Neutrino direction and flavor reconstruction from radio detector data using deep convolutional neural networks

Sigfrid Stjärnholm,[a,*] Oscar Ericsson[a,*] and Christian Glaser[a]

[a]*Uppsala University, Department of Physics and Astronomy,*
Lägerhyddsvägen 1, Uppsala, Sweden
E-mail: sigfrid.stjarnholm.2141@student.uu.se,
oscar.ericsson.7318@student.uu.se, christian.glaser@physics.uu.se

With the construction of RNO-G and plans for IceCube-Gen2, neutrino astronomy at EeV energies is at the horizon for the next years. Here, we determine the neutrino pointing capabilities and explore the sensitivity to the neutrino flavor for an array of shallow radio detector stations. The usage of deep learning for event reconstruction is enabled through recent advances in simulation codes that allow the simulation of realistic training data sets. A large data set of expected radio signals for a broad range of neutrino energies between 100 PeV and 10 EeV is simulated using NuRadioMC. A deep neural network is trained on this low-level data and we find a direction resolution of a few degrees for all triggered events. We present the model architecture, how we optimized the model, and how robust the model is against systematic uncertainties. Furthermore, we explore the capabilities of a radio neutrino detector to determine the flavor id.

---

*Presenter

PoS(ICRC2021)1055

1. Introduction

The radio technique is a cost-efficient way to obtain sensitivity to ultra-high energy (UHE) neutrinos around energies of 10^{18} eV. A sparse array of radio detector stations installed close to the surface of polar ice sheets allows instrumentation of huge volumes. The particle cascades initiated by neutrino interactions in the ice generate radio emission via the Askaryan effect. All neutrino interactions initiate a so-called hadronic shower stemming from the fragmentation of the target nucleon. In addition, in electron neutrino charged-current interactions (ν_e-CC), the outgoing electron (positron) initiates an additional electromagnetic shower along side the hadronic shower. In this work we will exploit the differences between ν_e-CC and non-ν_e-CC interactions to obtain flavor sensitivity.

We note that muon and tau leptons created by CC interactions of the respective neutrino also initiate secondary showers due to stochastic energy losses or decay. These additional showers are typically significantly displaced from the initial interaction but their signals might be picked up by other radio detector stations. A detection of the initial and secondary interaction provide another opportunity to measure the neutrino flavor which is studied in [1, 2].

With the construction of the discovery scale RNO-G detector in Greenland [3], the first observation of an UHE neutrino might be possible in the next years. At the same time, an order-of-magnitude more sensitive radio detector is currently being designed as part of the proposed IceCube-Gen2 [4]. These experimental developments make the development of reconstruction algorithms a priority. On the one hand, to be able to analyze the first detected neutrino, and on the other hand, to optimize future detector designs in terms of reconstruction capabilities. In this contribution we will use deep-learning techniques to extract the neutrino direction and flavor directly from raw experimental data. It constitutes the first end-to-end reconstruction from radio detector data.

The radio data is particularly suited for application of deep-learning methods. All information on the neutrino is compressed into a few nano-seconds long radio flash and it is very challenging to design algorithms 'by-hand' that extract all the information contained in the radio pulse. This is also why no end-to-end reconstruction algorithm has been developed thus far, although important progress has been made over the last years [5–8]. In contrast, deep-learning techniques naturally consider all available information.

