

Machine Learning for KM3NeT/ORCA

The KM3NeT Collaboration^{‡*}

[‡] <https://www.km3net.org/km3net-author-list-for-icrc-2019>

E-mail: michael.m.moser@fau.de

The KM3NeT research infrastructure is currently under construction in two locations in the Mediterranean Sea. The KM3NeT/ORCA water-Cherenkov neutrino detector off the French coast will instrument several megatons of seawater with photosensors. Its main objective is the determination of the neutrino mass ordering. For the first time, deep convolutional neural networks are employed to achieve three specific reconstruction and classification tasks that constitute a complete alternative analysis pipeline for KM3NeT/ORCA detector data. The values and uncertainties of the energy, direction, and the interaction vertex of the incident neutrinos are reconstructed. The neutrino interaction type is recognized as shower- or track-like, and the main backgrounds to the detection of atmospheric neutrinos are classified and suppressed. Performance comparisons to other machine-learning methods used and maximum-likelihood reconstruction algorithms previously developed for KM3NeT/ORCA are provided. It is shown that this image recognition technique yields competitive results and performance improvements for large-volume neutrino telescopes.

Corresponding authors: Steffen Hallmann[†], Michael Moser, Stefan Reck, and Thomas Eberl
*Erlangen Centre for Astroparticle Physics (ECAP),
Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany*

*36th International Cosmic Ray Conference -ICRC2019-
July 24th - August 1st, 2019
Madison, WI, U.S.A.*

*for collaboration list see PoS(ICRC2019)1177

[†]Presenter.

1. KM3NeT/ORCA

The KM3NeT research infrastructure is under construction at two sites in the Mediterranean Sea. The KM3NeT/ORCA detector is located about 40 km off-shore of Toulon in the south of France. Its main goal is to investigate atmospheric neutrinos with GeV energies, while KM3NeT/ARCA, located roughly 100 km south-east of Sicily, aims to discover astrophysical neutrinos. The main design principles, the scientific goals of the experiment and an overview of the software packages employed to generate detailed Monte Carlo (MC) simulation datasets can be found in [1].

KM3NeT/ORCA will register interactions of all neutrino flavours with energies roughly above 3 GeV. The indirect recognition of such events proceeds via the detection of Cherenkov photons induced by charged secondary particles created in neutrino-nucleon collisions. KM3NeT/ORCA consists of a three-dimensional matrix of photosensors that are housed in glass spheres, denoted as digital optical modules (DOMs), and arranged along vertical detection lines. These sensors instrument several megatons of water. A particle interaction event induces a time series of detected photons distributed spatially in the matrix. The different possible neutrino interactions can be categorised into two main classes of light distribution, namely elongated, track-like events induced mainly by muons, and shower-like events with spatially localised light emission.

All neutrino flavours can interact through the weak neutral current (NC) mediated by the exchange of a Z^0 boson. This interaction results in a particle shower composed mainly of hadrons, while the scattered neutrino escapes undetected. An interaction via the weak charged current (CC), with the exchange of a charged W boson, can also result in a hadronic shower at the interaction vertex. Additionally, a lepton of the same family as the interacting neutrino is created and carries a fraction of the energy of the incoming neutrino. Muons result in track-like events, while electrons, due to their short radiation length, induce electromagnetic cascades, i.e. shower-like events.

Main backgrounds are track-like events induced by atmospheric muons entering the detector from above and pure noise events induced by the radioactive decay of ^{40}K and the light emission of biological organisms. Typical event rates expected in the full KM3NeT/ORCA detector are mHz for neutrino detections and several 10 Hz for background events.

Three main tasks need to be accomplished to perform neutrino analyses. Firstly, neutrino-induced events must be recognized and separated from the large number of background events. Secondly, the recognized neutrino events must be classified according to their flavour, or as a proxy according to their event topology differentiated into track- and shower-like events. Finally, the properties of the interacting neutrino, such as its direction of incidence and energy, must be inferred. Until now for KM3NeT/ORCA, the event properties have been reconstructed using maximum-likelihood reconstruction algorithms. Based on the resulting reconstruction parameters, decision-tree based machine learning algorithms [2] have been used to suppress background and deduce the event topology in order to approximately distinguish between the flavours ν_e , ν_μ , ν_τ .

