**Background Reduction for Improved Sensitivity at the Enriched Xenon
Observatory through advancements in Barium-Tagging Simulations,
Prediction and Prevention of High-Voltage Breakdown in Liquid Xenon
Time-Projection Chambers, and Event Classification using Ensemble
Methods and Machine Learning**

Jacob G. Watkins, B.Sc.

A thesis submitted to the
Faculty of Graduate and Postdoctoral Affairs
in partial fulfilment of the requirements of
Master of Science in Particle Physics

Ottawa-Carleton Institute for Physics

Department of Physics

Carleton University

Ottawa, Ontario

April 2019

©Copyright

Jacob G. Watkins, 2019

Abstract

The field of neutrino physics is abound with experiments looking to detect the hypothetical neutrinoless double beta decay ($0\nu\beta\beta$). Such an observation would prove that the neutrino is Majorana in nature, implying that lepton number is not a conserved quantity. A measurement of the half-life of the interaction would determine the absolute mass scale of neutrinos. This work presents advancements that have been made in three research campaigns that are intended to improve the capabilities of two $0\nu\beta\beta$ experiments – namely nEXO which is currently in the R&D stage, and its prototype, EXO-200 which finished data-collection in December 2018. First, an existing method for predicting breakdown voltages in liquid dielectrics through pattern recognition is applied to a liquid xenon time (LXe) projection chamber (TPC) that mimics the EXO-200 TPC. A proposed extension to this method, which aims to incorporate the statistical variations of breakdown voltage in liquid dielectrics, is outlined with comments on how future high voltage breakdown measurements should be taken in the laboratory. Advancements in applying machine learning ensemble methods to event discrimination in EXO-200 and nEXO are also presented, along with the basics of a new formalism for applying statistical inference to all classification algorithms. Finally, results of simulations of two barium-tagging concepts for nEXO are summarized. For the electrostatic grabbing probe method, the probability of ion capture has been mapped throughout a simulated TPC. For the capillary method, it has been shown that the image charge force created by charged particles on the capillary is sufficiently resisted by fluid drag forces to allow 100% transmission of ions through a capillary across which is an applied pressure difference.

Acknowledgements

Thanks to my undergraduate honours supervisor, Thomas Koffas, for introducing me to EXO a few years ago and for introducing me to *SIMION* software which I have continued to put to use since my honours project. Thanks to my master's thesis supervisor, Razvan Gornea, for guiding me in my research, providing invaluable opinions regarding physics results, and for setting me up with three unique and exciting research projects ... also, for tolerating my unorthodox working habits. As I'm rushing to finish this on time, I'll end my list there and get back to writing.

In theory there is no difference between theory and practice; In practice there is.

- Yogi Berra

Hence, experimental physics.

Contents

1 Introduction	1
2 Neutrinos & Neutrinoless Double Beta Decay	4
2.1 Historical Context	5
2.2 Neutrino Oscillations	12
2.3 Mass Hierarchy of Neutrinos	16
2.4 Neutrinoless Double Beta Decay	18
2.5 Energy Spectra of $\beta\beta$ -Decays	22
3 Searches for Neutrinoless Double Beta Decay	25
3.0 Counting Experiments	26
3.0.1 Sensitivity	26
3.0.2 Discovery Potential	28
3.1 Current Status of $0\nu\beta\beta$ Experiments	30
3.1.1 Large Liquids Scintillator Experiments	31
3.1.2 Bolometric Detectors	32
3.1.3 Germanium Diode Detectors	33
3.1.4 Tracking Calorimeters	35
3.1.5 Time Projection Chambers	36
3.2 The Enriched Xenon Observatory	39
3.2.1 Xenon as a $0\nu\beta\beta$ Source	39
3.2.2 EXO-200	42
3.2.3 nEXO	50
3.3 Background Limitations and Background Reduction	51
4 High Voltage Breakdown Studies	54
4.1 High Voltage Breakdown at EXO-200	54
4.2 Motivation for HV in TPCs	55

4.3 Breakdown Origin Hypotheses and miniEXO Tests	59
4.3.1 miniEXO	59
4.3.2 Breakdown Hypotheses	61
4.3.3 Past High Voltage tests with miniEXO and COMSOL Simulations	62
4.4 Pattern Recognition for Gap-Breakdown Prediction	66
4.4.1 The SVM Algorithm	66
4.4.2 Application of Pattern Recognition Prediction to miniEXO	70
4.4.3 A Statistical Extension to the Pattern Recognition Model	76
4.5 The HV Lab at Carleton	79
5 EXO-200 Data Analysis with Machine Learning	88
5.1 Current State of EXO-200 Data Analysis	89
5.1.1 Signal & Background Model & Variables Considered for Event Discrimination	89
5.2 Event-by-Event Classification with EXO-200 Data using Ensemble Methods.	97
5.2.1 Boosting	98
5.2.2 Decision Trees	100
5.2.3 Boosted Decision Trees Applied to EXO-200 Fit Model	102
5.2.4 Diversity Pruning	108
5.2.5 Stacking BDT with Naïve Bayes and KNN	111
5.3 Confidence Intervals for General Event-by-Event Classification Algorithms	117
5.4 Confidence Intervals for Classification Algorithms at EXO-200	125
5.5 Discussion Regarding the Choice of Metric for Classifiers.	130
6 Barium-Tagging Simulations	135
6.1 Sensitivity with Barium-Tagging for nEXO	135
6.2 The Electrostatic Probe	139
6.2.1 The Ba-Tagging Line with the Electrostatic Probe	140
6.2.2 COMSOL Model of Electroprobe	141
6.2.3 COMSOL Results for Single Ion Attraction with Electroprobe	145

6.3 Capillary Method of Ba ⁺ Extraction from a LXe TPC	159
6.3.1 The Barium-Tagging Line for the Capillary Method	159
6.3.2 Design of Flow Conditions in Capillary	162
6.3.3 Capillary Simulations in <i>COMSOL</i> and <i>SIMION</i>	167
7 Closing Remarks	171
A The Reynolds Number	174
B Increased Electric Field Strength due to Acrylic Struts	178

List of Tables

2.1	Best measurements and 3σ range of neutrino mixing parameters as of 2018. Table adapted from [23]. Values in brackets correspond to the inverted hierarchy while those outside brackets correspond to the normal hierarchy.	16
3.1	Summary of the present status of $0\nu\beta\beta$ studies. The tabulated lower limits for the half-lives correspond to the best limits that have been published for each isotope. $Q\beta\beta$ denotes the total energy released in $\beta\beta$ decay.	31
3.2	Candidate isotopes for $0\nu\beta\beta$ and their natural isotopic abundance, Q-value, and phase-space factors for both $\beta\beta$ modes. See (2.2.5)	40
3.3	Nuclear matrix elements of various $0\nu\beta\beta$ candidates according to the IBM-2 model and the QRPA model	41
3.4	EXO-200 results after 122 kg·yr exposure. Results are taken from an article published by the EXO collaboration in Nature, in 2014 [42]	48
3.5	Updated EXO-200 results from February 2018 [29]. Results combine previous exposure of 122 kg·yr with an additional 55.6 kg·yr of exposure in the upgraded EXO-200 detector. New analysis techniques involving boosted decision trees were implemented on the entire exposure to obtain the summarized results.	49
4.1	Complete list of training geometries used in pattern recognition algorithm for predicting the breakdown voltage of miniEXO. The Simulated Voltages column refers to the voltages to which the cathodes were set in COMSOL simulations. The breakdown voltage for hemispheres at 3 mm spacing is N/A as the system did not discharge by the maximum potential of -75 kV at that the HV supply could reach. . . .	74
4.2	Breakdown voltage of miniEXO as concluded by experiment, the non-weighted SVM classification algorithm, and the weighted SVM algorithm.	75
5.1	Pairwise Jensen-Shannon Divergence for collection wire risetime. The Jensen-Shannon Divergence is a symmetric metric so the lower triangular entries are omitted. The diagonal entries are always equal to zero and are also omitted.	92
5.2	Pairwise Jensen-Shannon Divergence for energy.	92
5.3	Pairwise Jensen-Shannon Divergence for number of collection wires	93

5.5	Correlations between parameters for $0\nu\beta\beta$ events	93
5.6	Correlations between parameters for $2\nu\beta\beta$ events	93
5.7	Correlations between parameters for ^{60}Co events.	94
5.8	Correlations between parameters for ^{232}Th events	94
5.9	Correlations between parameters for ^{238}U events	94
5.10	Correlations between parameters for ^{137}Xe events	94
5.11	Signal and background model for fitting in EXO-200 data analysis [38].	97
5.12	An example of bootstrap confidence intervals with nominal coverage of 68% applied to a dataset where the true event count of each event type (True column) is non-zero. We see that for this particular event distribution, we cannot distinguish $12\ 0\nu\beta\beta$ FROM ZERO. The <i>Observed</i> column denotes the number of events from the dataset that were classified as a given event type by the BDT algorithm before applying the MLE correction.	125
5.13	An example of bootstrap confidence intervals with nominal coverage of 68% applied to a dataset where the true event count of $0\nu\beta\beta$ is zero. Coverage studies show that nominal and true coverage may not be equivalent for rare events (i.e. $0\nu\beta\beta$ events in this case).	126
5.14	An example of bootstrap confidence intervals with nominal coverage of 68% applied to a dataset where the true event count of multiple event types is zero.	126
5.15	True coverage of various bootstrap confidence interval techniques applied to the event distribution of Test set 0. These numbers were obtained by creating confidence intervals for 438 unique datasets following the event distribution of Test Set 0. Each set used 10000 resamples during the bootstrap procedure.	127
5.16	True coverage, using percentile bootstrap, for nominal coverage of 68% with datasets respecting the event distribution of Test Set A	127
5.17	Mean and median biases of the binned MLE method for three unique event distributions. The median bias tends to be smaller than the mean bias for all datasets.	129

6.1	Properties of transitions of interest in $^{136}_{56}\text{Ba}^+$ tagging system. The wavelengths shown correspond to vacuum conditions [21][77].	138
6.2	Parameters for COMSOL simulation of electroprobe extraction of Ba ⁺ in a TPC. *The surface floating probe surface takes on the potential of the cathode in the simulation. ** To avoid the difficulty of simulating the electroprobe's mesh passing through the cathode's mesh, the electroprobe begins the simulation partially submerged in the LXe.	143
6.3	Description of capture regions for various rest times in Scenario A. A rest time of 3 s corresponds to the capture region in Figure 6.11. The bolded region in the table marks the sub-regions the capture region for which a time-dependence on the capture radius has been observed. Ranges of z-coordinate are used for some entries to conserve space. *Because the maximum initial radial coordinate of Ba ⁺ ions was 7.0 mm, any capture radii labelled as 7.0 mm may actually be greater than 7.0 mm.	154
6.4	Description of capture regions for various rest times in Scenario C. The bolded region in the table marks the sub-regions the capture region for which a time-dependence on the capture radius has been observed. The bolded region in the table marks the sub-regions the capture region for which a time-dependence on the capture radius has been observed. *Because the maximum initial radial coordinate of Ba ⁺ ions was 7.0 mm, any capture radii labelled as 7.0 mm may actually be greater than 7.0 mm.	156
6.5	Capture regions for the sub-regions of Scenario C with sub-regions defined in Figure 6.12. The different capture radii for equal z-coordinates demonstrates an angular dependence in capture probability for this scenario. Ranges of z-coordinate are used for some rows to conserve space. *Refer to Figure 6.12 for definitions of <i>Inner Region</i> and <i>Outer Region</i> headings. **Because the maximum initial radial coordinate of Ba ⁺ ions was 7.0 mm, any capture radii labelled as 7.0 mm may actually be greater than 7.0 mm.	157

List of Figures

2.1	A chart of all of the fundamental particles predicted by the SM, including the three quark families (purple), the three lepton families (green), the gauge bosons (orange-red), and the scalar boson – the Higgs (yellow). In 2012 the last, as of then unobserved particle, namely the Higgs boson, was discovered to complete the observation of all particles in the SM.	4
2.2	Detection mechanism used by Cowan and Reines. They observed inverse beta decay (predicted by applying crossing symmetry to Fermi’s beta decay) by the detection of a positron and a neutron in succession after the collision of an antineutrino with a proton. The neutron undergoes neutron capture in Cadmium, resulting in the emission of photons from the excited cadmium nucleus. The positron annihilates with one of the electrons that are bound to the Cadmium nuclei in the detection volume, resulting in the emission of photons.	10
2.3	Solar neutrino flux as a function of energy. Above the figure are ranges of energies to which various experiments & detection schemes were sensitive during their lifetimes. Adapted from [9]	12
2.4	Candidate mass hierarchies for the mass eigenstates of the neutrinos and the flavour compositions of the mass eigenstates. The normal hierarchy has m_3 significantly more massive than m_2 while the inverted hierarchy has m_3 significantly less massive than m_1 . In either scheme, m_1 & m_2 have a much smaller separation from each other than from m_3 . The degenerate case is one in which the mass differences are much smaller than the absolute masses of the neutrinos and in which the splitting can be ignored for most calculations. It can have a normal or inverted hierarchy as long as this condition is met. Cosmological studies, most recently the Planck telescope, have essentially ruled out this possibility regardless of the nature of neutrinos with an upper limit of the smallest neutrino mass of $<0.12\text{eV}$ [18]. However, such results are model dependent. KATRIN will be able to rule this out definitively	17
2.5	Double beta decay.	19

2.6	Mass parabola for isobars of 136. Single beta decay of ^{136}Xe is expressly forbidden while double beta decay is energetically favourable. The parabola is computed using the semi-empirical mass formula.	20
2.7	Neutrinoless double beta decay ($0\nu\beta\beta$). The demarcation of the vertices as SM Vertices implies the strictly left-handed weak current of the SM. The image on the right is adopted from [17].	21
2.8	Feynman diagram for $0\nu\beta\beta$ with Majorana mass term shown explicitly. The charge conjugating effect of the mass term as well as the flip of handedness allows both weak vertices to be of the typical left-handed SM type. Figure from [17].	22
2.9	<i>Left:</i> $2\nu\beta\beta$ energy distribution (broad) v.s. $0\nu\beta\beta$ (delta function distribution). <i>Right:</i> Overlap between energy distributions of $2\nu\beta\beta$ and $0\nu\beta\beta$ for different energy resolutions. EXO-200 currently operates with a resolution of 1.6% near the Q-value (see <i>Figure 3.13</i>). In these plots the Q-value is denoted by T_0 . Figure adapted from [21]	24
3.1	Projected sensitivities for several running and proposed $0\nu\beta\beta$. The bands in each plot correspond to effective Majorana masses of the neutrino with smearing due due to uncertainties in the nuclear matrix elements of various models. Adapted from [22].	28
3.2	The least detectable signal within the ROI of any counting experiment for a given expected Poisson background. Figure from [21].	30
3.3	Schematic of the KamLAND-Zen $0\nu\beta\beta$ large liquid scintillator experiment	31
3.4	A schematic of the bolometer employed at CUORE.	33
3.5	GERDA detection principle. Electrons from beta-producing events are collected on electrodes in the Ge-76 diodes. Light is detected in LAr to tag gamma-producing events.	34
3.6	The working principle of SuperNEMO. Thin foils of various $\beta\beta$ emitters are placed between tracking calorimeters and scintillators. The calorimeters and scintillators record the total energy of events while the calorimeter enables vertex locating and	

complete event reconstruction. Ultimately this enables measurements of the $0\nu\beta\beta$ angular distribution cross section..	35
3.7 Operating principle of a typical TPC. Events occurring within the detection volume (1) are tagged via scintillation light and charge collection. In this diagram, the scintillation light is collected by a PMT (4) but in principle the type of light measuring apparatus must be chosen carefully according to the experimental requirements (e.g. wavelength, radiopurity, quantum efficiency, etc...). Charge is collected by a wire grid, specially designed panels, etc.. (3). The constant drift field is maintained with a cathode (2'), an anode (2''), and several field-shaping rings. Figure adapted from [43]	36
3.8 Location of EXO-200 experiment at WIPP.	39
3.9 Schematics of EXO-200 TPC [37]. <i>Right:</i> APDs are read out in gangs. The top inset shows a zoomed view of the relative configuration of the wire planes and the APD plane. The bottom inset illustrates the angle between the wires planes.	43
3.10 <i>Left:</i> Principle of simultaneous light and charge detection in the EXO-200 TPCs. The 2D cross-sectional view of the wire planes demonstrates results of SIMION simulations for determining the potential at which electric transparency of the v-wires is achieved. Signals from the wires are read out in triplets to reduce radioactivity caused by excessive use of material. <i>Right:</i> View of induction and collection wires crossed at 60° in the EXO-200 TPC before installation to the cryostat. Images adapted from [37].	43
3.11 EXO-200 TPC (left) prior to installation in its copper vessel (right). Images adapted from [37].	44
3.12 Artist's interpretation of EXO-200 detector. Image adapted from [37]	45
3.13 <i>Left:</i> Computer model of copper guide tube wrapping around the copper vessel. <i>Right:</i> Various calibration source locations denoted by S2, S5, S8, S11.	46
3.14 <i>Left:</i> Graphical depiction of event-by-event anticorrelation between ionization and scintillation, and rotated energy axes in EXO-200. The energy is rotated to minimize the variance at the 2615keV peak of the ^{228}Th calibration source. <i>Right:</i> A	

comparison of energy spectra for scintillation only, ionization only, and rotated channel. The plots shown here only apply to SS events. Adapted from [40].	47
3.15 Projected sensitivity of EXO-200 after finishing data collection in December 2018. Figure adapted from [43]	49
3.16 Depth of SNO site in comparison to EXO-200 site at WIPP and other underground experiment locations. Operating nEXO at the SNO site would reduce the muon flux by more than two orders of magnitude relative to the flux at the EXO-200 site at WIPP. This reduces scattering backgrounds as well as the ^{137}Xe contribution (muon capture followed by decay) to the fits shown in the previous section. Figure adapted from [47].	50
3.17 <i>Left:</i> Example of a cut on energy ensuring 88% signal survival in the ROI. The height of the green line to the right of the cut indicates the unsuppressible background of $2\nu\beta\beta$ while the red lines indicate the signal distribution. <i>Right:</i> Signal-to-background ratio (3.3.1) as a function of the $0\nu\beta\beta$ half-life. Plots taken from [32] . . .	52
4.1 Expected improvements in energy resolution and SS fraction with increased cathode voltage at nEXO	56
4.2 Drift velocity for varied drift fields at EXO-200, measured with ^{228}Th calibration data and I alpha-sources from low-background runs, compared to drift velocity in Xe measured by other collaborations. LXe temperature: Gushchin-165K, Sorensen-177K, Akerib-175K, Miller-163K, Aprile-182K, EXO-200-167K. Adapted from [44] . . .	58
4.3 miniEXO (1) and one of the EXO-200 TPCs (2). Adapted from [43].	59
4.4 CAD model of TPC housed in the cryostat previously operated in Bern, now operating at Carleton. In this model a TPC other than the HV TPC is depicted. All details apart from the TPC are unchanged when the depicted TPC is swapped for the HV TPC. 1: TPC in copper chamber. 2: Steel base-plate with attached level meters for monitoring LXe level. 3: Cryogenic-camera for monitoring HV tests. 4: Reserved space. This space is planned to eventually house a Barium-tagging extraction device (see <i>Chapter 7</i>). 5: Location of indium seal connecting copper volume in which TPC is housed and the reserved space currently housing cryogenic cameras. 6: Cryostat	

	manifold comprised of several feedthroughs to components within the inner volume (e.g. connection to the cryo-camera).	61
4.5	Discharges in miniEXO captured by cryogenic cameras. <i>Top row:</i> The prompt flash of light during breakdown seems to occur on or near a strut in all cases. <i>Bottom Row:</i> There is a phosphorescent glow on the acrylic struts after discharges, presumably induced from LXe UV scintillation light during the spark. Figure adapted from [43]. . .	63
4.6	Electric field strength in miniEXO as determined by electrostatic COMSOL simulations. <i>Top:</i> Fields near acrylic combs. <i>Bottom:</i> Fields away from acrylic combs. The acrylic comb clearly amplifies the field strength by a small factor of a few percent	64
4.7	Data collected during HV tests with EXO-200 during a shutdown. The hexagons on the left of the image show event locations as measured by various wire planes (of which there are 4 – 2 for each TPC) and APD planes (2 – 1 for each TPC). The middle plots show the duration of signals in various channels for the corresponding hexagon on the left. Note the corner signature in the APD plane for the spark observed in this image. The corners are located near acrylic struts.	65
4.8	CAD model of SLAC test cell. <i>Top:</i> This particular CAD depicts the hemisphere electrode geometry. The upper hemisphere electrode can be swapped for other electrodes. <i>Bottom:</i> Feedthrough of voltage supply for upper electrodes. (<i>These images were provided by Peter Rowson and Knut Skarpaas of SLAC</i>).	71
4.9	Wire harness geometry operated in the SLAC test cell. Measurements in left image are given in inches. The holder is made of stainless steel while the wire is phosphor-bronze and has a diameter of 20 microns. <i>These images were provided by Peter Rowson and Knut Skarpaas of SLAC</i>)	72
4.10	XEBRA’s HV test cell. <i>Bottom middle:</i> CAD model used for COMSOL simulations of XEBRA’s Rogoski electrodes. <i>Bottom Right:</i> Electric field strength as simulated by COMSOL with a cathode potential of -75kV. (<i>All images were provided by Dr. Lucie Tvrznikova</i>).	73

4.11	CAD model of miniEXO simulated in <i>COMSOL</i> . This particular model was created by Tamer Tolba. The simulations relevant to the pattern recognition algorithm were completed by the author. The grid coordinates are displayed in centimetres. This model was built by Tamer Tolba who kindly provided the file to me for use in my thesis work.	75
4.12	Weibull distribution parameterization of breakdown voltages in XEBRA. Best fit parameters for describing the Weibull are provided on the left of the figure. Provided by Lucie Tvrznikova.	77
4.13	Carleton HV Group's cleanroom lab as of September 2016.	81
4.14	<i>Red circle</i> : Loose screw causing the purity monitor to connect virtual ground to ground. The Teflon cover prevents contact between the cryostat and the outer vessel after sealing.	82
4.15	<i>Top Left</i> : Burn on wire connected to ^{241}Am purity monitor. <i>Top middle</i> Connection to purity monitor was partially severed. <i>Top right</i> : Feedthrough for wires connecting to components of the inner volume. After it was replaced, the connection to the purity monitor had to be threaded through this feedthrough to reach the readout connection on the cryostat manifold. <i>Bottom Left</i> : <i>Inner volume</i> readout connection on for the purity monitor (coaxial connection). <i>Bottom middle</i> : Outer connection of purity monitor's readout. <i>Bottom right</i> : Showing purity monitor wires connected to the readout after successful soldering and current tests.	83
4.16	<i>Left/Middle</i> : Seal for the cryostat. The seal is a single piece of indium wire, and the only break is where the two ends meet. The butterfly cross is recommended to prevent leaks at this break. <i>Right</i> : Seal cavity prior to addition of indium wire.	84
4.17	Lowering the TPC into the inner volume of the cryostat. Lead bricks coated in blue paint can be seen on the lid. Because the lid doesn't have an even weight distribution, these bricks were used to balance the lid in such a way as to ensure an even distribution of pressure when the lid came into contact with the indium seal. ...	85
4.18	Leak check line connecting RGA to cryostat manifold. <i>Top left/Top Middle</i> : Turbo pump mounted on cryostat manifold and connected to a scroll pump for backing.	