One of the largest challenges of using deep-learning is that large and realistic training datasets are required. Their generation was enabled through the development of the NuRadioMC code over the last years [9]. NuRadioMC allows for a fast simulation of arbitrary detector layouts. A new level of precision in the calculation of the Askaryan radiation has been reached by combining a library of microscopically simulated shower profiles with the ARZ model [10] that calculates the Askaryan radiation from the shower profile with a precision of 3% compared to a microscopic simulation. In particular, it enables a proper treatment of the LPM effect that reduces that cross section of high-energy electrons (positrons) which manifests in electromagnetic showers as an overall elongation of the shower profile and a charge excess profile characteristic of multiple sub-showers. We will use this signature later to obtain flavor sensitivity as all non-ν_e-CC interactions are not impacted by the LPM effect.
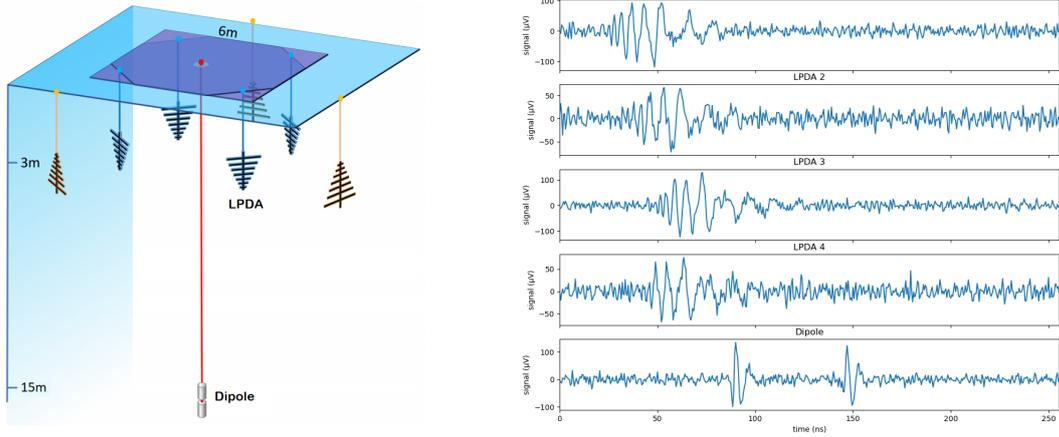
Figure 1: (left) Schematic image of the shallow detector station design. (right) Example neutrino event. The first four traces are from the LPDA antennas, while the last trace is from the dipole antenna. For each trace, 512 samples are taken with a time interval of $\Delta t = 0.5$ ns. The dipole channel clearly shows the DnR signature of a time delayed direct and reflected-off-the-surface signal path.

2. Simulation of Radio Detector Data

In this work, we focus on a shallow station design as depicted in Fig. 1 left. This station design is based on ARIANNA experience [11], foreseen for the ARIANNA-200 detector [12] as well as part of the IceCube-Gen2 detector and, with slight variations, RNO-G [3, 4]. For the direction and flavor reconstruction, the method presented here will utilize the 4 downward pointing LPDA antennas as well as the dipole antenna. In Fig. 1 right, an example event is presented with a very high signal-to-noise ratio (SNR) that is not representative for the data set.

The simulation of the radio emission from neutrino events was done using the NuRadioMC code [9]. The framework simulates the neutrino interaction, the production of the Askaryan emission, the propagation of the signal in ice, and the detector response. This yields a complete and realistic model for what a neutrino interaction will look like in a radio detector. Training data sets with a few million of triggered events were created with neutrino energies ranging from 10^{17} eV to 10^{19} eV, random incoming directions, random interaction vertices in a cylindrical volume around the detector station located at the South Pole, i.e., the site of the proposed IceCube-Gen2 detector. The trigger condition is a majority logic trigger on the four downward facing LPDAs where the sensitivity was optimized by restricting the bandwidth in the trigger [13], and where the trigger threshold was adjusted to give a low-level trigger rate of 100 Hz on thermal noise[1]. This can lead to signal amplitudes of as low as 2 times the RMS noise over the full bandwidth of 80 MHz to 800 MHz.

Three data sets with different Askaryan emission models and interaction channels were created. The first data set (*Alvarez2009 had*) with 8.3×10^6 events contains only hadronic showers (i.e. non-ν_e-CC neutrino interactions) and the Askaryan emission was calculated using a frequency domain parameteriztion from [15, 16]. The second data set (*ARZ had*) with 4.1×10^6 events also contains

---

[1]The low level trigger rate of 100 Hz can be reduced to manageable mHz rates by a real-time thermal noise rejection on the radio detector station [14]

only hadronic showers but this time the Askaryan emission is calculated based on a shower library and the ARZ model as described earlier. The third data set (*ARZ em + had*) with 5×10^6 events also uses the ARZ model but contains only ν_e-CC interactions that produce both a hadronic and electromagnetic shower. The trigger probability increases quickly with neutrino energy and an effort was made to simulate more low energy neutrinos to obtain a uniform energy distribution in $\log_{10}(E)$ of triggered events, however, as we see later, the event distribution is still biased towards higher energies.

