics beyond the SM. But, the weak interaction between neutrinos and electrons/nuclei makes their detection exceedingly difficult, posing a significant challenge in the field of physics research. The SNO+ (Sudbury Neutrino Observatory Plus) experiment (SNO+ 2021a) was established in this context, aiming to measure the properties of neutrinos through precise detection of neutrino events. Located in the SNOLAB underground laboratory in Canada, SNO+ builds upon and upgrades the hardware facilities of the original SNO experiment, which was awarded the 2015 Nobel Prize in Physics. Its goal is to detect neutrino interactions using specially researched and optimized scintillators, thereby providing a high-precision data platform for neutrino physics research. Through meticulous design and deep shielding measures, the experiment effectively reduces radiation backgrounds, offering an ideal environment for exploring this subtle yet crucial particle.

1.2 The Potential of Deep Learning in High-Energy Physics

As a frontier field of artificial intelligence technology, deep learning (DL) has achieved groundbreaking results in areas such as computer vision and natural language processing, thanks to its powerful nonlinear modeling capabilities and advantages in automatic feature extraction. Similar to these fields, data from large-scale high-energy physics (HEP) experiments also exhibit characteristics such as complexity and significant backgrounds levels. By employing deep learning, physicists can automatically extract complex features from raw measurement data, capturing nonlinear relationships between the data, thereby significantly enhancing the accuracy and efficiency of data analysis. For instance, models based on self-attention mechanisms can learn the spatiotemporal distribution information of photomultiplier tube (PMT) hit data in SNO+, which has the potential to enable more accurate reconstruction of the neutrino events. This not only helps physicists better understand the various particle interactions occurring within the detector but also provides great potential for exploring new physical phenomena.

1.3 Bidirectional Advancement Between HEP and DL

There is an increasingly significant bidirectional advancement between HEP and deep learning. HEP experiments provide unique data and testing platforms for the development of deep learning, meanwhile deep learning offers advanced data processing tools for HEP experiments. On one hand, HEP experiments can use high-precision simulation data based on physical principles to effectively train the DL models and validate their predictive results of DL models, therefore enhancing the models' reliability and interpretability. For instance, by embedding prior physical knowledge into the model and incorporating constraints such as energy and momentum conservation into the loss function, the model's predictions can better align with known physical laws. On the other hand, deep learning models, with their powerful nonlinear modeling capabilities, excel at handling huge amounts of data with complex non-linear space-time relationship and random noisy fluctuations produced by HEP experiments. This bidirectional advancement not only propels the development of HEP experiments but also opens new directions in the evolution of deep learning.

1.4 Main Contributions

The main contributions of this paper include: First, we have constructed the ATHANOR analysis pipeline, achieving full automation from data preprocessing and model construction and training to result analysis and visualization. This tool is designed for HEP researchers who are not specialists in deep learning, enabling them to easily utilize and modify deep learning frameworks for data analysis, thus promoting the integrated development of deep learning and high-energy physics. Second, we have embedded physics-based constraints into the framework's design and model training process, improving the model's physical consistency and predictive accuracy. Finally, we have applied the ATHANOR pipeline to the interaction position reconstruction task in SNO+, achieving preliminary results that validate the functionality of the pipeline.

1.5 Paper Structure

The organization of this paper is as follows: Section 2 introduces the SNO+ detector and its data characteristics, analyzing the main challenges in data analysis; Section 3 details the construction of the ATHANOR deep learning analysis pipeline; Section 4 presents the application and experimental results of ATHANOR in the SNO+ interaction reconstruction task; and Section 5 discusses the limitations of the current model and outlines future research directions.

The SNO+ Detector and Data Characteristics

2.1 Overview of the SNO+ Detector

The SNO+ detector, Figure 1, is located in the SNO-LAB deep underground laboratory, buried approximately two kilometers beneath the Earth's surface. This depth effectively shields cosmic rays and environmental radiation, providing ideal conditions for precise neutrino measurements. The core of the SNO+ detector consists of a spherical acrylic vessel with a diameter of 12 meters, filled with about 780 tonnes of high-purity liquid scintillator (linear alkylbenzene, LAB). To enhance the optical properties and light yield of the scintillator, 2 tonnes of dissolved PPO (2,5-diphenyloxazole) and a small amount of Bis-MSB (1,4-bis(2-methylstyryl)benzene) are added. The inclusion of PPO and Bis-MSB significantly increases the numbers of photons that can be detected, thereby enhancing the experiment's signal response capability (SNO+ 2021b). This enhanced sensitivity enables the SNO+ experiment to detect particle interactions, especially neutrinos, with greater precision, ushering in a new phase of data collection. In the future, SNO+ plans to add ^{130}Te to the liquid scintillator to search for neutrinoless double-beta decay, which, if observed, could prove that neutrinos are Majorana particles and is unambiguous new physics signal (SNO+ 2022).