	<i>Top Right:</i> Scroll pump. The red cable mounted to the steel beam connects to our helium source that was used for leak checking. <i>Bottom Left/Bottom Middle:</i> Connection between RGA and inner volume. <i>Bottom Right:</i> Analog pressure gauge connected to scroll pump. The line exhausts to the cleanroom air when the turbo is turned off. It then exhausts to the RGA once the turbo is switched on. We used a particle counter to confirm that exhausting the scroll to the room did not damage our cleanroom conditions.	86
4.19	Mass spectrum as measured by RGA sweep during leak checking with a Helium source. <i>Top left:</i> Mass spectrum. <i>Top right:</i> Mass spectrum with pedestal subtracted. Both spectra show a lack of a peak near the helium mass cluster (i.e. near 4 amu). <i>Bottom:</i> Helium and Argon partial pressures in the leak line. Both levels satisfy UHV conditions	87
5.1	Distributions, in various parameters, of event types included in the EXO-200 signal & background fit model. Risetime (top left), energy (top right), number of collection wires (bottom left), and standoff distance (bottom right).	92
5.2	Pseudocode for AdaBoost. In the final step for the strong learner, the sign of the weighted average – the summation – is taken. This is to classify events as “-1-like” events or as “1-like” events. In practice the -1 and the 1 are placeholders for events of a particular nature such as beta-producing or gamma-producing events or, in the case of this analysis, as any of $0\nu\beta\beta$, $2\nu\beta\beta$, ^{60}Co , ^{232}Th , ^{238}U , ^{137}Xe . The above algorithm must be slightly modified for the latter 6 class case, however [66]. In this algorithm ε_t denotes the error of the t^{th} weak learner, α_t denotes the weight of the t^{th} weak learner in the weighted average, h_t denotes the prediction of the t^{th} weak learner, and D_t denotes the distribution of data points that is updated based on the rate of misclassification.	99
5.3	Illustration of weighted average procedure in AdaBoost. Three weak learners acting as binary classifiers are combined to create a more complex decision boundary with superior classification accuracy to any of the individual learners.	100

5.4	Graphical representation of the logic of a BiDT. This particular tree would take two input variables, x & y , to make its cuts in classifying events as signal or background. The thresholds of the cuts are denoted by c , d , e which are learned thresholds on parameters, chosen to optimize discrimination	101
5.5	Learning curves for a BDT trained on EXO-200's 6-class signal/background fit model. The dotted line denotes the performance of a single decision tree. Evidently the boosting procedure improves discriminating power	104
5.6	The classification matrix for the optimized BDT. Each entry denotes the probability that a true event, given by the row index, is classified as a particular event-type, given by the column index. The classifier was trained on MC events generated using Geant-4 and the result of this matrix are from testing on a MC dataset. Below the matrix is the result of testing the MC-trained BDT on events from a ^{232}Th calibration source. Bold entries denote correct classification rates, red entries indicate the class as which each true event type is most often misclassified, and green entries indicate the class as which each true event type is most rarely misclassified	105
5.7	Histograms of the bootstrapped MLE for the correct classification rate of ^{232}Th using MC data (left) and calibration data (right). The bootstrap procedure was repeated 5000 times for both MC and calibration data. A Gaussian fit is superimposed on each histogram.	106
5.8	Learning curves for performance of BDT on MC ^{232}Th (left) and calibration ^{232}Th (right).	106
5.9	Learning curves for a MC test set with (red) and without (blue) diversity pruning. There is a modest improvement after pruning (green circle)	109
5.10	Error of each tree plotted in order of addition to the overall forest (left). Evidently the error rate of a given tree tends to increase the later it is constructed and added to the forest. Of course, the weights of each tree correspondingly decrease the later the tree is constructed (right).	110
5.11	Classification matrix of a Naïve Bayes' trained and tested on MC datasets (the same datasets as were used to train the BDT)	112

5.12	Learning curves of the KNN for various values of p in the Minkowski distance metric. The horizontal axis denotes the number of neighbours considered by the algorithm in the voting step while the vertical axis denotes the overall error on 1200 test points. The test points are evenly distributed among the six event types (i.e. 200 points per event type).	114
5.13	Classification matrix of a KNN trained and tested on MC datasets (the same datasets as were used to train the BDT). The number of neighbours used in the classification vote was 56.	114
5.14	Architecture of the stacked learner. The hidden layer includes a NB learner and a KNN learner. The final layer combines the outputs from these hidden layer algorithms with the parameters risetime, energy, number of triggered U-wires, and standoff distance as inputs for a BDT which makes the final prediction	116
5.15	Learning curves of a stacked BDT taking predictions from a KNN and a NB as input along with the measured energy, U-wire risetime, number of triggered U-wires, and standoff distance.	116
5.16	Classification matrix of a stacked BDT taking predictions from a KNN and a NB as input along with the measured energy, U-wire risetime, number of triggered U-wires, and standoff distance. The BDT includes 6000 trees.	117
5.17	Overlapping probability distribution functions in 1D. The red, vertical line denotes the optimal classification cut. All points to the left of the line are labelled as class a and all points to the right of the line are labelled as class b	118
5.18	A classification algorithm acts as a transformation that takes a physics event parameterized by d parameters to an integer that represents an event type. The transformation is illustrated here as a hashing transformation	120
5.19	The observation vector can be binned into a 1D histogram referred to as the <i>Observed Histogram</i>	122
5.20	<i>Left:</i> Purity of $0\nu\beta\beta$ column in the classification matrix corresponding to <i>Figure 5.5</i> as a function of the number of trees in the forest. <i>Right:</i> Zoomed view of the left plot corresponding to the flat region of the learning curve in <i>Figure 5.5</i>	133

6.1	Sensitivity of nEXO with and without Ba-tagging, compared to final EXO-200 sensitivity. Sensitivity is computed with a 90% confidence interval [22]. The red ribbon denotes the space occupied by the inverted hierarchy, and the blue ribbon denotes the normal hierarchy	137
6.2	The single electron system of $^{136}_{56}\text{Ba}^+$ with branching ratios of the emission/absorption spectrum.	137
6.3	Ba-tagging line for a LXe TPC scheme with an electrostatic probe as the extraction device. There are five steps in the line. From left to right they are: the TPC, attraction to the electroprobe, removal from the electroprobe, the ion trap, the two-laser system	140
6.4	CAD model of the TPC and probe simulated in COMSOL. Inner circular planes are for the purpose of a FEM meshing trick and do not affect physics results.	142
6.5	Saturation curve for Xenon. The leftmost orange point denotes the triple point. The red point denotes the conditions of LXe in the simulated TPC	142
6.6	Potential of surfaces in electroprobe simulation. There is an imposed, linear change in potential along the walls of the TPC (walls are continuous – a simplification of the field shaping rings) from the anode to the cathode which are held at 0kV and -8kV respectively. The electroprobe is set to a floating potential condition. It can be seen here that the surface of the probe takes of the potential of the cathode through which it passes during its descent	144
6.7	Electroprobe simulation scenarios. A scenario is defined by the position at which the electroprobe comes to rest in the TPC. The “height” parameter refers to the height at which the probe come to rest, with the anode being at a height of zero. The “distance from rings” parameter is the shortest distance from the probe’s stopping point to the rings. The TPC coordinates are given in centimetres.	146
6.8	Velocity maps for Ba^+ at various times during the motion of the probe. Each image corresponds to the map at a different cross-sectional time during the motion of electroprobe as defined by (6.2.2). <i>Top left</i> : probe moving at -1cm/s (time = 1s), <i>Top</i>	

	<i>Right: probe moving at speed of -1cm/s (time = 5.2s), Bottom Left: instant that probe comes to rest (time = 7.2s), Bottom right: probe has been at rest for 3s (time = 10.2s)</i>	147
6.9	Velocity map for Ba^{+} at instant that probe comes to rest (see equation 6.2.2) in <i>Scenario C. Top: Velocity maps showing cut planes for two bottoms plots. Bottom Left: Velocity for Cut Map A. Bottom Right: Velocity for Cut Map B.</i>	148
6.10	Capture region for <i>Scenario A</i> . Radial and longitudinal dimensions are not to scale . .	151
6.11	Convergence of ions towards the electroprobe. This snapshot shows the trajectories of each ion from time zero to the point at which the probe comes to rest (i.e. from 0 s to 7.2 s). There are three cases: 1) Some ions move too far vertically before converging towards the probe resulting in collisions with the sides of the probe (cyan/light blue-green points) 2) Others simply do not converge enough radially in the duration of the probe's descent (dark blue points). 3) Finally, many converge nicely to the tip of the probe (red and orange points).	152
6.12	Flow field after various rest times for <i>Scenario A</i>	153
6.13	Cut plane for splitting ion distribution into two regions in <i>Scenario C</i> . The image on the right is an overhead view of the image on the left. The overhead image is included to make it clear that the cut plane passes through the center of the probe. The cut plane breaks geometric symmetry producing different electric fields and flow fields on either side of the plane	155
6.14	Concept of a Ba-tagging line for a GXe TPC operating at 10bar. The above diagram is the general concept while the lower diagram depicts a setup that has already been tested at McGill University [85].	159
6.15	Ba-tagging line for the capillary extraction scheme.	161
6.16	Flow regimes in a fluid undergoing nucleate boiling in a long tube. The temperature of the tube versus the fluid temperature are depicted on the left of the figure for a typical fluid. Flow regimes and the type of heat transfer that dominates each regime are given on the right of the figure. Red arrows indicate the convergence of the gas	

	phase as bubbles are ejected from the wall and the divergence of a small annular liquid region that replaces the volume of escaping bubbles	162
6.17	Illustration of nucleate boiling at a rectangular interface between a liquid and a heated surface	163
6.18	Typical, experimentally determined flow types (compressible or incompressible) for various Mach numbers (left) and Reynolds number regions producing laminar or turbulent flow (right).	165
6.19	Plots of the Reynolds number for GXe (left) and for LXe (right) when the capillary is designed to displace the maximum mass of fluid that can be pumped at the McGill University RF-Funnel [86]	166
6.20	<i>Left:</i> Streamlines for the single-phase LXe flow with applied 0.75bar/m pressure gradient in 1 180um diameter capillary. Results are obtained from <i>COMSOL</i> FEM software with the <i>laminar flow</i> module. The analogous single-phase GXe version of this plot is identical apart from the fact that the flow speed is greater for GXe. <i>Right:</i> Streamlines near the outlet in single-phase GXe flow with the same applied BCs. The red arrows indicate the direction of GXe flow at the cross-section of the capillary's outlet.	168
6.21	<i>Left:</i> Ba^+ trajectories along the capillary in <i>SIMION</i> simulations. <i>Left:</i> Side view of ions trajectories in the capillary. Green lines mark trajectories. <i>Right:</i> Red dots mark positions of Ba^+ ions at the inlet and outlet of the capillary. Red dots are used to mark the beginning and ending points (notice that there aren't any in the left image, signifying no losses along the capillary).	169
6.22	Trajectories of Ba^+ ions in a capillary. Red dots mark the beginning and ending points of an ion's trajectory through the capillary. The farthest that any of them travel before colliding with the wall of the capillary is 14 cm.	170

Statement of originality and collaborative contributions

I hereby declare that the work contained in this Master's thesis constitutes my own original research and that I have not used any sources other than those stated.

The work summarized in *Chapter 4* concerns electrical breakdown in bulk liquid xenon (LXe) across which there is a potential difference. Simulation work that provided the quantitative motivation for operating a time projection chamber (TPC) at high voltage (HV), summarized in *Figure 4.1*, was completed by Liangjian Wen and Guofu Cao of *IHEP*. High voltage breakdown tests on the apparatuses depicted in *Figure 4.8* and *Figure 4.9* were completed and the corresponding data were provided by Peter Rowson and Knut Skarpaas of *SLAC*. High voltage breakdown tests on miniEXO (see *Figure 4.3*) were completed and the corresponding data were provided by Razvan Gornea and Sebastien Delaquis (then at *University of Bern*). High voltage tests in the apparatus depicted in *Figure 4.10* were completed and the corresponding data in *Figure 4.12* were provided by Lucie Tvrznikova (then at *UC Berkeley*). Results of the electric field simulations of miniEXO, depicted in *Figure 4.6*, and the corresponding CAD model were completed and provided by Tamer Tolba (then at *University of Bern*). These results initiated my own work in which I adopted a pattern recognition algorithm, developed by other authors studying electrical breakdown in systems with gaseous dielectrics [61][62], to predict the voltage at which dielectric breakdown will occur across LXe in a TPC. I combined the data from the above high voltage tests into a single dataset. Except for miniEXO, I built CAD models of each of the apparatuses, and simulated each of them in a finite element method (FEM) software called *COMSOL Multiphysics®*. I extracted several electric field parameters and geometric parameters from these simulations and paired them with the corresponding laboratory data. I used this dataset to train the pattern recognition algorithm. The algorithm that was developed by other authors uses a single breakdown voltage, explicitly disregarding the time-dependence of dielectric breakdown system – a fairly accurate approximation in gases but not in liquids. Thus, as a second part of my work on electrical breakdown in LXe, I proposed

a method and measurements for extending the pattern recognition algorithm to liquid dielectrics in which the statistical nature of the problem can be considered. This method makes use of an equivalence between static breakdown tests and dynamic breakdown tests that was first pointed out by other authors [53]. I used this equivalence to define a new method of labelling the threshold for breakdown in (4.4.8). Finally, I assisted in the commissioning of a cleanroom in which future high voltage breakdown studies will be conducted in miniEXO – likely employing the suggested statistical measurement approach from my thesis work. I constructed the cleanroom itself, installed a TPC in the cryostat, rewired and soldered severed electrical connections between the TPC components and the cryostat manifold, and sealed the TPC into the cryostat. Following this, I assisted in installing a temporary leak-check line (see *Figure 4.18*) to the cryostat manifold and confirmed that there were no leaks using an RGA and a helium source.

The work summarized in *Chapter 5* concerns the development of new classification algorithms for discrimination between event types in the EXO-200 fit model, and the development of a new formalism for making statistical inference about entire datasets using the output of such algorithms. Part of this work was initiated by Warren Cree in his Master’s thesis work [65]. He developed the original boosted decision tree algorithm that made use of three parameters – collection wire risetime, number of triggered collection wires, and standoff distance – and output a score between -1 (signal event) and 1 (background event) [65]. He selected these parameters based on a statistical analysis of their discriminating power between signal and background events. My work consisted of adding the energy parameter to his algorithm and producing a direct classification as one of the six event types in the EXO-200 fit model instead of a score between -1 and 1. The Monte Carlo (MC) data that was used for training and testing of the various classification algorithms that I developed were provided by Igor Ostrovskiy, Hank Richards, and Mike Jewell of the machine learning analysis team at EXO-200. The data itself was generated by various members of the EXO-200 MC team. I used the parameters determined by Warren Cree and the provided MC data to develop all algorithms whose results are summarized in *Figure 5.6*, *Figure 5.11*, and *Figure 5.16* via classification matrices. I myself completed all of the work done in building and tuning these classification algorithms, part from the

aforementioned work done by Warren Cree. I also completely developed the mathematical formalism, summarized in *Chapter 5 – 5.3*, for creating confidence intervals based on the output of any classification algorithm. In *Chapter 5 – 5.4*, I applied this approach to the EXO-200 MC data that was provided to me and conducted coverage studies to show that the confidence intervals built using the method in *Chapter 5 – 5.3* produce nominal coverage.

The work in *Chapter 6* concerns simulation studies of two ion extraction methods for the proposed Barium-tagging measurement at nEXO. All of the work summarized in *Chapter 6* is my own work, unless cited. My contributions to this chapter were the design and analysis of two simulations. I developed the CAD model of an electrostatic probe. I simulated the motion of this probe inside a liquid TPC using a moving mesh in the *COMSOL* FEM software. Previous attempts at simulating this motion by other authors failed to include electrostatics, fluid dynamics, and particle tracking effects in a single simulation and also required the removal of the cathode in the TPC in order to simulate the moving mesh. My simulations successfully incorporated all three physics effects, coupling the fluid motion to the motion of the moving electrostatic probe, and did not require removal of the cathode. I used these results to map regions in which capture of Ba^+ ions of the probe's tip are possible. I also completed simulations Ba^+ ions moving through a capillary at a fixed potential due to the drag force of laminar flow in LXe and GXe. I showed that the image charge force between Ba^+ ions and the capillary wall is divergent (i.e. the force is towards the wall of the capillary) but that drag force due to the flowing xenon is sufficiently strong to oppose this divergence in 100 % of cases.

Chapter 1

Introduction

It has been nearly a century since the conception of neutrinos in the mind of Wolfgang Pauli. Since then, they have proved elusive to physicists and there remains much to be learned about their properties. The most recent major discovery in neutrino physics was made around the turn of the millennium when two experiments, *SNO* and *Super-K*, observed oscillations between neutrino flavours [1][2]. Neutrino theory suggests that flavour oscillations are only possible if neutrino flavour states are super-positions of mass eigenstates of different masses. Thus, the results from *SNO* and *Super-K* showed that neutrinos are massive. However, oscillation experiments are sensitive to the difference in squared masses, and not to the masses themselves. Indeed, the absolute mass scale of the neutrinos is an open question. Furthermore, past oscillation experiments were not sensitive to the so-called *mass hierarchy* of neutrinos which describes how the neutrinos are ordered by mass

One approach to determining the absolute mass scale of the neutrinos is by observing a process called neutrinoless double beta decay ($0\nu\beta\beta$). This hypothetical interaction, if observed, would allow physicists to determine two properties of neutrinos. Firstly, observing the process at all requires that neutrinos be their own antiparticles (a.k.a. *Majorana* particles). Secondly, the lifetime of the decay is proportional to the square of the neutrino mass and so a measurement of the lifetime will allow physicists to determine the absolute mass scale of neutrinos. This will provide constraints on the open mass hierarchy question and possibly provide a definitive answer. Presently, there are several experiments, both currently collecting data or in the R&D stage, that are searching for $0\nu\beta\beta$ (see *Chapter 3*). One such experiment, the Enriched Xenon Observatory (EXO), is attempting to observe the interaction in ^{136}Xe . The current best estimate for the lower limit on the $0\nu\beta\beta$ half-life in ^{136}Xe is greater than 10^{26} years. Thus, if $0\nu\beta\beta$ is observed it will be the slowest process ever measured. Indeed, we expect only one or two $0\nu\beta\beta$ events per year in a tonne-scale ^{136}Xe experiment. For this reason, background reduction and

rejection capabilities are of the utmost importance in $0\nu\beta\beta$ experiments in general. The bulk of the work summarized in this thesis is devoted to, in broad terms, background reduction for EXO. The structure of the rest of this thesis is outlined below.

Chapter 2 provides a basic description of our current understanding of the theory and the properties of the neutrino. This is presented in a historical context, beginning nearly a century ago and culminating the recent discovery of neutrino oscillations. The implications of the existence of neutrino oscillations sets the stage for a discussion of the theory of neutrinoless double beta decay, with which the chapter concludes.

Chapter 3 is a discussion of counting experiments for the detection of $0\nu\beta\beta$. Various detection techniques for $0\nu\beta\beta$ counting experiments are presented along with the current best results on the limit of the Majorana neutrino mass for prominent experiments employing each technique. Of these experiments, the Enriched Xenon Observatory (EXO) experiment is considered in particular detail, as the majority of R&D presented in this thesis is devoted to, broadly speaking, background reduction for EXO. An overview of EXO's detector is given along with a list and descriptions of the measurements that it records for use in data analysis. The discussion continues by presenting the case for ^{136}Xe as a candidate $0\nu\beta\beta$ source. The chapter ends with a brief discussion of the background limitations for EXO in comparison to the other experiments presented in the chapter. These background limitations motivate the bulk of the work that is presented in the remainder of the thesis.

Chapter 4 presents results of a study of high voltage breakdown in liquid xenon. The use of high voltage in time projection chambers is motivated through improved measurement resolution. From there, a summary of past high voltage work conducted at EXO-200 and the University of Bern is given. These results motivate the current work in which data from several liquid xenon breakdown studies is combined through pattern recognition analysis to predict the breakdown voltage in a liquid xenon time projection chamber. The chapter ends by outlining a proposed method for extending the pattern recognition approach to account for statistical variations in breakdown voltage.

Chapter 5 summarizes advancements that have been made in applying classification algorithms to EXO-200 data as well as the beginnings of a formalism for statistical inference in the context of classification algorithms. We begin with a description of the current state of data analysis in the EXO-200 experiment. This includes a list and description of each of the measured parameters that are available for use in data analysis as well as a list of the various types of events (or classes) that are included in the signal & background fit model. Once this is established, a description of ensemble classification methods is given. These methods are then applied to EXO-200 Monte Carlo data. The chapter then generalizes beyond the scope of EXO-200 to provide a framework for creating confidence intervals for any classification algorithm with an arbitrary number of classes and variables. To the knowledge of myself and the rest of the machine learning analysis team at EXO, this is a new development in the field. The chapter ends by applying this new technique to various EXO-200 Monte Carlo datasets.

