Deep neural networks for energy and position reconstruction in EXO-200

View the article online for updates and enhancements.
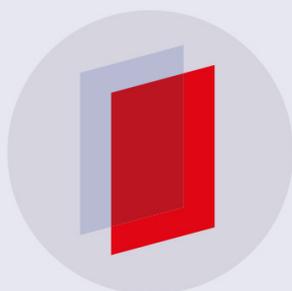
IOP ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

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

Deep neural networks for energy and position reconstruction in EXO-200

S. Delaquis,[a,1] M.J. Jewell,[b] I. Ostrovskiy,[c,2] M. Weber,[b] T. Ziegler,[d] J. Dalmasson,[a,b]
L.J. Kaufman,[a,e] T. Richards,[c] J.B. Albert,[e] G. Anton,[d] I. Badhrees,[f,3] P.S. Barbeau,[g]
R. Bayerlein,[d] D. Beck,[h] V. Belov,[i] M. Breidenbach,[a] T. Brunner,[j,q] G.F. Cao,[k] W.R. Cen,[k]
C. Chambers,[l] B. Cleveland,[m,n] M. Coon,[h] A. Craycraft,[l] W. Cree,[f] T. Daniels,[o] M. Danilov,[i,4]
S.J. Daugherty,[e] J. Daughhetee,[p] J. Davis,[a] A. Der Mesrobian-Kabakian,[m] R. DeVoe,[b]
J. Dilling,[q] A. Dolgolenko,[i] M.J. Dolinski,[r] W. Fairbank Jr.,[l] J. Farine,[m] S. Feyzbakhsh,[s]
P. Fierlinger,[t] D. Fudenberg,[b] R. Gornea,[f,q] G. Gratta,[b] C. Hall,[u] E.V. Hansen,[r] D. Harris,[l]
J. Hoessl,[d] P. Hufschmidt,[d] M. Hughes,[c] A. Iverson,[l] A. Jamil,[v] A. Johnson,[a] A. Karelin,[i]
T. Koffas,[f] S. Kravitz,[b,5] R. Krücken,[q] A. Kuchenkov,[i] K.S. Kumar,[w] Y. Lan,[q] D.S. Leonard,[x]
G.S. Li,[b] S. Li,[h] C. Licciardi,[m] Y.H. Lin,[r] R. MacLellan,[p] T. Michel,[d] B. Mong,[a] D. Moore,[v]
K. Murray,[j] O. Njoya,[w] A. Odian,[a] A. Piepke,[c] A. Pocar,[s] F. Retière,[q] A.L. Robinson,[m]
P.C. Rowson,[a] S. Schmidt,[d] A. Schubert,[b,6] D. Sinclair,[f] A.K. Soma,[c] V. Stekhanov,[i]
M. Tarka,[s] J. Todd,[l] T. Tolba,[k] V. Veeraraghavan,[c] J.-L. Vuilleumier,[y] M. Wagenpfeil,[d]
A. Waite,[a] J. Watkins,[f] L.J. Wen,[k] U. Wichoski,[m] G. Wrede,[d] Q. Xia,[v] L. Yang,[h] Y.-R. Yen[r]
and O.Ya. Zeldovich[i]

[a] SLAC National Accelerator Laboratory, Menlo Park, California 94025, U.S.A.

[b] Physics Department, Stanford University, Stanford, California 94305, U.S.A.

[c] Department of Physics and Astronomy, University of Alabama, Tuscaloosa, Alabama 35487, U.S.A.

*[d] Erlangen Centre for Astroparticle Physics (ECAP), Friedrich-Alexander-University Erlangen-Nürnberg,
Erlangen 91058, Germany*

[e] Physics Department and CEEM, Indiana University, Bloomington, Indiana 47405, U.S.A.

[f] Physics Department, Carleton University, Ottawa, Ontario K1S 5B6, Canada

*[g] Department of Physics, Duke University, and Triangle Universities Nuclear Laboratory (TUNL),
Durham, North Carolina 27708, U.S.A.*

[h] Physics Department, University of Illinois, Urbana-Champaign, Illinois 61801, U.S.A.

[i] Institute for Theoretical and Experimental Physics, Moscow, Russia

[1]Deceased.

[2]Corresponding author.

[3]Permanent position with King Abdulaziz City for Science and Technology, Riyadh, Saudi Arabia.

[4]Now at P.N.Lebedev Physical Institute of the Russian Academy of Sciences, Moscow, Russia.

[5]Now at Lawrence Berkeley National Laboratory, Berkeley, California, U.S.A.

[6]Now at OneBridge Solutions.

^{j}*Physics Department, McGill University, Montréal, Québec H3A 2T8, Canada*

^{k}*Institute of High Energy Physics, Beijing, China*

^{l}*Physics Department, Colorado State University, Fort Collins, Colorado 80523, U.S.A.*

^{m}*Department of Physics, Laurentian University, Sudbury, Ontario P3E 2C6, Canada*

^{n}*SNOLAB, Sudbury, Ontario P3Y 1N2, Canada*

^{o}*Department of Physics and Physical Oceanography, University of North Carolina Wilmington,
Wilmington, North Carolina 28403, U.S.A.*

^{p}*Department of Physics, University of South Dakota, Vermillion, South Dakota 57069, U.S.A.*

^{q}*TRIUMF, Vancouver, British Columbia V6T 2A3, Canada*

^{r}*Department of Physics, Drexel University, Philadelphia, Pennsylvania 19104, U.S.A.*

^{s}*Amherst Center for Fundamental Interactions and Physics Department, University of Massachusetts,
Amherst, MA 01003, U.S.A.*

^{t}*Technische Universität München, Physikdepartment and Excellence Cluster Universe,
Garching 80805, Germany*

^{u}*Physics Department, University of Maryland, College Park, Maryland 20742, U.S.A.*

^{v}*Department of Physics, Yale University, New Haven, Connecticut 06511, U.S.A.*

^{w}*Department of Physics and Astronomy, Stony Brook University, SUNY,
Stony Brook, New York 11794, U.S.A.*

^{x}*IBS Center for Underground Physics, Daejeon 34047, Korea*

^{y}*LHEP, Albert Einstein Center, University of Bern, Bern, Switzerland*

E-mail: iostrovskiy@ua.edu

ABSTRACT: We apply deep neural networks (DNN) to data from the EXO-200 experiment. In the studied cases, the DNN is able to reconstruct the relevant parameters — total energy and position — directly from raw digitized waveforms, with minimal exceptions. For the first time, the developed algorithms are evaluated on real detector calibration data. The accuracy of reconstruction either reaches or exceeds what was achieved by the conventional approaches developed by EXO-200 over the course of the experiment. Most existing DNN approaches to event reconstruction and classification in particle physics are trained on Monte Carlo simulated events. Such algorithms are inherently limited by the accuracy of the simulation. We describe a unique approach that, in an experiment such as EXO-200, allows to successfully perform certain reconstruction and analysis tasks by training the network on waveforms from experimental data, either reducing or eliminating the reliance on the Monte Carlo.

In memory of our dear friend and colleague Dr. Sébastien Delaquis.

Contents

1 Introduction

Deep Neural Networks (DNNs) are a new class of machine learning algorithms that are gaining increasing attention in experimental high energy particle physics. DNNs are made up of multiple hierarchical neural layers, as compared to "shallow" neural networks, which have one hidden layer. A general introduction into DNNs can be found, for example, in ref. [1]. Recent studies demonstrate the ability of DNNs to compete with or outperform more conventional approaches in detecting and localizing events in a complex detector [2], as well as discriminating signal events from background events [2–6]. In a typical application, one uses a Monte Carlo simulation (MC) to generate a large number of events for which the "truth" (i.e. the correct value of classification, or other desired quantity) is known. One then reconstructs the relevant parameters of the simulated events, such as energy and position, and trains the network with these quantities as inputs. The performance is evaluated on an independent set of MC events, not used during training. The performance of the algorithm is therefore contingent on the accuracy of both the MC and the reconstruction routines

used to extract input for the network. The latter requirement already led previous investigations (refs. [2, 4, 7]) to attempt to circumvent as much of the conventional reconstruction as possible, with ref. [2] using almost raw waveforms (with only noise filtering and gain calibration applied) as input to DNNs. It can also be said that an algorithm trained on MC can only exploit those features of the data that were explicitly modeled in the MC.

In this work, we investigate the possibility to avoid certain conventional reconstruction steps and use DNNs to extract the energy and event position directly from the waveforms. We also consider the situation where information about a relevant parameter is accessible from distinct detection channels. This is the case in EXO-200, where both scintillation light and free, drifting electrons, which are produced by particles interacting in the detector, are recorded independently. We show that it is possible to use one of these two signal channels, from which the information can be more easily and accurately extracted, as a source of truth information for reconstruction of the other channel. This avoids reliance on the MC of the other channel, which in the case of EXO-200, can contain systematic errors leading to imperfect modeling of the data. In two out of the three considered cases, actual detector calibration data is available, and we use it to evaluate and validate the performance of the DNNs.

The paper is organized as follows. The second section provides a brief overview of the EXO-200 detector and the existing conventional event reconstruction chain, to the extent necessary to understand the following sections. The third section describes the DNN algorithm developed for energy reconstruction of a point charge deposited on an individual charge collection wire. A special case of induction signals on nearby collection wires is also discussed. The fourth section is dedicated to the DNN reconstruction of arbitrary charge deposits on all charge collection wires. The fifth section deals with data-driven DNN reconstruction of an event position using scintillation light in the detector. The results are summarized and discussed in the last section.

2 Basics of EXO-200: detector, reconstruction, and simulation

2.1 The EXO-200 Detector

The details of the EXO-200 detector have been previously described elsewhere [8]. A brief description of the detector components relevant for these studies is included here for clarity.

The EXO-200 experiment is searching for neutrinoless double beta decay ($0\nu\beta\beta$), a theoretical nuclear decay that, if observed, would indicate that neutrinos are Majorana particles and could help constrain the absolute neutrino mass [9]. EXO-200 started taking data in Sep. 2011 and is still in operation. The current data is divided into two periods, defined as Phase I and Phase II, which are separated by a hiatus due to incidents at the host site. In addition, the start of Phase II was preceded by an electronics upgrade and a change to the electric field in the detector.

The EXO-200 detector consists of a cylindrical radiopure copper vessel filled with Liquid Xenon (LXe), enriched to 80.672(14)% in the isotope ^{136}Xe. The detector is ∼44 cm long and ∼40 cm in diameter and includes two identical back-to-back time projection chambers (TPCs) that share a cathode located at the center of the detector.

Particles interacting with the LXe deposit energy by producing both scintillation light (178 nm) and electron-ion pairs (ionization). Electrons are drifted from their initial location toward the anode

by a uniform bulk electric field. EXO-200 has operated at different bulk fields ranging from $20\,\mathrm{V\,cm^{-1}}$ to $615\,\mathrm{V\,cm^{-1}}$. For simplicity, this study only focuses on data taken during Phase I with a drift field of $380\,\mathrm{V\,cm^{-1}}$ [10].

The charge signal (number of electrons) is detected at the anodes of each TPC by a pair of wire planes crossed at a $60°$ angle, each of which lies in front of an array of large-area avalanche photodiodes (APDs) that detect scintillation light. The frontmost wire plane in each pair (V-Wires) serves both as a shielding grid and to detect induced signals as electrons are drifted from the interaction location to the second wire plane (U-Wires) where charge is collected. Each wire plane includes 38 readout channels with 9 mm pitch resulting in 152 total charge channels. Each APD plane consists of 37 readout channels for a total of 74 APD channels (four APD channels have since been disconnected from the front-end electronics due to large leakage currents).

Calibration of the detector is periodically performed by positioning γ-ray sources at several locations around the detector. For this publication, three sources are used (^{60}Co, ^{228}Th, and ^{226}Ra) to span the energy range of interest for the double beta decay search. These sources can be repeatedly deployed to fixed locations in a copper guide tube that wraps around the outside of the LXe TPC. The positions used for calibration (see figure 1) include two positions located around the edge of the cathode and separated by $90°$ (named S5 and S11) and two positions centered on the drift axis behind either anode plane (named S2 and S8).
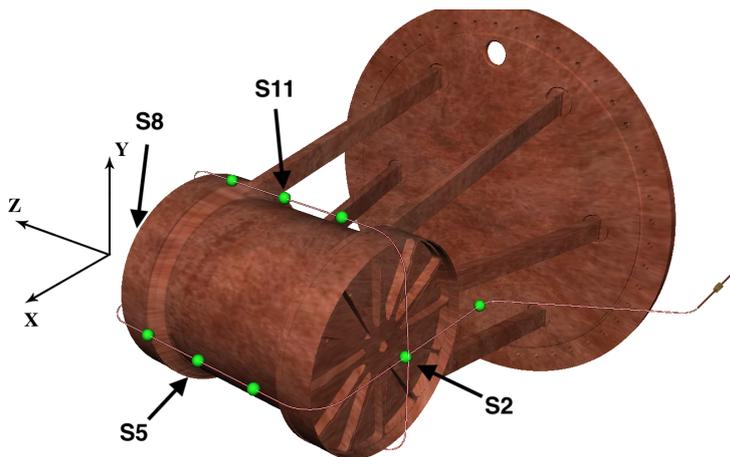
Figure 1. Calibration source positions around the detector vessel. Details in the text.

2.2 Conventional reconstruction in a nutshell

For each trigger in the EXO-200 detector, the digitized raw waveforms from each of the 152 charge channels and 74 light channels are stored. Each waveform consists of 2048 samples taken at 1 MS/s with 1024 samples occurring before the trigger and 1024 samples after the trigger. The standard analysis described in previous studies uses a conventional event reconstruction ("EXO-recon") algorithm to process and analyze the raw waveforms to determine event level information such as energy and position. The conventional reconstruction comprises three stages, which are briefly described below. A more detailed discussion of the conventional reconstruction can be found in [11].

An initial signal finding stage identifies channels containing a signal above a given noise threshold. For this, a matched filter is applied to each channel using a predefined signal template to search for signals above the noise. A second iteration of signal finding is applied to the waveforms after reshaping the signals found by the matched filter. This is used to disentangle waveforms with closely-spaced signals, which the matched filter is ill-suited to find.

