

Application of Deep Learning Methods to the Search for Neutrinoless Double Beta Decay with the EXO-200 Experiment

Anwendung von Deep Learning
Methoden zur Suche nach dem
Neutrinolosen Doppel-Betazerfall
mit dem EXO-200 Experiment

der Naturwissenschaftlichen Fakultät
der Friedrich-Alexander-Universität Erlangen-Nürnberg
zur Erlangung des Doktorgrades Dr. rer. nat.

vorgelegt von
Tobias Ziegler
aus Nürnberg

Als Dissertation genehmigt
von der Naturwissenschaftlichen Fakultät
der Friedrich-Alexander-Universität Erlangen-Nürnberg

Tag der mündlichen Prüfung: 02.07.2020

Vorsitzender des Promotionsorgans: Prof. Dr. Georg Kreimer

Gutachter: Prof. Dr. Gisela Anton

Prof. Dr. Christopher van Eldik

Abstract

Proving the Majorana nature of neutrinos would establish physics beyond the Standard Model of particle physics, by demonstrating that neutrinos are their own antiparticles. To date, the best candidate for this proof is the observation of the neutrinoless double beta decay. The EXO-200 experiment searches for the neutrinoless double beta decay in ^{136}Xe with an ultra-low background time projection chamber filled with liquid xenon. The current generation of experiments are sensitive to half-lives of this extremely rare decay of up to $\sim 10^{26}$ yr. The main challenge for any experiment that searches for the neutrinoless double beta decay is therefore to reduce background. Primarily, background reduction is achieved during measurement by evaluating the kinetic energy of the decay products but also by applying particle identification techniques.

In this thesis, deep learning based methods are adapted for data analysis in EXO-200 from approaches used in image recognition. These algorithms are developed in order to improve the sensitivity to the half-life of the neutrinoless double beta decay. A deep neural network is trained to reconstruct the kinetic energy deposited in the detector. In particular, this algorithm outperforms the traditional EXO-200 reconstruction in terms of energy resolution by 10 % (12 %) in Phase-I (Phase-II) of EXO-200 operation at the decay energy of the neutrinoless double beta decay. In an additional study, deep neural networks are developed to discriminate double beta decays from the dominant background interactions. The discrimination power of these algorithms exceeds those of other discriminators which utilize classical machine learning methods. In order to confirm a robust performance, the deep neural networks of both studies are validated on Monte Carlo simulated data and on measured data.

The deep learning based discriminator developed in this thesis contributes significantly to the most recent search for neutrinoless double beta decay of the EXO-200 experiment published in Phys. Rev. Lett. 123 (161802). This analysis outperforms other potential analysis configurations and provides the most stringent median half-life sensitivity of $5.0 \cdot 10^{25}$ yr at the 90 % confidence level. The half-life sensitivity is further increased by utilizing the energy reconstructed by the deep neural network. This represents the best analysis configuration and results in an improvement in sensitivity by ~ 35 % compared to the baseline analysis. These improvements highlight the value of deep learning based methods in complex data analyses for current and future experiments. Additional improvements represent a promising path toward a potential observation of the neutrinoless double beta decay.

Kurzzusammenfassung

Der Nachweis der Majorana-Natur von Neutrinos würde Physik jenseits des Standardmodells der Elementarteilchenphysik etablieren, indem sie zeigt, dass Neutrinos ihre eigenen Antiteilchen sind. Die Beobachtung des neutrinolosen Doppel-Betazerfalls ist bis heute der beste Kandidat für diesen Nachweis. Das EXO-200 Experiment sucht nach dem neutrinolosen Doppel-Betazerfall von ^{136}Xe mit einer hochreinen Zeitprojektionskammer, die mit flüssigem Xenon gefüllt ist. Die aktuelle Generation von Experimenten ist empfindlich auf die Halbwertszeit dieses extrem seltenen Zerfalls von bis zu $\sim 10^{26}$ yr. Die größte Herausforderung für jedes Experiment, das nach dem neutrinolosen Doppel-Betazerfall sucht, besteht daher darin, den Hintergrund zu reduzieren. In erster Linie wird dies während der Messung durch die Bestimmung der kinetischen Energie der Zerfallsprodukte, aber auch durch den Einsatz von Partikelidentifizierungstechniken erreicht.

In dieser Abschlussarbeit werden Deep Learning Methoden aus der Bilderkennung für die Datenanalyse in EXO-200 adaptiert. Diese Algorithmen werden entwickelt, um die Empfindlichkeit gegenüber der Halbwertszeit des neutrinolosen Doppel-Betazerfalls zu verbessern. Ein Deep Neural Network wird trainiert, um die im Detektor deponierte kinetische Energie zu rekonstruieren. Insbesondere übertrifft dieser Algorithmus die traditionelle EXO-200-Rekonstruktion in Bezug auf die Energieauflösung um 10 % (12 %) in Phase-I (Phase-II) von EXO-200 bei der Zerfallsenergie des neutrinolosen Doppel-Betazerfalls. In einer weiteren Studie werden Deep Neural Networks entwickelt, um Doppel-Betazerfälle von den dominanten Hintergrundereignissen zu trennen. Die Unterscheidungskraft dieser Algorithmen übertrifft die anderer Diskriminatoren, die klassische Machine Learning Methoden verwenden. Um eine robuste Funktion zu bestätigen, werden die Deep Neural Networks beider Studien anhand von Monte Carlo Simulationsdaten und Messdaten validiert.

Der in dieser Arbeit entwickelte Deep Learning basierte Diskriminator trägt wesentlich zur jüngsten Suche nach dem neutrinolosen Doppel-Betazerfall des EXO-200-Experiments bei, die in Phys. Rev. Lett. 123 (161802) veröffentlicht ist. Diese Analyse übertrifft andere potenzielle Analysekonfigurationen und bietet die höchste mittlere Sensitivität auf die Halbwertszeit von $5.0 \cdot 10^{25}$ yr bei einem Konfidenzniveau von 90 %. Die Sensitivität auf die Halbwertszeit wird durch die Verwendung der durch das Deep Neural Network rekonstruierten Energie weiter erhöht. Dies stellt die beste Analysekonfiguration dar und führt zu einer Verbesserung der Empfindlichkeit um ~ 35 % im Vergleich zur Basisanalyse. Diese Verbesserungen verdeutlichen den hohen Stellenwert von Deep Learning basierten Methoden in komplexen Datenanalysen für aktuelle und zukünftige Experimente. Weitere Verbesserungen stellen einen vielversprechenden Weg zu einer möglichen Beobachtung des neutrinolosen Doppel-Betazerfalls dar.

Contents

Abstract	v
Kurzzusammenfassung	vii
1 Introduction	1
2 The Neutrinoless Double Beta Decay	3
2.1 Theory of the neutrinoless double beta decay	4
2.1.1 Implications on theoretical physics	7
2.1.2 Experimental approach	10
2.2 The EXO-200 experiment	14
2.2.1 Conceptual design	14
2.2.2 Working principle	17
2.2.3 Overview of the detector Monte Carlo simulations	18
2.2.4 Overview of the event reconstruction and data analysis	19
3 Deep Learning	23
3.1 Machine learning	24
3.2 Deep neural networks	25
3.2.1 Multilayer neural networks	25
3.2.2 Network training	27
3.2.3 Convolutional neural networks	30
4 Energy Reconstruction with Deep Neural Networks	33
4.1 Goal and design of study	34
4.1.1 Motivation and training data preparation	34
4.1.2 Network architecture and training	37
4.2 Performance on Monte Carlo simulated data	39
4.2.1 Bias from a non-uniform training dataset	41
4.2.2 Comparison of DNN and traditional EXO-200 reconstruction	43
4.2.3 Mixed induction and collection signals	45
4.3 Performance on measured data	47
4.3.1 Combination of ionization and scintillation signatures	52
4.3.2 Background reduction in the ROI of $0\nu\beta\beta$	58
4.4 Summary	60

5	Signal-Background Discrimination with Deep Neural Networks	61
5.1	Goal and design of study	62
5.1.1	Motivation and training data preparation	62
5.1.2	Network architecture and training	64
5.2	Performance on Monte Carlo simulated data	67
5.2.1	Comparison of DNN and EXO-200 BDT	68
5.2.2	DNN validation	71
5.3	Performance on measured data	74
5.4	Alternative discriminator concepts	77
5.5	Summary	83
6	Improving the $0\nu\beta\beta$ half-life sensitivity with Deep Learning	85
6.1	Analysis concept	86
6.1.1	Likelihood fit model	86
6.1.2	Analysis configurations	88
6.1.3	Background model	89
6.1.4	Systematic uncertainties	90
6.2	Results	99
6.2.1	Published $0\nu\beta\beta$ half-life sensitivity	99
6.2.2	Published $0\nu\beta\beta$ half-life data limit	101
6.2.3	$0\nu\beta\beta$ half-life sensitivity with the DNN energy	104
6.3	Summary	107
7	Conclusion and Outlook	109
A	List of Figures	113
B	List of Tables	115
C	Bibliography	117
	Acknowledgements/Danksagung	125

Chapter 1

Introduction

Since the discovery of the Higgs boson [1–3], all fundamental particles predicted by the Standard Model of particle physics have been observed. However, there is definitive proof of physics beyond the Standard Model to explain dark matter [4] and neutrino oscillations [5–7]. After several decades of research [8], some fundamental properties of neutrinos remain unknown. The experimental observation of neutrino oscillations established that they carry a non-zero mass. This raises the question of their absolute mass scale. In addition, it remains unclear if neutrinos are Majorana particles [9] which would imply that they are their own antiparticles. This contradicts the current description but would advance our understanding about the asymmetry of matter and antimatter in the Universe [10, 11].

To date, the search for neutrinoless double beta decay [12, 13] is the most promising candidate to probe the Majorana nature of neutrinos. A variety of experiments search for this hypothetical decay in different nuclei. The EXO-200 experiment operated a time projection chamber filled with liquid xenon to search for this decay in the isotope ^{136}Xe . It completed data acquisition in 2018 and is among the most sensitive experiments in this field of research [14]. The current generation of experiments reach sensitivities on the half-life of this decay of up to $\sim 10^{26}$ yr [15]. Because the decay occurs extremely rarely, if it exists at all, the main challenge of experimental approaches is to reduce background in order to avoid false positive observations. In operation, the most important measure to unambiguously identify the neutrinoless double beta decay is the kinetic energy of the decay electrons. For the decay, a characteristic decay energy is expected, whereas for the standard double beta decay [16], a fraction of the energy is transferred to the neutrinos and remains undetected. To reduce background from this decay channel as well as from other interactions in the detector, a precise energy measurement is crucial. Furthermore, background can be suppressed by particle identification techniques that discriminate candidate double beta decays from background interactions.

Machine learning and deep learning techniques in particular demonstrated superhuman performance in several domains over the last years [17, 18]. Especially the availability of increased computing power together with large public datasets have boosted this development. Today, we encounter machine learning systems in our everyday life, for example when using search engines. Also, in particle physics, these techniques revolutionize the approach to data analysis [19]. The advantages of these methods become apparent for information-rich and data-intensive experiments. Since a few years, the number of deep learning applications in particle physics experiments has increased drastically and is expected to become an even more essential part in the toolkit of future experiments.

Within the scope of this thesis, deep learning methods are adapted for data analysis in the EXO-200 experiment. These algorithms are utilized to improve the analysis of key parameters in order to increase the sensitivity to the half-life of neutrinoless double beta decay. The approach is motivated by the recent success of deep learning in the field of computer vision [20]. Because experimental data in EXO-200 is conceptually similar to images, deep neural networks are applicable naturally. The application of deep learning techniques in particle physics experiments in general represents a promising path to increase the discovery potential of present and future experiments.

In this thesis, Chapter 2 presents the theoretical background of neutrinoless double beta decay and its implication for particle physics. In addition, the experimental aspects for a successful search for this decay are presented with a focus on the EXO-200 experiment and its functionality. Chapter 3 introduces the basic concepts of deep neural networks and how these particular techniques are associated to the general field of machine learning algorithms. Chapter 4 presents a novel approach for reconstruction of the ionization charge energy in EXO-200 with deep neural networks. Chapter 5 presents deep neural networks that discriminate double beta decays from the dominant background interactions. In Chapter 6, the detailed analysis concept of searches for neutrinoless double beta decay in EXO-200 is presented. The application of the developed deep neural networks presented in the previous chapters for the search of this decay is presented and is compared to traditional analysis concepts. Chapter 7 provides a summary of the results presented in this thesis and an outlook on potential further improvements by deep learning based analyses in particle physics experiments.

Chapter 2

The Neutrinoless Double Beta Decay

Contents

2.1	Theory of the neutrinoless double beta decay	4
2.1.1	Implications on theoretical physics	7
2.1.2	Experimental approach	10
2.2	The EXO-200 experiment	14
2.2.1	Conceptual design	14
2.2.2	Working principle	17
2.2.3	Overview of the detector Monte Carlo simulations	18
2.2.4	Overview of the event reconstruction and data analysis	19

This chapter presents a brief overview of the physics related to the neutrinoless double beta decay. The theoretical requirements to allow this hypothetical decay are discussed and its implications to particle physics beyond an observation solely. Also, aspects that contribute to a successful experimental approach are presented.

The EXO-200 experiment is currently among the most sensitive experiments that search for the neutrinoless double beta decay. Its conceptual design and the working principle of the detector is discussed. Further, a short overview of the detector Monte Carlo (MC) simulations, the event reconstruction and the higher level data analysis is given.

2.1 Theory of the neutrinoless double beta decay

The double beta decay ($2\nu\beta\beta$) is a second-order decay weak interaction that is allowed in the Standard Model of particle physics (SM). It was first suggested by M. Goeppert-Mayer in 1935 [16]. Its general decay scheme is shown in Equation 2.1 and the corresponding Feynman diagram is shown in Figure 2.2a.

$$\begin{aligned} {}_Z^A X &\rightarrow {}_{Z+2}^A Y + 2e^- + 2\bar{\nu}_e & (2.1) \\ {}_{54}^{136} \text{Xe} &\rightarrow {}_{56}^{136} \text{Ba} + 2e^- + 2\bar{\nu}_e \quad (\text{example}) . \end{aligned}$$

This decay can occur in some unstable nuclei that cannot undergo single beta decay (β) due to energetic constraints given by the semi-empirical mass formula [21] shown in Equation 2.2. For these nuclei, the decay products of β decays are energetically disfavored whereas those of double beta decays ($\beta\beta$) are energetically favored as is shown in Figure 2.1. More details on the formula are given in Ref. [21].

$$E_B = a_V A - a_S A^{2/3} - a_F (N - Z)^2 \cdot A^{-1} - a_C Z^2 \cdot A^{-1/3} + \delta \cdot a_P A^{-1/2} . \quad (2.2)$$

Under certain conditions, nuclei could decay via the the neutrinoless double beta decay ($0\nu\beta\beta$) which is a decay beyond the Standard Model. It was first suggested by W.H. Furry in 1939 [12]. In this hypothetical decay, no neutrinos are emitted as is shown in Equation 2.3 and in the Feynman diagram in Figure 2.2b.

$$\begin{aligned} {}_Z^A X &\rightarrow {}_{Z+2}^A Y + 2e^- & (2.3) \\ {}_{54}^{136} \text{Xe} &\rightarrow {}_{56}^{136} \text{Ba} + 2e^- \quad (\text{example}) . \end{aligned}$$

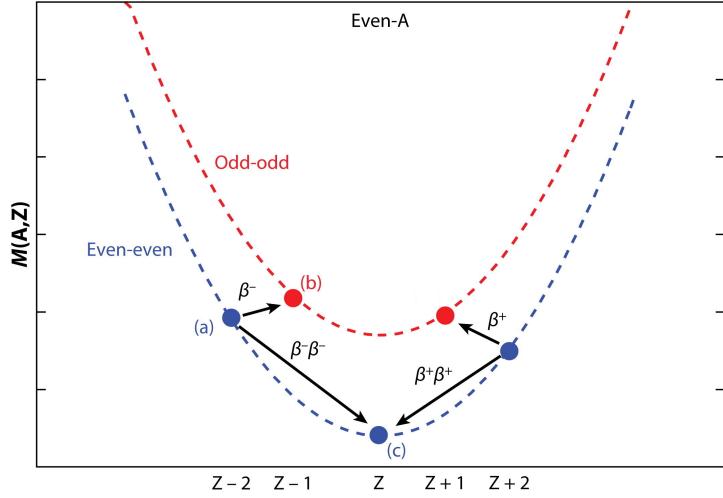


Figure 2.1: Mass parabolas for nuclear isobars with even mass number A. Even-even nuclei (blue) have lower masses compared to odd-odd nuclei (red). The β decay from (a) to (b) is energetically forbidden whereas the $\beta\beta$ decay from (a) to (c) is allowed. Taken from [22].

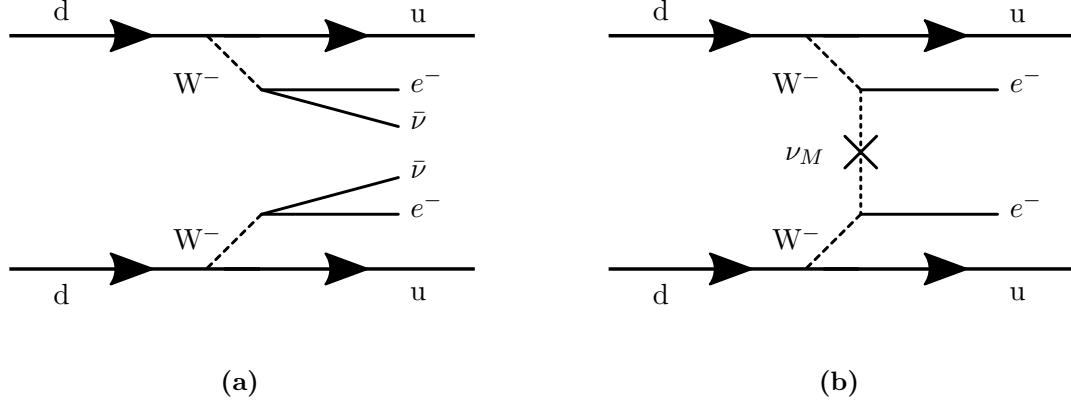


Figure 2.2: Feynman diagrams at tree-level of the $2\nu\beta\beta$ decay (a) and the $0\nu\beta\beta$ decay (b).

To make this decay possible, neutrinos must be Majorana particles [9] which implies that neutrinos are their own antiparticles. However, neutrinos, as well as the other fermions, are described as Dirac particles in the Standard Model. In this description, particles are distinct from their antiparticles. Furthermore, to allow $0\nu\beta\beta$ decay, neutrinos are required to have a non-zero mass which allows them to change their helicity.

Several experiments [5–7] observed neutrino flavor oscillations which demonstrates that at least some light neutrinos are massive particles. Measurements confirm only three light active neutrino flavors with $m_\alpha < m_Z/2$, where m_Z is the mass of the Z boson [23]. However, additional light sterile or heavy neutrinos may exist [4]. The existence of neutrino oscillations implies that the neutrino flavor eigenstates $|\nu_\alpha\rangle$, with $\alpha = e, \mu, \tau$, participate in weak interactions. The flavor eigenstates are different from the propagating mass eigenstates $|\nu_i\rangle$, with $i = 1, 2, 3$. The flavor eigenstates can be expressed as a linear combination of the mass eigenstates:

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle, \quad (2.4)$$

where U is the (3×3) PMNS neutrino mixing matrix [24, 25]:

$$U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \times \text{diag} \left(1, e^{i\frac{\alpha_{21}}{2}}, e^{i\frac{\alpha_{31}}{2}} \right). \quad (2.5)$$

Here, the abbreviations represent $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$ with the angles $(\theta_{12}, \theta_{13}, \theta_{23})$. The matrix further contains the phases $(\delta, \alpha_{21}, \alpha_{31})$. δ is the Dirac-phase that is responsible for potential CP violation in the lepton sector. If neutrinos are Dirac particles, the Majorana phases α_{21} and α_{31} vanished [4, 26].

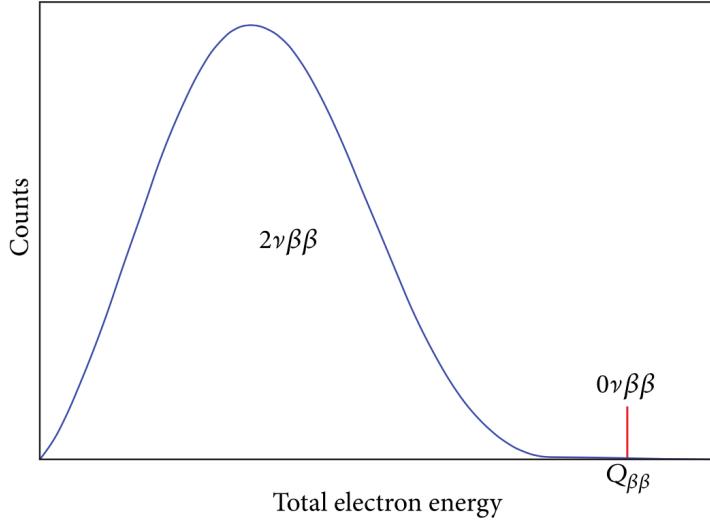


Figure 2.3: Sum kinetic energy spectra of emitted electrons in $2\nu\beta\beta$ and hypothetical $0\nu\beta\beta$ decay. Taken from [28].

The release of energy from a double beta ($\beta\beta$) decay, either $2\nu\beta\beta$ or $0\nu\beta\beta$, is given by the Q value:

$$Q = M_i - M_f - 2m_e , \quad (2.6)$$

where M_i and M_f are the energy levels of the initial and final nuclei. For $2\nu\beta\beta$ decays, the final state is a leptonic four-body state where the sum of the kinetic energy of both emitted electrons forms a continuous spectrum from zero to the Q value. The recoil energy of the final nucleus is negligible. However, for $0\nu\beta\beta$ decays, this spectrum is monoenergetic at the Q value since no neutrinos are emitted as is shown in Figure 2.3 [27].

While the $2\nu\beta\beta$ decay conserves both baryon and lepton numbers separately, the $0\nu\beta\beta$ decay is explicitly a leptogenic process by creating two out-going leptons ($\Delta L = +2$) and thus violates the lepton number conservation. While the standard $0\nu\beta\beta$ decay mechanism is the exchange of a light Majorana neutrino, several other potential mechanisms exist that could mediate this decay [29]. Nevertheless, the “black-box” Schechter-Valle theorem ensures that any observation of $0\nu\beta\beta$ decay also inevitably proves the Majorana nature of neutrinos, regardless of the exact decay mechanism [13]. In the standard mechanism, the parent nucleus emits a pair of virtual W^- bosons which exchange a Majorana neutrino. Conceptually, during the exchange of a light neutrino, an antineutrino is emitted from one W^- boson and is absorbed as a neutrino by the other boson. To be absorbed by the electroweak current with a flipped helicity, neutrinos must be massive Majorana neutrinos. Thus, they have a non-zero negative helicity component on the order of $\mathcal{O}(m/E)$. The amplitude of the $0\nu\beta\beta$ decay is related to the effective Majorana neutrino mass:

$$\begin{aligned} m_{\beta\beta} &= \left| \sum_i m_i U_{ei}^2 \right| = \\ &= \left| m_1 c_{12}^2 c_{13}^2 + m_2 s_{12}^2 c_{13}^2 e^{i\alpha_{21}} + m_3 s_{13}^2 e^{i(\alpha_{31} - 2\delta)} \right| , \end{aligned} \quad (2.7)$$

where m_i are the masses of the three neutrino mass eigenstates, following the syntax introduced in Equation 2.5. The rate $\Gamma^{0\nu}$ for this mechanism which is the inverse of the half-life of the $0\nu\beta\beta$ decay is given by [30]:

$$\Gamma^{0\nu} = \left(T_{1/2}^{0\nu} \right)^{-1} = G_{0\nu}(Q, Z) \cdot |M_{0\nu}|^2 \cdot m_{\beta\beta}^2, \quad (2.8)$$

where $G_{0\nu}(Q, Z)$ is the phase space factor and $M_{0\nu}$ is the nuclear matrix element. From Equation 2.8 it follows that a measurement of the half-life $T_{1/2}^{0\nu}$ for a given nucleus can be translated directly to a measurement of the effective Majorana mass $|m_{\beta\beta}|$. This measure is independent of the nucleus. Accordingly, a lower limit on the half-life translates to an upper limit on the neutrino Majorana mass. These predictions are subject to substantial systematic uncertainties arising mostly from calculating nuclear matrix elements $M_{0\nu}$. Depending on the calculation, the results differ by factors up to ~ 3 from each other [4]. Furthermore, the value of the axial coupling constant g_A affects the nuclear matrix element by $M_{0\nu} \sim g_A^{-2}$. Measurements of $2\nu\beta\beta$ decay rates indicate a quenching of the vacuum value $g_A = -1.2723 \pm 0.0023$ [4]. If this quenching occurs in $0\nu\beta\beta$ decay, which is not known to date [26], the $0\nu\beta\beta$ decay is consequently delayed which results in an increased half-life. Nevertheless, an observation of the $0\nu\beta\beta$ decay would inevitably prove the Majorana nature of neutrinos.

2.1.1 Implications on theoretical physics

The discovery of neutrino oscillations was a first hint for new physics beyond the Standard Model [5–7]. These results confirmed that at least two neutrino mass eigenstates have non-zero mass and also indicate that the description of neutrinos in the Standard Model must be extended. To date, neutrino oscillations are confirmed by measurements of the squared mass differences (Δm_{21}^2 , $|\Delta m_{31}^2|$) and the mixing angles ($\sin^2 \theta_{23}$, $\sin^2 \theta_{12}$, $\sin^2 \theta_{13}$). However, oscillation experiments are neither sensitive to the sign of Δm_{31}^2 nor to the absolute neutrino mass scale. This ambiguity results in two possible hierarchies of the mass eigenstates m_i which are shown in Figure 2.4. Possible realizations in nature are the normal mass hierarchy with $m_1 < m_2 < m_3$ and the inverted hierarchy with $m_3 < m_1 < m_2$. Equation 2.8 directly links a $0\nu\beta\beta$ half-life to the effective Majorana mass $m_{\beta\beta}$ and thus to the neutrino oscillation parameters. This implies that an observation of $0\nu\beta\beta$ decay could constrain or identify which neutrino mass hierarchy is realized in nature. In addition, next-generation experiments might also be able to answer this question [31].

The sum of the three mass eigenstates $m_{\text{tot}} = \sum_i m_i$ can be constrained by cosmological observations. While these constraints are model-dependent, the Planck experiment currently places the most stringent limit on $m_{\text{tot}} < 0.12 \text{ eV}$ at the 95 % confidence level (CL) [33]. In parallel, earth-based experiments study the endpoint spectrum of β decays which provides information on the effective neutrino mass of the electron antineutrino. Since the electron neutrino is a mixture of the three mass eigenstates, its mass corresponds to $m_{\beta}^2 = \sum_i m_i^2 |U_{ei}^2|$. The KATRIN experiment measures the endpoint spectrum of Tritium ${}^3\text{H}$ and aims to improve the best existing limit of 2.05 eV [34, 35] with a sensitivity of $m_{\beta} < 0.2 \text{ eV}$ [36]. Figure 2.5 shows the constraints on the effective neutrino Majorana mass $m_{\beta\beta}$ as a function of the lightest

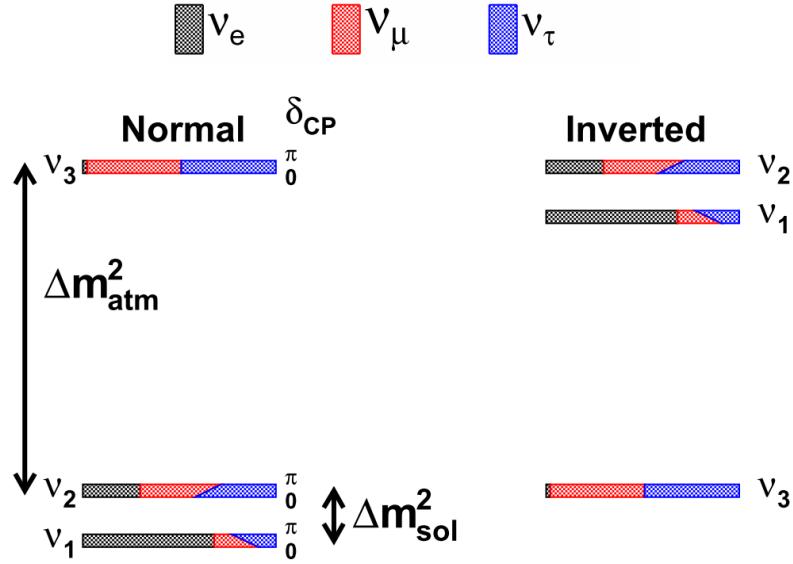


Figure 2.4: Shown are the neutrino masses and their mixing for both possible neutrino hierarchies. The normal hierarchy is shown left and the inverted one right. The flavor content in the mass eigenstates is denoted by color as a function of the CP violating phase δ . Taken from [32].

neutrino eigenstate m_{\min} (left), the sum m_{tot} (center), and the effective electron neutrino mass m_β (right) for both possible mass hierarchies.

From our current understanding, the Big Bang produced equal parts of matter and antimatter. As the Universe expands and cools, matter and antimatter would annihilate into electromagnetic radiation. But this contradicts the matter-dominated nature of the Universe. As an observation of $0\nu\beta\beta$ decay would imply lepton number violation and could give rise to leptogenesis, this could help to understand the asymmetry between matter and antimatter in the Universe.

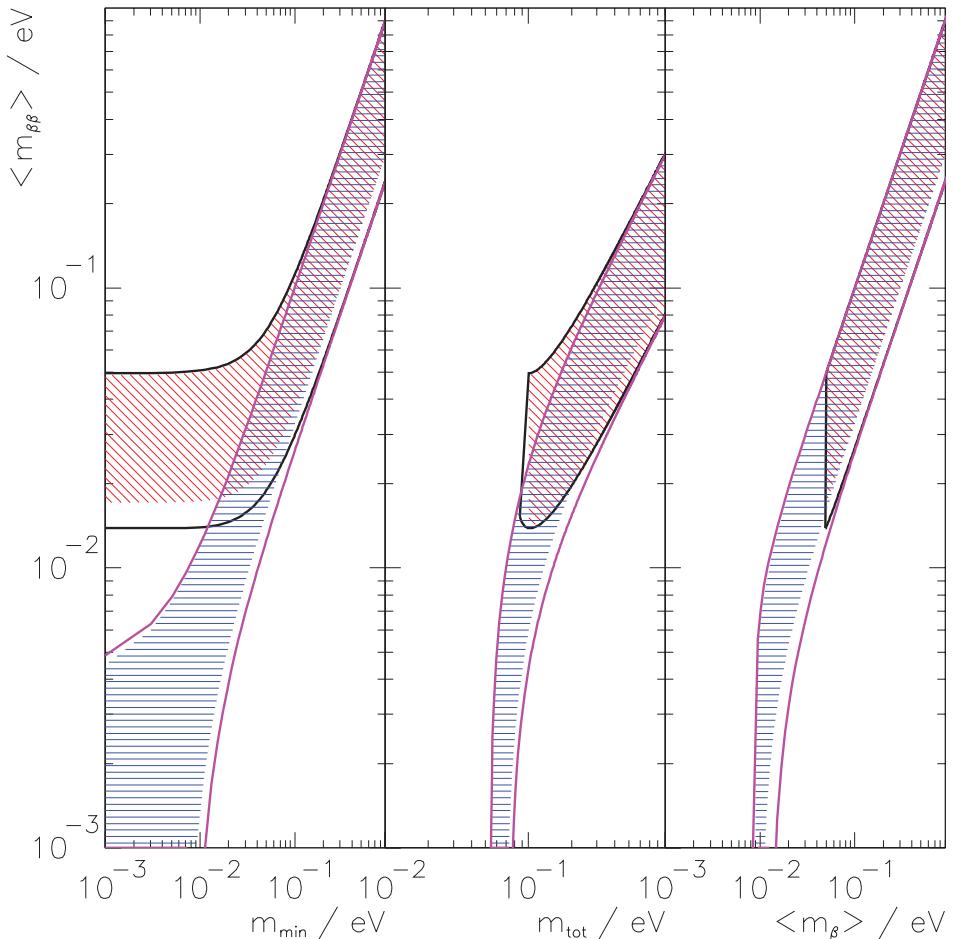


Figure 2.5: Predictions on the effective Majorana mass $m_{\beta\beta}$ as a function of the lightest neutrino mass m_{min} (left), the sum m_{tot} (center), and the effective neutrino mass m_{β} (right) in the case of the normal (blue) and the inverted (red) mass hierarchy. The hatched areas correspond to the allowed regions due to unknown Majorana phases. The solid lines take uncertainties on the oscillation parameters into account. Taken from [4].

2.1.2 Experimental approach

The observation of a single $0\nu\beta\beta$ decay would already prove the Majorana nature of neutrinos. However, assigning a candidate event unambiguously to a $0\nu\beta\beta$ decay as opposed to a $2\nu\beta\beta$ decay or any other background interaction is extremely challenging. In order to probe an exemplary Majorana neutrino mass of $m_{\beta\beta} = 50\text{ meV}$, the corresponding $0\nu\beta\beta$ half-life is about $T_{1/2}^{0\nu} = \sim 10^{26} - 10^{27}\text{ yr}$. This exceeds the age of the Universe of 10^9 yr by many orders of magnitude. Under these assumptions, to observe a single $0\nu\beta\beta$ decay per year, an experiment must operate about 100 kg of a potential $0\nu\beta\beta$ decay isotope. Experimental inefficiencies will further increase the required amount of decay isotope. The figure of merit of different experimental concepts of $0\nu\beta\beta$ searches is the sensitivity to $m_{\beta\beta}$ and thus to the $0\nu\beta\beta$ half-life of the chosen $\beta\beta$ isotope. Sticking to the example of an idealized experiment with no background, the experiment could place a lower limit on the $0\nu\beta\beta$ half-life or an upper limit limit on the Majorana neutrino mass, provided no $0\nu\beta\beta$ decay was observed. This limit depends on several parameters:

$$m_{\beta\beta} \propto \sqrt{\frac{1}{\epsilon}} \cdot \left(\frac{1}{Mt}\right)^{1/2}, \quad (2.9)$$

where Mt is the exposure of the mass M of the $\beta\beta$ isotope over a time t and ϵ is the signal detection efficiency.

Background control and suppression are among the key ingredients for a successful $0\nu\beta\beta$ experiment. The background level b in an experiment is given by:

$$b = c \cdot Mt \cdot \Delta E, \quad (2.10)$$

assuming a proportional accumulation of background with a rate c . The background rate is usually given in units of $\text{kg}^{-1}\text{yr}^{-1}\text{keV}^{-1}$. ΔE denotes the energy window around the Q value that is defined by the energy resolution. As opposed to the background-free example in Equation 2.9, the corresponding experiment with a background level b can place an upper limit on $m_{\beta\beta}$ which depends on:

$$m_{\beta\beta} \propto \sqrt{\frac{1}{\epsilon}} \cdot \left(\frac{c\Delta E}{Mt}\right)^{1/4}. \quad (2.11)$$

If the background near the Q value is not uniform in energy, the background rate c will depend on the energy window ΔE . In addition, it can depend significantly on the mass. This is the case for external background that is shielded by outer parts of the decay mass. Even with these simplified examples, the drastic impact of background on the success of $0\nu\beta\beta$ experiments becomes clear and motivates any efforts to reduce backgrounds.

For the search for $0\nu\beta\beta$ decay, $2\nu\beta\beta$ decays reflect an intrinsic background that cannot be suppressed systematically. The only difference in the signatures of both decays is the energy deposited in the detector, because the emitted neutrinos from a $2\nu\beta\beta$ decay leave the detector without interaction. The energy distributions of both decays are shown in Figure 2.3. This implies that the energy measurement must be sufficiently precise to be able to distinguish $2\nu\beta\beta$ from candidate $0\nu\beta\beta$ decays. For current experiments that operate time projection chambers or

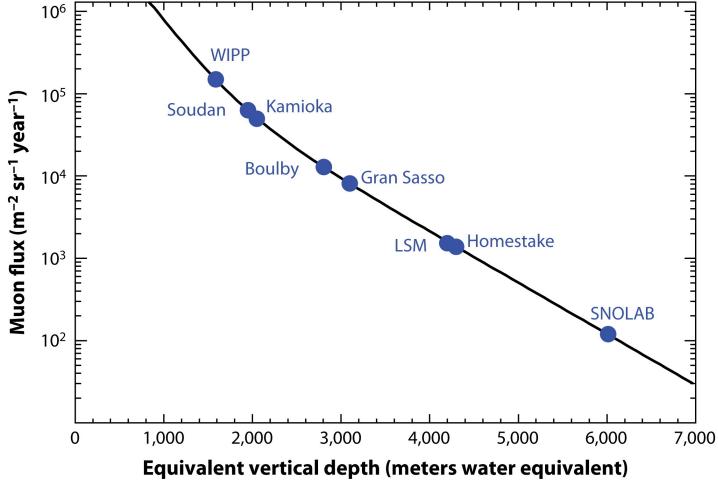


Figure 2.6: Cosmic muon flux as a function of the vertical depth given in meters water equivalent (mwe) for several underground laboratories. The solid line shows a parametrization with a flat overburden. Taken from [22].

crystal based detectors, the energy resolution is sufficient so that $2\nu\beta\beta$ decays do not contribute significantly to the background level. Nevertheless, any improvement to energy resolution results in a narrower energy window ΔE around the Q value and thus improves the signal-to-noise ratio.

Further aspects to reduce the background contamination include careful material selection of radiopure detector components. Detector concepts that allow particle identification, i.e. discriminating $\beta\beta$ decays from background γ and α interactions during data analysis, outperform simple calorimeter concepts in terms of background suppression. The ability to tag the daughter ion of $2\nu\beta\beta$ and candidate $0\nu\beta\beta$ decays provides another powerful tool to reject most of the remaining backgrounds. Together with a decent energy resolution, such an experiment can effectively be operated background-free. Since the technical realization is challenging, it was not utilized in $0\nu\beta\beta$ experiments so far. However, it is being developed as a potential upgrade for the next-generation experiment nEXO [37, 38].

Different experimental concepts have different capabilities to reject or suppress background. Due to the high interaction probability of α particles, they interact mostly close to material surfaces. Most detector concepts can handle this type of background by e.g. event position reconstruction. Background from atmospheric muons is suppressed passively by placing the experiment underground. The muon flux for different underground facilities is shown in Figure 2.6. Remaining muon interactions are usually monitored by active muon veto systems surrounding the actual $0\nu\beta\beta$ detector. While muon interactions can be suppressed successfully with these measures, cosmogenic activation from neutrons that result from muon interactions represents a risk of increasing the radioactivity of the detector components. Remaining background originate mostly from trace radioactive impurities in the detector components and from natural radioactivity external to the detector. The latter contribution is usually reduced by passive shields that absorb this background.

Besides controlling the background level in a real experiment, another key requirement is a high signal detection efficiency. As indicated by Equation 2.11, improving the sensitivity to $m_{\beta\beta}$ by doubling the signal efficiency ϵ is equivalent to using four times the isotope mass [39]. To achieve a high signal detection efficiency, homogeneous detector concepts are favored where the source material is at the same time the detection medium. These concepts do not lose efficiency from e.g. geometric acceptance or absorption in the source material itself.

The $0\nu\beta\beta$ decay could in principle be observed in 35 isotopes. However, several aspects need to be taken into account in order to choose an isotope where an observation is more likely and technically feasible. From theoretical aspects (see Equation 2.8) isotopes are favored where the phase space factors and nuclear matrix elements maximize the $0\nu\beta\beta$ decay rate for a particular Majorana neutrino mass. From background control aspects, isotopes with high Q values are favorable since most natural backgrounds have energies below ~ 3 MeV. As discussed in the example above, large masses of the $0\nu\beta\beta$ isotope are needed to achieve a competitive $0\nu\beta\beta$ sensitivity. This already eliminates some rare and expensive isotopes. Another desirable aspect of isotopes is enrichment in order to produce concentrated masses of the $0\nu\beta\beta$ isotope at a reasonable price. Especially for detector concepts without excellent energy resolution, isotopes with a slow $2\nu\beta\beta$ decay rate of the corresponding isotope are preferred to constrain this background contribution. For detector concepts with excellent energy resolution, this aspect is not relevant as the background contribution of $2\nu\beta\beta$ decays is negligible.

The various aspects contributing to the design of a successful experiment allow for different detector concepts each with their intrinsic advantages and disadvantages. The key parameters of selected current and next-generation experiments are shown in Table 2.1. More details on next-generation experiments are discussed in Ref. [40]. Some of the current and next-generation experiments are:

CUORE deploys TeO_2 -crystals in an ultra-low temperature cryostat and pursues a bolometric detector concept where the source material equals the detector medium. An advantage is the high natural abundance of the $0\nu\beta\beta$ isotope ^{130}Te which mitigates the need of isotope enrichment. Deposited energy in a crystal is measured as a temperature increase which allows for an excellent energy resolution. However, since the crystals cannot be grown very large, many of these crystals are organized in strings as opposed to a single monolithic detector volume [41].

GERDA deploys Ge-crystals enriched in ^{76}Ge that are immersed in a cryostat of liquid argon. The advantages of this detector concept are an excellent energy resolution and low background levels. However, similar to the CUORE experiment, the crystal-based concept does not allow for a monolithic detector [15]. After an intermediate upgrade, the next-generation experiment **LEGEND** will deploy ~ 1000 kg of ^{76}Ge [42].

KamLAND-Zen searches for the $0\nu\beta\beta$ decay in ^{136}Xe with the previously existing detector infrastructure of the KamLAND experiment [43]. The experiment consists of a liquid scintillator loaded with xenon which is enriched in ^{136}Xe . While the detector concept suffers from modest energy resolution, it benefits from a high deployed isotope mass and a straightforward way to scale the monolithic detector to even larger mass [44].

NEXT will use a time projection chamber of high-pressure gaseous xenon enriched in ^{136}Xe to search for the $0\nu\beta\beta$ decay in ^{136}Xe . Key advantages of this concepts are the excellent energy resolution and the topological event information by measuring both the ionization and scintillation signatures. With these features, backgrounds can be rejected effectively [45].

SNO+ will search for the $0\nu\beta\beta$ decay in ^{130}Te with the existing detector infrastructure of the SNO detector [46]. The detector concept is a liquid scintillator loaded with natural tellurium. Similar to the KamLAND-Zen experiment, the detector concept suffers from modest energy resolution. However, it offers the possibility to deploy much source material and a monolithic design [47].

EXO-200 and **nEXO** uses and will use a time projection chamber filled with liquid xenon enriched in ^{136}Xe . Key advantages are the multiparameter analysis concept that is based on the measurement of ionization and scintillation signatures. Furthermore, the high density of liquid xenon in a monolithic detector concepts provides an intrinsic self-shielding by the xenon from external background [14, 37, 48]. More details are discussed in Section 2.2.

Table 2.1: List of current and next-generation $0\nu\beta\beta$ experiments. The energy resolution is given at the Q value (FWHM). The reported $0\nu\beta\beta$ half-life limits and sensitivities are given at the 90 % confidence level (CL). For next-generation experiments, the $0\nu\beta\beta$ half-life sensitivity denotes the median 3σ discovery sensitivity after 5 yr of livetime. Numbers for next-generation experiments are taken from [40].

Current Experiment	Nucleus	Isotopic Exposure [kg yr]	Energy Resolution [keV]	Half-life Limit [10^{25} yr]	Reference
CUORE	^{130}Te	24.0	7.4	1.5	[41]
EXO-200	^{136}Xe	234.1	66	3.5	[14]
GERDA	^{76}Ge	82.4	3.3	9	[15]
KAMLAND-Zen	^{136}Xe	126	270	10.7	[44]
Next-generation Experiment	Nucleus	Isotopic Mass [kg]	Energy Resolution [keV]	Half-life Sensitivity [yr]	Important Publications
NEXT	^{136}Xe	1367	12	$5.3 \cdot 10^{25}$	[45]
SNO+	^{130}Te	$7.9 \cdot 10^3$	135	$4.8 \cdot 10^{26}$	[47]
LEGEND	^{76}Ge	873	3	$4.5 \cdot 10^{27}$	[42]
nEXO	^{136}Xe	$4.5 \cdot 10^3$	58	$4.1 \cdot 10^{27}$	[37, 48]

2.2 The EXO-200 experiment

2.2.1 Conceptual design

The EXO-200 experiment was located in the Waste Isolation Pilot Plant (WIPP) in New Mexico (USA) and completed data taking in December 2018. To date, the experiment is among the most sensitive experiments to search for the hypothetical $0\nu\beta\beta$ decay, in particular in ^{136}Xe . The experiment was operated in two phases; a Phase-I from 2011 to 2014 and, after upgrades to the detector, a Phase-II from 2016 to 2018. The main components of the EXO-200 detector are shown in Figure 2.7. The detector itself is a single phase time projection chamber (TPC) filled with liquid xenon (LXe) enriched to 80.6 % in ^{136}Xe [49]. Key advantages of the EXO-200 design are the monolithic detector where the source material is identical to the detection medium, and the complementary measurement of ionization and scintillation signals. The signal formation is explained in detail in Section 2.2.2. To operate the TPC at LXe conditions at about -110°C , it is surrounded by a thermal bath of cryofluid which is housed in a double-walled vacuum-insulated cryostat that shields it thermally. This again is surrounded by a lead shield in order to suppress external background [50, 51].

