

OPEN ACCESS

Event reconstruction for KM3NeT/ORCA using convolutional neural networks

To cite this article: S. Aiello *et al* 2020 *JINST* **15** P10005

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



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

RECEIVED: April 20, 2020

REVISED: July 10, 2020

ACCEPTED: August 11, 2020

PUBLISHED: October 8, 2020

Event reconstruction for KM3NeT/ORCA using convolutional neural networks

S. Aiello,^a A. Albert,^{bb,b} S. Alves Garre,^c Z. Aly,^d F. Ameli,^e M. Andre,^f G. Androulakis,^g M. Anghinolfi,^h M. Anguita,ⁱ G. Anton,^j M. Ardid,^k J. Aublin,^l C. Bagatelas,^g G. Barbarino,^{m,n} B. Baret,^l S. Basegmez du Pree,^o M. Bendahman,^p E. Berbee,^o A. M. van den Berg,^q V. Bertin,^d S. Biagi,^r A. Biagioni,^e M. Bissinger,^j M. Boettcher,^s J. Boumaaza,^p M. Bouta,^t M. Bouwhuis,^o C. Bozza,^u H. Brânzaş,^v R. Bruijn,^{o,w} J. Brunner,^d E. Buis,^x R. Buompane,^{m,y} J. Busto,^d B. Caiffi,^h D. Calvo,^c A. Capone,^{z,e} V. Carretero,^c P. Castaldi,^{aa} S. Celli,^{z,e,bc} M. Chabab,^{ab} N. Chau,^l A. Chen,^{ac} S. Cherubini,^{r,ad} V. Chiarella,^{ae} T. Chiarusi,^{aa} M. Circella,^{af} R. Cocimano,^r J. A. B. Coelho,^l A. Coleiro,^l M. Colomer Molla,^{l,c} R. Coniglione,^r P. Coyle,^d A. Creusot,^l G. Cuttone,^r A. D'Onofrio,^{m,y} R. Dallier,^{ag} M. De Palma,^{af,ah} I. Di Palma,^{z,e} A. F. Díaz,ⁱ D. Diego-Tortosa,^k C. Distefano,^r A. Domi,^{h,d,ai} R. Donà,^{aa,aj} C. Donzaud,^l D. Dornic,^d M. Dörr,^{ak} D. Drouhin,^{bb,b} T. Eberl,^{j,1} A. Eddyamoui,^p T. van Eeden,^o D. van Eijk,^o I. El Bojaddaini,^t D. Elsaesser,^{ak} A. Enzenhöfer,^d V. Espinosa Roselló,^k P. Fermani,^{z,e} G. Ferrara,^{r,ad} M. D. Filipović,^{al} F. Filippini,^{aa,aj} L. A. Fusco,^l O. Gabella,^{am} T. Gal,^j A. Garcia Soto,^o F. Garufi,^{m,n} Y. Gatelet,^l N. Geißelbrecht,^j L. Gialanella,^{m,y} E. Giorgio,^r S. R. Gozzini,^c R. Gracia,^o K. Graf,^j D. Grasso,^{an} G. Grella,^{ao} D. Guderian,^{bd} C. Guidi,^{h,ai} S. Hallmann,^j H. Hamdaoui,^p H. van Haren,^{ap} A. Heijboer,^o A. Hekalo,^{ak} J. J. Hernández-Rey,^c J. Hofestädt,^j F. Huang,^{aq} W. Idrissi Ibsalikh,^{m,y} G. Illuminati,^c C. W. James,^{ar} M. de Jong,^o P. de Jong,^{o,w} B. J. Jung,^o M. Kadler,^{ak} P. Kalaczyński,^{as} O. Kalekin,^j U. F. Katz,^j N. R. Khan Chowdhury,^c G. Kistauri,^{at} F. van der Knaap,^x E. N. Koffeman,^{o,w} P. Kooijman,^{w,be} A. Kouchner,^{l,au} M. Kreter,^s V. Kulikovskiy,^h R. Lahmann,^j G. Larosa,^r R. Le Breton,^l O. Leonardi,^r F. Leone,^{r,ad} E. Leonora,^a G. Levi,^{aa,aj} M. Lincetto,^d M. Lindsey Clark,^l T. Lipreau,^{ag} A. Lonardo,^e F. Longhitano,^a D. Lopez-Coto,^{av} L. Maderer,^l J. Mańczak,^c K. Mannheim,^{ak} A. Margiotto,^{aa,aj} A. Marinelli,^m C. Markou,^g L. Martin,^{ag} J. A. Martínez-Mora,^k A. Martini,^{ae} F. Marzaioli,^{m,y} S. Mastroianni,^m S. Mazzou,^{ab} K. W. Melis,^o G. Miele,^{m,n} P. Migliozi,^m E. Migneco,^r P. Mijakowski,^{as} L. S. Miranda,^{aw} C. M. Mollo,^m M. Morganti,^{an,bf} M. Moser,^{j,1} A. Moussa,^t R. Muller,^o M. Musumeci,^r L. Nauta,^o S. Navas,^{av} C. A. Nicolau,^e B. Ó Fearraigh,^{o,w} M. Organokov,^{aq} A. Orlando,^r G. Papalashvili,^{at} R. Papaleo,^r C. Pastore,^{af} A. M. Păun,^v G. E. Păvălaş,^v C. Pellegrino,^{aj,bg} M. Perrin-Terrin,^d P. Piattelli,^r

¹Corresponding author



C. Pieterse,^c K. Pikounis,^g O. Pisanti,^{m,n} C. Poirè,^k V. Popa,^v M. Post,^w T. Pradier,^{aq}
 G. Pühlhofer,^{ax} S. Pulvirenti,^r O. Rabyang,^s F. Raffaelli,^{an} N. Randazzo,^a A. Rapicavoli,^{ad}
 S. Razzaque,^{aw} D. Real,^c S. Reck,^j G. Riccobene,^r M. Richer,^{aq} S. Rivoire,^{am} A. Rovelli,^r
 F. Salesa Greus,^c D. F. E. Samtleben,^{o,ay} A. Sánchez Losa,^{af} M. Sanguineti,^{h,ai}
 A. Santangelo,^{ax} D. Santonocito,^r P. Sapienza,^r J. Schnabel,^j J. Seneca,^o I. Sgura,^{af}
 R. Shandize,^{at} A. Sharma,^{az} F. Simeone,^e A. Sinopoulou,^g B. Spisso,^{ao,m} M. Spurio,^{aa,aj}
 D. Stavropoulos,^g J. Steijger,^o S. M. Stellacci,^{ao,m} M. Taiuti,^{h,ai} Y. Tayalati,^p E. Tenllado,^{av}
 T. Thakore,^c S. Tingay,^{ar} E. Tzamariudaki,^g D. Tzanetatos,^g V. Van Elewyck,^{l,au}
 G. Vannoye,^h G. Vasileiadis,^{am} F. Versari,^{aa,aj} S. Viola,^r D. Vivolo,^{m,n} G. de Wasseige,^l
 J. Wilms,^{ba} R. Wojaczyński,^{as} E. de Wolf,^{o,w} D. Zaborov,^{d,bh} S. Zavatarelli,^h A. Zegarelli,^{z,e}
 D. Zito,^r J. D. Zornoza,^c J. Zúñiga,^c N. Zywucka^s

^aINFN, Sezione di Catania, Via Santa Sofia 64, Catania, 95123 Italy

^bIN2P3, IPHC, 23 rue du Loess, Strasbourg, 67037 France

^cIFIC - Instituto de Física Corpuscular (CSIC - Universitat de València), c/Catedrático José Beltrán, 2, 46980 Paterna, Valencia, Spain

^dAix Marseille Univ, CNRS/IN2P3, CPPM, Marseille, France

^eINFN, Sezione di Roma, Piazzale Aldo Moro 2, Roma, 00185 Italy

^fUniversitat Politècnica de Catalunya, Laboratori d'Aplicacions Bioacústiques, Centre Tecnològic de Vilanova i la Geltrú, Avda. Rambla Exposició, s/n, Vilanova i la Geltrú, 08800 Spain

^gNCSR Demokritos, Institute of Nuclear and Particle Physics, Ag. Paraskevi Attikis, Athens, 15310 Greece

^hINFN, Sezione di Genova, Via Dodecaneso 33, Genova, 16146 Italy

ⁱUniversity of Granada, Dept. of Computer Architecture and Technology/CITIC, 18071 Granada, Spain

^jFriedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen Centre for Astroparticle Physics, Erwin-Rommel-Straße 1, 91058 Erlangen, Germany

^kUniversitat Politècnica de València, Instituto de Investigación para la Gestión Integrada de las Zonas Costeras, C/Paraninfo, 1, Gandia, 46730 Spain

^lUniversité de Paris, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France

^mINFN, Sezione di Napoli, Complesso Universitario di Monte S. Angelo, Via Cintia ed. G, Napoli, 80126 Italy

ⁿUniversità di Napoli "Federico II", Dip. Scienze Fisiche "E. Pancini", Complesso Universitario di Monte S. Angelo, Via Cintia ed. G, Napoli, 80126 Italy

^oNikhef, National Institute for Subatomic Physics, PO Box 41882, Amsterdam, 1009 DB Netherlands

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

^qKVI-CART University of Groningen, Groningen, the Netherlands

^rINFN, Laboratori Nazionali del Sud, Via S. Sofia 62, Catania, 95123 Italy

^sNorth-West University, Centre for Space Research, Private Bag X6001, Potchefstroom, 2520 South Africa

^tUniversity Mohammed I, Faculty of Sciences, BV Mohammed VI, B.P. 717, R.P. 60000 Oujda, Morocco

^uUniversità di Salerno e INFN Gruppo Collegato di Salerno, Dipartimento di Matematica, Via Giovanni Paolo II 132, Fisciano, 84084 Italy

^vISS, Atomistilor 409, Măgurele, RO-077125 Romania

^wUniversity of Amsterdam, Institute of Physics/IHEF, PO Box 94216, Amsterdam, 1090 GE Netherlands