Here we present for the first time the application of convolutional neural networks (CNNs) to detailed MC simulations of the KM3NeT/ORCA water-Cherenkov neutrino detector with the goal to provide a complete analysis pipeline starting from low-level detector data. To this end, a Keras-based software framework, called "OrcaNet" ¹, has been developed that simplifies the usage

¹<http://github.com/ViaFerrata/OrcaNet>

of neural networks for neutrino telescopes. Similar investigations and developments have already been done in the context of KM3NeT/ORCA [3], [4].

2. Convolutional neural networks and data preprocessing

CNNs form a specialized class of deep neural networks. They are typically used in domains, where the input can be expected to be image-like, i.e. in image or video classification. Therefore, several changes are made in the architecture of CNNs compared to fully connected neural networks. The neurons in a convolutional layer are arranged in three dimensions, called width, height and depth. One of the main differences between convolutional layers and fully connected layers is that the neurons inside a convolutional layer are only connected to a local region of the input volume. In the convolution process, the local connection region of the input of a neuron is multiplied with a weight matrix and the result is passed on to the next layer. In order to set up a basic convolutional neural network, the convolutional layers are stacked. After the last convolutional layer, the multi-dimensional output is projected onto one dimension. A small fully-connected network can then be added, in order to connect the outputs of the last convolutional layer to the output neurons of the full network.

Detailed MC simulations of the detector response have been produced for three distinct types of triggered data, namely atmospheric muons, random noise and neutrinos. For neutrinos, all three possible charged current interactions on nucleons and nuclei (ν_e^{CC} , ν_μ^{CC} , ν_τ^{CC}) have been simulated, while NC interactions are represented by ν_e^{NC} only, since the detector signature is identical for all flavours. A signal in a photomultiplier (PMT) with a measured time range, for which the anode voltage stays above a tunable threshold, is called a hit. The time resolution is of the order of nanoseconds. In addition to the signal hits induced by the interactions of neutrinos and atmospheric muons in the sensitive detector volume, simulated background hits due to random noise are added to each event such that the simulated triggered data matches the real conditions as closely as possible. After the hit simulation, several trigger algorithms that rely on causality conditions are applied. These can be coincident hits on the same DOM during a certain amount of time. After the triggers have fired to record an event, all hits that have led to the triggering decision are labelled as triggered hits. Since the trigger algorithm is not fully efficient at identifying all signal hits, a larger time window, about $3 \mu\text{s}$, than the one defined by the triggered hits is saved for further analysis. For each hit in a simulated event, the PMT identifier, and thus the relative coordinate, of the hit PMT in a DOM is recorded. Additionally, the starting point when the PMT signal surpasses the threshold and the Time-over-Threshold (ToT) of each individual PMT are stored. The ToT value is not used as input for the CNNs. In order to feed this four-dimensional event data to a CNN, the input hit distribution must be pixelized. The spatial pixelization is defined in a way that exactly one DOM fits into one bin. In the case of the full KM3NeT/ORCA detector, this results in a $11 \times 13 \times 18$ (XYZ) pixel grid. The information, however, which PMT in a DOM has been hit, is lost. This is fixed by adding a PMT identifier dimension to the pixel grid, resulting in a XYZP grid. Since one DOM holds 31 PMTs, the final spatial shape of such an image is $11 \times 13 \times 18 \times 31$ (XYZP). The indispensable piece of information still missing in these images is the time at which a hit has been recorded. This information is added as an additional dimension, such that the final image of an event is five-dimensional: XYZTP. Background hits are discriminated by selecting the time range




Figure 1: Time distribution of about 3000 ν_{μ}^{CC} signal hits centred on the mean time of the triggered hits for single events. The dashed lines in black and red visualize possible time cuts.

in which most signal hits are found. Investigating the distribution of the deviation from the mean of the triggered hit times in individual events, as depicted in Fig. 1 for ν_{μ}^{CC} events, shows that it is asymmetric and that the relevant time range can be reduced significantly for the image generation binning. Since the time range covered by triggered hits is different for each event, the time range selection is defined relative to the mean time of the triggered hits for each event. As can be seen in Fig. 1, only a small fraction of the signal hits is cut, as indicated by the dashed black lines.