The TPC is double-sided with respect to the common cathode at the center and is shown schematically in Figure 2.8. It has a length of $\sim 44\text{ cm}$ and a diameter of $\sim 40\text{ cm}$. An electric field is applied in both sub-TPCs. A picture of one sub-TPC is shown in Figure 2.9. Field shaping rings ensure a homogeneous electric field in the TPC. Both sub-TPCs are equipped with identical detector systems, each consisting of two crossed wire planes for ionization detection and an array of 234 large area avalanche photodiodes (APDs) for scintillation light detection [52]. The scintillation light is collected and detected simultaneously by both APD arrays since the cathode is optically transparent ($\sim 90\%$) to xenon scintillation light at 178 nm. To reduce the

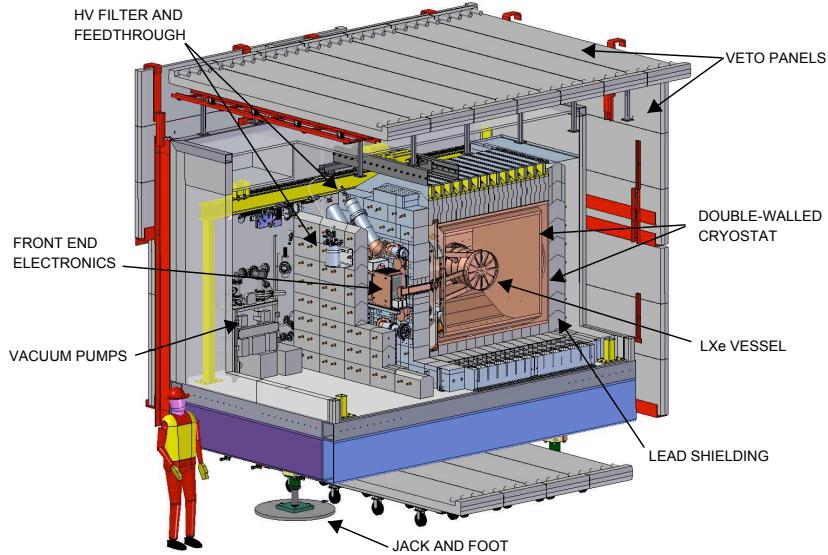


Figure 2.7: Conceptual design of the EXO-200 experiment. Taken from [50].

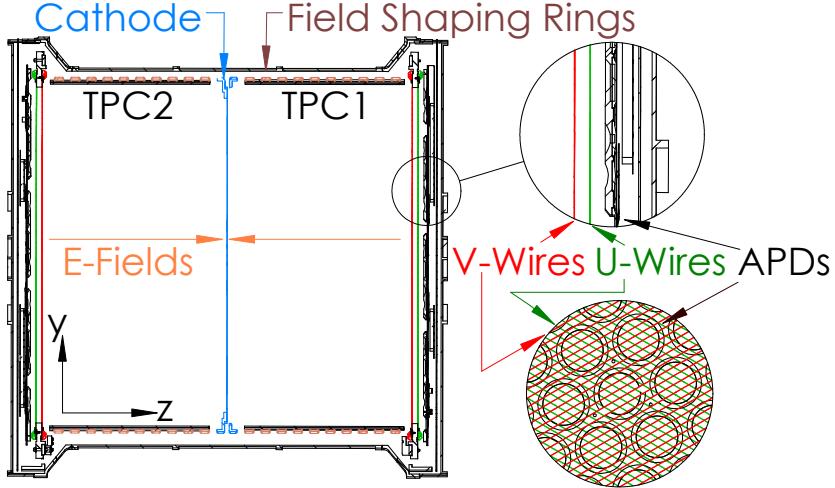


Figure 2.8: Schematic view of the double-sided EXO-200 TPC. Both sub-TPCs share a common cathode (blue) and are each equipped with detector systems consisting of a V-wire plane (red), a U-wire plane (green), and an array of APDs (black). Field shaping rings (brown) ensure a homogeneous electric field between the cathode to the crossed wire planes. Taken from [53].

number of readout channels, the APD array is grouped in 37 APD gangs. To increase the yield of optical light, the inner lateral sides of the TPC are covered with PTFE tiles which act as reflectors. For ionization charge detection, an electric field guides the secondary charges to the corresponding wire planes. The cathode is biased at -8 kV in Phase-I of EXO-200 operation and at -12 kV in Phase-II to improve the energy resolution. The first wire plane sees induced charge and shields the second wire plane which acts as the anode. These wires also see induced charge but more importantly they collect the drifting charge. The wires in both planes are segmented with a pitch of 3 mm and are read out in 38 wire triplets per plane, resulting in an effective wire pitch of 9 mm . The signals are read out by charge-sensitive preamplifiers. The wire planes are offset by 6 mm from each other and the APD array is offset by 6 mm from the anode plane. The induction wire plane is referred to as V-wire plane and the anode plane as U-wire plane. The electric drift field in the bulk volume is 380 V cm^{-1} (567 V cm^{-1}) in Phase-I (Phase-II) and the electric field between both wire planes is 778 V cm^{-1} . In this arrangement, the V-wire plane does not collect charge but is completely transparent to it [51].

Great effort is taken to reduce background in the detector [50, 51]. The background level is a key parameter for any $0\nu\beta\beta$ experiment (see Section 2.1.2). Several aspects contribute to accomplish a low background contamination. To reduce muon induced backgrounds, the experiment is located underground at WIPP with an $1624^{+22}_{-21}\text{ mwe}$ overburden [55] as is shown in Figure 2.6. With this overburden, the muon flux is reduced by about four orders of magnitude compared to that observed at sea level [22, 55]. Remaining muons that interact in the EXO-200 experiment are actively suppressed by a muon veto system located at four sides around the detector which is shown in Figure 2.7. Radiation originating from sources external to the

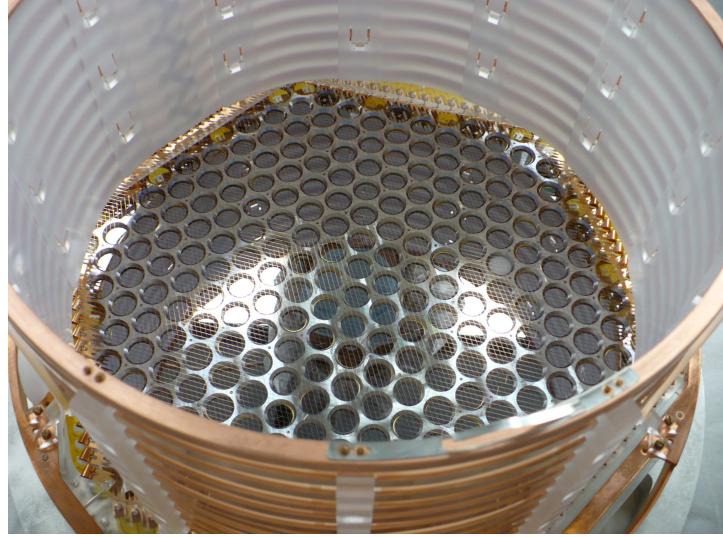


Figure 2.9: Picture of one of the two sub-TPCs installed in the EXO-200 TPC. The white PTFE tiles inside the field shaping rings reflect the scintillation light emitted in the xenon volume. Below the crossed wire planes, the empty platter is equipped with APDs in the final assembly [54].

experiment, e.g. from the surrounding rock, is passively shielded by layers of lead, copper, and coolant. Radioactive impurities in the internal detector components were minimized prior to construction by a careful material selection process [56, 57] and by minimizing the mass of the vessel components. Impurities in the LXe itself are reduced by permanently re-circulating it through heated getters outside the cryostat. Besides the background reduction aspects in detector construction, the EXO-200 detector design features multiple measures to suppress and reject background in data analysis. For example, α particles interact mostly close to the material surfaces and show a high light-to-charge ratio [58]. Using the 3D position reconstruction capability of EXO-200 (see Section 2.2.4) and by measuring both the scintillation and ionization signals (see Section 2.2.2 and 4.3.1), most this background can be rejected. Furthermore, the monolithic detector design naturally reduces the background contamination in the central LXe in the TPC by absorbing it in the outer LXe. This process is referred to as self-shielding. Interactions of γ particles can be distinguished to some extent from $2\nu\beta\beta$ and candidate $0\nu\beta\beta$ decays based on their number of localized energy deposits (see Section 2.2.4). Additional information like the spatial distributions and the detailed event topology further help to discriminate between penetrating γ particles in contrast to $2\nu\beta\beta$ and candidate $0\nu\beta\beta$ decays. Such a discriminator is discussed in Chapter 5.

2.2.2 Working principle

Particle interactions in the TPC deposit energy in the LXe. γ particles interact with xenon atoms and produce an energetic photoelectron while $2\nu\beta\beta$ and candidate $0\nu\beta\beta$ decays emit two energetic electrons. Either way, these electrons gradually deposit their kinetic energy E_0 in the LXe by ionizing and exciting surrounding xenon atoms along their paths. This results in N_i electron-ion Xe^+ pairs and in N_{ex} excited xenon atoms Xe^* . In addition, some energy is lost in undetectable channels such as atomic motion or “heat”. The energy transfer of the electrons can thus be written as [14, 58]:

$$E_0 = N_i E_i + N_{\text{ex}} E_{\text{ex}} + N_i \epsilon , \quad (2.12)$$

where E_i and E_{ex} denote the mean energy required to ionize and to excite a xenon atom. Some energy is transferred to sub-excitation electrons with a mean kinetic energy ϵ .

Through intermediate processes, the xenon ions can also be converted to excited xenon atoms. The excited atoms form excited xenon dimers Xe_2^* which de-excite to the ground state by emitting UV-scintillation light with a peak wavelength of 178 nm in LXe [59]. Since the scintillation light is produced from the identical intermediate conversion, it is independent of the initial process of either ionization or excitation. The scintillation light has two decay components which arise from the de-excitation of either the singlet or the triplet state. However, since both decay times are on the order of nanoseconds, they cannot be distinguished in EXO-200.

In the presence of an electric field, only a fraction r of electron-ion pairs may recombine while the remaining charge is separated spatially and may be detected as $N_q = (1 - r) N_i$ electrons or xenon ions. Then, assuming a recombination efficiency of one, the number of detectable photons increases to $N_\gamma = N_{\text{ex}} + r N_i$. Assuming no photon reduction processes, the mean energy to produce a scintillation photon is [58, 60]:

$$W_\gamma = \frac{E_0}{N_\gamma} = \frac{W_i}{1 + N_{\text{ex}}/N_i} , \quad (2.13)$$

where $W_i = E_0/N_i$ denotes the mean energy to ionize a xenon atom. Due to the specific density of electron-ion pairs along the track of an ionizing particle, the fraction r of recombination further depends on the type of the ionizing particle. While it is $W_\gamma^{(\gamma)} = 21.6 \text{ eV}$ for γ particles, it is $W_\gamma^{(\alpha)} = 17.9 \text{ eV}$ for α particles [58].

Both the number of photons N_γ as well as the number of secondary electrons N_q are measures for the total energy E_0 deposited by an ionizing particle. However, event-to-event fluctuations in the recombination fraction translate to fluctuations in the numbers of photons and secondary electrons, with both numbers being anti-correlated. This leads to a degradation of the energy resolution of either signature [61]. By combining both complementary measurements, these fluctuations cancel and improve the energy measurement drastically as is shown in Figure 2.10. The analytical approach to combine both measurements in EXO-200 is discussed in Section 4.3.1.

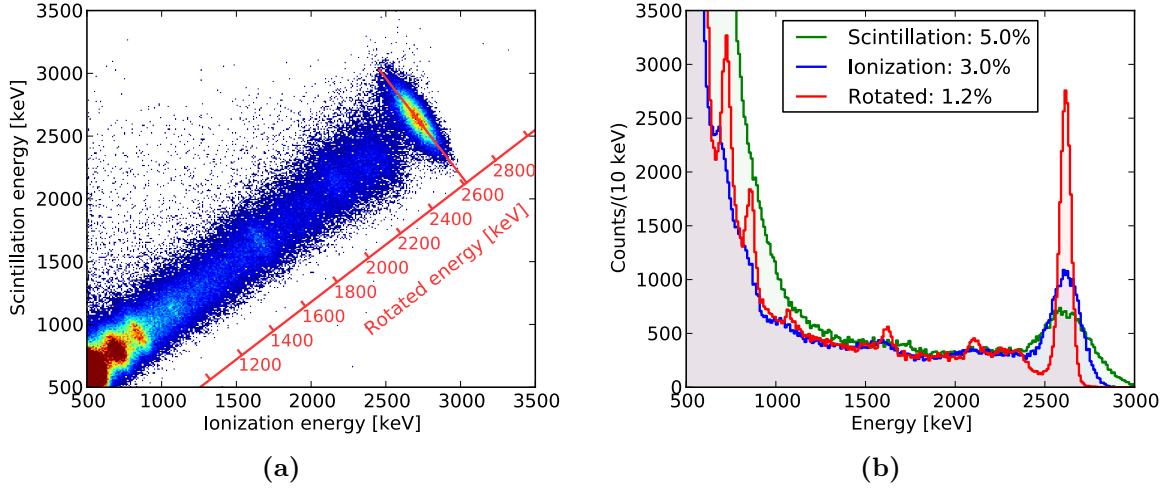


Figure 2.10: Calibration data of a ^{228}Th source. The anti-correlation of the scintillation and ionization measurements is observed in (a). This recombination fluctuations cancel when combining both signatures which leads to an improvement in energy resolution (b). Taken from [62].

2.2.3 Overview of the detector Monte Carlo simulations

A detailed understanding and modeling of the detector is essential for the analysis approach pursued by EXO-200. The Monte Carlo (MC) simulation consists of two separated stages and produces electronic signals of all detector subsystems similar to those recorded in measured data.

The first stage is performed using the GEANT4 software package, version 4.9.3p02 [63, 64]. The detailed detector geometry of the TPC and its internal components is implemented from the CAD model, along with its surroundings. In this detector model, particle interactions and energy deposits are simulated. The output of this simulation stage is a list of energy deposits in the LXe volume.

The second stage calculates the response of the electronics to the energy deposits modeled in the first stage. Therefore, a 2D model of the detector is implemented in this simulation to model the drift of secondary electrons in the bulk volume. To improve the precision of the simulation, a full 3D model is implemented near the wire planes. In the 2D model, the U- and V-wires are parallel to each other and the V-wires lie directly above the U-wires as is shown in Figure 2.11. The path of each energy deposition is sampled to cubic voxels with $0.2\text{ }\mu\text{m}$ edges and is tracked in a finite element simulation. The weighting potential $\phi(\mathbf{x})$ and the electric field $\mathbf{E}(\mathbf{x})$ are used to calculate the induced signal at each time step for every affected channel using the Shockley-Ramo theorem [65, 66]. Diffusion of the secondary electrons is taken into account in the simulation. The diffusion depends on the orientation of the electric field. The transverse diffusion coefficients D_T is expected to exceed the longitudinal diffusion coefficients D_L by $D_T \simeq 10D_L$ [67]. The transverse diffusion σ_{D_T} of secondary electrons is given

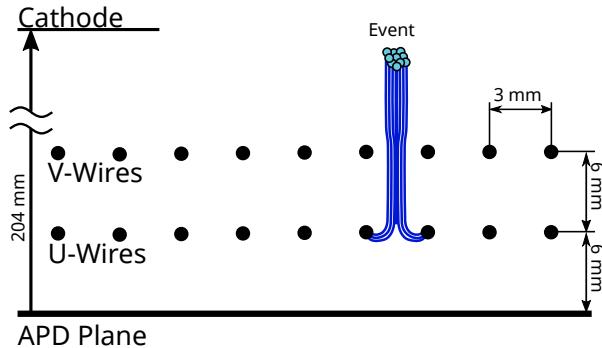


Figure 2.11: Geometry of the wire planes showing the drift of ionization charges. As opposed to the implementation in the 2D model shown here, in reality, the wire planes are crossed by 60°. Taken from [68].

by:

$$\sigma_{D_T} = \sqrt{D_T t_d} , \quad (2.14)$$

where $t_d = d/v_d$ denotes the time required to drift a distance d with a drift velocity v_d . The drift velocity is set by the external electric field E via $v_d = \mu E$, with μ being the electron mobility in LXe [58]. The transverse diffusion coefficient is measured in-situ in EXO-200 [53]. Raw waveform traces are calculated from the induced signals with the transfer functions of the corresponding electronics. In order to make the simulated waveforms more realistic, randomized measured noise traces are added to them.

The APD response is simulated in a simplified fashion to avoid tracing single photons. A parametrized light response is used to determine the light yield on both APD planes. This response function takes the position of the interactions into account as well as geometric and optical factors. The number of photons per APD plane is divided equally over all channels which is then used to calculate raw waveforms with the known transfer function. The simulation does not include the recombination fluctuations in LXe that are discussed in Section 2.2.2 since key parameters are still subject to current research.

2.2.4 Overview of the event reconstruction and data analysis

The event reconstruction in EXO-200 consists of multiple stages. First, a signal finding algorithm is used to identify signals in the raw waveform traces. This applies a matched filter in Fourier space to the waveforms. Because the detector subsystems use different settings in the front-end electronics, individual transfer functions are used to optimally find potential signals. The filtered trace $y(t)$ is given by:

$$y(t) = \mathcal{F}^{-1} [X(f)H^*(f)] , \quad (2.15)$$

where \mathcal{F} is the Fourier transform (FT), $X(f)$ is the FT of the original waveform trace $x(t)$, and $H^*(f)$ is the complex conjugate of the FT of the transfer function $h(t)$ [51, 69]. Potential signals on the filtered trace are found by identifying amplitudes that exceed a threshold determined

on an event basis from the mean absolute deviation. To improve the ability to find multiple signals close in time, found signals in the matched filter stage are “unshaped” to obtain the original charge deposit $q(t)$:

$$q(t) = \mathcal{F}^{-1} [H^{-1}(f)X(f)] , \quad (2.16)$$

where $H^{-1}(f)$ is the inverse transfer function. This unshaped trace is sensitive to nearby charge deposits which are identified by a peak-search algorithm that finds additional signals.

After finding signals, the corresponding signal parameters are extracted. This is done by minimizing a χ^2 function that is based on the template model for the specific detector subsystem. The results from the signal finding stage serve as initial fit values and to set the fit window. For U-wire channels, the signal finding and parameter extraction stages are repeated twice; once with the collection and once with the induction U-wire template. By comparing the χ^2 of both fit scenarios, signals may be tagged as induction-like, i.e. they do not contain net energy. As of the EXO-200 $0\nu\beta\beta$ search in 2019 [14], these signals tagged as induction-like further pass a stage where a combined fit of collection and induction templates is performed to recover potentially missed energy on the U-wires. This improvement was triggered by the results presented in Section 4.2.3 and published in Ref. [70].

Signals on adjacent channels are grouped together if they are close in time. These bundles are assumed to belong to a single physical event. Signals on U-wires are linked to signals on V-wires by grouping all bundles in their most probable configuration. The resulting groups are called charge clusters. In this clustering process, all potential combinations of U- and V-wire bundles are matched. The best configuration is selected by minimizing a cost function that is defined based on probability density functions (PDFs). These PDFs describe the time difference and the amplitudes between U- and V-wire bundles and whether the resulting 2D position is contained within the detector volume. The resulting charge clusters are associated to APD bundles. The full 3D position is reconstructed with the measured drift velocity in LXe [53] and the time difference between charge and scintillation clusters. It is noted that multiple charge clusters can be associated to a single scintillation cluster, e.g. for γ particles that interact multiple times in the detector via the Compton scattering process. On the other hand, multiple scintillation clusters cannot be associated to a charge cluster.

The orientation of the coordinate system is shown in Figure 2.13a. The event location (x, y) is calculated from the coordinates associated with the wire plane coordinates (u, v) while the z coordinate is calculated from the time difference of charge and scintillation clusters:

$$x = \begin{cases} v - u & \text{for } z > 0 \\ u - v & \text{else} \end{cases} \quad (2.17)$$

$$y = \frac{1}{\sqrt{3}}(u + v) \quad (2.18)$$

$$|z| = \begin{cases} d_U - d_{VU} \cdot \frac{\Delta t}{t_{\text{coll}}} & \text{for } \Delta t < t_{\text{coll}} \\ d_U - d_{VU} - v_{\text{drift}}(\Delta t - t_{\text{coll}}) & \text{else} \end{cases} . \quad (2.19)$$

While the cathode is located at $z = 0$, d_U denotes the distance of the U-wire plane to the cathode and d_{VU} the distance between both wire planes of a sub-TPC. $\Delta t = t_U - t_{\text{APD}}$ denotes

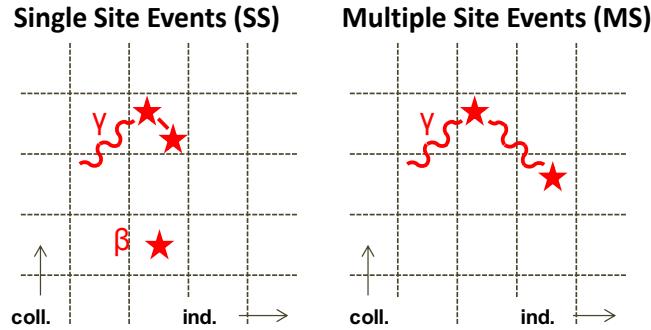


Figure 2.12: Illustration of example interactions categorized as “single-site” (SS) and “multi-site” (MS). Taken from [54].

the measured time difference of associated charge and scintillation clusters. v_{drift} denotes the drift velocity in the bulk volume which is $1.71 \text{ mm}\,\mu\text{s}^{-1}$ in Phase-I [53] and $1.83 \text{ mm}\,\mu\text{s}^{-1}$ in Phase-II. t_{coll} is the drift time between both wire planes which is $2.940 \mu\text{s}$ ($2.887 \mu\text{s}$).

Thanks to the monolithic design of the EXO-200 detector along with good position and energy resolution, many background events can be distinguished from $2\nu\beta\beta$ and candidate $0\nu\beta\beta$ decays. As was discussed in Section 2.2.1, the dominant background sources are γ interactions originating from remaining trace impurities in the detector materials. Typical interactions of γ particles and $\beta\beta$ decays are shown schematically in Figure 2.12. Events are labelled “single-site” (SS) if all energy is reconstructed in a single location within dimensions of $\sim 2\text{--}3 \text{ mm}$. This is observed for most $2\nu\beta\beta$ decays. On the other hand, events are labelled “multi-site” (MS) if there are multiple localized energy deposits in the detector. In the main energy range of EXO-200 analysis ($1000\text{--}3000 \text{ keV}$), γ particles mostly interact via the Compton scattering process and usually interact multiple times in the LXe. Thus, these events are categorized as MS. Interactions of γ particles may be categorized as SS if they interact via the photoelectric effect, if multiple Compton scattering processes occur too close to each other to be separated by the event reconstruction, or if energy deposits are missed by event reconstruction. $2\nu\beta\beta$ decays and γ interactions may produce additional localized energy deposits due to bremsstrahlung and thus are categorized as MS. For example, the fraction of SS events of $0\nu\beta\beta$ decays is predicted in MC simulations to be $\sim 75\%$ whereas that of γ particles is $\sim 15\%$ at $E_\gamma = Q$.

To monitor the TPC’s response to energy deposits, in particular from γ interactions, radioactive calibration sources are deployed near the TPC on a regular basis. Therefore, a dedicated guide tube allows to inject sources into the coolant volume near the TPC to various source locations that are shown in Figure 2.13b. For regular calibrations, a ^{228}Th source is deployed at source position S5 near the cathode in order to cover both sub-TPCs. Additional calibration sources like ^{137}Cs , ^{60}Co , and ^{226}Ra and additional source positions are used during extended calibration campaigns. The regular calibration measurements are used to evaluate the energy response and the electron lifetime in LXe which is discussed in Section 4.3.

The gains of individual U-wire channels are monitored periodically by dedicated charge injection measurements. A precisely known amount of charge is injected to individual wires

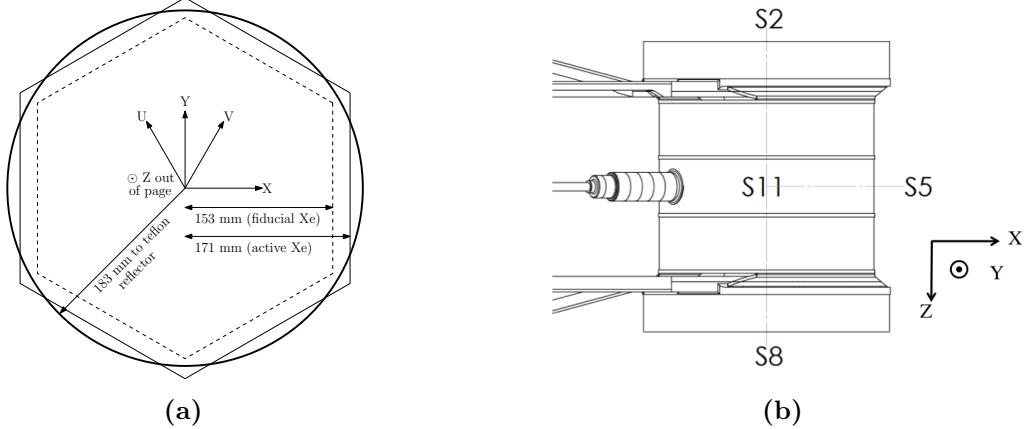


Figure 2.13: (a) shows the coordinate systems (u, v, z) and (x, y, z), projections of both the hexagonal active and fiducial LXe and of the PTFE panels. (b) shows the location of calibration sources around the TPC. The cathode is located at the center and is biased via the HV feedthrough. Taken from [51].

and their response is measured in units of electrons per ADC count. In addition, the gains are calculated from ^{228}Th source calibration data. Both measurements are highly correlated and result in a mean gain of ~ 380 e/ADC with a variation of 30 % between the channels. The V-wire channels are also monitored with dedicated charge injection measurements. The response of the APDs is monitored periodically using an external laser. However, the variation is absorbed in an APD light map $f(r, \phi, z)$ that corrects the amount of collected light as a function of the location of the energy deposit. The light yield dominantly depends on the event position due to solid angle effects and due to variations in the APD response [51].

Besides calibrations and corrections, several event selection criteria are applied in the analysis of low-background (LB) physics data. These include a dead-time after the muon veto system is triggered or an event in the TPC was tagged as muon in order to reject muon induced events. The time difference of consecutive events is required to be at least 0.1 s in order to remove fast β - α coincidences that occur from ^{220}Rn and ^{222}Rn radioactive impurities as well as other correlated decays [14]. Pile-up events with multiple scintillation clusters are rejected. Also, events with high light-to-charge ratios are rejected which removes background from α particles (see Section 2.2.2 and Ref. [58]). In addition, all particle interactions are required to be associated with a full 3D position. For some charge clusters, in particular those with low energy deposits, the position reconstruction fails. As of the EXO-200 $0\nu\beta\beta$ search in 2019 [14], this selection cut is relaxed to only require at least 60 % of the total energy to have an associated 3D position. This cut is discussed in detail in Section 6.1.4. The event position of all particle interactions is required to be reconstructed in a fiducial volume. By definition, it extends to a distance of 10 mm from the cylindrical PTFE, the cathode plane, and the V-wire plane. In addition, motivated by the profile of the wire planes, the energy deposits are required to lie in a hexagon in the xy plane with an apothem of 162 mm. The fiducial volume selection criterion is applied to exclude regions of the TPC near the material surfaces with increased discrepancies in MC simulations. The energy threshold for $0\nu\beta\beta$ searches is 1000 keV [14, 51].

Chapter 3

Deep Learning

Contents

3.1	Machine learning	24
3.2	Deep neural networks	25
3.2.1	Multilayer neural networks	25
3.2.2	Network training	27
3.2.3	Convolutional neural networks	30

Machine Learning is a popular concept in artificial intelligence which is introduced in this chapter. While many different approaches exist within machine learning, the state of the art concept is deep learning. This concept has proven to be applicable to a variety of problems with superior performance.

This chapter introduces deep neural networks which are used in deep learning. The concept of these networks is discussed based on multilayer neural networks. These networks are also utilized to discuss the general training procedure for deep neural networks. Finally, convolutional neural networks are introduced which have proven to outperform other approaches especially in the field of image recognition. Besides theoretical concepts, also practical considerations are outlined.

3.1 Machine learning

Currently, the most popular field in artificial intelligence is machine learning (ML). This sub-field covers algorithms that are able to correctly interpret input data without explicit instructions how to process it. Instead, the network is trained to do so implicitly by pursuing specific objectives [71]. A widely used definition of algorithms belonging to machine learning is:

A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E .

(T.M. Mitchell)

Based on this formal definition, many different methods and techniques exist to construct machine learning algorithms. While common choices of tasks T , performance measures P , and experiences E are discussed here, a more detailed overview is given in Ref. [71].

The task T describes how to process a given input. Common tasks include classification, regression, and synthesis. In classification, an algorithm f processes an input $\mathbf{x} \in \mathbb{R}^n$ and assigns it to a category k . Mathematically, this corresponds to $f : \mathbb{R}^n \rightarrow \{1, \dots, K\}$ where K denotes the number of different categories. In addition, more advanced classification tasks exist that map the input to multiple categories or to a probability distribution over all categories. A popular application in industry includes object recognition in images. In regression tasks, an algorithm f produces a continuous output $f : \mathbb{R}^n \rightarrow \mathbb{R}$. The difference to classification tasks is the natural ordering provided by \mathbb{R} while the categories in classification tasks are not ordered logically. For synthesis tasks, new samples are generated by an algorithm, i.e. $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$. For example, this includes the generation of images from random numbers according to the reference dataset used in training. In other applications, a corrupted input $\tilde{\mathbf{x}} \in \mathbb{R}^n$ is processed to an output $\mathbf{x} \in \mathbb{R}^n$ to remove noise [71] or to refine the accuracy of simulated samples [72].

The performance measure P denotes the metric that is used to evaluate the ability of an algorithm to perform the task T . While different metrics exist, the exact choice of suitable measure depends on the specific task. The ability to process new data is usually evaluated on an independent validation dataset. The choice of performance measure affects the behaviour of the trained algorithm. For example, for regression tasks, the effect of rare but large errors in contrast to many medium-sized errors in the prediction is a tradeoff that depends on the specific application.

The experience E is used to improve the performance measure P on a task T in training. Common types of training procedures are supervised and unsupervised learning. The latter concept aims at finding useful patterns in an unlabelled training dataset. During unsupervised training, the algorithm learns the underlying probability distribution that describes the training dataset. This knowledge can be used for data synthesis or to cluster data with similar attributes. In contrast to this concept, in supervised learning each training sample \mathbf{x} is associated to a target label y . The algorithm predicts \hat{y} from the input \mathbf{x} . During training, a loss function $L(y, \hat{y})$ is used to quantify and optimize the performance measure. Classification and regression tasks are usually trained via supervised learning.

Machine learning algorithms consist of a models which are fed with samples during training in order to learn to correctly interpret it according to some objective function. In a successful application, the model has learnt a representation of the training data and is able to also correctly interpret new, unseen data. Many different types of model concepts exist, such as support vector machines, decision trees, and neural networks [71]. Classical machine learning techniques require to manually engineer features from the data based on prior knowledge. However, neural networks are capable to learn these features by themselves during training with little pre-processing required. The concept of neural networks is inspired by biological neural networks. In computer science, neural networks consist of a collection of neurons, also called units, that are aggregated into layers. Units of different layers can be connected to each other which implies a flow of information between these units. Historically, “shallow” neural networks were used with only a few hidden layers which denote the layers between the input and output layers. While there is theoretical prove that a shallow neural network is sufficient to represent any function [73], this approach is not practicable. Deep learning (DL) represents a class of advanced concepts that stack multiple layers on top of each other. With this hierarchical approach, complex representations (features) can be extracted by combining simpler features captured in preceding layers. This capability renders deep learning techniques very powerful and flexible for numerous applications. Accordingly, these neural networks are called deep neural networks. This is the state of the art approach to solve many real-world problems like object and speech recognition in a wide variety of fields in research and industry [17, 18, 74–77].

3.2 Deep neural networks

3.2.1 Multilayer neural networks

The basic concept of deep neural networks is multilayer neural networks. A simple example with two hidden layers is shown in Figure 3.1. The nodes represent mathematical operations and the lines represent weighted connections between the nodes. Since each node is connected to every node in the previous and the next layer, these layers are called fully connected layers. The output $y_k(\mathbf{x})$ of a node k in the first hidden layer is computed for an input feature vector \mathbf{x} according to:

$$y_k(\mathbf{x}) = \sigma(z_k) = \sigma \left(\sum_{j=1}^K w_{kj}^{(1)} x_j + b_k^{(1)} \right), \quad (3.1)$$

where $w_{kj}^{(1)}$ denotes the weights of the connection between a node j in the input layer and the node k in the first hidden layer. $b_k^{(1)}$ denotes the bias. z_k is further processed by a non-linear activation function σ . Figure 3.2 shows the workflow to calculate Equation 3.1 for a single unit. For the example neural network shown in Figure 3.1, the sum spans over $K = 4$ input features while Equation 3.1 is evaluated for all nodes $k = 1, \dots, 3$ in the first hidden layer. Currently, the most common non-linear activation function is the rectified linear unit (ReLU):

$$\text{ReLU}(z) = \max(0, z). \quad (3.2)$$

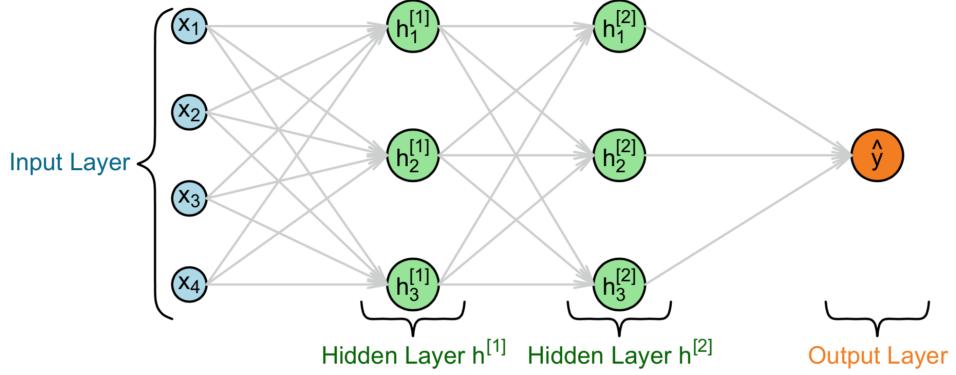


Figure 3.1: Illustration of a simple neural network that consists of four input features, two hidden layers with three nodes each and an output layer with a single node. The connections represent weights. Taken from [78].

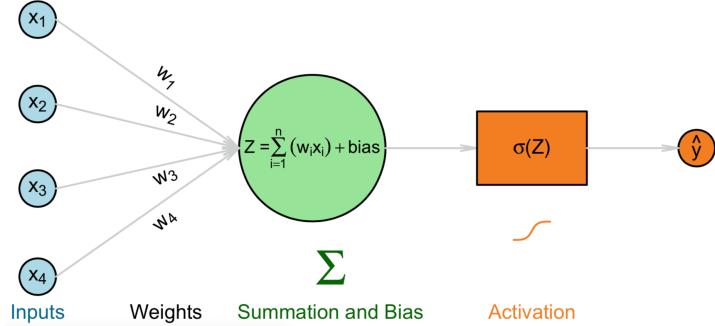


Figure 3.2: Illustration of a unit in a hidden layer of a neural network. The output of the node is calculated based on the output of the previous layer, the weights of their connections, and a bias term. Taken from [78].

Sigmoidal functions are commonly used for the output layer of classification tasks, while for regression tasks the identity $\sigma(z) = z$ is used here.

Stacking multiple layers, the output $y_k(\mathbf{x})$ of a node k in the second hidden layer can be written by stacking Equation 3.1. This is because the nodes in the first hidden layer $k \in K^{(1)}$ are connected to those in the second hidden layer $k \in K^{(2)}$ according to:

$$y_k(\mathbf{x}) = \sigma^{(2)} \left(\sum_{j=1}^{K^{(2)}} w_{kj}^{(2)} \sigma^{(1)} \left(\sum_{i=1}^{K^{(1)}} w_{ji}^{(1)} x_i + b_j^{(1)} \right) + b_k^{(2)} \right) . \quad (3.3)$$

Generally, the output $x_k^{(n)}$ of a node k in layer n depends on the output of the previous layer $(n-1)$ and can be calculated recursively based on the input $x^{(1)}$ by:

$$x_k^{(n)}(\mathbf{x}^{(n-1)}) = \sigma^{(n-1)} \left(\sum_{j=1}^{K^{(n-1)}} w_{kj}^{(n-1)} x_j^{(n-1)} + b_k^{(n-1)} \right) . \quad (3.4)$$

Evaluating this operation for every node in every layer successively is called forward propagation. Here, the information flows only in one direction through the network, i.e. from the input to the output layer. This type of networks is called feed-forward networks. Other concepts exist, e.g. with cyclic connections [71]. Introducing an additional virtual unit in each layer with a fixed activation of 1, the bias vectors \mathbf{b} can be absorbed in the weights vectors \mathbf{w} . This simplifies the notation so that the weights \mathbf{w} represents both vectors [79]. Using Equation 3.4, neural networks can be described as a set of parametric non-linear functions which map an input vector \mathbf{x} to an output \hat{y} . For a fixed set of weights \mathbf{w} , these networks produce a deterministic output for a given input.

3.2.2 Network training

The performance of a network is improved in a training process. Training in this context represents adjusting the weights \mathbf{w} in such a way that the performance $E(\mathbf{w})$ improves with respect to a loss function. For regression tasks, a common choice of loss function is the mean squared error over a set of N samples:

$$L_2(\mathbf{w}) = \frac{1}{2N} \sum_{n=1}^N \|\hat{y}(\mathbf{x}_n, \mathbf{w}) - y_n\|^2, \quad (3.5)$$

where $\hat{y}(\mathbf{x}_n, \mathbf{w})$ is the output of the network calculated based on Equation 3.4. y_n is the target which is provided in supervised training.

The goal of the training is to find a set of weights \mathbf{w} that optimizes the performance measure $E(\mathbf{w})$. During training, small changes to the weights $\mathbf{w} \rightarrow \mathbf{w} + \delta\mathbf{w}$ induce a change in the performance $\delta E \simeq \delta\mathbf{w}^T \nabla E(\mathbf{w})$. The optimum set of weights is then given where the gradient of $E(\mathbf{w})$ with respect to \mathbf{w} vanishes which implies that the performance measure would not improve by further changes to \mathbf{w} :

$$\nabla E(\mathbf{w}) = 0, \quad (3.6)$$

where the corresponding point $E(\mathbf{w})$ is neither a maximum nor a saddle point. Several local minima in the parameter space fulfill this condition. However, the goal is to find the global minimum. Since it cannot be found analytically in most scenarios, the optimization is solved iteratively starting from a set of initial weights $\mathbf{w}^{(0)}$:

$$\mathbf{w}^{(\tau+1)} = \mathbf{w}^{(\tau)} + \Delta\mathbf{w}^{(\tau)}, \quad (3.7)$$

where τ denotes the iteration step and $\Delta\mathbf{w}^{(\tau)}$ the weight update. In practice, many algorithms calculate the weight update at each iteration step based on the gradient of the loss function $\nabla E(\mathbf{w})$ which points towards the largest gradient. Thus, Equation 3.7 becomes:

$$\mathbf{w}^{(\tau+1)} = \mathbf{w}^{(\tau)} - \eta \nabla E(\mathbf{w}^{(\tau)}), \quad (3.8)$$

where $\eta > 0$ denotes the learning rate and where the weights are updated in direction of the negative gradient. This approach is known as gradient descent. However, more complex and powerful optimization schemes exist [78].

The gradients are usually calculated using the error backpropagation scheme. This means, calculating the gradient of $E(\mathbf{w})$ with respect to each weight w_{ij} . The gradient is then used to update the weight w_{ij} in order to improve the performance $E(\mathbf{w})$. Using the chain-rule, the gradient can be written as:

$$\frac{\partial E}{\partial w_{ji}} = \frac{\partial E}{\partial z_j} \frac{\partial z_j}{\partial w_{ji}}, \quad (3.9)$$

where z_j denotes the weighted sum affecting unit j as is discussed in Equation 3.1. Using this equation and including the bias in the weight vector, the last term in Equation 3.9 simplifies to $\frac{\partial z_j}{\partial w_{ji}} = x_i$. To simplify the notation, the error $\delta_j \equiv \frac{\partial E}{\partial z_j}$ is introduced. Provided the mean squared error is used as the loss function, as defined in Equation 3.5, the error for a unit in the output layer is given by:

$$\delta_k = \hat{y}(\mathbf{x}_k, \mathbf{w}) - y_k. \quad (3.10)$$

Here, the identity was used as activation function of the output layer. In general, the gradients are calculated backwards through the network from the output layer to the input layer. Thus, the gradient for a unit j in a hidden layer depends on the succeeding units in the network by:

$$\delta_j \equiv \frac{\partial E}{\partial z_j} = \sum_k \frac{\partial E}{\partial z_k} \frac{\partial z_k}{\partial z_j}, \quad (3.11)$$

where the sum is performed over all units k in the next layer that are connected to unit j . We denote the derivative of the activation function $\sigma'(z_j)$. The backpropagation formula is then given by:

$$\delta_j = \sigma'(z_j) \sum_k w_{kj} \delta_k. \quad (3.12)$$

With this description, the errors are first evaluated for the units in the output layer via Equation 3.10. Then, using these errors, those of the hidden layers can be evaluated recursively via Equation 3.12.

During training, all available N samples are evaluated by the network with fixed weights in the forward pass by recursively evaluating Equation 3.4 and comparing the outputs to the target values. Then, by recursively evaluating the backpropagation formula in Equation 3.12 and 3.9, the weights w_{ji} are updated. The weight update can be evaluated using the gradient descent introduced in Equation 3.8 in order to achieve an improved average performance E over N samples:

$$w_{ji}^{(\tau+1)} = w_{ji}^{(\tau)} - \eta \sum_n \frac{\partial E_n}{\partial w_{ji}^{(\tau)}}. \quad (3.13)$$

For large sets of training data, the set is usually split into small batches, and the weights are updated based on the performance on every batch. This concept is called mini-batch training. Mathematically, this corresponds to replacing N in Equation 3.13 by the batch size and looping over all batches recursively. A loop over all batches is called an epoch.

Prior to training, the weights of the network are initialized. If all units were initialized with the same weights, they are identical and can therefore not learn meaningful representations. This is mitigated by breaking their symmetry by initializing the weights with random numbers. In practice, different initialization algorithms can lead to different performances and can accelerate

the training process. A common choice is Glorot initialization [80], where the weights i in layer j are initialized with random numbers from a uniform distribution U :

$$w_{ji}^{(0)} \sim U \left[-\sqrt{\frac{6}{n_{j-1} + n_j}}, \sqrt{\frac{6}{n_{j-1} + n_j}} \right], \quad (3.14)$$

where n_{j-1} and n_j are the numbers of units in layer $(j-1)$ and in layer j .

Since neural networks can contain a large amount of trainable weights, they can converge to states that perfectly interpret the training data. However, they may fail to generalize on unseen validation data. This problem is called overfitting. It is promoted by larger networks in terms of number of trainable weights. It can be prevented or constrained by increasing the number of samples in the training dataset. In addition, regularization techniques exist to eliminate or reduce overfitting. Common techniques are for example the artificial extension of the training set by data augmentation or the regularization of the network during training. For example, with the dropout technique, only a fraction $p < 1$ of the units in a layer are used in training while the other units are disabled. The set of disabled units is determined randomly for every batch. With this technique, the network learns a more robust representation of the training data by relying less on individual units. This results in better generalization capabilities on unseen data [81]. Other regularization techniques add an additional penalty to the loss function during training:

$$L'(\mathbf{w}) = L(\mathbf{w}) + \lambda R(\mathbf{w}). \quad (3.15)$$

In practice, common penalty functions $R(\mathbf{w})$ include the $L_1 = \sum \|w\|_1$ norm and the $L_2 = \sum \|w\|_2^2$ norm where the sums are performed over all weights in a network. The magnitude of the penalty function is controlled by the hyperparameter λ . While both penalty functions have slightly different effects, they both promote small weight values. This constrains the impact of individual, potentially large weights and thus reduces overfitting. The optimum regularization technique depends on the specific application, network architecture, and training dataset.

In practice, a dataset is usually divided into a training and a validation set. The training set is utilized in weight optimization. Being statistically independent from the training set, the validation set is used to monitor the network performance and evaluates the generalization capabilities of the network. A degraded performance on validation data compared to training data indicates overfitting. After training, the set of weights of the neural network is usually restored to the state where the performance on validation samples was best.

3.2.3 Convolutional neural networks

Besides fully connected layers, there are other concepts to process the output from a previous layer. One of these is the convolutional layer which was proposed in Ref. [82]. It became popular after achieving superior results in image recognition compared to other approaches [17].