Following the signal finding stage is the parameter estimation stage, where each of the identified signals are analyzed to extract parameters relevant for the analysis. This is done by fitting each of the signals to their respective signal models. A χ^2 fit is used to extract both the amplitude and the timing of each signal. Additionally, every charge signal is classified as either "induction" or "collection", which attempts to determine whether the charge deposit terminated its track on the channel. This classification is done using additional information about the waveform, such as the pulse timing, pulse integral, nearest neighbor amplitude, and the χ^2 of the fit of the pulse to templates that describe the shape of induced and collected signals. A set of one dimensional (1D) cuts on these parameters is then applied to classify signals as induction or collection.

Finally, the set of signals and signal parameters for each event is combined to determine event topology and event energy. First, a "bundling" algorithm is applied to separately group U, V, and APD channels based on their timing and amplitude characteristics. A clustering algorithm is then used to group together the bundles of different channel types into the most likely configuration based on a set of probability density functions (PDFs). This results in "clusters," which roughly represent individual charge deposits in the detector. A cluster containing signals of all 3 types has a well defined 3D position and energy and is defined as a "3D cluster". The light and charge energy of an event can be estimated from the APD and U-Wire amplitudes respectively. The Z-position comes from the timing difference between the U-Signals and the APD signals and the U- and V-position from the energy weighted combination of the U- and V-Signal amplitudes.

Following event reconstruction, corrections are applied to the reconstructed parameters to account for known effects in the detector. First, each channel is corrected for its individual gain to normalize the response of each channel to the same energy deposit. These gains are measured using calibration data, and the correction accounts for known time dependencies. Next, a correction is applied to the charge energy to remove the effects of a finite electron lifetime in the detector. This removes the Z-dependent bias introduced from electrons in LXe capturing on electronegative impurities as they drift, which attenuates the detected charge signal exponentially with the drift time. During Phase I the APD channels experienced large time varying correlated noise that degraded energy resolution. To account for this, a denoising algorithm was applied to the light response to optimally estimate the light energy with knowledge of the correlated noise among APD channels and the position dependent light response [12]. Electronic upgrades prior to the start of Phase II significantly reduced this noise component and the need for denoising in Phase II. Finally, in order to improve the energy reconstruction and optimize the energy resolution, a linear combination of the light and charge energy is used to calculate a "rotated" energy of an event. The rotated energy has a better energy resolution than either charge or light alone due to the anti-correlation between the light and charge signals arising from the electron-ion recombination process in LXe [13].

2.3 Monte Carlo simulation

To understand the detector response to energy deposits in the detector volume, a Monte Carlo simulation is employed. A detailed description of the simulation package is found in [11], but a brief description of the details relevant to this paper are included below.

The simulation is divided into two independent components. First, using GEANT4 [14, 15], the EXO-200 detector geometry is modeled and the physical interactions in the detector are implemented. GEANT4 returns a list of charge depositions within the detector volume from the relevant ionizing interactions. These depositions are sampled with cubic voxels with $0.2\,\mu m$ edges.

The set of depositions is then passed into the second component of the simulation, where the electronic response of the system is modeled to produce waveforms similar to those recorded during data taking. For charge waveforms on the U/V planes, the signal generation step tracks each charge deposit through the detector using a 3D finite element simulation of the electric field to realistically model the path of the electrons. Diffusion is also included to correctly reproduce the observed channel multiplicity [16]. The effects of a finite electron lifetime are also optionally incorporated to model the Z-dependent response. At every time step during the drift, the charge induced on each readout channel by the drifting electrons is determined from the Shockley-Ramo theorem [17, 18] which produces a signal in each affected channel. The signals are then shaped using the known transfer functions of the readout electronics for each channel and then digitized. To model the APD response, the current simulation uses a parameterized light response function to give the light yield on each APD plane based on the position of the energy deposit in the detector. The light response is then evenly distributed among all channels on the given plane and assumed to be a step function. The un-shaped signal is then transformed with the known transfer function. This does not include the effect of recombination fluctuations between the light and charge yield in LXe. For this reason, the light simulation is used only to approximately simulate the light reconstruction threshold and is not used to determine the Monte Carlo energy.

Finally, in order to make simulated waveforms more realistic, randomized noise waveforms recorded periodically during data taking are added to each of the simulated waveforms on both the charge and light channels. This creates MC waveforms for each channel with a noise spectrum similar to the true spectrum observed in data. The absolute noise varies for each channel, but on average the broadband noise observed on the U-Wire waveforms is ~800 electrons (which at $380\,V\,cm^{-1}$ corresponds to ~16 keV). Figure 2 shows the average power spectral density of U-Wire noise traces.

3 Energy reconstruction of point-like charge deposition events

The energy resolution of EXO-200 is fundamental for separating potential $0\nu\beta\beta$ decay from background events. Reconstructing the deposited energy from the waveforms of the charge and light channels is limited by an irreducible error due to noise and statistical fluctuations. In the case of the charge channels, the induction signals on U-Wires can also contribute to the reconstruction errors, if not identified and corrected for. The irreducible error sets the limit to the best possible data analysis. The energy reconstruction from a single waveform is directly linked to this error and is thus a good estimator for possible improvements in an existing data analysis. In this study, we
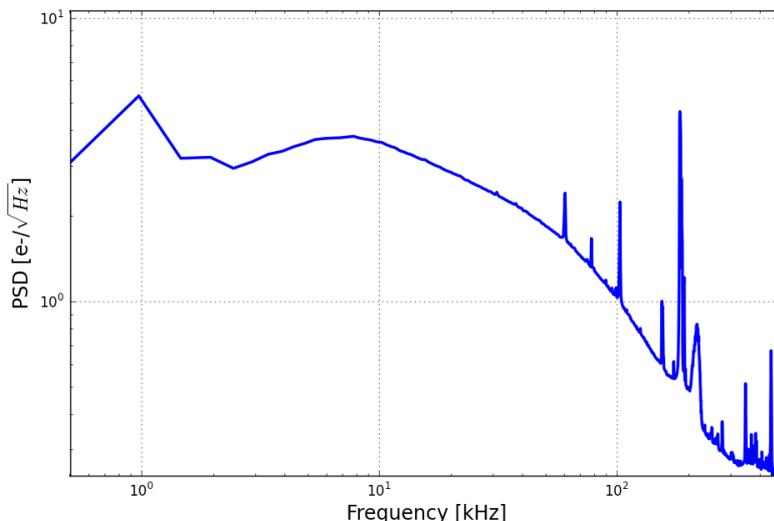
Figure 2. Average power spectral density (PSD) of U-Wire noise traces in the noise library used for adding noise to the waveforms of the MC simulation.

present a method of reconstructing the energy of a point-like charge deposition from the waveform of a single U-Wire using a DNN.

This study is split in three steps. First, a DNN is trained with MC data. Second, the accuracy of charge collection reconstruction is evaluated and the DNN performance is compared with the EXO-200 energy reconstruction algorithm (EXO-recon) on a different set from the same MC distribution. Lastly, the fraction of mis-reconstructed induction events is compared between the DNN and the EXO-recon.

3.1 Network architecture and training on Monte Carlo data

We have built a DNN that takes a waveform as input and predicts the energy of the charge deposition. To train the DNN, we used a supervised learning method. Supervised learning requires feeding not only the input to the DNN, but also comparing the network's output to the true value. However, the true energy of a waveform collected with EXO-200 is unknown *a priori*. Therefore, we used MC, where the true energy is known, to simulate waveforms on a single wire channel. As described in section 2.3, the MC waveforms contain noise samples from actual EXO-200 data. The MC set used in this study contains ~3×10^5 waveforms generated from point like charge deposits artificially generated with the detector MC. One third of the waveforms represent charge collection events on a single U-Wire where the point-like deposit terminates its track. These events have their energy distributed uniformly between 0 MeV and 3 MeV. The remaining waveforms represent induction signals on the two neighboring U-Wires, which do not collect charge but sense an induced signal produced by the drifting charge before it terminates its trajectory. Although some charge is induced on every U-Wire channel on the wire plane, the two neighboring channels are the most likely to have a signal above the noise due to their proximity to the collecting channel. These events are marked as induction signals and have zero true energy associated with them. Examples of charge collection

and induction waveforms are shown in figure 3. While a full EXO-200 waveform contains 2048 time entries, only the central 1000 entries — 500 before and after the trigger — are retained for this study. This reduces the computational load while keeping sufficient data for the baseline and noise estimation.
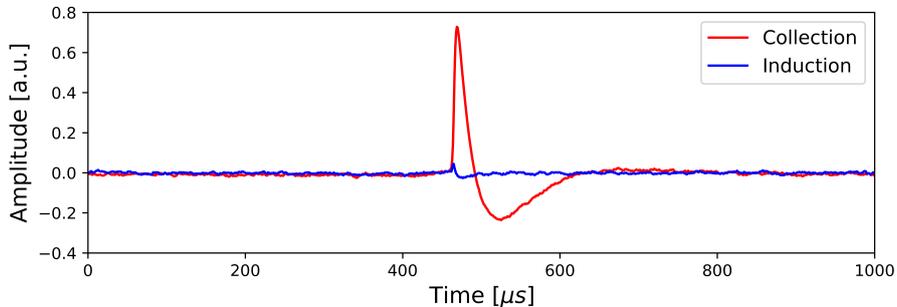
Figure 3. The plot shows two waveforms of the Monte Carlo data; in red a collection waveform and in blue an induction waveform.

When training the DNN, we split the data into three independent sets. First, the training set, containing 80 % of all MC events, is used to train multiple different DNNs. Second, a validation set (10 % of all MC), is used to select the best DNN, based on its performance over this set. Third, a test set (10 % of all MC) is used for evaluating the performance of the selected DNN. This procedure guarantees that the DNN doesn't overfit either the training set or the validation set. Overfitting is defined here as the production of an analysis which corresponds too closely or exactly to a particular set of data, and may therefore fail to fit additional data or predict future observations reliably [19].

A DNN can be parametrized by many parameters such as the number of hidden layers, number of neurons per layer, activation functions, layer type, and regularization techniques. To our knowledge, there are no universal algorithms for choosing an optimal network architecture. Typically, the number of neurons in the hidden layers is chosen to be less than in the input and more than in the output layers. We trained several different network architectures containing only fully connected layers and chose the one that gave the best performance. The architecture of the DNN with the best performance is shown in figure 4. It is built from 5 fully connected layers. A neuron i in a layer l connects with a certain weight vector $w_{[i]}^{[l]}$ to all neurons in the following layer. The number of neurons per layer gradually decreases from 1000 (the number of samples per waveform) in the input layer, to 1 in the output layer, a scalar predicting the energy of the charge deposition. From layer to layer, the number of neurons decreases by a factor of about 4 to 256, then to 64, then to 16, then to 4, and then to 1. Except for the input layer, all layers use the rectifier linear unit (ReLU). This means that an output of a neuron in those layers is given by the following function of its input:

$$ReLU(x) = \max(0, x) \tag{3.1}$$

where x is the weighted sum of the outputs of all neurons from a preceding layer that connect to this neuron. ReLU is currently the most commonly used activation function [1].

To train the network we use mini-batch gradient descent to minimize the loss function L. In this approach, the gradient of the loss is computed using a subset of the training data, called a
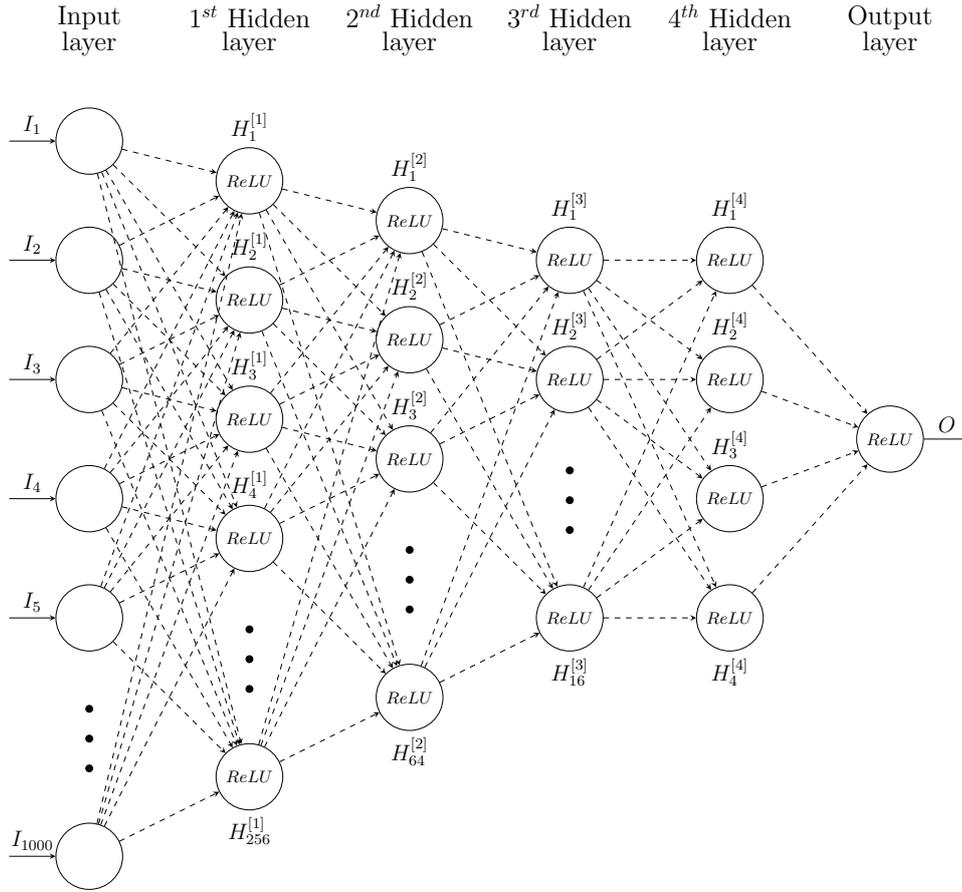
Figure 4. The architecture of the DNN used for the U-Wire study. The input layer represents the input waveform with 1000 data points. It is fed to the first hidden layer, which has 256 neurons with rectified linear unit (ReLU) activation functions. The number of neurons per layer is gradually decreased from 256 to 64, then to 16, then to 4, and then to 1 in the output layer. The DNN is trained to output the true energy in the output layer. The model contains 273 817 trainable parameters.