After both the spatial as well as the temporal binning are applied, the resulting images are five-dimensional: XYZTP. In order to define and train the CNNs, the software package TensorFlow² has been used in conjunction with the Keras³ high-level neural networks application programming interface. TensorFlow, however, does not support convolutional layers with more than three convolutional dimensions, since five dimensional inputs are not a usual case in computer vision. Hence, the five dimensions of the XYZTP images need to be reduced to four dimensions, such that one can use three-dimensional convolutional layers. To this end, one image dimension is summed up, i.e. the information of individual PMTs in a DOM is discarded, such that the resulting image is only four-dimensional (XYZT). A second image of the same event is then fed to the network (XYZP) that recovers the information which PMT in a DOM has been hit, but discards its hit time. Since these images only differ in the 4th dimension, i.e. the channel dimension of a convolutional layer, the images can be stacked in this dimension. Significant gains in performance for all CNN applications in this work could be observed when using this stacking method, as compared to just supplying a single XYZT image.

Using a Nvidia Tesla V100 GPU with CUDA 10, it has taken of the order of one to two weeks to fully train the CNNs in this work.

3. Event classification

3.1 Background rejection

An essential part of the KM3NeT/ORCA reconstruction pipeline is the background classifier, which discriminates atmospheric muons and random noise from neutrino-induced events. For

²<http://www.tensorflow.org>

³<http://keras.io>

this purpose, the currently employed classification algorithm is based on a Random Forest (RF) method [2]. The inputs of the RF are high-level observables (features), mainly determined from likelihood-based track and shower reconstruction algorithms. The CNN network architecture for the background classifier is based on the three-dimensional convolutional blocks introduced above, with two additional fully-connected layers at the end. The output layer of the CNN is composed of two neurons, such that the network only distinguishes between neutrino and non-neutrino events. As outlined in Sec. 2, XYZP and XYZT images, that are stacked, are used as input to the network. For both event images, a time cut has been defined along the lines introduced above. The total amount of




Figure 2: Fraction of correctly selected neutrino events with energies in the range 1 GeV to 100 GeV versus fraction of residual atmospheric muon contamination after applying the CNN background classifier (blue) and the RF classifier (orange). Error bars are statistical only. The neutrino events have been generated according to the Honda [5] atmospheric neutrino flux model.

MC events used for the training is split equally into the three event classes: neutrinos, atmospheric muons, and random noise. Hence the final class balance is 1/3 neutrino events and 2/3 non-neutrino events. The event dataset is then split into 75% training, 2.5% validation, and 22.5% test events. In total, the training data contains about 42.6×10^6 events. In order to compare the CNN performance directly to the RF background classifier, only the events from the test dataset are used. To quantify the performance of the CNN background classifier, the metric shown below is used to investigate the number of remaining atmospheric muon events for a given threshold value p . The atmospheric muon contamination, and the neutrino efficiency, are defined as: $C_\mu(p) = N_\mu(p)/N_{\text{total}}(p)$ and $\nu_{\text{eff}}(p) = N_\nu(p)/N_{\nu,\text{total}}$. Here, N_μ is the number of remaining atmospheric muon or random noise events, while N_{total} is the total number of remaining events, depending on the threshold probability p . Regarding the neutrino efficiency, $N_\nu(p)$ is the total number of neutrinos in the dataset as a function of the threshold value and $N_{\nu,\text{total}}$ is the number of neutrinos in the dataset, without applying any threshold. Fig. 2 depicts neutrino efficiency ν_{eff} versus atmospheric muon contamination C_μ for the CNN and the RF classifier. The CNN classifier performs better than the RF classifier. In particular, in the regime in which most atmospheric muons are suppressed, the neutrino efficiency is still significantly higher.

3.2 Separation between track- and shower-like neutrino event topologies

The architecture of the track-shower CNN and its input is essentially very similar to the background classifier. The time cut for the hit selection is tighter than for the background classifier. Since background events have already been rejected by the background classifier, the data presented to the track-shower classifier are already mostly neutrino interactions. The event dataset is balanced