^xTNO, Technical Sciences, PO Box 155, Delft, 2600 AD Netherlands

- ^yUniversità degli Studi della Campania “Luigi Vanvitelli”, Dipartimento di Matematica e Fisica, viale Lincoln 5, Caserta, 81100 Italy
- ^zUniversità La Sapienza, Dipartimento di Fisica, Piazzale Aldo Moro 2, Roma, 00185 Italy
- ^{aa}INFN, Sezione di Bologna, v.le C. Berti-Pichat, 6/2, Bologna, 40127 Italy
- ^{ab}Cadi Ayyad University, Physics Department, Faculty of Science Semlalia, Av. My Abdellah, P.O.B. 2390, Marrakech, 40000 Morocco
- ^{ac}University of the Witwatersrand, School of Physics, Private Bag 3, Johannesburg, Wits 2050 South Africa
- ^{ad}Università di Catania, Dipartimento di Fisica e Astronomia, Via Santa Sofia 64, Catania, 95123 Italy
- ^{ae}INFN, LNF, Via Enrico Fermi, 40, Frascati, 00044 Italy
- ^{af}INFN, Sezione di Bari, Via Amendola 173, Bari, 70126 Italy
- ^{ag}Subatech, IMT Atlantique, IN2P3-CNRS, Université de Nantes, 4 rue Alfred Kastler - La Chantrerie, Nantes, BP 20722 44307 France
- ^{ah}University of Bari, Via Amendola 173, Bari, 70126 Italy
- ^{ai}Università di Genova, Via Dodecaneso 33, Genova, 16146 Italy
- ^{aj}Università di Bologna, Dipartimento di Fisica e Astronomia, v.le C. Berti-Pichat, 6/2, Bologna, 40127 Italy
- ^{ak}University Würzburg, Emil-Fischer-Straße 31, Würzburg, 97074 Germany
- ^{al}Western Sydney University, School of Computing, Engineering and Mathematics, Locked Bag 1797, Penrith, NSW 2751 Australia
- ^{am}Laboratoire Univers et Particules de Montpellier), Place Eugène Bataillon - CC 72, Montpellier Cédex 05, 34095 France
- ^{an}INFN, Sezione di Pisa, Largo Bruno Pontecorvo 3, Pisa, 56127 Italy
- ^{ao}Università di Salerno e INFN Gruppo Collegato di Salerno, Dipartimento di Fisica, Via Giovanni Paolo II 132, Fisciano, 84084 Italy
- ^{ap}NIOZ (Royal Netherlands Institute for Sea Research) and Utrecht University, PO Box 59, Den Burg, Texel, 1790 AB, the Netherlands
- ^{aq}Université de Strasbourg, CNRS IPHC UMR 7178, 23 rue du Loess, Strasbourg, 67037 France
- ^{ar}International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia
- ^{as}National Centre for Nuclear Research, 02-093 Warsaw, Poland
- ^{at}Tbilisi State University, Department of Physics, 3, Chavchavadze Ave., Tbilisi, 0179 Georgia
- ^{au}Institut Universitaire de France, 1 rue Descartes, Paris, 75005 France
- ^{av}University of Granada, Dpto. de Física Teórica y del Cosmos & C.A.F.P.E., 18071 Granada, Spain
- ^{aw}University of Johannesburg, Department Physics, PO Box 524, Auckland Park, 2006 South Africa
- ^{ax}Eberhard Karls Universität Tübingen, Institut für Astronomie und Astrophysik, Sand 1, Tübingen, 72076 Germany
- ^{ay}Leiden University, Leiden Institute of Physics, PO Box 9504, Leiden, 2300 RA Netherlands
- ^{az}Università di Pisa, Dipartimento di Fisica, Largo Bruno Pontecorvo 3, Pisa, 56127 Italy
- ^{ba}Friedrich-Alexander-Universität Erlangen-Nürnberg, Remeis Sternwarte, Sternwartstraße 7, 96049 Bamberg, Germany
- ^{bb}Université de Strasbourg, Université de Haute Alsace, GRPHE, 34, Rue du Grillenbreit, Colmar, 68008 France
- ^{bc}Gran Sasso Science Institute, GSSI, Viale Francesco Crispi 7, L'Aquila, 67100 Italy
- ^{bd}University of Münster, Institut für Kernphysik, Wilhelm-Klemm-Str. 9, Münster, 48149 Germany
- ^{be}Utrecht University, Department of Physics and Astronomy, PO Box 80000, Utrecht, 3508 TA Netherlands
- ^{bf}Accademia Navale di Livorno, Viale Italia 72, Livorno, 57100 Italy

^{bg}INFN, CNAF, v.le C. Berti-Pichat, 6/2, Bologna, 40127 Italy

^{bh}NRC “Kurchatov Institute”, A.I. Alikhanov Institute for Theoretical and Experimental Physics, Bolshaya Cheremushkinskaya ulitsa 25, Moscow, 117218 Russia

E-mail: michael.m.moser@fau.de, thomas.eberl@fau.de

ABSTRACT: The KM3NeT research infrastructure is currently under construction at 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. This work aims at demonstrating the general applicability of deep convolutional neural networks to neutrino telescopes, using simulated datasets for the KM3NeT/ORCA detector as an example. To this end, the networks are employed to achieve reconstruction and classification tasks that constitute an alternative to the analysis pipeline presented for KM3NeT/ORCA in the KM3NeT Letter of Intent. They are used to infer event reconstruction estimates for the energy, the direction, and the interaction point of incident neutrinos. The spatial distribution of Cherenkov light generated by charged particles induced in neutrino interactions is classified as shower- or track-like, and the main background processes associated with the detection of atmospheric neutrinos are recognized. Performance comparisons to machine-learning classification and maximum-likelihood reconstruction algorithms previously developed for KM3NeT/ORCA are provided. It is shown that this application of deep convolutional neural networks to simulated datasets for a large-volume neutrino telescope yields competitive reconstruction results and performance improvements with respect to classical approaches.

KEYWORDS: Cherenkov detectors; Large detector systems for particle and astroparticle physics; Neutrino detectors; Performance of High Energy Physics Detectors

ARXIV EPRINT: [2004.08254](https://arxiv.org/abs/2004.08254)

Contents

1	Introduction	1
2	The KM3NeT/ORCA experiment	3
2.1	Layout of the detector	3
2.2	Monte Carlo simulations and trigger algorithms	4
3	Convolutional neural networks	4
4	Data pre-processing	7
4.1	Spatial binning	7
4.2	Temporal Binning	7
4.3	Multi-image Convolutional Neural Networks	8
5	Main network architecture	9
6	Background classifier	11
6.1	Image generation	11
6.2	Network architecture	12
6.3	Preparation of training, validation and test data	13
6.4	Performance and comparison to Random Forest classifier	13
7	Event topology classifier	16
7.1	Image generation	17
7.2	Network architecture	18
7.3	Preparation of training, validation and test data	18
7.4	Performance and comparison to Random Forest classifier	19
8	Event regression	22
8.1	Image generation	22
8.2	Network architecture and loss functions	23
8.3	Preparation of training, validation and test data	25
8.4	Loss functions and loss weights	25
8.5	Energy reconstruction performance	26
8.6	Direction reconstruction performance	29
8.7	Vertex reconstruction performance	30
8.8	Error estimation	31
9	Conclusion	32

1 Introduction

Precision measurements of the fundamental properties of neutrinos are one of the opportunities that might allow us to discover and understand the physics that exists beyond the established Standard Model of particle physics.

The detection of neutrinos, both for fundamental particle physics and high-energy astrophysics, can be achieved with the deep-sea and photon-detection technology that has been developed by the ANTARES [1] and KM3NeT [2] Collaborations for very-large-volume water-Cherenkov detectors.

KM3NeT/ORCA, the low-energy detector of KM3NeT, addresses the determination of a still unknown, but fundamental parameter of neutrino physics: the neutrino mass ordering. The experiment focuses on the measurement of the energy- and zenith-angle-dependent oscillation patterns of cosmic-ray-induced neutrinos with a few-GeV energy that originate in the atmosphere and traverse the Earth [3].

The power to distinguish between the two different mass orderings is linked to the detection of an excess or deficit of neutrino events in different regions of these oscillation patterns. This sensitivity increases with better energy and zenith-angle resolution and flavour identification for the interacting neutrinos, and finer control of systematic effects that influence the measurement. Therefore, one of the most important goals in the analysis of KM3NeT/ORCA data is the development and characterisation of neutrino event reconstruction and classification algorithms that improve these resolutions.

The neutrino detection principle of water- or ice-based large-scale Cherenkov detectors relies on the detection of Cherenkov photons induced by charged secondary particles created in a neutrino interaction with the target material. 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, generically referred to as a *hadronic system*, while the scattered neutrino escapes undetected. An interaction via the weak charged current (CC), with the exchange of a W^+ or W^- boson, also often results in a hadronic shower at the interaction vertex. Additionally, a lepton of the same flavour as the interacting neutrino is created, which carries a fraction of the incoming neutrino energy.

A muon neutrino or muon antineutrino CC interaction, $\bar{\nu}_\mu^{\text{CC}}$,¹ results in an outgoing muon in the final state. From now on the term ‘neutrino’ refers always to both neutrinos and antineutrinos, if not stated otherwise. The muon appears as a track-like light source in the detector, and can therefore be identified with good confidence, depending on its track length. The visible trajectory of the muon is determined by its energy loss, and in water it amounts to about roughly 4 m per GeV of muon energy for the relevant energy regime of a few GeV.

At the energy ranges considered in KM3NeT/ORCA, all neutrino-nucleon NC, ν_e^{CC} , and ν_τ^{CC} interactions, with the exception of roughly 18% of tau leptons decaying into muons, create a particle shower of a few meters length, that appears as an elongated, but localised, light source compared to the typical distance scales between the detector elements (9–20 m, cf. section 2.1). This event type is referred to as shower-like. The outgoing electron from a ν_e^{CC} event initiates an electromagnetic shower, a cascade of $e^\pm$ -pairs, while the hadronic system, typically at the neutrino interaction vertex, develops into a hadronic shower with large event-to-event fluctuations and a

¹The symbol $\bar{\nu}$ denotes both, neutrino and antineutrino, at the same time.

possibly complex structure of hadronic or electromagnetic sub-showers, depending on the decay modes of individual particles in the shower.

Although an electromagnetic shower consists of many $e^\pm$ -pairs with rather short path lengths (about 36 cm radiation length in water, see section 33.4.2 in ref. [4]) and overlapping Cherenkov cones, the small pair opening angle preserves the Cherenkov angle peak of the total angular light distribution. This results in a single Cherenkov ring projected onto the plane perpendicular to the shower axis. Similarly, each hadronic shower particle with energy above the Cherenkov threshold will produce a Cherenkov ring. Therefore, hadronic showers show a variety of different signatures due to the various possible combinations of initial hadron types, their momenta and the diversity of their hadronic interactions in the shower evolution.

While electromagnetic cascades show only negligible fluctuations in the number of emitted Cherenkov photons and in the angular light distribution, hadronic cascades show significant intrinsic fluctuations in the relevant few-GeV energy range. These intrinsic fluctuations of hadronic cascades and the resulting limitations for the energy and angular resolutions have been studied in detail in ref. [5].

Dedicated reconstruction algorithms for track-like and shower-like events have been developed for KM3NeT/ORCA based on maximum-likelihood methods. Additionally, a machine-learning algorithm, based on Random Forests [6], has been employed successfully to classify track-like, shower-like, and background events. These algorithms, their implementation and performance are described in the KM3NeT Letter of Intent [2].

In the last few years, significant progress has been made in the machine-learning community due to the advent of deep-learning techniques. A particularly successful deep-learning concept is that of a deep neural network. Specialised neural network model architectures have been designed for individual use cases. In the field of computer vision, Convolutional Neural Networks (CNNs) have led to a strong increase in image recognition performance. From 2010 to 2016, the error rates in e.g. the popular ImageNet image classification challenge improved by a factor of 10 [7, 8].

Since the data of many high-energy physics experiments can be interpreted in a way similar to typical images in the computer vision domain, these techniques have already been exploited by several experiments [9–12]. As an example, the classification performance of neutrino interactions in the NOvA experiment has been significantly improved by employing CNNs compared to classical reconstruction tools [13].

In this paper, we present for the first time the application of CNNs to detailed Monte Carlo simulations of a large water-Cherenkov neutrino detector with the goal to provide a comprehensive reconstruction pipeline for KM3NeT/ORCA, starting from data at the level of the data-acquisition system. For this purpose, a Keras-based [14] software framework, called *OrcaNet* [15], has been developed that simplifies the usage of neural networks for neutrino telescopes. Here, we apply this framework to neutrino event reconstruction and classification in KM3NeT/ORCA, and compare the results achieved with the algorithms described in ref. [2]. We note that also these algorithms continue to be developed and improved in KM3NeT.

The paper is organised as follows. Section 2 introduces the KM3NeT/ORCA detector, the Monte Carlo simulation chain used to generate training and validation data for the CNNs and introduces relevant aspects of the trigger algorithms. Section 3 gives a brief introduction to the main functional features of the employed CNNs, while Section 4 details the developed pre-

processing chain that creates suitable input images from the Monte Carlo simulation data. Section 5 provides an overview of the general network architecture that is shared by all CNNs that have been designed for the reconstruction and classification tasks, which together define the analysis pipeline for KM3NeT/ORCA. The concepts and performance of these specific CNNs, as well as exemplary comparisons to their counterpart algorithms, are explained in the next sections. Section 6 explains the background classifier, section 7 the event topology classifier used to distinguish track-like and shower-like events, while section 8 introduces event regression and its respective uncertainties, i.e. the reconstruction of the direction, energy, and vertex of the incident neutrinos. Section 9 summarises and concludes the paper.