mini-batch, which is chosen randomly during each minimization step. Therefore, it takes more than one step to go through all the training data; the corresponding number of steps is called an epoch. We use an Adaptive Moment Estimation algorithm, or Adam [20], to compute parameter updates at each step. The loss is defined as $L = C + \lambda \cdot R$, where C is the mean square error, and R is the regularization term. The mean square error is defined as:

$$C = \frac{1}{m} \sum_{j=1}^{m} \left(y_j - \widehat{y_j} \right)^2 \tag{3.2}$$

where m, y_j, and $\widehat{y_j}$ are the size of the mini-batch, the true energy, and the predicted energy of event j, respectively. The regularization parameter $\lambda = 10^{-3}$ defines the contribution of the regularization term (R) to the loss (L). Regularization was implemented using weight decay, also known as $L2$:

$$R = \frac{1}{N} \sum_{l}^{Lr} \sum_{i}^{N_l} \left(w_{[i]}^{[l]} \right)^2 \tag{3.3}$$

where N, Lr, N_l, and $w_{[i]}^{[l]}$ are the number of weights in the entire DNN, the number of layers, the number of weights in the l^{th} layer, and the weight vector i of layer l, respectively. The network is set up and trained using the TensorFlow framework [21]. The framework contains implementation of mini-batch gradient descent, Adam optimizer, and other relevant algorithms.
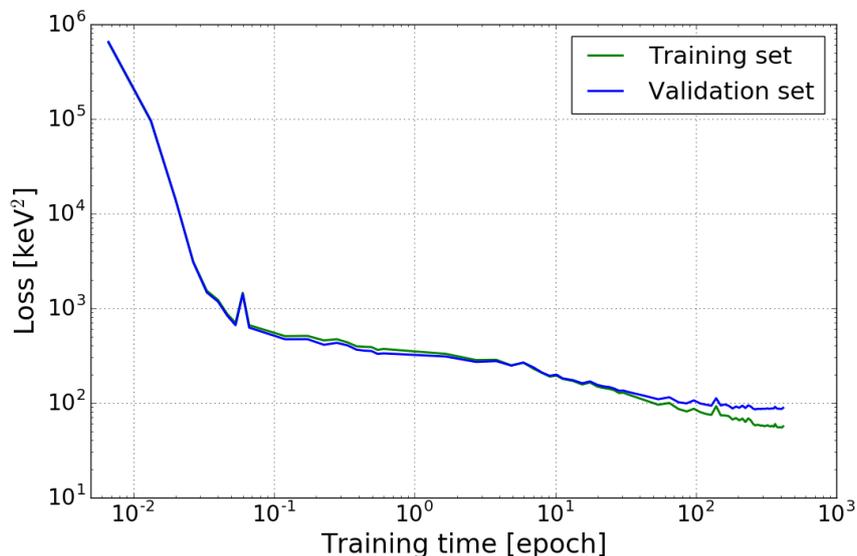

Figure 5. The learning curve shows the decreasing loss during training of the DNN. Thus, the DNN is learning to predict the energy from a waveform.

Training was done in several steps, gradually increasing the mini-batch sizes and the number of training steps per mini-batch size. During the training process, the loss was computed for the training and validation sets (see figure 5). The separation between the performance on the validation set and the training set, occurring at a training time of about 10^2 epochs, is a sign of overfitting. However, the loss on the validation set remains stable after 2×10^2 training epochs. To reduce overfitting, we tried to increase λ. However, with $\lambda > 10^{-3}$, L was dominated by R, and the performance on C was penalized. To reduce the overfitting, substantially increasing the size of the training set could help, which was not done in this study.

3.2   Testing the network on Monte Carlo data

With the trained DNN, the energies of the waveforms in the test set were predicted. Figure 6 shows the true energy versus the predicted energy of both the DNN and EXO-recon. For energies below ~50 keV (see zoom), the DNN can still partially predict the energy while EXO-recon begins to fail. Furthermore, the residuals of the energy predicted by the DNN show fewer outlier points. The distribution of the residuals (see figure 7) is 10 % narrower for the DNN compared to EXO-recon. Hence, we conclude that the DNN performs slightly better on the MC test set than EXO-recon does. The width of the distributions closely corresponds to the average U-Wire noise (see section 2.3), indicating that both algorithms approached a limit of performance.
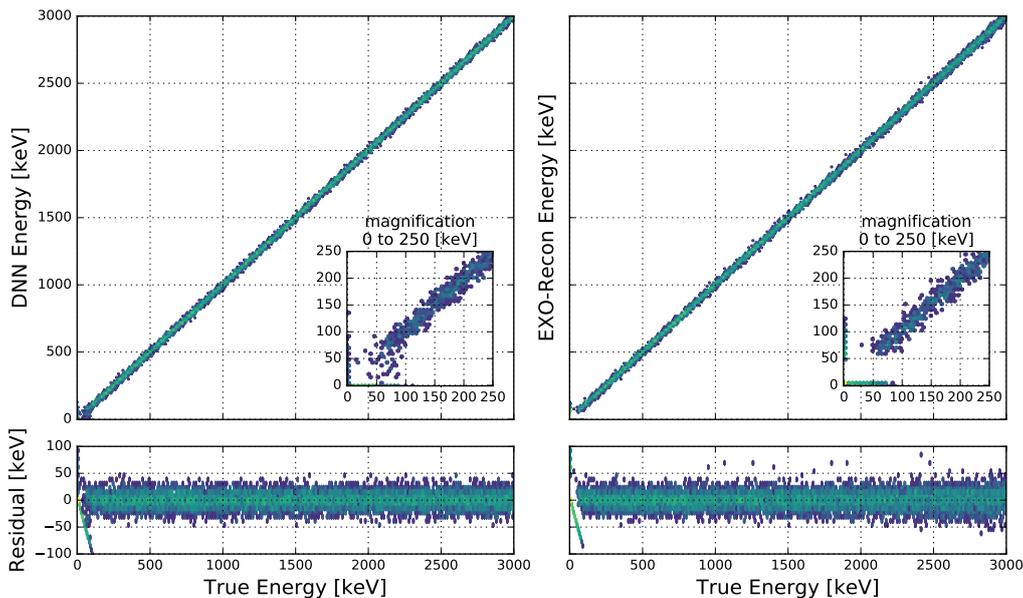
Figure 6. Predicted/reconstructed energy for the test set with both methods, DNN and EXO-recon.
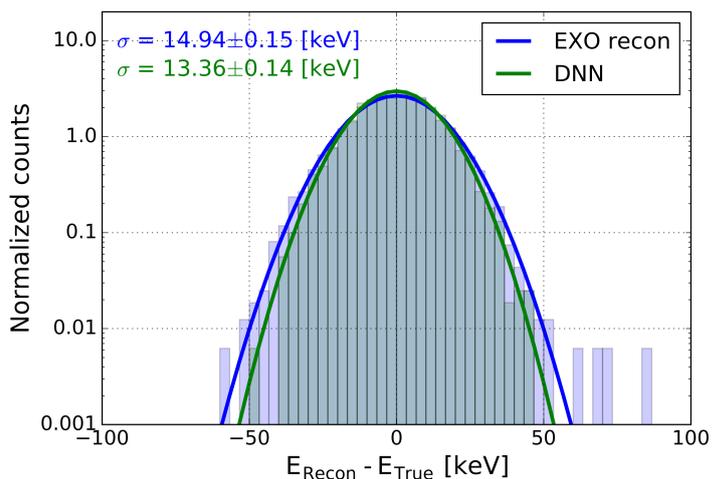
Figure 7. Distribution of the residuals, predicted/reconstructed energy minus true energy, for both methods. Events with zero true energy (the induction events) are not included.

3.3 Induction on collection wires

Upon further inspection of the reconstruction performance, a small fraction of events was found that have a zero true energy but a non-zero reconstructed energy. These are the cases when an induction signal is mistakenly reconstructed as a charge collection. In earlier EXO-200 analyses, mis-reconstructed induction signals led to a noticeable reduction of reconstruction efficiency [22]. The conventional reconstruction algorithm was then adapted to include a dedicated procedure to discriminate induction signals. This is done by applying a set of 1D cuts to parameters extracted from the found U-Wire signals, as mentioned in section 2.2.

The test MC set contains $\sim 3 \times 10^4$ simulated signals, with each collection signal accompanied by two induction signals on the neighboring U-Wires. In this MC set, EXO-recon initially identifies 113 collection signals on the set of waveforms associated with true induction signals. Of these 113 found signals 94 are correctly classified by the induction classification step described above. The remaining 19 signals remain mis-classified as collection signals as they fail at least one of the included cuts. This gives a mis-reconstruction rate of $\sim 17\%$, where the mis-reconstruction rate is defined as the number of true induction signals mis-classified as collection, divided by the total number of signals initially found by EXO-recon on induction channels. The rate determined here agrees within the statistical uncertainty with the previously published value of $\sim 23\%$ [11].

We considered developing a standalone DNN specifically for binary classification of induction and collection signals prior to energy reconstruction. However, the DNN described in this section assigns very low values of reconstructed energy to induction signals, allowing one to discriminate them with an induction signal rejection threshold. The threshold could be set to zero or another value well below the average noise level (section 2.2) and still perform competitively.[1] This implies that the network learned to sufficiently recognize induction signals as a by-product of the main task. To perform a direct comparison with the EXO-recon, we applied the DNN to the 113 (94 + 19) events mentioned above. We find that six of these events are assigned a non-zero reconstructed energy. Figure 8 shows the fraction of mis-reconstructed induction signals as a function of the threshold position for the DNN approach. The EXO-recon fraction is also shown for comparison. EXO-recon does not apply an energy cut to identify induction signals, so the fraction is not a function of threshold. The DNN approach slightly outperforms that of EXO-recon. This check only focused on signals which EXO-recon found and mis-classified initially. In addition there are in total 41 signals which are mis-classified by the DNN but not included in the above subset. Of these signals, 21 were found by EXO-recon during the induction matched filter stage of reconstruction and were correctly classified as induction. The remaining 20 signals are under the EXO-recon threshold and are not identified as either collection or induction. Since the DNN does not distinguish between induction and noise a direct comparison of these signals is difficult. To fully investigate the potential of the DNN to differentiate induction and collection signals one needs to create a neural network specifically for signal classification, which is not done in this study. Additionally, it is important to note that both algorithms use a simulated shape and amplitude distribution of induction signals. The accuracy of these algorithms when applied to experimental data will depend on the precision of the simulation, and any mis-modelings may result in a higher mis-reconstruction rate than seen in MC.

4 Charge-only energy reconstruction

4.1 Network topology and training

In this study, we investigate the energy reconstruction of events in the detector using all available charge collection wires. To do so, we combine the waveforms of all individual U-Wire channels to form an image. The pixel values of the resulting image correspond to the waveforms' amplitudes at a particular time. The study is performed with events regardless of the number of charge deposits

---

[1]This does not mean that a given induction signal was below the noise threshold. Rather, a detectable induction signal was recognized and assigned energy very close to zero (true value).
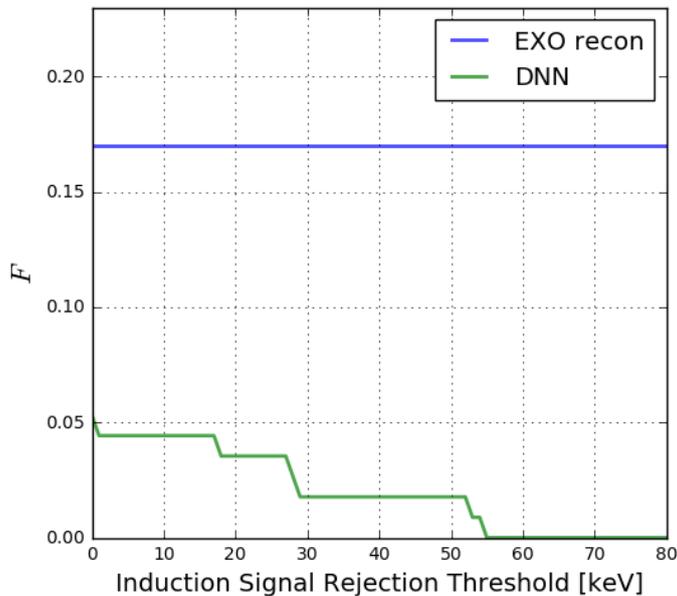
Figure 8. Fraction, F, of induction signals that were assigned a non-zero reconstructed energy by the EXO-recon and by the DNN. For the DNN case, the fraction is shown as a function of the energy cut.

and the multiplicity of the events. With the multiplicity, one can classify the events as single-site (SS) or multi-site (MS) events. Preprocessing consists of subtracting a channel-dependent baseline obtained by averaging the first 800 μs and applying a channel-dependent gain correction. In order to save computational resources, the waveform's first and last 512 entries are cropped. The resulting image of the event consists of 76 channels, each having 1024 time entries, is shown in figure 9.
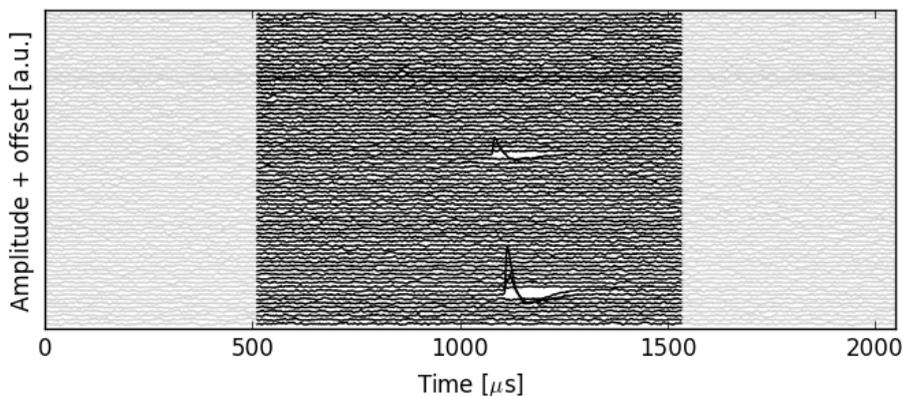
Figure 9. Example waveforms on the collection wires showing multiple deposits in the detector. For clarity, the signals on each wire are offset vertically. Due to computational reasons, the first and last 512 time bins (light gray) are omitted, when being fed to the neural network.