In Chapter 6 a method for rejecting all non-double-beta events within EXO's detector, called Barium-Tagging, is presented. Advancements in simulations of two unique ion extraction schemes for the barium-tagging branch of the nEXO collaboration are summarized. The first is an electrostatic grabbing probe that can be inserted into a liquid xenon time projection chamber and attract charged particles to its tip. The second method applies a pressure gradient across a capillary to suck particles from the liquid xenon volume. Heat transfer is used to boil the liquid, enabling separation of charged particle from the gaseous xenon.

Chapter 2

Neutrinos & Neutrinoless Double Beta Decay

At its most fundamental level, this work is concerned with the properties of neutrinos – the subject of this chapter. However, before dealing with them directly, we begin by placing neutrinos in their historical context as well as within the current Standard Model of Particle Physics (SM).

The SM is a framework that attempts to describe the fundamental, sub-atomic particles of the universe, tabulated in *Figure 2.1*, as well as the interactions between them. While the SM has enjoyed remarkable success in predicting the existence of sub-atomic particles – the last missing piece, the Higgs boson, being observed at the LHC in 2012 – a number of the properties of these particles continue to elude observation. The neutrinos remain particularly mysterious and have proven to cause trouble for the SM. Indeed, the neutrinos interact with matter only via the weak interaction which makes them notoriously difficult to detect and it is for this reason that many of their fundamental properties remain unknown.

Figure 2.1 A chart of all of the fundamental particles predicted by the SM, including the three quark families (purple), the three lepton families (green), the gauge bosons (orange-red), and the scalar boson – the Higgs (yellow). In 2012 the last, as of then unobserved particle, namely the Higgs boson, was discovered to complete the observation of all particles in the SM.

For example, the SM assumed, for decades, that the neutrino was massless (though there is no physical law requiring their masslessness). Recently, however, results from SNO [1] and Super-Kamiokande [2] showed that this is not the case. Indeed, both experiments observed neutrino oscillations, showing definitively that neutrinos must be massive (the explanations leading to this conclusion will be presented further on in this chapter). These observations immediately implied new open questions about neutrinos. For example, these experiments proved that neutrinos must be massive, they did not measure the masses of the neutrinos directly. Rather, they were only sensitive to the squares of the mass differences between the neutrinos. This being the case, the mass hierarchy (discussed later) and the absolute masses of the neutrinos are still unknown.

That the neutrinos have mass also calls into question their nature – (i.e. is the origin of their mass a Dirac term or a Majorana term?) – when combined with other neutrino properties. A particular process, $0\nu\beta\beta$, was proposed by W.H. Furry in 1939 [3]. If such a process can occur in nature, then the neutrino must be of Majorana nature – that is, it must be its own anti-particle. Furthermore, such a process would reveal the absolute mass scale and the mass hierarchy of the neutrinos if it were to be observed.

In the remaining sections of this chapter, the open questions regarding the masses of the neutrinos as well as their nature will be discussed. To provide the motivation for searching for $0\nu\beta\beta$, we begin with the birth of the neutrino in the mind of Pauli and proceed through the history of discoveries around the neutrino. This provides context for experimental $0\nu\beta\beta$ searches. The chapter ends with a description of $0\nu\beta\beta$ and how it may be observed in an experiment.

2.1 Historical Context

By 1927 (we will see momentarily why we start here), beta decay of nuclei had long since been observed. As early as the year 1900, Henri Becquerel measured the charge to mass ratio of the emitted “beta” particle and discovered that it was an electron. In the following year, Rutherford

and Soddy showed that the process involved the transmutation of one elemental nucleus to another with the emission of this electron. The process was thought to be described by,

$$A \rightarrow B + e^{-} \quad 2.1.0$$

where A is the original nucleus and B is the final nucleus of elemental identity one space to the right of the identity of A on the periodic table.

It turned out that (2.1.0) was an incorrect equation for describing beta decay. It can be shown that, if (2.1.0) were correct, one would expect the energy, E , of the electron to be fixed for a decay in which A is stationary, and to be given by,

$$E = \frac{m_A^2 - m_B^2 + m_e^2}{2m_A} \quad 2.1.1$$

where m refers to the mass of the particle of identity denoted by the subscript. It is well known by now that this is not what is observed in experiment. The first experiment to contradict (2.1.1) was carried out in 1927 by Wooster & Ellis [91]. They reported the observation of a continuous energy distribution of the emitted electron. This result seemed to contradict conservation of energy as a third particle carrying away the energy difference between the energy of the emitted electron and (2.1.1) was never observed.

In 1930, Wolfgang Pauli proposed what turned out to be the correct explanation for the continuous energy spectrum of the electron emitted in beta decay. He conjectured that there actually *was* a third particle, which he called the *neutron*, on the right hand side of (2.1.0) and that it had not been detected on account of being electrically neutral. That is, he proposed that (2.1.0) should be replaced by,

$$A \rightarrow B + e^{-} + \bar{\nu} \quad 2.1.2$$

where ν denotes his hypothesized neutral particle (the bar notation is inserted in hindsight to account for the fact that it is an antineutrino).

The name *neutron* was eventually used for another particle when, in 1932, James Chadwick announced his discovery of a massive neutral particle in a letter submitted to *Nature*, entitled *Possible Existence of a Neutron* [4]. While this particle was not Pauli's neutral particle, it turned out to be useful in the discovery of his particle. Indeed, in 1933 Fermi developed a theory of beta decay making use of both the neutron, recently discovered by Chadwick, and Pauli's neutral particle which Fermi, now that the name neutron referred to Chadwick's particle, renamed the *neutrino*. His theory assumed that nuclei were made up of protons and the newly discovered neutrons and that beta decay proceeded by the following interaction,

$$n \rightarrow p + e^{-} + \bar{\nu} \quad 2.1.3$$

where n denotes a neutron in the original nucleus, p denotes a proton in the final nucleus, and where the entire process describes the decay of a neutron in the original nucleus into a proton in the final nucleus (which successfully described the observed change of identity in the nuclei that underwent beta decay) with the emission of an electron and a neutrino.

In the following decades Fermi's theory enjoyed remarkable predictive success as beta decay energy spectra predicted by his theory were in excellent agreement with the measured energy spectra of electrons emitted in decays of many nuclei. This success had many physicists convinced of the existence of the neutrino and, indeed, many physicists used his theory to develop theories of their own.

One such theory of great importance in the history and in the field of neutrinos is the mechanism of thermonuclear fusion in the Sun that Bethe & Critchfield developed throughout the 1930s [6]. Arthur Eddington had originally proposed thermonuclear fusion as the primary source of stellar energy in 1920 making use of Einstein's famous equation from special relativity, however, the theoretical tools required for a complete description of the process did not exist until Fermi published his theory of beta decay. Once he had done so, Bethe & Critchfield put his theory to use in order to formalize Eddington's hypothesis.

Bethe and Critchfield described the complete description of thermonuclear fusion in the Sun by way of eleven interactions. The interactions show how hydrogen clouds (protons) interact to create alpha particles. The interactions are the following,

$$p + p \rightarrow d + e^+ + \nu_e \quad 2.1.4$$

$$p + p + e^- \rightarrow d + \nu_e \quad 2.1.5$$

$$d + p \rightarrow {}^3\text{He} + \gamma \quad 2.1.6$$

$${}^3\text{He} + p \rightarrow \alpha + e^+ + \nu_e \quad 2.1.7$$

$${}^3\text{He} + {}^3\text{He} \rightarrow \alpha + p + p \quad 2.1.8$$

$${}^3\text{He} + \alpha \rightarrow {}^7\text{Be} + \gamma \quad 2.1.9$$

$${}^7\text{Be} + e^- \rightarrow {}^7\text{Li} + \nu_e \quad 2.1.10$$

$${}^7\text{Li} + p \rightarrow \alpha + \alpha \quad 2.1.11$$

$${}^7\text{Be} + p \rightarrow {}^8\text{B} + \gamma \quad 2.1.12$$

$${}^8\text{B} \rightarrow {}^8\text{Be}^* + e^+ + \nu_e \quad 2.1.13$$

$${}^8\text{Be}^* \rightarrow \alpha + \alpha \quad 2.1.14$$

and the total set of reactions is often referred to as *the proton-proton chain*.

The calculations for the amount of energy released in these interactions modelled the measured energy generated within the Sun remarkably well. They calculated an energy production of 2.2 erg/g·sec, where an *erg* is equal to 10^{-7} J, versus the then presently observed value of 2.0 erg/g·sec [6].

It is interesting to note that, even though they used Fermi's theory of beta decay, Bethe & Critchfield left the neutrino out of their work. Perhaps they were hedging their bets until a direct experimental confirmation of the existence of neutrinos had been published. They would have to wait for nearly two decades for such a confirmation when, in 1954, Cowan & Reines conducted an experiment in which they successfully observed neutrinos interacting with matter [5].

While the success of Fermi's theory in explaining particle decays had many physicists convinced of the existence of neutrinos, Cowan & Reines maintained the opinion that neutrinos must be directly observed interacting with matter before Fermi's theory of beta decay could be accepted. They put Fermi's theory to the test by making use of crossing symmetry and the fact that (2.1.3) implies the existence of the following interaction under the rules of the SM,

$$\bar{\nu} + p \rightarrow n + e^+ \quad 2.1.15$$

where $\bar{\nu}$ denotes an anti-neutrino, and where the interaction is referred to as inverse beta decay. The difference between (2.1.3) & (2.1.15) is clearly that one cannot observe *any* of the particles on the right hand side of the latter unless a neutrino has interacted with matter. In this sense it is a direct observation of neutrinos whereas (2.1.3) is passive.

The reader may have noticed that there is still a neutral particle on the RHS of (2.1.15) and that this presents essentially the same difficulty as was found in the original problem. The difference, of course, is that neutrons interact with matter far more often than do neutrinos and therefore all particles on the RHS of (2.1.15) are readily observed. Indeed, the detection mechanism employed by Cowan & Reines for doing this was beautifully simple and can be seen in the figure below.

Figure 2.2: Detection mechanism used by Cowan and Reines. They observed inverse beta decay (predicted by applying crossing symmetry to Fermi's beta decay) by the detection of a positron and a neutron in succession after the collision of an antineutrino with a proton. The neutron undergoes neutron capture in Cadmium, resulting in the emission of photons from the excited cadmium nucleus. The positron annihilates with one of the electrons that are bound to the Cadmium nuclei in the detection volume, resulting in the emission of photons.

We now return to Bethe & Critchfield's mechanism for describing thermonuclear fusion as an energy source in the Sun. In 1968, long since Cowan and Reines made the first detections of neutrinos, Davis, in his so-called Homestake experiment, attempted to take advantage of the by then known interactions of neutrinos with various particles of matter to put Bethe and Critchfield's theory of thermonuclear fusion to the test.

The relevance of neutrinos in such a test is the fact that most particles produced in the proton-proton chain interact frequently on their way to the surface of the Sun. Thus, all information about their origins is lost by the time that they reach Earth. However, neutrinos interact only weakly and very infrequently and most of them make it to Earth without interacting with matter at all. This means that very little information about their origins is lost on their journey towards detectors on Earth. Thus, now that physicists had been directly measuring neutrino interactions for about a decade, it was natural to use the neutrino products in the proton-proton chain to observe the inner workings of the Sun. Using Bethe and Critchfield's model, the neutrino flux could be calculated and since neutrinos only interact weakly, they would easily propagate towards Earth for measurement. Davis and Bahcall used chlorine detectors to take advantage of the measurement interaction proposed by Pontecorvo and Alvarez: $^{37}\text{Cl} + \nu_e \rightarrow ^{37}\text{Ar} + e^-$. Interestingly, Davis & Bahcall only measured one third of the neutrino flux

expected by the proton-proton chain – a strange result considering the success of the model in reproducing the measured energy production in the Sun. This severe discrepancy, whose measurement won Davis the Nobel Prize in 2005, came to be known as the *Solar Neutrino Problem*. It ushered in a new era of neutrino physics after initially being met with scepticism, if not indifference.

The scepticism was founded primarily on two points. One was that Davis' experiment required him to be able to successfully measure one Argon atom per interaction in a 615-tonne tank of cleaning fluid. Perhaps his detector was simply not sensitive enough to measure the total flux. The second criticism was related to the energy threshold of the detection mechanism. Davis' experiment was sensitive to high energy neutrinos produced in the *proton-proton chain* and the flux of high energy neutrinos is sharply dependent on the assumptions of the so called *Standard Solar Model* (SSM). Well, Bahcall predicted the high energy neutrino flux using the SSM and the proton-proton chain. Therefore, perhaps the intricate details of the SSM were slightly incorrect, producing a sharp and incorrect effect in the tuning of the flux of high energy neutrinos. The first criticism was refuted when the Kamioka detector and the Super-Kamioka detector also reported deficits in the solar neutrino flux. The second criticism was refuted by the results of several Gallium-based experiments. A detector based on the inverse beta-decay of Gallium-71 was able to probe low energy neutrinos (down to 0.233 MeV). At such energies, the solar neutrino flux is essentially independent of the details of the SSM. Therefore, if gallium detectors also observed neutrino deficits the Solar Neutrino Problem would have to be taken seriously. Alas, several gallium experiments reported deficits in the solar neutrino flux and the second criticism of the Solar Neutrino Problem had been refuted.

The energy threshold for detection in various detection media are shown in *Figure 2.3*. In it we can see that the gallium experiments probed much lower energy ranges than previous experiments.

With this result, physicists could no longer remain indifferent to the observed discrepancy. Something needed to give beyond the SSM. That something was a beautifully simple solution proposed by Bruno Pontecorvo – neutrino oscillation.

Figure 2.3: Solar neutrino flux as a function of energy. Above the figure are ranges of energies to which various experiments & detection schemes were sensitive during their lifetimes. Adapted from [9].

2.2 Neutrino Oscillations

As we have seen, the Solar Neutrino Problem was originally not particularly troublesome to most physicists. Davis had to extract 33 argon atoms from a tank of 615 tonnes of tetrachloroethylene. This is a miniscule signal in comparison to the size of the tank and, as explained in the previous section, most of his peers agreed that he had simply missed some argon atoms. Even if he hadn't, the calculations performed by Bahcall for the expected neutrino flux as a function of energy relied heavily on the SSM which had yet to be experimentally confirmed as the correct process of fusion interactions in the Sun. However, Bethe and Critchfield also used the SSM along with their p-p chain and, as stated previously, predicted the energy generation in the Sun with great accuracy and so the Solar Neutrino Problem could not be entirely discarded. Indeed, over time and with further experiments reporting similar results, the discrepancy became a dilemma in the field. The last blow to the criticism of the SSM came with the results of gallium experiments.

In 1968, Pontecorvo proposed a phenomenon of neutrinos that made use of newly discovered neutrino flavours. Indeed, in 1962, researchers at BNL reported the first observations of muon-neutrinos [10] and, although the existence of tau-neutrinos was not confirmed until 2001 by the DONUT experiment [11], it was generally assumed that they existed. Pontecorvo's proposed mechanism turned out to be the correct solution to the Solar Neutrino Problem. He suggested that neutrinos propagate as eigenstates of the free-particle Hamiltonian – i.e. as the mass eigenstates ν_1, ν_2, ν_3 – but that they interact with matter as weak eigenstates (a.k.a flavour states, in the case of neutrinos). These flavour eigenstates, ν_e, ν_μ, ν_τ , are not states of definite mass but are linear combinations of mass eigenstates whose coefficients oscillate in a complicated manner. If this were true, then the electron neutrinos to which Davis and Bahcall's experiment was sensitive could have oscillated into other flavour states. Now, their experiment was based on observing a *charged* weak current on a target which is made of only proton and neutrons. Any muon and tau neutrinos that were produced through oscillations would only interact with this target through the *neutral* weak current. However, their experiment was not sensitive to neutral current interactions, so any such events would have been missed by their detectors.

We can illustrate such a mixing/oscillation process by assuming a simpler case, such as that used by Griffiths [4], in which there are only two neutrino flavours. He starts by defining two orthogonal mass states as linear combinations of the observable flavour states,

$$\nu_1 = \cos(\theta) \nu_e + \sin(\theta) \nu_\mu \quad 2.2.1$$

$$\nu_2 = -\sin(\theta) \nu_e + \cos(\theta) \nu_\mu \quad 2.2.2$$

Assuming a neutrino starts out in a ν_e flavour state, following Griffiths' derivation, one arrives at the following form for the probability of the neutrino having oscillated into the ν_μ flavour state after travelling some distance, L ,

$$P_{\nu_e \rightarrow \nu_\mu} = \sin^2(2\theta) \sin^2\left(\frac{m_2^2 - m_1^2}{4E} L\right) \quad 2.2.3$$

where E is the energy of the neutrino, and θ is the mixing angle between the states.

The three flavour case is more mathematically tedious but the concepts remain the same, apart from the addition of a CP-violating Dirac phase that cannot be removed as in the two-neutrino case. In the three flavour case, the flavour states are defined by,

$$\nu_i = \sum_{m=1}^3 U_{ij}^* \nu_j \quad 2.2.4$$

where ν_i is the i^{th} flavour state, ν_j is the j^{th} mass eigenstate and U is the so-called Pontecorvo-Maki-Nakagawa-Sakata mixing (U_{PMNS}). This rotation matrix is defined by,

$$U_{\text{PMNS}} = \begin{pmatrix} c_{13}c_{12} & c_{13}s_{12} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{\frac{i\varphi_{21}}{2}} & 0 \\ 0 & 0 & e^{\frac{i\varphi_{31}}{2}} \end{pmatrix} \quad 2.2.5$$

where c_{ij} (s_{ij}) is the cosine(sine) of a rotation angle, θ_{ij} , connecting the i^{th} flavour state to the j^{th} mass state. The constant, δ , is a CP violating phase. If the neutrino turns out to be of Dirac type then this is the only CP violating phase. For the case in which the neutrino is of Majorana nature, there are two additional CP violating phases, φ_{21} , φ_{31} . These are included along the diagonal of the rightmost matrix in (2.1.16). For the case in which the neutrino is of Dirac nature, these two phases are equal to zero and the rightmost matrix becomes the identity (this is the reason for writing it as a separate matrix).

Now, if the reader is following the timeline of discoveries in neutrino physics, they will notice that by the time the Solar Neutrino Problem started to be taken seriously, Pontecorvo had already proposed neutrino oscillations. However, this proposed phenomenon was very much like Pauli's neutrino; it was a useful mathematical tool for explaining experimental results, but nobody had ever directly observed the phenomenon. Thus, many experimentalists began to develop and construct detectors that could search for oscillations directly. The two major experiments that eventually confirmed the existence of oscillations between weak eigenstates of neutrinos were observed directly in the early 2000s by

the Super-Kamiokande and SNO experiments. By directly we mean that *changes* in neutrino flavour were detected as opposed to earlier experiments which were just designed to detect neutrinos. The Super-Kamiokande experiment did this via the detection of Cerenkov light in water. This detection mechanism is sensitive to all three neutrino flavours and is described by neutrino-electron scattering,

$$\nu + e^- \rightarrow \nu + e^- \quad 2.2.7$$

where the electrons are bound to water molecules. After scattering has occurred, these electrons produce Cerenkov light. The orientation of the conical Cerenkov profile serves as an excellent discriminator for backgrounds. Assuming that neutrinos did not oscillate, only 45% of the expected flux was observed. However, detection of electron neutrinos was 6.5 times more efficient than the detection of other flavours under (2.2.7). Therefore, applying this correction and using Homestake data for determining which fraction of the detected events were electron-neutrinos, SuperK observed a flux that was very close to the expected Solar flux. Still, the conditions of the Homestake experiment were different and such a comparison came with non-negligible systematic error. The case for the existence of oscillations was, therefore, flawed.

SNO was the experiment that finally observed all flavours separately with a so-called two-in-one detector. This detector employed heavy water as a detection medium. The advantage of this medium is that it can detect charge current weak interactions which is dominated by an electron-neutrino interaction (2.2.8) and a neutral current weak interaction (2.2.9) in which all three neutrino flavours may participate (mechanism (2.2.7) is, of course, also possible in heavy water),

$$\nu_e + d \rightarrow p + p + e^- \quad 2.2.8$$

$$\nu + d \rightarrow p + n + \nu \quad 2.2.9$$

where d denotes a heavy water nucleus (neutron and proton). SNO measured a flux of 1.76 ± 0.11 for the CC channel and a flux of 5.09 ± 0.65 for the NC channel (both in units of

millions of neutrinos per square centimetre per second) [1]. Therefore, they concluded that 3.33 ± 0.66 electron-neutrinos produced in the p-p chain had oscillated into other flavours. It is worth noting that the p-p chain combined with the SSM at the energy threshold of SNO would have predicted a total flux of 5.05 ± 0.40 . Therefore, in one sweep, SNO had confirmed that neutrinos oscillate and had provided one of the best confirmations of the SSM.

As described earlier, that neutrinos oscillate suggest that they have mass. Neutrino oscillation experiments are, however, only sensitive to the differences between the squared masses of the mass eigenstates, as suggested by (2.2.3).

Parameter	best-fit	3σ
$\Delta m_{21}^2 [10^{-5} \text{ eV}^2]$	7.37	6.93 – 7.96
$\Delta m_{31(23)}^2 [10^{-3} \text{ eV}^2]$	2.56 (2.54)	2.45 – 2.69 (2.42 – 2.66)
$\sin^2 \theta_{12}$	0.297	0.250 – 0.354
$\sin^2 \theta_{23}, \Delta m_{31(32)}^2 > 0$	0.425	0.381 – 0.615
$\sin^2 \theta_{23}, \Delta m_{32(31)}^2 < 0$	0.589	0.384 – 0.636
$\sin^2 \theta_{13}, \Delta m_{31(32)}^2 > 0$	0.0215	0.0190 – 0.0240
$\sin^2 \theta_{13}, \Delta m_{32(31)}^2 < 0$	0.0216	0.0190 – 0.0242

Table 2.1: Best measurements and 3σ range of neutrino mixing parameters as of 2018. Table adapted from [23]. Values in brackets correspond to the inverted hierarchy while those outside brackets correspond to the normal hierarchy.