Convolutional layers exploit the structure of the input data by enforcing only local connectivity between the units of subsequent layers. In this way, the spatial correlation of adjacent units is captured. Conceptually, for one-dimensional data, this means that the input of a layer is only a subset of n units from the previous layer. Furthermore, this so-called receptive field is enforced to be spatially contiguous. Analogously for two-dimensional data, like images, the dimension of the receptive field is $n \times m$ units. Many applications use quadratic receptive fields with $n = m$. In addition, the weights are shared between the different receptive fields to ensure translational invariance. This concept is motivated by the aim to capture the same feature at any location of the input feature map. An example of two-dimensional data is shown in Figure 3.3. Equation 3.1 is evaluated for each 3×3 receptive field (blue shaded) of the 4×4 input feature map (blue). The 3×3 grid which contains the trainable weights is called kernel. The output feature map (green) represents the response between the input and the kernel. The output feature map is produced by shifting the kernel over the input feature map and evaluating Equation 3.1 at every location. Mathematically, convolutional layers perform a discrete convolution of the input feature map with the kernel. In common applications, multiple kernels are applied to an input feature map in a single convolutional layer. This is done because each kernel extracts individual information from the input feature map. This allows to capture different types of features at the same time.

The spatial dimensions of the output feature maps are usually smaller than those of the input feature maps as is shown in Figure 3.3. However, for many practical applications, the spatial dimensions are supposed to remain unchanged. To accomplish this, the input feature maps are usually padded with zeros to artificially increase their spatial dimensions. This concept is shown in Figure 3.4.

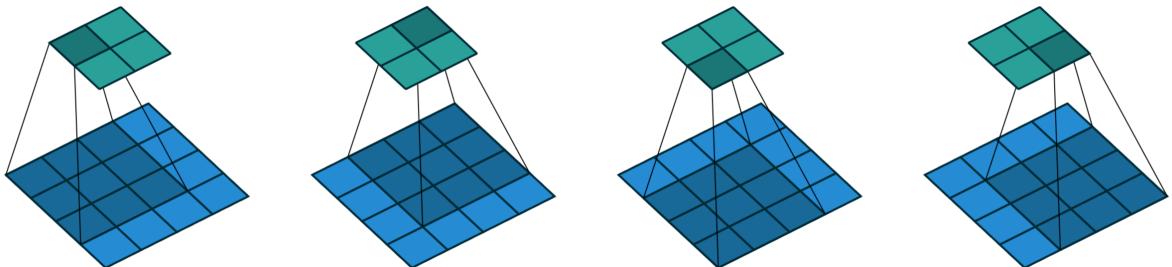


Figure 3.3: Convolution of a 4×4 input feature map (blue) with a 3×3 kernel (blue shaded). The resulting output feature map is shown in green. Taken from [83].

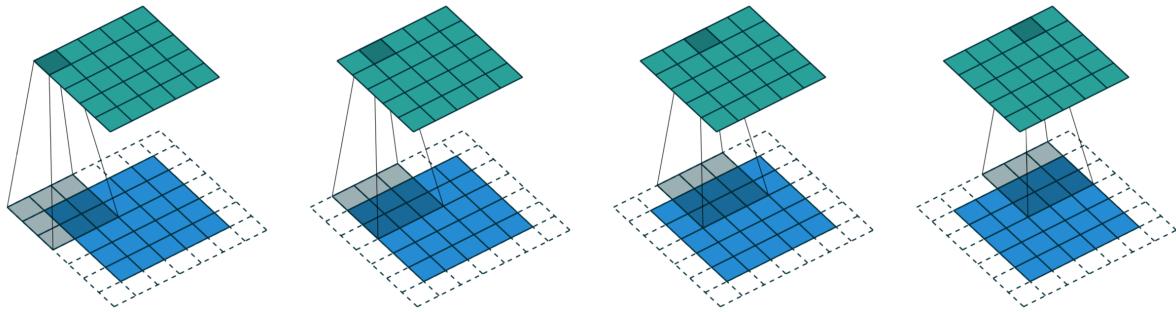


Figure 3.4: Convolution of a 5×5 input feature map (blue) with a 3×3 kernel (blue shaded). The input feature maps are padded with zeros (dashed) in order to maintain identical spatial dimensions of input and output feature maps. The resulting output feature map is shown in green. Taken from [83].

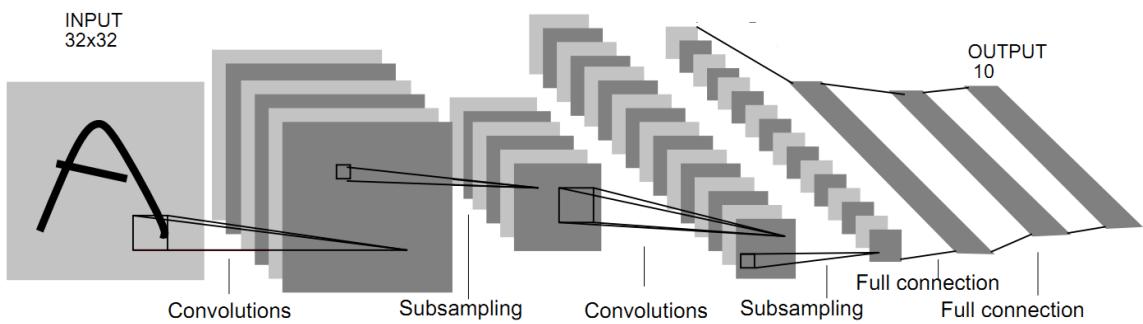


Figure 3.5: Illustration of a simple convolutional neural network. The input (left) is processed by a series of convolutional layers where each kernel in every layer produces a separate feature map. Subsampling reduces the spatial dimensions of the feature maps. The convolutional part is followed by a series of fully connected layers that lead to the final output layer (right). Modified from [85].

Normally, several convolutional layers are stacked hierarchically. Higher-level features can then be captured by combining lower-level features of previous layers. The output feature maps of the final convolutional layer are flattened to a one-dimensional vector which serves as input to a series of fully connected layers which are discussed in Section 3.2.1. These layers perform the final reasoning by combining all individual feature maps. This architecture is called convolutional neural network and an example is shown in Figure 3.5. To date, several variations and extensions of the basic convolutional layers exist to further boost the performance of this concept [77, 84].

In training, the weights which define the kernels in the convolutional layers are adjusted. The training approach is in principle identical to the one discussed for fully connected layers in Section 3.2.2. Besides enforcing explicit knowledge about local correlation in the input data, the convolutional layers also drastically reduce the amount of trainable weights. For example, in a fully connected layer, an image of 100×100 pixels would correspond to 10 000 units that are all connected to the subsequent layer with a distinct trainable weight. With a hidden layer of only 10 units, the number of trainable weights is on the order of 10^5 . On the other hand, a convolutional layer with 10 distinct kernels of 3×3 receptive fields would correspond to 100 trainable weights. In general, the number of trainable weights for convolutional layers is independent of the input size but is determined only by the kernel size and the number of kernels.

A common approach for designing a convolutional neural network is to increase the number of feature maps in subsequent layers. At the same time, the spatial dimensions of the feature maps are reduced. The motivation of this approach is to increase the complexity of the representation and to down-sample the dimensions gradually. This reduction is acceptable in most applications since the exact location of features is not crucial. Rather, the relation between higher-level features becomes important. One way to reduce the dimensions is to increase the stride s . For unit strides ($s = 1$), the kernel is shifted over the input feature map one step at a time. For larger strides ($s > 1$), the step size is increased so that some locations are skipped. For example, in Figures 3.3 and 3.4, a stride of $s = 2$ corresponds to skipping every second step. Another concept is pooling, where a non-linear function is applied to each feature map. For example, a 2×2 pooling layer together with a stride $s = 2$ evaluates the non-linear function in non-overlapping partitions of the feature maps. While there are several non-linear functions used in practice, the most common one is maximum pooling. This pooling operation takes the maximum value of each receptive field. Intuitively, this can be interpreted by keeping the most prominent response to a kernel and discarding the others in the receptive field of the 2×2 pooling layer. In this example, the spatial dimensions are reduced by 75 %. Using pooling layers, the performance of the network is invariant to small shifts and distortions which also improves its robustness [20].

Chapter 4

Energy Reconstruction with Deep Neural Networks

Contents

4.1	Goal and design of study	34
4.1.1	Motivation and training data preparation	34
4.1.2	Network architecture and training	37
4.2	Performance on Monte Carlo simulated data	39
4.2.1	Bias from a non-uniform training dataset	41
4.2.2	Comparison of DNN and traditional EXO-200 reconstruction	43
4.2.3	Mixed induction and collection signals	45
4.3	Performance on measured data	47
4.3.1	Combination of ionization and scintillation signatures	52
4.3.2	Background reduction in the ROI of $0\nu\beta\beta$	58
4.4	Summary	60

This chapter presents a new approach for the reconstruction of the ionization energy in EXO-200 with a deep neural network (DNN). After discussing the potential of this method, the components and decisions made for training a DNN are presented. The DNN is designed such that it meets the requirements specific to EXO-200 data analysis. The performance of the DNN based reconstruction is evaluated and compared to the traditional energy reconstruction EXO-200. This is done on Monte Carlo (MC) simulated and measured data from calibration sources.

Parts of the results presented in this chapter are published in “Deep Neural Networks for Energy and Position Reconstruction in EXO-200” in Journal of Instrumentation (JINST) [70] and were presented at the Neutrino Conference (2018) [86].

4.1 Goal and design of study

4.1.1 Motivation and training data preparation

A precise energy measurement is a key requirement for the identification of potential $0\nu\beta\beta$ decays. The energy is the only experimental measure to distinguish $2\nu\beta\beta$ from $0\nu\beta\beta$ decays. The signature of a $0\nu\beta\beta$ signal is an excess of events at the Q value in the energy spectrum. In contrast, the electrons emitted in $2\nu\beta\beta$ decays produce a continuous energy spectrum as discussed in Section 2.1. Due to the different spectral shapes, a moderate energy resolution would suffice to distinguish $2\nu\beta\beta$ from $0\nu\beta\beta$ events as is shown in Figure 2.3. However, other background which originates mostly from γ particle interactions also produce events near the Q value. The dominant components arise from long-lived radionuclides of the ^{238}U and the ^{232}Th decay chains. They would mask a potential observation of $0\nu\beta\beta$. To address this issue, precise energy measurements are needed to reduce the background contributions at the Q value. For a more detailed discussion of the background components in EXO-200 see Section 2.1.2.

As discussed in Section 2.2, ionizing radiation can deposit energy in the TPC. The produced secondary ionization electrons are drifted toward the wire grid via a homogeneous electric drift field that is applied between the wire and the cathode planes. The potentials of the V- and U-wire planes are set so that the V-wire plane is transparent to drifting charge. This charge produces an induction signal in the V-wire planes and is collected at the U-wire plane. Due to energy conservation, the ionization and scintillation signatures are anti-correlated (see Figure 2.10). The collected charge is delivered to the front-end electronics as current and is converted into a voltage signal by a multi-stage amplification and signal shaping scheme. Afterwards, it is converted into a 12-bit digital waveform by an analog-to-digital converter (ADC) sampled at a rate of 1 MHz. After being triggered, the data acquisition (DAQ) records 2048 samples from all channels. The samples are recorded symmetrically around the triggered sample. The triggering module handles different trigger schemes, e.g. from the APDs, wires, or external triggers. The U-wire waveforms of an example event are shown in Figure 4.1.

The energy deposit of primary particles is simulated with the GEANT4 software package [63, 64]. The detailed geometry of the EXO-200 detector is implemented in this simulation that is used to model the interaction of particles in the detector. Ionizing charge is drifted toward the wire planes and the resulting charge signals are determined using the Shockley-Ramo theorem [65, 66]. Raw waveforms of the wire plane channels are generated in the simulation by considering the readout electronics and by adding measured randomized noise waveforms. For more details see Section 2.2.3 and Ref. [51].

For training a DNN to reconstruct the ionization energy, several design choices were made. First, the training data is produced in the EXO-200 software framework for two classes of events, γ events and $0\nu\beta\beta$ events. The MC simulations takes the correlation of both primary electrons emitted in $0\nu\beta\beta$ decay into account in terms of energy and angular distributions [87]. The training set is populated with $0\nu\beta\beta$ and γ events to equal parts. 95 % of the dataset are for training and the remaining 5 % for validation. The energy region on which the DNN is trained determines the region it can be applied to later. Therefore, second, the energy of simulated primary particles is restricted to 1000–3000 keV based on the main energy region

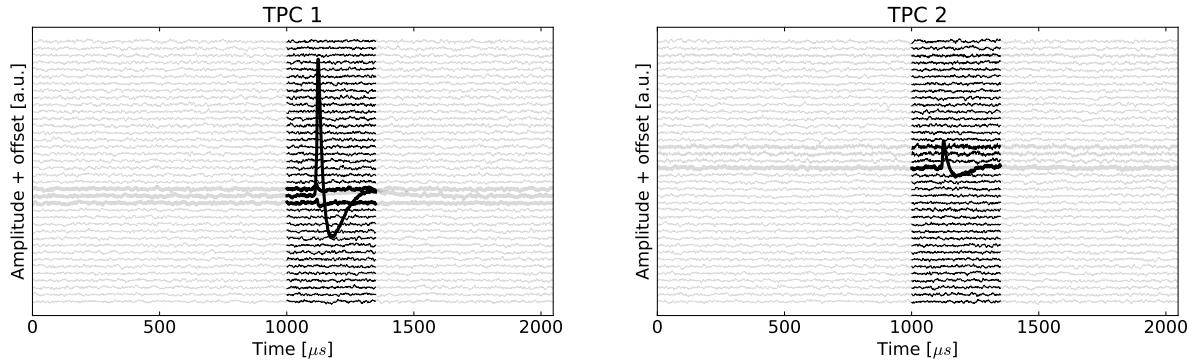


Figure 4.1: U-wire waveforms of an example event showing the channel amplitudes as a function of the time. The channels are offset for clarity. The waveforms are cropped in time (light area) before being used as inputs of the DNN to reduce computational costs. For clarity, the channels that contain any signal found by EXO-200 reconstruction are shown in bold.

that EXO-200 uses for $0\nu\beta\beta$ decay searches [14]. γ particles and $0\nu\beta\beta$ decays are simulated from a uniform distribution in energy and Q value in this range. Third, the spatial distribution of primary particles is simulated uniformly in the detector volume. The correlation of the spatial coordinates and the energy of the training dataset are shown in Figure 4.2 for both event classes. Instead of using uniformly distributed primary particles, spatial and energy distributions of realistic particle sources, like γ calibration sources, could be used for training. However, this could introduce biases in the DNN’s prediction as will be discussed in Section 4.2.1. Some event selection cuts are applied to the training data: Events are required to have at least one reconstructed charge cluster. Further, all clusters need to have a well-reconstructed 3D position as well as to be contained in a fiducial volume that is looser than that used for physics analyses. In particular, the event selection does not cut multi-site events (MS). More details on event selection cuts are discussed in Section 2.2.4.

The inputs for the DNN are grayscale images. They are built by arranging neighboring channels next to each other. The channel amplitudes of U-wire waveforms are encoded as pixel values. The images are cropped from 2048 samples to the samples in the region 1000–1350 in order to reduce computational effort. This is indicated also in Figure 4.1. In MC simulations, events are triggered by APD signals at sample 1024, so any signals on the U-wires are still fully contained in the cropped image. Preprocessing further includes subtraction of a constant baseline of each individual channel and correction for slightly different measured gains of the electronics of individual channels that is discussed in Section 2.2.4.

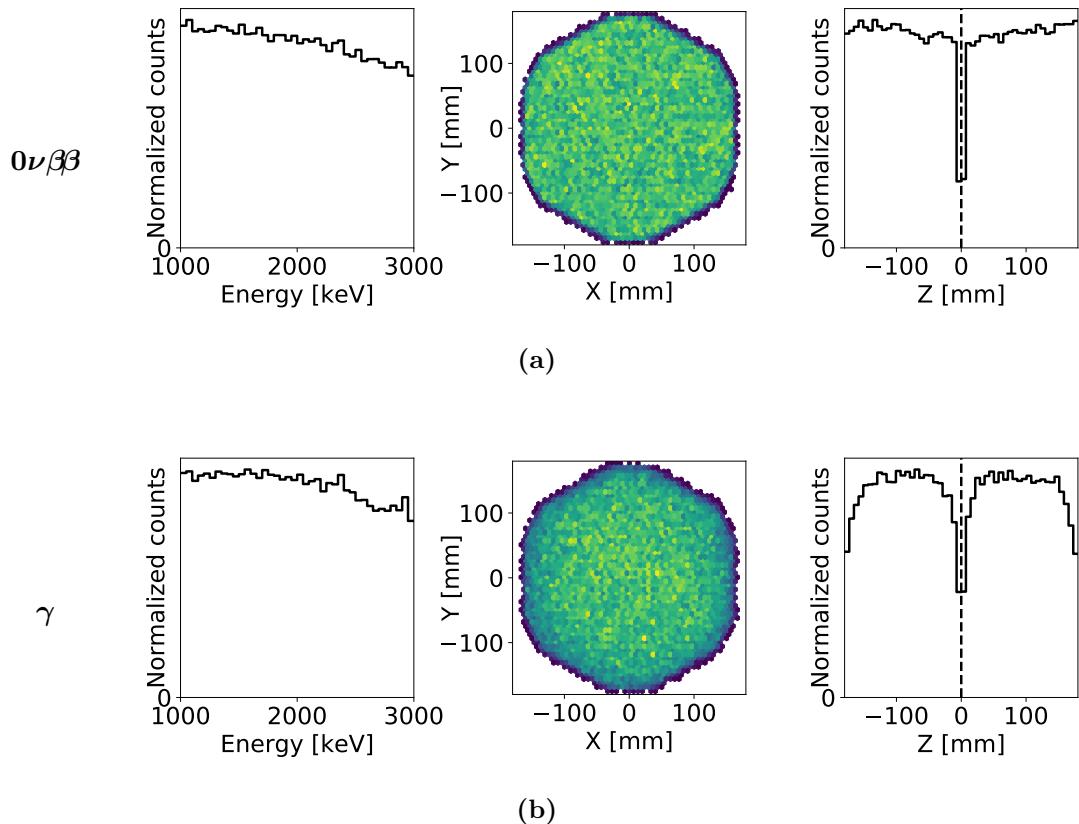


Figure 4.2: The distribution of deposited energy (left), the correlation of the x and y spatial coordinates (center), and the distribution of z coordinates (right) are shown for $0\nu\beta\beta$ events (a) and for γ events (b). The combined dataset is used in training and validation of the DNN. The primary particles are emitted uniformly in the detector volume in MC simulations. $0\nu\beta\beta$ decays have a random Q value and γ particles a random energy.

4.1.2 Network architecture and training

The two U-wire images of both sub-TPCs are fed to the DNN separately. Each image has dimensions of 350×38 pixels. The DNN architecture is shown in Figure 4.3. As the physics contained in both images is identical, they pass through the same layers with the same mathematical operations applied to them. This approach reduces the number of trainable weights. The architecture contains nine convolutional layers with 16, 16, 32, 32, 64, 64, 128, 128 and 256 kernels. The feature maps of the last convolutional layer are flattened to a one-dimensional vector. This vector serves as input to two successive fully connected layers with 32 and 8 units while the outputs of both network branches are added. The resulting output is then mapped directly to the final output unit that represents the energy of the event that is deposited in the TPC. The kernel sizes of the initial four convolutional layers are (5, 3), and (3, 3) for the remaining layers. To maintain the dimensions after the convolutional operation, zero padding is applied to each layer which surrounds the image with zeros. In order to reduce the dimension of the feature maps intentionally, the convolutional part includes five maximum pooling layers with pool sizes of (4, 1), (4, 2), (2, 2), (2, 2) and (2, 2). Throughout the network, $\text{ReLU}(x) = \max(x, 0)$ is used as non-linear activation function. Trainable weights are initialized randomly using the Glorot algorithm [80] introduced in Equation 3.14. The DNN architecture is implemented with the Keras library [88] using the Tensorflow backend [89].

During training, the DNN target variable in MC simulations consists of the summed energy that is deposited on the wire channels. Energy that is lost on inactivate parts of the detector, like the PTFE reflectors, is discarded as it cannot be reconstructed. The Adam optimizer [90] is used to minimize the mean squared error L_2 of predicted \hat{y}_i and target values y_i over a batch of n samples:

$$L_2 = \frac{1}{2n} \sum_i^n (y_i - \hat{y}_i)^2. \quad (4.1)$$

Every 10 epochs, the learning rate of the optimizer is reduced by a factor 2. An additional L_2 regularization term weighted with 0.01 is applied to each layer (see Section 3.2.2). As EXO-200 was operated in two phases with slightly different conditions (e.g. electric field), a DNN is trained for each phase individually to optimally exploit all features available in data. However, training data production, DNN architecture, and DNN training procedure are identical for both DNNs. The training curve of the DNN used for Phase-II is shown in Figure 4.4.

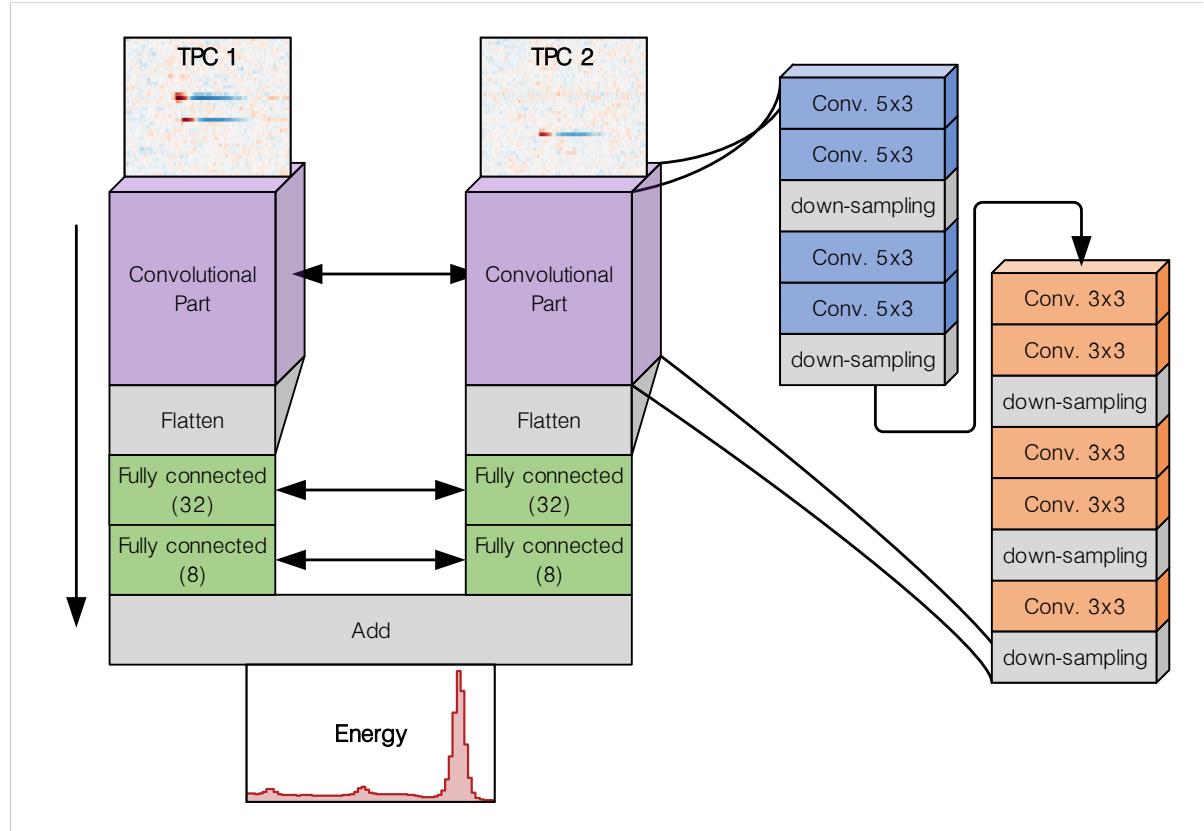


Figure 4.3: The network architecture used for the DNN. The input consists of two grayscale images with dimensions of 350×38 pixels. After the convolutional part, the final feature maps are transformed into a one-dimensional array and are fed to fully connected layers. Both network branches which contain information from a single TPC share the same network weights. The units of both branches are added before the final output layer. The output layer has a single unit corresponding to the total deposited energy in that event.

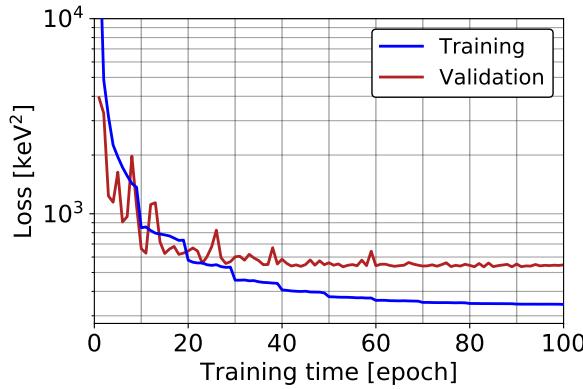


Figure 4.4: The performance measure L_2 during training (blue) as a function of the epoch of the DNN used for Phase-II data. Validation (red) is done with an independent part of the dataset to monitor potential overfitting. No overfitting is observed here. The reduction of the learning rate every 10 epochs results in a step-like behaviour in the curves that is best visible for the training curve.

4.2 Performance on Monte Carlo simulated data

The performance of the trained DNN is tested on the exposure of the EXO-200 detector with γ particles emitted from a radioactive ^{228}Th calibration source in MC simulations. If not stated otherwise, throughout this chapter, all calibration data is MC simulated or recorded from a calibration source located at source position S5 (see Figure 2.13b). Similar to training data, this simulation does not include xenon impurities, i.e. an infinite electron lifetime of secondary electrons is assumed. The ^{228}Th calibration source is suited for evaluation of the energy reconstruction because it produces a prominent peak at 2615 keV. This peak corresponds to the full absorption of the γ particle emitted from the ^{208}Tl decay. This energy is of particular interest for EXO-200 reconstruction and calibration efforts due to its proximity to the Q value of ^{136}Xe at 2458 keV. The correlation of the predicted energy of the DNN and the true MC energy is shown in Figure 4.5. This shows a good agreement between the DNN energy reconstruction and the target variable. However, a few outlier events exist, especially below the diagonal line where the DNN predicts lower energy compared to the true energy deposited in the detector. This corresponds to events, where the DNN misses some energy. Two dominant effects contribute to this. First, these events may have low energy deposits that are below the reconstruction threshold. Second, these events may contain complex mixtures of collection and induction signals where the DNN partially fails to disentangle both signal types and thus cannot predict the deposited energy correctly. The latter events will be discussed in more detail in Section 4.2.3. Above the diagonal line, there is no significant accumulation of events. For these events the DNN would predict more energy than there was actually deposited in the detector. The residuals of the DNN prediction to the true MC energy do not show any energy-dependent features. Due to electronics noise and limitations of the DNN reconstruction, the residuals are broadened.

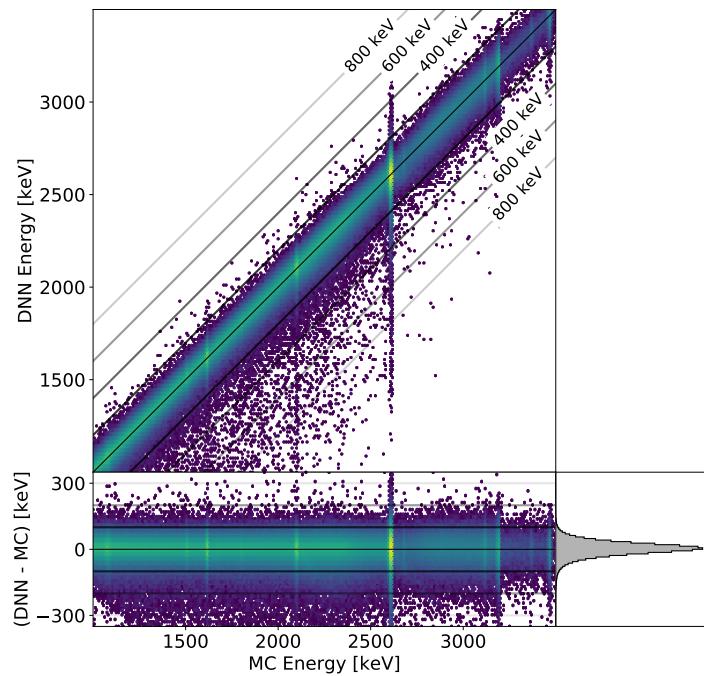


Figure 4.5: Predicted energy of the DNN as a function of the true MC energy deposited in the detector from a simulated ^{228}Th calibration source. The lower left panel shows the residuals. The residuals essentially follow a normal distribution shown in the right panel on linear scale. The color shows the intensity on a logarithmic scale. To guide the eye, solid lines indicate certain residual values in the left panels.

4.2.1 Bias from a non-uniform training dataset

In addition to the training approach discussed in Section 4.1.1, here, the effect of a biased training dataset on the DNN performance is discussed. For this approach, events from a ^{228}Th calibration source in MC simulations are used to train the DNN. The corresponding energy and spatial distributions are shown in Figure 4.6. Opposed to those of the baseline training dataset (see Figure 4.2), the spatial distribution is biased towards large x coordinates and the energy spectrum has distinct peaks. The most prominent one is the full absorption peak of ^{208}Tl at 2615 keV.

Validating a DNN trained on ^{228}Th MC simulated data yields promising results at the ^{208}Tl full absorption peak. However, Figure 4.7 shows this DNN and another DNN trained on data with a uniform energy distribution. Both DNNs are then used to reconstruct MC simulated events of the radioactive ^{228}Th calibration source, but with an MC energy spectrum that is broadened by an artificial energy resolution of 90 keV. This resolution is guided by the charge-only energy resolution of $\sigma/E \approx 3.5\%$ observed at the ^{208}Tl full absorption peak in Phase-I of EXO-200 operation. Comparing both spectra, the DNN trained on ^{228}Th data shows undesired effects: The predicted energy of events near the ^{208}Tl peak is shifted toward the peak as is shown in the lower panel of Figure 4.7. This can be understood by the large excess of ^{208}Tl peak events in the training spectrum of ^{228}Th compared to the peak's vicinity (see Figure 4.6). With this spectral feature, the average loss contribution in training introduces a bias that favors events near the peak. Thereby, the resolution seems to be improved because the peak is narrowed artificially. The DNN trained on a uniform energy spectrum is not biased toward certain energies and guarantees a valid energy reconstruction over the entire energy range without introducing systematic biases.

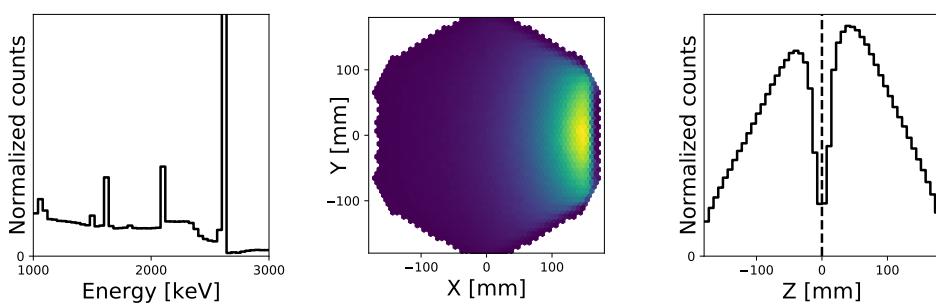


Figure 4.6: The distribution of deposited energy (left), the correlation of the x and y spatial coordinates (center), and the distribution of z coordinates (right). The ^{228}Th calibration source in MC simulations is located near the cathode outside the TPC.

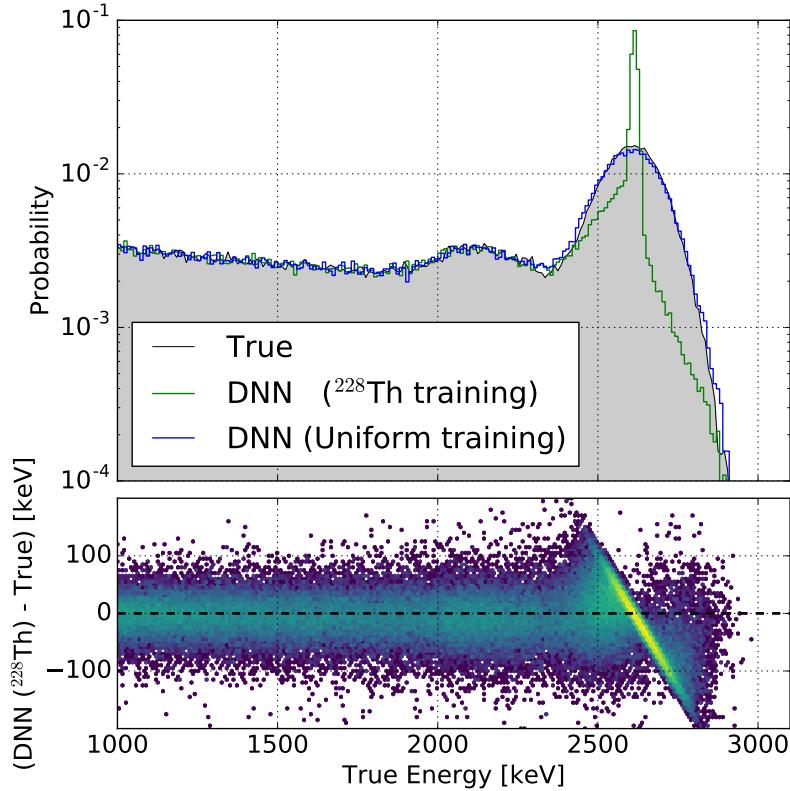


Figure 4.7: The MC energy spectrum (gray) as reconstructed by two DNNs that were trained on different datasets (top). The MC simulated events are generated from a ^{228}Th calibration source whose energy spectrum is broadened by 90 keV. One DNN was trained on the same kind of MC simulated events but without broadening the MC energy (green) while the other DNN was trained on a uniform energy spectrum (blue). The residuals to the true MC energy are shown in the lower panel the the DNN trained on MC simulated ^{228}Th events. The color denotes the intensity on a logarithmic scale. Taken from [70].

4.2.2 Comparison of DNN and traditional EXO-200 reconstruction

Figure 4.8 shows the reconstructed energy spectra of a ^{228}Th calibration source in MC simulations for single-site (SS) and multi-site (MS) events separately. For the ^{228}Th spectra, the spectral features, e.g. the ^{208}Tl full absorption peak at 2615 keV, the corresponding single and double escape peaks are reproduced by both the DNN and the traditional EXO-200 reconstruction methods. To compare both reconstruction methods, the energy resolution at the ^{208}Tl full absorption peak is evaluated by fitting the peak region with the function:

$$\begin{aligned} f(E) &= S(E, A_1, \mu, \sigma) + B(E, A_2, \mu, \sigma) \\ &= A_1 \cdot \exp\left(-\frac{(E - \mu)^2}{2\sigma^2}\right) + A_2 \cdot \left(1 - \text{erf}\left(\frac{E - \mu}{\sqrt{2}\sigma}\right)\right). \end{aligned} \quad (4.2)$$

In this simplified approach, the peak is modeled by a Gaussian distribution and the background contribution near the peak by a Gaussian error function. Both functions share the mean μ and standard deviation σ parameters. The resulting energy resolution values are shown in Table 4.1. The energy resolution of the DNN outperforms that of the traditional EXO-200 reconstruction in both SS and MS events. A major contribution to the resolution is the electronics noise that is added to the simulated waveforms. Besides energy resolution, the agreement in the region between the Compton shoulder and the full absorption peak is significantly improved by the DNN (see Figure 4.8a). This effect will be discussed in more detail in Section 4.2.3. Similarly, for simulated $0\nu\beta\beta$ decays, the DNN shows a superior energy resolution of $(0.65 \pm 0.01)\%$ for SS events.

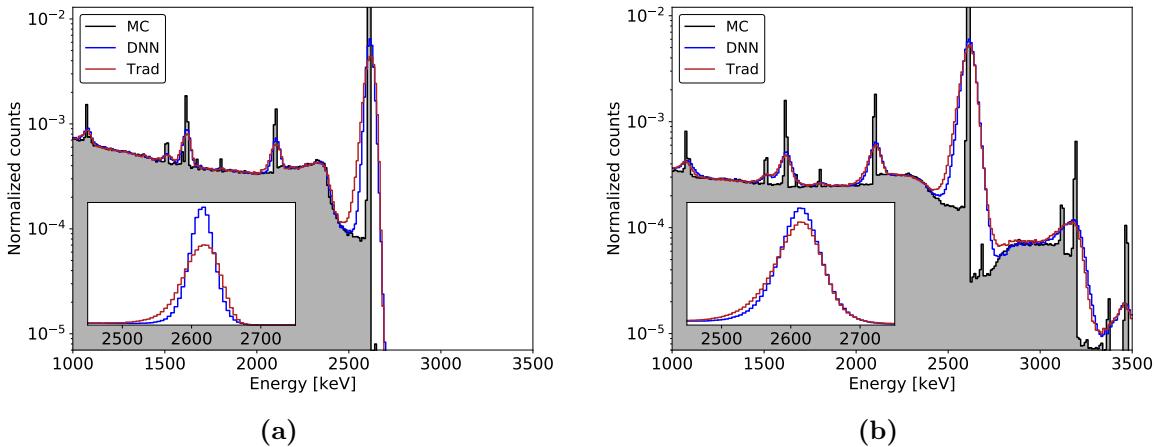


Figure 4.8: Energy spectra reconstructed from MC simulated events of a radioactive ^{228}Th calibration source. (a) shows SS events and (b) shows MS events. The spectra of energy deposited in the detector is known in MC simulations is shown in shaded gray. The spectra are reconstructed by the DNN (blue) and the traditional EXO-200 reconstruction (red). The insets show a zoom to the ^{208}Tl full absorption peak at 2615 keV on a linear scale.

Table 4.1: Energy resolution σ/E observed at the ^{208}Tl γ line at 2615 keV for a ^{228}Th calibration source in MC simulations. The resolution is determined for the DNN and the traditional EXO-200 reconstruction in all events, and in single-site (SS) and multi-site (MS) events separately.

Event type	Energy resolution σ/E [%]		Relative improvement [%]
	DNN	Trad.	
all	1.07 ± 0.03	1.19 ± 0.02	10.1 ± 0.4
SS	0.70 ± 0.01	0.96 ± 0.03	27.1 ± 1.0
MS	1.17 ± 0.02	1.24 ± 0.02	5.6 ± 0.2

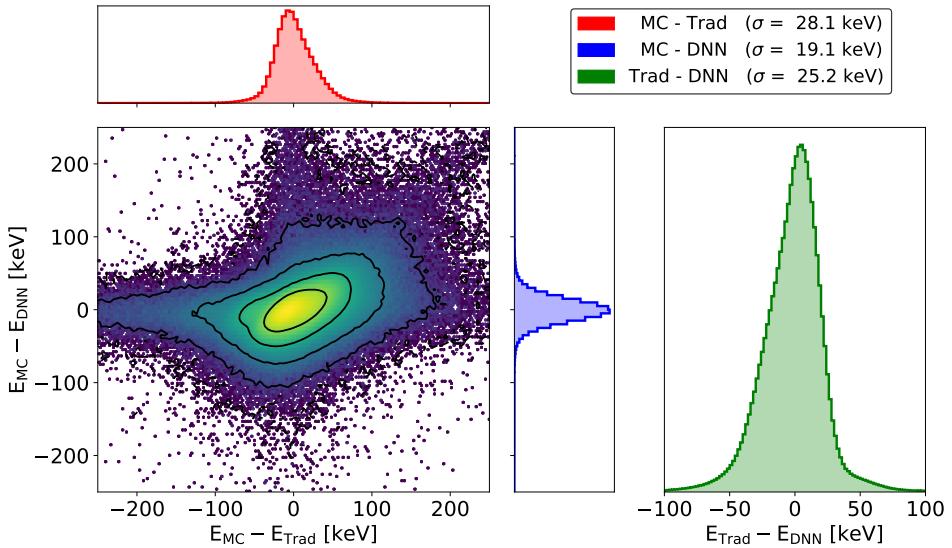


Figure 4.9: The correlation of residuals of E_{Trad} and E_{DNN} with respect to the true MC energy deposited in the detector E_{MC} for simulated $0\nu\beta\beta$ events. The intensity is denoted by color on a logarithmic scale. The projections are shown for both axes (red, blue) and to the diagonal line (green) on a linear scale. Adapted from [70].

The correlation between the residuals of both reconstruction methods with respect to the true energy deposited in the detector for MC simulated $0\nu\beta\beta$ events is shown in Figure 4.9. The tilted contour lines represent a positive correlation of these residuals. This correlation indicates an effect from the electronics noise on the waveforms. This noise affects the energy estimation of both methods in a similar way which results in a positive correlation of their residuals.

The energy resolution σ/Q of DNN and traditional EXO-200 reconstruction for simulated $0\nu\beta\beta$ SS events is shown in Figure 4.10 as a function of the event position in the radial plane R and the height Z together with their projections. The true position of the energy deposit is known in MC simulations. The detector is split into equally sized slices in Z and into slices in R such that all segments contain the same volume. Besides the improvement in energy resolution

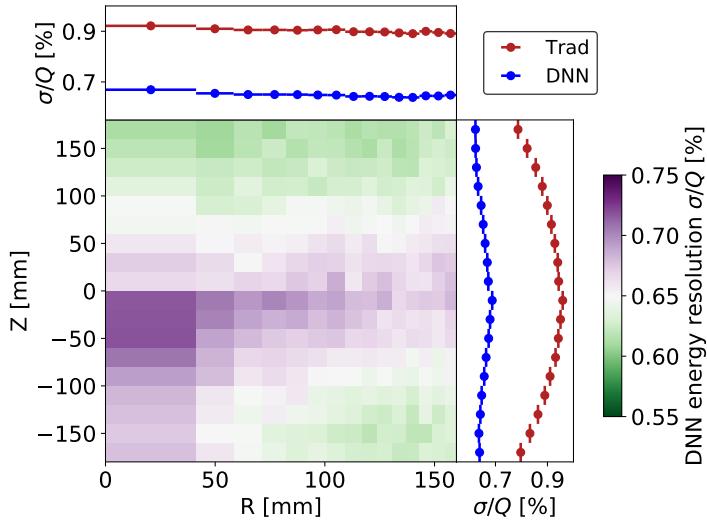


Figure 4.10: The center panel shows the energy resolution for simulated $0\nu\beta\beta$ SS events as a function of the radius R and the height Z reconstructed by the DNN. The top and right panels show the projections for the DNN (blue) and, in addition, for the traditional EXO-200 (red) reconstruction.

with the DNN reconstruction, its variation over R is on the same level compared to that of the traditional EXO-200 reconstruction. In Z , both methods show a trend toward worse resolution in direction to the cathode. This trend is understood to arise from an increasing extent of the drifting charge cloud due to diffusion for events with longer drift paths. Thereby, the fraction of events that deposit energy on multiple rather than on single U-wires increases. With multiple active wires, the electronics noise contributions add up which leads to a worse energy resolution of events near the cathode. The asymmetry of the energy resolution in Z with respect to the cathode at $Z = 0$ is not understood thoroughly but is observed by both reconstruction methods. This suggests that it is an intrinsic feature of the EXO-200 TPC design rather than a drawback of the DNN energy measurement.

4.2.3 Mixed induction and collection signals

Besides the improved energy resolution on MC simulated events, the DNN shows an improved agreement in the region between the Compton shoulder and the ^{208}Tl full absorption peak as is shown in Figure 4.8a. This region is of particular interest to EXO-200 as it contains the Q value of ^{136}Xe . The improved agreement is understood to arise from two main processes. First, the DNN is able to reconstruct smaller energy deposits that are below the reconstruction threshold of the conventional method. Second, the DNN can better disentangle complex signatures of collection signals that are superimposed by induction signals [70]. Figure 4.11 shows a raw waveform that is fitted by a template function for either a collection, induction, or mixed signal. Until the $0\nu\beta\beta$ search in 2018 [91], EXO-200 reconstruction applied template fits to raw data for either collected or induced charge signals. Based on the comparison of χ^2 of both fits,

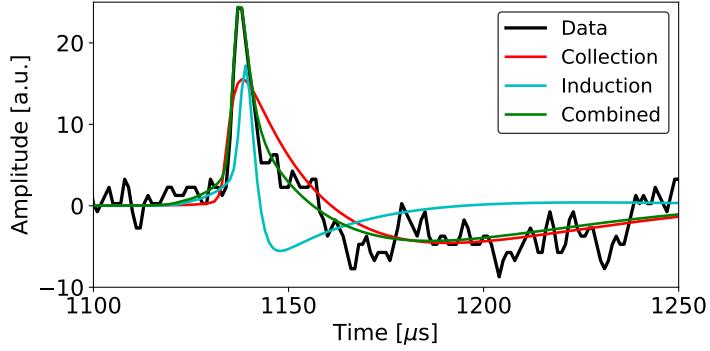


Figure 4.11: Example signal on U-wires. The graph shows a template fit from either collected (red) or induced (blue) charge only. The improved reconstruction applies a template fit of combined signals (green) which best matches the signal.

the scenario with the lower performance was discarded. If flagged as a pure induction signal, the reconstructed energy is set to $E_{\text{Trad}} = 0$. This binary nature in the event reconstruction process led to missing energy. Based on the results in Ref. [70], EXO-200 reconstruction was improved by applying template fits of both signal types to allow for entangled signals of collected and induced charge in wires. This reconstruction was used in the most recent $0\nu\beta\beta$ search in 2019 [14] and recovered energy that was previously missed for mixed signals. Figure 4.12 shows the impact on MC simulated events of the ^{208}Tl full absorption peak at $E = 2615\text{ keV}$. The traditional EXO-200 reconstruction with (red) and without (orange) improved template fit is shown. The proper modeling of mixed signals recovers some energy that was previously flagged as induction and thus missed (gray). However, the DNN (blue) still outperforms the improved reconstruction.