The image is fed to a deep convolutional neural network consisting of six successive convolutional layers. The first convolutional layer consists of 16 kernels of size 5 (time-wise) × 3 (channel-wise) pixels. Kernel corresponds here to a neuron that at any moment is only connected to a part of the

input image, called receptive field. The size of the receptive field is defined by the kernel's size. Each kernel is convolved (cross-correlated) with the input image. The result of the convolution is stored as a two-dimensional array called a feature map. The 16 feature maps from the first convolutional layer are fed as input to the second convolutional layer, which has 32 kernels of the same size as in the first convolutional layer. The following convolutional layers have kernel sizes of 3 (time-wise) × 3 (channel-wise) pixels. The number of kernels gradually increases to 64, 128, and 256 in the last two layers. After the first two convolutional layers, the dimensions of the feature maps are decreased by maximum pooling by a factor of 4 (time-wise) and a factor of 2 (channel-wise). This means that an output of one layer is divided into a set of non-overlapping rectangles, and for each such rectangle only the maximum pixel value is retained to produce a smaller input image for the next layer. Maximum pooling by a factor of 2 in both dimensions is used for the subsequent layers. The network's topology is shown in figure 10. Zero padding, or
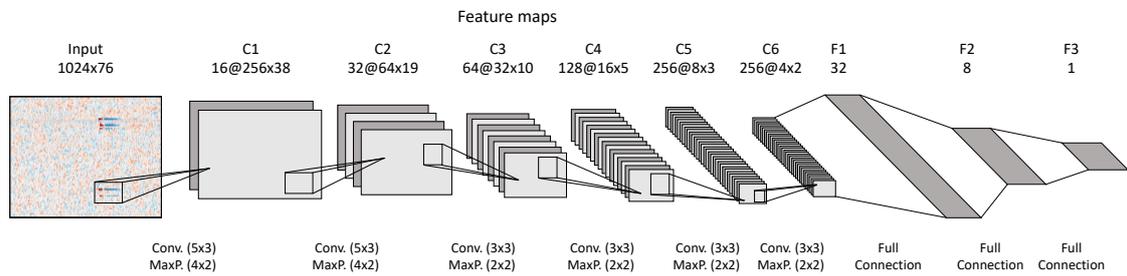

Figure 10. Depiction of the network topology used for the charge-only energy reconstruction. The input image has dimensions of 1024×76 pixels that are fed to the network. The convolutional part consists of alternating layers of convolutions (Conv.) and Maximum pooling (MaxP.) resulting in feature maps with particular dimensions as given in the figure. The dimensions of the feature maps decrease after each convolutional layer while the number of kernels increases. The final feature maps are transformed into a one-dimensional array and are fed to fully connected layers where the last layer has a single unit corresponding to the deposited energy in that image.

surrounding the two-dimensional input image with zeros, is applied to each layer. This allows the output to have the same dimension as the input. Following the convolutional layers are three fully connected layers consisting of 32, 8, and a final unit. The final unit represents the deposited energy of the image corresponding to that event. Throughout the network, the ReLU activation function is used. Initial values of trainable parameters are set using the Glorot algorithm [23]. The network is implemented and trained using the Keras framework [24]. The framework contains implementation of all relevant algorithms.

The neural network is trained with the Adam optimizer by minimizing the mean square error. An additional $L2$ regularization term weighted with 10^{-2} is applied in each layer. MC events are generated using the standard EXO-200 MC framework discussed in section 2.3. For training and validating the neural network, MC events from a gamma ray source positioned in the center of the detector are used. The gamma ray source produces an energy deposition that is uniform in energy. This is motivated by the initial but unsuccessful approach utilizing a ^{228}Th MC source for training, see section 4.2. The noise on the wire channels is sampled from a real noise library (see section 2.3). MC events that are not fully 3D reconstructed within the standard EXO-200 framework are rejected.

Additionally, events that have charge deposits outside the fiducial volume [11] are rejected. The true energy is defined as the sum of all pixelated charge deposits in the MC. The network is trained on events with a total deposited energy larger than 500 keV. For training and validation, about 750 000 MC events from a flat energy distribution in batches of 16 samples are used, with 95 % used for training and the remaining 5 % for validation. Figure 11 shows the learning curve of the network over 100 epochs.
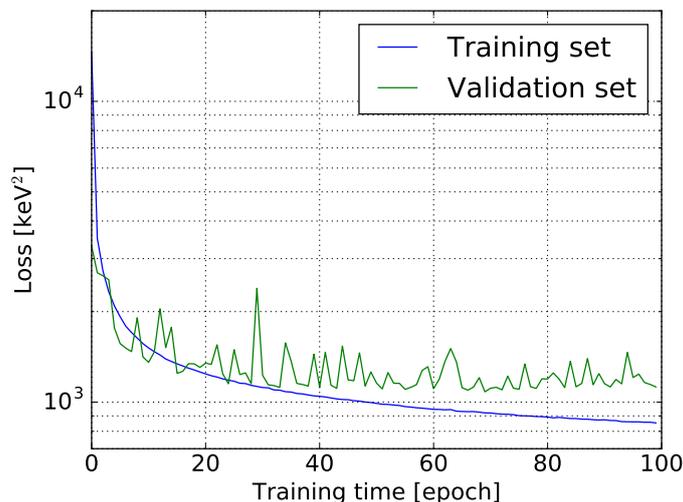

Figure 11. The learning curve shows a decreasing loss during training for both the training set and the validation set. Thus, the convolutional neural network is learning to predict the energy deposited on the charge collection wires.

4.2 Testing on Monte Carlo data

MC-based checks are performed for the exposure of the EXO-200 detector with gamma rays emitted from a radioactive ^{228}Th source at position S5. The correlation between the reconstructed energy from the neural network and the true MC energy is shown in figure 12a. The energy reconstruction by the DNN works as expected. However, there is a small number of outlier events that are also present in the standard reconstruction (cf. figure 12b). The residuals to the true energy of both methods do not show any energy-dependent features other than broadening. The correlation between the residuals of both reconstruction methods with respect to the true MC energy is shown in figure 13. The positive correlation of these residuals indicates an impact due to the noise on the waveforms that affects the energy estimation of both methods in a similar way.

Comparing the reconstructed ^{228}Th energy spectra to the MC input spectrum, the spectral components, e.g. the full absorption peak at 2614.5 keV from the decay of ^{208}Tl and the single and double escape peaks, are reproduced by both reconstruction methods. Both methods yield a similar energy resolution at the full absorption peak at 2.6 MeV (1.29(3) % for EXO-recon and 1.22(2) % for DNN) where a major contribution is electronics noise on the waveforms. The energy resolution of SS-only events at that peak is 1.15(3) % (EXO-recon) and 0.94(1) % (DNN) whereas for MS-only events it is 1.36(4) % (EXO-recon) and 1.38(2) % (DNN) at that peak. The energy axis is calibrated in a simplified way compared to the conventional EXO analysis. Here, we apply a calibration factor

(a) DNN

(b) EXO-recon
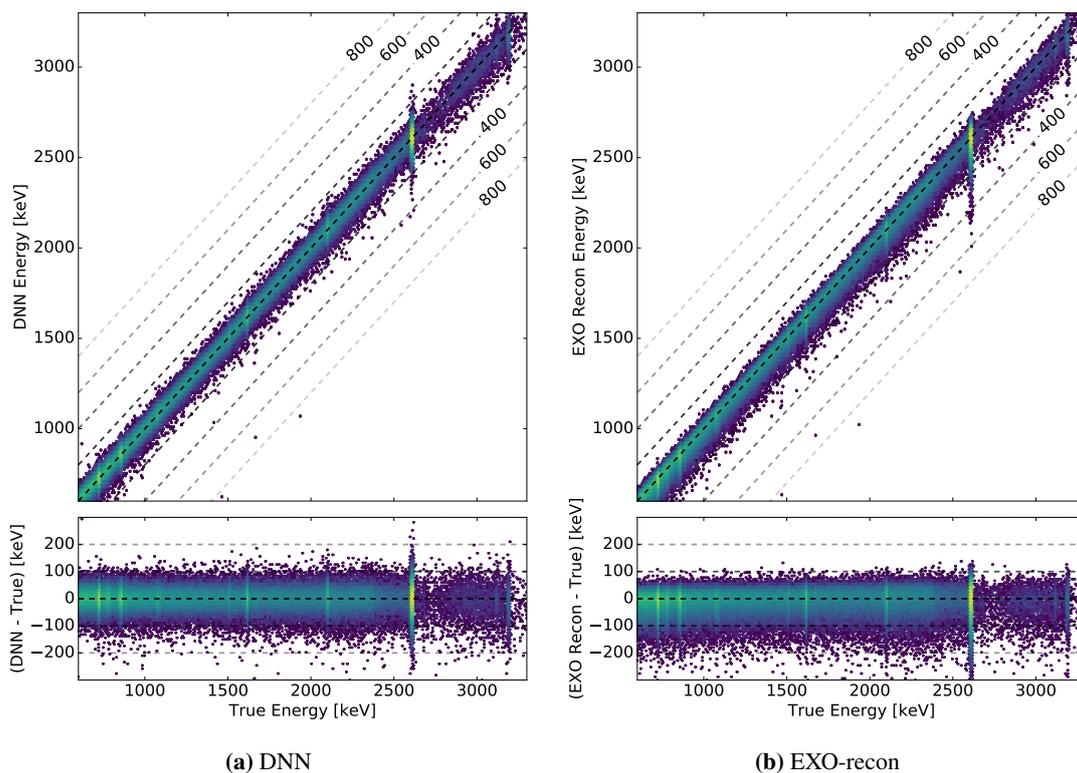
Figure 12. Reconstructed energy of the neural network **(a)** and the conventional analysis **(b)** as function of the true MC energy. The lower panels show the corresponding residuals. The color denotes the intensity on a logarithmic scale. Dashed lines indicate a certain residual value as given in the figure.

for both SS and MS events separately, matching the ^{208}Tl full absorption peak towards the nominal value of 2614.5 keV. The reconstructed charge energy spectra for both the neural network and the conventional reconstruction are shown in figure 14 as well as the spectra for SS-only and MS-only events. The gray shaded spectra show the MC true energies. One noticeable difference between the neural network and the conventional reconstruction is the improved agreement in the region between the Compton shoulder and the full absorption peak. Two main processes contribute to this. First, the neural network is able to reconstruct smaller energy deposits that are below the reconstruction threshold of the conventional method. Second, collection signals superimposed by induction signals are partially flagged as pure induction signal by the conventional reconstruction whereas the neural network outperforms this by an improved capability to disentangle mixed signals.

Initially, in a manner analogous to testing, training was performed on MC events by gamma rays emitted from a radioactive ^{228}Th calibration source. Both training and validation results were promising. Testing the network's performance on MC events from a ^{228}Th calibration source whose energy spectrum was broadened with an arbitrary, artificial resolution of 3.5 % (guided by the charge-only energy resolution from EXO-200) revealed undesired issues. The reconstructed energy of events from the full absorption peak of the ^{208}Tl decay at around 2.6 MeV is shuffled toward the nominal value of the peak, shown in figure 15 (green). Intuitively, this can be explained
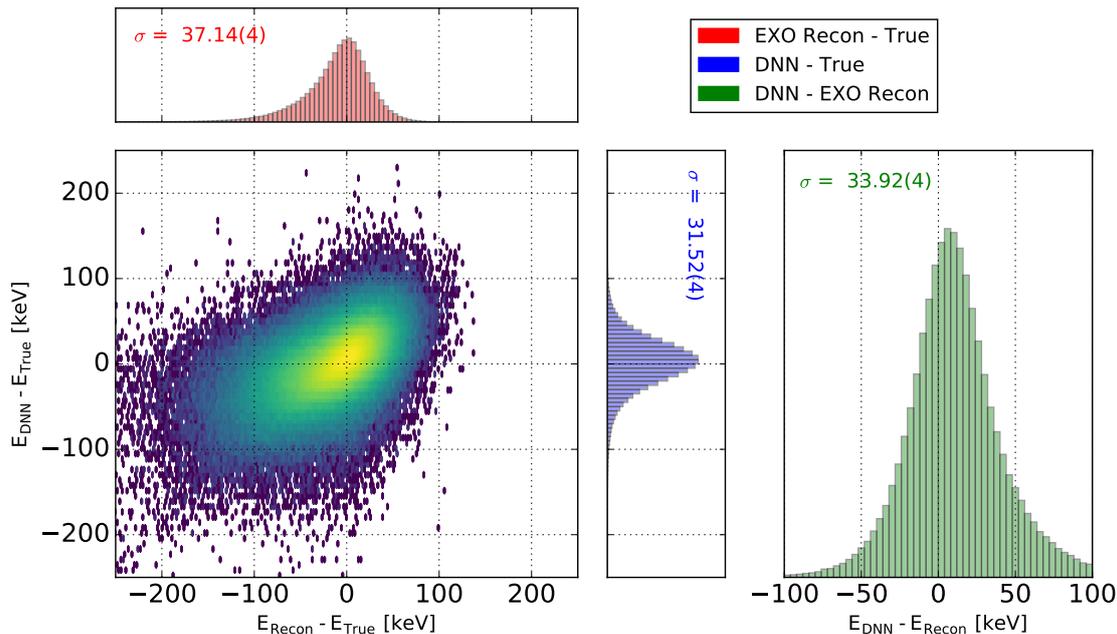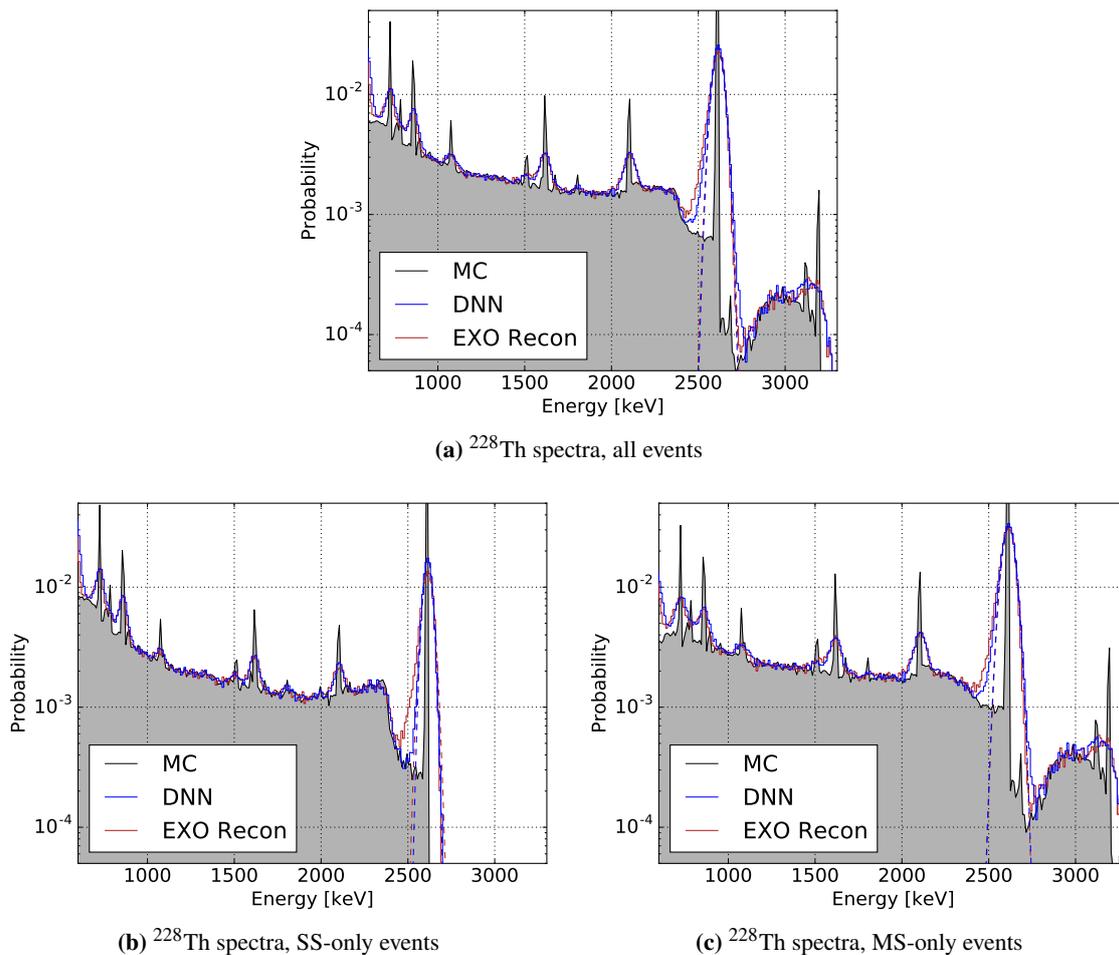
Figure 13. The two dimensional distribution shows the residuals of the energy from the conventional reconstruction, E_{Recon}, and from the neural network, E_{DNN}, with respect to the true MC energy E_{True}. The intensity is denoted by color on a logarithmic scale. The positive correlation indicates an impact due to the noise on the waveforms that affects the energy estimation of both methods in a similar way. The projections are shown for both axes. On the right, the distribution of residuals between both reconstruction methods is shown.