Oscillation experiments are only sensitive to Δm_{ij}^2 , so, although oscillations prove that neutrinos are massive, the absolute mass scale of the neutrinos remains unknown. Presently, the measured energy distribution of the electron points to the electron-neutrino mass being extremely small [18]. Even detectors of state of the art resolution still measure emitted electron energies up to the value given by (2.1.1). Indeed, it remains unknown to this day and it has only been recently observed that neutrinos have mass at all, as mentioned earlier. Experiments such as KATRIN, based in Germany, are aiming to measure the endpoint of the beta decay spectrum. KATRIN is expected to be sensitive enough to set an upper limit for the lowest neutrino mass at 0.1 meV.

2.3 Mass Hierarchy

In section 2.2 it was stated that oscillation experiments are sensitive not to absolute neutrino masses, but to the differences in the squares of their masses. That is, oscillation experiments have observed that neutrinos are indeed massive but, as of yet, have not observed the absolute masses of the mass eigenstates. Additionally, past oscillation experiments have not been able to determine the order of the mass eigenstates from lightest to heaviest – the mass hierarchy. The positive value of $7.37^{+0.59}_{-0.44} \cdot 10^{-5} eV^2$ for Δm_{21}^2 of *Table 2.1* signifies that the second mass eigenstate is heavier than the first. However, neutrino oscillation experiments have, as of yet, not been sensitive to the sign of Δm_{31}^2 (hence the multiple options for its value in *Table 2.1*). As a consequence of this limitation, which of the following cases is true is unknown: that $m_3 > m_2$ or that $m_3 < m_1$. The former case is referred to as the *normal* mass hierarchy, while the latter is referred to as the *inverted* mass hierarchy. We call the former case the normal hierarchy because it mimics the mass hierarchy of the charged leptons – the electron, the muon, and the tauon. Of course, the numbering of mass eigenstates is arbitrary so the reader may conclude that the ordering is irrelevant. However, it is known that the tau-neutrino is mostly made up of the m_3 mass eigenstate, and the electron-neutrino is mostly made up of the m_1 mass eigenstate. Thus, if $m_3 > m_2$ then the mass hierarchy of the flavour eigenstates of the neutrinos mimics that of the charged leptons. The two possibilities for the mass hierarchy are neatly summarized graphically in *Figure 2.4*.

Figure 2.4: Candidate mass hierarchies for the mass eigenstates of the neutrinos and the flavour compositions of the mass eigenstates. The normal hierarchy has m_3 significantly more massive than m_2 while the inverted hierarchy has m_3 significantly less massive than m_1 . In either scheme, m_1 & m_2 have a much smaller separation from each other than from m_3 . The degenerate case is one in which the mass differences are much smaller than the absolute masses of the neutrinos and in which the splitting can be ignored for most calculations. It can have a normal or inverted hierarchy as long as this condition is met. Cosmological studies, most recently the Planck telescope, have essentially ruled out this possibility regardless of the nature of neutrinos with an upper limit of the smallest neutrino mass of $<0.12\text{eV}$ [18]. However, such results are model dependent. KATRIN will be able to rule this out definitively.

The colour scale is a representation of the flavour composition of the mass eigenstates and is the same for both candidate hierarchies (refer to caption for explanation of third hierarchy).

In order to determine the correct hierarchy for a theory of neutrinos, oscillation experiments must turn to long-baseline experiments where the interaction of neutrinos with matter makes it possible to determine the signs of all squared mass differences [13]. Still, even with long-baseline approaches, oscillations experiments will never be able to determine the absolute mass scale of neutrinos. For this, physicists must look to alternative approaches. One approach to answering this question, and possibly determining the hierarchy simultaneously, is the measurement of a hypothetical interaction called neutrinoless double beta decay ($0\nu\beta\beta$). This interaction is considered in the next section.

2.4 Neutrinoless Double Beta Decay

One process of particular interest in the field of neutrino physics is double beta decay, in which two neutrons simultaneously decay to one proton each, by each emitting an anti-electron neutrino and an electron (see *Figure 2.5*). The double beta decay interaction was first

formalized mathematically by Maria Goeppert-Mayer in 1935 [14]. Shortly thereafter, in 1937, Ettore Majorana showed that all predictions of beta decay theory would remain unchanged if the neutrino were its own antiparticle [15]. Particles that are their own particles are now usually referred to as Majorana particles and henceforth we will refer to them as such. The combination of these two ideas allowed W.H. Furry, in 1939, to point out that, if neutrinos were in fact Majorana particles, then double beta decay should have a mode in which no neutrinos are emitted [3]. This is referred to as $0\nu\beta\beta$ while Goeppert-Mayer's interaction is often referred to as two-neutrino double beta decay ($2\nu\beta\beta$), as two neutrinos are emitted from the original nucleus.

It is important to clarify that the double beta decay interaction is not two separate beta decays and, in general, the branching ratio for this process is greatly outweighed by that of two single beta decays.

Figure 2.5: Double beta decay.

However, there exist nuclei in which single beta decay is energetically forbidden while double beta decay is energetically favourable. One such nucleus is ^{136}Xe . The mass parabola for isobars of 136, shown in Figure 2.6, demonstrates that double beta decay of ^{136}Xe to ^{136}Ba is energetically favourable while single beta decay to ^{136}Cs is forbidden. This allows for a clean detection of the former decay process.

Figure 2.6: Mass parabola for isobars of 136. Single beta decay of ^{136}Xe is expressly forbidden while double beta decay is energetically favourable. The parabola is computed using the semi-empirical mass formula.

Now that we can imagine nuclei for which it is realistic to consider the detection of double beta decay due to suppression of single beta decay, it is worth looking into the hypothetical neutrinoless mode. By crossing symmetry, the emission of an antineutrino is essentially the same process as the absorption of a neutrino. That is, if one of the following decay processes is allowed, then it is implied that the other is allowed as well,

$$n \rightarrow p + e + \bar{\nu}_e \quad 2.4.1$$

$$n + \nu_e \rightarrow p + e \quad 2.4.2$$

One can thus propose a modification to the Feynman diagram for double beta decay illustrated in *Figure 2.5*. This modification produces what is called $0\nu\beta\beta$ and the corresponding Feynman diagram is shown in *Figure 2.7*.

Figure 2.7: Neutrinoless double beta decay ($0\nu\beta\beta$). The demarcation of the vertices as *SM Vertices* implies the strictly left-handed weak current of the SM. The image on the right is adopted from [17].

The modification that produces the interaction of *Figure 2.7* takes advantage of the equivalence between (2.4.1) and (2.4.2). In the modification we assume that one neutron decays by the process of (2.4.1), emitting an antineutrino. The second neutron then absorbs the neutrino and decays by (2.4.2). In order for this two part decay to be analogous to double beta decay, the neutrino must remain virtual such that the decay of each neutron occurs simultaneously. Furthermore, the second nucleus should not be able to absorb the neutrino unless the neutrino is its own antiparticle. If it is not its own antiparticle, then the neutrino in (2.4.2) would be an antineutrino and could not be absorbed in the second nucleus due to the left-handedness of the weak charged current interaction.

However, if we assume that neutrinos are their own antiparticles then, under the Majorana formalism, the antineutrino emitted on the left is emitted with mostly right-handed helicity but still a small left-handed component. The nucleus at the second vertex is able to absorb the left-handed component. Incidentally, if this occurs, it would suggest that lepton number is not a strictly conserved quantity in nature.

This mixing of right-handed and left-handed antineutrinos and neutrinos is addressed by the Majorana mass term in the Lagrangian of neutrinos under the Majorana formalism,

$$\mathcal{L}_M \sim m_M \bar{\nu}_L \nu_L^C \quad 2.4.5$$

where ν_L^C accounts for the absorption of a right-handed antineutrino and where $\bar{\nu}_L$ accounts for the creation of a left-handed neutrino. This is quite clearly what is needed as the mass insertion

of $0\nu\beta\beta$ in order to account for charge conjugation and the flip in handedness. Here it has been assumed that there is no Dirac mass component in the Lagrangian. If the Lagrangian is generalized to contain Dirac mass terms in addition to the Majorana mass terms, $0\nu\beta\beta$ becomes more complex. However, because a charge conjugation must occur, observation of $0\nu\beta\beta$ must be proportional to the Majorana mass term. Indeed, the process cannot occur with only a Dirac mass as Dirac mass terms only account for a flip in handedness in the Majorana formalism and not for charge conjugation (unless one is willing to add right-handed weak currents as an addition to the SM). In general, any $0\nu\beta\beta$ process will include an odd number of Majorana mass insertions and an odd number of Dirac mass insertions. This is shown elsewhere [17].

The mass insertions are shown explicitly in the following Feynman diagram,

Figure 2.8: Feynman diagram for $0\nu\beta\beta$ with Majorana mass term shown explicitly. The charge conjugating effect of the mass term as well as the flip of handedness allows both weak vertices to be of the typical left-handed SM type. Figure from [17].

Showing the mass insertion explicitly demonstrates the amplitude of the $0\nu\beta\beta$ process is proportional to the effective Majorana mass. Indeed the half-life of the process is defined by,

$$\frac{1}{T_{1/2}^{0\nu}} = G^{0\nu} |M^{0\nu}|^2 \frac{\langle m_\nu \rangle^2}{m_e^2} \quad 2.4.5$$

where $G^{0\nu}$ is the phase space factor as is known exactly for a given Q-value (i.e. the energy released in the decay of the ^{136}Xe nucleus), $M^{0\nu}$ is the model-dependent nuclear matrix element, and $\langle m_\nu \rangle$ is the effective Majorana neutrino mass defined by,

$$\langle m_\nu \rangle = \left| \sum_{i=1}^3 U_{ei}^2 m_i \right| \quad 2.4.6$$

where U denotes the PMNS matrix and where m_i denotes the i th mass eigenvalue. The e subscript on the PMNS matrix ensures that only the electron-neutrino mixing components are included.

This inverse proportionality between Majorana mass and half-life is what motivates the use of passive counting experiments in the search for $0\nu\beta\beta$ – one simply waits and counts the rate of $0\nu\beta\beta$ events. If any events are observed we can conclude that the neutrino is Majorana in nature and we can also estimate the masses of the neutrinos and therefore determine the hierarchy. Such searches, in particular the Enriched Xenon Observatory (EXO) are discussed in Chapter 3.

2.5 Energy Spectra of $\beta\beta$ -Decays

Both $0\nu\beta\beta$ and $2\nu\beta\beta$ are extremely rare processes (see *Table 3.1*). Because of their rarity, it is a hopeless endeavour for the experimentalist to discriminate between the two based on detecting, or not detecting, neutrinos after the decay. Instead, the experimentalist must consider the difference in the energy spectra of the two types of $\beta\beta$ -decay.

In $2\nu\beta\beta$, the two outgoing neutrinos carry away some of the decay energy while the outgoing electrons carry the rest. The result is that the sum of the energies of the two electrons, which are easily detected compared to the neutrinos, follows a continuous distribution where the maximum energy is just below the Q -value – the energy released in the decay. The shape of the distribution can be described by the Primakoff-Rosen approximation [16],

$$\frac{dN_{2\nu}}{dK} \sim K(Q - K)^5 \left(1 + 2K + \frac{4}{3}K^2 + \frac{K^3}{3} + \frac{K^4}{30} \right) \quad 2.5.1$$

where K is the sum of the kinetic energy of the two outgoing electrons, and Q is the Q -value of the decay in units of electron mass.

For the $0\nu\beta\beta$ mode of the decay, the two outgoing electrons carry away all of the decay energy and the distribution of the sum of their kinetic energies is a delta function at the Q-value,

$$\frac{dN_{0\nu}}{dK} \sim \delta(K - Q) \quad 2.5.2$$

Due to finite detector resolution, both of these distributions will be smeared in practice and will overlap. The overlap between (2.5.1) and (2.5.2) is shown in *Figure 2.8* after applying various energy resolutions.

Figure 2.9: Left: $2\nu\beta\beta$ energy distribution (broad) v.s. $0\nu\beta\beta$ (delta function distribution). Right: Overlap between energy distributions of $2\nu\beta\beta$ and $0\nu\beta\beta$ for different energy resolutions. EXO-200 currently operates with a resolution of 1.6% near the Q-value (see *Figure 3.13*). In these plots the Q-value is denoted by T_0 . Figure adapted from [21].

Chapter 3

Searches for Neutrinoless Double-Beta Decay

In *Chapter 2* we began by placing the neutrino in its historical context from its inception in the mind of Pauli, to the theories describing its interactions with matter put forward by Fermi and Pontecorvo. We saw how these theories were tested and confirmed experimentally from experiments by Cowan & Reines, to the discovery of neutrino oscillations at SNO. Together, these two theories describe what we know about neutrinos.

We ended the chapter by discussing the open questions that surround the neutrino – things not explained by Fermi and Pontecorvo – after the discovery at SNO. In particular, we saw that the mass scale & mass hierarchy and, finally, the nature (Majorana or Dirac) of the neutrino are open questions whose answers may lie beyond the SM. It was seen that the hypothetical process $0\nu\beta\beta$, if observed, would answer all of these questions simultaneously. Its existence would directly show that neutrinos are Majorana particles, while the half-life of the process would provide a measurement for the effective Majorana mass, ultimately giving an implication of the correct structure of the mass hierarchy.

In this chapter, we consider the experimental side of $0\nu\beta\beta$. We will begin by describing the most direct and most common type of experiment its detection – counting experiments. Of the many counting experiments currently searching for $0\nu\beta\beta$, we will consider the Enriched Xenon Observatory (EXO) – a collaboration consisting of two experiments, namely EXO-200 and nEXO – in particular. We will end the chapter with a discussion of the background limitations of EXO compared to those of other counting experiments and list approaches to eliminating or minimizing these limitations.

3.0 – Counting Experiments

A simple method for observing $0\nu\beta\beta$ is the *counting experiment* – a passive type of experiment. The objective is to monitor a volume of some source medium which is expected to undergo a decay interaction of interest. In our case, the interaction of interest is $0\nu\beta\beta$. The detector that houses the source medium attempts to observe and record all such decays (i.e. the detector *counts* the number of events that occur with their origin within the source medium).

Discoveries in counting experiments are made by detecting an excess, with some chosen significance level, in the total number of events that would be expected if the discovered event did not exist. In a counting experiment, the experimenter typically defines some region of interest (ROI) in the observed parameter space in which the largest fraction of signal is expected and where the signal to background ratio is favourable. The ROI is refined by making cuts that suppress backgrounds. More complex analysis techniques can be used to expand the ROI, minimizing the fraction of signal events that are left out, without increasing the signal to background ratio. In other words, better analysis techniques can improve discriminating power between signal and background. Some of these analyses are summarized in *Chapter 5*.

A counting experiment that is looking for new physics should be designed to push the limit that is placed on an event rate by current physics. In order to understand the extent to which the design of a counting experiment can do this, one must understand the sensitivity and discovery potential of the design.

3.0.1 – Sensitivity

In the context of a counting experiment, *sensitivity* is defined as the lower limit that will be placed on an event rate if a discovery of the event is not declared.

To get this value, the experimenter will define a region of interest (ROI) in parameter space, where the parameters include any/all measured quantities in the experiment. This ROI will be defined by some signal-to-background ratio.

Using the Feldman-Cousins prescription [20] for estimating the upper limit on the total number of recorded signal events in the ROI, the sensitivity is defined by the value of s_{up} (the upper limit the number of signal events observed) such that,

$$0.5 = \sum_{s'=0}^{s_{up}} R(N = s' + b|b, s) \quad 3.0.1$$

where $R(N|b, s)$ is the ratio of two likelihoods if one adopts the Feldman-Cousins prescription for constructing confidence intervals, and N is the observed number of events in the ROI given an expected signal, s , and a fixed background, b . For the case of counting experiments the PDFs in the ratio, R , are described by the Poisson distribution with a fixed background, b ,

$$p(N = s' + b|s, b) = \frac{(s + b)^{s'+b} e^{-s-b}}{(s' + b)!} \quad 3.0.1$$

and where the ratio itself is given by,

$$R = \frac{p(N|b, s)}{p(N|b, s_{best})} \quad 3.0.2$$

The parameter s_{best} is then determined as the value for s' that maximizes $p(N|b, s)$ while requiring the physical constraint defined by,

$$s_{best} = \max(0, N - b) \quad 3.0.3$$

In practice, b , is not known exactly but is also estimated. This produces additional uncertainties in the calculation of (3.0.1), the intricacies of which are described elsewhere [21][22].

In other words, (3.0.1) tells us the median value of s_{up} that would be measured in a series of identical experiments. The choice of using the median, if not somewhat intuitive, is convention.

Now, in the context of $0\nu\beta\beta$ experiments, s_{up} is inversely proportional to the lower limit on the half-life (i.e. if we see more events, then the half-life is shorter and vice versa) and the lower limit on the half-life tends to be what is quoted for a measurement of sensitivity as opposed to s_{up} . The relationship between the upper estimate and the lower limit for the half-life is given by,

$$T_{\frac{1}{2}}^{0\nu} = \epsilon \frac{\ln 2}{s_{up}} \frac{f m N_A}{M} t \quad 3.0.4$$

where ϵ is the detection efficiency of the counting experiment, f is the enrichment fraction of the isotope of interest, m is the total mass of the source volume, N_A is Avogadro's number, M is the molar mass of the isotope of interest, and t is the running time of the experiment. Upon substitution into (2.2.4),

$$\langle m_{\beta\beta} \rangle_{up} = \sqrt{\frac{s_{up} m_e^2 M}{\epsilon f m N_A t G^{0\nu} |M^{0\nu}|^2 \ln(2)}} \quad 3.0.5$$

Figure 3.1: Projected sensitivities for several running and proposed $0\nu\beta\beta$. The bands in each plot correspond to effective Majorana masses of the neutrino with smearing due to uncertainties in the nuclear matrix elements of various models. Adapted from [22].

Some of the $0\nu\beta\beta$ experiments that are represented in the figure above are discussed in section 3.1. There we briefly consider their detection mechanisms and resulting backgrounds in the ROI.

3.0.2 – Discovery Potential

In the context of counting experiments, *discovery potential* is defined as the smallest observable event rate for which a discovery can be claimed at a specified level of significance. In other words, given a particular background rate in the ROI and for some chosen acceptable level of significance, what is the minimum amount of signal events (i.e. the *least detectable*

signal) that must be detected in order to claim that we have, in fact, seen additional events above the expected background rate? How unlikely must it be for all events recorded in the ROI to have occurred from backgrounds? If the observed number of events in the ROI is denoted by $N = s+b$ and the least detectable signal by s_L , then s_L is determined by satisfying the following equations,

$$3 \times 10^{-7} = 1 - \sum_{N=0}^{N_{cut}-1} P(N|H_0 : s = 0) \quad 3.0.6$$

$$0.5 = \sum_{N=0}^{N_{cut}} P(N|H_i : s_L) \quad 3.0.7$$

where (3.0.6) is satisfied by varying N_{cut} and (3.0.7) is satisfied by substituting this N_{cut} into the equation and determining s_L . In both equations it is assumed that the expected number of background events is known.

For rare event searches, a Poisson distribution can be used for the form of P in the preceding expressions,

$$P(N|\mu = s + b) = \frac{\mu^N}{N!} e^{-\mu} \quad 3.0.8$$

Under this prescription, the discovery potential can be plotted as a function of the number of expected background events.

Figure 3.2: The least detectable signal within the ROI of any counting experiment for a given expected Poisson background. Figure from [21].

It should be noted that the background number refers to background events that fall within the ROI. In general, more background events can occur outside of the ROI as these are easily tagged as non-signal events by definition. As we will see further ahead, the discriminating power in the ROI can be improved through various approaches. The essential effect of each approach is to lower the value of b in the ROI and/or to expand the ROI in the parameter space without increasing the background thereby improving the signal to background ratio.

In a real counting experiment the background itself is not known but is also estimated. In broad terms this estimate is made by fitting to the observed data that falls outside of the ROI. It is assumed that no signal exists in this region and that any fitting will produce an estimate for the background processes only. This result is then applied to counts in the ROI by applying models to extrapolate the background fit into the ROI.

3.1 – Current Status of $0\nu\beta\beta$ Counting Experiments

Presently there are several counting experiments that are either collecting data or that are still in the R&D phase, that are attempting to detect $0\nu\beta\beta$. Between these collaborations, various sources and detector types are used. These choices affect the signal to noise ratio of each the ROI from experiment to experiment, which ultimately affects the half-life to which each experiment is sensitive. To understand the benefits and the difficulties of using particular candidate isotopes and detector designs, it is instructive to consider each of the main

approaches. We will briefly look into each of these approaches and compare their limiting half-lives based on background rejection capabilities. The EXO collaboration will be considered in more detail.

Isotope	$Q_{\beta\beta}$ (keV)	Isotopic abundance (%)	$T_{1/2}(0\nu\beta\beta)$ (10^{25} yr)	$T_{1/2}(2\nu\beta\beta)$ (10^{21} yr)
^{76}Ge	2039	7.61	> 5.3 [26]	$1.65^{+0.14}_{-0.12}$
^{82}Se	2995	8.73	> 0.036 [27]	0.092 ± 0.007
^{100}Mo	3034	9.63	> 0.11 [27]	0.0071 ± 0.0004
^{130}Te	2528	34.17	> 0.40 [25]	0.69 ± 0.13
^{136}Xe	2479	8.87	> 11 [24]	2.19 ± 0.06

Table 3.1 : Summary of the present status of $0\nu\beta\beta$ studies. The tabulated lower limits for the half-lives correspond to the best limits that have been published for each isotope. $Q_{\beta\beta}$ denotes the total energy released in $\beta\beta$ decay.

3.1.1 – Large Liquid Scintillators

KamLAND-Zen is the only currently operating large liquid scintillator experiment looking for $0\nu\beta\beta$. The design of this type of experiment consists of a volume of liquid scintillator loaded with a $\beta\beta$ candidate. The particular design of KamLAND –Zen is shown below where a balloon filled with liquid scintillator houses a second, smaller balloon filled with ^{136}Xe (a $0\nu\beta\beta$ candidate) dissolved into the liquid scintillator.