3. Deep learning architecture

We use a neural network that is inspired by the VGG model [17] and consists of *convolutional blocks*, followed by *dense layers*. We decided to use a convolutional neural network to identify and characterize the signal pulses that are invariant in time (cf. Fig. 1). Because the relative time differences of the signals between antennas carry information, we apply the convolution filters independently to each antenna, i.e., the input shape is (5, 512, 1) where 5 are the 5 antennas, 512 is the number of samples, and the last dimension is unity as we only available information from each antennas is the voltage amplitude (unlike three color channels in the case of images). We use four convolutional blocks that each have three convolutional layers with padding enabled, a filter size of (1,5) and ReLU activation. In each convolutional block, the number of filters is doubled from 32 to 64 to 128. Each convolutional block ends with an average pooling layer of size 4, which reduces the size of the trace by a factor of 4 in order to decrease the dimensionality of the traces. The network is flattened and complemented by several dense layers.

The models developed for flavor and direction reconstruction are very similar, but it was found that two slight variations proved to reconstruct each property with better accuracy. The network used for the directional reconstruction uses 6 dense layers with a decreasing number of nodes of 2048, 2048, 1024, 512, 128. The final output is a 3-dimensional Cartesian coordinates pointing in the neutrino direction. The network used for the flavor reconstruction uses just two dense layers with 512 nodes each. The final output are two nodes with a softmax activation suitable for this binary classification task. All networks are implemented using Keras/tensorflow. The optimizer was Adam with a learning rate of 0.0001. A Nvidia Quadro RTX 6000 GPU was used to perform the training.

4. Flavor Reconstruction

A radio detector is not directly sensitive to the neutrino flavor but differences in the event signatures of ν_e-CC interactions compared to all other interaction channels can be used to measure the flavor for a subset of the events. Interactions other than ν_e-CC induce a hadronic particle shower stemming from the fragmentation of the target nucleon. For ν_e-CC interactions, the outgoing electron (positron) initiates an additional electromagnetic shower that is effected by the LPM effect along side the hadronic shower.

A convolutional neural network capable of classifying interactions based on said features was developed. The only information available to the network are the voltage measurements by five antennas (as depicted in Fig. 1), and of course the true event category during training. The total

available dataset consisted of $8.0 \cdot 10^6$ simulated events where half of them are ν_e-CC interactions. The other half are interactions the only produced a hadronic showers. The neural network was trained using 95% of the data set. A fraction of 2.5% was used for validation and another 2.5% for an independent test data set.

The main results are presented in the form of a confusion matrix of class-separated prediction accuracies in Fig. 2 left. We find that almost 90% of non-ν_e-CC interactions are classified correctly whereas 66% of ν_e-CC are classified correctly. The largest confusion stems from ν_e-CC interactions that were classified as non-ν_e-CC. More insight can be gained by inspecting the accuracy as a function of signal-to-noise ratio[2] in four separate energy intervals of size $\Delta \log E_\nu = 0.5$, which is presented in Fig. 2 right. For ν_e-CC events we find that the accuracy increases with SNR and energy, whereas the accuracy of non-ν_e-CC events shows almost no dependence on the neutrino energy nor the SNR.

The discrepancy in accuracy between the two flavor classes is largely attributed to the LPM effect. ν_e-CC interactions with neutrino energies in the upper end of the considered interval produce Askaryan emissions which show clear signs of LPM elongation of the electromagnetic shower. At the lower end of the neutrino energy interval (close to 10^{17} eV), LPM elongation of the electromagnetic showers is less pronounced. In fact, the features of radio pulses produced by ν_e-CC interactions at lower neutrino energies are more reminiscent of those the network learned to associate with hadronic showers, i.e., non-ν_e-CC events. This is evident from the sub 50% accuracy on ν_e-CC events for nearly all signal-to-noise ratios in the lowest energy bin of Fig. 2 right. As previously mentioned, hadronic shower events are not subject to the LPM effect to the same extent. The shape of the measured radio signals from non-ν_e-CC events thus remain relatively unchanged across the entire considered energy range.