by the network's smaller loss contribution during training, when events from around the peak are reconstructed. Compared to events around this peak, the abundance of ^{208}Tl peak events is larger by orders of magnitude (cf. figure 14). Thus, on average, the contribution is smaller when reconstructing events toward the peak than away from it. This mimics a seemingly improved resolution at the ^{208}Tl peak. Hence, from this point onwards, a training spectrum uniform in energy was utilized to improve the network's capability of reconstructing events from different energy distributions. Using this training approach, the broadened spectrum is reconstructed correctly without introducing systematic effects as shown in figure 15 (blue).

4.3 Testing on calibration data

Performance checks with real data events are done with ^{228}Th events of a calibration source at position S5 from EXO-200 data taken in Phase I. The event selection is analogous to the MC event selection introduced in section 4.1. Additionally, events with multiple scintillation clusters (pile-up of light signals) are discarded to ensure that the charge signal is in the central 1024 µs window that is used in this study. The correlation between the reconstructed energy from the neural network and the conventional reconstruction is shown in figure 16. The difference in reconstructed energy between both methods does not show any energy dependent features other than broadening.

The time variation of the LXe purity affects the electron lifetime and thus the detector's response to a given initial energy deposit. During Phase I of EXO-200, a mean electron lifetime of 4.5 ms

(a) ^{228}Th spectra, all events

(b) ^{228}Th spectra, SS-only events

(c) ^{228}Th spectra, MS-only events

Figure 14. Energy spectra reconstructed from MC events from a ^{228}Th calibration source with both the conventional reconstruction (red) and the neural network (blue) are shown in **(a)**. The MC spectrum is shown in gray. **(b)** shows the spectra for SS-only events and **(c)** shows spectra for MS-only events. A major contribution to the energy resolution is the electronics noise that is added to the MC signal. For SS-only events, the neural network performs noticeably better in the region between the Compton shoulder and the full absorption peak.

was achieved [11]. This value is used for generating MC events. However, the electron lifetime varies between different runs used for testing the network on real data. For consistency, the results discussed here are produced from runs with an electron lifetime of about 4.5 ms. However, we do not see any loss in performance of the network on events with a very different electron lifetime.

The neural network yields competitive results compared to the conventional reconstruction, seen in figure 17. The energy resolution at the 2.6 MeV full absorption peak of the ^{208}Tl decay is similar for both methods (3.93(4) % for EXO-recon and 3.84(3) % for DNN). The energy resolution of SS-only events at that peak is 3.53(7) % (EXO-recon) and 3.44(6) % (DNN) whereas for MS-only events it is 4.08(4) % (EXO-recon) and 4.02(3) % (DNN) at that peak. The energy axis is calibrated following the procedure introduced in section 4.1 for the MC case.
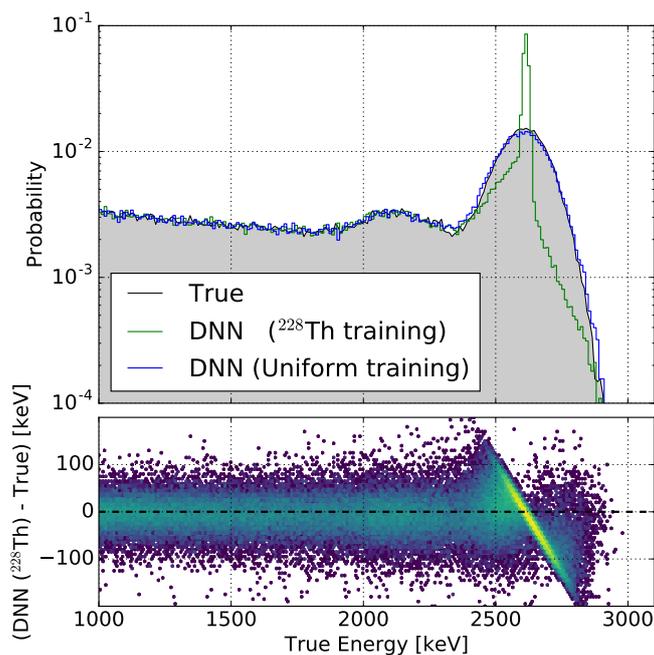
Figure 15. The MC spectrum (gray) reconstructed by two different neural networks (top). The MC events are generated from a ^{228}Th calibration source whose energy spectrum is broadened with an arbitrary, artificial resolution of 3.5 %. One network was trained with a ^{228}Th MC source (green) and the other network (blue) was trained on a uniform energy spectrum. The energy residuals of the network trained on the ^{228}Th MC spectrum (shown in figure 14a in gray) to the true MC energy is shown in the lower plot, where the color denotes the intensity on a logarithmic scale. The residuals indicate an anti-correlation that leads to an artificial improvement in resolution at the peak of the ^{208}Tl decay at 2.6 MeV.

The network's performance on events of the ^{228}Th calibration source when located at other positions is also evaluated. Even though the training is performed on a MC spectrum uniform in energy, the events are not uniformly distributed over the detector volume but are concentrated near the cathode for training and for source position S5. Thus, the position distribution of calibration events from position S5 is similar to the one from the MC gamma ray source. To test the network's performance on highly different event distributions, events from ^{228}Th runs at positions S2 and S8 are evaluated. These positions are centered on the drift axis behind either anode plane. The resulting energy spectra for the conventional reconstruction and the neural network are shown in figure 18. Again, the energy resolution at the 2.6 MeV peak from the ^{208}Tl decay is similar for the two methods. We conclude that the position distribution does not have a significant impact on the network's performance in reconstructing energy.

Due to the anti-correlation of ionization and scintillation in LXe, the energy resolution can be improved by combining both signatures. This is illustrated in figure 19. For the purpose of comparing the network's performance with the conventional EXO reconstruction, we utilize the scintillation signal assessed by the conventional EXO reconstruction for both methods. To improve the energy resolution, denoising is applied to the scintillation channel as introduced in [12]. The optimal energy variable is found by determining the linear combination of charge and light signals
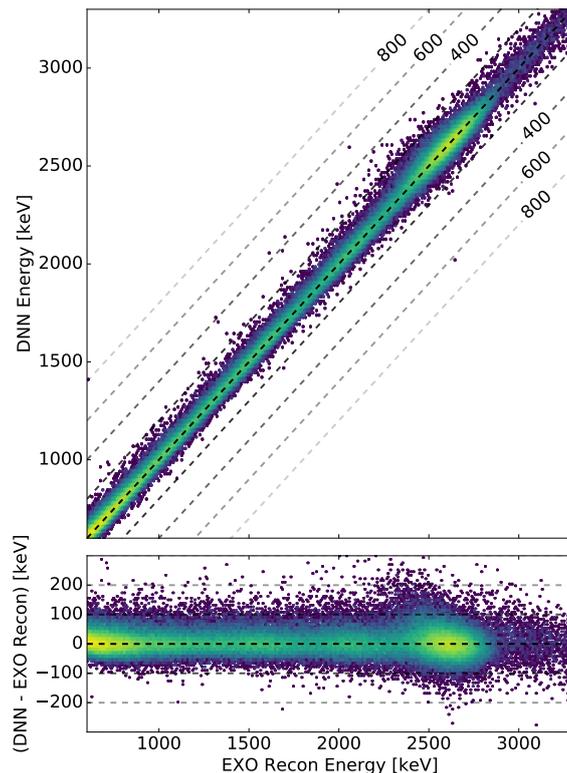
Figure 16. Energies of experimental events from a ^{228}Th calibration source at position S5 reconstructed by the neural network as function of the energy reconstructed by the conventional analysis. The corresponding residuals are shown in the lower panel. The color denotes the intensity on a logarithmic scale. Dashed lines indicate a certain residual as given in the figure. The residual of both methods does not show systematic energy dependent features.
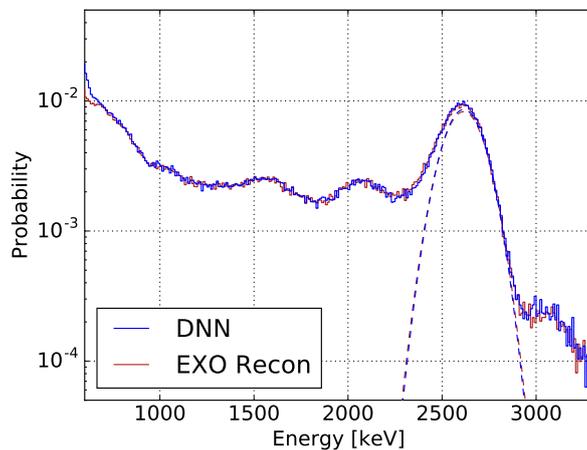
Figure 17. Charge energy spectra reconstructed from data events from a ^{228}Th calibration source at position S5 with both the conventional analysis (red) and the neural network (blue). The energy axis is calibrated as discussed in section 4.1. The resolution is dominated by the recombination fluctuation between the signatures in the light and the charge channel.

(a) ^{228}Th source at S2   **(b)** ^{228}Th source at S8
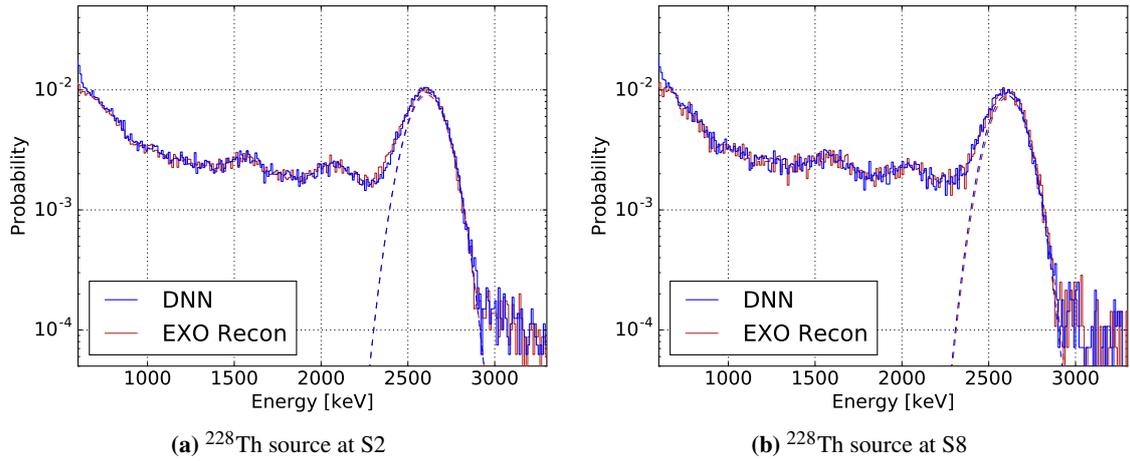
Figure 18. Energy spectra reconstructed from data events from a ^{228}Th calibration source at calibration positions S2 **(a)** and S8 **(b)** with both the conventional analysis (red) and the neural network (blue). The resolution is dominated by recombination fluctuations between the light and the charge channel.