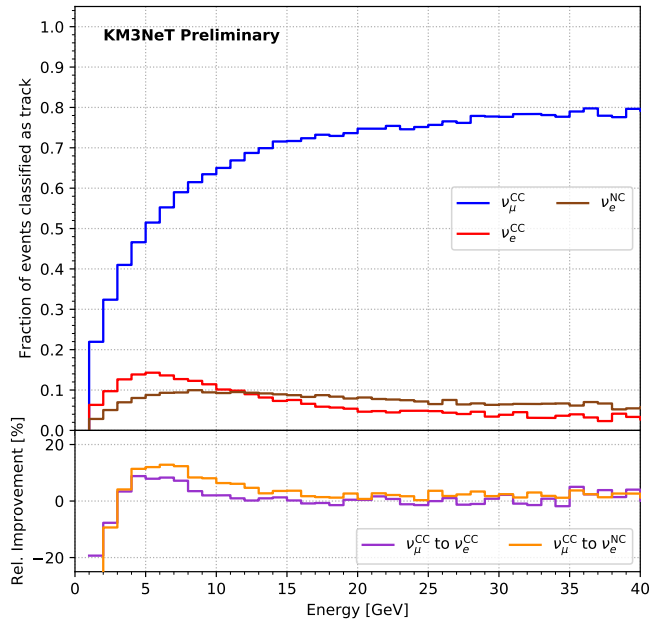


Figure 3: Fraction of events classified as track-like, i.e. with a track probability > 0.5 for different interaction channels (*top panel*) versus the true MC neutrino energy. The power of the CNN classifier to discriminate between ν_μ^{CC} events and showering $\nu_e^{\text{CC}}/\nu_e^{\text{NC}}$ channels, defined as the difference of the fractions classified as track-like, is improved with respect to the RF classifier above the energy turn-on at 3 GeV (*bottom panel*).

such that 50% of the events are track-like (ν_μ^{CC}) and 50% are shower-like. The shower class consists of 50% ν_e^{CC} and 50% ν_e^{NC} . Additionally, the dataset has been balanced in a way that the ratio of track-like to shower-like events is always one and is independent of neutrino energy. This rebalanced dataset is then split into 70% training, 6% validation, and 24% test events. In total, the training dataset contains about 13×10^6 events. Fig. 3 (top panel) shows the fraction of events classified as track-like as a function of neutrino energies in the range 1 GeV to 40 GeV and for different neutrino interaction channels. An event is accepted as track-like if its probability to be track-like is 0.5 or above. It can be seen that the separation power between track-like and shower-like events depends significantly on the neutrino energy. The reason is that the hit distribution of low energetic ν_μ^{CC} events can be similar to that of shower-like events, since the outgoing muon decays after propagating only a short distance in the detector. The bottom panel shows the relative improvement in separation power for different interaction channel combinations when comparing the CNN and RF classifier performance. The performance gain is more important for neutrino energies below 15 GeV and amounts to several percent, in particular in the important low-energy, efficiency-turn-on regime.

4. Event regression

The network architecture of the CNN for the event regressor is identical to the one for the background and track-shower classifier, except for the fully-connected layers after the convolutional layers. The properties to be reconstructed by the network are energy, direction, and vertex position of the initial neutrino interaction. For direction and vertex, the CNN outputs are two arrays with three elements each, containing the X, Y and Z components of the reconstructed direction and vertex position vector. Consequently, the output of the network consists of seven reconstructed floating-point numbers, denoted as \vec{y}_{reco} .

4.1 Energy and direction regression

Fig. 4 (left) shows the median relative error (MRE) as the median of the normalised residual distribution $|(E_{\text{reco}} - E_{\text{true}})/E_{\text{true}}|$ of the reconstructed energy versus the true MC energy for pre-selected ν_e^{CC} events. The performance of the CNN event regressor and of the maximum-likelihood method agree very well. This is likely due to the fact that both methods come close to the intrinsic energy reconstruction performance limit as dictated by light yield fluctuations of the hadronic particle cascade [6].




Figure 4: *Left:* median relative energy reconstruction error versus true MC energy for ν_e^{CC} events for the CNN event regressor (blue) and the ORCA maximum-likelihood shower reconstruction algorithm (orange). *Right:* median error in units of radians of the zenith angle and neutrino direction reconstruction versus true MC energy for ν_e^{CC} events; for the CNN event regressor (green and blue line) and the ORCA maximum-likelihood reconstruction algorithm (orange and red lines).