2 The KM3NeT/ORCA experiment

The KM3NeT research infrastructure is under construction at two sites in the Mediterranean Sea. The KM3NeT/ORCA detector is located at a depth of 2450 m about 40 km off-shore of Toulon in the south of France. Its main goal is to detect atmospheric neutrinos with GeV energies (3–40 GeV), while KM3NeT/ARCA, located south-east of Sicily, aims to investigate astrophysical neutrinos. The main design principles and scientific goals of the experiment can be found in the KM3NeT Letter of Intent [2].

2.1 Layout of the detector

The detector volume of KM3NeT/ORCA will be instrumented with 115 Detection Units (DUs), which are vertical, string-like structures anchored to the seabed and held upright by a buoy at the top of the DU. Currently, the first six DUs have been installed and are operational. Each DU holds 18 Digital Optical Modules (DOMs). The DOMs contain 31 photomultiplier tubes (PMTs) with a diameter of 3" each. The PMTs are used to measure two quantities, the arrival time and the time range that the anode output signal remains above a tunable threshold (time-over-threshold, ToT) with a time resolution on the nanosecond scale. The ToT can be used as a proxy for the amount of light registered by the PMT. The vertical spacing between the DOMs on a single DU is on average 9 m, while the average horizontal distance between the DUs is about 20 m, so that the DUs can be contained in a cylinder with roughly 120 m radius. Each DU is 200 m in height with DOMs starting about 40 m above the sea floor. The distribution of the DUs on the sea floor is shown in figure 1 (left). This results in a total instrumented volume of about six megatons of seawater.

There are two main sources of background in KM3NeT/ORCA, namely atmospheric muons reaching the detector from above and random optical background due to beta decays of ^{40}K in seawater and bioluminescence. The atmospheric muon flux has been measured with the first DU of KM3NeT/ORCA and has been found to be compatible with expectations over the entire depth range considered [16]. Optical background, which is dominated by decays of ^{40}K , accounts for about 7 kHz of uncorrelated single photon noise per PMT with a rate of two-fold coincidences of about 500 Hz per DOM [16].

The atmospheric muon background can be reduced significantly by requiring the reconstructed vertex position to be inside or close to the instrumented detector volume. In addition, predominantly atmospheric muons enter the detector from above and therefore the atmospheric muon background

can be further reduced by discarding events for which the direction of the emerging particle trajectory is reconstructed as downwards.

2.2 Monte Carlo simulations and trigger algorithms

Detailed Monte Carlo (MC) simulations of the detector response have been produced for three distinct types of triggered data, namely atmospheric muons, random noise and neutrinos. A detailed introduction to the KM3NeT simulation package and the trigger algorithms can be found in ref. [2].

For neutrinos, ν_e^{CC} and ν_μ^{CC} interactions on nucleons and nuclei in seawater have been simulated, while all NC interactions are represented by ν_e^{NC} , since the detector signature is identical for all flavours. Charged-current interactions of ν_τ are neglected for simplicity, as the resulting detector signatures for the different decay modes of the tau lepton are very similar to either ν_e^{CC} or ν_μ^{CC} .

The distance between DUs for the simulations employed in this work is on average 23 m and hence 3 m larger than for the simulations used in ref. [2]. The vertical inter-DOM spacing is set to 9 m on average, which was identified as optimal for the determination of the neutrino mass ordering in ref. [2]. The highest level of simulated data consists of a list of *hits*, i.e. time stamp, ToT, and identifier, for all PMTs in the detector. In addition to the *signal* hits induced by the interactions of neutrinos and by atmospheric muons in the sensitive detector volume, simulated *background* hits due to random noise are added such that the simulated triggered data matches the real conditions as closely as possible.

After hits have been simulated, several trigger algorithms that rely on causality conditions are applied. Once a trigger has fired to define an *event*, all hits that have fired the trigger, including signal and background hits, are labelled as *triggered* hits. Since the trigger algorithm is not fully efficient in identifying all signal hits, a larger time window than the one defined for the triggered hits is saved for further analysis. Assuming a triggered event with t_{first} as the time of the first triggered hit and t_{last} as the time of the last triggered hit, all photon hits in each PMT in a time window $[t_{\text{first}} - t_{\text{marg}}, t_{\text{last}} + t_{\text{marg}}]$ are recorded, where t_{marg} is defined by the maximum amount of time that a photon propagating in water would need to traverse the whole detector. As a result, the total time window of triggered neutrino events in KM3NeT/ORCA is about 3 μs .

A summary with detailed information about the simulated data for KM3NeT/ORCA is shown in table 1.

3 Convolutional neural networks

This section introduces the concepts and nomenclature used in the description of the networks that have been developed and used in this work.

Convolutional neural networks [19, 20] form a specialised class of deep neural networks. Generally, neural networks are used in order to approximate a function $f(x)$, which maps a certain number of inputs $x_i \in X$ to some outputs $y_i \in Y$. The goal is then to find an approximation $\hat{f}(x)$ to the function $f(x)$ that describes the relationship between the inputs x_i and the outputs y_i . Neural networks are based on the concept of artificial neurons that are arranged in layers. For a fully connected neural network, each neuron in a layer is connected to all neurons in the previous layer.

Table 1. List of Monte Carlo simulations for a KM3NeT/ORCA detector composed of 115 DUs. The first column reports the simulated event type. The neutrino simulations comprise neutrino and antineutrino interactions of the indicated type. The second column specifies the power law used to simulate the energy spectrum of the interacting neutrinos. A reweighting scheme is used in this work, where appropriate, to simulate an atmospheric neutrino flux model [17]. N_{trig} is the number of events that remain after triggering. Atmospheric muons have been simulated with the MUPAGE package [18]. Random noise events have been simulated conservatively with a 10 kHz single rate per PMT with additional n -fold coincidences (600 Hz two-fold, 60 Hz three-fold, 7 Hz four-fold, 0.8 Hz five-fold and 0.08 Hz six-fold [16]). Time-varying increases of the hit rate due to bioluminescence in seawater have not been simulated.

Event type	E_{gen} spectrum	N_{trig} [10^6]	Energy range
Atmospheric muon	-	65.2	-
Random noise	-	23.3	-
ν_e^{NC}	E^{-1}	1.1	1-5 GeV
$\bar{\nu}_e^{\text{NC}}$	E^{-3}	3.7	3-100 GeV
ν_e^{CC}	E^{-1}	1.5	1-5 GeV
$\bar{\nu}_e^{\text{CC}}$	E^{-3}	4.4	3-100 GeV
ν_μ^{CC}	E^{-1}	1.7	1-5 GeV
$\bar{\nu}_\mu^{\text{CC}}$	E^{-3}	8.3	3-100 GeV

Stacking multiple layers of neurons can be interpreted as multiple functions that are acting on the input X in a chain. For a two-layer network and thus two functions $f^{(1)}$ and $f^{(2)}$ (the $\hat{f}$ symbol of $\hat{f}$ is neglected from this point on) this is:

$$f(x) = f^{(2)}(f^{(1)}(x)). \quad (3.1)$$

Here, $f^{(1)}$ refers to the first layer in the network and $f^{(2)}$ to the second. The first layer of a neural network is called the *input* layer, the intermediate layers are called the *hidden* layers and the last layer is called the *output* layer.

In order to learn the relationship between (X, Y) , learnable weights are used for each neuron. If a single neuron has inputs x_i , then each x_i has a weight w_i associated to this input. Additionally, a single, learnable *bias* parameter is added in order to increase the flexibility of the model to fit the data. This process in a neuron, consisting of the weights w_i and the bias b , shows a linear response:

$$f_\Sigma = \sum_{i=1}^n w_i x_i + b. \quad (3.2)$$

However, many physical processes in nature are inherently nonlinear. To account for this, the output of the transfer function can be wrapped in another, nonlinear function. Additionally, it can be shown that a nonlinear, two-layer neural network can approximate any function (Chap. 6.4.1. in ref. [20]). The most commonly used, nonlinear function is the *rectified linear unit* (ReLU):

$$f_{\text{ReLU}}(x) = \max(0, x). \quad (3.3)$$

Such functions are called *activation* functions or layers.

The weights of each neuron get updated iteratively during a training process. For this purpose, one needs to define a so-called *cost* or *loss function*, which measures the distance between the output of the neural network $f(x) = y_{\text{reco}}$ and the ground truth y_{true} . This can for example be done by measuring the mean squared error and minimising $(y_{\text{true}} - y_{\text{reco}})^2$.

Typically, iterative gradient descent (Chap. 4.3. in ref. [20]) based optimisation algorithms are used that minimise the cost function until a low value is achieved. During this training process, the cost error is back-propagated using a back-propagation algorithm (Chap. 6.5. in ref. [20]), which allows for the tuning of the neural network's weights.

Convolutional neural networks are frequently 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 main concepts of convolutional neural networks are based on only locally and not fully connected networks and on parameter sharing between certain neurons in the network.

Typical input images to convolutional neural networks are two-dimensional (2D). However, since most images are coloured, they in fact are encoded by three dimensions (3D): width, height and channel. Here, the channel dimension specifies the brightness for each colour channel (red, green and blue) of the image. This three dimensional array is then used as input to the first convolutional layer.

Similar to the input layer, the neurons in a convolutional layer are also arranged in three dimensions, called width, height and depth. As already mentioned, 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. This is often called the *receptive field* of the neuron. The connections of this local area to the neuron are local in space (width, height), but they are always full along the depth of the input volume. Hence, for a $[32 \times 32 \times 3]$ image (width, height, channel) and receptive field size of 5 pixels, each neuron in the first convolutional layer is connected to a local $[5 \times 5 \times 3]$ (width, height, depth) patch of the input. To each of these connections, a weight is assigned, such that each neuron has a $[5 \times 5 \times 3]$ weight matrix, which is often called the *kernel* or *filter*. These weights are used in performing a dot product between the receptive field of the neuron and its associated kernel, also called the *convolution* process. Here, the total number of parameters for the single neuron would be $5 \cdot 5 \cdot 3 + 1$ (bias) = 76. Additionally, each neuron at the same depth level covers a different part of the image with its receptive field. For more information the reader is referred to refs. [19, 20].

For CNNs, an important assumption is that abstract image structures, such as edges, occur multiple times in the image. Under this assumption, the neurons at a certain depth can share their weights, which significantly reduces the number of parameters in the network. The number of parameters in CNNs can be further reduced with the aid of *pooling* layers, which reduce the dimensionality of the layer outputs by selecting or combining the neuron outputs [19, 20].

For the training of a neural network, the data are split up into a training, validation, and test dataset. The network is trained on the training dataset and validated by applying the network to the validation dataset. Once the training is finished, the network is applied to the test dataset to determine its performance. If the value of the loss function is significantly greater for the training dataset than for the validation dataset, this is called *overfitting* with respect to the training dataset. The smaller the size of the training dataset, the higher the probability that the network will focus

on the peculiarities of individual input images instead of generalising generic image features. In order to avoid overfitting, so-called regularisation techniques such as *dropout* layers have been developed [21]. In a dropout layer, inputs are randomly set to zero with a probability defined by the *dropout rate* δ . In order to set up a basic convolutional neural network, the convolutional layers are stacked. After the last convolutional layer, the multi-dimensional output array is reshaped without ordering into a one-dimensional array by a *flattening layer*. 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.

4 Data pre-processing

For each hit in a simulated event, the PMT identifier, i.e. the relative coordinate of the hit PMT in a DOM, is recorded. Additionally, the time at which the PMT signal crosses the discriminator threshold, and the measured ToT, are stored. However, the ToT value itself is not used as input for the CNNs. In order to feed this four-dimensional event data to a CNN, the hits can be binned into rectangular pixels, such that each image encodes three spatial dimensions (XYZ) and the time dimension T.