(a)   **(b)**
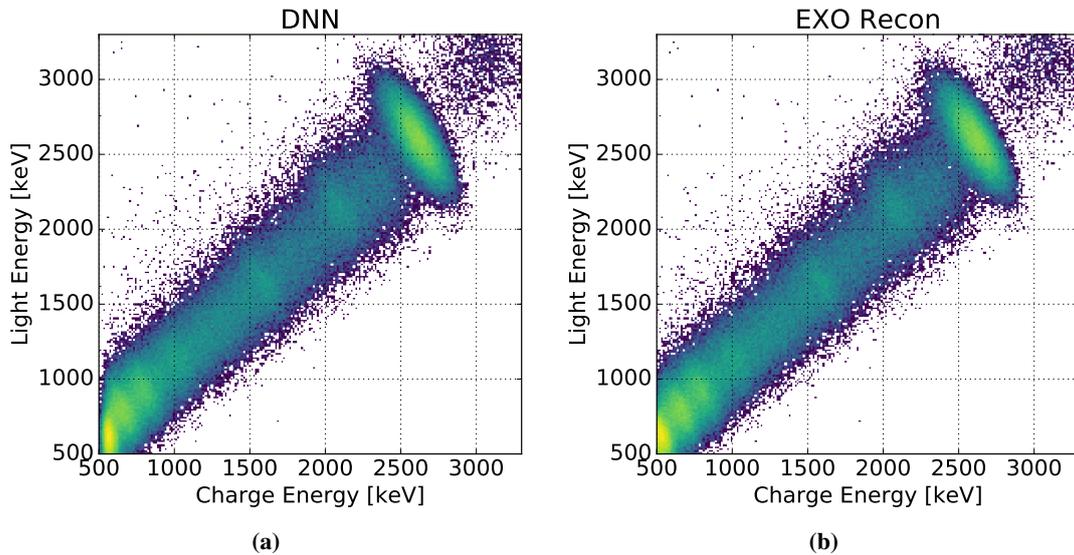
Figure 19. Anti-correlation between ionization and scintillation from a ^{228}Th calibration source. The ionization energy is reconstructed by the neural network (left) and by the conventional analysis (right) whereas the scintillation energy is given by the conventional analysis. The color denotes the intensity on a logarithmic scale. The prominent feature is the full absorption peak of ^{208}Tl.

that minimizes the energy resolution at the ^{208}Tl peak. This linear combination is applied to all events regardless of their energy [11]. This is done for SS and MS events separately. The resulting energy spectra are shown in figure 20 and yield comparable results for both reconstruction methods of the charge energy with an energy resolution at the ^{208}Tl peak of 1.70(2) % (EXO-recon) and 1.65(2) % (DNN). The energy resolution of SS-only events at that peak is 1.61(3) % (EXO-recon) and 1.50(2) % (DNN) whereas for MS-only events it is 1.75(2) % (EXO-recon) and 1.71(2) % (DNN) at that peak. This indicates a physically valid reconstruction by the neural network.

(a) ^{228}Th spectra, all events

(b) ^{228}Th spectra, SS-only events

(c) ^{228}Th spectra, MS-only events
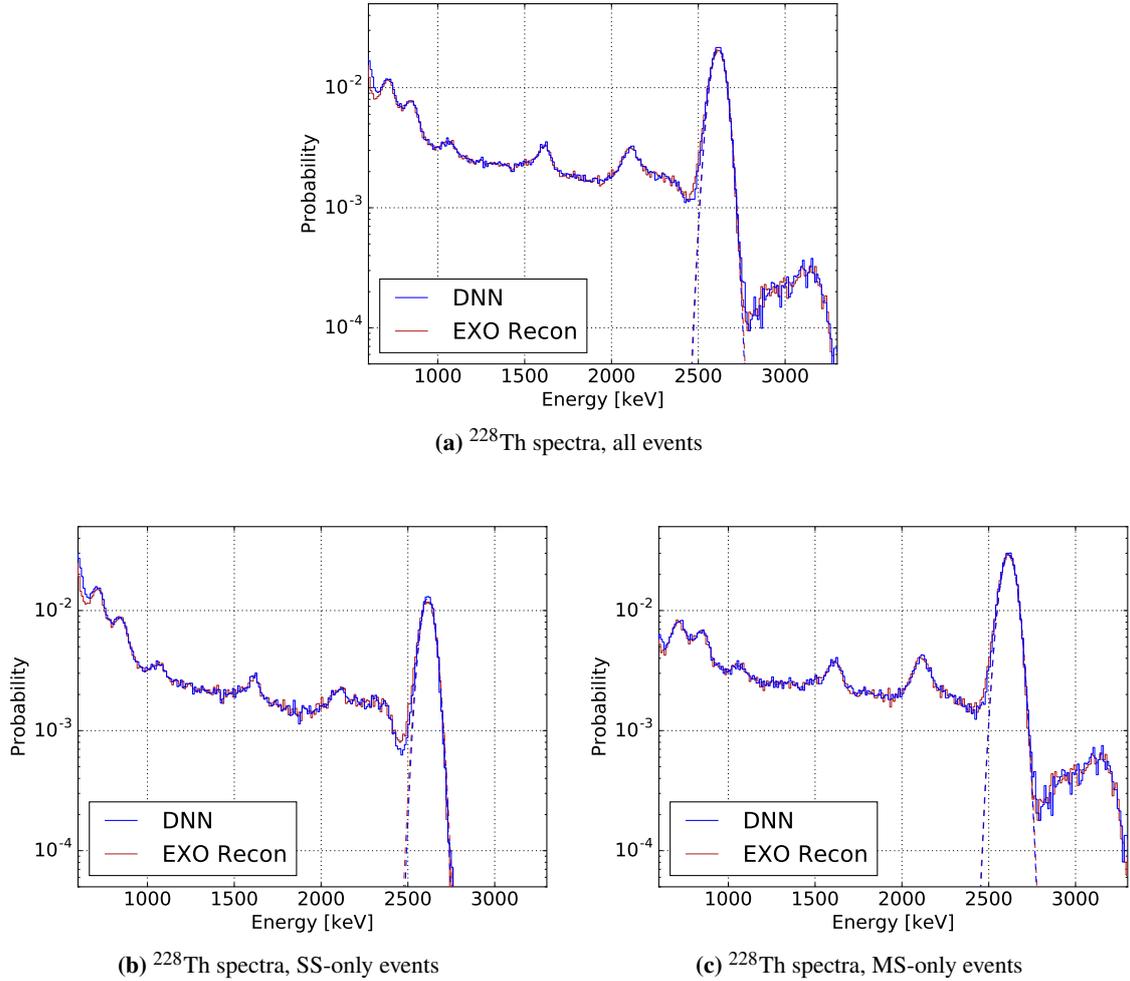
Figure 20. Reconstructed energy spectra after combining the ionization and the scintillation channel with the optimal linear combination from data events of a ^{228}Th calibration source at position S5 with both the conventional analysis (red) and the neural network (blue) are shown in **(a)**. **(b)** shows the spectra for SS-only events and **(c)** shows spectra for MS-only events. The energy axis is calibrated as discussed in section 4.1. The shapes of both spectra of the conventional analysis (red) and the one by the neural network (blue) agree and yield a similar resolution at the full absorption peak of the ^{208}Tl decay at 2.6 MeV. For SS-only events, the neural network performs noticeably better in the region between the Compton shoulder and the full absorption peak.

In order to probe different energies, the same procedure of combining the scintillation and ionization signatures was applied to events from a ^{226}Ra and a ^{60}Co calibration source at position S5. The corresponding spectra are shown in figure 21. The scintillation and ionization channels are combined with the same optimum linear combination as for ^{228}Th. The neural network yields competitive results for all tested calibration sources with the shapes of the spectra being in good agreement with the conventional reconstruction.

(a) ²²⁶Ra spectra   **(b)** ⁶⁰Co spectra
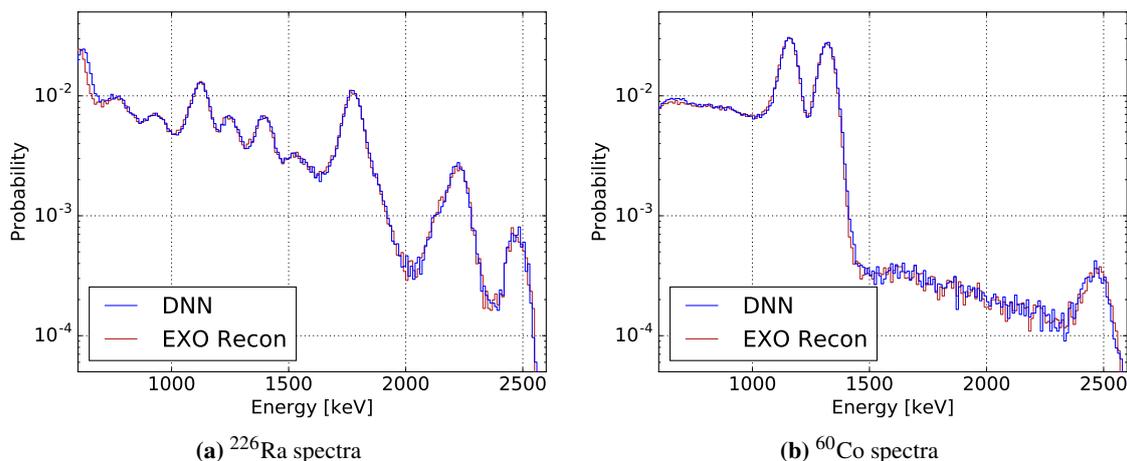
Figure 21. Reconstructed energy spectra from combining the ionization and the scintillation channel from data events from a ²²⁶Ra **(a)** and a ⁶⁰Co **(b)** calibration source at source position S5. The shapes of the spectra of the conventional analysis (red) and the one by the neural network (blue) agree.

## 5 Data driven scintillation light reconstruction

The previously discussed studies attempt to reconstruct the energy deposited by an event occurring in the detector through the amount of free electrons. To accomplish this task, deep neural networks are trained using MC simulations of such events. The accuracy and choice of simulated events will impact the performance of the reconstruction. The fact that EXO-200 registers interactions of ionizing radiation not only through the collection of electrons but also by detecting the scintillation light from re-combinations and excitation, allows for novel approaches of training deep neural networks that do not rely on MC simulations. Here we describe one such approach to reconstruct the position of an event from the distribution of scintillation light signals among the 74 APD channels using a deep neural network, which is trained on the information extracted from the collected ionization. The ability to reconstruct the event position from scintillation light is valuable where insufficient charge collection leads to missing information. This may happen when the position of an event is close to the Teflon reflector, and part or all of the ionization charge is lost and therefore below the detection threshold. This is also often the case for α-decay events which have a large light yield but small charge yield. In the conventional EXO-200 analysis such events are excluded since for them no full 3D position can be reconstructed and they may have the wrong charge/light ratio.

In total, the EXO-200 detector contains 468 APDs, which are read out in groups of 7, resulting in 37 channels in each half of the TPC. They are arranged to cover a maximal area at the end caps of the cylindrical detector and are positioned behind the U- and V-Wires (for detailed description of the detector design see ref. [8] and section 2.1). The conventional EXO-200 analysis extracts the X-Y position information of an event by combining the signal information from charge collection on the U-Wires and induction on the V-Wires. Together with the light signal and the known drift velocity, this enables a three dimensional position reconstruction with an accuracy of approximately $\sigma_{3D} = 3$ mm. The limitation of the position resolution is dominated by the 9 mm wire pitch between

charge readout channels. When the electron cloud is large enough to be collected on more than one neighboring wire, an amplitude weighted position is used. Since we are using this position information with finite resolution to train the network, it also sets the best performance of the model. However, due to larger noise on the APD channels, their coarser segmentation compared to wire channels, and other physical limitations due to light absorption and reflection on different materials, the position resolution extracted from the scintillation light is expected to be worse. This further validates the use of charge information as truth labels and avoids the requirement to simulate light propagation and detection in the detector.

5.1 Training data preparation

The input data for the deep neural network consists of the raw waveforms of all 74 APD channels. The only pre-processing done is the subtraction of the baseline for each waveform determined from the first 680 samples, being the first third of the full waveform, on an event by event basis. Furthermore, we crop the waveforms to 350 samples which guarantees that, independent of whether the event was triggered on U/V-Wires or APDs, the scintillation signal is fully contained in the frame. All waveform signals are stacked vertically in the order of their assigned channel number to form an image of 350×74 pixels. Figure 22 (left panel) shows an example image where the event position is close to the anode plane of the first TPC half defined as positive Z. It illustrates how the position information is encoded in the image: the APDs in proximity to the event location collect the major part of the light and therefore have the largest signal amplitude. APDs far away from the event location have much smaller signal amplitudes.
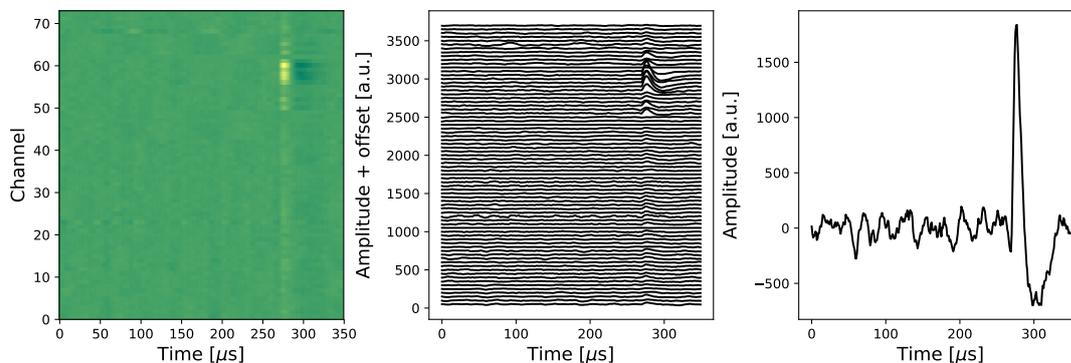

Figure 22. Example of the input image to the convolutional neural network (left panel). The vertical axis represents the APD channel and the horizontal axis represents the time in µs. The middle panel shows the waveforms of each APD channel while the right panel is the sum waveform over all channels.