Interactions that originate from the ^{232}Th decay chain represent a dominant background in low-background (LB) physics data. This background component presents the ^{208}Tl full absorption peak at $E = 2615\text{ keV}$ which is right above the Q value of ^{136}Xe . Missing a part of the deposited energy of these events in reconstruction represents a direct risk of increasing the background in the ROI around Q . Furthermore, since this effect is not modeled in the MC simulations, it worsens the spectral agreement between measured data and MC simulations. The impact of this effect on the reconstruction of measured data is discussed in Section 4.3.2.

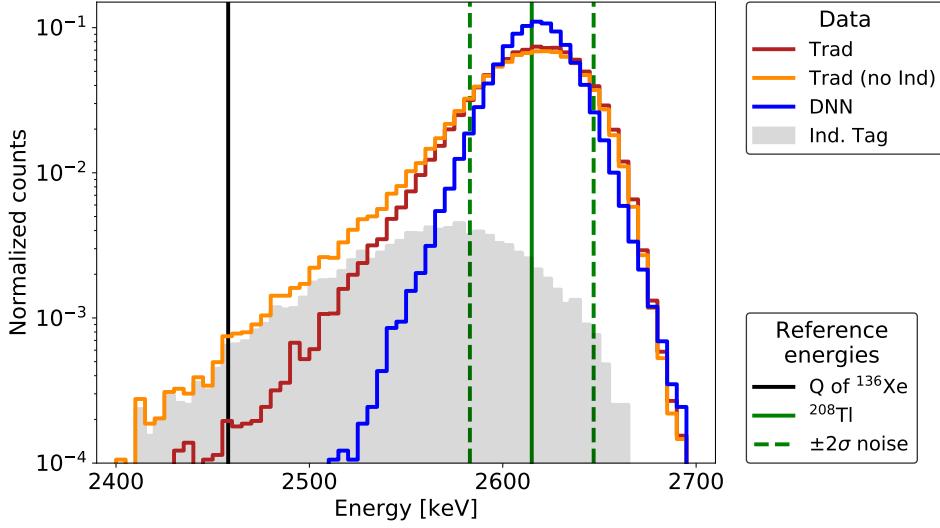


Figure 4.12: Energy spectra of MC simulated ^{208}Tl full absorption peak events which deposit 2615 keV (green solid) in the detector volume. The spectra reconstructed by the DNN (blue) and by the traditional EXO-200 reconstruction are shown with (red) and without (orange) improved mixed signal modeling. For reference, the expected broadening due to noise only ($\sigma \simeq 16$ keV [70], green dashed) and the Q value of ^{136}Xe (black) are shown. The distribution of events labelled as induction signals in EXO-200 reconstruction are shaded gray.

4.3 Performance on measured data

In order to apply the DNN on measured data, two aspects need to be taken into account. First, the liquid xenon in the detector contains electronegative impurities which lead to a finite electron lifetime of secondary electrons:

$$N_e(t) = N_{e,0} \cdot \exp\left(-\frac{t}{\tau_e}\right), \quad (4.3)$$

where $N_{e,0}$ is the initial number of secondary electrons, τ_e the electron lifetime in LXe, and t the drift time. The electron lifetime τ_e is measured regularly with source calibration data. To correct for the finite electron lifetime, all ionization signals are multiplied by a factor of $\exp(t/\tau_e)$ in EXO-200 reconstruction [51]. The time-average electron lifetime is about 4.5 ms (3.0 ms) for Phase-I (Phase-II), leading to a correction factor of 2.6 % (3.6 %) for maximum drift lengths from the cathode to the wires. Since the DNN does not predict the energy of individual ionization signals but the sum energy over all ionization signals that belong to a scintillation signal, this correction was adjusted. For events with single ionization signals, the same correction like in the traditional EXO-200 reconstruction was used. For events with multiple ionization signals, the correction factor is calculated from the most energetic cluster found by the traditional EXO-200 reconstruction. This approximation is necessary because the DNN predicts the total energy deposited in an event in contrast to the energy of individual charge clusters.

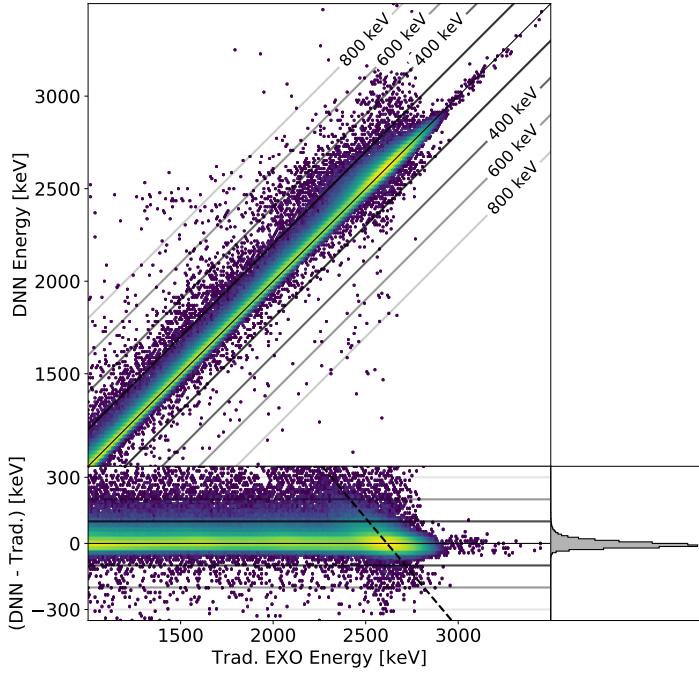


Figure 4.13: The energy predicted by the DNN as a function of the energy from the traditional EXO-200 reconstruction. The data represents measured SS events from a ^{228}Th calibration source. The lower left panel shows the residuals. The color denotes the intensity on a logarithmic scale. To guide the eye, solid lines indicate certain residual values in the left panels. The projection of residuals is shown in the right panel on a linear scale. The excess in positive residuals is best visible in the lower left panel for ^{208}Tl full absorption peak events. The dashed line represents $E_{\text{DNN}} = 2615 \text{ keV}$.

Second, in measured data there are several trigger strategies. For high energies, nearly all events are triggered by scintillation channels, while for lower energies, a significant fraction of events is triggered by ionization channels. Depending on the triggered channel, the signals on the U-wires used as input for the DNN are shifted compared to those in MC simulations. To account for this, the region of the selected window used by the DNN is shifted to ensure a fixed position relative to the scintillation signal. If necessary, the cropped waveforms are padded with zeros.

The correlation of E_{Trad} and E_{DNN} for measured data of a ^{228}Th calibration source is shown in Figure 4.13. As there is no true energy label in measured data, the reconstruction error where both methods show a large discrepancy cannot be assigned to either reconstruction method unambiguously. However, in general a good agreement on an event basis is observed as most events are distributed around the diagonal line. The residuals of both methods are unbiased over the entire energy range, i.e. are centered at zero. However, they show an asymmetric excess of events. The excess events are due to the traditional EXO-200 reconstruction missing some energy while the DNN reconstruction mostly recovers that energy as was discussed previously for MC simulated events.

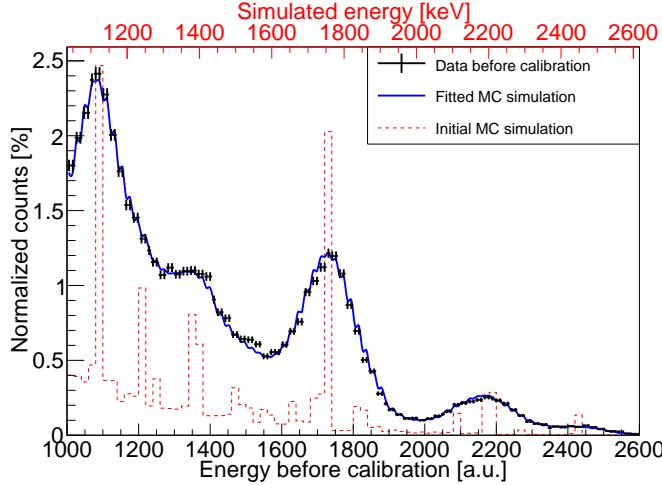


Figure 4.14: Example of a MC based fit (blue) to charge-only ^{226}Ra calibration data (black) recorded in Phase-II for SS events. The spectrum predicted in MC simulations is shown in red. The energy calibration and resolution functions are fitted simultaneously.

To compare the DNN and the traditional EXO-200 energy reconstruction quantitatively, the baseline EXO-200 calibration procedure is pursued. The energy scale is calibrated on a weekly basis from ^{228}Th source calibration data for SS and MS events separately. The linear calibration function is

$$E(E_0) = p_{0E} + p_{1E}E_0 , \quad (4.4)$$

where E_0 is the uncalibrated energy. The energy calibration $E(E_0)$ and the energy resolution $\sigma(E)$ functions are fitted simultaneously in order to smear the MC simulated spectra by:

$$\sigma^2(E) = \sigma_{\text{elec}}^2 + p_{1R}^2 E . \quad (4.5)$$

σ_{elec} denotes the electronic noise contribution and $p_{1R}^2 E$ represents statistical fluctuations. An example fit to ^{226}Ra data is shown in Figure 4.14.

The weekly energy resolution σ/E observed at the ^{208}Tl full absorption peak in SS events is shown in Figure 4.15. The energy resolution improves in Phase-II due to improvements to electronics and due to raising the cathode voltage from -8 kV to -12 kV . Comparing the weekly energy resolution of the DNN and the traditional EXO-200 reconstruction, the DNN outperforms it in almost every week. Without weighting the weeks by livetime, the average DNN energy resolution is 3.40 % (2.97 %) for Phase-I (Phase-II) whereas that of the traditional EXO-200 reconstruction is 3.46 % (3.06 %). The relative improvement of the DNN over the traditional EXO-200 reconstruction is 1.7 % (2.9 %) for Phase-I (Phase-II). The variation over time of both reconstruction methods is on a similar level and is dominated by variations in the conditions from operating the detector.

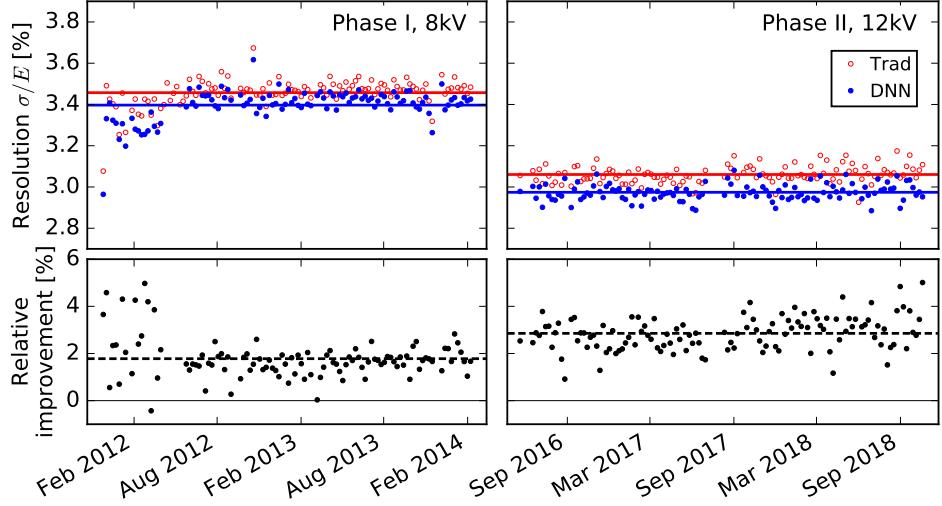


Figure 4.15: The weekly charge-only energy resolution at the ^{208}Tl full absorption peak for SS events is shown in the top panels. The left panels show Phase-I of EXO-200 operation and the right ones Phase-II. The cathode voltage is -8 kV (-12 kV) in Phase-I (Phase-II). The bottom panels show the relative improvement of the DNN over the traditional EXO-200 reconstruction. Horizontal lines in all panels indicate the mean values.

The weekly calibration of the energy scale is merged into an average calibration by smearing the MC simulated spectra by the weekly resolution parameters, weighted by the livetime in each week. Measured data from different calibration sources located at different source positions is combined to reflect positions near the cathode (S5, S11) and near the anodes (S2, S8) (see Figure 2.13b). Both classes are then weighted uniformly in the average calibration to account for spatial differences. To account for time and position-dependent variations, the calibration and resolution functions used for the average calibration are extended by a quadratic term [51]:

$$E(E_0) = p_{0E} + p_{1E}E_0 + p_{2E}E_0^2 \quad (4.6)$$

$$\sigma^2(E) = \sigma_{\text{elec}}^2 + p_{1R}^2E + p_{2R}^2E^2. \quad (4.7)$$

The average charge-only energy resolution of the DNN for SS events is $\sigma/Q = 3.33\%$ (2.85%) for Phase-I (Phase-II). Compared to the resolution of the traditional EXO-200 reconstruction, this represents an improvement of 1.5% (3.4%).

The calibrated SS charge-only energy spectra of ^{60}Co , ^{226}Ra , and ^{228}Th calibration sources in Phase-II of EXO-200 operation are shown in Figure 4.16 (top panel). The spectra by the DNN and the traditional EXO-200 (not shown) reconstruction agree well. Measured data and MC simulations show overall good spectral agreement within 20 % for both reconstruction methods (bottom panel). Below 1100 keV and in the region between the ^{208}Tl full absorption peak and its Compton edge, the simulation shows slightly worse agreement in both reconstruction methods. The second discrepancy is well known in EXO-200 data analysis. The DNN shows slightly improved agreement for ^{226}Ra source calibration data compared to the traditional EXO-200 reconstruction. The overall very similar spectral agreement between measured data

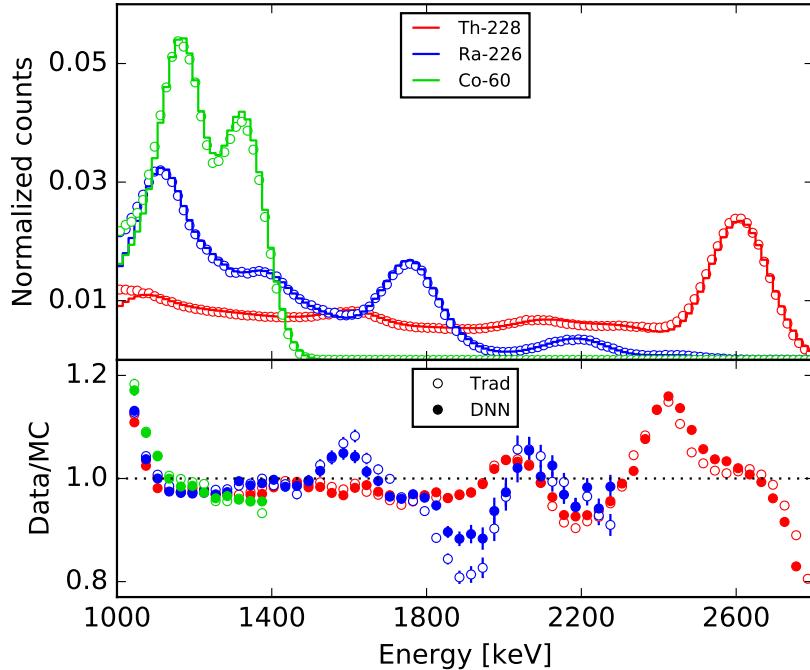


Figure 4.16: The top panel shows the calibrated charge-only energy SS spectra of calibration sources (^{60}Co , ^{226}Ra , ^{228}Th) reconstructed by the DNN in measured data (circles) and MC simulations (lines). The bottom panel shows the ratios between measured data and MC simulations for the DNN (solid) and the traditional EXO-200 (open) reconstruction.

and MC simulations for DNN and the traditional EXO-200 reconstruction confirms an unbiased ionization energy measurement of the DNN on data.

The DNN reconstruction in terms of spatial detector uniformity as a function of the event position is shown in Figure 4.17. It is shown as a function of the radial plane R and the height Z for ^{228}Th source calibration SS data. Additionally, the projections are shown. The event position is determined in EXO-200 reconstruction. The detector is split into slices in Z and R so that all segments contain the same volume. The energy resolution σ/E at the ^{208}Tl full absorption peak at 2615 keV shows a trend toward better resolution near the walls. In Z , the resolution is mostly flat over the detector. Overall, the variations observed over the detector volume are on the same level compared to the traditional EXO-200 reconstruction. Compared to the trend in Z that is observed in MC simulations, the effect of diffusion is subdominant in data. This is due to the anti-correlated fluctuation of ionization and scintillation signals which smears the energy measurement of both signatures. Improving the energy measurement by combining both channels will be discussed in the next section.

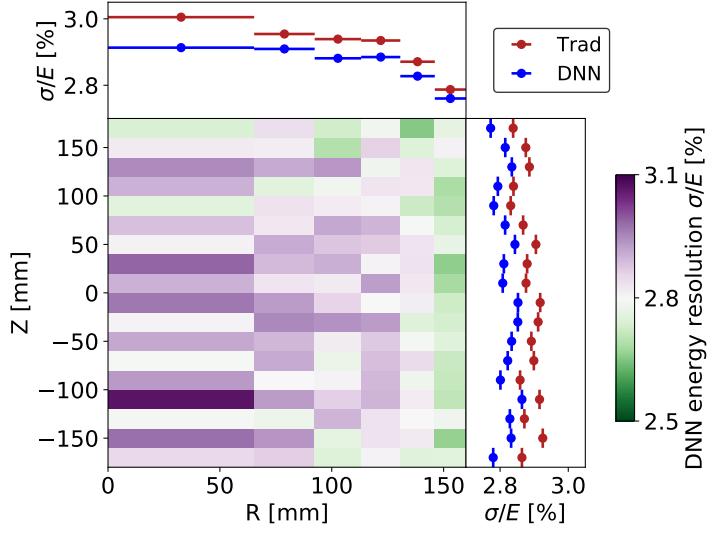


Figure 4.17: The energy resolution σ/E at the ^{208}Tl full absorption peak from calibration SS data recorded with a ^{228}Th source is shown as a function of the radius R and the height Z in the center panel reconstructed by the DNN. The top and right panels show the projections for the DNN (blue) and, in addition, for the traditional EXO-200 (red) reconstruction. As opposed to MC simulations, the true event position is unknown in measured data. Instead, the event position as reconstructed by EXO-200 is used.

4.3.1 Combination of ionization and scintillation signatures

The signatures of ionization charge and scintillation light fluctuate due to a varying fraction of recombination of secondary electrons in the liquid xenon. This effect is discussed in Section 2.2.2. By measuring both signatures and combining them, these fluctuations cancel and the energy measurement becomes more precise compared to those of the individual channels. The correlation in both channels is shown in Figure 4.18. The optimum energy variable is calculated by joining both channels to a “combined energy”

$$E_{\text{comb}} = E_{\text{I}} \cdot \cos \theta + E_{\text{S}} \cdot \sin \theta , \quad (4.8)$$

where E_{I} is the energy measured in the ionization channel and E_{S} is the scintillation channel. The angle θ is used to combine the two energy measurements E_{I} and E_{S} . The angle is chosen to minimize the energy resolution at the Q value.

The optimum angle θ is identified on a weekly basis from ^{228}Th source calibration data in a three-step scan. An initial coarse scan over θ is followed by two finer scans near the found optimum angle from the previous step. For each angle, a MC based fit to the combined energy spectrum determines the energy resolution at Q . Figure 4.19 shows an example MC based fit to the energy spectrum after combination by the corresponding optimum angle. This energy spectrum resolves peaks considerably better compared to the charge-only energy spectrum of corresponding data shown in Figure 4.14. For each scan over different angles, the energy

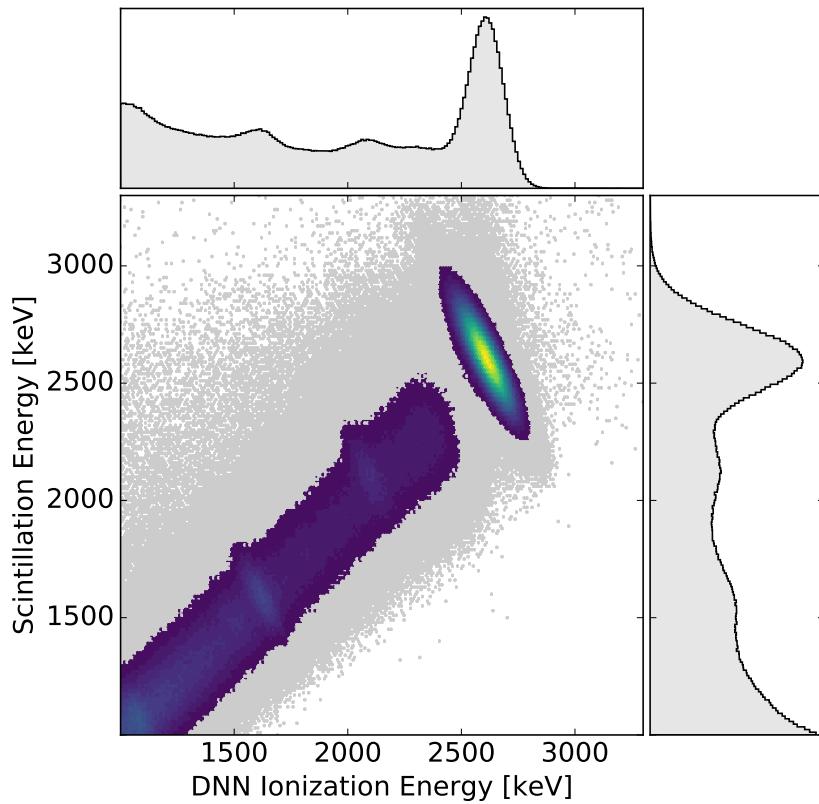


Figure 4.18: The center panel shows the correlation of scintillation energy and DNN ionization energy for ^{228}Th calibration data. The color denotes the intensity and is set to gray for low intensities for clarity. The top and right panels show the projections to both axes. By optimally combining both signatures, the anti-correlated fluctuations cancel and the energy resolution is improved significantly which is best visible at the ^{208}Tl full absorption peak at 2615 keV.

resolution at Q as a function of the angle is fitted by a quadratic function. Figure 4.20 shows the procedure for determining the optimum angle θ for an example week.

After determining the angles that optimally combine ionization and scintillation signals, the subsequent calibration procedure is identical to the one described in Section 4.3 for ionization signals. A linear calibration is applied on a weekly basis from ^{228}Th source calibration data. The weekly energy resolution σ/E observed at the ^{208}Tl full absorption peak in SS events measured with a ^{228}Th calibration source is shown in Figure 4.21. The improvement in energy resolution in Phase-II is due to improvements to electronics and to the raised electric field strength. The stability in energy resolution in Phase-II is improved by upgraded frontend electronics of the APDs that reduce the readout excess noise which caused the fluctuations in Phase-I. The weekly energy resolution of the DNN reconstruction outperforms that of the traditional EXO-200 reconstruction in almost every week. Without weighting the weeks by livetime, the average DNN energy resolution is 1.33 % (1.05 %) for Phase-I (Phase-II) whereas that of the traditional EXO-200 reconstruction is 1.37 % (1.12 %). This represents a relative improvement

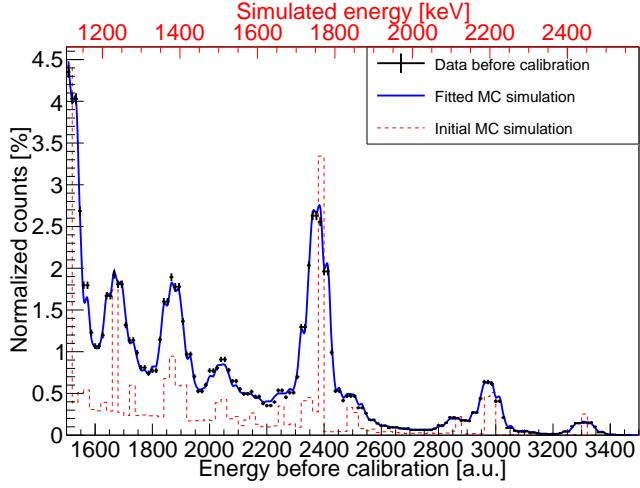


Figure 4.19: Example of a MC based fit (blue) to ^{226}Ra calibration data (black) recorded in Phase-II for SS events. The energy is calculated by combining the ionization and scintillation measurements. The spectrum predicted in MC simulations is shown in red. The energy calibration and resolution functions are fitted simultaneously.

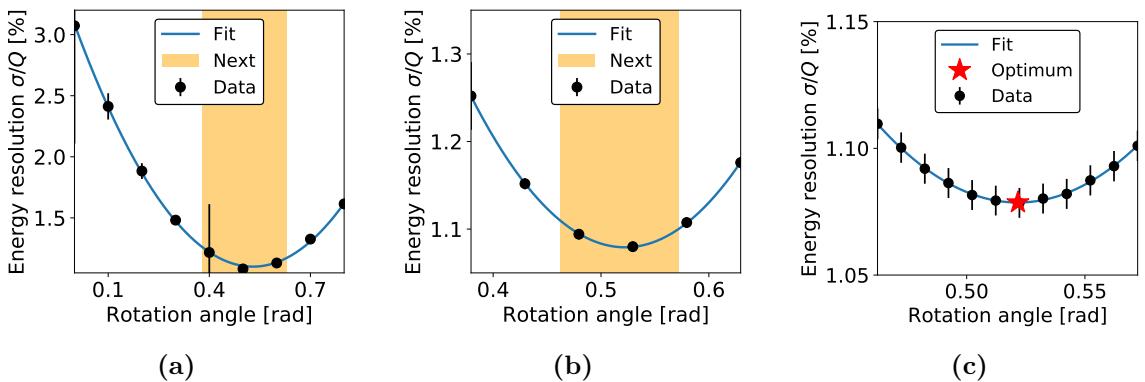


Figure 4.20: Combined energy resolution σ/Q as a function of the angle θ . After a coarse scan (a), the angle is scanned in the next step (b) near the optimum angle (yellow) determined by fitting a quadratic function (blue) to data (black). Finally, the best-fit angle (red star) of the fine scan (c) is used to combine ionization and scintillation energy measurements.

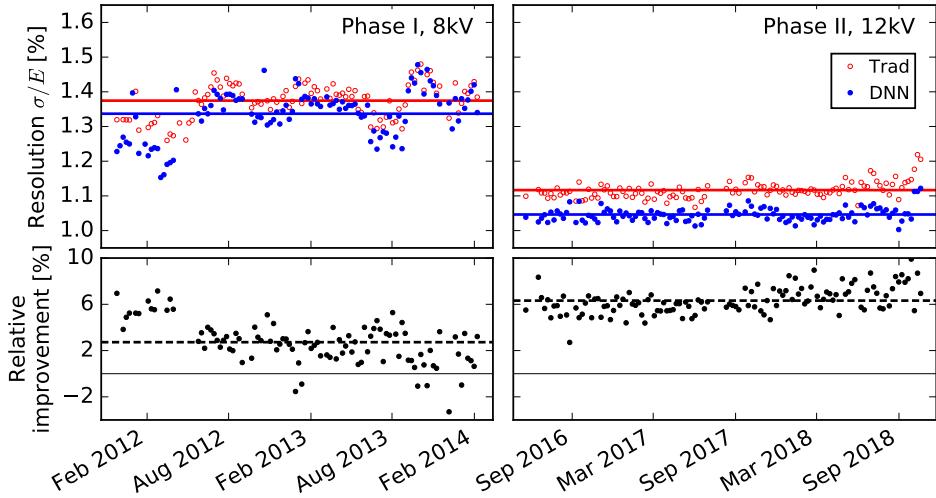


Figure 4.21: The top panels show the weekly energy resolution at the ^{208}Tl γ line for SS events after combining the ionization and scintillation signals. The left panels show Phase-I of EXO-200 operation and the right ones Phase-II. The cathode voltage is -8kV (-12kV) in Phase-I (Phase-II). The bottom panels show the relative improvement of the DNN over the traditional EXO-200 reconstruction. Horizontal lines in all panels indicate the mean values.

of the DNN reconstruction of 2.9 % (6.3 %) for Phase-I (Phase-II) over the traditional EXO-200 reconstruction. The weekly calibration from ^{228}Th calibration data is then used to determine an average calibration from all source calibration data following Equation 4.6 and 4.7. The average energy resolution of the DNN for SS events is $\sigma/Q = 1.22\%$ (1.01 %) for Phase-I (Phase-II). For comparison, the resolution of traditional EXO-200 reconstruction is 1.35 % (1.15 %). This represents a relative improvement in energy resolution of the DNN over the traditional EXO-200 reconstruction by 10 % (12 %) for Phase-I (Phase-II). The calibrated SS energy spectra of ^{60}Co , ^{226}Ra and ^{228}Th calibration sources in Phase-II of EXO-200 operation are shown in Figure 4.22 (top panel). The spectra by the DNN and the traditional EXO-200 (not shown) reconstruction agree well. Measured data and MC simulations show an overall good spectral agreement within 20 % for both reconstruction methods (bottom panel). The discrepancy at the low energy tail of the ^{208}Tl γ line is known for several EXO-200 analyses and is present in both reconstruction methods. The overall similar spectral agreement for both methods confirms an unbiased and robust energy measurement of the DNN.

The energy resolution σ/E as a function of energy is shown in Figure 4.23 for Phase-II of EXO-200 operation. In general, the resolution improves with increasing energy for both DNN and traditional EXO-200 reconstruction. The resolution is degraded in MS events compared to SS events. The resolution of the DNN is improved by about 2 % throughout the entire energy range over the traditional EXO-200 reconstruction. In SS events, the resolution from the DNN measurement outperforms that of the traditional EXO-200 reconstruction. In Phase-I, the improvement increases from 1 % at 1000 keV to 12 % at 3000 keV. In Phase-II, the corresponding improvement is from 4 % to 14 % in that energy range. A key contribution to the improvement is understood to be the improved disentangling of mixed collection and induction signals

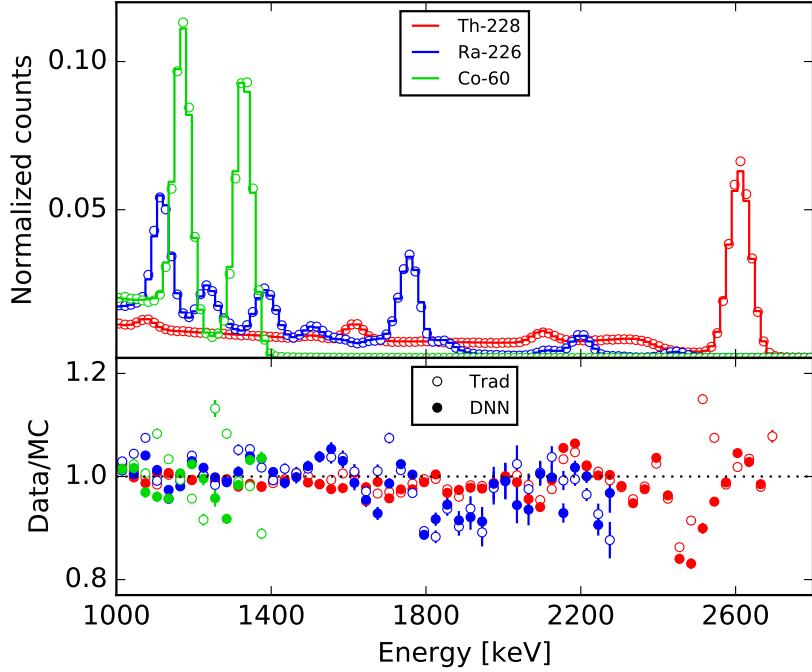


Figure 4.22: The top panel shows the SS energy spectra after combining the ionization and scintillation signals of calibration sources (^{60}Co , ^{226}Ra , ^{228}Th) reconstructed by the DNN in measured data (circles) and MC simulations (lines). The bottom panel shows the ratios of measured data and MC simulations for the DNN (solid) and the traditional EXO-200 (open) reconstruction.

with the DNN (see Section 4.2.3). Since induction signals on U-wires occur more often with increasing energy, the improvement in energy resolution becomes larger with increasing energy. At $Q = 2458$ keV, the improvement in energy resolution with the DNN is 10% (12%) in Phase-I (Phase-II) over the traditional EXO-200 reconstruction. The energy resolution at Q as a function of the date of submission is shown in Figure 4.24. It shows previously published analyses of EXO-200 together with the results of the DNN presented here. For a fair comparison, only analyses that use the MC based fit calibration procedure are shown. These include the $2\nu\beta\beta$ search published in 2014 [51] and the $0\nu\beta\beta$ searches published in 2014 [49], 2018 [91], and 2019 [14]. This comparison reveals the big improvement in energy resolution with the DNN reconstruction in both phases of EXO-200 operation.

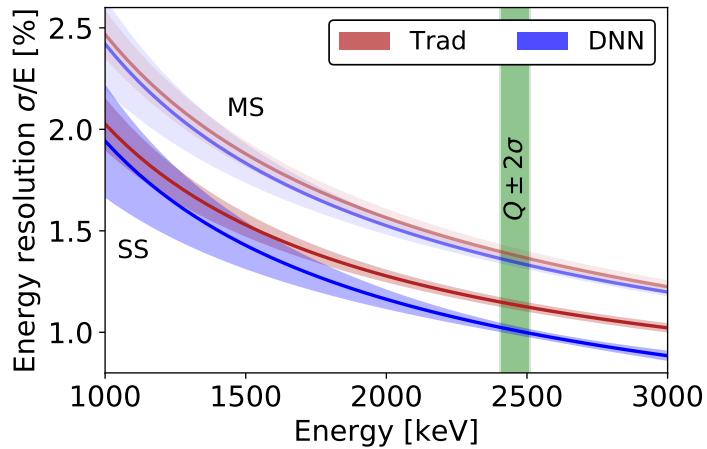


Figure 4.23: The energy resolution σ/E as a function of the energy is shown separately for SS (dark, bottom curves) and MS (light, top curves) events in Phase-II for the DNN (blue) and the traditional EXO-200 (red) reconstruction. The bands (shaded) indicate the fluctuations observed from weekly calibrations. The ROI, $Q \pm 2\sigma$, defined by the traditional EXO-200 reconstruction is indicated as green band. The ROI defined by the DNN is not shown separately as it is only slightly smaller.

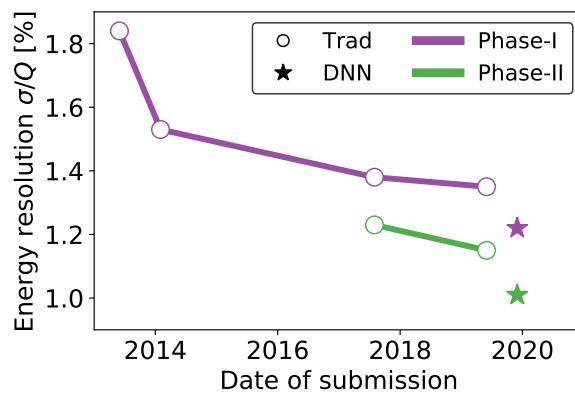


Figure 4.24: The energy resolution σ/Q at $Q = 2458$ keV for SS events is shown as a function of the date of submission. The results of the traditional EXO-200 reconstruction are shown (circles) and those of the DNN (stars). Results for Phase-I (purple) and Phase-II (green) are shown separately.

4.3.2 Background reduction in the ROI of $0\nu\beta\beta$

Besides improving the energy resolution, the DNN also improves the capability to disentangle complex signatures of mixed collection and induction signals compared to the traditional EXO-200 reconstruction [70]. This ability recovers energy in mixed signals that is partially missed by the traditional EXO-200 reconstruction. For MC simulations, this is discussed in Section 4.2.3.

A dominant background contribution in the ROI around the Q value of ^{136}Xe are events that originate from naturally occurring ^{232}Th decay chain components. These arise in the materials surrounding the liquid xenon. This contribution includes the ^{208}Tl γ line at 2615 keV. Missing a fraction of energy in reconstruction for these events presents an immediate risk of increasing the background contribution B in the ROI. Since the background contribution is related to the half-life sensitivity for a particular isotope exposure via $T_{1/2}^{0\nu} \propto 1/\sqrt{B}$ (see Equation 2.11), this directly affects the $0\nu\beta\beta$ half-life sensitivity of EXO-200.

The improvement in $0\nu\beta\beta$ half-life sensitivity from the reduced contribution of ^{232}Th decay chain events in the ROI can be estimated using ^{228}Th source calibration data. This data also presents the ^{208}Tl γ line that is expected to leak into the ROI if energy is missed in reconstruction. Figure 4.25 shows the corresponding SS energy spectra (top panel). The bottom panel shows the ratio of the DNN and the traditional EXO-200 reconstruction. The ROI is usually defined as the region of $Q \pm 2\sigma$. Due to improvements to energy resolution and to disentangling mixed collection and induction signals, this ratio falls below one in the ROI. This indicates a reduction of background events using the DNN over the traditional EXO-200 reconstruction. The relative improvement to the half-life sensitivity can be estimated by:

$$\frac{T_{1/2}^{\text{DNN}}}{T_{1/2}^{\text{Trad}}} = \sqrt{\frac{B^{\text{Trad}}}{B^{\text{DNN}}}} . \quad (4.9)$$

To assess the background in the ROI, the best-fit values of the background determined in the most recent EXO-200 $0\nu\beta\beta$ search are used [14]. The fit concept is discussed in Section 6.1. The background component specific to ^{232}Th , $B_{\text{Th-}232}^{\text{Trad}}$, is scaled by the relative reduction of events η^{DNN} of the DNN over the traditional EXO-200 reconstruction observed in ^{228}Th source calibration data. The contribution from other background components, being mostly events from the ^{238}U decay chain and from ^{137}Xe , are left unchanged. The estimated background using the DNN reconstruction is:

$$B^{\text{DNN}} = B^{\text{Trad}} - B_{\text{Th-}232}^{\text{Trad}} \cdot \eta^{\text{DNN}} , \quad (4.10)$$

with the best-fit value of number of events for the total background in the ROI $B^{\text{Trad}} = 32.3 \pm 2.3$ (30.9 ± 2.4) for Phase-I (Phase-II). The corresponding number of the ^{232}Th background contribution is $B_{\text{Th-}232}^{\text{Trad}} = 10.0$ (8.2) [14].

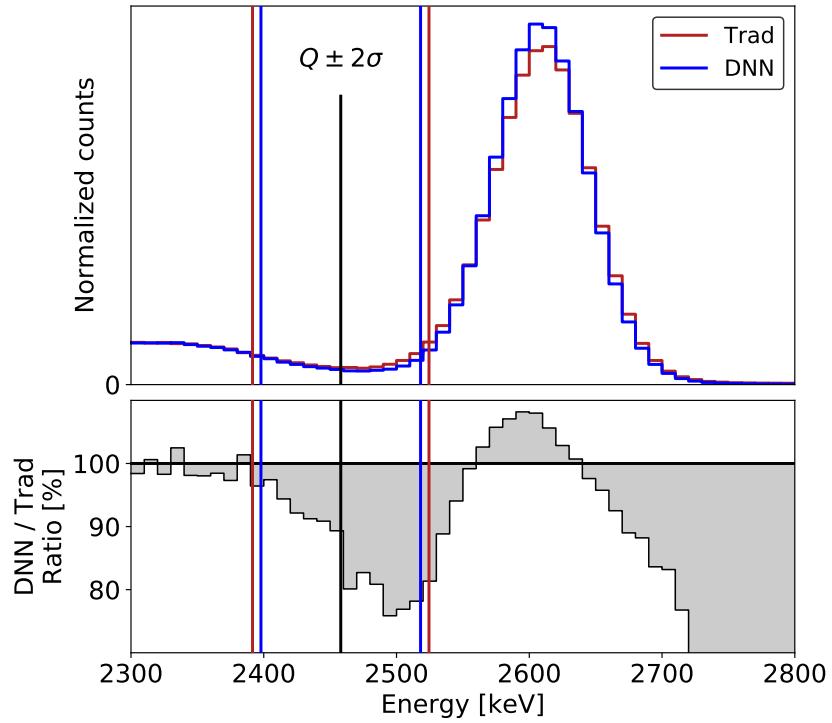


Figure 4.25: The top panel shows the SS energy spectra of ^{228}Th source calibration data recorded in Phase-I of EXO-200 operation for the DNN (blue) and the traditional EXO-200 (red) reconstruction. The lower panel shows the ratio of both methods. Colored vertical lines represent the ROI, $Q \pm 2\sigma$, around the Q value (black).

Table 4.2 shows the reduction of background events η^{DNN} in the ROI and the corresponding estimation on the improvement to half-life sensitivity. With this simplified measure for half-life sensitivity improvement, the DNN is better by 4.3% (2.9%) for Phase-I (Phase-II). The improvement is larger in Phase-I compared to Phase-II since the relative background contribution of ^{232}Th is larger. In addition, the energy resolution in Phase-I is worse which broadens the ROI around the Q value and the dominant ^{208}Tl γ line.

Table 4.2: Estimation of half-life sensitivity improvement from reduced ^{232}Th background levels in the ROI with the DNN over the traditional EXO-200 reconstruction. The numbers are evaluated for ^{228}Th source calibration SS data that also presents the ^{208}Tl γ line. “Fixed ROI” represents the results based on the ROI of the traditional EXO-200 reconstruction and “new ROI” those of the ROI determined from the corresponding energy resolution at the Q value.

Type	Resolution $\sigma/E @$		Reduction η of ^{232}Th events in		Sensitivity improvement
	2615 keV	2458 keV	fixed ROI	new ROI	
Phase-I	Trad	1.29 %	1.35 %	–	–
	DNN	1.16 %	1.22 %	13.7 %	26.1 %
Phase-II	Trad	1.11 %	1.15 %	–	–
	DNN	0.97 %	1.01 %	7.4 %	20.8 %

4.4 Summary

In this chapter, an alternative approach was presented for the reconstruction of the ionization energy in the EXO-200 detector. The new reconstruction method is based on a deep neural network (DNN). Several performance checks validate a robust energy reconstruction of the developed DNN. However, performing the training with ill-suited data can introduce a bias. These problems were mitigated successfully by a careful composition of the training dataset. In terms of spatial uniformity, the DNN shows very similar performance on MC simulations and on measured data compared to the traditional EXO-200 reconstruction. The stability in time of the DNN reconstruction shows a similar variation as that of the traditional EXO-200 reconstruction which implies that it is dominated by variations in the operation of the detector. These tests confirm the robust energy reconstruction by the DNN on both MC simulated and measured data.

The energy resolution observed with the DNN based energy measurement outperforms that of the traditional EXO-200 reconstruction for both MC simulated and measured data in both phases of EXO-200 operation. The improvements in reconstructing the ionization energy translate to improvements in the combined energy measurement of ionization and scintillation. Both energy measures, from ionization only and from the combination with the scintillation measurement, show good spectral agreement between measured data and MC simulations for source calibration data. The big improvement in energy resolution of 10 % (12 %) in Phase-I (Phase-II) with the DNN reconstruction shows the strong potential of DNN based methods for other complex analyses. The promising results presented here further pave the way for evaluation of a potential improvement to the $0\nu\beta\beta$ half-life sensitivity using the DNN based energy reconstruction. This study is presented in Chapter 6.

Chapter 5

Signal-Background Discrimination with Deep Neural Networks

Contents

5.1	Goal and design of study	62
5.1.1	Motivation and training data preparation	62
5.1.2	Network architecture and training	64
5.2	Performance on Monte Carlo simulated data	67
5.2.1	Comparison of DNN and EXO-200 BDT	68
5.2.2	DNN validation	71
5.3	Performance on measured data	74
5.4	Alternative discriminator concepts	77
5.5	Summary	83

This chapter presents several deep neural network (DNN) based concepts to discriminate double beta decays from the most dominant background in EXO-200 which is γ interactions. After discussing the design choices made to train the DNN for this task, the performance of the DNN is validated in Monte Carlo (MC) simulations. The DNN is compared to the boosted decision tree (BDT) based discriminator that was used in the EXO-200 $0\nu\beta\beta$ search in 2018 [91]. The spectral agreement between measured data and MC simulations is evaluated for signal and background events.

The study presented in this chapter contribute to the publication “Search for Neutrinoless Double-Beta Decay with the Complete EXO-200 Dataset” [14] published in Physical Review Letters (PRL).

5.1 Goal and design of study

5.1.1 Motivation and training data preparation

Having multiple parameters to discriminate candidate $0\nu\beta\beta$ decays from other interactions is a very important feature of the EXO-200 experiment. This implies that the EXO-200 experiment is not a simple counting experiment but pursues a multi-parameter analysis. Since most backgrounds originate from γ sources external to the xenon volume, the spatial distribution contains additional information on the signal type. In previous analyses, this information was captured via the so-called standoff distance r_{SD} which measures the minimum distance of an event to the closest material surface excluding the cathode. Furthermore, the detailed shape and topology of signals can be used to discriminate double beta ($\beta\beta$) decays from γ particle interactions. This information was first captured via a boosted decision tree (BDT) in Ref. [92] and was later used in the EXO-200 $0\nu\beta\beta$ search in 2018 [91] which led to an improvement in $0\nu\beta\beta$ half-life sensitivity of $\sim 15\%$. This large improvement from a signal-background discriminator motivates the evaluation of a deep neural network (DNN) based discriminator.