Figure 3.3: Schematic of the KamLAND-Zen $0\nu\beta\beta$ large liquid scintillator experiment.

The scintillator balloon is surrounded by a large volume of buffer oil which is used for reducing background particles from entering the scintillator volume. The oil is then surrounded by PMTs

for detecting the scintillation light, followed by a water shielding volume and a second layers of PMTs. The LLS design has low energy resolution in comparison to bolometers or Ge-diodes ($\sim 4\%$ in the ROI v.s. $\sim 0.2\%$). Nevertheless, due to the excellent shielding provided by the oil-buffer and water volumes, KamLAND-Zen has achieved a low background of $\sim 160 \text{ cts}/(\text{ton} \cdot \text{yr})$ leading to a half-life limit in ^{136}Xe of $T_{1/2}^{0\nu} > 1.07 \times 10^{26} \text{ yr}$ (90%), corresponding to an upper limit on the effective Majorana mass of $(0.061 \text{ eV} \leq \langle m_{\beta\beta} \rangle \leq 0.169 \text{ eV})$.

3.1.2 – Bolometric Detectors

Solid-state experiments are characterized by their high energy resolutions. This allows them to distinguish between event types to a high level of precision as particular interactions have a characteristic energy spectrum.

Some of the solid-state experiments that are searching for $0\nu\beta\beta$ include, CUORE, GERDA, and the Majorana Demonstrator. Each of these uses a solid state source of $\beta\beta$ -emitting isotopes.

The CUORE experiment uses TeO_2 crystals (using ^{130}Te) as the detection medium in an array of bolometers, which are high-resolution instruments for measuring temperature. In order to detect the small changes in temperature caused by single events, the bolometers must be cooled to near absolute zero temperatures (10 – 15 mK). As can be seen in *Table 3.2* later in this chapter, ^{130}Te is also a candidate $0\nu\beta\beta$ isotope. This is extraordinarily convenient in practice, as the detection medium and the source are the same and the bolometers' crystals can be grown from the $0\nu\beta\beta$ candidate. Another convenient characteristic of ^{130}Te as a source/detection medium is that its natural isotopic fraction is 34%. This reduces the need for enrichment and drastically reduces costs in comparison to $0\nu\beta\beta$ candidates with low isotopic fractions (see *Table 3.0*). In the figure below, the mechanism of the TeO_2 bolometers is shown.

Figure 3.4: A schematic of the bolometer employed at CUORE.

The purpose of the bolometers is to enable the detection of small changes in the temperature of the medium caused by single events. For example, the two electrons emitted in $0\nu\beta\beta$ have a characteristic combined energy (i.e. the Q-value of $0\nu\beta\beta$ in Te-130), which leads to a characteristic increase in temperature as measured by the bolometers. Combining this with CUORE's energy resolution of $\sim 0.2\%$, even the insuppressible $2\nu\beta\beta$ events are easily identified as background. Because the distribution of the sum of the kinetic energies of the electrons in $0\nu\beta\beta$ is sharply peaked at the Q-value, this allows for nearly complete background rejection using a tight cut around the Q-value. As a result, in theory, CUORE is able to set ever smaller upper limits on the Majorana neutrino mass by extending the duration of data collection – they don't need to worry about background counts overwhelming the parameter space around the Q-value as more data is collected. CUORE began taking data in 2017 and expects to be able to set a limit of $T_{1/2}^{0\nu} > 9.5 \times 10^{25} \text{ yr}$ (90 %), corresponding to an upper limit for the effective Majorana mass of $(0.05 \text{ eV} \leq \langle m_{\beta\beta} \rangle \leq 0.13 \text{ eV})$. Its present lower limit is $T_{1/2}^{0\nu} > 4.0 \times 10^{24} \text{ yr}$ [25].

3.1.3 – Germanium Diode Detectors

GERDA is another major experiment in the Majorana neutrino field. It makes use of bare ^{76}Ge semiconductor diodes as both the detector and source for double beta decay. It has an energy resolution of $\sim 0.2\%$, much like the bolometer experiments. The diodes are submersed in liquid argon for cooling to avoid thermal activation/leakage and for shielding. The LAr shield is supplemented by an outer shield of water. Within the ROI, GERDA has managed to obtain a

background index of $0.001 \text{ cts}/(\text{keV} \cdot \text{kg} \cdot \text{yr})$ producing field-leading measurement in ^{76}Ge of $T_{1/2}^{0\nu} > 5.3 \times 10^{25} \text{ yr}$ (90%) [26].

GERDA plans to continue taking data until an exposure of $100 \text{ kg} \cdot \text{yr}$ is obtained (exposure was $46.7 \text{ kg} \cdot \text{yr}$ as of beginning of 2018), corresponding to a half-life limit of $T_{1/2}^{0\nu} > 2 \times 10^{26} \text{ yr}$ (90 %), implying an upper limit for the effective Majorana mass of $(0.09 \text{ eV} \leq \langle m_{\beta\beta} \rangle \leq 0.29 \text{ eV})$.

Figure 3.5: GERDA detection principle. Electrons from beta-producing events are collected on electrodes in the Ge-76 diodes. Light is detected in LAr to tag gamma-producing events.

The Majorana Demonstrator, which began taking data in 2016, also uses Ge-76. Its main difference to GERDA is that the diodes are not submersed in LAr. Instead of LAr, the diodes are housed in vacuum cryostats shielded by a dense configuration of lead bricks and copper. This eliminates the leakage current drawback caused by ionization of LAr in the GERDA experiment. It has the same energy resolution as GERDA and is currently conducting radioassay studies to reduce backgrounds.

The Majorana collaboration and the GERDA collaboration plan to combine their efforts in a tonne-scale detector in the next generation of experiments. Combined, they are expecting a *total* background index of 3.5 cts/(ton · yr) in the ROI.

3.1.4 – Tracking Calorimeters

The largest tracking calorimeter experiment is SuperNEMO. The working principle for this experiment is shown schematically in the figure below.

Figure 3.6: The working principle of SuperNEMO. Thin foils of various $\beta\beta$ emitters are placed between tracking calorimeters and scintillators. The calorimeters and scintillators record the total energy of events while the calorimeter enables vertex locating and complete event reconstruction. Ultimately this enables measurements of the $0\nu\beta\beta$ angular distribution cross section.

The difficulty with this design is that electrons originating from the source foil lose a significant portion of their energy as they pass through the source foils themselves. Furthermore, they employ plastic scintillators whose energy resolutions are $\sim 3\text{--}4\%$, much worse than the resolution of bolometer and diode experiments. This significantly reduces the energy resolution. SuperNEMO reduces this problem by employing thin foils but it cannot be completely mitigated. Furthermore, the thin film solution inhibits scaling to larger mass experiments.

The main advantage of this kind of experiment is the modular design of the foils that allows for the simple installation of various $\beta\beta$ emitters. Additionally, the calorimeters allow for track reconstruction which can be used for background suppression. Not only this, but track reconstruction also allows for measurement of the angular distribution of events which can, in

principle, be used to compare competing theoretical models for $0\nu\beta\beta$ (e.g. IBM-2, QRPA, etc..). No other existing detector designs can do this, barring the development of new analysis techniques.

Presently, the NEMO collaboration SuperNEMO has given the best limit for the half-life of Mo-100 and Se-82 (see *Table 3.1*).

3.1.5 – Time Projection Chambers

In addition to large liquid scintillator experiments like KamLAND-Zen, liquid/gas $\beta\beta$ emitters can also be studied with time projection chambers (TPCs). A TPC is a particular kind of drift chamber in which the electric field is constant along one axis and zero along all other axes.

Figure 3.7: Operating principle of a typical TPC. Events occurring within the detection volume (1) are tagged via scintillation light and charge collection. In this diagram, the scintillation light is collected by a PMT (4) but in principle the type of light measuring apparatus must be chosen carefully according to the experimental requirements (e.g. wavelength, radiopurity, quantum efficiency, etc...). Charge is collected by a wire grid, specially designed panels, etc.. (3). The constant drift field is maintained with a cathode (2'), an anode (2''), and several field-shaping rings. Figure adapted from [43].

A typical TPC consists of four components: a detection volume, a drift field, a charge readout plane, a light detection system.

The detection volume is usually a liquid or gas scintillator which also has good ionization properties (a fairly low W-value with low thermal contribution to this value, for example). In order to detect ionizing particles passing through this volume, a drift field is created that drifts electrons and ions towards the charge readout/collection plane such as a copper wire mesh.

Because the drift field is constant, the location on the readout plane at which the electrons are collected provides the location of the event origin in the x,y -plane.

The drift field itself is created with a combination of a parallel cathode & anode fixed to some constant potential difference, and a series of so-called field-shaping rings. Ideally, one would have the cathode and anode connected by a continuous conductive sheet with a linear change in potential with the cathode end equal to the cathode potential and the anode end equal to the anode potential. This would ensure a constant electric field in the TPC volume. This can't be done, so the ideally continuous change in potential from cathode to anode is instead completed in stepwise fashion with the field shaping rings. The deviation from a constant field only becomes noticeable on a micro scale near the rings themselves.

The light detection must be tailored to the design criteria of a particular experiment as will be discussed in the next section. In general, detection of photons produced by particles passing through the scintillator volume can be used for obtaining two pieces of data. The first is that the combined energy of detected electrons and photons from a given event provides the total energy deposited by that event. This energy spectrum is characteristic of the event type and can be used for discrimination. The second use of the light signal is for 3D position locating of an event. This is done by measuring the difference in the arrival time of the first photon produced by an event (detected essentially instantly) and the arrival time of the electrons at the charge readout plane. Again, the constant drift field is useful here as it produces a constant drift velocity of charged particles in the fluid volume. It is therefore a simple calculation to determine the distance that the detected electrons drifted given the recorded time difference in light and charge signals.

One of the advantages of a TPC over liquid scintillators is the additional charge measurement. Indeed, there exists an anti-correlation between the light and charge signals produced by a given event and this anti-correlation can be used to improve the energy resolution. The EXO collaboration uses this anti-correlation to improve its energy resolution at EXO-200 from $\sim 3.5\%$ to $\sim 1.5\%$ of 1σ in the ROI. An example of the dramatic energy resolution improvement when moving from only collecting light signals to using a linear combination of charge and light can be

seen in *Figure 3.14* of the next section of this chapter. Pairing charge readout with prompt light collection from events also enables full 3D event reconstruction and clustering of events. This creates at least two more variables for use in analysis: 1) The ability to distinguish event types by cluster multiplicity, where multi-site events tend to be gamma-producing scatterers and where single-site events tend to be beta-producing. 2) External backgrounds tend to deposit most of their energy near the outer portion of the detection volume while events originating within the volume (e.g. $0\nu\beta\beta$) are uniformly distributed throughout. This allows for fiducialization of the detection volume in the analysis and a straightforward rejection of events outside of that region.

As charge collection is evidently an important capability of a TPC, a fundamental design criterion for TPCs is the purity of the drift volume. Since a TPC relies so heavily on charge readout, the drift volume must be free of electronegative impurities that can cause recombination of drifting electrons before they reach the readout plane. This makes noble elements the ideal candidates for a detection medium in a TPC. Coincidentally, ^{136}Xe is a noble element and a candidate for $0\nu\beta\beta$, meaning a TPC using ^{136}Xe can use it as both its source and its detection medium, greatly simplifying the design criteria.

The EXO collaboration employs ^{136}Xe TPCs in its search for $0\nu\beta\beta$. One of the experiments, namely EXO-200, is a smaller scale prototype for a future tonne-scale experiment known as nEXO, and operates as a liquid xenon (LXe) TPC using 200 kg of LXe enriched to 80.6 % in ^{136}Xe . The current half-life sensitivity of EXO-200 is 3.7×10^{25} yr which corresponds to a lower limit for the $0\nu\beta\beta$ half-life in ^{136}Xe of 1.8×10^{25} yr range for effective Majorana mass of $(0.147 \text{ eV} \leq \langle m_{\beta\beta} \rangle \leq 0.398 \text{ eV})$ [28]. As the R&D portion of this thesis is devoted to EXO-200 and nEXO, we discuss these experiments in more detail in the next sections of this chapter.

3.2 The Enriched Xenon Observatory

The *Enriched Xenon Observatory* (EXO) is a two-phase experiment made up of two sub-experiments, namely EXO-200 and nEXO. Both experiments are designed to operate TPCs filled with LXe enriched in the isotope ^{136}Xe . EXO-200 is currently taking data and will continue to do so until December 2018, at the Waste Isolation Pilot Plant (WIPP) in Carlsbad, New Mexico. It has most recently published a lower limit for the $0\nu\beta\beta$ half-life of ^{136}Xe , summarized in *Table 3.4* [28].

Figure 3.8: Location of EXO-200 experiment at WIPP.

3.2.1 – Xenon as a $0\nu\beta\beta$ Source

The first choice to be made in designing an experiment to search for $0\nu\beta\beta$ is which isotope to use as the source medium. In *Chapter 2* the mass parabola containing ^{136}Xe was shown and from this parabola it was clear that ^{136}Xe can undergo double beta decay. However, there are several other isobars for which double beta should be possible, some of which are listed in *Table 3.2*. It is natural, then, to question the choice of ^{136}Xe as a source. In this section, the favourable properties of LXe in a search for $0\nu\beta\beta$ are considered. In this section it is explained why EXO has opted to use LXe in particular as its source & detection medium. Once this foundation has been laid, I will provide further details on the current stage of EXO.

Isotope	Abundance [31] (%)	Q-value [30] (keV)	$G^{2\nu}$ [30] (10^{-21} yr^{-1})	$G^{0\nu}$ [30] (10^{-15} yr^{-1})
$^{48}\text{Ca} \rightarrow ^{48}\text{Ti}$	0.187	4272.26	40330	64.34
$^{76}\text{Ge} \rightarrow ^{76}\text{Se}$	7.83	2039.061	124.9	6.128
$^{82}\text{Se} \rightarrow ^{82}\text{Kr}$	8.73	2995.12	4139	26.35
$^{96}\text{Zr} \rightarrow ^{96}\text{Mo}$	2.80	3350.37	17676	53.37
$^{100}\text{Mo} \rightarrow ^{100}\text{Ru}$	9.67	3034.4	8579	41.28
$^{110}\text{Pd} \rightarrow ^{110}\text{Cd}$	11.72	2017.85	357.1	12.49
$^{116}\text{Cd} \rightarrow ^{116}\text{Sn}$	7.49	2813.5	7168	43.31
$^{124}\text{Sn} \rightarrow ^{124}\text{Te}$	5.79	2286.97	1434	23.44
$^{128}\text{Te} \rightarrow ^{128}\text{Xe}$	31.74	865.87	0.6971	1.524
$^{130}\text{Te} \rightarrow ^{130}\text{Xe}$	34.08	2526.97	3965	36.88
$^{136}\text{Xe} \rightarrow ^{136}\text{Ba}$	8.858	2457.83	3716	37.81
$^{148}\text{Nd} \rightarrow ^{148}\text{Sm}$	5.7	1928.75	842.3	26.19
$^{150}\text{Nd} \rightarrow ^{150}\text{Sm}$	5.6	3371.38	94470	163.5
$^{154}\text{Sm} \rightarrow ^{154}\text{Gd}$	22.75	1215.03	24.87	7.819
$^{160}\text{Gd} \rightarrow ^{160}\text{Dy}$	21.86	1729.69	502.6	24.79
$^{198}\text{Pt} \rightarrow ^{198}\text{Hg}$	7.163	1047.17	39.83	19.59
$^{232}\text{Th} \rightarrow ^{232}\text{U}$	100	842.15	29.33	36.12
$^{238}\text{U} \rightarrow ^{238}\text{Pu}$	99.2742	1144.98	37.78	87.16

Table 3.2: Candidate isotopes for $0\nu\beta\beta$ and their natural isotopic abundance, Q-value, and phase-space factors for both $\beta\beta$ modes. See (2.2.5).

In choosing which of the $0\nu\beta\beta$ candidates to use in a search of $0\nu\beta\beta$ there are several properties, not only physical, that should be considered. A fairly exhaustive list is given by Maneschg in his summary of a recent status of $0\nu\beta\beta$ experiments [32]: a high $Q_{\beta\beta}$ value (far from most natural background radiation levels in underground experiments), a convenient $G^{0\nu}$ factor (to reduce the half-life), a high natural isotopic fraction (to reduce the need for enrichment and costs), simple enrichment possibilities (in case of low natural isotopic abundance), low cost (i.e. large world supply), simplicity of detection in the medium, background suppression possibilities, good energy resolution, fast handling/operation of detectors.

Isotope	$M^{0\nu}$ (dimensionless)	
	IBM-2 [33]	QRPA [34]
$^{48}\text{Ca} \rightarrow ^{48}\text{Ti}$	2.82	0.871
$^{76}\text{Ge} \rightarrow ^{76}\text{Se}$	7.54	8.31
$^{82}\text{Se} \rightarrow ^{82}\text{Kr}$	6.01	7.47
$^{96}\text{Zr} \rightarrow ^{96}\text{Mo}$	4.56	4.38
$^{100}\text{Mo} \rightarrow ^{100}\text{Ru}$	6.80	8.70
$^{110}\text{Pd} \rightarrow ^{110}\text{Cd}$	6.52	9.28
$^{116}\text{Cd} \rightarrow ^{116}\text{Sn}$	4.99	6.51
$^{124}\text{Sn} \rightarrow ^{124}\text{Te}$	5.14	4.12
$^{128}\text{Te} \rightarrow ^{128}\text{Xe}$	6.60	7.34
$^{130}\text{Te} \rightarrow ^{130}\text{Xe}$	5.96	6.26
$^{136}\text{Xe} \rightarrow ^{136}\text{Ba}$	4.91	3.51

Table 3.3: Nuclear matrix elements of various $0\nu\beta\beta$ candidates according to the IBM-2 model and the QRPA model.

This list includes ten items worth consideration. Of these items, LXe is favourable for several. For example, LXe has a $Q_{\beta\beta}$ value of approximately 2.46 MeV (see *Table 3.2*) which is higher than the energy level of most naturally occurring background radiation in underground experiments (where most rare event searches operate for the sake of reducing background contamination). With respect to other candidates, we see that its natural isotopic abundance is poor. However, due to its characteristic inertness, it is easily enriched to >80% abundance through centrifugation.

With respect to other candidates, we see that the phase space factor is particularly favourable for ^{136}Xe . Unfortunately its nuclear matrix element is one of the most unfavourable, regardless of model choice (IBM-2 or QRPA), and the half-life has a quadratic dependence on this term (see *Table 3.3*).

Perhaps the primary motivation for using ^{136}Xe is that xenon makes an excellent detection medium in a TPC and that experiments have been using xenon TPCs for decades. Thus, there is already an obvious detector/source pairing that comes with the choice of ^{136}Xe as the $0\nu\beta\beta$ source. Indeed, noble elements are, in general, excellent candidates for detection media in drift

chambers. For example, a TPC should not contain electronegative media or we risk losing the drift electrons to recombination. Significant electron loss means we cannot locate events and indeed cannot observe events of interest that require both scintillation and ionization signals. Furthermore, through decades of experiments, ionizing particle interactions in xenon are already well understood.

As any $0\nu\beta\beta$ counting experiment is a low-background experiment, the high density of xenon is particularly favourable. Its high density ensures that photons are significantly attenuated over a short distance. This both reduces contamination from backgrounds of external origin penetrating the detector, and increases the likelihood of containment of events with origins within the detector. The former point directly reduces background in the ROI, improving sensitivity, while the latter point allows for better event reconstruction, which eventually leads to better background discrimination and improved sensitivity.

3.2.2 EXO-200

The EXO-200 detector is made-up of two identical LXe TPCs connected at their ends by a shared cathode (see *Figure 3.9*). The LXe is enriched to 80.6% in the isotope ^{136}Xe and operates at a temperature of 167 K and a pressure of 1.47 bar. The LXe itself must be kept highly pure to maximize electron lifetime in the drift region. This is achieved through continuous recirculation of LXe through heated SAES getters [35] using a novel, high-purity, magnetically driven piston pump [36]. Because of this continuous recirculation, only 175 kg of the 200 kg of xenon is ever in liquid form in the TPC while the rest is circulated in gaseous form. Of the 175 kg of LXe, 110 kg are in the active detection volume.

Acting as the anodes, each TPC has its own charge readout plane situated on opposite sides of the total LXe volume. Each charge readout plane consists of a U-wire plane and a V-wire plane (see *Figure 3.9* & *Figure 3.10*) oriented at 60° to each other.

Figure 3.9: Schematics of EXO-200 TPC [37]. *Right:* APDs are read out in gangs. The top inset shows a zoomed view of the relative configuration of the wire planes and the APD plane. The bottom inset illustrates the angle between the wires planes.

The U-wires are used to collect charge and act as the anode of the TPC while the V-wires, situated in front of the U-wires with respect to the cathode, are used to collect induction signals from passing charges. Drifting electrons are made to pass through the induction planes without being collected by rendering the plane electrically transparent, achieved by setting the potential of the induction wires to roughly twice the potential of the anode/collection wires [37]. The signals from the U,V-wires combined with the 60° offset of the planes allows for X,Y-coordinate reconstruction of event locations in the LXe.

Figure 3.10: *Left:* Principle of simultaneous light and charge detection in the EXO-200 TPCs. The 2D cross-sectional view of the wire planes demonstrates results of SIMION simulations for determining the potential at which electric transparency of the v-wires is achieved. Signals from the wires are read out in triplets to reduce radioactivity caused by excessive use of material. *Right:* View of induction and collection wires crossed at 60° in the EXO-200 TPC before installation to the cryostat. Images adapted from [37].

The wire planes are situated between the cathode the light readout planes. The wire planes were found to have an optical transparency of 95.8 % [37]. These APD planes each consist of

234 bare APD diodes (see *Figure 3.10*). The choice of bare APDs was made for their low radioactivity and high quantum efficiency at the scintillation wavelength of LXe. To further improve light collection, the field shaping rings were lined with VUV Teflon (see *Figure 3.11*) which diffuses scintillation light towards the APD planes. The time difference between light collection at the APD planes and charge collection at the wires planes is used to determine the location of an event along the longitudinal axis of the detector. Combined with the 2D position sensitivity of the wire planes, this provides for 3D event reconstruction.