5. Direction Reconstruction

We used a similar approach to determine the neutrino direction from the measured voltage traces. The neutrino direction has a complex dependence on the signal arrival direction, the signal polarization, the angle under which the shower is observed and the distance to the neutrino interaction vertex. Developing a traditional reconstruction algorithm is challenging but first promising results have been obtained [7] by employing the forward-folding technique [6] but so far the reconstruction is limited to hadronic showers. Here, we present the first end-to-end reconstruction where a CNN learned the complex dependencies just based on the training data.

We quantify the goodness of the reconstruction by the space angle difference between the predicted direction and the true direction, denoted here as $\Delta\Psi$. The objective function during the training of the CNN is the mean absolute error (MAE) of $\Delta\Psi$. We chose MAE over the typically used mean squared error (MSE) to reduce the influence of a rare number of events with poor reconstruction. We are interested to obtain a good reconstruction of the majority of the events and plan to develop quality cuts to filter out events with insufficient information for a direction reconstruction. We train the CNN described in Sec. 3 on 75% of the data set and use 20% and 5%

---

[2]The signal-to-noise ratio is defined as the maximum signal amplitude in any of the five antennas after noise was added divided by the noise RMS.
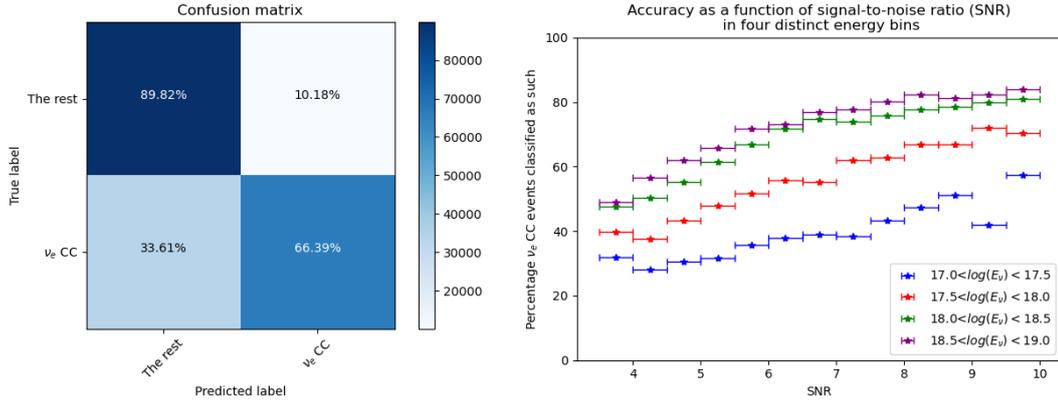
Figure 2: (left) Confusion matrices showing the category-separated performance of the final network on a test set consisting of 200,000 events split evenly between the categories. The main diagonal corresponds to accurate predictions as can be inferred from the axis labels. The color bar is scaled according to absolute number of events. (right) Predictive accuracy of the model on ν_e- CC events as a function of signal-to-noise ratio (SNR) within the intervals marked by the horizontal bars. The data set is split into four equidistant energy bins in $\log_{10}(E)$. The predictions were made on an independent test dataset unused during training and validation.

for validation and test respectively. We perform the training and evaluation separately on the three data sets.

The metric used for assessing the performance of the model is the 68 % quantile, denoted by σ_{68}. This quantity will be a stable estimator that is also immune to rare outlier events with bad reconstruction resolution. We present first the results for the data set of hadronic showers only that were simulated using the ARZ emission model in Fig. 3. The figure also shows the value of σ_{68} which is 4.1°. The expected distribution of this histogram is a Rayleigh distribution, given that the Cartesian coordinates predicted by the model are normally distributed. The figure shows that a Rayleigh distribution does not correspond well to the given distribution, indicated by the tail of the histogram.