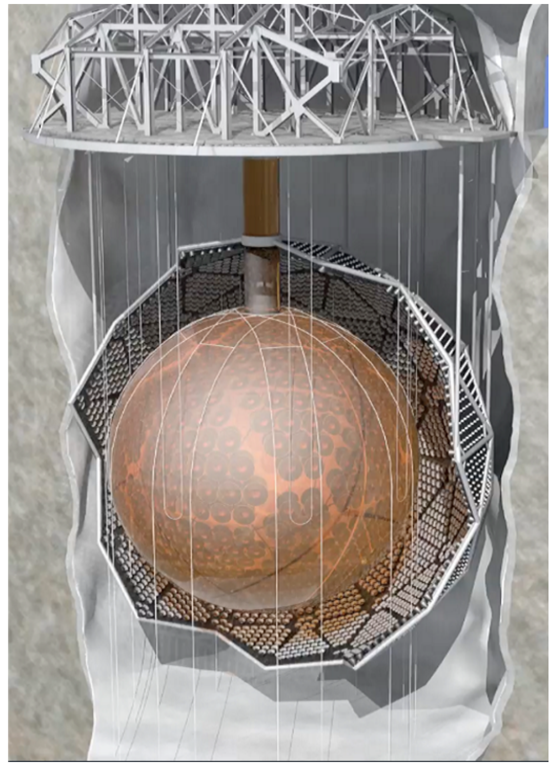

Figure 1: Surrounding SNO+ detector are 9,362 high-sensitivity, 8-inch inward-facing PMTs, distributed and precisely arranged on a geodesic sphere with a radius of 8.89 meters, forming a well-covered optical detection array (SNO+ 2021a). The spatial positions of each PMT have been precisely measured, providing crucial geometric information for the spatial reconstruction of interaction events. The outer layer of the detector is enveloped by approximately 7,000 tonnes of ultra-pure water, which not only further shields against environmental radioactive backgrounds but also serves as a medium for detecting Cherenkov radiation from high-energy charged particles. The detector is designed to maximize sensitivity to neutrino interactions and enhance reconstruction accuracy.

2.2 Working Principle and Scientific Goals

SNO+ detects neutrinos through their interactions with nuclei or electrons within the liquid scintillator. When a neutrino interacts with the medium, it deposits energy that excites the scintillator molecules, which then de-excite and emit photons following specific time profiles. The photons propagate through the medium (some of the photons would be absorbed, scattered, and/or reflected during the propagation), and eventually some photons would be detected by the PMTs. The photon propagation process would generate random fluctuation in the data. By analyzing the time, charge, and spatial distribution of the photons collected by the PMTs, it is possible to reconstruct physical quantities such as the vertex position, energy, and direction of neutrino events. These enable SNO+ to investigate a range of important physics topics, including neutrinoless double-beta decay, solar neutrino oscillations, geo-neutrinos, supernova neutrinos, and other particle decay processes. These studies are crucial for gaining a deeper understanding of neutrino mass, oscillation mechanisms, and exploring dark matter (SNO+ 2023).

2.3 Data Characteristics and Analytical Challenges

In SNO+ and many other HEP experiments, the measured data consists of long sequences of hit information, often containing noise, random fluctuations, and complex internal correlations. DL models, particularly those based on self-attention mechanisms (**?**) or selective structured state-space models (Gu and Dao 2024) , excel at handling long sequences by capturing global dependencies and learning intricate non-linear relationships. With training on large datasets, these models adapt to various experimental conditions, thereby achieving strong generalization capabilities. Furthermore, deep learning benefits from hardware accelerations, such as GPUs, which significantly enhance computational efficiency. These features make deep learning well-suited for the precise, efficient, and automated analysis of HEP data.