The last of the four TPC components discussed in section 3.2, the field cage (includes field-shaping rings), is made from highly radio-pure copper to minimize backgrounds in the detection volume. The TPCs themselves are housed in a copper vessel (see *Figure 3.11*) which is itself suspended in a large volume of HFE-7000, a cryogenic heat-transfer fluid.

Figure 3.11: EXO-200 TPC (left) prior to installation in its copper vessel (right). Images adapted from [37].

The HFE serves the purposes of cooling the copper vessel as well as for suppressing external, naturally occurring radioactive backgrounds that are inherent to underground experiments. The HFE is contained within a large, double-walled copper cryostat, made from the same radio-pure copper as the field shaping rings, which is itself covered by layers of lead shielding on four sides. The cryostat walls and the lead shield serve the purpose of reducing penetration of external radioactivity into the TPC. The inner walls of the cryostat are equipped with three heat-exchangers, for regulating the temperature of the LXe. The HFE's excellent heat transfer properties ensure that the heat is transferred uniformly between the cryostat walls and the TPC

vessel. Muon backgrounds are suppressed by four groups of plastic scintillators situated outside of the cleanroom (see *Figure 3.12*). The veto efficiency has been found to be 95% [38].

Figure 3.12: Artist's interpretation of EXO-200 detector. Image adapted from [37].

The complete EXO-200 detector just described (see *Figure 3.12*) is located 655m underground, corresponding to a 1585 m.w.e depth [38], at WIPP in order to reduce background contamination. Furthermore, all components of the detector were screened for use in a thorough radioassay campaign. This campaign is described elsewhere [39].

Interactions in the LXe deposit energy through light (scintillation) and charge (ionization) signals. Events producing α radiation are easily tagged with these two measurements due to their characteristic ratio for α particles. Unfortunately, β -producing (e.g. $0\nu\beta\beta$, $2\nu\beta\beta$) and γ -producing events (^{232}Th and ^{238}U chains) cannot be distinguished by this method. In this case, EXO-200 uses single-site (SS) and multi-site (MS) classification of events to distinguish between β -producing events and γ -producing events. SS events are defined by signals that produce

collections of charge on either 1 or 2 U-wires while MS events produce collection at more than 2 U-wires. Events that are β -producing are mostly SS events as electrons continuously deposit energy by Coulombic interactions, for example, and rarely produce additional charge clusters. By contrast γ -producing events are dominated by Compton scattering, photoelectric absorption, and pair production. The Compton interaction in particular is capable of producing additional charge clusters and so γ -producing events tend to dominate the fraction of MS events.

The total deposited energy (from all scintillation and ionization channels) is also used as an event discriminator in EXO-200. The detector's energy response is calibrated using ^{60}Co , ^{137}Cs and ^{228}Th .

Figure 3.13: *Left:* Computer model of copper guide tube wrapping around the copper vessel. *Right:* Various calibration source locations denoted by S2, S5, S8, S11.

The positional dependence of the detector's energy response is handled by placing the various calibration sources at several source positions outside of the copper vessel (see *Figure 3.13*). The positional dependence is caused by a combination of the solid angle of events with respect to the APDs and the various gains of each APD. Full calibration studies using each source at each position are conducted approximately every three months. Data is collected for each source/position pair for approximately three hours. Shorter calibration studies are conducted every two to three days using ^{228}Th at position S5. This particular short-term calibration is used to improve the energy resolution of the experiment by using the so-called *rotated energy*.

It has been found that scintillation and ionization signals can be combined and rotated to significantly improve energy resolution thereby increasing the discrimination power of energy measurements [40]. More details on the study of this phenomenon and discussions as to its origin can be found elsewhere [41]. The rotation is optimized to reduce the variance of the background 2615 keV gamma line, a primary background near the Q-value. The dramatic improvement is depicted in the figure below.

Figure 3.14: *Left:* Graphical depiction of event-by-event anticorrelation between ionization and scintillation, and rotated energy axes in EXO-200. The energy is rotated to minimize the variance at the 2615keV peak of the ^{228}Th calibration source. *Right:* A comparison of energy spectra for scintillation only, ionization only, and rotated channel. The plots shown here only apply to SS events. Adapted from [40].

EXO-200 chooses the rotation angle to minimize the variance of the known 2615 keV peak produced by the ^{228}Th calibration source placed at position S5.

The energy spectrum of events is rotated by taking advantage of the event-by-event anticorrelation between scintillation and ionization signals. The rotated energy, E_R , is defined by,

$$E_R = E_S \sin\theta + E_I \cos\theta \quad 3.2.0$$

where E_S is the scintillation signal of an event, E_I is the ionization signal of an event, and θ is the rotation angle. In general, θ is different between SS and MS events. While the ^{228}Th calibration is run every two to three days, the rotation angle is recalculated weekly using the calibration data from the week.

Thus far, the primary success of EXO-200, outside of being a successful proof of concept for the ton-scale nEXO, has been the discovery of $2\nu\beta\beta$ in ^{136}Xe with a measurement for the half-life of [42],

$$T_{\frac{1}{2}}^{2\nu} = (2.165 \pm 0.016(\text{stat}) \pm 0.059(\text{sys})) \times 10^{21} \text{ yr} \quad 3.2.1$$

making it the slowest natural process directly observed. This process is the main irreducible background for EXO and ultimately limits the running time for which improvements in sensitivity can be achieved for any given detector design. After its first two years of data collection EXO-200 did not claim a discovery of $0\nu\beta\beta$ in ^{136}Xe but, in a paper published in *Nature* in 2014, placed a limits on the process that are summarized in *Table 3.4* below .

Property	Value
$0\nu\beta\beta$ Half-Life Sensitivity	$1.9 \times 10^{25} \text{ yr}$
$0\nu\beta\beta$ Half-life lower limit	$1.1 \times 10^{25} \text{ yr}$
Effective Majorana Mass Upper Limit	190 – 450 meV

Table 3.4: EXO-200 results after 122 kg·yr exposure. Results are taken from an article published by the EXO collaboration in *Nature*, in 2014 [42].

Following this publication, the experiment was shut down for a year due to two safety incidents at WIPP [45][46]. During this period, new analysis techniques and hardware upgrades were developed and finally used during another year of data-taking once the experiment began running again. Employing new analysis techniques (namely *boosted decision trees*) on the original two years of data and the first year of data collected by the upgraded detector, EXO-200 has improved its measurements to the values summarized in *Table 3.5*.

Property	Value
$0\nu\beta\beta$ Half-Life Sensitivity	3.7×10^{25} yr
$0\nu\beta\beta$ Half-life lower limit	1.8×10^{25} yr
Effective Majorana Mass Upper Limit	147 – 398 meV

Table 3.5: Updated EXO-200 results from February 2018 [29]. Results combine previous exposure of 122 kg·yr with an additional 55.6 kg·yr of exposure in the upgraded EXO-200 detector. New analysis techniques involving boosted decision trees were implemented on the entire exposure to obtain the summarized results.

After completing data collection in December 2018, the summarized limits will be updated once more and EXO-200 is expected to reach a mass sensitivity illustrated in the figure below. Evidently the updated EXO-200 result is not expected to sweep out a large portion of the mass hierarchies that has not already been precluded by cosmological studies. In this sense, a discovery of $0\nu\beta\beta$ is perhaps unexpected. To sweep a significant portion of the mass hierarchies, the EXO collaboration will have to look to the ton-scale nEXO which is briefly discussed next.

Figure 3.15: Expected sensitivity of EXO-200 after finishing data collection in December 2018. The horizontal axis is the mass of the lightest neutrino, while the vertical axis is the effective Majorana mass. Figure adapted from [43].

3.2.3 nEXO

The next phase of EXO, namely nEXO, will be a multi-tonne TPC operating with LXe as the source medium just as in EXO-200. In general, nEXO is designed to be as similar to EXO-200 as possible such that the physics of the detector will already be well understood. The much larger source volume of nEXO will increase the potential to observe $0\nu\beta\beta$ over a fixed period of time according to (2.2.4). The proposed operating location is at the SNO site in Sudbury, Canada.

Figure 3.16: Depth of SNO site in comparison to EXO-200 site at WIPP and other underground experiment locations. Operating nEXO at the SNO site would reduce the muon flux by more than two orders of magnitude relative to the flux at the EXO-200 site at WIPP. This reduces scattering backgrounds as well as the ^{137}Xe contribution (muon capture followed by decay) to the fits shown in the previous section. Figure adapted from [47].

The depth of the SNO site reduces the expected muon flux by two orders of magnitude in comparison to EXO-200. This significantly reduces backgrounds that occur, for example, in the form of cosmogenically activated ^{60}Co and ^{137}Xe in EXO-200. Furthermore, the larger volume of LXe will benefit from self-shielding as most backgrounds inherent to underground experiments (e.g. Rn, U, Th) are quickly attenuated in the dense LXe and deposit most of their energy at the periphery of the LXe volume. This allows for rejection of these backgrounds by fiducialization of the source LXe. By the same effect, a larger source volume ensures more complete containment of events. That is, events are more likely to deposit their full energy in the

detector without escaping. This enables an improvement in event reconstruction efficiency, ultimately leading to better event discrimination.

Of course, any background due to LXe itself, such as $2\nu\beta\beta$, cannot be suppressed by an increase in LXe volume. The limitations that this places on nEXO are discussed in the next section of this chapter.

3.3 Background Limitations and Background Reduction

There is a limit for the half-life after which experiments with non-zero backgrounds can essentially no longer extend a measurement for the half-life. This is illustrated by considering the following parameterization of signal-to-background, as in [32],

$$R = 1 + \frac{B(2\nu\beta\beta)}{S(0\nu\beta\beta)} \quad 3.3.1$$

where a ratio of $R = 1$ suggests a perfect suppression of background in the region of interest (ROI - some region of parameter space close to the peak of the signal distribution), $B(2\nu\beta\beta)$ denotes the fraction of events in the ROI occurring from $2\nu\beta\beta$ background, and $S(0\nu\beta\beta)$ denotes the fraction of events in the ROI that are $0\nu\beta\beta$. There is a relationship between this number and the half-life of $0\nu\beta\beta$. Put simply, for a given cut for the ROI, longer half-lives will lead to larger values of R for experiments with non-zero backgrounds. This is fairly intuitive but the rate of growth in R requires some calculating. Such calculations are not done here but we will refer to [32] in which example calculations were made for several major $0\nu\beta\beta$ counting experiments. The results obtained from the calculations in are shown in the following figure.

Figure 3.17: *Left:* Example of a cut on energy ensuring 88% signal survival in the ROI. The height of the green line to the right of the cut indicates the unsuppressible background of $2\nu\beta\beta$ while the red lines indicate the signal distribution. *Right:* Signal-to-background ratio (3.3.1) as a function of the $0\nu\beta\beta$ half-life. Plots taken from [32].

In these calculations, the ROI is set by separating $0\nu\beta\beta$ from $2\nu\beta\beta$, with a cut that excludes 12 % of $0\nu\beta\beta$ events. It is shown in *Figure 3.17* at which values for the half-life that R begins to become a problem for the non-zero background experiments. The zero background experiments, of course, have no such limitations. They come with their own difficulties. It is, for example, much easier and cost-effective to scale a liquid TPC than it is to scale the listed solid state detectors.

At this point it should be noted that $2\nu\beta\beta$ is not the only event type that contributes to backgrounds around the ROI. However, it tends to be the dominant background and it is one that cannot be suppressed with shielding as the event type occurs within the bulk of the detector uniformly and is inherent to any $\beta\beta$ -emitter. In this sense it is the limiting background and so (3.3.1) is useful in this regard. Indeed, for tonne scale experiments such as nEXO, most other backgrounds will be heavily suppressed due to the self-shielding inherent to LXe. A simple cut on the fiducial volume will entirely eliminate these events from the ROI.

A second point worth considering is that the calculations from the figure above assume an 88% survival of $0\nu\beta\beta$ events using the resolution of the considered experiments as of April 28th, 2017. As analysis techniques become more complex, resolutions will improve and these results will change. For example, with improved discrimination capabilities, we may be able to expand the ROI from an 88% survival region to a 95% survival region without reducing the signal to

background ratio (3.3.1). Regardless, the fundamental problem for non-zero background experiments that is illustrated in *Figure 3.17* persists.

As we have already seen, the signal-to-background ratio of an experiment will limit its sensitivity and its discovery potential. Thus, for an experiment with some fixed resolution, there will be a value of the half-life for which simply extending the length of data taking will be useless. At this half-life the background becomes so large that no appreciable gain in sensitivity can be achieved by collecting more data. Because such a situation is inevitable for half-lives longer than $\sim 10^{29}$ yrs for nEXO, there is a large R&D effort going towards reducing the background in nEXO. The rest of this thesis is devoted to this effort and contributes through three approaches: safely ramping to higher voltages of the TPC's cathode to improve parameters' resolutions, using complex machine learning analysis to improve discrimination in our parameter space, and Barium-tagging to eliminate all backgrounds not related to double beta decay. Each approach will be considered in its own dedicated chapter.

Chapter 4

High Voltage Breakdown at EXO-200

In this chapter, we begin by motivating the use of HV in TPCs in terms of improved resolution of various measured parameters. From there, we summarize previous HV test results obtained at SLAC and Bern and the conclusions that were made in ruling out various HV-breakdown hypotheses. These conclusions motivated the original HV work done for this thesis. This new work involves a pattern recognition approach to predicting the probability of breakdown, across a gap filled with a liquid dielectric, in a system with a complex geometry.

Work that was done towards commissioning a lab at Carleton which will be used for further HV tests is also discussed. The purpose of this lab will be to continue from the results of HV tests conducted in Bern. Due to loss of a virtual ground connection during operation, several tests conducted in Bern were inconclusive. It is of primary interest to repeat these tests in the Carleton lab. The steps that have been completed in eliminating the loss of virtual ground and re-commissioning for future HV tests at Carleton are summarized.

4.1 HV-Breakdown at EXO-200

In its original design, the EXO-200 TPC was able to provide a potential difference of up to 75 kV across the drift volume of TPC. The TPC currently operates at a 12 kV potential difference. During commissioning of the detector, an oscilloscope coupled to the HV line detected spikes, or “glitches”, in the voltage beginning at an applied potential difference of 10 kV. Simultaneously, the LAAPDs located behind the readout planes seemed to record light signals within the TPC. It is generally believed that glitches signal imminent breakdown should voltage be ramped further or should the voltage be left at the glitch voltage for a long period of time. Thus, during its first phase of data-collection, EXO-200 operated at a potential difference of -8 kV. This phase ended when accidents at WIPP forced the experiment to be temporarily suspended [45, 46]. Before re-commissioning for the second phase of data collection, the HV voltage group of EXO conducted voltage ramping tests on the detector and found that the

glitches could be eliminated by so-called voltage training. This process involves slowly increasing the voltage, lowering it slightly at the onset of glitches, and then increasing again once the glitches dissipate. During these tests, the HV group found that the TPC could be operated safely up to at least a potential difference of 12 kV. Thus, EXO-200 currently operates at a 12 kV potential difference. However, because operating at higher voltages can improve data quality, studying the glitch phenomenon and its relationship to discharge/breakdown is a subject of great interest within the EXO collaboration. Indeed, the successor to EXO-200, nEXO, must operate at a much larger potential difference than 12 kV in order to produce the same drift field as EXO-200. It is important to understand whether this higher voltage will result in breakdown or if the breakdown voltage will scale with the increased size of nEXO relative to EXO-200.

4.2 Motivation for High Voltage in TPCs

That EXO-200 has been unable to ramp up its applied cathode voltage to its designed limit has hindered the experiment's sensitivity. The ability to safely (i.e. without inducing discharge that may damage electronic equipment) apply higher voltages to the cathode improves the resolution of several measured parameters. This ultimately leads to improved sensitivity and discovery potential due to the increased discriminating power between event types.

Energy resolution can be significantly improved by operating at higher voltages. This is the most useful variable for discriminating between $0\nu\beta\beta$ and $2\nu\beta\beta$ as techniques like barium-tagging cannot discriminate between the two event types. The expected improvement in resolution with voltage is shown graphically in *Figure 4.1*. Such an improvement is important as the various event types in the fit model of EXO-200 have characteristic energy spectra. Thus, by improving resolution, we can more easily distinguish between events of a different type based on energy cuts. This is also discussed in *Chapter 3* and in *Chapter 5*. The resulting reduction in overlap of the two energy distributions that produces improved resolution was shown previously in *Figure 2.8*.

Figure 4.1: Expected improvements in energy resolution and SS fraction with increased cathode voltage at nEXO as predicted by simulations of the EXO-200 detector using the *Geant 4* software package. The horizontal axis denotes the electric field in the drift volume of the TPC. The black/left vertical axis denotes energy resolution in terms of the ratio of the energy resolution at the Q-value to the Q-value itself. The black curves represent how this ratio is expected to decrease (i.e. improve) as voltage increases. Each black curve corresponds to a result with different light collection efficiency. The red/right vertical axis denotes the fraction of $0\nu\beta\beta$ events that are reconstructed as single-site events. The red curves represent how this fraction increases with increased voltage. Each red curve corresponds to results from a simulation in which a different diffusion coefficient of electrons was used. Incidentally, the diffusion coefficient itself is expected to decrease with an increased operating voltage (see explanation below (4.2.1)). Figure created by Liagjian Wen and Guofu Cao in MC studies using *GEANT 4*.

The increased drift velocity of electrons also tends to reduce the capture-cross section on oxygen impurities for a fixed LXe purity [44][51]. It is expected that the capture cross-section would be reduced by a factor of 0.6 for a given LXe purity when changing the drift field from 100 V/cm to 1000 V/cm (EXO-200 currently operates at ~ 600 V/cm) [44]. Additionally, larger operating voltages reduce the relative electronics noise in our measurements. Improvements such as this improve the single-site fraction of reconstructed $0\nu\beta\beta$ events – depicted by the red curve in *Figure 4.1*. This fraction is equal to the number of $0\nu\beta\beta$ events that are truly SS, divided by the total number of $0\nu\beta\beta$ events (i.e. the probability that a $0\nu\beta\beta$ event is truly SS). This is explained as follows: $0\nu\beta\beta$ events tend to be SS, however, no detector can provide a perfect reconstruction of a physics event. Thus, in our TPC, some $0\nu\beta\beta$ events that are truly SS events end up being misidentified as MS events. For a TPC, these imperfections in the measurement are reduced at higher operating voltage. This causes the measured distribution for site

multiplicity of $0\nu\beta\beta$ events to approach its true distribution. Most $0\nu\beta\beta$ events are SS in the true distribution and this is why the red curve increases with larger potential differences in *Figure 4.1*. This is an important measurement because, as was explained in *Chapter 3*, γ -producing background events tend to interact via Compton scattering and produce many charge clusters in the TPC. Thus, if we can successfully reconstruct more $0\nu\beta\beta$ events as SS events, it becomes easier to discriminate $0\nu\beta\beta$ from these γ -producing backgrounds. Incidentally, as is discussed next, a higher operating voltage makes it easier to distinguish charge clusters in events that are truly MS events – meaning some γ -producing events that would be identified as SS at low voltages can be correctly identified as MS events at high voltages. This effect is studied via MC studies and compared to calibration data for a ^{228}Th source in an EXO-200 publication [29]. This publication also discusses systematic error in the shape agreement between the MC and calibration data distributions. Of course, a similar comparison between calibration and MC could not be made for $0\nu\beta\beta$ events and we must rely on the MC study results of *Figure 4.1*. In summary, we improve the discriminating power between $0\nu\beta\beta$ and γ -producing events by both correctly identifying truly SS $0\nu\beta\beta$ events as SS, and identifying truly MS γ -producing events as MS.

Higher voltage also improves position resolution within the detection volume by resisting diffusion. It is explained in an EXO-200 publication that the radial variance of particles diffusing in a liquid can be described by [44],

$$\langle \Delta R(t)^2 \rangle = \langle \Delta x(t)^2 \rangle + \langle \Delta y(t)^2 \rangle = 4D_T t \quad 4.2.1$$

where ΔR is the radial distance of the particle with respect to its initial position (i.e. the original location of electron-ion pair creation) in the TPC at time, $t = 0$, Δx & Δy are the Cartesian coordinates of the particle with respect to its initial position in the TPC, and D_T is the transverse/radial diffusion coefficient in LXe. Evidently, then, the variance of the radial position is reduced if the particle is detected in a short amount of time. Since the drift velocity tends to increase in a TPC with an increased applied voltage on the cathode, the radial variance of diffusing particles can be reduced by increasing the cathode voltage. Moreover, it has been observed that the diffusion coefficient itself decreases with an increased drift field. This has not

been observed in EXO-200 tests but has been observed at higher operating voltages in other experiments [48][49]. Incidentally, if we again consider *Figure 4.1*, it is clear that an increased diffusion coefficient is expected to improve our ability to measure $0\nu\beta\beta$ events as SS events when they are truly SS events.

Figure 4.2: Drift velocity for varied drift fields at EXO-200, measured with ^{228}Th calibration data and α sources from low-background runs, compared to drift velocity in Xe measured by other collaborations. LXe temperature: Gushchin-165K, Sorensen-177K, Akerib-175K, Miller-163K, Aprile-182K, EXO-200-167K. Adapted from [44].

A reduction in radial variance allows simpler discrimination of, say, multiple Compton scattering sites – a parameter used in EXO-200’s data analysis. It also improves the resolution of a parameter called *standoff distance* that is used in data analysis at EXO-200 (see *Chapter 5* for details). Therefore, the reduction in radial variance brought about by an increased operating voltage improves the resolution of at least two spatial parameters that can be used for analysis at EXO.

Other effects that are qualitatively understood are the reduction of positive ion accumulation with stronger drift fields. This is an issue for intense calibration sources (e.g. ^{228}Th used for calibration in EXO-200) at lower drift fields. Stronger fields also tend to reduce the effect of fringe fields near field-shaping rings and other detector components. These fringe fields can cause charge build-up on dielectrics (e.g. acrylic struts) and severely distort the drift field. These

two effects are difficult to model in Monte Carlo simulations making it difficult to train complex learning algorithms for event discrimination.