This image is then fed to a convolutional neural network, which is described in detail below. The dataset used for this study consists of source calibration runs. During a calibration run, one of the four available radioactive sources (^{228}Th, ^{226}Ra, ^{60}Co, ^{137}Cs) is deployed to specific, well known locations outside of the TPC. Due to the relatively short interaction length of gamma rays in LXe, the event distribution is strongly biased towards the location of the source. We combine several runs from different source positions, taken over a period of 6 months, to acquire a sufficiently large training set. We use the event information extracted by the conventional EXO-200 reconstruction to

apply event selection cuts — namely the number of scintillation clusters, the event multiplicity and charge energy — and retrieve the true labels. Note that this information could also be determined by means of DNNs such as the ones described above. However, in this study we do not demonstrate this linkage due to concurrent development of both DNN algorithms. We select events where only a single energy deposit (SS) occurred and only one prompt scintillation signal was found. We apply an energy threshold of 800 keV, as measured by charge, and cut events with high scintillation and low charge yield to exclude the detection of α-particles, which exhibit a signal pattern different from high energy electrons that occur after γ-detections or β-decays in the TPC. The remaining set of events is further reduced in order to create a data set with uniform position and energy distribution. This step turned out to be necessary in order to improve the model's ability to generalize on new, unseen data, in particular for event positions that occur at low frequency in the training data. From the original event distribution with $P = (x, y, z, E)$ we build a 4 dimensional probability density function (PDF) using a Gaussian Kernel Density Estimation (KDE) [25]. We then loop through all events, draw a random number $\rho \in \mathcal{U}(0, 1)$ from a uniform distribution, and keep the event if

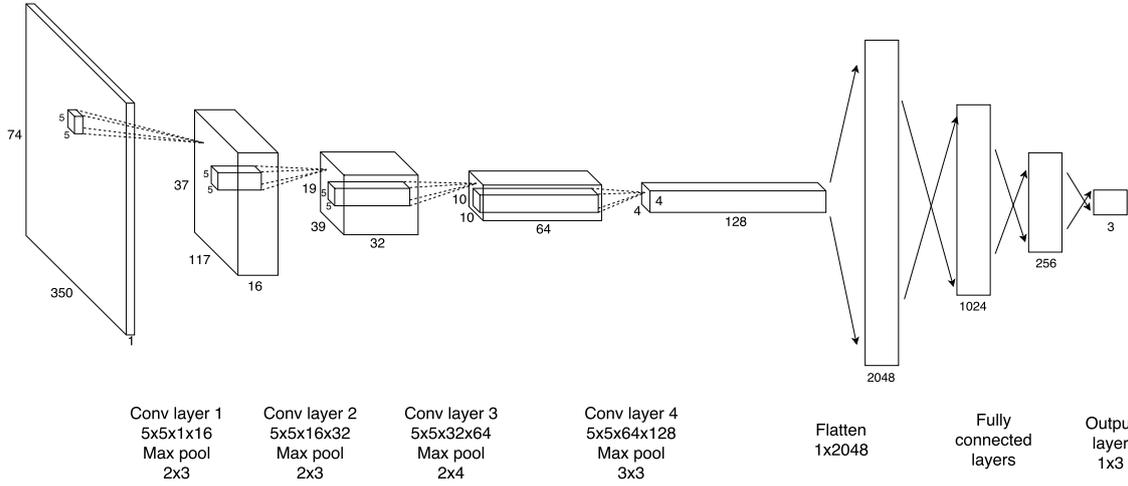$$\rho < \kappa \cdot \frac{PDF(P_{\min})}{PDF(P)} \tag{5.1}$$

where $P_{\min}$ is location of lowest statistics and κ a margin factor allowing to trade off some uniformity for an increased number of events. Figure 23 shows the distribution of events in the training set after applying the described procedure. In total, $\sim$70 000 events are kept, from which 1024 events (1.5 %) are randomly selected and split into a validation set. The events which are removed by the described procedure are stored in separate files and can be used to evaluate the trained model.
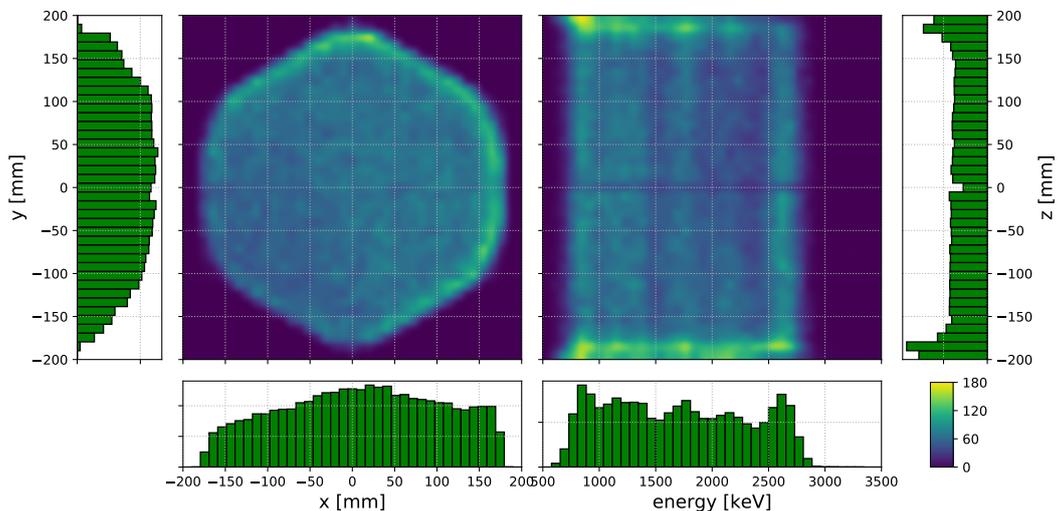
Figure 23. Event distribution of the training set. The distinct locations of the calibration sources outside of the TPC causes the detector volume to be illuminated non-uniformly. The process of discarding some events based on the original event distribution (see text) creates a more uniform distribution in all spatial dimensions and energy.

5.2 Network architecture

For this study we construct a convolutional neural network that takes event images like the one illustrated in figure 22 (left) as input and has three units in the output layer corresponding to the x-,y- and z-position of the event. Given the choice of selecting 350 samples per waveform and 74 waveforms in total containing the unique signature of the event position, a convolutional neural network seems most suitable to handle the 25 900 input features. The structure of the neural network is illustrated in figure 24. The first convolutional layer has 16 kernels of size 5×5 with a stride of 1 followed by a maximum pooling layer of size 2(channel wise)×3(time-wise). This pattern is repeated in layer 2 while increasing the number of kernels to 32. In the following two layers, the number of kernels with the same size is increased to 64 and 128 respectively with max pooling layers of size 2(channel-wise)×4(time-wise) and 3(channel-wise)×3(time-wise) in between. After the last pooling layer, the structure is transformed into a one-dimensional array, yielding 2048 units, which are fed to a densely connected network reducing the size of units per layer to 1024, 256 and finally 3 in the output layer. It consists of $\sim 2.4 \times 10^6$ trainable parameters which are initialized

| Conv layer 1 5x5x1x16 Max pool 2x3 | Conv layer 2 5x5x16x32 Max pool 2x3 | Conv layer 3 5x5x32x64 Max pool 2x4 | Conv layer 4 5x5x64x128 Max pool 3x3 | Flatten 1x2048 | Fully connected layers | Output layer 1x3 |

Figure 24. Architecture of the deep neural network for the scintillation light reconstruction. Four convolutional layers with increasing number of kernels and constant kernel size are followed by three densely connected layers with a final output layer consisting of three units.

using Glorot initializer. We use an Adam optimizer to minimize the loss function $L = C + \lambda \cdot R$, where

$$C = \frac{1}{3m} \sum_{i=1}^{m} \sum_{k=1}^{3} \left(y_i^k - \hat{y}_i^k \right)^2 \tag{5.2}$$

is the mean square deviation of the predicted and true position over all coordinates with m being the size of the mini batch. $\lambda \cdot R$ is the $L2$ regularization term where R is defined in Equation 3.3 adapted to the architecture described here. We have tuned the hyper-parameters such as the size of the neural network, learning rate ($\alpha = 2 \times 10^{-6}$) and regularization ($\lambda = 5 \times 10^{-3}$) for best performance by testing the model during training on the validation dataset. After training the network for 200 epochs, the improvement in the loss is $<0.6\,\mathrm{mm}^2$/epoch while the variance is sufficiently small as shown in figure 25. The remaining bias does not improve with increased size of the network

and can therefore be attributed to the finite position resolution and limited amount of events in the training data set. The loss may be further decreased by adding more data to the training set, but this is not demonstrated in this study. The value of the loss reaches $200\,\mathrm{mm}^2$ after 200 epochs, which translates to an accuracy of $\sigma_{3D} = 24.5\,\mathrm{mm}$. We deem training to be complete at this point and evaluate the accuracy of the model on an independent set of data.
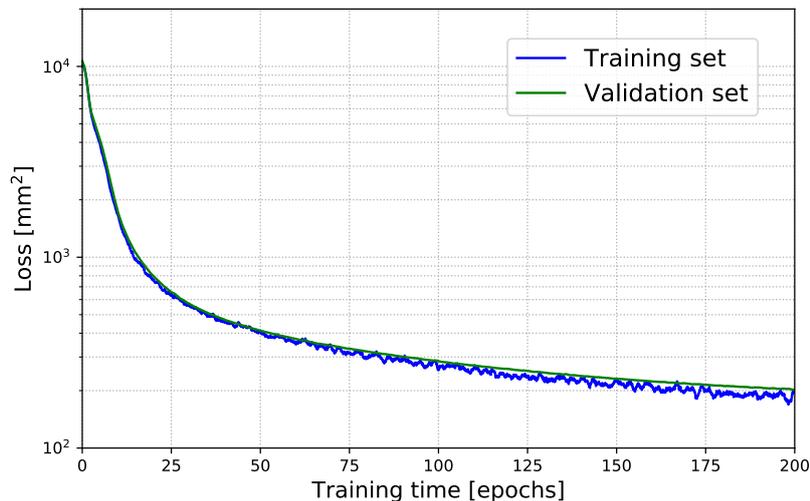
**Figure 25.** Training and validation loss as a function of training epochs.

5.3 Testing of the trained DNN

The accuracy of the trained DNN is tested using various sets of independent data. We utilize the events that were discarded during the process of creating a training set with uniform event distribution. Each test set contains the events of only one source run with specific source type and location. Therefore, the event distribution is biased towards the location of the source and differs from that of the training set. Figure 26 (a) shows the result of evaluating the model on a ^{228}Th source run at position S5 separated into x-, y- and z-coordinates. The green distributions correspond to the true position as determined by charge information whereas the red distributions show the predicted position from the scintillation signal. The bottom panels show the residuals defined as (predicted - true)/true. The total number of tested events is 13211 and the averaged distance between true and predicted event position in 3D space is $\overline{d} = 22.5\,\mathrm{mm}$. The averaged distances in x-, y- and z-coordinates separately are: $\overline{d_x} = 13.6\,\mathrm{mm}$, $\overline{d_y} = 11.3\,\mathrm{mm}$, $\overline{d_z} = 8.1\,\mathrm{mm}$. We have also evaluated the model on different source types, including ^{226}Ra and ^{60}Co, at various source positions and similar accuracies for the predicted event positions were determined. The result for a ^{228}Th run at S11 is shown in figure 26 (b). In this case, the averaged distance between the predicted and true event position is $22.4\,\mathrm{mm}$, while $\overline{d_x} = 11.8\,\mathrm{mm}$, $\overline{d_y} = 12.7\,\mathrm{mm}$ and $\overline{d_z} = 8.6\,\mathrm{mm}$. For both ^{228}Th runs at S5 and S11 the discrepancy between predicted and true position is largest in the coordinate for which the event distribution is close to detector edges and the Teflon reflector (x-coordinate for S5 and y-coordinate for S11). This discrepancy is further illustrated in figure 27 which shows the distribution of (predicted - true) event positions where the mean is non-zero in these cases. However, this shift

(a) ^{228}Th at S5

(b) ^{228}Th at S11
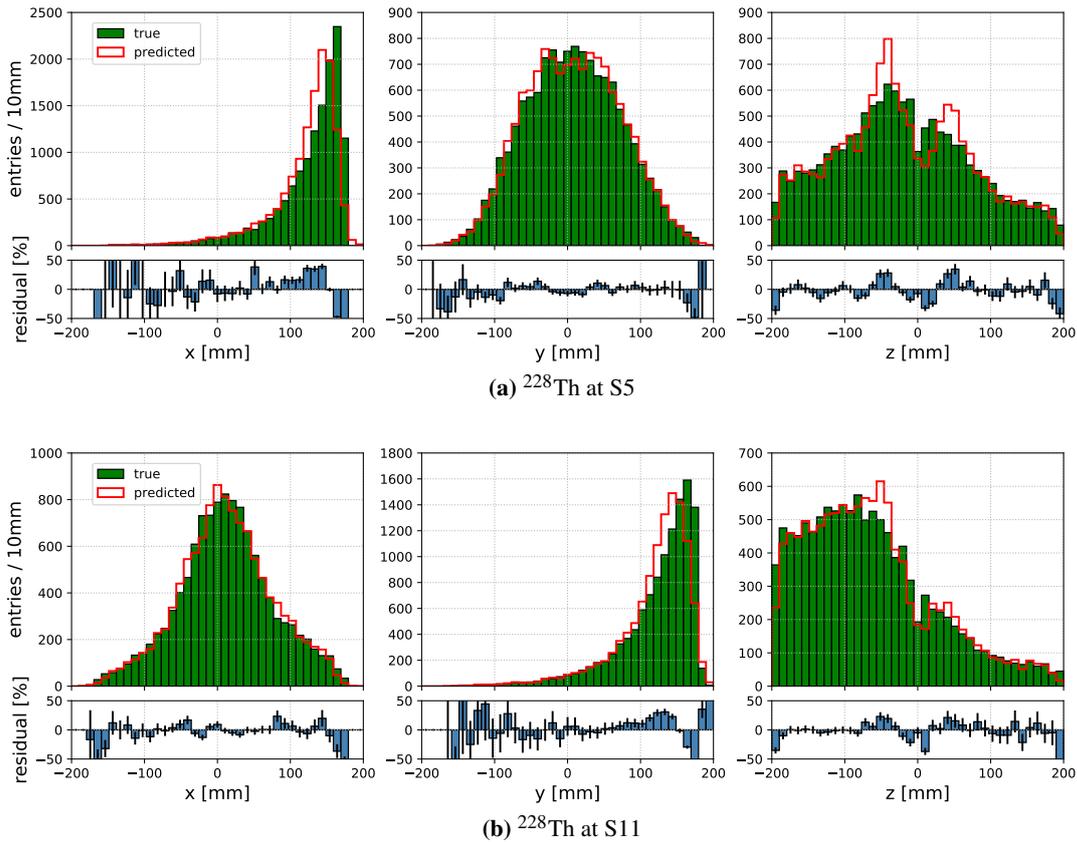
Figure 26. Evaluation of the trained model on an independent set of source data. Events removed from the training set during the process of generating a uniform event distribution (see text) are used to compare the predicted position by the model to the true position from charge information for a ^{228}Th run at S5 (a) and S11 (b).