4.1 Spatial binning

The number of pixels required to resolve the spatial coordinates of the individual PMTs inside the DOMs would be very large, and most bins would be empty due to the sparsely instrumented detector volume. Therefore, the pixelation is defined such that exactly one DOM fits into one bin, while some bins remain empty due to the detector geometry. In the case of the full KM3NeT/ORCA detector, this results in a $11 \times 13 \times 18$ (XYZ) pixel grid. An XY-projection of this pixel grid is shown in figure 1 (left), while figure 1 (right) depicts an event image used for the training of the CNN.

In this way, however, the important information regarding which PMT in a DOM has been hit would not be used, and this is corrected for 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). Such image types can also be found in classical computer vision tasks, e.g. as coloured videos. The only difference is that in a conventional video the Z-coordinate is replaced by the time and the PMT identifier is replaced by the red-green-blue colour information.

4.2 Temporal Binning

The indispensable piece of information still missing in these images is the time at which a hit has been recorded. This information can be added as an additional dimension, such that the final image of an event is five-dimensional: XYZTP.

The time resolution in KM3NeT/ORCA is of the order of nanoseconds [22]. As explained in section 2.2, the time length of an event is about $3 \mu\text{s}$, implying 3000 bins for the time dimension to reach nanosecond resolution. However, with each additional bin, the size of the event image gets larger and this leads to additional computations in the first layer of a CNN. Hence, the number of bins of the final image should be as low as possible, while still containing the relevant timing information.




Figure 1. Left: footprint of the KM3NeT/ORCA detector with the DU positions as used in the Monte Carlo simulations. The gray squares indicate the pixel grid chosen for image generation. Right: event image depicting the number of hits on each DU (XY-projection) induced by a ν_e^{CC} event (including random noise). The neutrino was up-going at an angle of roughly 45° , interacted below the detector, and had an energy of about 30 GeV.

Most hits that lie outside the time range of the triggered hits are background hits and not signal hits. Therefore, background hits can be discriminated against to some extent by selecting the time range in which most signal hits are found. Investigating the distribution of the time of the signal hits relative to the mean of the triggered hit times in individual events, as depicted in figure 2 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 the 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 figure 2, only a small fraction of the signal hits are removed, as indicated by the timecut defined by the dashed, black lines. A compromise needs to be found between the width of the timecut window and the number of time bins, implying a certain time resolution available to the network. The specific values of the timecuts used in this work are reported in the respective image generation sections of the presented CNNs. The timecut window is a parameter that can be further optimised with respect to the final performance of a trained neural network. In this work, no such parameter optimisation studies have been carried out for any of the presented CNNs, which implies that their performance for specific use cases can likely be improved further.

4.3 Multi-image Convolutional Neural Networks

After binning, the resulting images are five-dimensional: XYZTP. In order to train the neural networks presented in this paper, the deep learning framework TensorFlow [23] has been used in conjunction with the Keras [14] high-level neural network programming library. However, TensorFlow does not support convolutional layers which accept more than four dimensions as input, 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




Figure 2. Time distribution of ν_{μ}^{CC} signal hits relative to the mean time of the triggered hits calculated for each individual event. For this distribution, about 3000 ν_{μ}^{CC} events in the energy range from 3 GeV to 100 GeV have been used, cf. table 1. The dashed line in black indicates a possible timecut for the time range used to generate the CNN input images.

individual PMTs in a DOM is discarded, such that the resulting image is only four-dimensional (XYZT). However, a second image of the same event is then fed to the network (XYZP) that recovers the information regarding 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 depth dimension of a convolutional layer, the images can be stacked in this dimension. For example, an XYZP image of dimension $11 \times 13 \times 18 \times 31$ can be combined with an XYZT image of dimension $11 \times 13 \times 18 \times N_T$ into a single, stacked XYZ-T/P image of dimension $11 \times 13 \times 18 \times (N_T + 31)$. These images lack the information about the hit time for a specific PMT, if more than one hit has occurred on a DOM in an event.

Significant gains in performance for all CNN applications in this work were observed when using this stacking method, as compared to just supplying a single XYZT image. Furthermore, it will be demonstrated that networks with such input limitations can still match or outperform the KM3NeT/ORCA reconstruction algorithms as presented in ref. [2].

5 Main network architecture

Four-dimensional images that have been created from simulated events are fed as input to a CNN. All networks that have been designed for a specific task in this work share a common architecture. The CNNs consist of two main components: the convolutional part with the convolutional layers and a small fully-connected network in the end. The convolutional part consists of *convolutional blocks*, each of them containing a convolutional layer, a batch normalisation layer [24], an activation layer, and, optionally, a dropout [21] or a pooling layer. The batch normalisation layer usually enables a faster and more robust training process of deep neural networks by normalising, scaling and shifting the output of the convolutional layer. The scaling and shifting transform is controlled by learnable

parameters during the training process. Recent studies indicate that the batch normalisation method smoothes the optimisation landscape and induces a more predictive and stable behaviour of the gradients, allowing for faster training [25].

In the three-dimensional convolutional layer, the weights are initialised based on a uniform distribution whose variance is calculated according to ref. [26], while the biases are set to zero. Additionally, the kernel size is three ($3 \times 3 \times 3$), the stride, i.e. the step size in shifting the convolutional kernel, is one ($1 \times 1 \times 1$) and zero-padding (Chap. 9 in ref. [20]) is used. For the batch normalisation layers, the standard parameters from ref. [24] are used. After this, a ReLU activation layer is added. These three layers are found in all convolutional blocks that are used in this work. Additionally, optional maximum pooling and dropout layers are added. In the case of maximum pooling layers, zero-padding is not applied. A scheme of these convolutional blocks is shown in table 2.

Table 2. Scheme of a convolutional block used for all CNNs defined in this work.

Layer type	Properties
Convolution	kernel size ($3 \times 3 \times 3$), uniform initialisation [26], zero-padding
Batch normalisation	
Activation	ReLU
Maximum pooling	optional, no zero-padding
Dropout	optional

Furthermore, all models use the Adam gradient descent optimiser [27] with standard parameter values, in particular for the exponential decay rates of the first and second moment estimates, $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and the learning rate which is set to 10^{-3} . An exception is the parameter ϵ , a small constant for numerical stability, which is increased from its default value of 10^{-8} to 10^{-1} . A larger value of ϵ results in smaller weight updates after each training step. In our case, it has been observed that the network occasionally did not start to learn, depending on the random initialisation of the parameters. This could be fixed by changing the value of ϵ to 10^{-1} as suggested in ref. [28], while significant drawbacks, such as a slower training convergence due to smaller weight updates, have not been observed. The weights of the neural network are updated after one *batch* of images is passed through the network. This is known as batch gradient descent. The *batch size* is defined as the number of images contained in one batch. The batch size in the training for all presented CNNs is generally set to 64 and the learning rate, i.e. the step size in the Adam algorithm for the update of the weights is annealed exponentially.

The training of all presented networks has been executed at the TinyGPU cluster at the RRZE computing centre.² It consists of 32 nodes with 4 GPUs each. The GPUs are either Nvidia GTX1080, GTX1080Ti, RTX2080 Ti, or Tesla V100. All CNNs in this work have been trained with CUDA 10 [29]. In order to train the networks, an open-source software framework called *OrcaNet* [15] has been developed, which is intended as a high-level application programming interface on top of Keras [14], specifically suited to the large datasets that frequently occur in astroparticle physics.

²Regionales Rechenzentrum Erlangen.

6 Background classifier

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 this purpose, the employed classification algorithm is based on a Random Forest (RF) [6] method. The inputs of the RF are high-level observables (features), mainly determined from likelihood-based track and shower reconstruction algorithms. Details on the used RF, its methodology, event pre-selection requirements, and performance in rejecting atmospheric muons can be found in section 4.5 of ref. [2]. In the following, an alternative classifier based on CNNs is presented and its performance is compared to the RF classifier.

6.1 Image generation

As outlined in section 4.3, XYZ-T/P images are used as input to the network. For both event images, i.e. the XYZT and XYZP components of the stacked XYZ-T/P image, a timecut has been defined, as introduced in section 4.2. The signal hit time distribution of atmospheric muons, relative to the mean time of all triggered hits, is shown in figure 3. This distribution has a larger variance than for neutrino events shown in figure 2. The reason is that, on average, atmospheric muons traverse larger parts of the detector compared to the secondary particles of the GeV-scale neutrino interactions of interest. Hence, the timecut window for all event classes has been set conservatively based on atmospheric muon events, resulting in a width of 950 ns.




Figure 3. Time distribution of signal hits in atmospheric muon events relative to the mean time of the triggered hits calculated for each individual event. For this distribution, about 3000 atmospheric muon events have been used. The dashed black line indicates the timecut, set to an interval of $[-450 \text{ ns}, +500 \text{ ns}]$, that has been applied for the generation of the background classifier images.

The timecut for this distribution, as indicated in figure 3, has been set to an interval of $[-450 \text{ ns}, +500 \text{ ns}]$, keeping more early signal hits than late hits. The reason for this is that events which produce late hits are typically energetic and leave a longer trace in the detector. Hence, they

can be better reconstructed than less energetic atmospheric muons that only produce a few hits at the edge of the detector. Consequently, cutting away a few late hits has a small effect compared to discarding early hits from a low-energy atmospheric muon event, which already produces a low number of hits.

The number of time bins is set to 100, such that the XYZT images have dimensions of $11 \times 13 \times 18 \times 100$. The resulting time resolution of each time bin is 9.5 ns, which translates to about 2 m of photon propagation distance. Adding the information of the 31 PMTs per DOM, as described in section 4.3, yields final event images that have a dimension of $11 \times 13 \times 18 \times 131$.

6.2 Network architecture

The CNN network architecture for the background classifier is based on the three-dimensional convolutional blocks introduced in section 5, with two additional fully-connected layers, also called *dense* 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. An overview of the final network structure is shown in table 3.

Table 3. Network structure of the background classifier’s three-dimensional CNN model with XYZ-T/P input. No dropout is used due to the large training dataset of 42.6×10^6 training events.

Building block / layer	Output dimension
XYZ-T Input	$11 \times 13 \times 18 \times 100$
XYZ-P Input	$11 \times 13 \times 18 \times 31$
Final stacked XYZ-T + XYZ-P Input	$11 \times 13 \times 18 \times 131$
Convolutional block 1 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 2 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 3 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 4 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 5 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 6 (64 filters)	$11 \times 13 \times 18 \times 64$
Max pooling (2,2,2)	$5 \times 6 \times 9 \times 64$
Convolutional block 1 (128 filters)	$5 \times 6 \times 9 \times 128$
Convolutional block 2 (128 filters)	$5 \times 6 \times 9 \times 128$
Convolutional block 3 (128 filters)	$5 \times 6 \times 9 \times 128$
Convolutional block 4 (128 filters)	$5 \times 6 \times 9 \times 128$
Max pooling (2,2,2)	$2 \times 3 \times 4 \times 128$
Flatten	3072
Dense + ReLU	128
Dense + ReLU	32
Dense + Softmax	2

Initially, a CNN with three output neurons was tested, so that neutrinos, atmospheric muons and random noise events could be classified separately. However, it was observed that the three-class CNN performed slightly worse than the two-class CNN that distinguishes only neutrino events from all others. This is due to the fact that the network cannot prioritise neutrino versus non-neutrino classification in the three-class case. A mistakenly classified atmospheric muon, e.g., classified as a random noise event, has the same effect on the total loss as an atmospheric muon classified as a neutrino.

No regularisation techniques, such as dropout, are added to the network. The training dataset is large enough, cf. section 6.3, so that virtually no overfitting occurs, i.e. the training-phase loss is of the same order as the loss during the validation phase.

6.3 Preparation of training, validation and test data

For the training of the background classifier, the simulated data from table 1 is split into a training, validation and test dataset.