In summary, using a higher operating voltage tends to improve the resolution of parameters used in analysis. We also expect reduction of nuisance effects like fringe fields and charge accumulation which hinder data quality systematically.

4.3 Breakdown Origin Hypotheses and miniEXO Tests

4.3.1 miniEXO

In order to understand the glitch phenomenon at EXO-200, a HV R&D group was set up within the EXO-200 collaboration. Breakdown cannot be studied in EXO-200 itself, however, as any discharge poses the risk of severely damaging expensive detector components. Furthermore, repairing the detector could cause delays in data collection of up to several months due to, for example, the intense radiopurity screening that would need to be applied. Instead, the HV group decided to construct a smaller model of the EXO-200 TPCs in order to study the glitch & breakdown phenomena. This smaller TPC is referred to as *miniEXO* (see *Figure 4.3*).

Figure 4.3: miniEXO (1) and one of the EXO-200 TPCs (2). Adapted from [43].

The miniEXO detector is approximately 1/2-scale in comparison to EXO-200 in the radial axis (19cm diameter in comparison to EXO-200's 40 cm diameter). Because good radio-purity is inessential for HV tests, the field cage of miniEXO is made from stainless steel instead of copper like at EXO-200. It is possible that material choice could affect the value of the breakdown voltage. The acrylic struts used in miniEXO are taken from the unused stock of EXO-200 and have undergone the same out-gassing treatments. To simulate the EXO-200 cryostat, miniEXO was surrounded by a cylindrical stainless steel shroud (a.k.a *The Shroud of Bern-in*) a radial distance of 11 cm from the field cage. This is equal to the separation between the field-shaping rings and the cryostat in EXO-200. As sharp points can lead to glitches or breakdown, the shroud was polished and its ends rounded. The miniEXO TPC can be operated with or without the *Shroud of Bern-in*. The shroud can be seen in *Figure 4.3-1c* while the EXO-200 cryostat surrounds the TPC in *Figure 4.3-2c*. Image *1c* of *Figure 4.3* shows the rounded end of the *Shroud of Bern-in* quite clearly. The TPC is housed in a cryostat, a CAD model of which is depicted in *Figure 4.4*. The manifold on the lid of the cryostat consists of several feedthroughs for making various connections to detector components that are housed in the inner volume (e.g. cameras, purity monitors, etc).

Figure 4.4: CAD model of TPC housed in the cryostat previously operated in Bern, now operating at Carleton. In this model a TPC other than the HV TPC is depicted. All details apart from the TPC are unchanged when the depicted TPC is swapped for the HV TPC. **1:** TPC in copper chamber. **2:** Steel base-plate with attached level meters for monitoring LXe level. **3:** Cryogenic-camera for monitoring HV tests. **4:** Reserved space. This space is planned to eventually house a Barium-tagging extraction device (see *Chapter 7*). **5:** Location of indium seal connecting copper volume in which TPC is housed and the reserved space currently housing cryogenic cameras. **6:** Cryostat manifold comprised of several feedthroughs to components within the inner volume (e.g. connection to the cryo-camera).

4.3.2 Breakdown Hypotheses

One of miniEXO's primary objectives was to test various hypotheses for the source of glitches and breakdown in a TPC like EXO-200. Within the EXO collaboration, three hypotheses were suggested as the limiting cause of discharge/breakdown. One of these was that discharge originates in the voltage divider (*divider breakdown*). A second was that discharge occurs between a field-shaping ring and the grounded cryostat by jumping the 11 cm LXe gap to the cryostat wall (*gap breakdown*). The third hypothesis was related to the acrylic struts on which the field-shaping rings were mounted. In order to reduce out-gassing of the struts, they were

stored in dry nitrogen gas for several years during the commissioning and construction of EXO-200. Such a treatment reduces the natural spark suppression inherent to absorbed oxygen & water impurities. It was proposed that the lack of these impurities could change the electrical properties of the acrylic sufficiently to allow discharge to occur across its surface between field shaping rings. Breakdown across the surface of plastic dielectrics surrounded by liquid noble elements has been observed in published and unpublished HV work from the EXO HV group at SLAC and by other groups [50][51]. Recall that miniEXO used leftover struts from the EXO-200 stock, making miniEXO an excellent place to study the third hypothesis in particular. Of course, the other hypotheses can be studied but the results will not be directly applicable due to material differences for example.

4.3.3 Past High Voltage tests with miniEXO and COMSOL Simulations

A combination of COMSOL simulations of the electrostatics of miniEXO and experimental data collected while miniEXO was operating in Bern suggest that gap breakdown is the limiting source of discharge for LXe TPCs. Discharges originating in miniEXO were recorded at potential differences of 38kV and larger [43]. The discharges were recorded audibly, on an oscilloscope coupled to the anode, and with two cryogenic cameras placed inside of a cryostat in which miniEXO was housed. Images captured with these cameras suggest that discharge always occurs on or near an acrylic strut at the upper field-shaping ring (see *Figure 4.5*).

Figure 4.5: Discharges in miniEXO captured by cryogenic cameras. *Top row:* The prompt flash of light during breakdown seems to occur on or near a strut in all cases. *Bottom Row:* There is a phosphorescent glow on the acrylic struts after discharges, presumably induced from LXe UV scintillation light during the spark. Figure adapted from [43].

If gap breakdown is the limiting cause of breakdown then one would expect the discharge path between the field-shaping rings and the shroud to originate where the potential difference between the two is the largest. Electrostatic simulations of miniEXO completed by Tamer Tolba (see *Figure 4.6*) show that the largest field occurs at acrylic struts. This is in agreement with the flashes and afterglow recorded by the cryo-cameras.

Further evidence of gap breakdown near acrylic pieces was provided by tests on EXO-200. In 2014, during HV ramping tests, the event displays of EXO-200 glitches show a peak signal at corner APD gangs. These corner signatures suggest sparking near acrylic struts. This suggests that EXO-200 may breakdown similarly to how miniEXO breakdown, at least in terms of location and mechanism (the voltage, of course, could be different).

Figure 4.6: Electric field strength in miniEXO as determined by electrostatic COMSOL simulations. *Top:* Fields near acrylic combs. *Bottom:* Fields away from acrylic combs. The acrylic comb clearly amplifies the field strength by a small factor of a few percent.

Figure 4.7: Data collected during HV tests with EXO-200 during a shutdown. The hexagons on the left of the image show event locations as measured by various wire planes (of which there are 4 – 2 for each TPC) and APD planes (2 – 1 for each TPC). The middle plots show the duration of signals in various channels for the corresponding hexagon on the left. Note the corner signature in the APD plane for the spark observed in this image. The corners are located near acrylic struts.

The increase in electric field strength near the acrylic struts is expected as a consequence of continuity of the electric displacement at the boundary between dielectrics (LXe & acrylic in this case). In general, if we add a dielectric to the system and if it has a greater dielectric constant than LXe, the electric field will increase in the LXe in a gap-like geometry such as the one that we have in miniEXO (and vice versa). This is shown roughly in *Appendix B*. In this particular case, acrylic has a dielectric constant of approximately 3.5 in comparison to the 1.9 of LXe.

Thus, of the three proposed hypotheses, the results from miniEXO tests conducted in Bern tend to favour gap breakdown. Results from EXO-200 itself favour this hypothesis as well. Gap breakdown near acrylic struts is of particular interest given results from the cryo-cameras and COMSOL simulations of miniEXO, and HV ramping tests with EXO-200. Of course, if the current limiting cause is eventually understood and subsequently prevented, the other two mechanisms may become the limiting causes. Nevertheless, the current result suggests that gap breakdown in the TPC should be addressed, understood, and prevented first. Indeed, it is

difficult to collect any data related to the other hypotheses at all if gap breakdown tends to precede breakdown by these other mechanisms. In the next section, an approach to predicting gap breakdown voltages in complex systems is presented. A prediction is made on the miniEXO setup and the result is compared to the laboratory data.

4.4 Pattern Recognition for Gap-Breakdown Prediction

Results from COMSOL simulations and data collected while miniEXO was operated at Bern suggest that the gap breakdown mechanism is the limiting cause of electrical breakdown in TPCs like miniEXO. As miniEXO is designed to be similar to EXO-200, we assume that this holds true in EXO-200 as well. In this section we discuss an approach to predicting the voltage at which gap breakdown will occur in miniEXO. We compare this prediction to the true value at which breakdown occurred to quantify the predictive power of the algorithm.

4.4.1 The SVM Algorithm

Breakdown in dielectrics is not a newly observed phenomenon. There is a huge body of literature that attempts to explain the process by which it occurs [52-60]. Some approaches seek a first principles or phenomenological explanation and, while some of the processes that lead to dielectric breakdown are well understood through these approaches, the degree to which each process contributes in any given system remains a difficult question to answer.

An alternative approach to first principles derivations is using pattern recognition tools to discover unstable relationships between particular variables. By unstable, we mean that for certain values, the particular set of variables describing the HV system will lead to breakdown in the dielectric. This approach sacrifices a precise quantitative understanding of the underlying physics for an approximated mapping of unstable groups of parameters in the parameter space. This technique may appear new but it is essentially already employed in HV applications. Indeed, breakdown results from one apparatus are often extrapolated on a linear scale to estimate the voltage at which similar, larger-scale or smaller-scale apparatuses will discharge. Pattern recognition approaches are essentially doing the same thing but with more sophisticated and not necessarily linear interpolation or extrapolation. Of course, in order to obtain a meaningful mapping, one needs to at least understand which physical variables are

important in the process of electrical breakdown in liquid dielectrics. In this sense, we can't completely abandon the first principles approach. For example, a $\sim 250\text{kV/cm}$ breakdown field in SLAC tests corresponded to a $\sim 85\text{kV/cm}$ breakdown field in miniEXO tests. In order to accurately map this with a pattern recognition algorithm we need to understand why we would see such a result. A result like this may be expected if volume and surface area of electrodes were to be considered (i.e. the stressed area). For example, a larger dielectric volume and electrode surface area can lead to a greater build-up of space charge near the electrode. This increases the field locally and reduces the Fowler-Nordheim barrier at the surface of the electrode, allowing for field-emission of electrons in much greater quantities. Therefore, one should include the surface area of electrodes as a parameter in a pattern recognition algorithm.

In this work, we attempt to replicate an approach to predicting breakdown voltages that was originally proposed in [61]. This proposed algorithm reduces the problem to a binary classification problem where points in a parameter space are classified as breakdown points or stable points (the statistical aspect of breakdown is addressed later in this chapter as a proposed extension to the algorithm). A point is represented by a collection of physical parameters associated with a particular apparatus whose cathode is operating at some fixed voltage (anode HV is not considered here). Some example parameters are maximum electric field strength, cathode voltage, length of discharge path, and total energy contained in the dielectric volume. The full list of parameters is given in [61]. The fundamental principle of this approach is that we can collect breakdown data for a collection of simple geometries (e.g. Rogowski electrodes, parallel plates, etc ...) and then use this data to predict the breakdown voltage on complex geometries (e.g. a TPC).

The classification algorithm that the model adopts is the Support Vector Machine (SVM). While we have mostly adopted the SVM simply because we are attempting to reproduce a method that was used by other authors in which the SVM was used, the SVM is a natural choice of algorithm for this problem. The SVM employs the structural risk minimization principle (SRMP) where the predictive success of a model is weighted against its complexity. It does this by simultaneously minimizing the empirical risk of misclassification and the width of confidence intervals. The SRMP can be used on small datasets without problems with overfitting and poor

generalization. Thus, since it isn't feasible to collect breakdown data on a large number of different systems/apparatuses, the SVM is attractive for studying the HV problem. In the future, if the amount of collected breakdown data grow in size, it may be feasible to use another kind of algorithm whose performance improves with large datasets (e.g. neural networks).

In a typical HV problem there is a multidimensional non-linear relationship between the parameters of the apparatus (electric field parameters, volume, surface area, field gradients, work functions, etc...) and the breakdown voltage. The SVM aims to construct a hyperplane in multidimensional parameter space that separates distinct regions or classes. In our case these classes are labelled as breakdown points or stable points. Now, the algorithm assumes that the regions are linearly separable, which we have just suggested may not be the case in a HV problem. This is quickly remedied in the SVM by using the *kernel trick* where a kernel function is applied to the data before it is linearly separated. In this work, the radial basis function (RBF) kernel is used. The RBF maps points in parameter space like,

$$K: x_i \rightarrow e^{-\gamma \|x_i - x_c\|} = y_i \quad 4.4.1$$

where $\vec{x}$ represents a point in parameter space with coordinates x_i , x_c is some reference coordinate (often the mean of the i^{th} coordinate of all training points), and γ is a scale parameter that must be tuned for optimal performance. In this work, the mapped points, $\vec{y}$, are said to reside in the *feature space* while the unmapped points, $\vec{x}$, reside in the *parameter space*.

Application of the SVM approach consists of four stages. The first stage is collecting the training data, the second is training the algorithm, the third is collecting data for the apparatus for which we would like to predict the breakdown voltage (we call this the *target* geometry), and the fourth stage is predicting on the test data with the trained classification algorithm.

Collecting training data consists of its own sub-stages. We must first assemble a small number of geometries on which to conduct breakdown tests in the laboratory. For each of these geometries, the voltage on the cathode is ramped until breakdown is observed and this value is recorded. We label this point as the breakdown threshold and assume that all voltages below

this value are stable while all voltages above are unstable. In general this voltage is understood to be statistically distributed. This point is addressed later in this chapter when an extension to the base algorithm is proposed and partially implemented. Once the breakdown thresholds are obtained, we simulate each geometry with its cathode set to the geometry's experimental breakdown voltage in some FEM software in order to extract relevant electric field parameters for use in the classification algorithm (the software of choice in this work is *COMSOL*). The extracted set of parameters is paired with its corresponding experimental breakdown voltage in order to create a data point for training the classification algorithm. This data point is labelled as a *breakdown point*. The *COMSOL* process is repeated for a few voltages below the breakdown voltage in order to create points with which to train the SVM in the stable parameter space. These new points are labelled as stable data points. We apply a similar procedure with a few voltages above the breakdown voltage and label these as breakdown points. Thus, we end the first stage of the pattern recognition approach with a dataset consisting of several breakdown points and several stable points for each simple geometry that we have considered.

The second stage also involves two sub-stages. In the first stage we scale the input data so as not to potentially confuse the SVM with parameters of different orders of magnitude. Indeed we shouldn't expect each geometry to produce results that are all on the same scale (this is essentially the essence of the HV problem in the first place!). We apply the same normalization technique as the original author [61],

$$x'_i = \frac{x_i - x_{min}}{x_{max} - x_{min}} \quad 4.4.2$$

where x_i is defined as in (4.4.1).

The second sub-stage is straightforward as we simply feed our labelled training data into a binary SVM classifier. For implementing the SVM, we have adopted the SVC algorithm from *Python's sklearn* library and we keep the default parameters for training [92].

Once the algorithm is trained, it is able to take input points of similar structure to the training points and classify them as either breakdown or stable points. To get this data we move on to

stage three of the pattern recognition approach. In this stage there will be some target geometry for which we would like to predict the breakdown voltage. In general we don't want to induce breakdown in this geometry but for the purposes of testing the pattern recognition approach, breakdown was induced in the target geometry (i.e. miniEXO) for comparison with the SVM prediction. Thus, as with the training points, the voltage is ramped in the target geometry until discharge is observed. This point is labelled as the breakdown voltage for the target geometry. The target geometry is then simulated in *COMSOL* to extract relevant parameters. These parameters are fed into the trained SVM which then predicts a stable point or a breakdown point. The typical procedure here is to begin with a *COMSOL* simulation with the cathode set to well below the observed breakdown voltage. We then feed the results to the SVM for prediction. If the point is deemed stable by the SVM, the simulated cathode potential is increased and the procedure is repeated. This is repeated until the SVM finally predicts breakdown. The voltage is then compared to the experimentally observed voltage to quantify the error in prediction.

4.4.2 Application of Pattern Recognition Prediction to miniEXO

In this work, the objective was to predict the breakdown voltage of a TPC – miniEXO. As training geometries, we assembled data from high voltage tests conducted in LXe at SLAC (collected by Peter Rowson and Knut Skarpaas) and XEBRA (collected by Lucie Tvrznikova). The SLAC geometries were all operated in the same test cell where the cathode's electrode geometry could be switched out to create a new geometry (see *Figure 4.8*). We were able to use two different cathodes as training points, namely a micron-scale wire loop and a rounded cylinder (a.k.a *hemisphere*). The anode was a hemispheric in both cases. The geometry from XEBRA was a pair of Rogowski electrodes, with faces parallel to each other and with an adjustable spacing between them (see *Figure 4.10*). Experimental breakdown voltages for each of these geometries were collected prior to the beginning of this work. Thus, creation of training points required only the *COMSOL* simulation step of stage one of the pattern recognition procedure.

Figure 4.8: CAD model of SLAC test cell. *Top:* This particular CAD depicts the hemisphere electrode geometry. The upper hemisphere electrode can be swapped for other electrodes. *Bottom:* Feedthrough of voltage supply for upper electrodes. (These images were provided by Peter Rowson and Knut Skarpaas of SLAC).

Figure 4.9: Wire harness geometry operated in the SLAC test cell. Measurements in left image are given in inches. The holder is made of stainless steel while the wire is phosphor-bronze and has a diameter of 20 microns. *These images were provided by Peter Rowson and Knut Skarpaas of SLAC).*

Figure 4.10: XEBRA's HV test cell. *Bottom middle:* CAD model used for COMSOL simulations of XEBRA's Rogowski electrodes. *Bottom Right:* Electric field strength as simulated by COMSOL with a cathode potential of -75kV. (All images were provided by Dr. Lucie Tvrznikova).

Geometry	Breakdown Voltage (-kV)	Simulated Voltages (-kV)
Hemispheres (1mm Spacing)	25	5, 10, 15, 20 25, 30, 35
Hemispheres (2mm Spacing)	50	10, 20, 30, 40 45, 50, 55, 60
Hemispheres (3mm Spacing)	N/A	15, 30, 45 60, 73
Wire Harness (3mm Spacing)	4	1, 2, 3, 4, 5, 6
Wire Harness (10mm Spacing)	13	10, 11, 12, 13, 14, 15
Rogowski Electrodes (5mm Spacing)	55	15, 35, 55, 65, 70 75, 80, 85, 90

Table 4.1: Complete list of training geometries used in pattern recognition algorithm for predicting the breakdown voltage of miniEXO. The *Simulated Voltages* column refers to the voltages to which the cathodes were set in COMSOL simulations. The breakdown voltage for hemispheres at 3mm spacing is N/A as the system did not discharge by the maximum potential of -75kV that the HV supply could reach.

The CAD model for miniEXO is depicted in *Figure 4.11*. A small volume near one of the acrylic struts was used for parameter extraction for our trained SVM. The results of the SVM prediction are summarized in *Table 4.2*. We see that the SVM predicts a voltage of 250 kV versus the real voltage of 38 kV. Future work should aim to improve the resolution of the SVM prediction. This can be done by using smaller step-sizes in the training voltages summarized in *Table 4.1*. In order to improve accuracy beyond basic order of magnitude agreement, the consideration of the statistical nature of breakdown voltages is suggested. A particular method for incorporating this into the algorithm is outlined in the next section of this chapter. Pruning of the 40-parameter SVM model should also be considered as some parameters may be acting as nuisance parameters. Indeed, in the original paper where the SVM approach was considered, the optimal SVM was pruned down to 26 parameters [61]. A second set of authors that applied the method also found that parameter pruning improved performance [62].

Method	Voltage (-kV)	Breakdown
Experimental	38	Yes
SVM	10	No
SVM	20	No
SVM	30	No
SVM	38	No
SVM	50	No
SVM	70	No
SVM	100	No
SVM	250	Yes

Table 4.2: Breakdown voltage of miniEXO as concluded by experiment, the non-weighted SVM classification algorithm, and the weighted SVM algorithm.

Figure 4.11: CAD model of miniEXO simulated in *COMSOL*. This particular model was created by Tamer Tolba. The simulations relevant to the pattern recognition algorithm were completed by the author. The grid coordinates are displayed in centimetres. This model was built by Tamer Tolba who kindly provided the file to me for use in my thesis work.

4.4.3 A Statistical Extension to the Pattern Recognition Model

It is well known in the field of dielectric breakdown that any given system's breakdown voltage is statistically distributed. By contrast, the model which this work uses as a starting point makes the assumption of a single breakdown voltage. We have attempted to extend this method by using a distribution for the breakdown voltage. In the literature, the *Weibull* distribution has been widely adopted to parameterize the breakdown voltage distribution and it has been adopted here.

For a system described by a variable, t , which denotes the "time-to-failure" of the system, the distribution of failure as a function of time is described by a Weibull distribution if the rate of failure is proportional to a power of time. In the HV problem, there is a 2D parameterization with time-to-failure and voltage-to-failure. That is, the rate of failure depends on the applied voltage and the duration for which a voltage is applied. More details on the nature of the 1, 2, and 3 parameter Weibull distributions and methods for estimating their parameters can be found here [63].

Though it was not explicitly stated, we have actually attempted to incorporate the Weibull distribution parameterization of breakdown voltages to our training points in the previous section of this chapter. However, as breakdown distributions were lacking for some training geometries, we have in actuality implemented a hybrid of the original model with fixed breakdown voltages and the statistical extension. The fixed voltage parameterization was implemented for all geometries provided by SLAC and miniEXO tests from Bern, while the Weibull parameterization was used for XEBRA geometries. For XEBRA, we used the 10% cumulative voltage of the estimated Weibull distribution shown in Figure 4.12 (created and provided by Lucie Tvrznikova).

Figure 4.12: Weibull distribution parameterization of breakdown voltages in XEBRA. Best fit parameters for describing the Weibull are provided on the left of the figure. Figure provided by Lucie Tvrznikova.

The reason for the choice of 10% will become apparent below when the concept of *acceptably low probability of breakdown* is discussed. It is likely that these values are close the median of each geometry's respective Weibull distribution. One would expect a higher breakdown estimate for miniEXO when XEBRA's median is used as a training point as opposed to when its 10% cumulative voltage is used.