Different design choices are possible for this discriminator. The design choices of the baseline concept are made to pursue the following goals, while alternative concepts will be discussed in Section 5.4.

1. Maximizing the topological discrimination power between signal $\beta\beta$ decays and background γ particle interactions.
2. Being able to evaluate the DNN in both measured data and MC simulations as well as for signal and background events.
3. Keeping a reasonable level of spectral agreement between measured data and MC simulations in the DNN discriminator variable.

The first goal, maximizing the discrimination power of $\beta\beta$ and γ events, is achieved by training the DNN over the energy range of 1000–3000 keV. In particular, the discriminator does not focus on $0\nu\beta\beta$ events at $Q = 2458$ keV but is trained over a broad energy distribution. The signal $\beta\beta$ and background γ events are simulated to have the same spatial and energy distributions. That way, the discrimination is based on topological features only. By this design choice, the DNN does not know about differences in spatial distributions between $\beta\beta$ decays and γ interactions. While $\beta\beta$ decays occur uniformly in the detector volume in low-background (LB) physics data, γ interactions occur mostly close to material surfaces. The definition of single-site (SS) and multi-site (MS) events that is used as a binary classifier in EXO-200 data analysis is not known to the DNN explicitly. However, the DNN is expected to learn a highly correlated and continuous representation of this definition implicitly.

The second goal motivates the design of a general $\beta\beta$ over an explicit $0\nu\beta\beta$ discriminator. Discriminators of both designs can be validated on MC simulations. With a $\beta\beta$ discriminator, the signal class can be validated with measured $2\nu\beta\beta$ data. However, for a $0\nu\beta\beta$ discriminator, corresponding measured data is not available which can be used for validation since it is a hypothetical decay. For background events, measured data is available for validation from source calibration measurements.

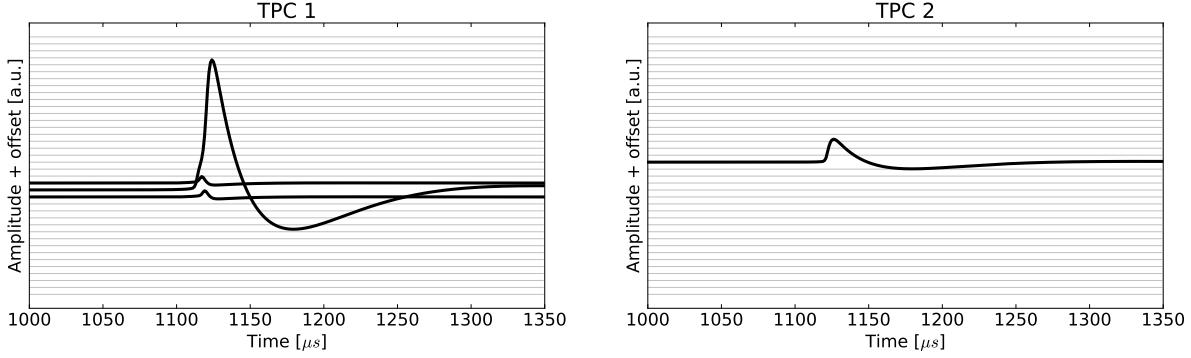


Figure 5.1: Re-generated U-wire waveforms of the identical example event shown in Figure 4.1. The channel amplitudes are shown as a function of time. The channels are offset for clarity. The gray channels do not contain signals and are exactly zero.

The third goal is to keep the spectral agreement between measured data and MC simulations on an acceptable level. This implies that any DNN concept that utilizes the raw V-wire signals is disfavored since they are not modelled in MC simulations to the precision of the U-wires. Nevertheless, alternative DNN concepts that use V-wires are discussed in Section 5.4. Also, due to this third goal, two DNN concepts are pursued in this study that differ by their input images:

DNN-Raw is fed with the identical images as the DNN used for energy measurement discussed in Chapter 4. An example event is shown in Figure 4.1. The inputs for this DNN concept are grayscale images built by arranging neighboring U-wire channels next to each other and encoding the raw amplitudes of the U-wire waveforms as pixel values for both TPCs separately. This approach has the advantage of not relying on the traditional EXO-200 reconstruction as it directly exploits the raw waveforms.

DNN-Recon is fed with images where the raw waveforms are replaced by those that are re-generated from the best-fit signals of collection and induction templates found by the traditional EXO-200 reconstruction. An example event for this concept is shown in Figure 5.1 where the identical event for the DNN-Raw concept is shown in Figure 4.1. The DNN-Recon concept is pursued in order to mitigate inaccuracies in modelling the raw U-wire signals in MC simulations. Further, the prediction of this DNN discriminator is guaranteed to be based only on features already known to the traditional EXO-200 reconstruction which supports the understanding of the DNN discriminators performance. On the other hand, the discrimination power may be limited by the precision of the traditional EXO-200 reconstruction. In particular, this applies to small energy deposits captured by DNN-Raw that are below the threshold of EXO-200 reconstruction. It also applies to entangled collection and induction signals where DNN approaches proved to outperform the traditional EXO-200 reconstruction (see Section 4.2.3).

In fact, the same dataset was used for training the discriminator DNNs as for the DNN used for energy measurement that is discussed in Chapter 4. The training data is produced in MC simulations for two classes of events: background-like, composed of γ particle interactions with

a uniform energy distribution in the range of 1000–3000 keV; and signal-like $0\nu\beta\beta$ decays with a random decay energy Q restricted to the same energy range. The locations of the interactions of both event classes are simulated uniformly in the detector volume to focus discrimination on the topological event characteristics only. The dataset is populated evenly with events of both classes and is divided into statistically independent datasets for training (90 %) and for validation (10 %).

5.1.2 Network architecture and training

The basic concept of the DNN architecture and the way the input data is processed and fed to it is equivalent to the approach described in Section 4.1.2. Again, the U-wire images of both TPCs are fed to the DNN separately where each image has a dimension of 350×38 pixels. Both images pass through the same layers with the same mathematical operations applied to them in order to reduce the number of free parameters of the DNN. Both TPC branches are concatenated immediately before the output of the DNN. The architecture that is used is inspired by the Inception architecture which was proposed by Google [84]. The concept of Inception is built on Convolutional Neural Networks by combining convolutional layers with different sized receptive fields in parallel. With this approach, the DNN is able to capture both small and big features at the same time. These parallel feature maps are merged before being fed to the subsequent layer.

The complete architecture consists of a stem followed by 10 Inception blocks. The stem consists of 3 convolutional layers with 32, 32 and 64 kernels and a (4, 2) maximum pooling layer before the last convolutional layer. The Inception part includes 3 maximum pooling layers after the 2, 4 and 6 Inception blocks with (2, 2), (2, 1) and (2, 1) pool sizes in order to reduce the spatial dimensions intentionally. To maintain the spatial dimensions in the convolutional operation, zero padding is applied to each layer. This operation surrounds the image with zeros. Also, batch normalization is applied with each convolutional layer. This technique acts as a regularizer that helps to prevent overfitting and renders the training more robust [93]. After the last Inception block, the spatial average of each feature map is calculated rather than applying a fully connected part. This design choice helps to prevent overfitting during training. By removing the spatial information, the DNN is more robust to spatial translation in the input image, i.e. in the detector volume [94]. The vectors which contain the averages over feature maps from both TPC branches are concatenated and fed to the last layer which is responsible for the final discrimination. Throughout the network, the $\text{ReLU}(x) = \max(x, 0)$ activation function is used. In the last classification layer the softmax activation function is used instead:

$$\sigma(\mathbf{z})_i = \frac{e^{z_i}}{\sum_{j=1}^K e^{z_j}} \quad \text{for } i = 1, \dots, K \text{ and } \mathbf{z} = (z_1, \dots, z_K). \quad (5.1)$$

The softmax function maps each component z_i of a non-normalized input vector \mathbf{z} to the interval $\sigma(\mathbf{z})_i \in (0, 1)$ while ensuring $\sum_{i=1}^K \sigma(\mathbf{z})_i = 1$. This activation function is commonly used for classification tasks as it allows to interpret the outputs of the DNN as a probability distribution over the predicted output classes. Trainable weights are initialized using the Glorot algorithm [80] discussed in Equation 3.14. The DNN architecture is implemented with the

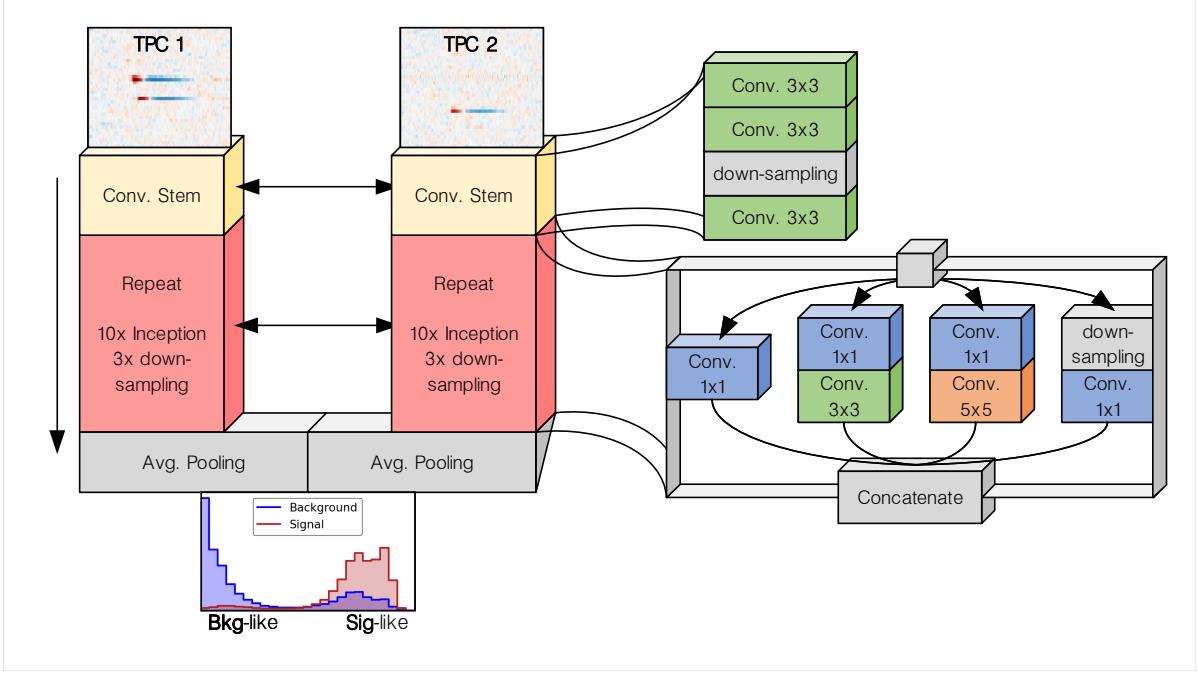


Figure 5.2: The network architecture used for the DNN-Raw and DNN-Recon concepts. The input consists of two images with a dimension of 350×38 pixels each. A stem containing 3 convolutional layer is followed by 10 Inception blocks. Each block includes parallel convolutional layers with receptive fields of different sizes in order to be sensitive to different sized features. The resulting feature maps are concatenated at the end of the block. The spatial dimensions of the feature maps are reduced at several steps to save computation time. The final feature maps are transformed into a one-dimensional array. Both branches which contain information from the individual TPCs share the same network weights. Before the final output, the units of both branches are concatenated. The last layer has two units which correspond to the probability of the event belonging to either signal or background class.

Keras library [88] using the Tensorflow backend [89]. The architecture is shown in Figure 5.2.

In MC simulations, the target variable is known. Background γ events are assigned to class 0 and signal $\beta\beta$ events to class 1. The dissimilarity between the target distribution p and the predicted distribution q can be described by the cross entropy $H(p, q)$:

$$H(p, q) = - \sum_k^K p(k) \log(q(k)) , \quad (5.2)$$

where K is the number of classes. In training, the Adam optimizer [90] is used to minimize the mean cross entropy L over a batch of N samples:

$$L = \frac{1}{N} \sum_i^N H(p_i, q_i) . \quad (5.3)$$

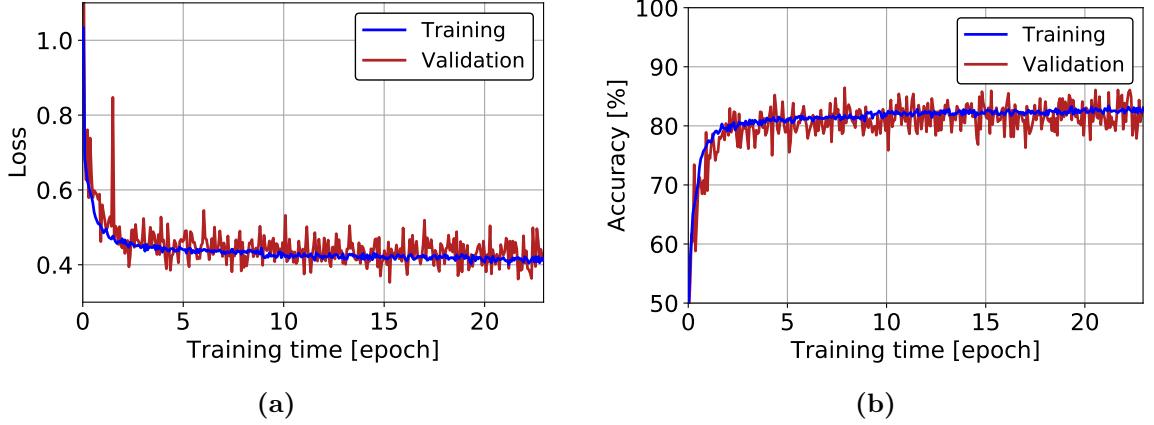


Figure 5.3: The loss curves of the training of DNN-Raw used for Phase-II (a) in validation (red) and training (blue) data. Both curves are used to monitor potential overfitting which is not observed here. In parallel, the accuracy (b) is monitored which reflects the fraction of correctly identified events.

For classifiers with only two classes ($K = 2$), the predicted and target values of both classes can be written as $p \in \{p_{y=1} = y, p_{y=0} = 1 - y\}$ and $q \in \{q_{y=1} = \hat{y}, q_{y=0} = 1 - \hat{y}\}$. Then, the mean cross entropy can be written as [71]:

$$L = -\frac{1}{N} \sum_i^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] . \quad (5.4)$$

Both DNN concepts (DNN-Raw and DNN-Recon) are trained on the identical events in MC simulations and with the identical DNN architecture and training procedure. The only difference is the input images. There are two DNNs trained for each DNN concept since EXO-200 was operated in two phases with slightly different conditions. This helps to optimally capture all features available in training data for the individual phases. However, the DNNs of both phases of EXO-200 operation are identical in terms of DNN architecture, training data production, and training directives. The training curves of DNN-Raw used for Phase-II are shown in Figure 5.3.

5.2 Performance on Monte Carlo simulated data

The performance of both DNN concepts, DNN-Raw and DNN-Recon, is validated in MC simulations. For this, the 10 % from the training dataset reserved for validation is used. In addition, MC simulations that correspond to measured data, like $2\nu\beta\beta$ or source calibration data, are used.

The distributions of the DNN-Raw discriminator from the validation dataset are shown in Figure 5.4. They are separated by the target variable into background (blue) and signal (red) class. Besides showing all events from the dataset (left), the corresponding distributions are shown for SS (center) and MS (right) events. Although, the SS/MS definition used in EXO-200 data analysis is already a binary classifier used to discriminate signal $\beta\beta$ from background γ events, the DNN can access additional discrimination power in both channels.

To infer the discrimination power of a binary classifier, a common measure is the receiver operating characteristic (ROC) curve [95, 96]. It shows the signal efficiency as a function of the background rejection power as the cutoff value is varied. DNN discriminators, like other continuous discriminators, can be converted to binary classifiers by introducing a cutoff value that maps the predictions to either class. A perfect discriminator produces a point in the ROC curve at (1, 1). This implies that there is no loss in signal efficiency while the background is fully rejected. A random discriminator would produce points along the bisector of the ROC curve where the loss in signal efficiency changes linearly with the background rejection. To compare ROC curves, the area under the curve (AUC) can be used to condense the discrimination power to a single number. While a perfect discriminator would have an AUC of 1, the random discriminator has an AUC of 0.5. However, condensing the performance of a discriminator to a single number represents a loss of information and may promote misinterpretation [97]. Since the ROC curve depends on the prevalence of both classes, here, it is only used to directly compare different discriminators based on the same data. Also, datasets with a high imbalance between the classes must be interpreted with high caution [97]. Various metrics exist to analyze and interpret discriminators and classifiers besides ROC analysis [98].

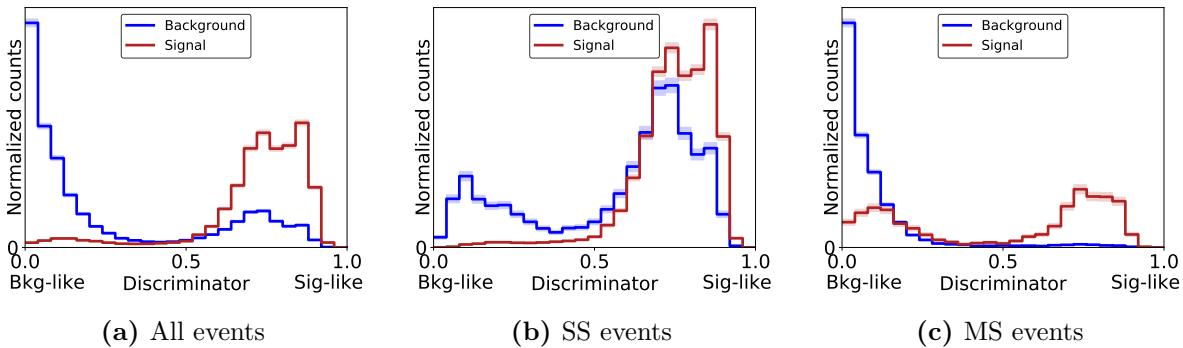


Figure 5.4: Distribution of the DNN-Raw variable in Phase-II for signal (red) and background (blue) validation events. (a) shows all events, (b) and (c) only those belonging to the SS and MS class, respectively.

The conversion from a continuous discriminator to a binary classifier causes a loss of information. It is no longer known how reliable the prediction was. If the DNNs were used as binary classifiers, a common concept is to select the cutoff value based on the Youden's J statistic [99]. Graphically, the optimum cutoff value can be chosen by maximizing the height above the chance line of a random discriminator.

The ROC analysis is only used as initial metric to evaluate the performance of different DNN discriminator concepts and to compare them to alternative approaches existing in EXO-200 data analysis. The final metric relevant for EXO-200 is the $0\nu\beta\beta$ half-life sensitivity that is discussed in Chapter 6.

5.2.1 Comparison of DNN and EXO-200 BDT

To evaluate their performances, both DNN concepts are compared to the BDT discriminator which was used in the EXO-200 $0\nu\beta\beta$ search in 2018 [91]. This BDT was only used in SS events as toy studies indicated no additional improvement to $0\nu\beta\beta$ half-life sensitivity when using it in MS events. The BDT comprises two variables which contain topological information and another variable which contains spatial information. Topological variables were chosen to describe the extent of signals in Z dimension via the rise time of U-wire signals and in U dimension via the number of affected U-wire channels. The spatial information is captured via the standoff distance [91].

In contrast to the training dataset, in low-background (LB) physics data, γ particles originate from sources external to the xenon. Thus, γ interactions predominantly occur close to material surfaces while $2\nu\beta\beta$ and hypothetical $0\nu\beta\beta$ decays occur uniformly distributed in the xenon volume. By design, the DNN discriminators do not capture this spatial difference but focus on topological discrimination only. Different DNN concepts which capture spatial information are discussed in Section 5.2.2. Thus, for a fair comparison of DNN and BDT discriminators, the DNN discriminator is combined with the standoff distance in order to also include spatial discrimination power. Both variables are wrapped in a BDT discriminator. Their distributions for LB data in MC simulations are shown in Figure 5.5. This approach is only used for qualitative performance checks and comparisons in this chapter. The half-life sensitivity to $0\nu\beta\beta$ for analysis configurations including a DNN discriminator is discussed in Chapter 6. Alternative DNN based concepts that directly exploit spatial information are discussed in Section 5.4.

Figure 5.6a shows the ROC curves of the DNN-Raw and the BDT discriminators. It shows LB physics data in MC simulations consisting of $0\nu\beta\beta$ decays and background events from the dominant background in the ROI of $0\nu\beta\beta$ which is from the decay chains of ^{232}Th and ^{238}U . The AUC of the DNN exceeds that of the BDT. However, there are points of the ROC curves where the BDT outperforms the DNN by also capturing spatial information. This deficit is compensated by the discriminator that combines the DNN-Raw and the standoff distance variable. This discriminator then covers naturally the pure DNN but also the BDT discriminator at all points of the ROC curve. This implies an improved discrimination of $\beta\beta$ and γ events with the DNN based approach over the BDT based approach and suggests a potential improvement to $0\nu\beta\beta$ half-life sensitivity.

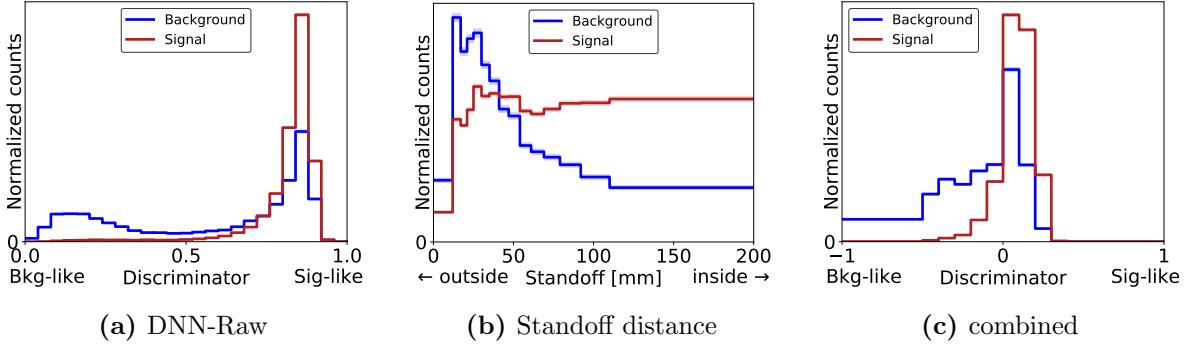


Figure 5.5: Distribution of the DNN-Raw variable in Phase-II for SS events (a). Signal events are simulated $0\nu\beta\beta$ decays (red) while background events (blue) represent the main γ background to the ROI of $0\nu\beta\beta$ from ^{232}Th and ^{238}U . The standoff distance contains information about the spatial distribution of both event classes (b). Combining both variables to a more powerful discriminator in a BDT allows to consider both topological and spatial information (c).

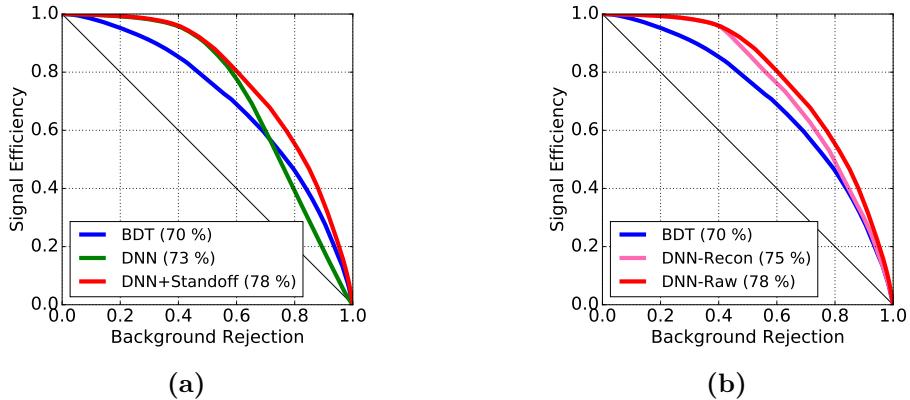


Figure 5.6: ROC curves of several discriminators for SS events from the dominant background in the ROI and $0\nu\beta\beta$ decays in MC simulations. The area under the curve is given in brackets. The ROC curve of DNN-Raw (green) improves after including spatial information via the standoff distance (red). It then completely covers the curve of the BDT (blue) that was used in the EXO-200 $0\nu\beta\beta$ search in 2018 [91] (a). Combined with the standoff distance, DNN-Recon (pink) also outperforms the BDT while showing a slightly degraded discrimination power compared to DNN-Raw (b).

Figure 5.6b shows the ROC curves comparing DNN-Raw and DNN-Recon, both combined with the standoff distance, and the BDT used in the EXO-200 $0\nu\beta\beta$ search in 2018 [91]. The DNN-Recon concept yields a slightly degraded discrimination power compared to DNN-Raw. This is expected since DNN-Recon is limited to the precision of EXO-200 reconstruction whereas DNN-Raw can fully exploit any information available in the raw waveforms. Both DNN discriminator concepts outperform the BDT discriminator at all points of the ROC curve which implies potential for improving the $0\nu\beta\beta$ half-life sensitivity with DNN discriminators.

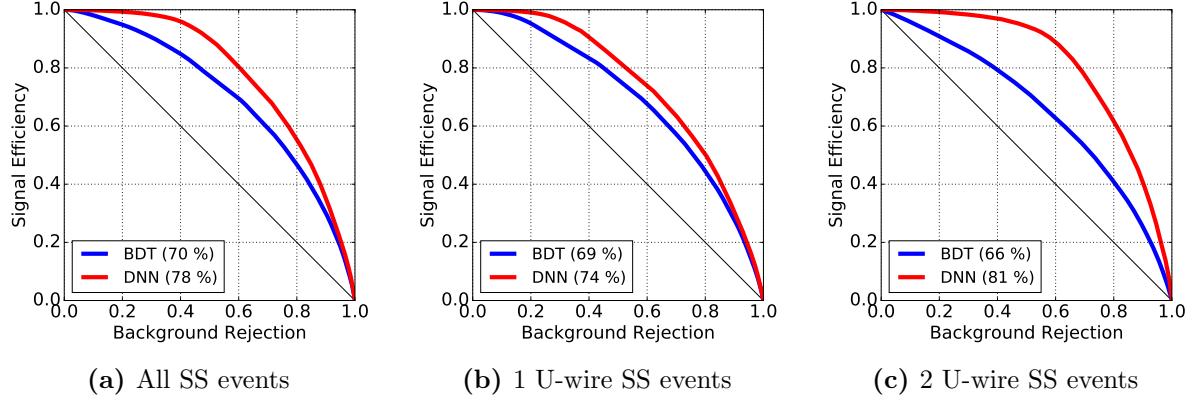


Figure 5.7: Comparison of ROC curves of BDT [91] and DNN-Raw combined with standoff distance for all (a), 1 U-wire (b), and 2 U-wire (c) SS events. The area under the curve is given in brackets.

In order to evaluate why the DNN discriminators in combination with the standoff distance variable outperform the BDT discriminator, the dataset of SS events is further divided into two distinct classes. Figure 5.7 shows the comparison of ROC curves of DNN-Raw and the BDT for all SS events (left), and for those affecting only 1 (center) or 2 (right) U-wires. Most of the improvement of the DNN based approach is seen for 2 U-wire SS events for which the BDT discriminator is less powerful. This becomes clear by the ROC curves and when comparing their AUC values. On the other hand, the DNN has an even improved discrimination power for 2 compared to 1 U-wires events. The inferior performance of the BDT on these events is understood from the fact that for each reconstructed signal an individual rise time is measured. But since the BDT only takes scalar values as input, two values are condensed to a single number by taking the maximum value. This represents a loss of information for the BDT discriminator. On the other hand, for these events, the DNN discriminator uses both signals and can therefore constrain the spatial extensions even better than for 1 U-wire events. Extensions to the BDT that focus on improving the discrimination power in particular for 2 U-wire SS events were discussed internally in EXO-200. However, they did not reach the discrimination power of the DNN based approaches.

Both DNN concepts, DNN-Raw and DNN-Recon, outperform the BDT discriminator used in the EXO-200 $0\nu\beta\beta$ search in 2018 [91] in terms of discrimination power. The big advantage of the DNN over the BDT is that it alleviates the need to manually engineer weak discriminator variables from the raw data. Instead, it allows for using the raw data directly as input. This is of particular benefit for complex event signatures which contain multiple and potentially overlapping signals. For these events, the BDT shows degraded performance whereas that of the DNNs even improves. The DNN-Recon concept has the advantage that its predictions are guaranteed to be completely based on information that is known to EXO-200 reconstruction. This makes it more comprehensible than the DNN-Raw concept. Instead of re-generating the input images for the DNN-Recon concept, the best-fit parameters from template fits to the raw waveform traces could be directly used as input to a DNN. This would bypass the need to re-generate the entire event and would reduce the input to a few numbers. However, the

approach presented here offers a natural solution of local correlation and to handle a varying numbers of signals in events. Further, criticism of DNNs being a black-box can be addressed by this concept more easily than for the DNN-Raw concept. As only certain signal types can be re-generated for the input to DNN-Recon, their effects on the DNN’s prediction can be evaluated easily. For example, this includes the effect of induction U-wire signals on the DNN prediction. Of course, the DNN’s discrimination power is limited to the precision of EXO-200 reconstruction which explains the slightly degraded performance compared to DNN-Raw.

5.2.2 DNN validation

To gain insight into the DNN, its performance is evaluated for different scenarios in MC simulations. These tests verify a robust performance by showing that the prediction is correlated to physically relevant topological features that distinguish $\beta\beta$ from γ events. Furthermore, potential biases of the DNN may be revealed by these tests.

The correlation of the DNN discriminator is evaluated as a function of the true event size. For this test, the spatial dimension is determined in MC simulations. The event size represents the maximum distance between energy deposits in liquid xenon that exceed the detection threshold of ~ 20 keV. $\beta\beta$ decays tend to have a smaller spatial extension than γ particle interactions. This can be understood from the fact that γ interactions produce a single photo-electron via photoelectric effect which contains the total deposited energy. On the other hand, the energy in $\beta\beta$ decays is partitioned over two electrons. Both electrons usually comprise angles of less than 180° and the range of electrons does not scale linearly with energy, leading to smaller spatial extensions of $\beta\beta$ over γ events on average. In addition, γ particles often interact multiple times in the detector volume via Compton scattering effect, leading to multiple spatially distinct energy deposits. Both $\beta\beta$ and γ events may produce additional small, separated energy deposits due to bremsstrahlung. The distribution of spatial extensions of $2\nu\beta\beta$ and ^{226}Ra calibration data in MC simulations is shown in the top panel in Figure 5.8 for all events, i.e. for SS and MS events. The hypothesis is that the DNNs capture the differences in spatial extension of $\beta\beta$ and γ events and correlate their predictions on these differences. The bottom panel in Figure 5.8 shows the fraction of correctly predicted events for DNN-Raw and DNN-Recon of both signal $2\nu\beta\beta$ and background ^{226}Ra events as a function of the true event size. The predictions of the DNN are converted into classifications with the cutoff value determined from the Youden’s J statistic [99]. Basically all signal events with small spatial extension and background events with large spatial extension are classified correctly. This confirms the hypothesis that the DNNs strongly correlate their predictions to the event size. The offset in the curves between DNN-Raw and DNN-Recon is expected to arise from the finite precision of EXO-200 reconstruction, especially for small energy deposits.

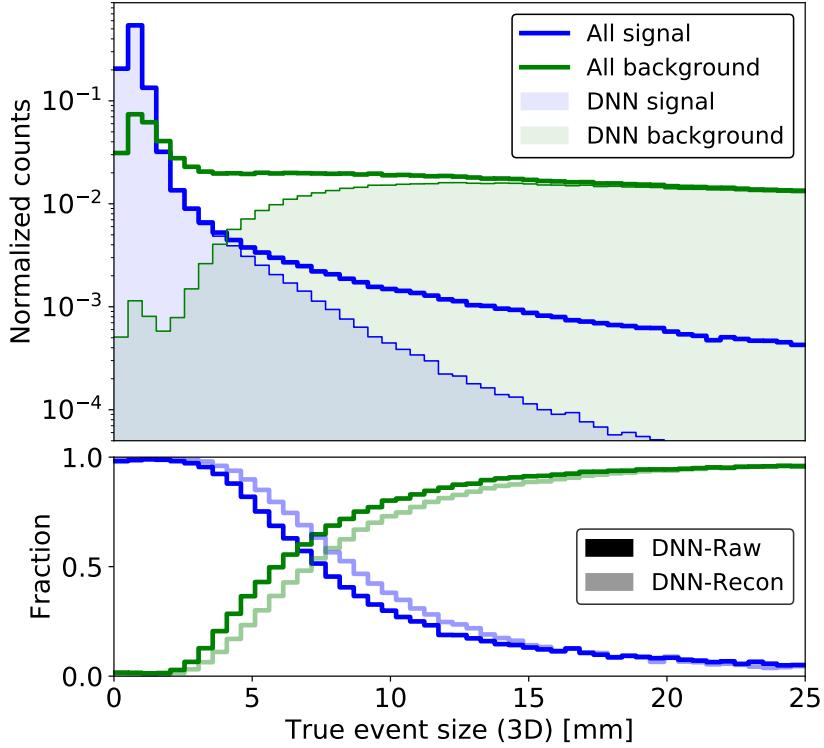


Figure 5.8: The top panel shows the distributions of event sizes which are known in MC simulations for signal $2\nu\beta\beta$ (blue) and background ^{226}Ra (green) events. The filled areas show the distributions that are classified correctly by DNN-Raw. The prediction cutoff value is determined based on the Youden's J statistic [99]. The bottom panel shows the fraction of correctly classified events for DNN-Raw (dark) and DNN-Recon (light).

By design, the DNN discriminators should not be biased spatially since they are trained on uniformly distributed events in the detector volume. In LB physics data, hypothetical signal $0\nu\beta\beta$ events are indeed distributed uniformly in the xenon volume, whereas the dominant background γ particles originate from the materials of the TPC vessel. Thus, background γ events mostly occur close to the detector walls. The top panel in Figure 5.9 shows the standoff distance distribution for signal $2\nu\beta\beta$ and background ^{226}Ra events in MC simulations. The bins are chosen to contain an equal xenon volume. The bottom panel shows the fraction of correctly classified events for DNN-Raw and the BDT used in the EXO-200 $0\nu\beta\beta$ search in 2018 [91]. The cutoff value is determined analogously to Figure 5.8. Unlike the BDT which explicitly uses spatial information, the DNN-Raw discriminator variable is independent of the standoff distance. This confirms that the DNN is not biased spatially.

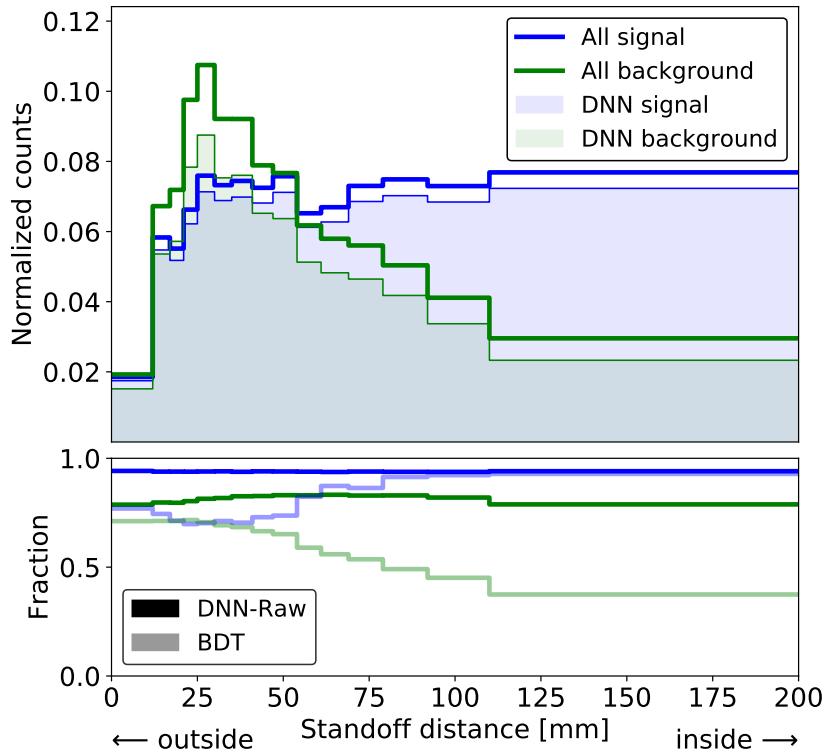


Figure 5.9: The top panel shows the standoff distance distributions of signal $2\nu\beta\beta$ (blue) and ^{226}Ra (green) events in MC simulations. The filled areas show the distributions that are classified correctly by DNN-Raw. The prediction cutoff value is determined based on the Youden's J statistic [99]. The bottom panel shows the fraction of correctly classified events for DNN-Raw (dark) and BDT [91] (light).

5.3 Performance on measured data

The DNN discriminators are applied to measured data. To account for different trigger strategies, the DNN input images are processed analogously to the approach discussed for the DNN energy measurement in Section 4.3. This implies, adjusting the window of the input image relative to the trigger. However, for the DNN discriminators there is no need to apply a correction that accounts for the finite electron lifetime.

For a successful use of the DNN discriminators in $0\nu\beta\beta$ searches, the spectral agreement between measured data and MC simulations must not be overly degraded. Any discrepancies in the agreement are taken into account as systematic uncertainties and thus decrease the sensitivity of a $0\nu\beta\beta$ analysis configuration. This is discussed in more detail in Section 6.1.4. The spectral agreement is evaluated for DNN-Raw and DNN-Recon on both signal and background data. Calibration source data of ^{228}Th , ^{226}Ra , and ^{60}Co is used to evaluate the spectral agreement on background γ data. Signal $\beta\beta$ data is not available from calibration measurements, but $2\nu\beta\beta$ data is recorded in LB physics data. However, this dataset is contaminated by background γ events. To extract the $2\nu\beta\beta$ events, a maximum likelihood (ML) model is fitted to the energy spectrum of LB physics data whereas the ROI around the Q value is blinded. More information on the fit is given in Section 6.1. Measured $2\nu\beta\beta$ data is extracted by subtracting all background components in LB data according to their best-fit values. While the number of available $2\nu\beta\beta$ events is limited, this background-subtracted $2\nu\beta\beta$ data still is the best proxy for signal $\beta\beta$ events recorded by EXO-200. Figure 5.10 shows the best-fit LB distributions of DNN-Recon (black) for SS (left) and MS (right) events. The best-fit background distribution (green) in MC simulations is subtracted from the LB physics data to build the background-subtracted $2\nu\beta\beta$ dataset (blue).

The spectral agreement between measured data and MC simulations is shown in Figure 5.11 for DNN-Raw (top) and DNN-Recon (bottom). Background-subtracted $2\nu\beta\beta$ data is used as proxy for signal data (black), and background γ data (red) consists of calibration data from a ^{226}Ra source located near the cathode outside the TPC. The agreement is shown for all (left), SS (center), and MS (right) events separately. The spectral agreement is improved in γ data compared to background-subtracted $2\nu\beta\beta$ LB data. In SS events, there is an overall shift in MC simulations compared to measured data toward lower discriminator values, i.e. toward more background-like predictions. This shift is visible as a slope in the ratio of measured data and MC simulations that is present in both DNN concepts. This trend represents a risk to the applicability in $0\nu\beta\beta$ searches by introducing a systematic bias when fitting the MC simulated spectra to LB physics data. Its impact on the $0\nu\beta\beta$ half-life sensitivity is discussed in Section 6.1.4. The binary SS/MS classification of EXO-200 reconstruction further introduces discrepancies in the DNN agreement between measured data and MC simulations. This is perceptible when omitting the SS/MS classification as this improves the spectral agreement. This effect is expected since the SS/MS classifier and the DNN discriminators are correlated by their design goal. Any discrepancies in the binary classifier are propagated to the spectral agreement of the DNN discriminator variables. Thus, the spectral agreement between measured data and MC simulations is improved when not discriminating SS from MS events.

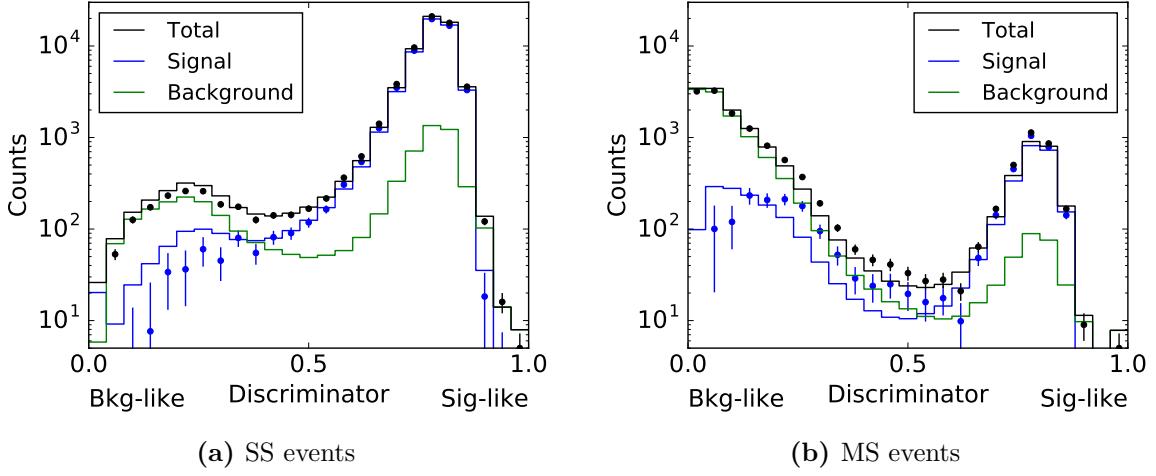


Figure 5.10: The distributions of the DNN-Recon discriminator for the best fit to LB physics data in SS events (left) and in MS events (right) in measured data (circles) and in MC simulations (lines) recorded in Phase-II of EXO-200 operation. Background-subtracted $2\nu\beta\beta$ events (blue) in LB physics data are extracted from the total dataset (black) by subtracting the best-fit background contributors in MC simulations (green).

DNN-Raw shows a degraded agreement between measured data and MC simulations compared to DNN-Recon. This is understood to arise from two dominant aspects. First, inaccuracies in modelling the raw pulse shapes in MC simulations affect only DNN-Raw which is trained on raw signals. For DNN-Recon, this problem does not exist because it is trained on signals found by EXO-200 reconstruction and the same analytical waveform template is used in reconstruction of measured and simulated waveforms. Second, inaccuracies in modelling complex effects in EXO-200 like induction on U-wires and diffusion at the raw waveform level directly affect the spectral agreement of DNN-Raw. For DNN-Recon, these effects are shielded to some degree by the precision of EXO-200 reconstruction while causing a slightly degraded discrimination power (see Figure 5.6b).

The discrepancies between measured data and MC simulations can be mitigated by profiling the DNN discriminator variables at the cost of losing discrimination power. This approach is beneficial as long as the impact of reduced systematic uncertainties outweighs the loss of discrimination power. Profiling the variables between 0, 0.3, 0.6 and 1 improves the spectral agreement, especially for the bins containing the majority of $2\nu\beta\beta$ events. Analogously to the fine binned scenario, the spectral agreement of DNN-Raw is worse than that of DNN-Recon. The impact of the spectral agreement to the $0\nu\beta\beta$ half-life sensitivity is discussed in more detail in Section 6.1.4.