In order to balance the datasets with respect to their class frequency, one could split the data into 50% neutrino and 50% non-neutrino events (25% atmospheric muons + 25% random noise). Considering that the neutrino sample has the lowest number of events (about 20.7×10^6), one would have to remove a significant fraction of the 23.3×10^6 generated random noise events, and of the 65.2×10^6 atmospheric muon events for a class-balanced data splitting. On the other hand, based on the RF background classifier, it can be expected that the final accuracy of the classifier should be close to 99%. Therefore, a balanced splitting of the data into 50% for each class is not necessary, in order to avoid a local minimum during the training process. The following data splitting is used: 1/3 neutrino events, 1/3 random noise events and 1/3 atmospheric muon events, and hence the final class balance is 1/3 neutrino events and 2/3 non-neutrino events. Using this data splitting and considering the number of MC events summarised in table 1, the size of the used training dataset is larger compared to a 50/50 split. The fractions of different neutrino flavours and interaction types is kept as indicated in table 1.

This rebalanced dataset is then split into 70% training, 3% validation, and 27% test events, which is a trade-off between maximizing the training dataset and retaining sufficient statistics for performance evaluations. Additionally, the events that have been removed to balance the dataset (mostly atmospheric muons) are added to the test dataset. In total, the training data contains about 43.5×10^6 events.

Using a Nvidia Tesla V100 GPU, it takes about a week to fully train this CNN background classifier. The time needed for the training scales more weakly than linear with the number of time bins, which can be increased to improve the time resolution of the input images.

6.4 Performance and comparison to Random Forest classifier

The performance of the CNN background classifier is evaluated using the training and validation cross-entropy loss [20] of a specific classifier, the *softmax classifier* [20], as a function of the number of epochs, and is shown in figure 4. An epoch is defined as one training process of the CNN, using the entire training event dataset. The training is stopped after approximately two epochs, as the validation loss shows no further significant improvement. At the end of the training no overfitting is observed.




Figure 4. Training and validation cross-entropy loss of the background classifier during the training. Each data point of the training loss curve is averaged over 250 batches, i.e. 250×64 event images.

In order to compare the CNN performance to the RF background classifier, the same test dataset is used. A pre-selection of these events is carried out to reduce the fraction of atmospheric muon and random noise events to a few percent of all triggered events. For atmospheric muons this is achieved by selecting only events for which the reconstructed particle direction is below the horizon, i.e. which are *up-going* events in the detector. Furthermore, the events must have been reconstructed with high quality by the KM3NeT/ORCA maximum-likelihood reconstruction algorithms for either track-like or shower-like events. Finally, events reconstructed by the maximum-likelihood-based algorithms as originating from outside of the instrumented volume of the detector are removed.

In total, the pre-selected test dataset consists of about 3.3×10^6 neutrino, about 6×10^4 atmospheric muon and about 4×10^4 random noise events. This selection is used for all of the following performance evaluations.

In order to get a first impression of the CNN-based background classifier, the distribution of the neutrino class probability is investigated for all three event classes (neutrinos, atmospheric muons, random noise), cf. figure 5.

Based on the results shown in figure 5, it can be seen that the rate of random noise events misclassified as neutrino-induced events is significantly lower than for atmospheric muons. Using the predicted probability for an event to be classified as neutrino-induced, a threshold value p can be set to remove background events.

In order to quantify the performance of the CNN background classifier, the metric shown below is used to investigate the fraction of remaining atmospheric muon and random noise events for a given threshold value p . The atmospheric muon or random noise contamination, and the neutrino efficiency, are defined as:

$$C_{\mu/\text{RN}}(p) = \frac{N_{\mu/\text{RN}}(p)}{N_{\text{total}}(p)}, \quad (6.1)$$




Figure 5. Distribution of the CNN neutrino probability for pre-selected atmospheric muon (blue), random noise (red), and neutrino (brown) events. All three distributions have been normalised to the area under each histogram.

$$\nu_{\text{eff}}(p) = \frac{N_{\nu}(p)}{N_{\nu, \text{total}}} . \quad (6.2)$$

Here, $N_{\mu/\text{RN}}$ is the number of atmospheric muon or random noise events, whose probability to be a neutrino-induced event is higher than p , while $N_{\text{total}}(p)$ accounts for the total number of events, after the same cut on p . Regarding the neutrino efficiency, $N_{\nu}(p)$ is the total number of neutrinos in the dataset whose neutrino probability is greater than the threshold value p , while $N_{\nu, \text{total}}$ is the number of neutrinos in the dataset, without applying any threshold.

Based on the results shown in figure 6 for the neutrino efficiency ν_{eff} as a function of the residual atmospheric muon contamination C_{μ} , it can be concluded that the CNN background classifier yields a higher neutrino efficiency, of the order of a few percent, for the same muon contamination compared to the RF background classifier.

Comparing different neutrino energy ranges, 1 GeV to 5 GeV in figure 7 and 10 GeV to 20 GeV in figure 8, it can be seen that the performance gap between the CNN and the RF classifier widens with increasing neutrino energy. A possible explanation may be that small details in the distribution of the measured hits increase in importance if the neutrino events are less energetic and thus produce less hits. Then, the limitations of the input images, which do not contain the full information about an event, may become more relevant compared to events with higher energies and many signal hits.

For random noise events, the performance of the CNN and the RF classifier are comparable. In particular, both methods achieve about 99% neutrino efficiency at 1% random noise event contamination. As expected, the suppression of atmospheric muon events is significantly more difficult.




Figure 6. Neutrino efficiency, ν_{eff} , versus atmospheric muon event contamination, C_μ , weighted with the Honda atmospheric neutrino flux [17] model. The CNN (RF) performance is depicted in blue (orange). All neutrinos from the pre-selected test dataset are included (1 GeV to 100 GeV).




Figure 7. Neutrino efficiency, ν_{eff} , versus atmospheric muon event contamination, C_μ , weighted with the Honda atmospheric neutrino flux [17] model. The CNN (RF) performance is depicted in blue (orange). Only neutrinos with a MC energy in the range of 1 GeV to 5 GeV have been used.

7 Event topology classifier

Similar to the background classifier, a RF is used in the current KM3NeT/ORCA analysis pipeline to separate track-like from shower-like neutrino events. This section introduces a CNN-based event topology classifier that distinguishes between these two event types.




Figure 8. Neutrino efficiency, ν_{eff} , versus atmospheric muon event contamination, C_μ , weighted with the Honda atmospheric neutrino flux [17] model. The CNN (RF) performance is depicted in blue (orange). Only neutrinos with a MC energy in the range of 10 GeV to 20 GeV have been used.

7.1 Image generation

The input event images for the CNN-based track-shower classifier are similar to the ones of the background classifier introduced in section 6.1, i.e. the input also consists of XYZ-T/P images. The timecut for the hit selection is tighter with respect to the background classifier. Since background events have already been rejected by the background classifier, the data presented to the track-shower classifier are mostly neutrino interactions which produce secondary particles that, on average, traverse smaller parts of the detector compared to atmospheric muon events. The event class that shows the broadest signal hit time distribution is the class of ν_μ^{CC} events due to the outgoing muon. Therefore, the timecut of the track-shower classifier is set based on these events. The time distribution of signal hits relative to the mean time of the triggered hits for ν_μ^{CC} events is shown in figure 2. Based on this distribution the timecut is set to an interval of $[-250 \text{ ns}, +500 \text{ ns}]$, as indicated by the dashed, black lines in figure 2.

Since the timecut interval is smaller than the one used for the background classifier, less time bins (60) are used for the time dimension. This implies a reduction of the time resolution of about 30% with respect to the background classifier, i.e. 12.5 ns per time bin. The light emission profile of hadronic and electromagnetic showers in the GeV range has an extension of at most a few metres, see figure 70 in ref. [2]. Since a muon of comparable energy induces the emission of Cherenkov radiation along a significantly greater path length, the reduced time binning still provides a sufficient resolution to distinguish a shower-like from a track-like event topology, while significantly speeding up the training of the CNN. Consequently, the XYZT images now have $11 \times 13 \times 18 \times 60$ pixels.

7.2 Network architecture

The network architecture, as depicted in table 4, is the same as in section 6.2, except for additional dropout layers. Since the size of the training dataset is significantly smaller than that for the background classifier, overfitting can be observed without any regularisation. Thus, dropout layers with a rate of $\delta = 0.1$ are added in every convolutional block and also in between the last two fully connected layers.

Table 4. Network structure of the track-shower classifier three-dimensional CNN model with XYZ-T/P input. The symbol δ specifies the dropout rate used in the respective convolutional block, cf. section 3.

Building block / layer	Output dimension
XYZ-T Input	$11 \times 13 \times 18 \times 60$
XYZ-P Input	$11 \times 13 \times 18 \times 31$
Final stacked XYZ-T + XYZ-P Input	$11 \times 13 \times 18 \times 91$
Convolutional block 1 (64 filters, $\delta = 0.1$)	$11 \times 13 \times 18 \times 64$
Convolutional block 2 (64 filters, $\delta = 0.1$)	$11 \times 13 \times 18 \times 64$
Convolutional block 3 (64 filters, $\delta = 0.1$)	$11 \times 13 \times 18 \times 64$
Convolutional block 4 (64 filters, $\delta = 0.1$)	$11 \times 13 \times 18 \times 64$
Convolutional block 5 (64 filters, $\delta = 0.1$)	$11 \times 13 \times 18 \times 64$
Convolutional block 6 (64 filters, $\delta = 0.1$)	$11 \times 13 \times 18 \times 64$
Max pooling (2,2,2)	$5 \times 6 \times 9 \times 64$
Convolutional block 1 (128 filters, $\delta = 0.1$)	$5 \times 6 \times 9 \times 128$
Convolutional block 2 (128 filters, $\delta = 0.1$)	$5 \times 6 \times 9 \times 128$
Convolutional block 3 (128 filters, $\delta = 0.1$)	$5 \times 6 \times 9 \times 128$
Convolutional block 4 (128 filters, $\delta = 0.1$)	$5 \times 6 \times 9 \times 128$
Max pooling (2,2,2)	$2 \times 3 \times 4 \times 128$
Flatten	3072
Dense + ReLU	128
Dropout ($\delta = 0.1$)	128
Dense + ReLU	32
Dense + Softmax	2

7.3 Preparation of training, validation and test data

In order to train the CNN track-shower classifier, only simulated neutrino events are used. The total neutrino dataset is rebalanced 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 such a way that the ratio of track-like to shower-like events is always one, independent of neutrino energy.

The rebalanced dataset is then split into three datasets with 70% training, 6% validation, and 24% test events. In total, the training dataset contains about 14×10^6 events.

Using a Nvidia Tesla V100 GPU, it takes about one and a half weeks to fully train this CNN track-shower classifier.

7.4 Performance and comparison to Random Forest classifier

The evolution of the cross-entropy loss [20] during the training is shown in figure 9. Even though




Figure 9. Training (grey lines) and validation (black circles) cross-entropy loss of the track-shower classifier during the training process. Each line element of the training loss represents an average over 250 batches, i.e. 250×64 event images.

the validation cross-entropy loss is within the fluctuations of the training loss curve at the end of the training, some minor overfitting occurs. The reason is that during the application of the trained network on the validation dataset, no neurons are dropped by the dropout layers, contrary to the training phase. Therefore, the validation loss should be lower than the average training loss, if no overfitting is observed. This can be seen by investigating the training and validation loss curves in the earlier stages of the training process, e.g. between epoch 0 and 3. Here, the validation loss is typically found at the bottom of the training loss curve.

The binned probability distribution for all used neutrino events with energies in the range of 1 GeV to 40 GeV to be classified as track-like is shown in figure 10. This energy range has been chosen since the classification performance saturates at about 40 GeV, as can be seen in figure 11 and further discussed below. The classified neutrino events have been selected according to the criteria described in section 6.4.