The median absolute error (ME) is defined as the median of the distribution of the residual distribution for the reconstructed direction. The ME for both the full direction and the zenith angle reconstruction is compared for the CNN event regressor and the maximum-likelihood method in Fig. 4 (right) for ν_e^{CC} events. For energies below a few GeV, the ME is similar for both reconstructions, while the difference in performance increases in favour of the maximum-likelihood method with increasing energy. However, further improvements of the CNN event regressor can be expected in the future since no optimisation has been done yet for this first comparison. Similar comparison results have been obtained regarding the reconstruction of track-like events.

4.2 Error estimate on the regression result

A neural network can be also used to estimate the uncertainties on the components of the reconstructed \vec{y}_{reco} . One possibility is to let the network predict the absolute residual of the mean absolute error loss function. Thus, an additional neuron is needed that yields the reconstructed uncertainty, σ_{reco} , for each of the reconstructed components of \vec{y}_{reco} . Consequently, the loss function that measures the quality of the reconstructed uncertainty value is: $L = \frac{1}{n} \sum_{i=1}^n (\sigma_{\text{reco}} - |y_{\text{true}} - y_{\text{reco}}|)^2$. Using this loss function L , the network learns to estimate the average absolute residual: $\sigma_{\text{reco}} \approx \langle y_{\text{abs}} \rangle$. Fig. 5 shows the standard deviation of the residual distribution of the reconstructed cosine of the zenith angle for ν_e^{CC} versus the fraction of events discarded by the CNN uncertainty estimator. For

all three energy intervals depicted, one can see that the zenith angle resolution improves significantly when the events that have the largest reconstruction errors are discarded as predicted by the CNN.

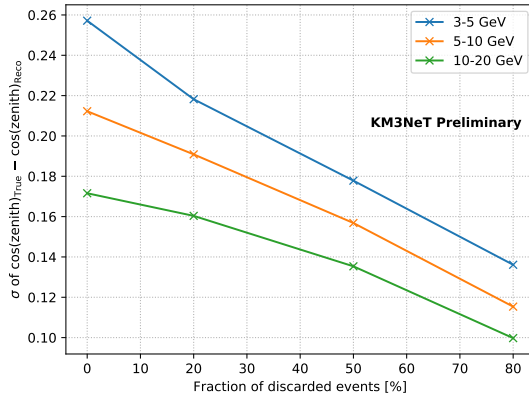


Figure 5: Standard deviation of the residual distribution of the reconstructed cosine of the zenith angle for ν_e^{CC} versus the fraction of events discarded by the CNN uncertainty estimator. The neutrino events have been generated in three different energy intervals as indicated by the three coloured lines and according to the Honda [5] atmospheric neutrino flux model.

5. Conclusion

In this work, the first application of deep convolutional neural networks to the reconstruction of simulated neutrino events in KM3NeT/ORCA has been reported. Suitable CNNs have been designed and tested. Large and detailed MC datasets have been suitably preprocessed to generate pixelized high-dimensional image input for the CNNs. All tasks of the hitherto standard event reconstruction and classification pipeline for KM3NeT/ORCA have been achieved using CNNs and performance comparisons are shown based on the same MC datasets. Though not yet fully optimized, the application of CNNs exhibits promising, and in many cases already better performances than the standard methods.

Acknowledgement: Part of this work has been supported by the ASTERICS and ESCAPE projects, funded by the European Commission Framework Programme Horizon 2020 Research and Innovation action under grant agreement n. 653477 and n. 824064, respectively.

References

- [1] S. Adrián-Martínez et al., *Letter of Intent for KM3NeT 2.0*, *J. Phys.* **G43** (2016), no. 8 084001, [[arXiv:1601.0745](#)].
- [2] L. Breiman, *Random forests*, *Machine Learning* **45** (Oct, 2001) 5–32.
- [3] C. D. Sio, *Machine Learning in KM3NeT*, *EPJ Web Conf.* **207** (2019) 05004.
- [4] C. Bozza et al., *pLISA: a parallel Library for Identification and Study of Astroparticles*, *Proceedings of Science*, in press (2019).
- [5] M. Honda et al., *Atmospheric neutrino flux calculation using the NRLMSISE-00 atmospheric model*, *Phys. Rev.* **D92** (2015) 023004, [[arXiv:1502.0391](#)].
- [6] S. Adrián-Martínez et al., *Intrinsic limits on resolutions in muon- and electron-neutrino charged-current events in the KM3NeT/ORCA detector*, *JHEP* **05** (2017) 008, [[arXiv:1612.0562](#)].