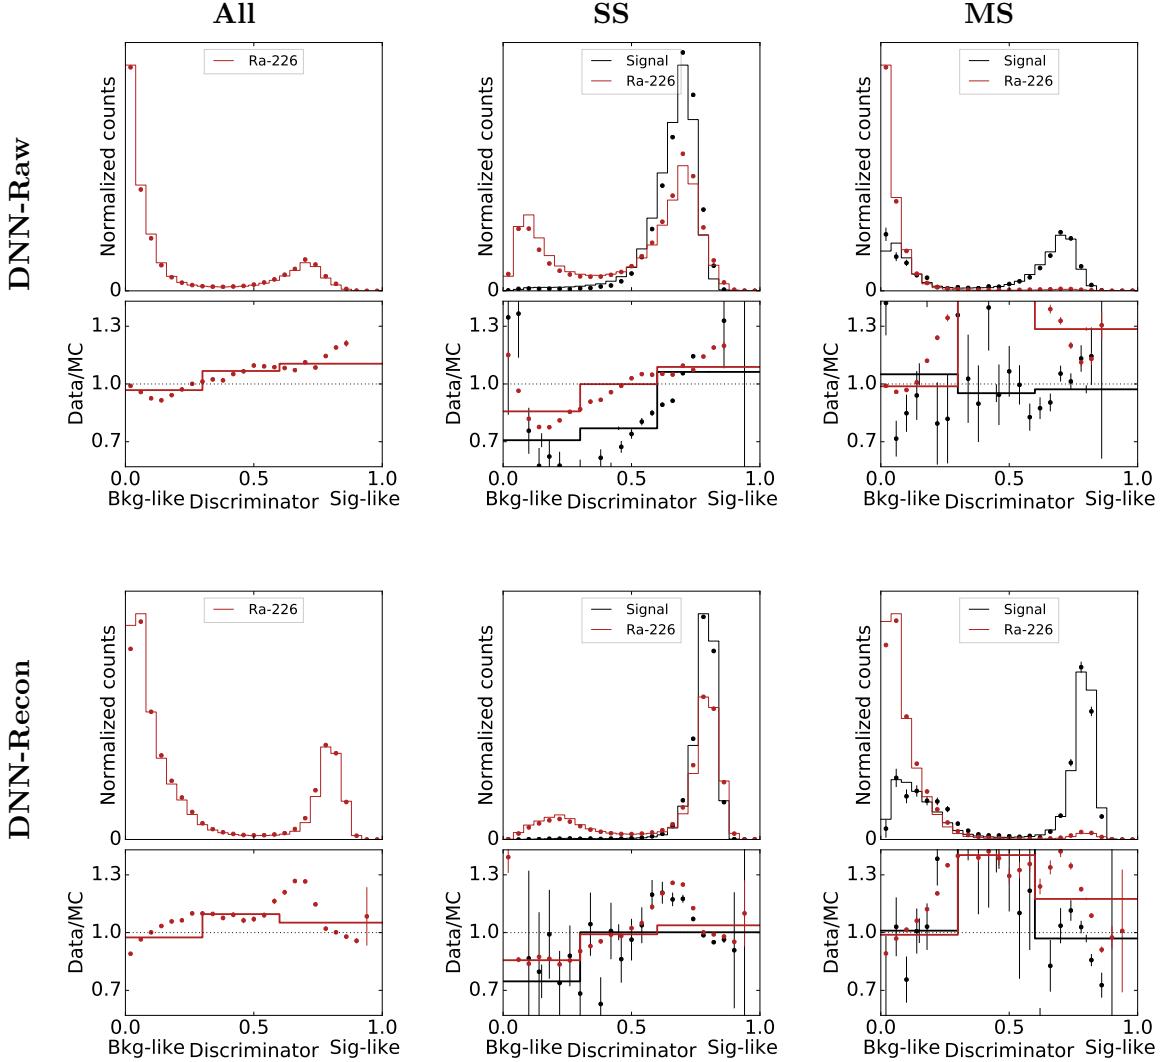


Figure 5.11: The spectral agreement between measured data and MC simulations is shown for DNN-Raw (top row) and DNN-Recon (bottom row). The top panels show the distributions in measured data and MC simulations while the bottom panels show their ratios. Background γ data (red) represents data from a ^{226}Ra calibration source located near the cathode outside the TPC. Signal $\beta\beta$ data (black) is from background-subtracted $2\nu\beta\beta$ LB data. The spectral agreement is shown individually for all events (left), SS events (center) and MS events (right). In addition, the bottom panels show the ratios of measured and MC simulated data when profiling the DNN discriminators between 0, 0.3, 0.6 and 1.

5.4 Alternative discriminator concepts

The concepts presented so far focused on topological discrimination only while ignoring any differences in spatial distributions between signal $\beta\beta$ and background γ events present in LB physics data. Although these concepts are valid, here, different concepts are discussed that to also capture spatial information directly.

In order to capture spatial information, the background γ data in the training dataset is replaced by γ particle interactions that are no longer distributed uniformly in the detector volume. The signal $\beta\beta$ data is kept unchanged, since in LB physics data, $2\nu\beta\beta$ and hypothetical $0\nu\beta\beta$ decays occur uniformly in the xenon volume. The new background data is shown in Figure 5.12 and consists of γ particles emitted in MC simulations by the materials of the TPC vessel. This represents the dominant background contribution in LB physics data, especially in the ROI of $0\nu\beta\beta$ [14]. Thus, the DNN is trained on a realistic spatial distribution with respect to the γ background in LB physics data. However, unlike LB physics data, the $\beta\beta$ and γ energy distributions used for training the DNN are still uniform in order to not introduce a bias to the discriminator in terms of in energy.

The new dataset is used to train a DNN on raw U-wire signals. The DNN architecture is identical to that of the baseline DNN concept shown in Figure 5.2. This DNN concept (DNN-U) is sensitive to differences in spatial distributions of γ and $\beta\beta$ events as is shown in Figure 5.13. Background γ events are more likely to be identified correctly near the material surface. The trend is inverse for signal $0\nu\beta\beta$ events. However, since this DNN is trained on U-wires only, the full 3D spatial information is not available to the DNN. Instead, only the U and Z coordinates of the event position are fully available to the DNN. While the V -wires are not known to the DNN, the V coordinate of the event position is still partially available since both wire planes are crossed by 60° and thus are not fully uncorrelated (see Figure 2.9). This explains why the DNN identifies background and signal events mostly uniformly as a function of the V position. The corresponding ROC curve is shown in Figure 5.14 (red). It improves compared to the baseline DNN concept (black) that was trained on a uniform spatial distribution of background γ events (see Figure 5.6a). However, it does not completely achieve the discrimination power of

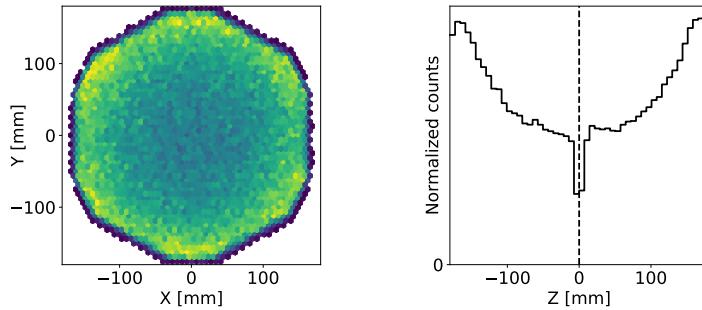


Figure 5.12: Correlation of the x and y spatial coordinates (left) and the distribution of z coordinates (right) of the new background dataset which represents γ particle interactions. The γ particles are emitted in MC simulations by the TPC vessel materials which represents the dominant γ background in the ROI of $0\nu\beta\beta$.

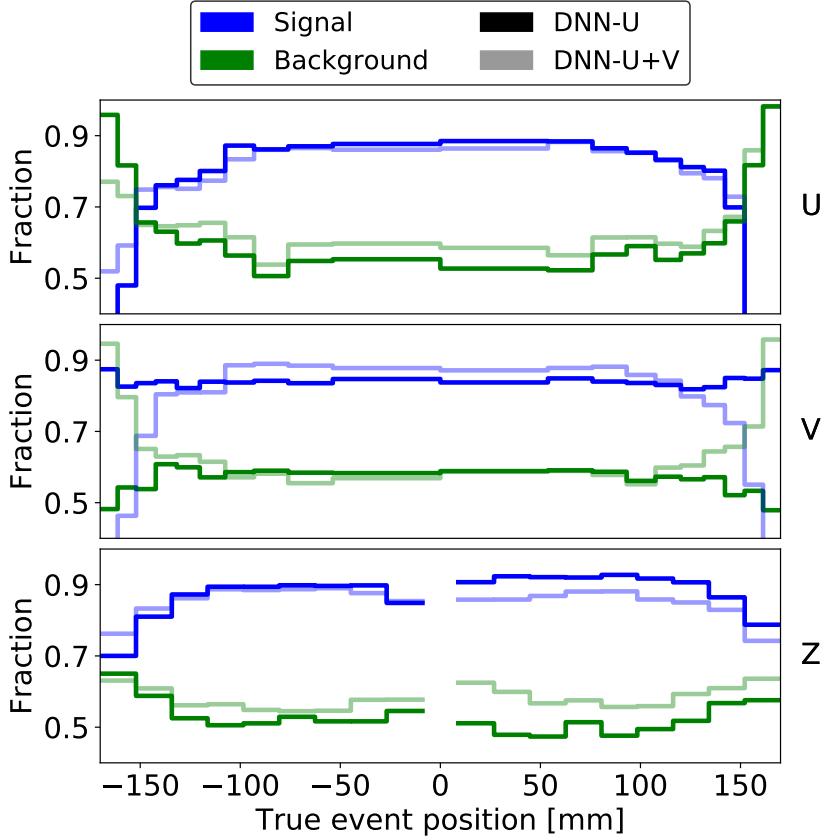


Figure 5.13: The panels show the fraction of correctly classified signal $0\nu\beta\beta$ (blue) and background γ (green) SS events as a function of the event position in the U (top), V (center) and Z (bottom) coordinates. Shown are curves for the DNNs trained on U -wires (dark) and on U - and V -wires (light). The prediction cutoff value is determined based on the Youden's J statistic [99]. Background γ events represent the main background in the ROI which occurs predominantly close to the material surfaces while signal $0\nu\beta\beta$ events are distributed uniformly in the xenon volume.

that DNN in combination with the standoff distance due to the partial loss of spatial information in V position.

The full spatial information can be captured when training a DNN on U - and V -wires at the same time (DNN- $U+V$). The input to this DNN- $U+V$ is 4 grayscale images consisting of 350×38 pixels each. An example input image is shown in Figure 5.15. Otherwise, the DNN architecture is kept unchanged to the one of the baseline DNN (see Figure 5.2). This concept, DNN- $U+V$, is now sensitive to the full 3D spatial information as is shown in Figure 5.13. The classification as a function of the U and Z coordinates of the event position follows the trend of DNN- U . In V position, DNN- $U+V$ is able to reject background γ events more efficiently toward the material surface while DNN- U was mostly unaffected by this spatial component. Using this additional spatial information, the ROC curve of DNN- $U+V$ (green) in Figure 5.14 outperforms that of DNN- U (red). Furthermore, it slightly outperforms that of the baseline DNN combined

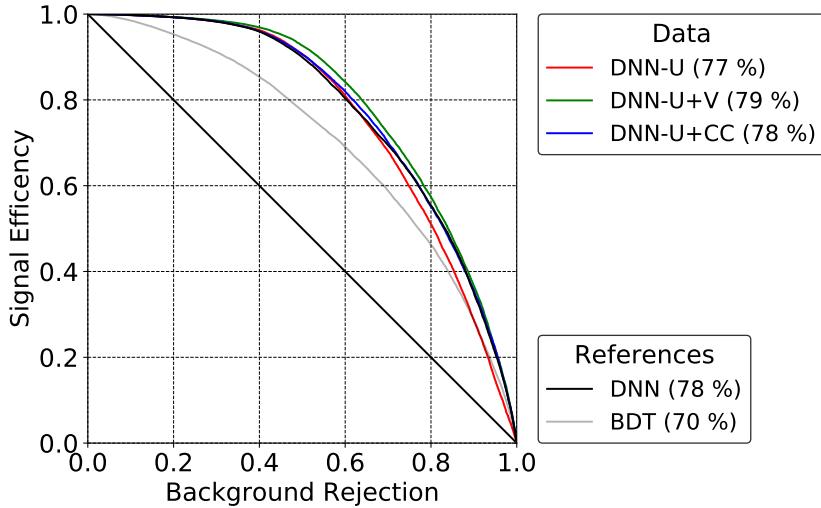


Figure 5.14: The ROC curves are shown for different DNN concepts from SS events of the main γ background to the ROI and signal $0\nu\beta\beta$ in MC simulations. For reference, the baseline DNN (black) combined with the standoff distance and the BDT (gray) used in the EXO-200 $0\nu\beta\beta$ search in 2018 [91] are shown. All concepts shown here capture spatial information and thus outperform the baseline DNN shown in Figure 5.6a. The DNN concepts discussed include those trained on U-wires only (DNN-U, red), on U- and V-wires (DNN-U+V, green), and on U-wires with additional information about the charge clusters of an event (DNN-U+CC, blue). The area under the curve is given in brackets.

with the standoff distance (black). This is due to not only adding additional spatial information with the raw V-wire waveforms but also adding additional topological information.

Using both U- and V-wires also has disadvantages since the raw V-wire waveforms are not modelled to the precision of the raw U-wire waveforms. This is due to the higher complexity of the signatures on V-wire channels. Consequently, the spectral agreement between measured data and MC simulations is degraded for the DNN-U+V concept. To quantify and compare the agreement to other DNN concepts, the error weighted average \bar{r} of the absolute residuals between measured data and MC simulations is used:

$$\bar{r} = \frac{\sum_i w_i \cdot |r_i|}{\sum_j w_j} \quad \text{with} \quad r_i = \left(\frac{\text{Data} - \text{MC}}{\text{MC}} \right)_i, \quad (5.5)$$

where the sums span over all bins i and j , and the uncertainty contributing to w_i is calculated assuming Poisson statistics. The spectral agreement of the DNN-U+V variable between measured data and MC simulations is shown in Figure 5.16b. The data represents interactions from γ particles emitted from a ^{228}Th source located near the cathode outside the TPC. Using Equation 5.5, the average residual is $(16.4 \pm 0.7)\%$. For comparison, it is $(8.8 \pm 0.4)\%$ for the baseline DNN concept (see Figure 5.16a). So, while improving the discrimination power when using both U- and V-wires, the spectral agreement is worsened at the same time. Insufficient spectral agreement limits the applicability of the discriminator for $0\nu\beta\beta$ searches as will be discussed in Chapter 6.

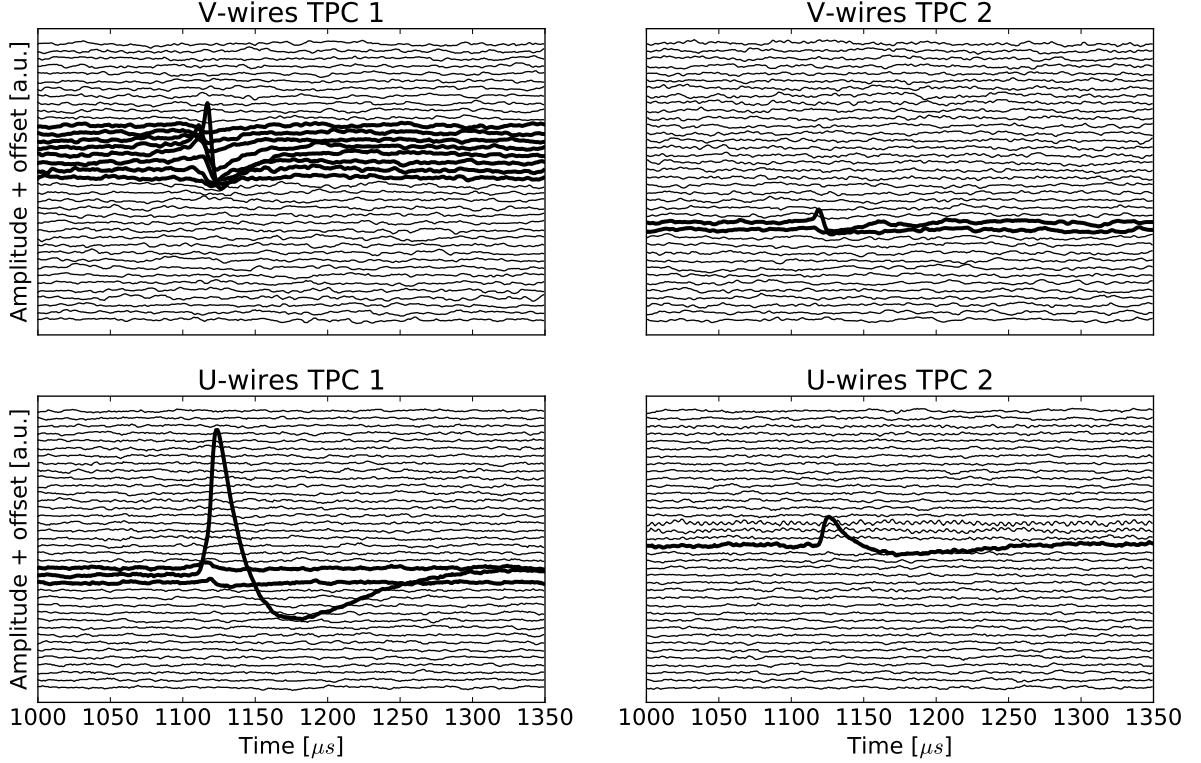


Figure 5.15: U-wire (bottom) and V-wire (top) waveforms of the identical example event shown in Figure 4.1 and 5.1. The channel amplitudes are shown as a function of time. The channels are offset for clarity and the channels are shown in bold that contain any signal found by EXO-200 reconstruction.

To bypass the limitations from using raw V-wires while still capturing the full 3D spatial information, another DNN concept is evaluated. This concept (DNN-U+CC) builds on the baseline DNN concept of training on uniform spatial γ and $\beta\beta$ distribution. The DNN architecture is extended by a second branch which combines topological and spatial information. The spatial information is fed to the DNN via a vector containing the 3D positions and energies (u, v, z, E) of all charge clusters found in EXO-200 reconstruction. The vector containing purely topological information and the additional vector are concatenated before passing through three fully connected layers with a decreasing number of neurons. Then, another DNN output predicts the signal-likeness of the event. The architecture of this DNN concept is shown in Figure 5.17. This DNN provides two separate predictions. One is fully based on topological information and the other combines topological and spatial information. To improve the training procedure, the topological DNN branch is pre-trained separately and its weights are frozen during training of the appended branch. The ROC curve of this DNN is shown in Figure 5.14 (blue). The discrimination power exceeds that of the DNN that uses U-wires (red) only which is due to exploiting the full 3D spatial information. Furthermore, it slightly outperforms the baseline DNN combined with the standoff distance (black). This could be due to combining topological and spatial information at an earlier, higher dimensional stage compared to combining both

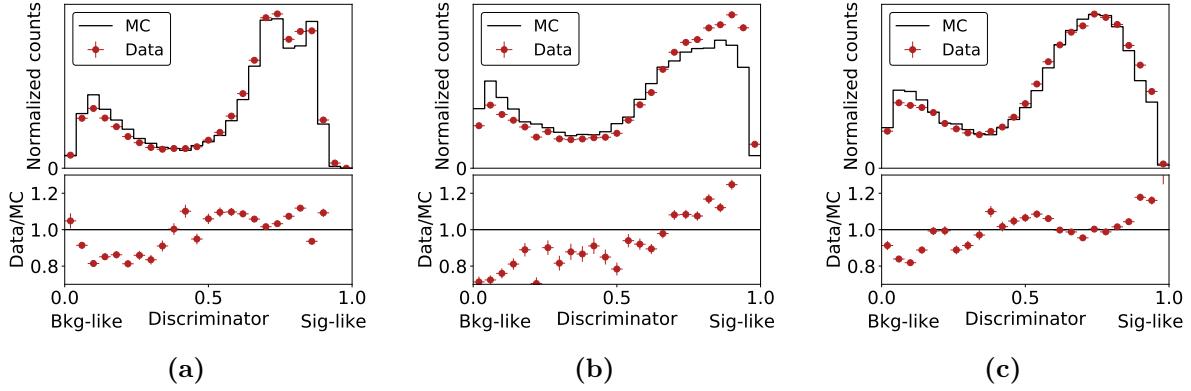


Figure 5.16: The top panels show the spectral agreement for SS events of different DNN concepts of a ^{228}Th calibration source located near the cathode outside the TPC in measured data (circles) and MC simulations (lines). The bottom panels show the corresponding ratios. (a) shows the baseline DNN concept, (b) the DNN concept trained on U- and V-wires (DNN-U+V), and (c) the DNN concept trained on U-wires together with charge cluster variables (DNN-U+CC).

pieces of information at the very end in another discriminator. However, it does not reach the discrimination power of DNN-U+V. First, this is due to relying on EXO-200 reconstruction compared to raw waveforms for spatial information. Second, using raw V-wires adds some additional topological information that is not available in this concept. The spectral agreement of this DNN concept is shown in Figure 5.16c with an average residual of $(6.0 \pm 0.4)\%$. This implies that with the DNN-U+CC concept, the full 3D spatial information can be exploited similar to the DNN trained on U- and V-wires while the spectral agreement is still on the level of the DNN concept trained on U-wires only.

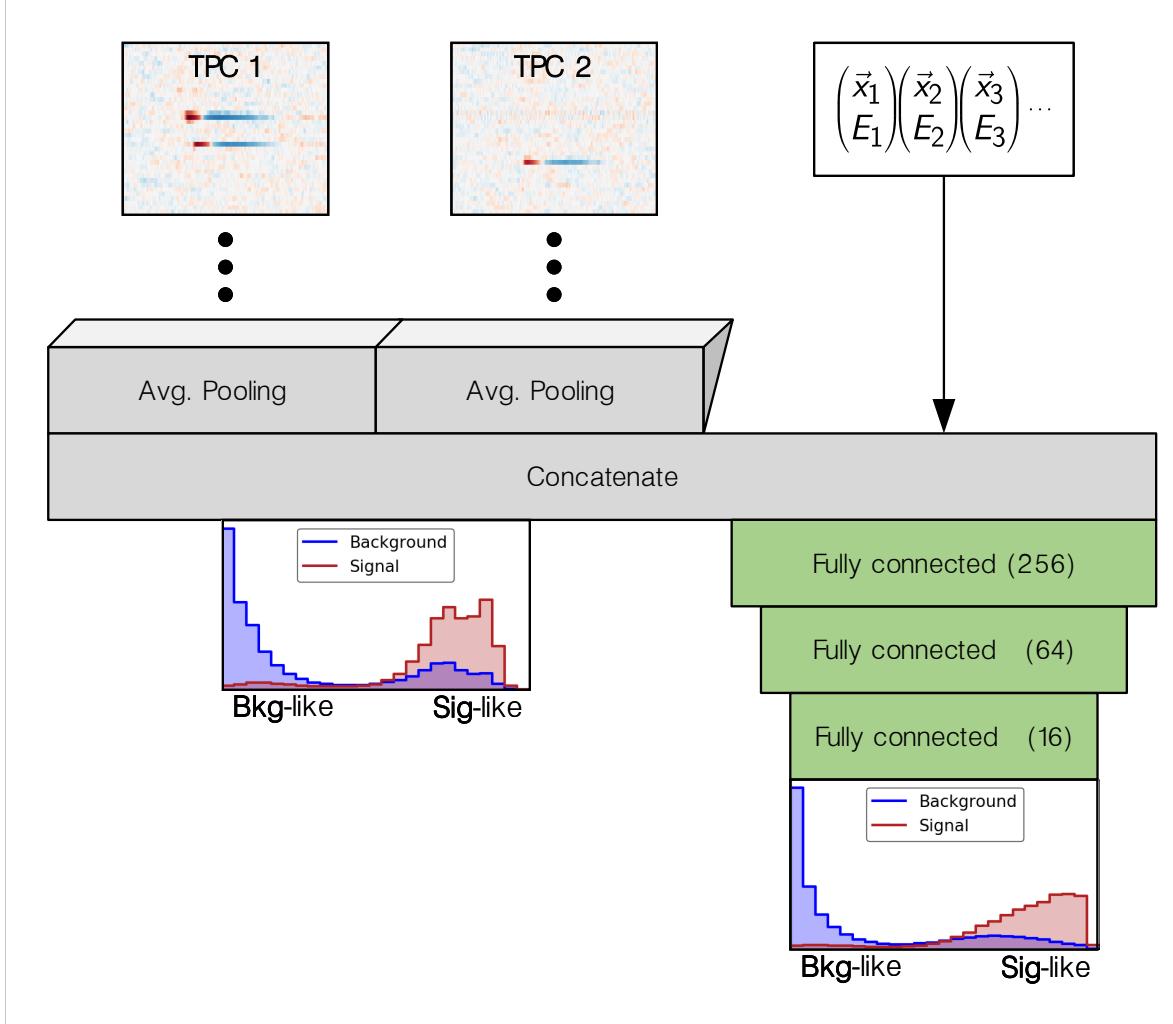


Figure 5.17: The DNN architecture is an extension of the one discussed in Figure 5.2. Additional spatial information is added to the vector that contains topological information. The additional information is the energy E and the coordinates (u, v, z, E) of every charge cluster as found by EXO-200 reconstruction. This extended vector is passed through three fully connected layers with decreasing numbers of units (256, 64 and 16). The final network output of two units represents the probability of the event belonging to either signal or background class.

5.5 Summary

This chapter presented a deep neural network based approach to discriminate signal $\beta\beta$ from background γ events in EXO-200. Two baseline concepts are evaluated; one using raw U-wire waveforms (DNN-Raw) and one using U-wire waveforms re-generated from signals found by EXO-200 reconstruction (DNN-Recon). Several checks in MC simulations validate a robust discrimination based on physically relevant features. By design, these concepts fully focus on topological features while they do not capture spatial differences between both event classes that are present in low-background physics data. This design choice is motivated by avoiding the need to assume a prior spatial background distribution. Then, the DNN can be applied to any physics analysis without introducing a potential bias from differences in spatial distribution in training and in physics data. Alternative DNN concepts were discussed which are able to also incorporate spatial features in the discriminator.

The presented DNN concepts outperform the discrimination power of the BDT that was used in the EXO-200 $0\nu\beta\beta$ search in 2018 [91]. Also, in contrast to the BDT, the DNNs are not limited to SS events but are trained independent from this definition. They can be applied to all events likewise. The spectral agreement of the discriminator variables between measured data and MC simulations was evaluated on γ source calibration data and on background-subtracted $2\nu\beta\beta$ data. In general, DNN-Recon shows better spectral agreement than DNN-Raw. Profiling the DNN variables between 0, 0.3, 0.6 and 1 improves the agreement at the cost of discrimination power.

The promising performance of the DNN discriminators presented here pave the way for a potential increase of the sensitivity to the half-life of $0\nu\beta\beta$ decay. This study will be presented in Chapter 6.

Chapter 6

Improving the $0\nu\beta\beta$ half-life sensitivity with Deep Learning

Contents

6.1	Analysis concept	86
6.1.1	Likelihood fit model	86
6.1.2	Analysis configurations	88
6.1.3	Background model	89
6.1.4	Systematic uncertainties	90
6.2	Results	99
6.2.1	Published $0\nu\beta\beta$ half-life sensitivity	99
6.2.2	Published $0\nu\beta\beta$ half-life data limit	101
6.2.3	$0\nu\beta\beta$ half-life sensitivity with the DNN energy	104
6.3	Summary	107

This chapter presents the detailed analysis concept of searches for neutrinoless double beta decay ($0\nu\beta\beta$) in EXO-200. This concept was used for the searches in Refs. [14, 49, 91, 100]. The results of the $0\nu\beta\beta$ search which utilizes a deep neural network (DNN) based discriminator are presented which is introduced in Chapter 5. These results contribute to the publication “Search for Neutrinoless Double-Beta Decay with the Complete EXO-200 Dataset” [14] published in Physical Review Letters (PRL). In addition, this chapter discusses analysis concepts that utilize the DNN based energy measurement introduced in Chapter 4 instead of the traditional EXO-200 energy measurement. Potential improvements to $0\nu\beta\beta$ half-life sensitivity gained by concepts that utilize DNN based variables are presented and compared to traditional approaches.

6.1 Analysis concept

6.1.1 Likelihood fit model

In EXO-200, the search for $0\nu\beta\beta$ decay is performed with a binned maximum likelihood (ML) fit. For computational reasons, the negative log-likelihood (NLL) model $\mathcal{L} = -\ln L$ is fitted to low-background (LB) physics data. It is based on the model built in Ref. [51] which was used for EXO-200 $0\nu\beta\beta$ searches in Refs. [14, 49, 91, 100]. The single-site (SS) and multi-site (MS) data is fitted simultaneously. The full model is

$$\mathcal{L} = \mathcal{L}_{\text{SS}} + \mathcal{L}_{\text{MS}} + G_{\text{constr}} , \quad (6.1)$$

where G_{constr} are Gaussian constraints. In the fit, \mathcal{L} is minimized. For each multiplicity $m \in \{\text{SS}, \text{MS}\}$, the likelihood function is

$$\mathcal{L}_m = -\ln L_m = \sum_i [\mu_i^m - (k_{\text{obs},i}^m \ln(\mu_i^m))] , \quad (6.2)$$

where $k_{\text{obs},i}^m$ is the number of events in bin i observed in LB physics data with multiplicity m . μ_i^m is the expected number of events from the fit model. The sum proceeds over all bins in all fit observables. The expected number of events μ_i^m in a bin i is calculated by:

$$\mu_i^m(\mathbf{s}, \mathbf{n}, N) = \int_{i^{\text{th}} \text{ bin}} F^m(\mathbf{s}, \mathbf{n}, N, \mathbf{y}) \, d\mathbf{y} , \quad (6.3)$$

where $\mathbf{s} = \{s_0, \dots, s_{N_{\text{PDF}}}\}$ denotes the SS fraction, $(\text{SS}/(\text{SS}+\text{MS}))$. $\mathbf{n} = \{n_0, \dots, n_{N_{\text{PDF}}}\}$ denotes the total number of events for each PDF that contributes to the model. The complete model is discussed in Section 6.1.3 and consists of N_{PDF} PDFs. N is an overall normalization parameter that carries the uncertainty on the signal detection efficiency. This uncertainty is discussed in Section 6.1.4. The integral in the i^{th} bin is performed for the fit observables \mathbf{y} that always include the energy E but can also include further observables [100]. Different fit configurations are discussed in Section 6.1.2. The integrand function F^m is defined as:

$$F^{\text{SS}}(\mathbf{s}, \mathbf{n}, N, \mathbf{y}) = N \sum_j^{N_{\text{PDF}}} n_j s_j f_j^{\text{SS}}(\mathbf{y}) \quad (6.4)$$

$$F^{\text{MS}}(\mathbf{s}, \mathbf{n}, N, \mathbf{y}) = N \sum_j^{N_{\text{PDF}}} n_j (1 - s_j) f_j^{\text{MS}}(\mathbf{y}) . \quad (6.5)$$

These functions sum over all PDFs $f_j^m(\mathbf{y})$ that are included in the fit model with the SS fraction s_j and the number of events n_j that are associated to PDF j . The PDF $f_j^m(\mathbf{y})$ itself is a function of the set of observables \mathbf{y} used in the fit.

Systematic uncertainties are included in the negative log-likelihood function \mathcal{L} by additional Gaussian constraints which are summarized by G_{constr} in Equation 6.1. Each constraint contributes the following term to \mathcal{L} :

$$G_{\text{constr}}(\boldsymbol{\rho}, \boldsymbol{\rho}_0, \boldsymbol{\Sigma}) = 0.5 (\boldsymbol{\rho} - \boldsymbol{\rho}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\rho} - \boldsymbol{\rho}_0) , \quad (6.6)$$

where $\boldsymbol{\rho}$ denotes the constrained parameters and $\boldsymbol{\rho}_0$ their expectation values. $\boldsymbol{\Sigma}$ is the covariance matrix:

$$\boldsymbol{\Sigma}(\boldsymbol{\rho}) = \begin{pmatrix} \text{Var}(\rho_1) & \text{Cov}(\rho_1, \rho_2) & \cdots \\ \text{Cov}(\rho_2, \rho_1) & \text{Var}(\rho_2) & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix}. \quad (6.7)$$

For uncorrelated parameters ρ_i in $\boldsymbol{\rho}$, Equation 6.6 simplifies to contributions of:

$$G_{\text{constr}}(\boldsymbol{\rho}, \boldsymbol{\rho}_0, \boldsymbol{\Sigma})_i = 0.5 \left(\frac{\rho_i - \rho_{0,i}}{\sigma_i} \right)^2, \quad (6.8)$$

where $\rho_{0,i}$ is the expected value and $\sigma_i = \sqrt{\boldsymbol{\Sigma}(\boldsymbol{\rho})_{ii}}$ is the error of ρ_i . In the most recent EXO-200 $0\nu\beta\beta$ search in 2019 [14], G_{constr} comprises five different constraints:

- The SS fractions \mathbf{s} are constrained to the values observed in Monte Carlo (MC) simulations for each PDF $s_{0,i}$ using the error σ_{SS} determined from source calibration data via $0.5((s_i - s_{0,i})/(\sigma_{\text{SS}} s_{0,i}))^2$. A 85 % correlation between the SS fraction errors of γ -like PDFs is included in the constraint motivated by the level of correlation observed in source calibration data [91].
- The overall normalization N is constrained to one with the error on the signal detection efficiency σ_N via $0.5((N-1)/\sigma_N)^2$.
- An additive normalization is constrained to zero for PDFs that contribute to the ROI. This constraint is used to explicitly take discrepancies into account that affect the ROI. The error on this normalization is determined from background model inadequacies and differences in spectral agreement between measured data and MC simulations.
- The measured activity of ^{222}Rn and its daughter nuclei is used to constrain the activity of three background fit components measured to 10 % precision. A correlation of 90 % between the components is used [51].
- The relative fractions of PDFs related to neutron capture are constrained within 20 % to the values observed in MC simulations [101].

More details on systematic uncertainties and their evaluations are discussed in Section 6.1.4. Since the energy scale is calibrated from γ calibration sources (E_γ), potential differences to the energy scale for β -like and $\beta\beta$ -like components (E_β) are accounted for in the ML fit by a scaling factor B :

$$E_\beta = B E_\gamma. \quad (6.9)$$

The factor B is optimized in the ML fit. However, it is found to be consistent with unity within uncertainty [14].

6.1.2 Analysis configurations

In contrast to simple counting experiments, EXO-200 uses an analysis concept based on a maximum likelihood fit of MC simulated spectra to LB physics data. By this concept, the MC simulated spectra of the different fit model components are used to model the background in the ROI of the $0\nu\beta\beta$ decay. This implies that the data outside the ROI helps to model the background in the ROI. This approach is superior to analysis concepts which focus only on the ROI of $0\nu\beta\beta$. In addition, EXO-200 data analysis distinguishes single-site (SS) from multi-site (MS) events in order to discriminate γ from $\beta\beta$ events, as discussed in Section 2.2.4. For example, the SS fraction of $0\nu\beta\beta$ decays is predicted in MC simulations to be $\sim 75\%$ whereas that of γ particles with $E_\gamma = Q$ is $\sim 15\%$. This implies that the purity of $\beta\beta$ -like components is increased in the SS data which improves the ability to identify candidate $0\nu\beta\beta$ decays in LB physics data. Moreover, the MS data helps to constrain the γ -like background components in the SS data since both spectra are connected by the SS fraction in Equations 6.4 and 6.5.

In previous analyses, different sets of observables \mathbf{y} were used to search for the $0\nu\beta\beta$ decay. The observables were chosen to reach the best sensitivity to the $0\nu\beta\beta$ half-life. The first search for $0\nu\beta\beta$ decay of EXO-200 in 2012 [100] was performed with a fit to the energy spectra only which were split into SS and MS events. An improved set of fit observables was used in the EXO-200 $0\nu\beta\beta$ search in 2014 [49] by using the energy E and the standoff distance r_{SD} as fit observables ($\mathbf{y} = (E, r_{\text{SD}})$). This configuration exploits additional knowledge about the spatial distributions of the individual model components. For example, the dominant γ background in the ROI of $0\nu\beta\beta$ decay originates from sources external to the xenon and thus occurs more frequently close to the material surfaces. In contrast, $2\nu\beta\beta$ and hypothetical $0\nu\beta\beta$ decays are distributed uniformly in the xenon volume. More details about the standoff distance are discussed in Section 5.2.2. Further improvement to the half-life sensitivity is achieved by not only exploiting spatial information in the fit but also detailed topological information beyond simple SS/MS classification. In the EXO-200 $0\nu\beta\beta$ search in 2018 [91], fit observables included the energy E and a BDT discriminator in SS events ($\mathbf{y}^{\text{SS}} = (E, \text{BDT})$, $\mathbf{y}^{\text{MS}} = E$). This BDT incorporates two topological variables and the standoff distance. More details on the BDT are discussed in Section 5.2.1. The topological variables discriminate signal $\beta\beta$ from background γ events based on the detailed signal shape. The BDT is not applied in MS data as toy studies indicated that using it there does not enhance the sensitivity [91]. The most recent EXO-200 $0\nu\beta\beta$ search in 2019 [14] includes three fit observables; energy E , standoff distance r_{SD} , and a DNN discriminator ($\mathbf{y} = (E, r_{\text{SD}}, \text{DNN})$). This fit concept is shown in Figure 6.1. The observables are used in both SS and MS data. Spatial and topological information is separated. Topological discrimination is achieved with the DNN discussed in Chapter 5. This analysis configuration proved to achieve the best $0\nu\beta\beta$ half-life sensitivity in toy studies among all configurations considered. The results of this analysis are discussed in Sections 6.2.1 and 6.2.2.

The $0\nu\beta\beta$ half-life sensitivities of analysis configurations where the traditional EXO-200 energy measurement E_{Trad} is replaced by the DNN energy E_{DNN} are discussed in Section 6.2.3. This DNN is introduced in Chapter 4.

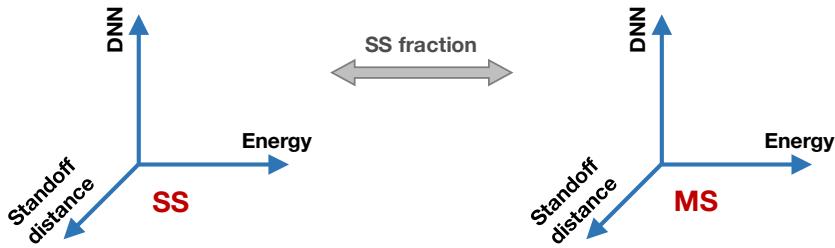


Figure 6.1: Fit configuration used in the EXO-200 $0\nu\beta\beta$ search in 2019 [14]. The observables energy, standoff distance, and a DNN discriminator are used in both SS and MS events while both event classes are connected by the SS fraction.

6.1.3 Background model

In MC simulations, an accurate model is necessary to understand the recorded LB physics data and to produce meaningful results. This is especially true for the analysis concept pursued by EXO-200 which relies on a maximum likelihood fit of the individual background components. It relies on precise modeling the spectral shape of individual background components. The model used to fit the LB physics data represents all significant components recorded in data. It comprises:

- $0\nu\beta\beta$: signal
- LXe decays
 - $2\nu\beta\beta$
 - ^{137}Xe
 - ^{135}Xe
- ^{232}Th decay chain
 - Near component: TPC vessel copper components
 - Remote component: inner cryostat
- ^{238}U -like (those presenting the ^{214}Bi line at 2.44 MeV)
 - Near component: TPC vessel copper components
 - Remote component: radon in the air gap
 - LXe component
 - ^{222}Rn decay chain in the inactive LXe
 - ^{214}Bi on the cathode surface
 - ^{214}Pb in the active LXe
- ^{60}Co decay
 - Near component: TPC vessel copper components
 - Remote component: source guide tube

- Other decays in the vessel copper components
 - ^{40}K
 - ^{65}Zn
- Neutron capture components
 - ^1H in coolant
 - ^{63}Cu and ^{65}Cu in the outer cryostat, inner cryostat, and TPC vessel
 - ^{136}Xe de-excitation γ

The energy threshold is 1 MeV for the most recent $0\nu\beta\beta$ search in 2019 [14]. In contrast to previous studies [51], ^{54}Mn is removed from the model since its γ line at 834 keV is more than 5σ below the energy threshold. The individual components are simulated using GEANT4 [63, 64]. The decay chains of ^{232}Th and ^{238}U are simulated assuming secular equilibrium.

Near components in the background model represent components arising from natural radioactive impurities in the TPC vessel. The vessel comprises the Cu vessel itself, but also materials in the TPC from the APD support materials, the wire plane support materials, the field rings, and the cathode plane materials, among others (see Figure 2.9). Remote components arise from either the inner cryostat, the outer cryostat, or the air gap between cryostat and lead wall which are shown in Figure 2.7. Additional potential remote components are highly correlated to the default ones used and are thus not included in the fit separately. This implies that the best-fit values for remote components also include contributions from those remote components. To account for this inaccuracy, the errors that arise from the exact choice of remote components are considered in the ML fit as systematic errors.

6.1.4 Systematic uncertainties

The $0\nu\beta\beta$ half-life sensitivity is affected by systematic uncertainties due to the low statistics nature of $0\nu\beta\beta$ searches. Systematic uncertainties which arise from imperfections in the MC simulation or detector understanding negatively affect the $0\nu\beta\beta$ sensitivity. They are incorporated as constraints in the maximum likelihood (ML) fit discussed in Section 6.1.1.

The $0\nu\beta\beta$ signal detection efficiency ϵ is raised in the EXO-200 $0\nu\beta\beta$ search in 2019 [14] to $(97.8 \pm 3.0)\%$ in Phase-I and to $(96.4 \pm 3.0)\%$ in Phase-II. The uncertainty on the signal detection efficiency σ_N is used to constrain the overall normalization N in the ML fit discussed in Section 6.1.1. The improvement in signal detection efficiency compared to previous searches is mostly caused by relaxing two event selection cuts. On the one hand, the minimum time required for subsequent events is relaxed from 1 s to 0.1 s. This event selection cut is introduced in order to reject time-correlated background events. On the other hand, events with a missing 3D position are not cut from the analysis if at least 60 % of the total event energy is reconstructed with a full 3D position. These events were cut in previous $0\nu\beta\beta$ searches. The relaxed 3D position selection cut only recovers MS events. Additional event selection cuts that did not change compared to previous $0\nu\beta\beta$ searches are discussed in Section 2.2.4.

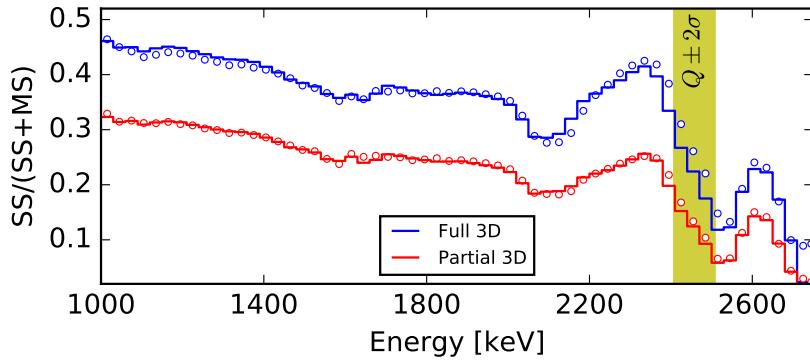


Figure 6.2: SS fraction, $SS/(SS+MS)$, as a function of DNN energy for source calibration data (circles) and MC simulations (lines) in Phase-II using a ^{228}Th source located near the cathode. The 3D cut criterion (“Full 3D”, blue) used in previous $0\nu\beta\beta$ searches [49, 91, 100] is compared to the relaxed criterion (“Partial 3D”, red) used in the $0\nu\beta\beta$ search in 2019 [14]. The ROI of $0\nu\beta\beta$ is highlighted in yellow.

The next paragraphs discuss different uncertainties that are considered in the ML fit. The resulting uncertainties are discussed for the fit configuration ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$) that is used in the $0\nu\beta\beta$ search in 2019 [14] and for fit configurations that are based on the DNN energy measurement E_{DNN} . The systematic uncertainties are evaluated with the events in the ROI around the Q value being blinded. This means that candidate $0\nu\beta\beta$ decays are not considered while improving the data analysis and evaluating systematic uncertainties. The statistical uncertainties on the systematic errors are not incorporated in the maximum likelihood (ML) fit and thus are omitted in this discussion.

SS fraction error

In previous analyses, all individual reconstructed charge clusters were required to have a reconstructed 3D position (“Full 3D”). Since the V-wires have an increased reconstruction threshold (~ 200 keV) compared to the U-wires (~ 90 keV) [14], some charge clusters are reconstructed well on the U-wires but lack corresponding V-wire signals. Thus, the z coordinate is reconstructed but the xy coordinates are missing. As of the $0\nu\beta\beta$ search in 2019 [14], the event selection criteria are relaxed. Now, only those events are removed from the analysis whose energy deposits with incomplete 3D position make up $>40\%$ of the total event energy while all other events are kept (“Partial 3D”). The SS fractions, $SS/(SS+MS)$, as a function of energy reconstructed by the DNN, E_{DNN} , with the full and partial 3D cut criteria are shown in Figure 6.2 for source calibration data of a ^{228}Th source located near the cathode. As the relaxed cut criterion only recovers MS events, the SS fraction is reduced. This relaxed cut criterion contributes to the increase in signal detection efficiency ϵ to nearly one. By applying a discriminator to MS events, as opposed to only SS events as in the $0\nu\beta\beta$ search in 2018 [91], the $0\nu\beta\beta$ half-life sensitivity improves from the additional information in MS events.

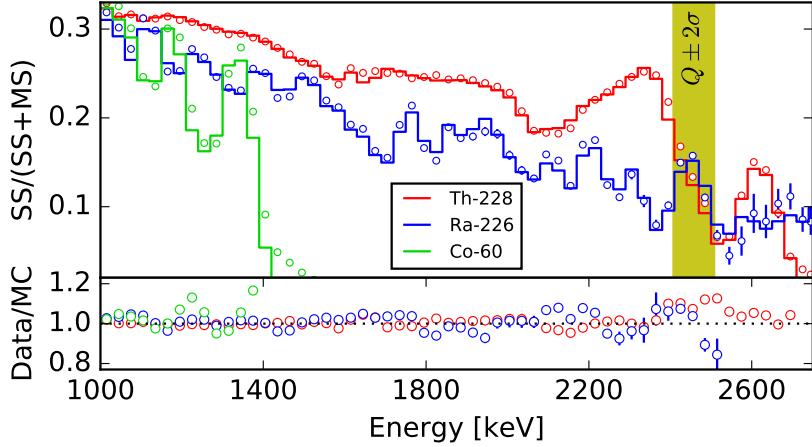


Figure 6.3: SS fraction, $SS/(SS+MS)$, as a function of DNN energy for calibration sources located near the cathode. Measured data (circles) and MC simulations (lines) are shown for ^{60}Co (green), ^{226}Ra (blue), and ^{228}Th (red) calibration sources in Phase-II. The ROI of $0\nu\beta\beta$ is highlighted in yellow. The bottom panel shows the ratios between measured data and MC simulations. Adapted from [14].