To gain more insides into what determines the accuracy in the angular reconstruction, the dependence of σ_{68} is shown as a function of the SNR and the neutrino energy in Fig. 4. A clear trend can be seen that the resolution improves with increasing SNR which comes at no surprise. More interestingly is the dependence on neutrino energy where an improved resolution with increasing neutrino energy is observed. This behaviour can partly be due to a correlation with neutrino energy and SNR. However, this correlation is not strong enough to explain the energy dependence. Also at higher energies, many events are still close to the trigger threshold because their interaction vertex is further away which compensates for the larger emitted signal strengths. The more likely explanation for the energy dependence is an unequal distribution of the training data set. As also shown in Fig. 4, the number of events also increases quickly with energy and we see an anti-correlation between resolution and number of events per bin. An improvement in the resolution of low energy events had little impact on the MSA due to the small event counts in that bin. Our hypothesis is that we can increase the performance at low energies significantly by providing more training data in this energy regime. Also other strategies such as weighting of
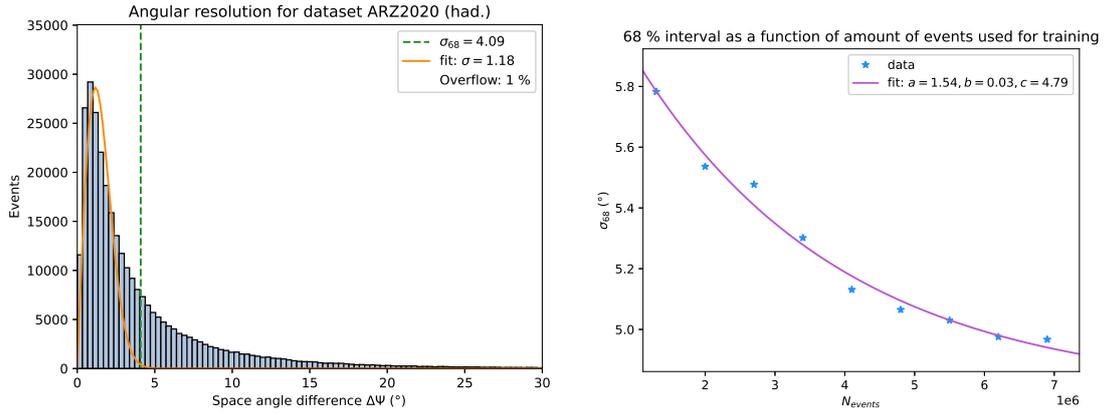
Figure 3: (left) Space angle difference histogram of *ARZ had* data set. The curve is fitted to a Rayleigh distribution. (right) The plot shows the 68 % interval value σ_{68} as a function of the amount of events used for training for the *Alvarez2009 had* data set.
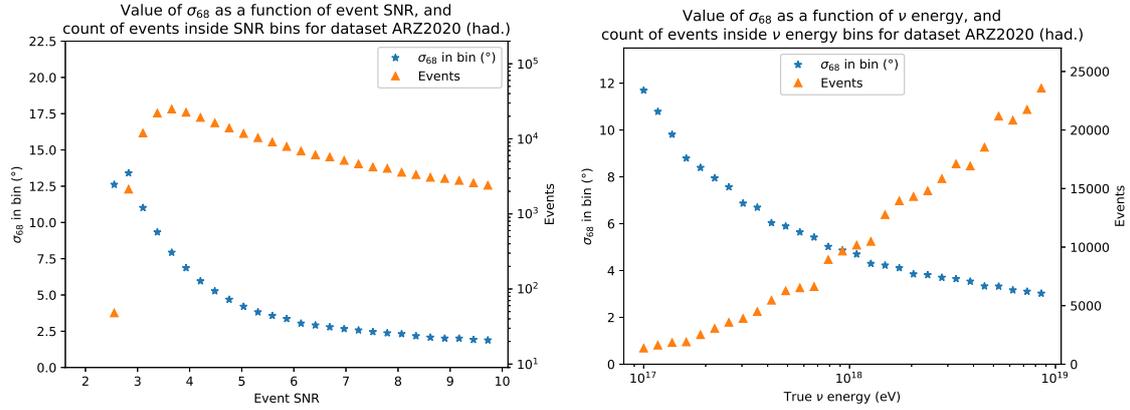
Figure 4: Values of σ_{68} in binned neutrino energy intervals (left), and binned event SNR intervals (right). The count of events in each bin is also shown as orange triangles and need to be read off on the right y-axis.

events might already help and will be explored in the future.