ATHANOR Deep Learning Analysis Pipeline

To address the complex data processing challenges presented by the SNO+ experiment, we developed ATHANOR, a DL analysis pipeline built upon PyTorch. Designed with efficiency, modularity, and scalability in mind, ATHANOR enables HEP researchers, including those without extensive deep learning experience, to process data effectively, construct models, and ensure reproducible results. The pipeline is structured as a series of Python packages, each corresponding to a specific component of the analysis workflow, and within each package, multiple modules implement specific functionalities. Figure 2 illustrates the workflow of the ATHANOR as applied in SNO+.

The Data Handler package preprocesses raw PMT data. It includes modules for data cleaning, normalization, and format conversion, transforming complex raw inputs into structured data suitable for deep learning models. By removing noise and outliers and incorporating precise PMT spatial positions, the Data Handler ensures data integrity and enables models to capture essential spatial correlations inherent in the detector geometry.

The Architecture Library provides a collection of advanced deep learning architectures implemented in PyTorch, including self-attention mechanisms, selective structured state-space models, and graph transformer networks (Yun et al. 2020). These architectures are chosen for their ability to handle the high-dimensional, complex, and noisy data of the SNO+ experiment. Some are directly available in PyTorch, while others are custom-defined in ATHANOR

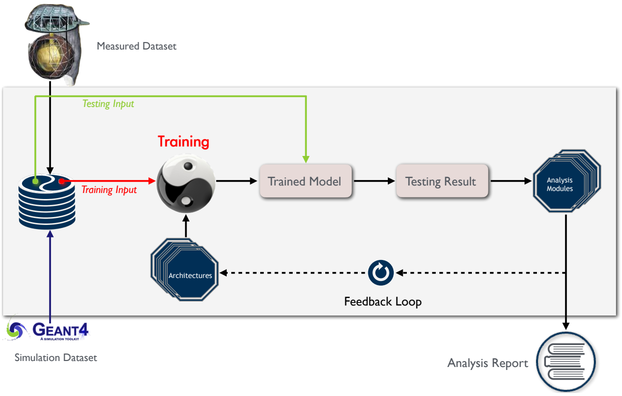

Figure 2: Measured data from the detector and simulation data generated by the GEANT4 (Agostinelli et al. 2003) based RAT software are preprocessed by the Data Handler package to prepare them for model training and evaluation. The preprocessed data is then divided into Training and Testing Datasets. Researchers select and adapt suitable architectures from ATHANOR's modular library. After training, the model is evaluated using the Testing Dataset, and the results are passed to the Analysis Modules for detailed performance assessment. If required, the Feedback Loop facilitates the redesign and optimization of the model until the desired performance is reached. Upon satisfactory evaluation, an Analysis Report is generated, ensuring transparency and reproducibility.

with specific classes and equations. Defining these architectures within ATHANOR allows researchers to easily modify them to better adapt to specific data analysis needs, offering a high degree of flexibility in model development.

The Model Constructor and Trainer package facilitates model building and training. Physical constraints can be integrated directly into the modeling process to enhance predictive accuracy and ensure consistency with known physical laws. The loss function balances data fidelity with adherence to physical principles:

$$\mathcal{L} = \mathcal{L}_{\text{data}} + \lambda \mathcal{L}_{\text{physics}}$$

where $\mathcal{L}_{\text{data}}$ is the data fidelity term, $\mathcal{L}_{\text{physics}}$ represents physical constraints, and λ balances their influence. Leveraging PyTorch's capabilities, this package utilizes GPU acceleration and parallel computing to efficiently handle large-scale datasets. Optimization strategies such as learning rate scheduling, regularization, and early stopping are employed to prevent overfitting and improve generalization.

After training, the Model Inference and Analysis package applies the trained models to new data for prediction and evaluates performance using metrics like mean squared error and mean absolute error. Visualization tools generate loss curves, metric trends, and three-dimensional event reconstructions, aiding in interpreting model performance and identifying areas for improvement.

ATHANOR offers both a command-line interface (CLI) and a graphical user interface (GUI), providing user-friendly access to the pipeline's functionalities. Researchers can perform data processing, model training, inference, and analysis without deep programming expertise. Configuration files manage essential parameters, including model architectures, training hyperparameters, and data paths, ensuring experimental reproducibility and facilitating easy adjustments to accommodate different analysis requirements.