is several times smaller than the standard deviation of these distributions. Possible explanations are the fact that events occurring near the Teflon reflector might get perturbed due to inhomogeneity in the drift field and charge loss. In such cases the true position deviates from the information retrieved from charge collection. Moreover, the proximity to the Teflon reflector causes more light to be reflected before detection by the APDs which complicates the reconstruction of the light distribution. In addition to the individual source runs, we have generated a test set consisting of events which are contained in a tight fiducial volume of $50\,\mathrm{mm} \times 50\,\mathrm{mm} \times 50\,\mathrm{mm}$ at position $x = 100\,\mathrm{mm}$, $y = 0\,\mathrm{mm}$ and $z = 100\,\mathrm{mm}$ with energies above $2400\,\mathrm{keV}$. The event distribution of this data set is manually constructed and hence is completely independent of the training data. The data set contains 160 events. The true positions of these events are shown as green dots in figure 28 while the predicted positions are represented as red circles. A very good agreement between true and predicted position is visible. The averaged distance is 13.0 mm, while the contributions from each coordinate are: $\overline{d_x} = 8.3\,\mathrm{mm}$, $\overline{d_y} = 4.8\,\mathrm{mm}$, $\overline{d_z} = 6.2\,\mathrm{mm}$. The smaller averaged distances as compared to the previous test sets is mainly attributed to the guaranteed large signal to noise ratio, due to the energy cut, and the distance of the selected events to the detector edges as described previously.
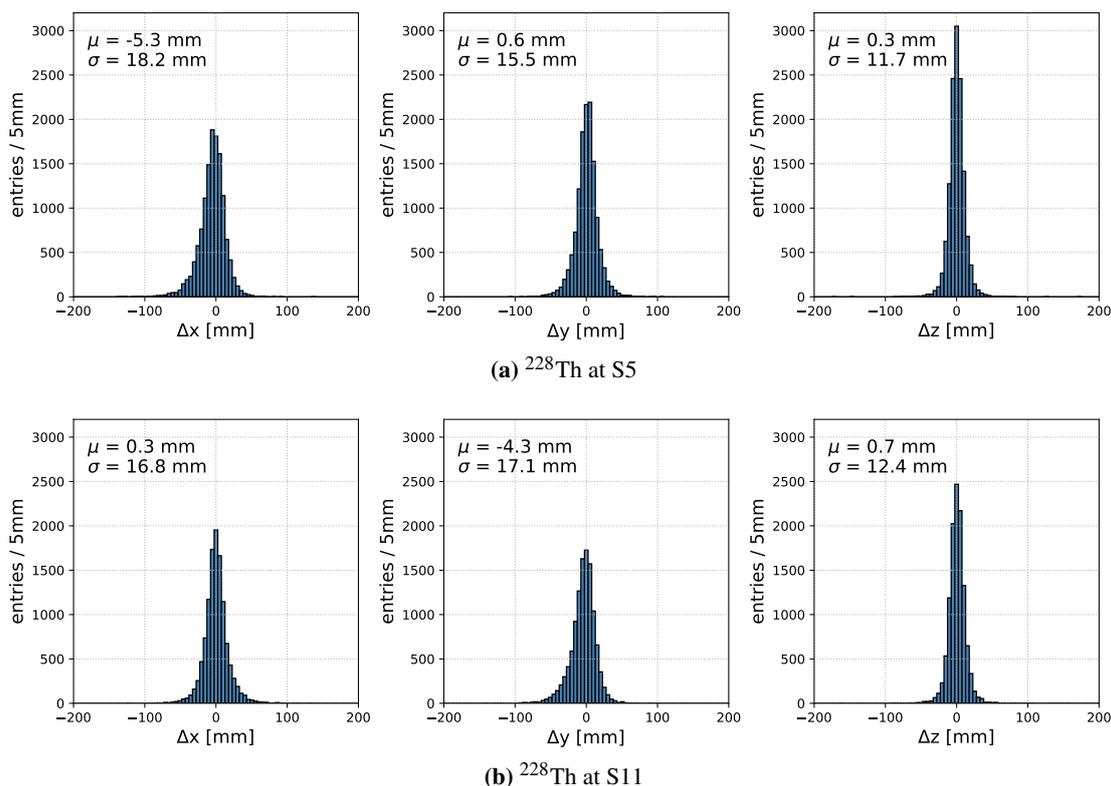
Figure 27. Distribution of difference in predicted and true event position for a ^{228}Th run at S5 (a) and S11 (b). The non-zero mean in X-coordinate for S5 and Y-coordinate for S11 indicate the larger discrepancy when the event distribution is shifted towards the edge of the detector and close to the Teflon reflector.

6 Conclusions

This work presents the first successful reconstruction of energy and position of events in EXO-200 with deep learning. We found that the energy resolution, as measured for a calibration gamma source, is slightly better than that of the conventional reconstruction.

Other works already pursued the reduction of reliance on conventional reconstruction and its possible imperfections, with ref. [2] using only the step of noise filtering before waveforms are used as input for a DNN. We continue this trend and push it a step forward by using, essentially, raw waveforms. Raw waveforms were used in this study to directly extract the energy and position of an event without applying any higher level routines, such as noise filtering.

More importantly, the cited DNN approaches have only been validated with MC. This means that they may not hold their promise when faced with real data. This work has demonstrated that this risk is real. We found a way to mitigate this risk in a particular case of EXO-200 charge energy reconstruction.

We validate the performance of reconstruction of charge energy and position of events with real calibration data. To our knowledge, this is the first time that the real detector data is used in validation of a DNN algorithm in high energy and neutrino physics. We believe that data driven validation is a necessary prerequisite for acceptance of any DNN based reconstruction.
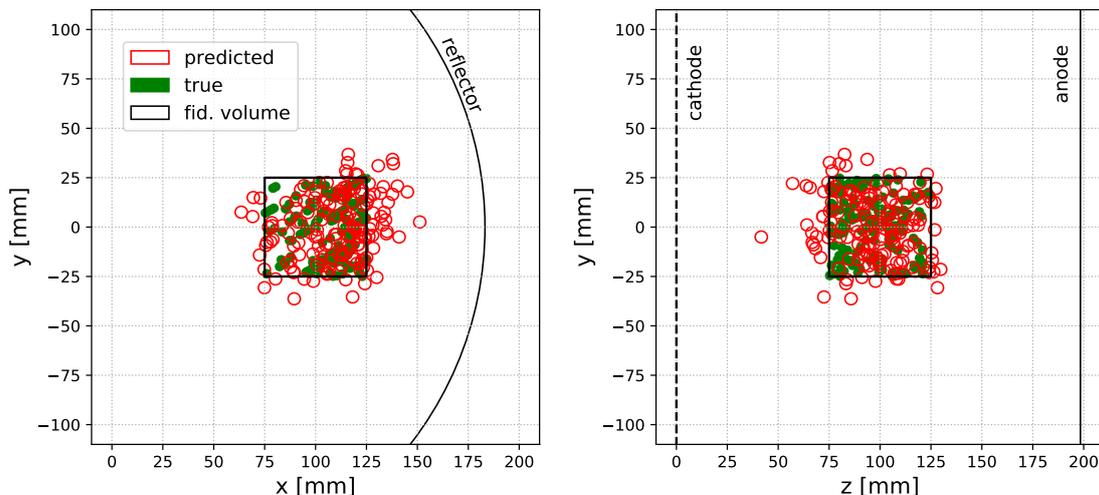
Figure 28. Evaluation of the trained model on an independent set of test data. The events are selected to be in a tight fiducial volume of $50\,\text{mm} \times 50\,\text{mm} \times 50\,\text{mm}$ at position $x = 100\,\text{mm}$, $y = 0\,\text{mm}$ and $z = 100\,\text{mm}$ and have energies above $2400\,\text{keV}$.

The position of events in EXO-200 was also successfully reconstructed using raw APD waveforms that did not undergo the denoising procedure. Due to the absence of a detailed MC simulation of light propagation and collection in EXO-200, we used the event position obtained from U- and V-Wire signals to provide the truth input during the DNN training. The validation of the position reconstruction was also done with the calibration data.

The full EXO-200 dataset spans several years. In order to use the approaches developed here on the full dataset one needs to also apply corrections for time-dependent effects, e.g., for variation of the finite electron lifetime in the detector. The methods to derive such corrections rely on the same conventional reconstruction routines and are not believed to pose a substantially different challenge to the DNN approach. Nevertheless, the effort of determining the time dependent corrections with DNNs has not been attempted. So this work does not demonstrate a successful reconstruction of the full EXO-200 dataset, only of a set of randomly chosen calibration data runs. Reconstructing the full EXO-200 dataset; performing higher level analysis of EXO-200, such as event identification and classification; and further reduction of the reliance of the DNNs on MC during training are the natural avenues for future work.

Lastly, in addition to the performance advantages of the DNN approaches discussed here and in the referenced publications, one can also notice the following. The DNN approaches described here required surprisingly little time to set up and begin to perform competitively. Partly, this is because the DNN approaches take advantage of already existing free but professionally developed and maintained software libraries, like TensorFlow and Keras. With continued scrutiny and advances, it may become possible that the DNN based reconstruction and data analysis will be employed as first choice of the next generation experiments in neutrino, dark matter, and high energy physics.

Acknowledgments

References

[1] Y. LeCun, Y. Bengio, and G. Hinton, *Deep learning*, *Nature* **521** (2015) 436.

[2] R. Acciarri et al., *Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber*, 2017 *JINST* **12** P03011 [`arXiv:1611.05531`].

[3] P. Baldi, P. Sadowski and D. Whiteson, *Searching for Exotic Particles in High-Energy Physics with Deep Learning*, *Nature Commun.* **5** (2014) 4308 [`arXiv:1402.4735`].

[4] A. Aurisano et al., *A Convolutional Neural Network Neutrino Event Classifier*, 2016 *JINST* **11** P09001 [`arXiv:1604.01444`].

[5] J. Renner et al., *Background rejection in NEXT using deep neural networks*, 2017 *JINST* **12** T01004 [`arXiv:1609.06202`].

[6] H. Qiao, C. Lu, X. Chen, K. Han, X. Ji and S. Wang, *Signal-background discrimination with convolutional neural networks in the PandaX-III experiment*, `arXiv:1802.03489`.

[7] D. Guest, J. Collado, P. Baldi, S.-C. Hsu, G. Urban and D. Whiteson, *Jet Flavor Classification in High-Energy Physics with Deep Neural Networks*, *Phys. Rev.* **D 94** (2016) 112002 [`arXiv:1607.08633`].

[8] M. Auger et al., *The EXO-200 detector, part I: Detector design and construction*, 2012 *JINST* **7** P05010 [`arXiv:1202.2192`].

[9] I. Ostrovskiy and K. O'Sullivan, *Search for neutrinoless double beta decay*, *Mod. Phys. Lett.* **A 31** (2016) 1630017 [`arXiv:1605.00631`].

[10] EXO collaboration, J.B. Albert et al., *Search for Neutrinoless Double-Beta Decay with the Upgraded EXO-200 Detector*, *Phys. Rev. Lett.* **120** (2018) 072701 [`arXiv:1707.08707`].

[11] EXO-200 collaboration, J.B. Albert et al., *Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector*, *Phys. Rev.* **C 89** (2014) 015502 [`arXiv:1306.6106`].

[12] EXO-200 collaboration, C.G. Davis et al., *An Optimal Energy Estimator to Reduce Correlated Noise for the EXO-200 Light Readout*, 2016 *JINST* **11** P07015 [`arXiv:1605.06552`].

[13] EXO-200 collaboration, E. Conti et al., *Correlated fluctuations between luminescence and ionization in liquid xenon*, *Phys. Rev.* **B 68** (2003) 054201 [`hep-ex/0303008`].

[14] J. Allison et al., *Geant4 developments and applications*, *IEEE Trans. Nucl. Sci.* **53** (2006) 270.

[15] S. Agostinelli et al., *GEANT4: A Simulation toolkit*, *Nucl. Instrum. Meth.* **A 506** (2003) 250.

[16] EXO-200 collaboration, J.B. Albert et al., *Measurement of the Drift Velocity and Transverse Diffusion of Electrons in Liquid Xenon with the EXO-200 Detector*, *Phys. Rev.* **C 95** (2017) 025502 [`arXiv:1609.04467`].

[17] W. Shockley, *Currents to Conductors Induced by a Moving Point Charge*, *J. Appl. Phys.* **9** (1938) 635.

[18] S. Ramo, *Currents induced by electron motion*, *Proc. Ire.* **27** (1939) 584.

[19] Wikipedia, *Overfitting — Wikipedia, The Free Encyclopedia*, https://en.wikipedia.org/wiki/Overfitting#cite_note-1 (2018).

[20] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, `arXiv:1412.6980`.

[21] M. Abadi et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, Software available from tensorflow.org, https://www.tensorflow.org/ (2015).

[22] EXO-200 collaboration, M. Auger et al., *Search for Neutrinoless Double-Beta Decay in ^{136}Xe with EXO-200*, *Phys. Rev. Lett.* **109** (2012) 032505 [`arXiv:1205.5608`].

[23] X. Glorot and Y. Bengio, *Understanding the difficulty of training deep feedforward neural networks*, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, Chia Laguna Resort Italy (2010), pg. 249, http://proceedings.mlr.press/v9/glorot10a/glorot10a.pdf.

[24] F. Chollet et al., *Keras*, https://github.com/keras-team/keras (2015).

[25] E. Parzen, *On Estimation of a Probability Density Function and Mode*, *Ann. Math. Statist.* **33** (1962) 1065.