About 25% of ν_{μ}^{CC} and $\bar{\nu}_{\mu}^{\text{CC}}$ events are identified as track-like events with a probability close to one. The correct identification of muon tracks increases with their length and hence with their energy. As the outgoing muon has on average higher energy for $\bar{\nu}_{\mu}^{\text{CC}}$ than for ν_{μ}^{CC} events, $\bar{\nu}_{\mu}^{\text{CC}}$ events have a higher probability to be identified as track-like.

The top panel of figure 11 shows the fraction of events classified as track-like as a function of the neutrino energy in the range of 1 GeV to 40 GeV. An event is accepted if its CNN probability




Figure 10. Distribution of the CNN track probability $P_{i,\text{track}}$ for pre-selected neutrino events of different flavours and interaction channels in the energy range of 1 GeV to 40 GeV. All distributions have been normalised to the area under each histogram.

to be a track-like event is 0.5 or above. The fraction of correctly classified ν_μ^{CC} and $\bar{\nu}_\mu^{\text{CC}}$ events rises strongly with neutrino energy. The reason is that the spatial and temporal distribution of hits induced by low-energy ν_μ^{CC} events is similar to that of shower-like events, since the outgoing muon stops at these energies after propagating only a few meters in the detector. Comparing ν_e^{CC} and $\bar{\nu}_e^{\text{CC}}$ with energies below 10 GeV, it can be seen that $\bar{\nu}_e^{\text{CC}}$ are more likely classified as track-like by the CNN classifier. This is again due to the fact that the outgoing charged lepton carries on average more energy in $\bar{\nu}_e^{\text{CC}}$ interactions with respect to ν_e^{CC} interactions. On the other hand, for energies below 10 GeV, NC interactions have a lower misclassification rate compared to ν_e^{CC} interactions, since there is no charged secondary lepton that could mimic the signature of a low-energy muon as induced in ν_μ^{CC} events. For energies above 10 GeV the opposite is true, since a localised and bright light source disfavours a ν_μ^{CC} event.

A measure for the separation power of the classifier is the difference between the fractions of recognised track-like events for ν_μ^{CC} ($f_{\text{track}}^{\nu_\mu^{\text{CC}}}$, blue line) and ν_e ($f_{\text{track}}^{\nu_e^{\text{CC/NC}}}$, red and brown lines) in the top panel of figure 11. In the bottom panel the ratio of these differences, as determined by the CNN and RF classifiers, i.e. the relative improvement Δf of the CNN with respect to the RF classifier, is shown for different channel combinations:

$$\Delta f = \left(f_{\text{track}}^{\nu_\mu^{\text{CC}}} - f_{\text{track}}^{\nu_e^{\text{CC/NC}}} \right)_{\text{CNN}} \bigg/ \left(f_{\text{track}}^{\nu_\mu^{\text{CC}}} - f_{\text{track}}^{\nu_e^{\text{CC/NC}}} \right)_{\text{RF}}. \quad (7.1)$$

It can be seen that the separation power between track-like and shower-like events depends on the neutrino energy. The CNN performs better than the RF in the energy range of roughly 3 GeV to 10 GeV, while the RF performs better than the CNN for energies below 2 GeV. As explained at the end of section 6.4, this could again be due to limitations in the input images that become more




Figure 11. Fraction of events classified as track-like, f_{track} , with a CNN track probability $P_{i,\text{track}} > 0.5$, for different interaction channels (top panel) versus the true MC neutrino energy. The relative improvement Δf , cf. eq. (7.1), of the CNN with respect to the RF classifier in discriminating between ν_{μ}^{CC} events and shower-like $\nu_e^{\text{CC}}/\nu_e^{\text{NC}}$ channels is shown in the bottom panel. Here, the contributions of neutrinos and antineutrinos are averaged for each flavour and interaction. The discrimination power is defined as the difference between the fractions of events classified as track-like and shower-like in the top panel.

relevant for events with only a few hits. Above 15 GeV the separation power of the two classifiers is roughly equal.

Since the performance comparison shown in figure 11 depends on the threshold value that is set for an event to be classified as track-like, another comparison metric independent of the threshold value is introduced. $P_{\text{track}}^{\nu_{\mu}}$ is defined as the probability that a classifier recognises a ν_{μ}^{CC} event as track-like, and similarly for ν_e^{CC} . Based on figure 10, the ability of a classifier to separate track-like and shower-like events can be estimated by calculating the energy-dependent correlation factor $c(E)$ of the two binned probability distributions $P_{i,\text{track}}^{\nu_{\mu}}(E)$ and $P_{i,\text{track}}^{\nu_e}(E)$. The *separability* $S(E)$ is then given by:

$$S(E) = 1 - c(E) = 1 - \frac{\sum_i P_{i,\text{track}}^{\nu_{\mu}}(E) \cdot P_{i,\text{track}}^{\nu_e}(E)}{\sqrt{\sum_i \left(P_{i,\text{track}}^{\nu_{\mu}}(E)\right)^2 \cdot \sum_i \left(P_{i,\text{track}}^{\nu_e}(E)\right)^2}}, \quad (7.2)$$

where the index i in $P_{i,\text{track}}^{\nu_{\mu}/\nu_e}$ is the i -th bin of the distribution $P_{\text{track}}^{\nu_{\mu}/\nu_e}$ as shown in figure 10.

The separability $S(E)$ as a function of binned neutrino energy for the CNN and RF classifier is shown in figure 12. It is evident that the CNN-based classifier performs better by an absolute value of up to about 20% for energies below 10 GeV. For these energies, the RF- and CNN-based

classifier show both a better separability of ν_μ^{CC} and ν_e^{NC} compared to ν_e^{CC} events, although this loss in separability is less pronounced and reverses above 7 GeV for the RF classifier, i.e. at slightly lower energies than for the CNN classifier.




Figure 12. Separability $S(E) = 1 - c(E)$ based on ν_μ^{CC} and ν_e^{CC} (ν_μ^{CC} and ν_e^{NC}) events as a function of true MC neutrino energy for the CNN represented by blue (brown) markers and the RF represented by red (green) markers (top panel). Only ν (not $\bar{\nu}$) have been used for the comparison. The relative improvement of the CNN with respect to the RF classifier in discriminating between ν_μ^{CC} events and shower-like $\nu_e^{\text{CC/NC}}$ channels is shown in the bottom panel. The error bars indicate statistical errors only.

8 Event regression

This section introduces the CNN-based *event regressor*. The designed CNN now reconstructs continuous properties of the interacting neutrinos such as the direction, energy, the position of the interaction, i.e. the vertex, and an estimate for the uncertainty of each of these inferred observables. In the reconstruction pipeline for KM3NeT/ORCA as described in ref. [2], these reconstruction tasks are handled by two maximum-likelihood-based algorithms optimised for track-like and shower-like events, respectively.

The major difference of the CNN for the regression task with respect to the classifiers in the previous sections is that its output is a continuous instead of a categorical variable.

8.1 Image generation

For the event regressor, the event images are created in a similar way to the image generation for the background and the track-shower classifier. Thus, both XYZT and XYZP images are produced

and then stacked in the last dimension. Since the background classifier has already been applied to the data, at this final stage of the reconstruction pipeline neutrinos remain predominantly as input to the event regressor. Hence, exactly the same images as specified in section 7 are used as input for the CNN event regressor.

8.2 Network architecture and loss functions

The network architecture of the CNN for the event regressor is similar to the one used for the background and track-shower classifier. However, for this work, it has been found that dropout layers in the CNN event regressor increase fluctuations in the training loss, even with a small dropout rate of $\delta = 0.05$. Furthermore, the validation loss of the event regressor CNNs with dropout layers has proven to be significantly worse than for CNNs without dropout. Thus, no dropout layers are included. Similar observations concerning the effect of dropout in a regression problem have been reported in ref. [30].

The fully-connected layers after the convolutional layers are also different. The properties that should be reconstructed by the network are the energy, the direction, and the vertex position of the initial neutrino interaction. For the direction and the vertex, the CNN output consists of two arrays with three elements each, containing the components of the vertex position vector and the direction cosines, respectively. Consequently, the output of the network consists of seven reconstructed floating-point numbers, denoted as $\vec{\xi}_{\text{reco}}$, one component for the energy, three for the direction and three for the vertex position. The final fully-connected output layer therefore consists of seven neurons. The CNN architecture up to this point is shown in table 5.

Since the aim of this network is to reconstruct continuous variables, a categorical cross-entropy loss [20], as for the background and track-shower classifier, cannot be used. Instead, other cross-entropy-based loss functions can be employed, e.g. the mean squared error (MSE) or the mean absolute error (MAE) [20], which are defined as:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (\xi_{\text{reco}} - \xi_{\text{true}})^2, \quad (8.1)$$

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |\xi_{\text{reco}} - \xi_{\text{true}}|. \quad (8.2)$$

The main difference between the MSE and the MAE loss is that the MSE loss penalises outlier events that are badly reconstructed with a significantly larger loss value compared to the MAE loss. Since many detectable events are not fully contained in the KM3NeT/ORCA detector, these events are on average reconstructed more poorly. Thus, the MAE loss is used for the energy, direction and vertex estimation, so that the network focuses less on outlier events.

Using a neural network, it is also possible to estimate the uncertainties on the seven components of $\vec{\xi}_{\text{reco}}$. One possibility is to let the network predict the absolute residual of the MAE loss. Thus, one needs an additional neuron that yields the estimated uncertainty, σ_{reco} , for each of the reconstructed components of $\vec{\xi}_{\text{reco}}$. Consequently, the loss function that measures the quality of the estimated uncertainty value is:

$$L = \frac{1}{n} \sum_{i=1}^n (\sigma_{\text{reco}} - |\xi_{\text{true}} - \xi_{\text{reco}}|)^2. \quad (8.3)$$

Table 5. Network structure of the CNN event regressor with XYZ-T/P input, but without the additional fully-connected network for the uncertainty estimation. In the convolutional blocks and the last fully-connected layers no dropout layer is used.

Building block / layer	Output dimension
XYZ-T Input	$11 \times 13 \times 18 \times 60$
XYZ-P Input	$11 \times 13 \times 18 \times 31$
Final stacked XYZ-T + XYZ-P Input	$11 \times 13 \times 18 \times 91$
Convolutional block 1 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 2 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 3 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 4 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 5 (64 filters)	$11 \times 13 \times 18 \times 64$
Convolutional block 6 (64 filters)	$11 \times 13 \times 18 \times 64$
Max pooling (2,2,2)	$5 \times 6 \times 9 \times 64$
Convolutional block 1 (128 filters)	$5 \times 6 \times 9 \times 128$
Convolutional block 2 (128 filters)	$5 \times 6 \times 9 \times 128$
Convolutional block 3 (128 filters)	$5 \times 6 \times 9 \times 128$
Convolutional block 4 (128 filters)	$5 \times 6 \times 9 \times 128$
Max pooling (2,2,2)	$2 \times 3 \times 4 \times 128$
Flatten	3072
Dense + ReLU	128
Dense + ReLU	32
Dense + Linear	7

Using this loss function L , the network learns to estimate the average absolute residual:

$$\sigma_{\text{reco}} \approx \langle |\xi_{\text{true}} - \xi_{\text{reco}}| \rangle. \quad (8.4)$$

Assuming that the residuals are normally distributed, the estimated uncertainty value σ_{reco} can be converted to a standard deviation by multiplying by $\sqrt{\frac{\pi}{2}}$ [31].

In general, it can be expected that features that are learned by the convolutional part of the network and that are useful for the reconstruction are also crucial for the inference of the uncertainty on the reconstruction. Therefore, both inference tasks are handled by using the same CNN. The CNN, however, must be prevented from focusing too much on the uncertainty estimation in the learning process, since its main goal is the prediction of $\tilde{\xi}_{\text{reco}}$ and not of the uncertainties. For this purpose, a second fully-connected network with three layers is added after the convolutional output, but the gradient update during the back-propagation stage is stopped after the first fully-connected layer of this sub-network, as shown in figure 13. The specific structure of the fully-connected network for the uncertainty reconstruction estimation is shown in table 6.