We also analyzed the two other data sets and found that we also obtain a good resolution for the ν_e-CC data set. At neutrino energies below 3×10^{17} eV the performance is similar to the non-ν_e-CC data set with the ARZ Askaryan model and at higher energies the angular resolution is about one degree worse. This result underlines the strengths of a deep-learning approach: The network learned to more complex dependencies in LPM effected showers and a resolution almost as good as for the simpler hadronic case was obtained. Interestingly, the performance of the non-ν_e-CC data set with the simpler *Alvarez2009* Askaryan model lies between the two other data sets.

When training the neural network, the amount of training events used can play a large role in the performance of the model. In Fig. 3 right, the dependence of σ_{68} with respect to the amount of training events used is presented using the *Alvarez2009 had* data set. The graph clearly shows an increase in the performance as the amount of training events used is increased. It was empirically found that a function of the following form $\sigma_{68} = ae^{-bN_{\text{events}}} + c$ can describe the relation between σ_{68} and the amount of training data. Assuming that this function resembles the true dependence, a

7

fit to the data can be used to estimate the performance on an even larger dataset.

6. Conclusions and Outlook

We presented the first end-to-end reconstruction of the direction and flavor of neutrinos measured with a shallow radio detector. We trained and optimized a deep neural network to extract the neutrino direction and flavor directly from raw data. We find an angular resolution of a few degrees with a strong energy dependence that is likely due to an unequal energy distribution of the training data set. A comparison with a traditional reconstruction technique (that was restricted to hadronic showers only) found an angular resolution of $3°$ with little energy dependence [7] which indicates that the deep-learning reconstruction can be improved further.

The full advantage of using deep learning is visible in a good angular resolution also for ν_e-CC interactions where the LPM effect leads to significant variations in the shower development and the resulting radio signals. Traditional reconstruction techniques either ignored this interaction channel or found a significant worsening in angular resolution. A deep neural network is capable of learning these complex dependencies from the training data set and to still provide a good angular resolution of a few degrees.

We exploited the differences between ν_e-CC and non-ν_e-CC interactions to obtain flavor sensitivity. The network was able to identify 90% of non-ν_e-CC interactions correctly. The accuracy of ν_e-CC identification is energy and SNR dependent and also reaches up to 90% for large enough SNR and high neutrino energies where the LPM effect is strong enough.

Acknowledgements

References

[1] D. García-Fernández, C. Glaser, and A. Nelles *Phys. Rev. D* **102** (2020) 083011.

[2] C. Glaser and others *PoS* **ICRC2021** (these proceedings) 1231.

[3] **RNO-G** Collaboration, J. A. Aguilar *et al. JINST 16 P03025* (2021) .

[4] S. Hallmann for the IceCube-Gen2 collaboration *PoS* **ICRC2021** (these proceedings) 1183.

[5] **ARIANNA** Collaboration, A. Anker *et al. JINST* **15** no. 09, (2020) P09039.

[6] C. Glaser *et al. European Physical Journal* **C79** no. 6, (2019) 464.

[7] S. Barwick for the ARIANNA collaboration *PoS* **ICRC2021** (these proceedings) 1151.

[8] C. Welling, P. Frank, T. A. Enßlin, and A. Nelles *JCAP* **04** (2021) 071.

[9] C. Glaser *et al. European Physical Journal* **C80** (2020) 77.

[10] J. Alvarez-Muñiz *et al. Physical Review D* **101** (2020) 083005.

[11] **ARIANNA** Collaboration, A. Anker *et al. JCAP* **03** (2020) 053.

[12] S. Barwick for the ARIANNA collaboration *PoS* **ICRC2021** (these proceedings) 1067.

[13] C. Glaser and S. W. Barwick *JINST* **16** no. 05, (2021) T05001.

[14] A. Anker for the ARIANNA collaboration *PoS* **ICRC2021** (these proceedings) 1074.

[15] J. Alvarez-Muñiz *et al. Astroparticle Physics* **32** (2009) 100 – 111.

[16] J. Alvarez-Muñiz *et al. Astroparticle Physics* **35** (2012) 287 – 299.

[17] K. Simonyan and A. Zisserman *arXiv:1409.1556* (2015) .