A key feature of ATHANOR is the integration of physical constraints within the architectures and loss functions. By embedding physical laws directly into the modeling process, we ensure that predictions are not only data-driven but also physically consistent, enhancing reliability and interpretability. This approach allows models to respect known physical principles, which is crucial in HEP applications and contributes to a deeper understanding of neutrino properties.

To lower the technical barrier for users unfamiliar with deep learning, we have fine-tuned a Llama-based large language model to provide assistance within ATHANOR. This feature helps users learn how to operate the pipeline, troubleshoot issues during model training, and obtain answers to common questions, further enhancing the pipeline's accessibility to the HEP community.

Application of ATHANOR in SNO+

To validate ATHANOR, we employed it to develop a regression model consisting of a self-attention mechanism (**?**) combined with a neural network, using Monte Carlo simulation data from the SNO+ experiment for interaction position reconstruction. Position reconstruction is a challenging yet important task in experimental data analysis, it is the basis of all the other reconstruction tasks and can be used to reject backgrounds. The simulation data was generated using the RAT framework, accurately modeling neutrino interactions and detector responses to simulate data as exactly the same with real experimental data as possible. The figure-of-merit for evaluating the interaction position reconstruction performance is the distance between the predicted (reconstructed) position and the mc-true (simulated) position, or called the reconstruction error. The smaller the reconstruction error, the better the performance. During model training, we incrementally expanded the dataset to assess ATHANOR's performance scalability, ensuring training and test sets remained independent to prevent overfitting. Data augmentation further enhanced model generalization across different experimental conditions. The training process utilized techniques such as learning rate scheduling and early stopping to optimize performance.

ATHANOR outperformed traditional methods like maximum likelihood estimation (MLE) in event position reconstruction. While MLE tends to perform well in certain energy regions but struggles in others, ATHANOR consistently delivered superior results across all energy ranges (0 to 5 MeV in our study). It achieved high accuracy throughout the entire detector volume, including challenging edge regions where MLE often falters. Figure 3 illustrates the advantage of deep learning in capturing complex, non-linear relationships inherent in HEP data analysis. Moreover, ATHANOR's GPU acceleration and parallel computing capabilities provided a faster reconstruction speed compared to MLE. Comparative analysis showed that ATHANOR outperformed traditional MLE methods in both accuracy and efficiency. Similar performance has been achieved with other DL studies by SNO+ collaborators (Mark and Cal 2024). Leveraging DL's capabilities and GPU acceleration, ATHANOR efficiently processed complex, high-dimensional data, automatically learning intricate patterns to achieve superior reconstruction.

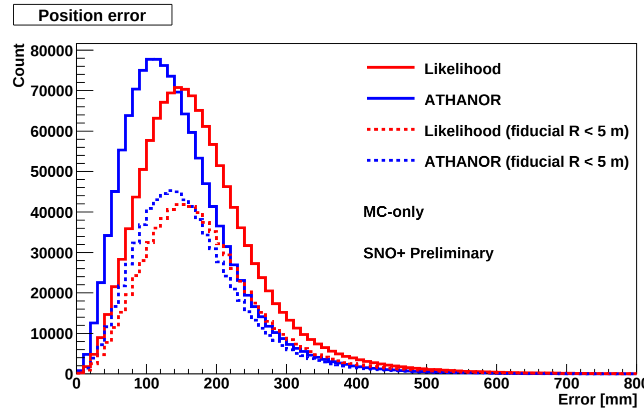

Figure 3: Comparing the position reconstruction performances of ATHANOR and MLE. MLE is a type of classic and commonly used method for interaction reconstruction in HEP. ATHANOR achieves smaller reconstruction error than MLE does. For all events in the detector volume (solid lines) and the events in fiducial core volume (dotted lines), ATHANOR has better performance.

Discussion

We have introduced ATHANOR, a deep learning analysis pipeline that effectively addresses the challenges of high-dimensional, complex data in HEP. Applied to interaction position reconstruction in SNO+, ATHANOR outperforms traditional methods, demonstrating deep learning's great potential in this domain. Its modular design ensures adaptability to various experimental conditions and data characteristics, making it a valuable tool for physicists.