Figure 13. Scheme of the CNN architecture used for the event regressor. The properties of the convolutional layers and of the fully-connected network for the reconstruction of $\vec{\xi}_{\text{reco}}$ is depicted in table 5 and in table 6.

Table 6. Layout of the fully-connected uncertainty reconstruction network, which is part of the event regressor CNN with XYZ-T/P input. The input to this dense network is the flattened output of the last convolutional layer of the main CNN.

Building block / layer	Output dimension
Dense + ReLU	128
Dense + ReLU	64
Dense + ReLU	32
Dense + Linear	7

8.3 Preparation of training, validation and test data

In order to train the CNN event regressor, again only simulated neutrino events are used. Similar to the track-shower classifier, the dataset consists of 50% track-like events (ν_{μ}^{CC}) and 50% shower-like events (25% ν_e^{CC} , 25% ν_e^{NC}). Additionally, the total dataset has been split again into a training, validation and test dataset. As for the track-shower classifier, the training data makes up 70%, the validation dataset 6% and the test dataset 24% of the total rebalanced dataset. The track and shower classes have not been balanced to be independent of neutrino energy, cf. section 7.3, in order to get a larger training dataset compared to the track-shower classifier. In total, the training dataset then contains about 14×10^6 events.

The MC energy of ν_e^{NC} events is set to the visible energy during the training, since the network cannot infer how much energy is carried away by the outgoing neutrino in the neutral current interaction.

Using a Nvidia Tesla V100 GPU, it takes about one and a half weeks to fully train this CNN event regressor.

8.4 Loss functions and loss weights

Before training the CNN event regressor, suitable *loss weights* need to be defined in order to scale the losses to the same magnitude. For example, the neutrino energy in the dataset ranges from 1 GeV to 100 GeV, while the components of the direction vector range from -1 to 1 . Thus, the impact of the

energy loss would be overwhelmingly large compared to the loss that is contributed by the direction reconstruction. Consequently, one needs to scale the individual losses of the components of $\vec{\xi}_{\text{reco}}$ and their estimated uncertainties with a loss weight, such that all single losses are of the same order of magnitude at the start of the training. However, it has been observed that the CNN does not start to learn during the first few epochs of the training, if the individual losses of the energy, direction and vertex are of the same scale. The reason is that direction and vertex contribute six out of seven variables to be learned so that their impact on the total loss is dominant. Hence, any improvement with respect to the energy reconstruction, i.e. in the discrimination of background and signal hits, is marginalised by random statistical fluctuations in the loss for the direction or vertex components.

It is expected that the first concept the neural network will learn during the training is to distinguish signal and background hits. At the same time, the recognition of signal hits is a prerequisite for the direction and vertex reconstruction. Thus, the loss weights have been set in such a way that the scale of the energy loss at the start of the training is about two times larger than the scale of the three combined directional losses. For the same reason, the loss weights of the vertex reconstruction are set in such a way that they have the same scale as the directional losses. Indeed, it could be observed that the energy loss always converges significantly earlier during the training than the loss for the direction or vertex. Since the scale of the individual losses changes dramatically after the network has started to learn a property such as the energy or the direction, the loss weights need to be tuned during the training.

As a general strategy, it has proven useful to do this after the energy loss has sufficiently converged. How often and by how much the individual loss weights are retuned is a parameter optimisation process. For this work, the loss weights of the direction and vertex components have been increased by a factor of three after the energy loss convergence, and, at the end of the training, these loss weights have been doubled again to investigate if further improvements would occur, which was not the case. Increasing the weights for the direction and vertex losses did not significantly worsen the energy loss.

For the loss weights of the uncertainty outputs none of this matters, since their gradients are stopped after the first fully-connected layer of the uncertainty reconstruction sub-network. Thus, the loss weights for these variables just need to be set such that the scale of the individual losses is about the same.

8.5 Energy reconstruction performance

In this section the neutrino energy reconstruction performance of the CNN event regressor is investigated. The study and the comparisons to the maximum-likelihood-based reconstruction algorithm, described in ref. [2] and called *SOR* in the following, focus on shower-like events. More details on the event generation and event characterisation in terms of expected numbers of PMT and DOM hits per GeV of shower energy can be found in section 4.2.2 in ref. [2].

Events have been selected according to the criteria detailed in section 6.4. In particular, only events reconstructed as up-going by the *SOR* algorithm are selected for the comparison and, additionally, the vertex position, direction, and energy of true ν_e^{CC} events must have been reconstructed with high confidence by *SOR*.

The reconstructed energy versus the true MC energy for reconstructed ν_e^{CC} events is shown in figure 14. As can be seen in the figure, the energy reconstruction tends to underestimate




Figure 14. Energy reconstructed by the CNN regressor versus true MC neutrino energy for reconstructed ν_e^{CC} events. The distribution has been normalised to unity in each true energy bin.

the true MC energy of the neutrino, in particular for energies above 30 GeV. This effect has been found to be even more evident for SOR, so that a correction function has been applied to the reconstructed energy. This correction function has been derived from MC simulations and depends on several reconstructed quantities, such as the energy, the zenith angle, and the interaction inelasticity (section D.1 in ref. [32]). So as to allow for a comparison between the SOR algorithm and the CNN event regressor, a similar but simpler correction function, only using the reconstructed energy, has been derived and applied to the energy reconstruction of the CNN event regressor for ν_e^{CC} events. The resulting distributions can be compared in figure 15.




Figure 15. Reconstructed and corrected energy versus true MC neutrino energy for reconstructed ν_e^{CC} events. The distribution has been normalised to unity in each true energy bin. Left: CNN event regressor. Right: maximum-likelihood-based reconstruction algorithm SOR for shower-like events as described in ref. [2].

In order to allow for quantitative comparisons, the following two performance metrics are used: the *median relative error* (MRE) and the *relative standard deviation* (RSD). Both metrics are calculated

for each MC energy bin. The MRE is defined as:

$$\text{MRE} = \text{Median} \left(\left| \frac{E_{\text{reco}} - E_{\text{true}}}{E_{\text{true}}} \right| \right), \quad (8.5)$$

where E_{reco} (E_{true}) stands for the reconstructed (true MC) energy.

For the RSD, the standard deviation σ of the reconstructed energy distribution for each MC true energy bin is calculated and then divided by the energy of the bin:

$$\text{RSD} = \frac{\sigma}{E_{\text{true}}}. \quad (8.6)$$

For shower-like events, the energy reconstruction performance is intrinsically limited by light yield fluctuations of the hadronic particle cascade [5]. Assuming that both the CNN event regressor and the SOR algorithm come close to this intrinsic resolution limit, it is to be expected that the MRE and RSD performances for both methods agree very well, as can be seen in figure 16 (left) and figure 16 (right), respectively. For the MRE, the CNN shows lower values than the SOR algorithm for energies close to the boundaries of the investigated energy range. This is due to the fact that the CNN learns the limits of the simulated neutrino energy spectrum, i.e. 1 GeV and 100 GeV. Therefore, the CNN, contrary to the SOR algorithm, never reconstructs energies outside of this energy range, which biases the MRE for the CNN towards lower values. For the high-energy boundary, this effect can be easily avoided by extending the simulated energy range. For the low-energy boundary, this is not easily possible and biases the comparison with the SOR algorithm, albeit only in an energy range that is close to the detection energy threshold of KM3NeT/ORCA. On the other hand, it is notable that the RSD for the CNN event regressor improves by a few percent with respect to the SOR algorithm also for energies far from the boundaries. In the case of ν_e^{NC} events,




Figure 16. Left: median relative error (MRE) on the reconstructed energy versus true MC energy for ν_e^{CC} events for the CNN event regressor (blue) and SOR (orange). Right: relative standard deviation (RSD) on the reconstructed energy versus true MC energy for the CNN event regressor (blue) and SOR (orange).

the energy reconstruction performance is limited by the non-detectable energy that is carried away by the outgoing neutrino. For ν_e^{NC} events, as for charged-current events, the energy reconstruction performance of the CNN event regressor and SOR is similar.

For track-like ν_μ^{CC} events, the RSD on the reconstructed energy distribution is almost flat at 25% to 28% in the true energy range from 10 GeV to 100 GeV. The MRE on the reconstructed

energy is flat at 15% in the range from 10 GeV to 30 GeV and then rises slowly to 20% at 80 GeV true neutrino energy, since the detectable fraction of light from the muon decreases as the path length of the muon increases with energy.

8.6 Direction reconstruction performance

The measurement of not only the energy but also of the zenith angle of up-going neutrinos is key to achieve the scientific goals of KM3NeT/ORCA. The direction reconstruction performance of the CNN event regressor is investigated in the same way as for the energy reconstruction. The X, Y and Z components of the reconstructed direction vector are converted to spherical azimuth and zenith coordinates. Since the DUs, cf. section 2.1, are approximately vertical structures, the azimuth angle is defined by the projection of the direction vector onto a plane roughly perpendicular to the DUs, while the zenith angle defined here corresponds to the elevation angle with respect to the horizon, i.e. to $\pi/2$ minus the angle enclosed by the direction vector pointing back to the particle's source and the DUs. The zenith angle reconstructed by the CNN event regressor versus the true zenith angle for up-going ν_e^{CC} events is depicted in figure 17 (left). The distribution shows a clear correlation between the true and reconstructed zenith angle without bias.

The median absolute error (ME) is defined as the median of the distribution of $\xi_{\text{reco}} - \xi_{\text{true}}$. In figure 17 (right) the ME for the azimuth and the zenith angle of the reconstructed direction is compared for the CNN event regressor and SOR. The comparison is done for ν_e^{CC} events and as a function of the true MC neutrino energy. For energies below 3 GeV, the ME is very similar for




Figure 17. Left: reconstructed zenith angle versus true zenith angle for up-going ν_e^{CC} events. A zenith angle of 90° corresponds to a vertically up-going particle. Right: median absolute error of the zenith and azimuth angle reconstruction for the CNN event regressor and the shower reconstruction algorithm SOR versus true MC energy for ν_e^{CC} events.

both reconstructions, while the difference in performance increases in favour of the SOR algorithm with increasing energy. At 20 GeV, the ME for the reconstruction of the neutrino zenith (azimuth) angle is about 15% (25%) larger for the CNN event regressor than for SOR. The reason for this worse performance of the CNN might be connected to the rather coarse time binning chosen for the image generation. Since the same input images as for the track-shower classifier have been used for the event regressor, the timing resolution of hits on individual DOMs is 12.5 ns. While this is

sufficient for the track-shower classifier as explained in section 7.1, it is more than a factor of two worse than the full time resolution available and exploited by the SOR algorithm. Moreover, time is treated only as a channel input in the image generation. The corresponding comparison for ν_e^{NC} events shows the same general trend, but the performance gap is smaller since the reconstruction is intrinsically limited by the interaction kinematics.

For track-like ν_μ^{CC} events, the median error on the CNN-reconstructed neutrino direction (zenith angle) decreases from 12° (7°) at 5 GeV to 7° (4°) at 20 GeV, and becomes essentially flat at 4° (2.5°) for true energies above 50 GeV.

8.7 Vertex reconstruction performance

For the comparison of the vertex reconstruction performance the distance between the reconstructed and the true MC vertex point is investigated based on the longitudinal and perpendicular components of the residual vector with respect to the direction of the incoming neutrino. The CNN event regressor is trained such that it reconstructs the neutrino interaction vertex, while the SOR algorithm reconstructs the brightest point of the Cherenkov light emission, which is shifted by an energy-dependent displacement of the order of meters.

The perpendicular versus longitudinal distance of the reconstructed vertex with respect to the true ν_e^{CC} interaction point is shown in figure 18 (left) for the CNN event regressor, while the right plot shows the same distribution for the SOR algorithm. The above mentioned displaced reconstruction of the brightest emission point can be seen for the SOR algorithm in figure 18 (right). The CNN-based vertex resolution is of the order of one meter with a small offset in perpendicular direction in the investigated energy range up to 100 GeV. For the SOR algorithm, the resolution in the longitudinal direction is similar, while it is better in the perpendicular direction. A correlation of longitudinal and perpendicular errors can be seen for the SOR algorithm, which is absent for the CNN reconstruction.