The error related to the SS fraction, $SS/(SS+MS)$, is estimated by comparing the predicted fraction in MC simulations to that observed in source calibration data. Figure 6.3 shows the agreement in source calibration data between measurement and MC simulations with the partial 3D cut criterion for the DNN energy measurement discussed in Chapter 4. The corresponding figure for the traditional EXO-200 reconstruction is shown in Ref. [14]. The error is determined over the full energy range of 1–3 MeV. It reflects the error weighted average of the absolute residuals between measured data and MC simulations as introduced in Equation 5.5. To account for spatial variations in the detector volume, the average error observed in ^{228}Th source calibration data at four different source locations is used. These contain locations near the anode planes (S2, S8) and two locations near the cathode (S5, S11) which are shown in Figure 2.13b.

The SS fraction error is 5.8% (4.6%) in Phase-I (Phase-II) for the partial 3D cut criterion with the traditional EXO-200 energy measurement. This observable is used in the $0\nu\beta\beta$ search in 2019 [14]. The corresponding errors for the DNN energy are shown in Table 6.1 for both 3D cut criteria. With the full 3D cut criterion, the SS fraction error for the DNN is 5.4% (5.2%) in Phase-I (Phase-II). With the partial 3D cut criterion, the error increases slightly to 6.2% (5.5%). The SS fraction error is incorporated in the ML fit as constraint (see Section 6.1). Comparing the SS fraction error of the DNN and the traditional EXO-200 energy, the DNN shows a slightly increased error. The reason for this could be the approach used for correcting for the finite electron lifetime which is only an approximation in MS events. More details are discussed in Section 4.3. To bypass this approximation, E_{DNN} could be only used in SS events, where the correction is applied properly, and E_{Trad} in MS events, where the correction can be applied for each charge cluster individually.

Table 6.1: SS fraction, $SS/(SS+MS)$, error for the DNN energy E_{DNN} from a ^{228}Th calibration source at different locations near the cathode (S5, S11) and near the anodes (S2, S8). Errors are shown for Phase-I and Phase-II and for the full and partial 3D cut criteria. The total error is the average error from different locations and is incorporated as constraint in the ML fit discussed in Section 6.1.1.

Location	ID	Full 3D [%]		Partial 3D [%]	
		Phase-I	Phase-II	Phase-I	Phase-II
Cathode	S5	3.5 ± 0.1	3.0 ± 0.1	3.2 ± 0.1	2.6 ± 0.1
	S11	5.6 ± 0.8	6.8 ± 0.3	5.1 ± 0.8	7.5 ± 0.2
Anode	S2	4.9 ± 0.7	5.8 ± 0.3	6.6 ± 0.7	5.7 ± 0.3
	S8	7.2 ± 1.2	5.2 ± 0.3	9.6 ± 1.1	6.1 ± 0.3
Total		5.4	5.2	6.2	5.5

Table 6.2: Locations of remote background components that are alternated in order to evaluate the impact from the exact choice of components to the ROI background.

Source	Default location		Alternative location
^{232}Th	inner cryostat	\longleftrightarrow	HFE
^{238}U	air gap	\longleftrightarrow	inner cryostat
^{60}Co	guide tube	\longleftrightarrow	inner cryostat

Background model error

The error arising from the choice of background components presented in Section 6.1.3 and their detailed locations is evaluated for remote components that contribute to the ROI. The error is estimated by alternating the components from the default locations to their most-plausible alternatives as shown in Table 6.2. The energy distributions of these background components are shown in Figure 6.4. The expected number of events in the ROI from the fit to the LB physics data blinded around the ROI of the alternative model is compared to the corresponding fit with the default fit model. The relative changes in the numbers of ROI events from alternating individual components are added in quadrature. This systematic error from the choice of the background model is considered as contribution to the ROI normalization error discussed in Section 6.1.1.

The background model error for the fit configuration ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$) used in the $0\nu\beta\beta$ search in 2019 [14] with the traditional EXO-200 energy E_{Trad} is 4.0% (4.6%) in Phase-I (Phase-II). The corresponding background model errors using the DNN energy E_{DNN} are shown in Table 6.2 for both phases of EXO-200 operation. For the fit configurations $\mathbf{y} = E_{\text{DNN}}$, the background model error is 6.2% (4.6%) in Phase-I (Phase-II). For the fit configuration $\mathbf{y} = (E_{\text{DNN}}, r_{\text{SD}}, \text{DNN})$, the corresponding error is 4.1% (4.1%) in Phase-I (Phase-II). As

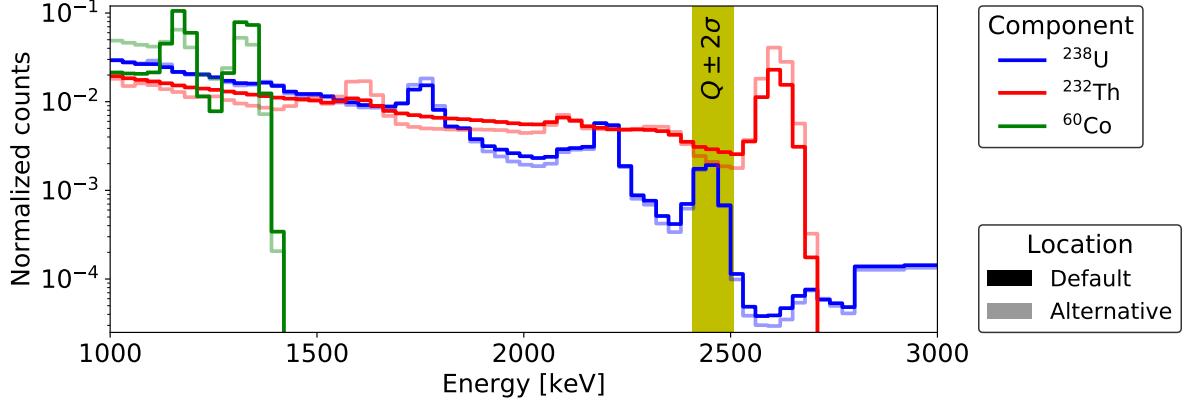


Figure 6.4: Energy spectra of remote background components contributing to the ROI (^{232}Th : red, ^{238}U : blue, ^{60}Co : green) from the default (dark) and their most plausible alternative (light) locations. The ROI of $0\nu\beta\beta$ is highlighted in yellow. The locations used for the individual components are shown in Table 6.2.

Table 6.3: Background model errors for different fit configurations using the DNN energy measurement E_{DNN} in Phase-I and Phase-II. The errors arise from the choices of location for far background components that are not known precisely. The total error is the quadratic sum of the individual components and contributes to the ROI normalization error in the ML fit discussed in Section 6.1.1.

Source	1D Fit [%]		3D Fit [%]	
	Phase-I	Phase-II	Phase-I	Phase-II
^{238}U	1.2	1.0	1.1	0.7
^{232}Th	6.1	4.5	3.8	4.0
^{60}Co	0.1	0.0	1.3	0.2
Total	6.2	4.6	4.1	4.1

expected from Figure 6.4, the error arising from the location of the remote ^{232}Th component dominates the background model error in all configurations. On the other hand, the choice of the ^{60}Co far location only has a negligible impact on the number of ROI events since very few events contribute to the ROI. Comparing the energy measurements E_{DNN} and E_{Trad} , the background model errors are on the same level. This confirms the robustness of the DNN energy measurement.

Spectral shape error

The discrepancies in spectral agreement between measured data and MC simulations contribute to the ROI normalization error. In MC simulations, a large number (~ 1000) of toy datasets is drawn from the best-fit background model. The toy data is generated after correcting the PDFs by the residual differences between measured data and MC simulations. $\beta\beta$ -like SS toy data is corrected by the spectral disagreement observed in background-subtracted $2\nu\beta\beta$ data while γ -like toy data is corrected by the discrepancy observed in source calibration data. Details on the background-subtracted $2\nu\beta\beta$ data are discussed in Section 5.3. Each toy data realization is then fitted with the nominal PDFs. The relative bias between the drawn and the best-fit numbers of events in the ROI determines the spectral shape error. For fit configurations that use multiple observables, the error is evaluated for each observable individually. To determine the spectral shape error, the biases in ROI counts observed in each observable are added in quadrature. To further consider spatial differences, the error is evaluated on calibration data from sources located near the cathode (S5) and near the anode (S2). The final spectral shape error is estimated as the average error over both source locations.

The spectral agreement using the DNN energy measurement E_{DNN} is shown in Figure 6.5 for different calibration sources located near the cathode outside the TPC. The spectral agreements between measured data and MC simulations of the DNN-Recon discriminator variable (see Chapter 5) and the standoff distance are shown in Figure 6.6. It shows γ -like events from a ^{226}Ra source located near the cathode outside the TPC and background-subtracted $2\nu\beta\beta$ data. In addition, the distributions of $0\nu\beta\beta$ decay in MC simulations are shown. The spectral shape error for the fit configuration ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$) used in the $0\nu\beta\beta$ search in 2019 [14] with the traditional EXO-200 energy measurement E_{Trad} is 5.8% (4.4%) in Phase-I (Phase-II). The errors with the DNN energy measurement, E_{DNN} , are shown in Table 6.4 for both phases of EXO-200 operation. The errors given for “Placebo” are calculated from toy data drawn from nominal PDFs, i.e. without correcting by spectral differences of measured data and MC simulations. This error arises from fluctuations in MC simulations and from the low statistics nature of the $0\nu\beta\beta$ search. The errors are dominated by the error observed in the energy observable and increase as the source position approaches the anode. For the 1D fit configuration ($\mathbf{y} = E_{\text{DNN}}$), the spectral shape error is 2.3% (3.4%) in Phase-I (Phase-II). For the 3D fit configuration ($\mathbf{y} = (E_{\text{DNN}}, r_{\text{SD}}, \text{DNN})$), it is 4.9% (5.9%) in Phase-I (Phase-II). Comparing the measurements E_{DNN} and E_{Trad} , the spectral shape error is slightly improved (degraded) in Phase-I (Phase-II). Since two thirds of the observables are identical in the 3D fit configuration, this small change in spectral shape error further confirms the robustness of the DNN energy measurement.

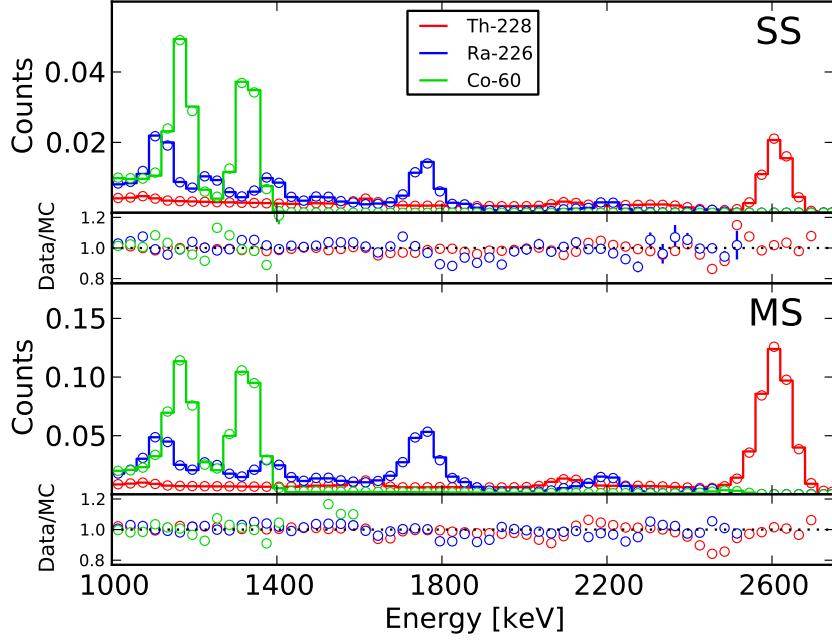


Figure 6.5: Spectral agreement in E_{DNN} for measured calibration data (circles) and MC simulations (lines) in Phase-II using ^{60}Co (green), ^{226}Ra (blue), and ^{228}Th (red) sources located near the cathode for SS (top half) and MS (bottom half) events. The bottom panels show the ratios between measured data and MC simulations. Adapted from [14].

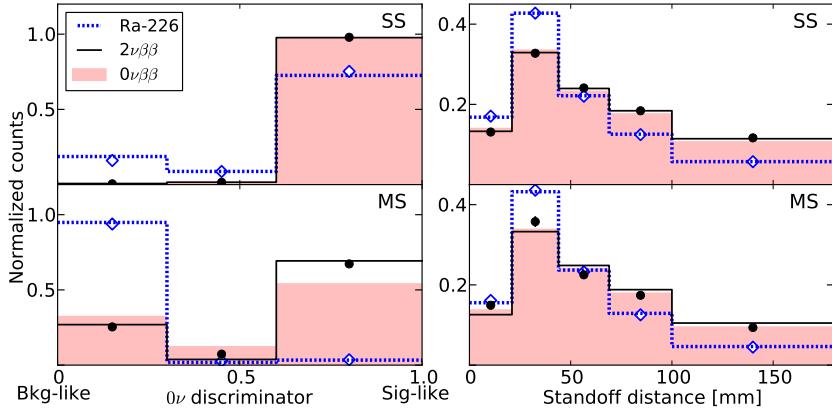


Figure 6.6: Distributions of measured data (circles) and MC simulations (lines) in Phase-II for the DNN-Recon variable (left) and the standoff distance (right) for SS (top) and MS (bottom) events. Data from a ^{226}Ra (blue) calibration source located near the cathode represents γ -like events and background-subtracted $2\nu\beta\beta$ data (black) represents $\beta\beta$ -like events. The distributions of $0\nu\beta\beta$ decay predicted in MC simulations are shown in red. Taken from [14].

Table 6.4: Spectral shape errors for different fit configurations using the DNN energy measurement E_{DNN} in Phase-I and Phase-II. The errors arise from differences in measured data and MC simulations that are evaluated for each fit observable individually and added in quadrature. The total error represents the average of the errors observed from calibration sources located near the cathode and anode. It contributes to the ROI normalization error in the ML fit discussed in Section 6.1. “Placebo” is discussed in the text and does not contribute to the total error.

Source	1D Fit [%]		3D Fit [%]	
	Phase-I	Phase-II	Phase-I	Phase-II
Placebo	0.3	1.4	0.2	2.9
Cathode	0.1	3.6	2.8	4.6
Anode	4.4	3.2	7.0	7.1
Total	2.3	3.4	4.9	5.9

For the $0\nu\beta\beta$ search in 2019 [14], DNN-Recon is used as topological discriminator. Even though DNN-Raw suggests an improved discrimination power as discussed in Section 5.2.1, the spectral agreement between measured data and MC simulations is degraded for this DNN concept (see Section 5.3). In an analysis configuration of 10 bins uniform in the discriminator variable, the resulting spectral shape error is found to be $\sim 30\%$ for DNN-Raw and $\sim 10\%$ for DNN-Recon in Phase-II. Since this error is considered in the ML fit, the resulting $0\nu\beta\beta$ half-life sensitivity is lower for DNN-Raw even though DNN-Raw outperforms DNN-Recon in terms of discrimination power. Moreover, the extent of spectral disagreement for DNN-Raw leads to unstable ML fit results. The spectral shape error can be mitigated by profiling the DNN-Recon variable between 0, 0.3, 0.6 and 1 at the cost of discrimination power. While toy studies indicated no significant loss in discrimination power for this fit configuration, the spectral shape error decreases from $\sim 10\%$ to 4.4% in Phase-II. Finally, the spectral error proved to be the critical argument to use DNN-Recon over DNN-Raw for the $0\nu\beta\beta$ search in 2019 [14] in a fit configuration with three bins.

Table 6.5: Summary of systematic error contributions. The 1D fit configuration represents $\mathbf{y} = E_{\text{DNN}}$ and the 3D fit configuration represents $\mathbf{y} = (E_{\text{X}}, r_{\text{SD}}, \text{DNN})$ with E_{X} denoting either E_{DNN} or E_{Trad} . The errors for E_{Trad} in the 3D fit configuration are taken from Ref. [14]. The indented errors are the individual components of the ROI normalization error which are added in quadrature.

Error ↓	Fit configuration →	Phase-I [%]			Phase-II [%]		
		E_{DNN}		E_{Trad}	E_{DNN}		E_{Trad}
		1D	3D	3D	1D	3D	3D
SS fraction		6.2		5.8		5.5	4.6
Overall normalization			3.1	3.1		3.1	3.1
ROI normalization		6.8	6.6	7.1		5.8	7.2
Spectral shape agreement		2.3	4.9	5.8		3.4	5.9
Background model		6.2	4.1	4.0		4.6	4.1
Other [91]			1.5	1.5		1.2	1.2

Summary

The individual contributions to the error on the ROI normalization and on the overall normalization are summarized in Table 6.5 together with the uncertainty on the SS fraction. The uncertainties for the fit configuration ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$) used in the $0\nu\beta\beta$ search in 2019 are taken from Ref. [14]. The uncertainties evaluated for the DNN energy, E_{DNN} , are given for the 1D and 3D fit configurations. The contribution “Other” to the ROI normalization error is not re-evaluated for the EXO-200 $0\nu\beta\beta$ search in 2019 [14] but is taken from that in 2018 [91] since it is unchanged. It accounts for spatial and temporal variations in energy resolution. Similar to the spectral shape error, these variations are used to generate toy data to estimate the error related to energy resolution. The corresponding error is 1.5% (1.2%) for Phase-I (Phase-II). The improvement in Phase-II arises from a lower variation in the energy resolution in time as is shown in Figure 4.21. This contribution is not re-evaluated for E_{DNN} since its variations are observed to be on a similar level as those of E_{Trad} as is discussed in Section 4.3.

Overall, for the 3D fit configuration, there is no overly increase in the error on the ROI normalization. This is achieved by reducing the impact of discrepancies between measured data and MC simulations in the DNN discriminator variable by profiling it to three bins. At the same time, the error from the choice of background model is reduced when exploiting additional spatial and topological information in the ML fit. Comparing both energy measurements, E_{DNN} and E_{Trad} , the resulting uncertainties for the 3D fit configuration are on a similar level. The evaluation of the systematic uncertainties has shown that using the DNN energy measurement does not impair the systematic uncertainties significantly. This implies that the improvements in energy resolution with the DNN energy can translate to improvements in $0\nu\beta\beta$ half-life sensitivity.

6.2 Results

6.2.1 Published $0\nu\beta\beta$ half-life sensitivity

The search for $0\nu\beta\beta$ decay represents a low statistics search that is prone to statistical fluctuations of the background. Thus, a $0\nu\beta\beta$ half-life limit strongly depends on the particular realization in nature of background in the ROI and is therefore not the optimum metric to compare the performance of different analysis configurations and experiments. The $0\nu\beta\beta$ half-life sensitivity is a more robust metric. It represents the median $0\nu\beta\beta$ half-life limit when performing the experiment in many statistical realizations in MC simulations. The $0\nu\beta\beta$ half-life (limit) at the 90 % confidence level (CL) is given by:

$$T_{1/2}^{0\nu} = \ln(2) \frac{N_A/M \cdot f \cdot m \cdot L}{n/\epsilon} , \quad (6.10)$$

where $N_A = 6.022 \cdot 10^{23} \text{ mol}^{-1}$ is the Avogadro number, $M = 0.135\,514 \text{ kg mol}^{-1}$ is the ^{136}Xe molar mass [49], $f = 80.6\%$ is the enrichment fraction [49] of the LXe in EXO-200 with a mass in the fiducial volume of $m = 92.40 \text{ kg}$ [49]. L is the detector livetime, n is the best-fit (upper limit) number of $0\nu\beta\beta$ decays, and ϵ is the signal detection efficiency in the fiducial volume. After quality cuts, the detector livetime L in the fiducial volume is 117.4 kg yr (116.7 kg yr) in Phase-I (Phase-II). The signal detection efficiency ϵ is $(97.8 \pm 3.0)\%$ ($(96.4 \pm 3.0)\%$).

To estimate the $0\nu\beta\beta$ half-life sensitivity, the experiment is performed ~ 1000 times in MC simulations. Toy data is generated using the best-fit background model from a fit that uses energy as the only observable ($\mathbf{y} = E_{\text{Trad}}$) while the ROI is blinded. To allow a fair comparison between different analysis configurations (see Section 6.1.2), the identical background model from the energy only fit configuration is used for all potential analysis configurations. For each experiment in MC simulations, the true number of $0\nu\beta\beta$ decays is profiled from $0 - 50$ which is sufficient to cover the 90 % CL. The profiles of representative pseudo-experiments are shown for the 1D and 3D fit configurations with E_{Trad} in Figure 6.7a in Phase-II. The limits at the 1σ CL and 90 % CL represent the intersections of the profile-likelihood curve with $\Delta\mathcal{L} = 0.5$ and $\Delta\mathcal{L} \simeq 1.35$, respectively. $\Delta\mathcal{L}$ is defined as:

$$\Delta\mathcal{L} = \mathcal{L} - \mathcal{L}_{\text{best}} , \quad (6.11)$$

where $\mathcal{L}_{\text{best}}$ denotes the best-fit negative log-likelihood. The improvement of the 3D fit configuration is reflected by a decreased upper limit n on the number of $0\nu\beta\beta$ decays. To account for different livetimes L and signal detection efficiencies ϵ in both phases of EXO-200 operation, the phases are combined by adding the profiles $n/\epsilon L$, instead of only the profiles n . Each pseudo-experiment in Phase-I is randomly assigned to one of Phase-II. A representative example is shown in Figure 6.7b.

The $0\nu\beta\beta$ half-life sensitivity is evaluated using the upper limit $n/\epsilon L$ at the 90 % CL in Equation 6.10. The resulting sensitivities from pseudo-experiments are shown in Figure 6.8 for the baseline 1D ($\mathbf{y} = E_{\text{Trad}}$, green) and the 3D fit configurations ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$, blue) for both phases of EXO-200 operation and for their combination. Using more information in the ML fit in addition to the energy observable improves the sensitivity. Analogously, the

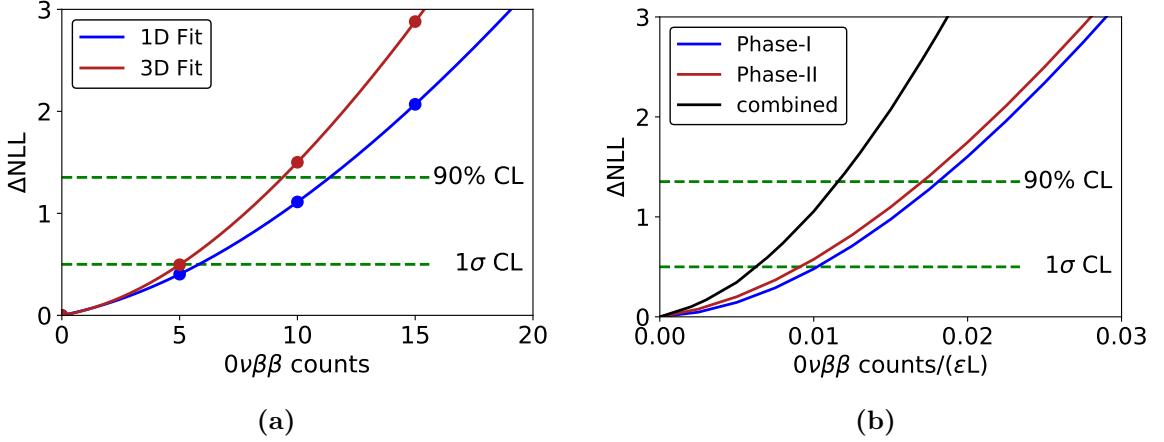


Figure 6.7: The profile-likelihood for $0\nu\beta\beta$ counts of representative pseudo-experiments in MC simulations from the baseline 1D ($\mathbf{y} = E_{\text{Trad}}$, blue) and the 3D ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$, red) fit configuration is shown in (a). The points represent fits to the MC toy data and the lines represent interpolations between the points. The corresponding curves scaled by livetime L and signal detection efficiency ϵ for the 3D fit configuration in Phase-I (blue) and Phase-II (red) are shown in (b). The combined curve of Phase-I+II (black) is obtained by adding the scaled profiles of both phases. The horizontal dashed lines (green) represent the 1σ CL and 90 % CL.

Table 6.6: Summary of median $0\nu\beta\beta$ half-life sensitivities at the 90 % CL for individual phases and their combination on profile level and for the baseline 1D ($\mathbf{y} = E_{\text{Trad}}$) and the 3D ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$) fit configurations. All numbers in 10^{25} yr.

Fit	Phase-I	Phase-II	Phase-I+II
1D	2.5	2.7	3.9
3D	3.3	3.4	5.0

median $0\nu\beta\beta$ half-life sensitivity is evaluated using the median upper limit $n/\epsilon L$ at the 90 % CL. The resulting median sensitivities are shown in Table 6.6 for the baseline 1D ($\mathbf{y} = E_{\text{Trad}}$) and the 3D ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$) fit configurations. The combined median half-life sensitivity on the $0\nu\beta\beta$ decay using the DNN discriminator is $5.0 \cdot 10^{25}$ yr with the complete dataset of EXO-200 operation. Using additional event information besides energy and ss/ms discrimination with the standoff distance r_{SD} and the DNN discriminator leads to a $\sim 25\%$ improvement in sensitivity over the baseline analysis configuration [14]. This increase proves the strong potential of a multi-parameter analysis as pursued by EXO-200. This fit configuration outperforms the one used in the previous EXO-200 $0\nu\beta\beta$ search in 2108 [91] with $\mathbf{y}^{\text{SS}} = (E_{\text{Trad}}, \text{BDT})$ and $\mathbf{y}^{\text{MS}} = E_{\text{Trad}}$ and achieves a 10 % improvement in $0\nu\beta\beta$ half-life sensitivity.

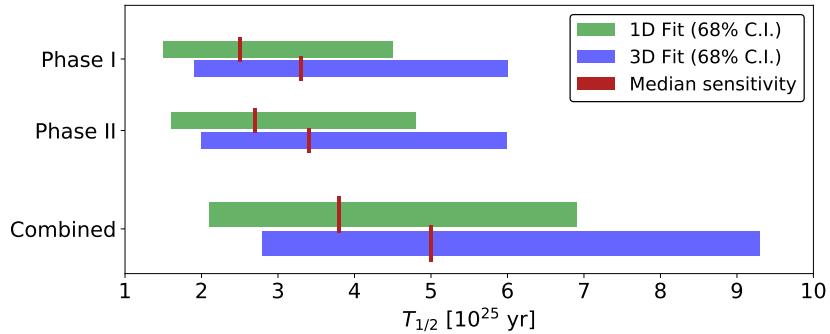


Figure 6.8: $0\nu\beta\beta$ half-life sensitivities at the 90 % CL from pseudo-experiments in MC simulations for Phase-I (top), Phase-II (center) and their combination (bottom). The median sensitivities are shown in red. The 68 % confidence intervals (CI) are shown for the baseline 1D ($\mathbf{y} = E_{\text{Trad}}$, green) and the 3D ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$, blue) [14] fit configurations.

6.2.2 Published $0\nu\beta\beta$ half-life data limit

The events in the ROI are un-blinded after evaluating the systematic uncertainties (see Section 6.1.4) and the median $0\nu\beta\beta$ half-life sensitivities of different potential analysis configurations considered for this $0\nu\beta\beta$ search (see Section 6.2.1). ML fits (see Equation 6.1) are performed to the complete datasets of Phase-I and Phase-II separately including all systematic uncertainties. The primary analysis in the EXO-200 $0\nu\beta\beta$ search in 2019 [14] is the 3D fit configuration with $\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$ selected based on the best median $0\nu\beta\beta$ half-life sensitivity. The best-fit results are shown in Figure 6.9.

The figure shows the low-background (LB) physics data (points) and best-fit (blue) for Phase-I and Phase-II separately. Dashed lines represent different background contributions, the filled area the $2\nu\beta\beta$ decay, and the purple line the $0\nu\beta\beta$ decay. The projections to the energy observable in SS events are shown in the left panels for Phase-I (top) and Phase-II (right). Studentized residuals between LB physics data and best-fit model are shown beneath the SS energy spectra. The projections of residuals are shown which typically follow normal distributions. Deviations that affect the ROI are taken into account by the systematic uncertainties discussed in Section 6.1.4. The width of the energy bins is 15 keV below 2800 keV and 30 keV above. The insets in the left panels show the region near the ROI in SS events where the $Q \pm 2\sigma$ region is highlighted in red. The bottom right panels show the energy spectra in MS events in Phase-I and Phase-II. The top right panels show projections to the DNN-Recon discriminator for events near the ROI (2395–2530 keV) in SS (left) and MS (right) events for Phase-I (top) and Phase-II (bottom). The best-fit numbers of background events in the ROI, $Q \pm 2\sigma$, are 32.3 ± 2.3 in Phase-I and 30.9 ± 2.4 in Phase-II. This corresponds to background rates in the fiducial LXe of $(1.7 \pm 0.2) \cdot 10^{-3} \text{ kg}^{-1} \text{ yr}^{-1} \text{ keV}^{-1}$ in Phase-I and $(1.9 \pm 0.2) \cdot 10^{-3} \text{ kg}^{-1} \text{ yr}^{-1} \text{ keV}^{-1}$ in Phase-II. In LB physics data, 39 (26) events were recorded in Phase-I (Phase-II) in this energy range. This implies an over-fluctuation of events in the ROI in Phase-I, while in Phase-II, there is a slight under-fluctuation of events [14].

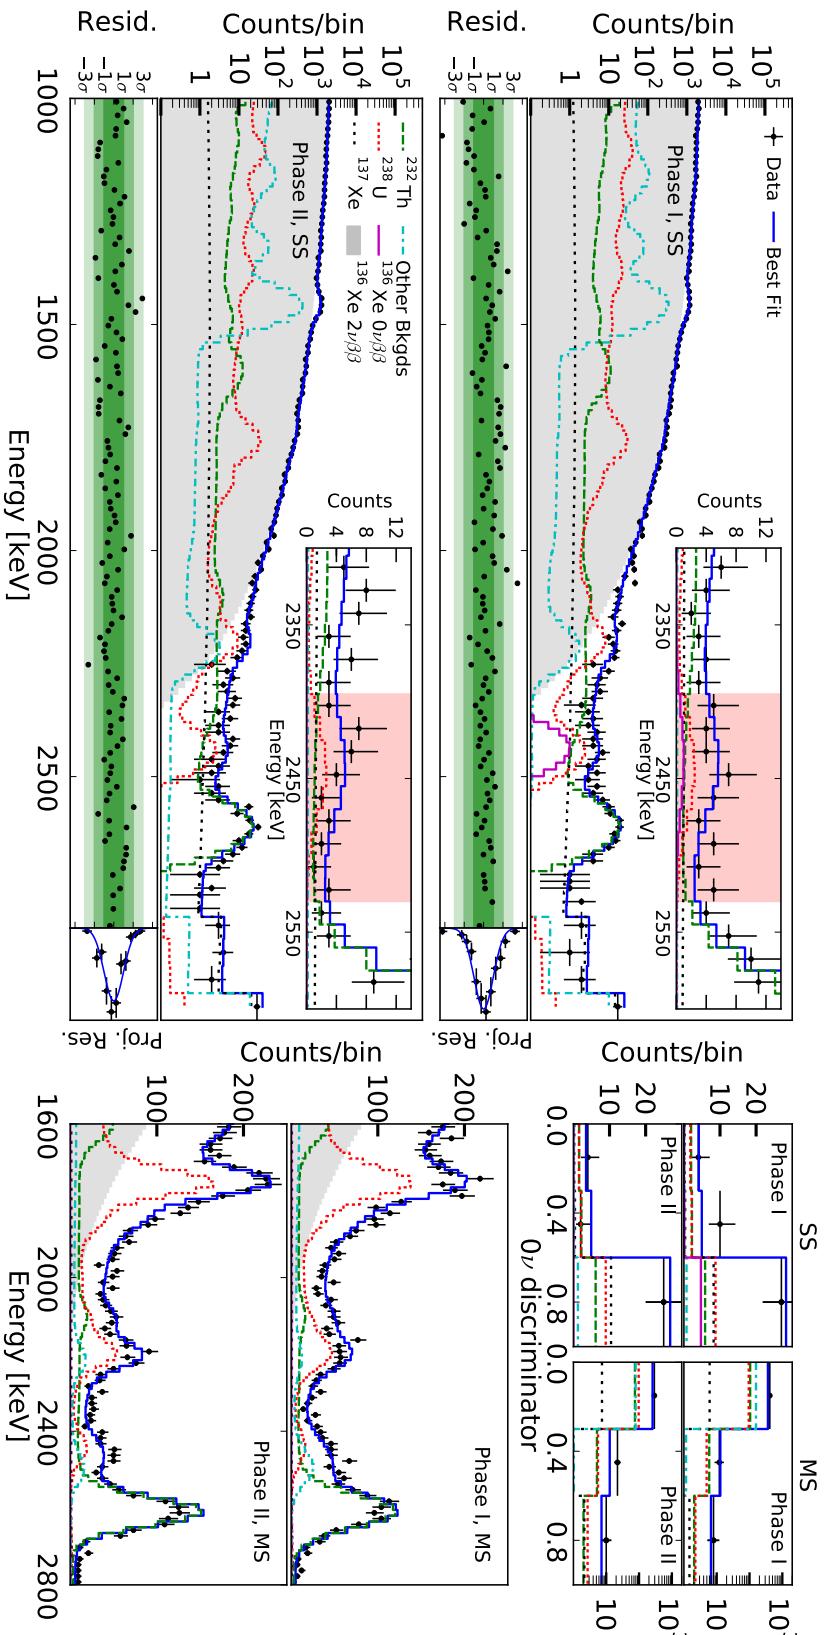


Figure 6.9: Best-fit model (blue) from the 3D fit configuration ($\mathbf{y} = (E_{\text{Trad}}, r_{\text{SD}}, \text{DNN})$) to low-background physics data (points). Individual background components are shown by dashed lines and $2\nu\beta\beta$ decays by filled area. Projections to the energy observable in SS events (left) are shown for Phase-I (top) and Phase-II (bottom). The residuals of best-fit model and LB physics data are shown below. The insets show the regions near the ROI of $0\nu\beta\beta$ decay which is highlighted in red. The best fits in MS events are shown on the bottom right panels in Phase-I (top) and Phase-II (bottom). Projections to the DNN discriminator for events near the ROI are shown in the top right panels separated in SS (left) and MS (right) events in Phase-I (top) and Phase-II (bottom). Taken from [14].

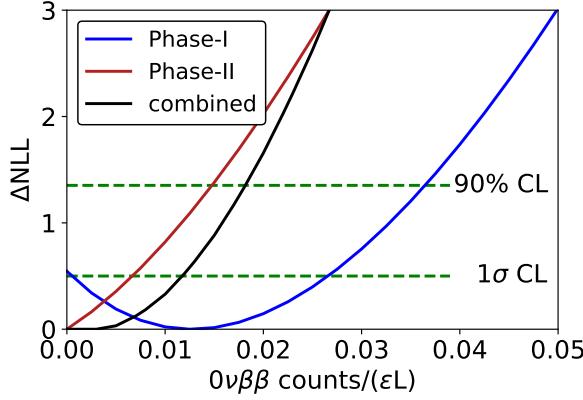


Figure 6.10: Profile-likelihood curves for the number of $0\nu\beta\beta$ decays from the fit to LB physics data (see Figure 6.9). The number of events is scaled by livetime L and signal detection efficiency ϵ for Phase-I (blue) and Phase-II (red). The combined curve of Phase-I+II (black) is obtained by adding the scaled profiles of both phases. The horizontal dashed lines (green) represent the 1σ CL and 90 % CL.

With the completed dataset of EXO-200, no statistically significant evidence for $0\nu\beta\beta$ decay is observed. Due to the over-fluctuation of ROI events in Phase-I, the best-fit number of $0\nu\beta\beta$ decays exceeds zero (7.1), which however is not significant. The best fit in Phase-II is consistent with zero $0\nu\beta\beta$ decays. The lower limit on the $0\nu\beta\beta$ half-life at the 90 % CL is evaluated by a profile likelihood scan that is shown in Figure 6.10. The lower limit on the $0\nu\beta\beta$ half-life at the 90 % CL is $T_{1/2}^{0\nu} > 1.7 \cdot 10^{25}$ yr for Phase-I and $T_{1/2}^{0\nu} > 4.3 \cdot 10^{25}$ yr for Phase-II. When combining both phases, a lower limit of $T_{1/2}^{0\nu} > 3.5 \cdot 10^{25}$ yr at the 90 % CL can be placed on the half-life of $0\nu\beta\beta$ decay in ^{136}Xe .

Using the nuclear matrix elements of Refs. [102–106], the phase space factor of Ref. [107], and Equation 2.8, the combined lower limit $T_{1/2}^{0\nu}$ corresponds to an upper limit on the Majorana neutrino mass of $m_{\beta\beta} < 93\text{--}286$ meV. The range of upper limits arises from different nuclear matrix element calculations.

The evolution of all EXO-200 $0\nu\beta\beta$ search results so far is shown in Figure 6.11. The lower limit on the $0\nu\beta\beta$ half-life falls below the median sensitivity after the first $0\nu\beta\beta$ search in 2012 [100]. This is due to an over-fluctuation of events in the ROI while accumulating additional exposure in subsequent searches. The projected evolution of the half-life sensitivity is shown based on the first $0\nu\beta\beta$ search. The sensitivity is expected to improve from an increased exposure via $\propto \sqrt{Mt}$ as is introduced in Equation 2.11. The subsequent $0\nu\beta\beta$ searches exceed this projection due to improvements to data analysis. This illustrates the potential of multi-parameter searches over simple counting experiments.

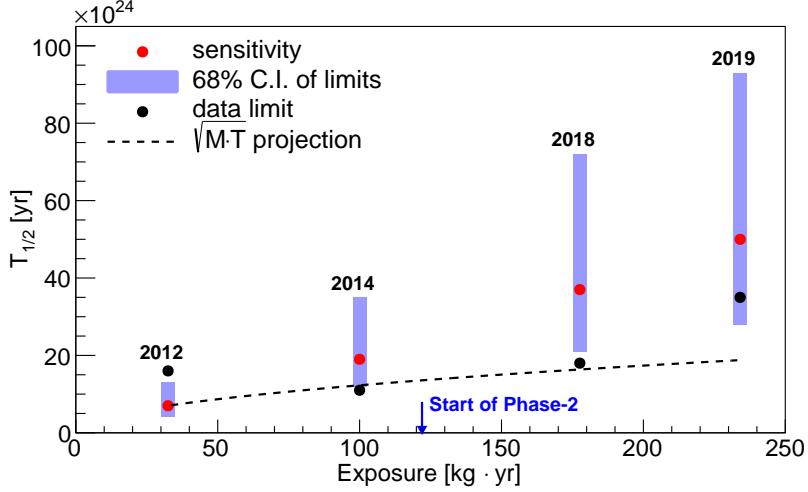


Figure 6.11: The evolution of EXO-200 $0\nu\beta\beta$ search results. The median $0\nu\beta\beta$ half-life sensitivities (red) and the 68 % confidence intervals (blue) are shown. The corresponding data limits are shown in black. The EXO-200 $0\nu\beta\beta$ searches include those in 2012 [100], 2014 [49], 2018 [91] and the one discussed here in 2019 [14]. The projection of half-life sensitivity from increased exposure only via $\propto \sqrt{Mt}$ is shown by dashed line (see Equation 2.11). Taken from [54].

6.2.3 $0\nu\beta\beta$ half-life sensitivity with the DNN energy

Since the ROI events are blinded for analysis, the half-life limit on $0\nu\beta\beta$ decay cannot be evaluated for the DNN energy measurement E_{DNN} which is discussed in Chapter 4. These events are only un-blinded for official $0\nu\beta\beta$ searches of the entire collaboration. ML fits (see Equation 6.1) are performed to the complete blinded datasets of Phase-I and Phase-II separately for the 1D ($\mathbf{y} = E_{\text{DNN}}$) and the 3D ($\mathbf{y} = (E_{\text{DNN}}, r_{\text{SD}}, \text{DNN})$) fit configurations. The ML fits include all systematic uncertainties discussed in Section 6.1.4. The best-fit results for the 3D fit configuration are shown in Figure 6.12. The figure shows the LB physics data (points) and best-fit model (blue) in SS events for Phase-I (top half) and Phase-II (bottom half). Dashed lines represent different background contributions and the filled area represents the $2\nu\beta\beta$ decay. The $0\nu\beta\beta$ decay is excluded from the fit model (see Section 6.1.3), since the SS events near the Q value of the $0\nu\beta\beta$ decay are blinded in the region of 2347–2569 keV. The left panels show the projections to the DNN energy observable. Studentized residuals between LB physics data and best-fit model are shown beneath the corresponding energy spectra. The projections of residuals are shown which typically follow normal distributions. The width of the energy bins is 15 keV below 2800 keV and 30 keV above. The right panels show projections to the DNN-Recon discriminator in SS events.

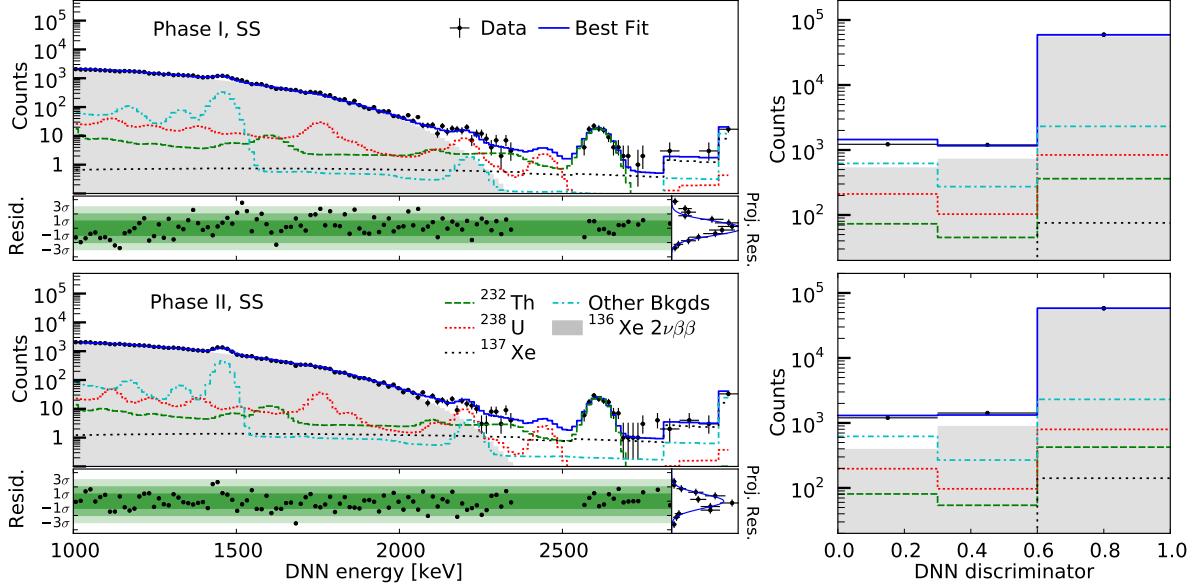


Figure 6.12: Best-fit model (blue) from the 3D fit configuration ($\mathbf{y} = (E_{\text{DNN}}, r_{\text{SD}}, \text{DNN})$) to LB physics data (points) for Phase-I (top half) and Phase-II (bottom half). Individual background components are shown by dashed lines and $2\nu\beta\beta$ decays by filled area. The left panels show projections to the energy observable in SS events while the residuals of best-fit model and LB physics data are shown below. The right panels show projections to the DNN discriminator observable in SS events. Since the region around the Q value of $0\nu\beta\beta$ decay is blinded for analysis, no data from that region is included in the ML fit.

The $0\nu\beta\beta$ half-life sensitivity is estimated for fit configurations that use the DNN energy measurement E_{DNN} . The procedure to determine the half-life sensitivity for these configurations follows exactly the one discussed in Section 6.2.1. The half-life sensitivity is evaluated using pseudo-experiments in MC simulations where the systematic uncertainties are considered that are discussed in Section 6.1.4. Analogously to the comparison of the 1D and 3D fit configurations based on E_{Trad} , here, the identical best-fit background model is used. It is determined in a 1D fit configuration with $\mathbf{y} = E_{\text{Trad}}$ while the ROI is blinded.

The $0\nu\beta\beta$ half-life sensitivities at the 90 % CL are shown in Figure 6.13 for both phases of EXO-200 operation separately and for their combination. The results based on the traditional EXO-200 energy measurement E_{Trad} are those also shown in Figure 6.8. The resulting median $0\nu\beta\beta$ half-life sensitivities for analyses using the DNN energy measurement E_{DNN} are shown in Table 6.7. The corresponding results with the traditional EXO-200 energy measurement E_{Trad} are given in Table 6.6. The 1D fit configuration with E_{DNN} achieves a $\sim 10\%$ improvement in median $0\nu\beta\beta$ half-life sensitivity over the corresponding fit configuration with E_{Trad} . For the 3D fit configuration, the corresponding improvement is $\sim 8\%$. The best analysis configuration with $\mathbf{y} = (E_{\text{DNN}}, r_{\text{SD}}, \text{DNN})$ achieves a $\sim 35\%$ improvement in median $0\nu\beta\beta$ half-life sensitivity over the baseline analysis with $\mathbf{y} = E_{\text{Trad}}$.