Future work will explore advanced model architectures better suited to HEP data, such as Graph Transformer Networks (GTNs) and Physics-Informed (PI) models, to further enhance performance and interpretability. Incorporating additional physical constraints and developing physics-guided interpretability methods will address the "black box" nature of deep learning models, providing deeper insights into their internal mechanisms and fostering greater trust in their predictions. Expanding training data diversity and employing techniques like self-supervised learning will enhance generalization capabilities, ensuring robustness across a wider range of experimental scenarios.

References

Agostinelli, S.; Allison, J.; Amako, K.; Apostolakis, J.; Araujo, H.; Arce, P.; Asai, M.; Axen, D.; Banerjee, S.; Barrand, G.; Behner, F.; Bellagamba, L.; Boudreau, J.; Broglia, L.; Brunengo, A.; Burkhardt, H.; Chauvie, S.; Chuma, J.; Chytracek, R.; Cooperman, G.; Cosmo, G.; Degtyarenko, P.; Dell'Acqua, A.; Depaola, G.; Dietrich, D.; Enami, R.; Feliciello, A.; Ferguson, C.; Fesefeldt, H.; Folger, G.; Foppiano, F.; Forti, A.; Garelli, S.; Giani, S.; Giannitrapani, R.; Gibin, D.; Gómez Cadenas, J.; González, I.; Gracia Abril, G.; Greeniaus, G.; Greiner, W.; Grichine, V.; Grossheim, A.; Guatelli, S.; Gumplinger, P.; Hamatsu, R.; Hashimoto, K.; Hasui, H.; Heikkinen, A.; Howard, A.; Ivanchenko, V.; Johnson, A.; Jones, F.; Kallenbach, J.; Kanaya, N.; Kawabata, M.; Kawabata, Y.; Kawaguti, M.; Kelner, S.; Kent, P.; Kimura, A.; Kodama, T.; Kokoulin, R.; Kossov, M.; Kurashige, H.; Lamanna, E.; Lampén, T.; Lara, V.; Lefebure, V.; Lei, F.; Liendl, M.; Lockman, W.; Longo, F.; Magni, S.; Maire, M.; Medernach, E.; Minamimoto, K.; Mora de Freitas, P.; Morita, Y.; Murakami, K.; Nagamatu, M.; Nartallo, R.; Nieminen, P.; Nishimura, T.; Ohtsubo, K.; Okamura, M.; O'Neale, S.; Oohata, Y.; Paech, K.; Perl, J.; Pfeiffer, A.; Pia, M.; Ranjard, F.; Rybin, A.; Sadilov, S.; Di Salvo, E.; Santin, G.; Sasaki, T.; Savvas, N.; Sawada, Y.; Scherer, S.; Sei, S.; Sirotenko, V.; Smith, D.; Starkov, N.; Stoecker, H.; Sulkimo, J.; Takahata, M.; Tanaka, S.; Tcherniaev, E.; Safai Tehrani, E.; Tropeano, M.; Truscott, P.; Uno, H.; Urban, L.; Urban, P.; Verderi, M.; Walkden, A.; Wander, W.; Weber, H.; Wellisch, J.; Wenaus, T.; Williams, D.; Wright, D.; Yamada, T.; Yoshida, H.; and Zschiesche, D. 2003. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3): 250–303.

Gu, A.; and Dao, T. 2024. Mamba: Linear-Time Sequence Modeling with Selective State Spaces. In *First Conference on Language Modeling*.

Mark, A.; and Cal, H. 2024. Machine learning for fast event reconstruction in the SNO+ scintillator phase.

SNO+. 2021a. The SNO+ experiment. *Journal of Instrumentation*, 16(08): P08059.

SNO+. 2021b. Optical calibration of the SNO+ detector in the water phase with deployed sources. *Journal of Instrumentation*, 16(10): P10021.

SNO+. 2022. Improved search for invisible modes of nucleon decay in water with the SNO+ detector. *Physical Review D*, 105(11): 112012.

SNO+. 2023. Evidence of Antineutrinos from Distant Reactors Using Pure Water at SNO+. *Physical Review Letters*, 130(9): 091801.

Yun, S.; Jeong, M.; Kim, R.; Kang, J.; and Kim, H. J. 2020. Graph Transformer Networks. arXiv:1911.06455.