Figure 18. Reconstructed neutrino interaction point with respect to the true MC vertex for ν_e^{CC} events. Shown is the perpendicular component versus the longitudinal component of the distance vector. Left: CNN event regressor; right: SOR algorithm.

8.8 Error estimation

The CNN is used to reconstruct the regression uncertainty of all reconstructed quantities, as explained in section 8.2. For the following performance investigation the event pre-selection, based on the current KM3NeT/ORCA reconstruction pipeline, is not applied, i.e. all triggered events of the simulated test dataset are used. In order to evaluate the goodness of the uncertainty reconstruction for any event quantity, e.g. energy or direction, the events are binned according to their estimated uncertainty value, σ_{reco} . For the event distribution in each bin, the standard deviation, σ_{true} , of the residuals of the reconstructed and true MC values $|\xi_{\text{reco}} - \xi_{\text{true}}|$ is calculated.

The distribution of σ_{true} versus σ_{reco} is shown in figure 19 (left) for the energy reconstruction of ν_e^{CC} events. Even though σ_{reco} is estimated for all triggered events and not only for those pre-selected as explained in section 8.5, the CNN event regressor estimation of the reconstruction uncertainty is clearly correlated to the true uncertainty. However, for events that are difficult to reconstruct for the CNN event regressor, and that consequently show large values of σ_{true} , also the uncertainty estimation is more difficult and a slight underestimation is visible.

The distribution of σ_{true} versus σ_{reco} is shown in figure 19 (right) for the Z component of the reconstructed direction vector in ν_e^{CC} events. In the range of $0.15 < \sigma_{\text{reco}} < 0.30$ a slight underestimation in the uncertainty reconstruction can be seen.




Figure 19. Standard deviation of the difference of true MC and reconstructed values versus the uncertainty estimated by the CNN event regressor. Left: for the energy reconstruction of ν_e^{CC} events. Right: for the Z component of the reconstructed direction vector for ν_e^{CC} events.

Since there is a clear correlation between the predicted and the true uncertainty, it is possible to discriminate badly reconstructed events based on the estimated uncertainty value σ_{reco} . In order to demonstrate this, the zenith-angle reconstruction for shower-like events is investigated. The standard deviation of the residual distribution of the reconstructed cosine of the zenith angle for ν_e^{CC} events versus the fraction of events discarded by the CNN uncertainty estimator is depicted in figure 20. For all three energy intervals depicted, one can see that the zenith-angle resolution improves significantly when the events are discarded, that have the largest reconstruction errors as predicted by the CNN.




Figure 20. 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.

9 Conclusion

In this work, the first application of deep convolutional neural networks to the reconstruction and classification of simulated neutrino events in the neutrino detector KM3NeT/ORCA has been reported. Detailed MC datasets have been suitably pre-processed to generate pixelated high-dimensional image input for the training of CNNs. CNNs have been designed to separate neutrino events from background, and track-like from shower-like neutrino events. In addition, CNNs have been used to reconstruct the direction, energy, and interaction point of the neutrinos, and to estimate the uncertainty on each of the reconstructed quantities. Comparisons to the performance of classical machine-learning approaches to classification and maximum-likelihood reconstructions show the competitive and in some cases already superior performance of the employed CNNs, though there are still many aspects that can possibly be optimised. In the next step, the developed CNNs will be tested for robustness with respect to real data acquisition conditions, e.g. for the effect of missing optical modules and imperfect calibrations, before they will be applied to data taken with the KM3NeT/ORCA detector.

Acknowledgments

The authors gratefully acknowledge the compute resources and support provided by the Erlangen Regional Computing Center (RRZE). Significant parts of this work are based on the doctoral thesis “Sensitivity studies on tau neutrino appearance with KM3NeT/ORCA using Deep Learning Techniques” by M. Moser.

The authors acknowledge the financial support of the funding agencies: Agence Nationale de la Recherche (contract ANR-15-CE31-0020), Centre National de la Recherche Scientifique (CNRS),

Commission Européenne (FEDER fund and Marie Curie Program), Institut Universitaire de France (IUF), LabEx UnivEarthS (ANR-10-LABX-0023 and ANR-18-IDEX-0001), Paris Île-de-France Region, France; Shota Rustaveli National Science Foundation of Georgia (SRNSFG, FR-18-1268), Georgia; Deutsche Forschungsgemeinschaft (DFG), Germany; The General Secretariat of Research and Technology (GSRT), Greece; Istituto Nazionale di Fisica Nucleare (INFN), Ministero dell'Università e della Ricerca (MUR), PRIN 2017 program (Grant NAT-NET 2017W4HA7S) Italy; Ministry of Higher Education, Scientific Research and Professional Training, Morocco; Nederlandse organisatie voor Wetenschappelijk Onderzoek (NWO), the Netherlands; The National Science Centre, Poland (2015/18/E/ST2/00758); National Authority for Scientific Research (ANCS), Romania; Ministerio de Ciencia, Innovación, Investigación y Universidades (MCIU): Programa Estatal de Generación de Conocimiento (refs. PGC2018-096663-B-C41, -A-C42, -B-C43, -B-C44) (MCIU/FEDER), Severo Ochoa Centre of Excellence and MultiDark Consolider (MCIU), Junta de Andalucía (ref. SOMM17/6104/UGR), Generalitat Valenciana: Grisolia (ref. GRISOLIA/2018/119) and GenT (ref. CIDEAGENT/2018/034) programs, La Caixa Foundation (ref. LCF/BQ/IN17/11620019), EU: MSC program (ref. 713673), Spain.

References

- [1] ANTARES collaboration, *ANTARES: the first undersea neutrino telescope*, *Nucl. Instrum. Meth. A* **656** (2011) 11 [[arXiv:1104.1607](#)] [[INSPIRE](#)].
- [2] KM3NET collaboration, *Letter of intent for KM3NeT 2.0*, *J. Phys. G* **43** (2016) 084001 [[arXiv:1601.07459](#)] [[INSPIRE](#)].
- [3] E. Akhmedov, S. Razzaque and A. Smirnov, *Mass hierarchy, 2-3 mixing and CP-phase with Huge Atmospheric Neutrino Detectors*, *JHEP* **02** (2013) 082 [Erratum *ibid.* **07** (2013) 026] [[arXiv:1205.7071](#)] [[INSPIRE](#)].
- [4] PARTICLE DATA GROUP collaboration, *Review of Particle Physics*, *Phys. Rev. D* **98** (2018) 030001 [[INSPIRE](#)].
- [5] 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:1612.05621](#)] [[INSPIRE](#)].
- [6] L. Breiman, *Random Forests*, *Machine Learning* **45** (2001) 5.
- [7] O. Russakovsky et al., *ImageNet Large Scale Visual Recognition Challenge*, *Int. J. Comput. Vis.* **115** (2015) 211.
- [8] F.-F. Li and D. Jia, *ImageNet: Where have we been? Where are we going?*, webinar, http://image-net.org/challenges/talks_2017/imagenet_ilsvrc2017_v1.0.pdf.
- [9] D. Guest, K. Cranmer and D. Whiteson, *Deep Learning and its Application to LHC Physics*, *Ann. Rev. Nucl. Part. Sci.* **68** (2018) 161 [[arXiv:1806.11484](#)] [[INSPIRE](#)].
- [10] I. Shilon et al., *Application of Deep Learning methods to analysis of Imaging Atmospheric Cherenkov Telescopes data*, *Astropart. Phys.* **105** (2019) 44 [[arXiv:1803.10698](#)] [[INSPIRE](#)].
- [11] M. Erdmann, J. Glombitza and D. Walz, *A Deep Learning-based Reconstruction of Cosmic Ray-induced Air Showers*, *Astropart. Phys.* **97** (2018) 46 [[arXiv:1708.00647](#)] [[INSPIRE](#)].
- [12] ICECUBE collaboration, *Deep Learning in Physics exemplified by the Reconstruction of Muon-Neutrino Events in IceCube*, *PoS ICRC2017* (2018) 1057 [[INSPIRE](#)].

- [13] A. Aurisano et al., *A Convolutional Neural Network Neutrino Event Classifier*, **2016 JINST** **11** P09001 [[arXiv:1604.01444](#)] [[INSPIRE](#)].
- [14] F. Chollet et al., *Keras*, <https://keras.io>.
- [15] M. Moser and S. Reck, *OrcaNet*, <https://github.com/KM3NeT/OrcaNet>.
- [16] KM3NeT collaboration, *Dependence of atmospheric muon flux on seawater depth measured with the first KM3NeT detection units: The KM3NeT Collaboration*, *Eur. Phys. J. C* **80** (2020) 99 [[arXiv:1906.02704](#)] [[INSPIRE](#)].
- [17] M. Honda, M. Sajjad Athar, T. Kajita, K. Kasahara and S. Midorikawa, *Atmospheric neutrino flux calculation using the NRLMSISE-00 atmospheric model*, *Phys. Rev. D* **92** (2015) 023004 [[arXiv:1502.03916](#)] [[INSPIRE](#)].
- [18] G. Carminati, A. Margiotta and M. Spurio, *Atmospheric MUons from PArametric formulas: A Fast GEnerator for neutrino telescopes (MUPAGE)*, *Comput. Phys. Commun.* **179** (2008) 915 [[arXiv:0802.0562](#)] [[INSPIRE](#)].
- [19] Stanford University, *Convolutional Neural Networks for Visual Recognition*, lecture notes (2016), <http://cs231n.github.io/convolutional-networks>.
- [20] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*, MIT Press (2016).
- [21] N. Srivastava et al., *Dropout: A Simple Way to Prevent Neural Networks from Overfitting*, *J. Mach. Learn. Res.* **15** (2014) 1929.
- [22] KM3NeT collaboration, *Characterisation of the Hamamatsu photomultipliers for the KM3NeT Neutrino Telescope*, **2018 JINST** **13** P05035 [[INSPIRE](#)].
- [23] M. Abadi et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, <http://tensorflow.org>.
- [24] S. Ioffe and C. Szegedy, *Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift*, in proceedings of 32nd International Conference on International Conference on Machine Learning (ICML'15) **37** (2015) [[arXiv:1502.03167](#)] [[INSPIRE](#)].
- [25] S. Santurkar et al., *How Does Batch Normalization Help Optimization?*, in proceedings of 32nd Conference on Neural Information Processing Systems (NeurIPS'18) (2018).
- [26] K. He, X. Zhang, S. Ren and J. Sun, *Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification*, in proceedings of 2015 IEEE International Conference on Computer Vision (ICCV'15) (2015) [[arXiv:1502.01852](#)] [[INSPIRE](#)].
- [27] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, in proceedings of International Conference on Learning Representations (ICLR'15) (2015) [[arXiv:1412.6980](#)] [[INSPIRE](#)].
- [28] *Tensorflow API documentation*, https://www.tensorflow.org/versions/r1.15/api_docs/python/tf/train/AdamOptimizer.
- [29] J. Nickolls et al., *Scalable Parallel Programming with CUDA*, *Queue* **6** (2008) 40.
- [30] P. Baldi, J. Bian, L. Hertel and L. Li, *Improved Energy Reconstruction in NOvA with Regression Convolutional Neural Networks*, *Phys. Rev. D* **99** (2019) 012011 [[arXiv:1811.04557](#)] [[INSPIRE](#)].
- [31] R.C. Geary, *The ratio of the mean deviation to the standard deviation as a test of normality*, *Biometrika* **27** (1935) 3.
- [32] J. Hofestädt, *Performance studies for the ORCA neutrino detector*, doctoral thesis, Friedrich-Alexander-Universität Erlangen–Nürnberg (FAU) (2017).