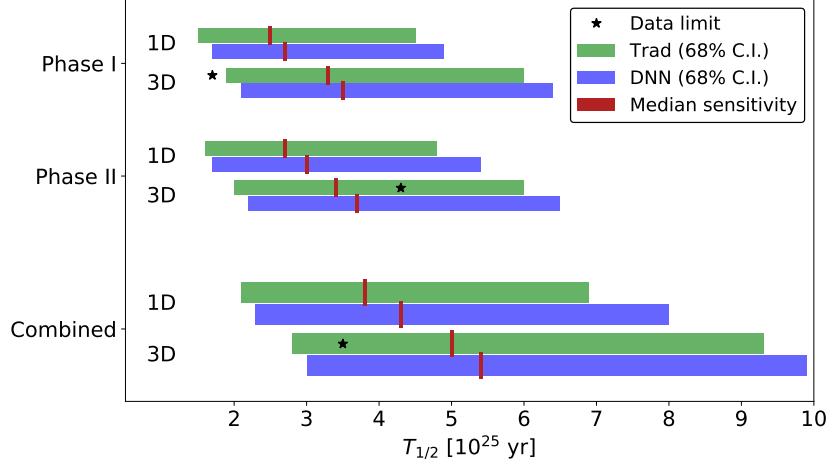


Figure 6.13: $0\nu\beta\beta$ half-life sensitivities at the 90 % CL are shown from pseudo-experiments in MC simulations for Phase-I (top), Phase-II (center) and their combination (bottom). The median half-life sensitivities are given in red and the data limits, where available from Ref. [14], in black. Half-life sensitivities for 1D ($\mathbf{y} = E_X$) and 3D ($\mathbf{y} = (E_X, r_{SD}, DNN)$) fit configurations are shown where E_X denotes either E_{DNN} or E_{Trad} . The bands reflect the 68 % CI for analyses with E_{Trad} (green) and with E_{DNN} (blue), respectively.

Table 6.7: Summary of median $0\nu\beta\beta$ half-life sensitivities at the 90 % CL for individual phases and their combination on profile level and for the 1D ($\mathbf{y} = E_{DNN}$) and 3D ($\mathbf{y} = (E_{DNN}, r_{SD}, DNN)$) fit configurations. All numbers in 10^{25} yr. The corresponding results for analysis configurations with E_{Trad} are shown in Table 6.6.

Fit	Phase-I	Phase-II	Phase-I+II
1D	2.7	3.0	4.3
3D	3.5	3.7	5.4

6.3 Summary

This chapter discussed the analysis concept pursued by EXO-200 for $0\nu\beta\beta$ searches in ^{136}Xe which is based on a maximum likelihood (ML) approach. The energy spectra in SS and MS events are fitted simultaneously while being linked to each other by the SS fraction. Additional fit observables besides energy can be used to exploit more information in the ML fit. These may include spatial or detailed topological information. Several systematic uncertainties are considered in the ML fit where the dominant contributors are re-evaluated for the analysis configurations discussed. This multi-parameter analysis approach is superior to simple counting experiments.

The deep neural network (DNN) discriminator discussed in Chapter 5, in particular DNN-Recon, is used in the EXO-200 $0\nu\beta\beta$ search in 2019 [14] as a topological discriminator of $\beta\beta$ decays and background γ interactions. In addition, the standoff distance is used to incorporate spatial information in the ML fit. Together with energy, this 3D fit approach provides the best sensitivity among all configurations considered, leading to a median $0\nu\beta\beta$ half-life sensitivity of $5.0 \cdot 10^{25}$ yr at the 90 % CL. The corresponding lower limit on the half-life of $0\nu\beta\beta$ in ^{136}Xe is $T_{1/2}^{0\nu} > 3.5 \cdot 10^{25}$ yr at the 90 % CL. Compared to the baseline analysis using only the energy, this represents an improvement by $\sim 25\%$ in $0\nu\beta\beta$ half-life sensitivity. This illustrates the high potential of the multi-parameter analysis concept pursued by EXO-200 for $0\nu\beta\beta$ searches.

The DNN energy measurement, which is discussed in Chapter 4, shows an improved energy resolution compared to the traditional EXO-200 reconstruction. These improvements translate to an improved median $0\nu\beta\beta$ half-life sensitivity after re-evaluating all significant contributions to systematic uncertainties. The increase in median $0\nu\beta\beta$ half-life sensitivity is $\sim 10\%$ when replacing the traditional EXO-200 energy observable in the baseline analysis by the DNN based observable. Using the DNN energy in the 3D fit configuration together with the standoff distance and the DNN discriminator provides a $\sim 35\%$ increase in median $0\nu\beta\beta$ half-life sensitivity relative to using the traditional EXO-200 energy only. This analysis configuration where two of three fit observables are based on DNNs provides the most stringent sensitivity. The large improvement in $0\nu\beta\beta$ half-life sensitivity from a mostly DNN driven analysis increases the probability to observe the $0\nu\beta\beta$ decay.

Chapter 7

Conclusion and Outlook

The neutrinoless double beta decay [12, 13] offers the most promising opportunity to answer fundamental questions about the nature of neutrinos. An observation of this decay would establish physics beyond the Standard Model of particle physics by proving that neutrinos are their own antiparticles [9]. This would have further implications on our understanding of neutrinos. The observation of the neutrinoless double beta decay would potentially help to identify the neutrino mass hierarchy and to understand the asymmetry between matter and antimatter in the Universe [10, 11]. A variety of different nuclei can undergo standard double beta decay [16]. These isotopes can be utilized to search for neutrinoless double beta decay. Nevertheless, the key performance requirements are shared among different experimental approaches. Current experiments provide a sensitivity to the half-life of this extremely rare decay of up to $\sim 10^{26}$ yr [15]. This illustrates why the main challenge for any experiment is to reduce the background. To achieve this in data analysis, a precise measurement of the kinetic energy of the decay products and good particle identification capabilities are required.

The EXO-200 experiment operated an ultra-low background time projection chamber filled with liquid xenon [14]. It completed data acquisition in 2018 and is among the most sensitive experiments that search for the neutrinoless double beta decay. A key advantage of this detector concept is its monolithic design that shields inner xenon from external background. In addition, a multi-parameter analysis allows to discriminate single from multiple localized particle interactions and enables strong background rejection capabilities. The analysis is based on measuring the ionization charge and scintillation light that are emitted from energy deposits in the xenon. These complementary measurements provide an optimized energy measurement as well as a full 3D position reconstruction.

The subject of this thesis was to improve the half-life sensitivity to the neutrinoless double beta decay of the EXO-200 experiment by applying deep learning methods in data analysis. Two main studies contributed to this objective. First, an energy reconstruction method based on a deep neural network was developed that directly exploits raw ionization charge signals. In particular, this algorithm outperforms the traditional EXO-200 energy reconstruction and achieves an improvement in energy resolution. Parts of this study are published in Ref. [70] by the EXO-200 collaboration. Second, deep neural networks were developed to discriminate double beta decays from background interactions. Exploiting the ionization charge signals, the deep neural network approach outperforms other discriminators that utilize classical machine learning techniques. This study contributed to the most recent search for neutrinoless double beta decay published by the EXO-200 collaboration in Ref. [14].

The deep neural network that is used to reconstruct the ionization energy from raw data was designed to optimally exploit the EXO-200 detector design and especially its two-sided symmetry. It was shown that a biased training dataset results in a biased performance of the deep neural network. This highlights the importance of a well-balanced training dataset which was used for the baseline deep neural network. The deep neural network is validated on source calibration data in both Monte Carlo simulations and measured data. Analogously to the traditional EXO-200 reconstruction, the ionization charge energy reconstructed by the deep neural network was combined with the complementary scintillation energy to optimize the energy measurement. The improvements of the deep neural network translated to improvements in the combined energy measurement. The deep learning based reconstruction achieved an energy resolution of 1.22% (1.01%) for single-site events at the decay energy of ^{136}Xe in Phase-I (Phase-II) of EXO-200 operation. This energy resolution represents an improvement of 10% (12%) compared to the traditional EXO-200 reconstruction. The robustness of the deep neural network reconstruction was validated. It was found that the variations in time and over the detector volume are on the same level as those observed with the traditional EXO-200 reconstruction. Furthermore, the similar spectral agreement between measured data and Monte Carlo simulations confirmed a robust performance of the deep neural network approach.

In an additional study, deep neural networks were developed to discriminate double beta decays from the dominant background in the EXO-200 experiment which represents γ particle interactions. In this study, a similar network architecture was used like in the energy reconstruction study that exploits detector symmetries. Also, to not introduce a bias to the network, particular focus was laid on the design of the training dataset. Besides using the raw ionization signals (DNN-Raw), another concept was discussed that exploits the re-generated signals from those found by the traditional EXO-200 reconstruction (DNN-Recon). Both deep neural network concepts were validated on measured data and Monte Carlo simulations. Both concepts outperformed the boosted decision tree (BDT) based discriminator that was used in the EXO-200 search for neutrinoless double beta decay in 2018 [91]. It was shown that both deep neural network concepts correlate the identification of double beta decays and background interactions to physically relevant parameters. This validates their robust performance. While the DNN-Raw concept slightly outperformed the DNN-Recon concept in terms of discrimination power, it showed a degraded spectral agreement between measured data and Monte Carlo simulations. Alternative concepts of deep neural networks that are able to capture additional spatial information were discussed. These deep neural networks exploited the full 3D spatial information while not impairing the spectral agreement between measured data and Monte Carlo simulations. However, this alternative approach requires to assume a prior spatial distribution of background. To search for neutrinoless double beta decay, trade studies favored the DNN-Recon discriminator concept among the other approaches.

A potential improvement of the half-life sensitivity to the neutrinoless double beta decay of the EXO-200 experiment was evaluated for the developed deep learning based analysis methods. EXO-200 pursues a multi-parameter analysis based on a maximum likelihood fit of the Monte Carlo simulated spectra to measured data. In this analysis concept, additional information besides the event energy can be exploited to further constrain background. The DNN-Recon concept was used in the most recent EXO-200 search for neutrinoless double beta

decay in ^{136}Xe to discriminate double beta decays from background interactions [14]. This search exploited a 3D maximum likelihood fit of energy, DNN-Recon discriminator and spatial information. This analysis configuration results in a median half-life sensitivity of $5.0 \cdot 10^{25}$ yr at the 90 % confidence level which provides the most stringent sensitivity to the neutrinoless double beta decay among all other analysis configurations within EXO-200. This represents an improvement in sensitivity by ~ 25 % compared to only utilizing the event energy. After unblinding, no statistically significant evidence for the neutrinoless double beta decay is observed in the complete dataset of Phase-I and Phase-II of EXO-200. This corresponds to a lower half-life limit of $3.5 \cdot 10^{25}$ yr at the 90 % confidence level or to an upper limit on the effective Majorana neutrino mass of $m_{\beta\beta} < 93 - 286$ meV. In order to replace the energy measurement of the traditional EXO-200 reconstruction with the one provided by the deep neural network, all significant systematic uncertainties that contribute to the maximum likelihood fit were re-evaluated. The resulting median half-life sensitivity to the neutrinoless double beta decay in an energy-only analysis configuration was improved by ~ 10 % compared to the traditional EXO-200 reconstruction. Likewise, the median half-life sensitivity is improved by ~ 35 % when using the deep neural network based energy measurement in a 3D analysis configuration in combination with the deep neural network based discriminator and a spatial variable. In this analysis configuration, two of three fit observables are provided by deep neural networks. It was shown that the remaining spatial information can also be captured successfully by the deep neural network based discriminator. This highlights the big potential of deep learning based analyses.

A fully deep learning based search for neutrinoless double beta decay is within reach. The remaining spatial information in the fit could also be provided by a separate deep neural network which reconstructs the full 3D position. This has already been successfully demonstrated where a deep neural network outperformed the traditional position reconstruction in EXO-200 [108]. In addition to the pure proof of concept, this analysis configuration is expected to increase the sensitivity to the half-life of this hypothetical decay. As was discussed for the DNN-Raw discriminator concept, a precise understanding and modeling of the detector is crucial for low background experiments like EXO-200. This argument is reinforced for deep learning based analyses that directly exploit low-level information from raw signals instead of high-level variables. Any discrepancies between measured data and Monte Carlo simulations negatively impact the success of such analyses. This type of discrepancies can probably be mitigated in the near future by generative neural networks, like generative adversarial networks that are able to refine the Monte Carlo simulations [72]. The increasing number of deep learning based applications over the last years provides a bright prospect to an increasing impact to future experiments. The nEXO experiment is a next-generation experiment that builds on the success of the EXO-200 experiment [37]. It will achieve a half-life sensitivity to the neutrinoless double beta decay of $\sim 10^{28}$ yr [48]. In addition, studies in Monte Carlo simulations indicate an improvement in half-life sensitivity of ~ 30 % by utilizing deep neural networks [109]. Enhancing the current sensitivities by two orders of magnitude, nEXO will fully cover the phase space of the inverted neutrino mass hierarchy. The nEXO experiment will be a big step toward a potential observation of the neutrinoless double beta decay and so to uncover the fundamental nature of neutrinos.

Appendix A

List of Figures

2.1	Semi-empirical mass parabolas	4
2.2	Feynman diagrams of $2\nu\beta\beta$ and $0\nu\beta\beta$ decay	5
2.3	Sum kinetic energy spectra of $2\nu\beta\beta$ and $0\nu\beta\beta$ decay	6
2.4	Neutrino mass hierarchies	8
2.5	Neutrino mass hierarchy bands as a function of Majorana neutrino mass	9
2.6	Cosmic muon flux	11
2.7	Conceptual design of the EXO-200 experiment	14
2.8	Schematic view of the EXO-200 LXe TPC	15
2.9	Picture of the EXO-200 TPC	16
2.10	Anti-correlation in scintillation and ionization signatures	18
2.11	Wire planes with charge drift	19
2.12	Illustration of interactions categorized as SS and MS	21
2.13	EXO-200 coordinate system and calibration source locations	22
3.1	Illustration of a simple neural network	26
3.2	Illustration of a unit in a neural network	26
3.3	Illustration of a convolutional layer	30
3.4	Illustration of a convolutional layer with padding	31
3.5	Illustration of a convolutional neural network	31
4.1	Example U-wire waveforms	35
4.2	Training data distributions in MC simulations	36
4.3	Network architecture	38
4.4	Training loss curve	39
4.5	Correlation of DNN predicted and MC true energy	40
4.6	^{228}Th data distributions in MC simulations	41
4.7	Reconstructed ^{228}Th spectra with artificial broadening in MC simulations	42
4.8	Reconstructed ^{228}Th spectra in MC simulations	43
4.9	Correlation of residuals of DNN and traditional EXO-200 reconstruction methods	44
4.10	Energy resolution over the detector volume for $0\nu\beta\beta$ decays in MC simulations	45
4.11	Example mixed signal on U-wires	46
4.12	Reconstructed ^{208}Tl peak in MC simulations	47
4.13	Correlation of DNN and traditional EXO-200 energy measurement	48

4.14 Example MC based fit to charge-only energy spectrum	49
4.15 Weekly charge-only energy resolution at the ^{208}Tl γ line	50
4.16 Charge-only spectral agreement between measured data and MC simulations . .	51
4.17 Charge-only energy resolution at the ^{208}Tl γ line over the detector volume . .	52
4.18 Anti-correlation of scintillation and DNN ionization signatures	53
4.19 Example MC based fit to combined energy spectrum	54
4.20 Energy resolution as a function of combination angle	54
4.21 Weekly combined energy resolution at the ^{208}Tl γ line	55
4.22 Combined spectral agreement between measured data and MC simulations . .	56
4.23 Energy resolution as a function of energy	57
4.24 Energy resolution as a function of date of submission	57
4.25 Reconstructed ^{208}Tl peak in ^{228}Th calibration data	59
 5.1 Example re-generated U-wire waveforms	63
5.2 Network architecture	65
5.3 Training loss and accuracy curves	66
5.4 Distributions of DNN-Raw for validation data in MC simulations	67
5.5 Distributions of DNN-Raw for LB physics data in MC simulations	69
5.6 ROC curves for different discriminators	69
5.7 ROC curves for DNN-Raw and BDT in SS events	70
5.8 Correlation of DNN classifications with event size in MC simulations	72
5.9 Correlation of classifications with standoff distance in MC simulations	73
5.10 Distributions of DNN-Recon for LB physics data	75
5.11 Spectral agreement between measured data and MC simulations of DNNs . .	76
5.12 Alternative background training data distributions in MC simulations	77
5.13 Correlation of DNN classifications with event position in MC simulations . .	78
5.14 ROC curves for alternative DNN discriminator concepts	79
5.15 Example U- and V-wire waveforms	80
5.16 Spectral agreement between measured data and MC simulations of DNNs . .	81
5.17 Extended network architecture	82
 6.1 3D fit configuration for maximum likelihood analysis	89
6.2 SS fraction for full and partial 3D selection cut	91
6.3 SS fraction as a function of energy	92
6.4 Energy spectra of remote background model components in MC simulations . .	94
6.5 Spectral agreement in energy between measured data and MC simulations . .	96
6.6 Spectral agreement in DNN and r_{SD} between measured data and MC simulations	96
6.7 Profile-likelihood curves for number of $0\nu\beta\beta$ decays in toy MC simulations . .	100
6.8 $0\nu\beta\beta$ half-life sensitivities for analysis configurations with E_{Trad}	101
6.9 Best-fit maximum likelihood curves to LB physics data with E_{Trad}	102
6.10 Profile-likelihood curves for number of $0\nu\beta\beta$ decays in LB physics data	103
6.11 Evolution of EXO-200 results over $0\nu\beta\beta$ searches	104
6.12 Best-fit maximum likelihood curves to LB physics data with E_{DNN}	105
6.13 $0\nu\beta\beta$ half-life sensitivities for analysis configurations with E_{DNN} and with E_{Trad}	106

Appendix B

List of Tables

2.1	Current and next-generation $0\nu\beta\beta$ experiments	13
4.1	Energy resolution at ^{208}Tl γ line in MC simulations	44
4.2	Reduction of ^{232}Th background in the ROI of $0\nu\beta\beta$ decay	60
6.1	SS fraction error	93
6.2	Location of remote background model components	93
6.3	Background model error	94
6.4	Spectral shape error	97
6.5	Summary of systematic errors	98
6.6	Median $0\nu\beta\beta$ half-life sensitivities for analysis configurations with E_{Trad}	100
6.7	Median $0\nu\beta\beta$ half-life sensitivities for analysis configurations with E_{DNN}	106

Appendix C

Bibliography

- [1] P. W. HIGGS. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (1964), pp. 508–509. DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508).
- [2] G. AAD et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716 (2012), pp. 1–29. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020).
- [3] S. CHATRCHYAN et al. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Physics Letters B* 716 (2012), pp. 30–61. DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021).
- [4] M. TANABASHI et al. “Review of Particle Physics”. In: *Phys. Rev. D* 98 (2018), p. 030001. DOI: [10.1103/PhysRevD.98.030001](https://doi.org/10.1103/PhysRevD.98.030001).
- [5] Y. FUKUDA et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562).
- [6] Q. R. AHMAD et al. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 89 (2002), p. 011301. DOI: [10.1103/PhysRevLett.89.011301](https://doi.org/10.1103/PhysRevLett.89.011301).
- [7] K. EGUCHI et al. “First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance”. In: *Phys. Rev. Lett.* 90 (2003), p. 021802. DOI: [10.1103/PhysRevLett.90.021802](https://doi.org/10.1103/PhysRevLett.90.021802).
- [8] C. L. COWAN et al. “Detection of the Free Neutrino: a Confirmation”. In: *Science* 124 (1956), pp. 103–104. DOI: [10.1126/science.124.3212.103](https://doi.org/10.1126/science.124.3212.103).
- [9] E. MAJORANA. “Teoria simmetrica dell elettrone e del positrone”. Italian. In: *Il Nuovo Cimento* 14 (1937), pp. 171–184. DOI: [10.1007/BF02961314](https://doi.org/10.1007/BF02961314).
- [10] M. FUKUGITA and T. YANAGIDA. “Baryogenesis without grand unification”. In: *Physics Letters B* 174 (1986), pp. 45–47. DOI: [10.1016/0370-2693\(86\)91126-3](https://doi.org/10.1016/0370-2693(86)91126-3).
- [11] V. KUZMIN, V. RUBAKOV, and M. SHAPOSHNIKOV. “On anomalous electroweak baryon-number non-conservation in the early universe”. In: *Physics Letters B* 155 (1985), pp. 36–42. DOI: [10.1016/0370-2693\(85\)91028-7](https://doi.org/10.1016/0370-2693(85)91028-7).
- [12] W. H. FURRY. “On Transition Probabilities in Double Beta-Disintegration”. In: *Phys. Rev.* 56 (1939), pp. 1184–1193. DOI: [10.1103/PhysRev.56.1184](https://doi.org/10.1103/PhysRev.56.1184).

- [13] J. SCHECHTER and J. W. F. VALLE. “Neutrinoless double- β decay in $SU(2) \times U(1)$ theories”. In: *Phys. Rev. D* 25 (1982), pp. 2951–2954. DOI: [10.1103/PhysRevD.25.2951](https://doi.org/10.1103/PhysRevD.25.2951).
- [14] G. ANTON et al. “Search for Neutrinoless Double- β Decay with the Complete EXO-200 Dataset”. In: *Phys. Rev. Lett.* 123 (2019), p. 161802. DOI: [10.1103/PhysRevLett.123.161802](https://doi.org/10.1103/PhysRevLett.123.161802).
- [15] M. AGOSTINI et al. “Probing Majorana neutrinos with double- β decay”. In: *Science* 365 (2019), pp. 1445–1448. DOI: [10.1126/science.aav8613](https://doi.org/10.1126/science.aav8613).
- [16] M. GOEPPERT-MAYER. “Double Beta-Disintegration”. In: *Phys. Rev.* 48 (1935), pp. 512–516. DOI: [10.1103/PhysRev.48.512](https://doi.org/10.1103/PhysRev.48.512).
- [17] O. RUSSAKOVSKY et al. “ImageNet Large Scale Visual Recognition Challenge”. In: *International Journal of Computer Vision* 115 (2015), pp. 211–252. DOI: [10.1007/s11263-015-0816-y](https://doi.org/10.1007/s11263-015-0816-y).
- [18] D. SILVER et al. “Mastering the game of Go with deep neural networks and tree search”. In: *Nature* 529 (2016), pp. 484–489. DOI: [10.1038/nature16961](https://doi.org/10.1038/nature16961).
- [19] A. RADOVIC et al. “Machine learning at the energy and intensity frontiers of particle physics”. In: *Nature* 560 (2018), pp. 41–48. DOI: [10.1038/s41586-018-0361-2](https://doi.org/10.1038/s41586-018-0361-2).
- [20] Y. LECUN, Y. BENGIO, and G. HINTON. “Deep learning”. In: *Nature* 521 (2015), pp. 436–444. DOI: [10.1038/nature14539](https://doi.org/10.1038/nature14539).
- [21] C. F. v. WEIZSÄCKER. “Zur Theorie der Kernmassen”. German. In: *Zeitschrift für Physik* 96 (1935), pp. 431–458. DOI: [10.1007/BF01337700](https://doi.org/10.1007/BF01337700).
- [22] R. SAAKYAN. “Two-Neutrino Double-Beta Decay”. In: *Annual Review of Nuclear and Particle Science* 63 (2013), pp. 503–529. DOI: [10.1146/annurev-nucl-102711-094904](https://doi.org/10.1146/annurev-nucl-102711-094904).
- [23] R. JONES et al. “Precision Electroweak Measurements on the Z resonance.” In: *Physics Reports* 427 (2006), pp. 257–454. DOI: [10.1016/j.physrep.2005.12.006](https://doi.org/10.1016/j.physrep.2005.12.006).
- [24] B. PONTECORVO. “Inverse beta processes and nonconservation of lepton charge”. In: *Sov. Phys. JETP* 7 (1958), pp. 172–173.
- [25] Z. MAKI, M. NAKAGAWA, and S. SAKATA. “Remarks on the Unified Model of Elementary Particles”. In: *Progress of Theoretical Physics* 28 (1962), pp. 870–880. DOI: [10.1143/PTP.28.870](https://doi.org/10.1143/PTP.28.870).
- [26] J. ENGEL and J. MENÉNZ. “Status and future of nuclear matrix elements for neutrinoless double-beta decay: a review”. In: *Reports on Progress in Physics* 80 (2017), p. 046301. DOI: [10.1088/1361-6633/aa5bc5](https://doi.org/10.1088/1361-6633/aa5bc5).
- [27] C. GIUNTI and C. W. KIM. *Fundamentals of neutrino physics and astrophysics*. Oxford university press, 2007. DOI: [10.1093/acprof:oso/9780198508717.001.0001](https://doi.org/10.1093/acprof:oso/9780198508717.001.0001).
- [28] S. DELL’ORO et al. “Neutrinoless double beta decay: 2015 review”. In: *Adv. High Energy Phys.* 2016 (2016), p. 2162659. DOI: [10.1155/2016/2162659](https://doi.org/10.1155/2016/2162659).
- [29] W. RODEJOHANN. “Neutrino-less double beta decay and particle physics”. In: *International Journal of Modern Physics E* 20 (2011), pp. 1833–1930. DOI: [10.1142/S0218301311020186](https://doi.org/10.1142/S0218301311020186).

[30] M. DOI, T. KOTANI, and E. TAKASUGI. “Double Beta Decay and Majorana Neutrino”. In: *Progress of Theoretical Physics Supplement* 83 (1985), pp. 1–175. DOI: 10.1143/PTPS.83.1.

[31] S. ADRIÁN-MARTÍNEZ et al. “Letter of intent for KM3NeT 2.0”. In: *Journal of Physics G: Nuclear and Particle Physics* 43 (2016), p. 084001. DOI: 10.1088/0954-3899/43/8/084001.

[32] X. QIAN and P. VOGEL. “Neutrino Mass Hierarchy”. In: *Prog. Part. Nucl. Phys.* 83 (2015), pp. 1–30. DOI: 10.1016/j.ppnp.2015.05.002.

[33] N. AGHANIM et al. “Planck 2018 results. VI. Cosmological parameters”. In: *arXiv e-prints* (2018), arXiv:1807.06209.

[34] V. N. ASEEV et al. “Upper limit on the electron antineutrino mass from the Troitsk experiment”. In: *Phys. Rev. D* 84 (2011), p. 112003. DOI: 10.1103/PhysRevD.84.112003.

[35] C. KRAUS et al. “Final results from phase II of the Mainz neutrino mass search in tritium β decay”. In: *The European Physical Journal C - Particles and Fields* 40 (2005), pp. 447–468. DOI: 10.1140/epjc/s2005-02139-7.

[36] J. WOLF. “The KATRIN neutrino mass experiment”. In: *Nucl. Instrum. Methods Phys. Res. A* 623 (2010), pp. 442–444. DOI: 10.1016/j.nima.2010.03.030.

[37] S. A. KHARUSI et al. “nEXO Pre-Conceptual Design Report”. In: *arXiv e-prints* (2018), arXiv:1805.11142.

[38] C. CHAMBERS et al. “Imaging individual barium atoms in solid xenon for barium tagging in nEXO”. In: *Nature* 569 (2019), pp. 203–207. DOI: 10.1038/s41586-019-1169-4.

[39] J. J. GOMEZ-CADENAS et al. “The Search for neutrinoless double beta decay”. In: *Riv. Nuovo Cim.* 35 (2012), pp. 29–98. DOI: 10.1393/ncr/i2012-10074-9.

[40] M. AGOSTINI, G. BENATO, and J. A. DETWILER. “Discovery probability of next-generation neutrinoless double- β decay experiments”. In: *Phys. Rev. D* 96 (2017), p. 053001. DOI: 10.1103/PhysRevD.96.053001.

[41] C. ALDUINO et al. “First Results from CUORE: A Search for Lepton Number Violation via $0\nu\beta\beta$ Decay of ^{130}Te ”. In: *Phys. Rev. Lett.* 120 (2018), p. 132501. DOI: 10.1103/PhysRevLett.120.132501.

[42] N. ABGRALL et al. “The Large Enriched Germanium Experiment for Neutrinoless Double Beta Decay (LEGEND)”. In: *AIP Conf. Proc.* 1894 (2017), p. 020027. DOI: 10.1063/1.5007652.

[43] A. GANDO et al. “ ^{7}Be solar neutrino measurement with KamLAND”. In: *Phys. Rev. C* 92 (2015), p. 055808. DOI: 10.1103/PhysRevC.92.055808.

[44] A. GANDO et al. “Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen”. In: *Phys. Rev. Lett.* 117 (2016), p. 082503. DOI: 10.1103/PhysRevLett.117.082503.

[45] J. J. GOMEZ-CADENAS. *Neutrinoless double beta decay*. XLV International Meeting on Fundamental Physics. 2017. URL: <https://indico.cern.ch/event/589204/contributions/2413789/attachments/1453832/2242882/GranadaDBM.pdf>.

[46] J. BOGER et al. “The Sudbury Neutrino Observatory”. In: *Nucl. Instrum. Methods Phys. Res. A* 449 (2000), pp. 172–207. DOI: 10.1016/S0168-9002(99)01469-2.

[47] S. ANDRINGA et al. “Current Status and Future Prospects of the SNO+ Experiment”. In: *Adv. High Energy Phys.* 2016 (2016), p. 6194250. DOI: 10.1155/2016/6194250.

[48] J. B. ALBERT et al. “Sensitivity and discovery potential of the proposed nEXO experiment to neutrinoless double- β decay”. In: *Phys. Rev. C* 97 (2018), p. 065503. DOI: 10.1103/PhysRevC.97.065503.

[49] J. B. ALBERT et al. “Search for Majorana neutrinos with the first two years of EXO-200 data”. In: *Nature* 510 (2014), pp. 229–234. DOI: 10.1038/nature13432.

[50] M. AUGER et al. “The EXO-200 detector, part I: Detector design and construction”. In: *Journal of Instrumentation* 7 (2012), P05010. DOI: 10.1088/1748-0221/7/05/P05010.

[51] J. B. ALBERT et al. “Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector”. In: *Phys. Rev. C* 89 (2014), p. 015502. DOI: 10.1103/PhysRevC.89.015502.

[52] R. NEILSON et al. “Characterization of large area APDs for the EXO-200 detector”. In: *Nucl. Instrum. Methods Phys. Res. A* 608 (2009), pp. 68–75. DOI: 10.1016/j.nima.2009.06.029.

[53] J. B. ALBERT et al. “Measurement of the drift velocity and transverse diffusion of electrons in liquid xenon with the EXO-200 detector”. In: *Phys. Rev. C* 95 (2017), p. 025502. DOI: 10.1103/PhysRevC.95.025502.

[54] G. LI. *Recent results from EXO-200 experiment*. WIN2019. The 27th International Workshop on Weak Interactions and Neutrinos. 2019. URL: https://agenda.infn.it/event/13938/contributions/89104/attachments/64106/77403/201906_WIN_EXO.pdf.

[55] J. ALBERT et al. “Cosmogenic backgrounds to $0\nu\beta\beta$ in EXO-200”. In: *Journal of Cosmology and Astroparticle Physics* 2016 (2016), pp. 029–029. DOI: 10.1088/1475-7516/2016/04/029.

[56] D. LEONARD et al. “Systematic study of trace radioactive impurities in candidate construction materials for EXO-200”. In: *Nucl. Instrum. Methods Phys. Res. A* 591 (2008), pp. 490–509. DOI: 10.1016/j.nima.2008.03.001.

[57] D. LEONARD et al. “Trace radioactive impurities in final construction materials for EXO-200”. In: *Nucl. Instrum. Methods Phys. Res. A* 871 (2017), pp. 169–179. DOI: <https://doi.org/10.1016/j.nima.2017.04.049>.

[58] E. APRILE and T. DOKE. “Liquid xenon detectors for particle physics and astrophysics”. In: *Rev. Mod. Phys.* 82 (2010), pp. 2053–2097. DOI: 10.1103/RevModPhys.82.2053.

[59] J. JORTNER et al. “Localized Excitations in Condensed Ne, Ar, Kr, and Xe”. In: *The Journal of Chemical Physics* 42 (1965), pp. 4250–4253. DOI: 10.1063/1.1695927.

[60] T. DOKE et al. “Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles”. In: *Japanese Journal of Applied Physics* 41 (2002), pp. 1538–1545. DOI: 10.1143/jjap.41.1538.

- [61] E. CONTI et al. “Correlated fluctuations between luminescence and ionization in liquid xenon”. In: *Phys. Rev. B* 68 (2003), p. 054201. DOI: [10.1103/PhysRevB.68.054201](https://doi.org/10.1103/PhysRevB.68.054201).
- [62] L. YANG. “Status and Prospects for the EXO-200 and nEXO Experiments”. In: *Journal of Physics: Conference Series* 888 (2017), p. 012032. DOI: [10.1088/1742-6596/888/1/012032](https://doi.org/10.1088/1742-6596/888/1/012032).
- [63] S. AGOSTINELLI et al. “Geant4a simulation toolkit”. In: *Nucl. Instrum. Methods Phys. Res. A* 506 (2003), pp. 250–303. DOI: [10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).
- [64] J. ALLISON et al. “Geant4 developments and applications”. In: *IEEE Transactions on Nuclear Science* 53 (2006), pp. 270–278. DOI: [10.1109/TNS.2006.869826](https://doi.org/10.1109/TNS.2006.869826).
- [65] W. SHOCKLEY. “Currents to Conductors Induced by a Moving Point Charge”. In: *Journal of Applied Physics* 9 (1938), pp. 635–636. DOI: [10.1063/1.1710367](https://doi.org/10.1063/1.1710367).
- [66] S. RAMO. “Currents Induced by Electron Motion”. In: *Proceedings of the IRE* 27 (1939), pp. 584–585. DOI: [10.1109/JRPROC.1939.228757](https://doi.org/10.1109/JRPROC.1939.228757).
- [67] R. ROBSON. “A Thermodynamic Treatment of Anisotropic Diffusion in an Electric Field”. In: *Australian Journal of Physics* 25 (1972), pp. 685–694.
- [68] S. SCHMIDT. “Analyses of the Electron Motion and Recombination in the EXO-200 Experiment”. MA thesis. Friedrich-Alexander-Universität Erlangen-Nürnberg, 2017.
- [69] D. O. NORTH. “An Analysis of the factors which determine signal/noise discrimination in pulsed-carrier systems”. In: *Proceedings of the IEEE* 51 (1963), pp. 1016–1027. DOI: [10.1109/PROC.1963.2383](https://doi.org/10.1109/PROC.1963.2383).
- [70] S. DELAQUIS et al. “Deep Neural Networks for Energy and Position Reconstruction in EXO-200”. In: *JINST* 13 (2018), P08023. DOI: [10.1088/1748-0221/13/08/P08023](https://doi.org/10.1088/1748-0221/13/08/P08023).
- [71] I. GOODFELLOW, Y. BENGIO, and A. COURVILLE. *Deep Learning*. MIT Press, 2016.
- [72] M. ERDMANN et al. “Generating and Refining Particle Detector Simulations Using the Wasserstein Distance in Adversarial Networks”. In: *Computing and Software for Big Science* 2 (2018), p. 4. DOI: [10.1007/s41781-018-0008-x](https://doi.org/10.1007/s41781-018-0008-x).
- [73] K. HORNIK. “Approximation capabilities of multilayer feedforward networks”. In: *Neural Networks* 4 (1991), pp. 251–257. DOI: [10.1016/0893-6080\(91\)90009-T](https://doi.org/10.1016/0893-6080(91)90009-T).
- [74] G. E. HINTON, S. OSINDERO, and Y.-W. TEH. “A Fast Learning Algorithm for Deep Belief Nets”. In: *Neural Computation* 18 (2006), pp. 1527–1554. DOI: [10.1162/neco.2006.18.7.1527](https://doi.org/10.1162/neco.2006.18.7.1527).
- [75] D. CIREŞAN, U. MEIER, and J. SCHMIDHUBER. “Multi-column Deep Neural Networks for Image Classification”. In: *arXiv e-prints* (2012), arXiv:1202.2745.
- [76] D. YU and L. DENG. *Automatic Speech Recognition: A Deep Learning Approach*. Springer London, 2015. DOI: [10.1007/978-1-4471-5779-3](https://doi.org/10.1007/978-1-4471-5779-3).
- [77] K. HE et al. “Deep Residual Learning for Image Recognition”. In: *arXiv e-prints* (2015), arXiv:1512.03385.
- [78] A. GHATAK. *Deep Learning with R*. Springer Singapore, 2019. DOI: [10.1007/978-981-13-5850-0](https://doi.org/10.1007/978-981-13-5850-0).

[79] C. BISHOP. *Pattern Recognition and Machine Learning*. Springer New York, 2006.

[80] 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*. Vol. 9. 2010, pp. 249–256.

[81] N. SRIVASTAVA et al. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. In: *Journal of Machine Learning Research* 15 (2014), pp. 1929–1958.

[82] Y. LECUN, L. BOTTOU, and Y. BENGIO. “Reading checks with multilayer graph transformer networks”. In: *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*. Vol. 1. 1997, pp. 151–154.

[83] V. DUMOULIN and F. VISIN. “A guide to convolution arithmetic for deep learning”. In: *arXiv e-prints* (2016), arXiv:1603.07285.

[84] C. SZEGEDY et al. “Going deeper with convolutions”. In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015, pp. 1–9. DOI: 10.1109/CVPR.2015.7298594.

[85] Y. LECUN et al. “Gradient-Based Learning Applied to Document Recognition”. In: *Proceedings of the IEEE* 86 (1998), pp. 2278–2324. DOI: 10.1109/5.726791.

[86] T. ZIEGLER. *Deep Neural Networks for Energy and Position Reconstruction in EXO-200*. 2018. DOI: 10.5281/zenodo.1300564.

[87] F. BOEHM and P. VOGEL. “Double Beta Decay”. In: *Physics of Massive Neutrinos*. 2nd ed. Cambridge University Press, 1992, pp. 159–216. DOI: 10.1017/CBO9780511622571.008.

[88] F. CHOLLET et al. *Keras*. <https://keras.io>. 2015.

[89] M. ABADI et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. 2015.

[90] D. P. KINGMA and J. BA. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).

[91] J. B. ALBERT et al. “Search for Neutrinoless Double-Beta Decay with the Upgraded EXO-200 Detector”. In: *Phys. Rev. Lett.* 120 (2018), p. 072701. DOI: 10.1103/PhysRevLett.120.072701.

[92] J. B. ALBERT et al. “Search for $2\nu\beta\beta$ decay of ^{136}Xe to the 0_1^+ excited state of ^{136}Ba with EXO-200”. In: *Phys. Rev. C* 93 (2016), p. 035501. DOI: 10.1103/PhysRevC.93.035501.

[93] S. IOFFE and C. SZEGEDY. “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift”. In: *Proceedings of the 32nd International Conference on Machine Learning*. Vol. 37. 2015, pp. 448–456.

[94] M. LIN, Q. CHEN, and S. YAN. “Network in network”. In: *arXiv e-prints* (2013), arXiv:1312.4400.

[95] J. P. EGAN. *Signal detection theory and ROC-analysis*. New York : Academic Press, 1975.

[96] K. A. SPACKMAN. “Signal Detection Theory: Valuable Tools for Evaluating Inductive Learning”. In: *Proceedings of the Sixth International Workshop on Machine Learning*. 1989, pp. 160–163. DOI: [10.1016/B978-1-55860-036-2.50047-3](https://doi.org/10.1016/B978-1-55860-036-2.50047-3).

[97] T. FAWCETT. “An introduction to ROC analysis”. In: *Pattern Recognition Letters* 27 (2006), pp. 861–874. DOI: [10.1016/j.patrec.2005.10.010](https://doi.org/10.1016/j.patrec.2005.10.010).

[98] T. G. DIETTERICH. “Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms”. In: *Neural Comput.* 10 (1998), pp. 1895–1923. DOI: [10.1162/089976698300017197](https://doi.org/10.1162/089976698300017197).

[99] W. J. YOUDEN. “Index for rating diagnostic tests”. In: *Cancer* 3 (1950), pp. 32–35. DOI: [10.1002/1097-0142\(1950\)3:1<32::AID-CNCR2820030106>3.0.CO;2-3](https://doi.org/10.1002/1097-0142(1950)3:1<32::AID-CNCR2820030106>3.0.CO;2-3).

[100] M. AUGER et al. “Search for Neutrinoless Double-Beta Decay in ^{136}Xe with EXO-200”. In: *Phys. Rev. Lett.* 109 (2012), p. 032505. DOI: [10.1103/PhysRevLett.109.032505](https://doi.org/10.1103/PhysRevLett.109.032505).

[101] J. B. ALBERT et al. “Searches for double beta decay of ^{134}Xe with EXO-200”. In: *Phys. Rev. D* 96 (2017), p. 092001. DOI: [10.1103/PhysRevD.96.092001](https://doi.org/10.1103/PhysRevD.96.092001).

[102] J. BAREA, J. KOTILA, and F. IACHELLO. “ $0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements in the interacting boson model with isospin restoration”. In: *Phys. Rev. C* 91 (2015), p. 034304. DOI: [10.1103/PhysRevC.91.034304](https://doi.org/10.1103/PhysRevC.91.034304).

[103] N. L. VAQUERO, T. R. RODRÍGUEZ, and J. L. EGIDO. “Shape and Pairing Fluctuation Effects on Neutrinoless Double Beta Decay Nuclear Matrix Elements”. In: *Phys. Rev. Lett.* 111 (2013), p. 142501. DOI: [10.1103/PhysRevLett.111.142501](https://doi.org/10.1103/PhysRevLett.111.142501).

[104] J. ENGEL, F. ŠIMKOVIC, and P. VOGEL. “Chiral two-body currents and neutrinoless double- β decay in the quasiparticle random-phase approximation”. In: *Phys. Rev. C* 89 (2014), p. 064308. DOI: [10.1103/PhysRevC.89.064308](https://doi.org/10.1103/PhysRevC.89.064308).

[105] J. MENENDEZ et al. “Disassembling the Nuclear Matrix Elements of the Neutrinoless beta beta Decay”. In: *Nucl. Phys.* A818 (2009), pp. 139–151. DOI: [10.1016/j.nuclphysa.2008.12.005](https://doi.org/10.1016/j.nuclphysa.2008.12.005).

[106] M. T. MUSTONEN and J. ENGEL. “Large-scale calculations of the double- β decay of ^{76}Ge , ^{130}Te , ^{136}Xe , and ^{150}Nd in the deformed self-consistent Skyrme quasiparticle random-phase approximation”. In: *Phys. Rev. C* 87 (2013), p. 064302. DOI: [10.1103/PhysRevC.87.064302](https://doi.org/10.1103/PhysRevC.87.064302).

[107] J. KOTILA and F. IACHELLO. “Phase-space factors for double- β decay”. In: *Phys. Rev. C* 85 (2012), p. 034316. DOI: [10.1103/PhysRevC.85.034316](https://doi.org/10.1103/PhysRevC.85.034316).

[108] J. LINK. “Event Reconstruction for EXO-200 using Deep Learning”. MA thesis. Friedrich-Alexander-Universität Erlangen-Nürnberg, 2019.

[109] Z. LI et al. “Simulation of charge readout with segmented tiles in nEXO”. In: *JINST* 14 (2019), P09020. DOI: [10.1088/1748-0221/14/09/p09020](https://doi.org/10.1088/1748-0221/14/09/p09020).

Acknowledgements/Danksagung

I would like to thank the following persons who supported me during this thesis:

Mein Dank gilt all den Personen, die mich während dieser Arbeit unterstützt haben und zu ihrem Gelingen beigetragen haben:

Prof. Dr. **Gisela Anton** für die Betreuung während der vergangenen Jahre und besonders für die Beratung in strategischen Fragen. Auch möchte ich mich für die Möglichkeit bedanken, dass ich Konferenzen auf drei Kontinenten besuchen und dort unsere Ergebnisse vorstellen konnte.

Dr. **Thilo Michel** für viele "Hast du kurz eine Minute?" meinerseits, und zwar meistens genau dann, wenn du eigentlich schon weg sein musstest. Da es dir in unserem Büro sowieso besser gefällt, darfst du gerne meinen Arbeitsplatz übernehmen!

Johannes Link und **Federico Bontempo** für das Engagement während ihrer Abschlussarbeiten, für unsere Diskussionen, und für ihre Ruhe trotz eines manchmal gestressten Betreuers. Auf jeden Fall hatten wir immer super Musik.

Den verrückten Kollegen aus Büro 208 und 209 für die vergangenen Jahre am ECAP. Wir hatten eine lustige Zeit, in der neben der Arbeit auch noch genug Zeit blieb um Chili Pflanzen zu züchten. Oder war es vielleicht doch eher andersrum?

Michael Jewell for your commitment and for answering my many questions. It was a though spring but finally we got it working. Thank you!

Dr. Igor Ostrovskiy, **Dr. Gaosong Li** and **Dr. Caio Licciardi** for your guidance and valuable discussions over the last years.

The **EXO-200** and **nEXO Collaborations** for their hard work, their willingness to listen to my presentations, and their feedback. Without all their previous work done, this thesis would not have been possible.

Uni Bayern e.V. für die reibungslose Unterstützung während dieser Arbeit.

Sebastian Schmidt, **Jonas Hammer** und **Theresa Palm** für die Bereitschaft diese Seiten zu lesen, zu korrigieren und einigermaßen leserlich werden zu lassen. Ein großes Dankeschön an euch!

Meiner Familie für die ständige Unterstützung. Was soll ich groß dazu sagen. Ohne euch wäre es einfach nicht möglich gewesen.

Theresa Palm für den Rückhalt ganz besonders auch während den stressigen Phasen. Dafür, dass du jede Geschichte und jeden Plot x-mal ertragen hast. Danke das du für mich da bist!