In the HV research field, the strengths of dielectrics are typically determined via static tests in which a particular voltage is applied for some varying period of time until breakdown is observed, or via dynamic tests in which voltage is typically ramped until breakdown is observed. Dynamic tests tend to be preferred as data points can be collected more rapidly than in the static case. However, most information regarding the time parameter is lost in such tests. For a collaboration like EXO, static tests are perhaps more relevant as the experiment operates at a fixed potential for years. Unfortunately it is not feasible to conduct static tests where data points are collected at rates on the order of years.

Fortunately there is a solution, described in [53], which can be used to determine static breakdown properties by conducting dynamic tests. The authors begin with the distribution for

the cumulative probability of breakdown for a static test, P_s , described by the cumulative two-parameter Weibull distribution,

$$P_s = 1 - e^{-C t^a E^b} \quad 4.4.4$$

where t is the time over which a static field, E , is applied and a, b, c are the Weibull parameters to be estimated for a particular system. These parameters are described in [53].

The experimenter is free to choose the rate of voltage ramping in an experiment. Supposing, as in [53] that the simple case of ramping at a constant rate is employed, then the field, E , at an instant of time, t , can be expressed as,

$$E = \dot{E} t \quad 4.4.5$$

where $\dot{E}$ is the linear rate of change of the field, or the ramping rate, and the initial field is zero. By using this simple case, the authors have successfully merged static and dynamic experiments as follows: The cumulative probability that breakdown does not occur (i.e. *survival probability*) at a fixed field is first expressed as the product of the survival probabilities over each small increment of time up to some time, t ,

$$P_s = 1 - e^{-C \int_{t'=0}^t a t'^{a-1} E^b dt'} \quad 4.4.6$$

Substituting the instantaneous field described by (4.4.5) into (4.4.6) finally produces the cumulative probability, P_d , analogous to (4.4.4) but for a dynamic test,

$$P_d = 1 - e^{-C \frac{a}{a+b} t^{a+b} \dot{E}^b} \quad 4.4.7$$

It is clear that the dynamic cumulative distribution depends on the same parameters a, b, c as does the static cumulative distribution. There is therefore complete equivalence between static and dynamic tests, at least under the linear voltage/field ramping case. It should be noted, however, that (4.4.7) cannot be used to estimate a & b separately. Instead, a constraint must be added to the optimization problem. This can be achieved by taking data at two separate ramping rates.

In general we suggest that the Weibull distribution extension be implemented by a collaboration like nEXO as follows: Once a , b , c have been estimated for each training geometry to be considered by the pattern recognition algorithm, if we then fix the lifetime, τ , of our target experiment/geometry we can, for each training geometry, solve for the maximum field, E_{max} , for which the probability of catastrophic breakdown is below some acceptable threshold probability, α . Rearranging (4.4.4) the expression becomes,

$$E_{max} = \left(\frac{\ln\left(\frac{1}{1-\alpha}\right)}{c\tau^a} \right)^{\frac{1}{b}} \quad 4.4.8$$

Thus, extensions of the original pattern recognition problem will still give binary classification but the regions into which points are separated will correspond to: *acceptably low probability of breakdown* v.s. *unacceptably high probability of breakdown* instead of *stable* v.s. *breakdown*, where the choice of α defines *acceptably low*. We emphasize now that τ is fixed for all training and test geometries. It corresponds to the lifetime of the target geometry/experiment. The parameters a , b , c , are in general different for each geometry – determined by fitting a Weibull distribution to the breakdown distribution – which ultimately means that E_{max} varies by geometry.

4.5 The HV Lab at Carleton University

To extend the pattern recognition model of the previous section, we need to take further data where we attempt to construct the Weibull distribution for each geometry considered by the algorithm (both training and target geometries). At Carleton we have taken steps to commission a HV laboratory in order continue the HV tests that were completed in Bern. Due to time constraints and various administrative setbacks, the results of these tests are not included in this work. In what follows, the steps that have been taken towards commissioning miniEXO , and the cleanroom in which it is housed at Carleton University, are outlined.

Two of the primary objectives of HV tests that were conducted in Bern were to determine the correct breakdown hypothesis and to determine if glitch patterns could be used to predict

imminent breakdown. The solution to the first objective appears to be that gap breakdown is the limiting cause of discharge in TPCs like EXO-200 and miniEXO. As stated previously, we would like to return to these tests to construct a Weibull parameterization. In addition to this, we would also like to revisit the glitch pattern tests as a flaw in the Bern tests meant that the glitch tests were inconclusive in predicting breakdown. Indeed, during Phase II tests at Bern, the virtual ground was connected to true ground after the cryostat was sealed [43]. At the time it was suspected that the copper support for one of the purity monitors had come into contact with the grounded cryostat wall. Redoing HV tests with the short fixed is considered a priority for future HV tests at the Carleton lab. Indeed, the possibility of monitoring glitch patterns to anticipate imminent breakdown is a subject of enormous interest within the EXO HV group and the HV field in general. Furthermore, the pattern recognition approach and glitch monitoring should really be viewed as complementary to each other. The former allows us to operate in a safe region of parameter space while the latter warns us when the experiment is one of the rare cases in the “safe” region of parameter space that ends up breaking down. For example, while ramping, we may apply an algorithm that has been trained to recognize glitch patterns that precede breakdown. The algorithm could then signal to us to slow ramping or to reduce the voltage. We would then wait for stability and begin ramping again. The warning from glitch monitoring could even allow us to raise the acceptably low probability of breakdown, α , that was described in the previous section of this chapter, as we will have time to respond and lower the applied voltage. However, this kind of glitch monitoring has never been done before and it may turn out that there simply isn’t enough time to respond once a problematic glitch pattern begins – breakdown may be inevitable.

The HV laboratory at Carleton is housed in a cleanroom as the HV tests require conditions such that the LXe can be kept at high purity similar to at EXO-200 (it is generally believed that impurities quench breakdown). Construction of this cleanroom began in August 2016 and was completed in September 2016. The steel frame of the cleanroom is covered in thick, flexible plastic sheets to keep out debris and the ceiling is equipped with four HEPA filters (see *Figure 4.13*). The room is also setup to have a pressure gradient such that air flows out the entrance after coming in through the HEPAs. This prevents accumulation of dust and debris in the lab. To

further reduce the presence of dust, the steel frame and the plastic sheets were cleaned with acetone and anhydrous ethanol. In addition, all equipment was scrubbed with acetone followed by anhydrous ethanol before being brought inside. Special items which had small crevices that were difficult to clean were sprayed with compressed air before undergoing the acetone/ethanol treatment.

Figure 4.13: Carleton HV Group's cleanroom lab as of September 2016.

When the cryostat was opened at Carleton, we were able to confirm that one of the purity monitors had indeed been accidentally grounded which, by extension, grounded the entire virtual ground line, including the glitch detector. This turned out to be due to a loose screw (see *Figure 4.14*). Presumably the screw was lost in one of the many wire feedthroughs on the lid of the cryostat and then fell out after sealing.

Figure 4.14: Red circle: Loose screw causing the purity monitor to connect virtual ground to ground. The Teflon cover prevents contact between the cryostat and the outer vessel after sealing.

It was also noticed that one of the wires connecting to a purity monitor had been severed. The Teflon insulator around the wire appeared to be burned, leading to the suspicion that a large discharge could have been what severed the wire. The cable was replaced with Teflon-insulated, single-thread, silver-plated copper wire. Simple tests with a voltmeter confirmed a current connection between the purity monitor and the readout on the cryostat manifold (see *Figure 4.15*).

Figure 4.15: *Top Left:* Burn on wire connected to ^{241}Am purity monitor. *Top middle:* Connection to purity monitor was partially severed. *Top right:* Feedthrough for wires connecting to components of the inner volume. After it was replaced, the connection to the purity monitor had to be threaded through this feedthrough to reach the readout connection on the cryostat manifold. *Bottom Left:* Inner volume readout connection on for the purity monitor (coaxial connection). *Bottom middle:* Outer connection of purity monitor's readout. *Bottom right:* Showing purity monitor wires connected to the readout after successful soldering and current tests.

After making these repairs, we began preparations for resealing the inner volume of the cryostat housing miniEXO. In order to ensure no leaks through this seal, the residual indium from operation in Bern had to be scraped out of the seal's cavity before being replaced by a new indium wire. This was achieved over two painstaking afternoons by scraping with sharpened acrylic sheets! Once finished scraping, the cavity had to be cleaned with anhydrous ethanol to remove any debris that may have come from the scraping process or from the air. The rest of the cryostat was then cleaned with anhydrous ethanol as well. Finally, a single piece

of indium wire was deposited in the seal cavity and the inner volume was sealed by lowering the TPC, connected to the cryostat lid, into the lower half of the inner volume. The lowering process is depicted in *Figure 4.17*. The final position of the HV TPC in the inner volume is shown in *Figure 4.16*.

Figure 4.16: *Left/Middle:* Seal for the cryostat. The seal is a single piece of indium wire, and the only break is where the two ends meet. The butterfly cross is recommended to prevent leaks at this break. *Right:* Seal cavity prior to addition of indium wire.

Figure 4.17: Lowering the TPC into the inner volume of the cryostat. Lead bricks coated in blue paint can be seen on the lid. Because the lid doesn't have an even weight distribution, these bricks were used to balance the lid in such a way as to ensure an even distribution of pressure when the lid came into contact with the indium seal.

A system was assembled to connect an RGA to the cryostat manifold (see *Figure 4.18*) and we were able to scan the inner chamber and leak check all flanges with a helium source. Sweeping the mass spectrum with the RGA we concluded that there were no leaks (see *Figure 4.19*). The cryostat was then lowered into the outer vessel and the first successful cooldown of the Carleton system was subsequently achieved in September 2018 (I did not participate in the cooldown effort). The next major step in commissioning will be to obtain xenon for filling miniEXO and conducting HV tests.

Figure 4.18: Leak check line connecting RGA to cryostat manifold. **Top left/Top Middle:** Turbo pump mounted on cryostat manifold and connected to a scroll pump for backing. **Top Right:** Scroll pump. The red cable mounted to the steel beam connects to our helium source that was used for leak checking. **Bottom Left/Bottom Middle:** Connection between RGA and inner volume. **Bottom Right:** Analog pressure gauge connected to scroll pump. The line exhausts to the cleanroom air when the turbo is turned off. It then exhausts to the RGA once the turbo is switch on. We used a particle counter to confirm that exhausting the scroll to the room did not damage our cleanroom conditions.

Figure 4.19: Mass spectrum as measured by RGA sweep during leak checking with a Helium source. *Top left:* Mass spectrum. *Top right:* Mass spectrum with pedestal subtracted. Both spectra show a lack of a peak near the helium mass cluster (i.e. near 4 amu). *Bottom:* Helium and Argon partial pressures in the leak line. Both levels satisfy UHV conditions.

Chapter 5

EXO-200 Data Analysis with Machine Learning

This chapter summarizes the work that has been completed by the author as a member of EXO-Deep. EXO-Deep is an analysis group of EXO where we are applying modern machine learning techniques to attempt to match or outperform the traditional analysis methods that are being used for analysis at EXO-200. Although the name *EXO-Deep* implies, and is derived from, *deep learning* (or *feature learning*), this group does not solely employ deep learning algorithms/feature learning algorithms for data analysis at EXO-200. Indeed, this thesis does not deal with any deep learning algorithms and any reference to or mention of deep learning or feature learning should only be understood as a reference to the name of the analysis group unless it is explicitly stated otherwise.

EXO-deep has already shown that we can compete with the traditional methods in energy and light reconstruction [64]. Our focus has now shifted to direct event classification of the various event types in the EXO-200 signal & background fit model. The use of machine learning algorithms for this task is motivated primarily by the fact that they can learn directly from empirical results without requiring the researcher make any assumptions about the properties of the data itself. While assumptions can provide fairly accurate models, we are nevertheless usually not imposing the correct model/representation of the true distribution of the data. Any deviation of this model from the true model results in a reduced ability of the method to make accurate and precise inferences about the properties of the measured data. In a machine learning approach, we only require a large dataset for the classification algorithm to consider. The algorithm then, in an iterative fashion, learns about the underlying properties of the data in the parameter space. Because it only ever references the data itself to update its beliefs about the space, there is no loss of information that arises from making modelling assumptions prior to analysis.

In this chapter we will discuss a particular event-by-event classification algorithm that has been developed for analysis of EXO-200 data. From there we will discuss a method for taking the output of event-by-event classification algorithms and creating confidence intervals for the overall event count from each type being considered by the algorithm. We are unaware of any other particle physics analysis that has been able to construct such intervals without compromising some of the discriminating power of the output of their classification algorithm. Thus, this newly developed method will permit modern, powerful classifiers to be used for publishing physics results independently of traditional methods for the first time.

5.1 Current State of EXO-200 Data Analysis

5.1.1 Signal & Background Model and Variables Considered for Event Discrimination

Before discussing the development of classification methods for data analysis at EXO-200, we must be familiar with the measurements taken for each physics event that occurs within the detection volume of the TPC. We must also consider which types of events – signal and backgrounds – to include in our fit model in the region of parameter space relevant to $0\nu\beta\beta$. These event types become classes in a classification algorithm.

In *Chapter 3*, the EXO-200 TPC was discussed and the various measurements of light and charge collection were listed. These measurements can be used to create at least four variables for data analysis – energy, standoff distance, number of U-wires triggered in coincidence, and risetime of collection wire signals. Energy is obtained by summing the scintillation and ionization signals. The physical significance of these parameters has been discussed in *Chapter 3* so we will provide a short summary here. As was explained in *Chapter 3* the energy parameter is characteristic to an event type in the EXO-200 fit model as they all originate from radioactive decays of characteristic lifetimes. Furthermore, the scintillation and ionization signals for an event depositing energy in the LXe are anti-correlated, so a rotation in the scintillation-ionization space can be used to achieve an improved resolution for this parameter. That is, the energy parameter can be pre-processed before being fed to a learning algorithm. The standoff distance is, loosely, the shortest distance between the location of an event in the drift volume of EXO-200 and any detector component. This allows for discrimination between

ionization/scintillation events that are caused by particles penetrating the detector from external origin and depositing energy near the outskirts of the LXe (e.g. ^{60}Co , ^{232}Th , ^{238}U), and ionization/scintillation events induced by β -decays that are inherent to LXe and which are distributed uniformly throughout the volume (e.g. $0\nu\beta\beta$, $2\nu\beta\beta$, ^{137}Xe). We are able to use this parameter thanks to the 3D event reconstruction capabilities of the EXO-200 detector, summarized in *Chapter 3*. The number of collection wires parameter is related to charge cluster multiplicity of an event depositing energy in the detection volume. The γ -producing backgrounds in EXO-200 tend to produce several charge clusters in the LXe due to Compton scattering, for example, while β -decays tend to produce only one or two charge clusters as they continuously deposit energy through Coulombic interactions like ionization, for example (i.e. continuous paths of ionization versus disjoint clusters like what Compton scattering produces). Thus, γ -producing events can be distinguished from β -decays by counting the number of coincident U-wire signals for a particular event. Finally, the risetime of signals on the U-wires can be used for event discrimination as it varies according to the energy of an event. We don't go into the mathematics of signal analysis in this work but simply state that the correlation exists - it thus serves as a kind of redundancy to the energy measurement. More detailed information about how this measurement is made, as well as agreement between calibration and MC for U-wire signals, can be found in a recent EXO-200 publication [29].

The inclusion of these four parameters for discriminating between event types was justified by Warren Cree in his thesis work [65]. He considered the separability of γ -producing and β -decay distributions using single parameters and determined that including risetime, standoff distance, and number of U-wires in a boosted decision tree score (boosted decision trees are explained in section 5.2 of this chapter), plotting this score against energy and employing multivariate fit to the background model improved the sensitivity of EXO-200 analysis by 15% compared to the then current analysis approach of fitting only standoff distance versus energy [29][65]. Other parameters, such as the induction wire signal, were also considered in Warren Cree's work but, as he explains, they were eventually discarded due to poor shape agreement between calibration data and Monte Carlo data or due to low separability between event types. The

induction wire signal, for example, due to a combination of poor agreement between Monte Carlo and calibration data, and due to low separability [65].

While γ -producing events tend to have low separability between themselves when considering each of the four parameters separately, we see that the four parameters correlate differently among the different γ -producing event types (the correlations between energy and number of U-wires is particularly interesting as is explained and discussed later). The same is true for β -producing event types. This suggests that a classifier can extend beyond discriminating only between β -producing and γ -producing event types, to discriminating directly between all event types. The distributions of the parameters, the single parameter separabilities, and the parameter correlations for each event type are summarized in the figures and tables below. All of these figures and tables correspond to MC datasets that were obtained based on simulations of the EXO-200 detector in *Geant 4*. The validity of the MC model has been studied and verified in several EXO-200 publications [29][37-40][93]. The method for validating MC is primarily based on shape agreement between the distributions of parameter in MC versus the distributions for various calibration sources such as ^{228}Th (described in *Chapter 3*), and ^{60}Co , ^{137}Cs , and ^{226}Ra which are used less frequently than ^{228}Th [29]. In the case of cosmogenic activation events (e.g. ^{137}Xe), Monte Carlo is compared with the measured muon flux at the EXO-200 detector [38]. The separabilities are measured by the Jensen-Shannon Divergence which has a value of zero when distributions are identical and a value of 1 when they don't overlap at all. The Jensen-Shannon Divergence is a symmetric metric so the lower triangular entries are omitted. The diagonal entries are always equal to zero and are also omitted.

Figure 5.1: Distributions, in various parameters, of event types included in the EXO-200 signal & background fit model. Risetime (top left), energy (top right), number of collection wires (bottom left), and standoff distance (bottom right).

	$0\nu\beta\beta$	$2\nu\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
$0\nu\beta\beta$	N/A	0.00124	0.31821	0.13845	0.14280	0.14027
$2\nu\beta\beta$	N/A	N/A	0.32393	0.14308	0.14691	0.01770
^{60}Co	N/A	N/A	N/A	0.04669	0.04458	0.21568
^{232}Th	N/A	N/A	N/A	N/A	0.00090	0.06998
^{238}U	N/A	N/A	N/A	N/A	N/A	0.07362
^{137}Xe	N/A	N/A	N/A	N/A	N/A	N/A

Table 5.1: Pairwise Jensen-Shannon Divergence for collection wire risetime.

	$0\nu\beta\beta$	$2\nu\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
$0\nu\beta\beta$	N/A	0.74046	0.06236	0.19685	0.41757	0.17360
$2\nu\beta\beta$	N/A	N/A	0.47048	0.61307	0.08597	0.36378
^{60}Co	N/A	N/A	N/A	0.12838	0.20196	0.04695
^{232}Th	N/A	N/A	N/A	N/A	0.37683	0.09089
^{238}U	N/A	N/A	N/A	N/A	N/A	0.14259
^{137}Xe	N/A	N/A	N/A	N/A	N/A	N/A

Table 5.2: Pairwise Jensen-Shannon Divergence for energy.

	$0\nu\beta\beta$	$2\nu\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
$0\nu\beta\beta$	N/A	0.00091	0.95986	0.63755	0.63347	0.39064
$2\nu\beta\beta$	N/A	N/A	0.96404	0.64433	0.64008	0.38294
^{60}Co	N/A	N/A	N/A	0.98112	0.98102	0.97663
^{232}Th	N/A	N/A	N/A	N/A	0.00234	0.14918
^{238}U	N/A	N/A	N/A	N/A	N/A	0.13195
^{137}Xe	N/A	N/A	N/A	N/A	N/A	N/A

Table 5.3: Pairwise Jensen-Shannon Divergence for number of collection wires.

	$0\nu\beta\beta$	$2\nu\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
$0\nu\beta\beta$	N/A	0.00004	0.16522	0.05162	0.03212	0.00026
$2\nu\beta\beta$	N/A	N/A	0.16651	0.05221	0.03237	0.00030
^{60}Co	N/A	N/A	N/A	0.03761	0.06159	0.15510
^{232}Th	N/A	N/A	N/A	N/A	0.00355	0.04540
^{238}U	N/A	N/A	N/A	N/A	N/A	0.02700
^{137}Xe	N/A	N/A	N/A	N/A	N/A	N/A

Table 5.4: Pairwise Jensen-Shannon Divergence for standoff distance.

	Risetime	Energy	Number of Collection Wires	Standoff Distance
Risetime	N/A	0.01381	0.39182	0.04475
Energy	N/A	N/A	-0.06393	-0.00848
Number of Collection Wires	N/A	N/A	N/A	0.057428
Standoff Distance	N/A	N/A	N/A	N/A

Table 5.5: Correlations between parameters for $0\nu\beta\beta$ events.

	Risetime	Energy	Number of Collection Wires	Standoff Distance
Risetime	N/A	0.00850	0.40136	0.045298
Energy	N/A	N/A	-0.03463	-0.01178
Number of Collection Wires	N/A	N/A	N/A	0.06567
Standoff Distance	N/A	N/A	N/A	N/A

Table 5.6: Correlations between parameters for $2\nu\beta\beta$ events.

	Risetime	Energy	Number of Collection Wires	Standoff Distance
Risetime	N/A	0.02964	0.13060	0.00536
Energy	N/A	N/A	0.04975	0.07037
Number of Collection Wires	N/A	N/A	N/A	0.05552
Standoff Distance	N/A	N/A	N/A	N/A

Table 5.7: Correlations between parameters for ^{60}Co events.

	Risetime	Energy	Number of Collection Wires	Standoff Distance
Risetime	N/A	0.13141	0.27264	0.02025
Energy	N/A	N/A	0.29081	-0.01666
Number of Collection Wires	N/A	N/A	N/A	-0.00929
Standoff Distance	N/A	N/A	N/A	N/A

Table 5.8: Correlations between parameters for ^{232}Th events.

	Risetime	Energy	Number of Collection Wires	Standoff Distance
Risetime	N/A	0.026851	0.26183	0.01484
Energy	N/A	N/A	0.05269	-0.01900
Number of Collection Wires	N/A	N/A	N/A	-0.00577
Standoff Distance	N/A	N/A	N/A	N/A

Table 5.9: Correlations between parameters for ^{238}U events.

	Risetime	Energy	Number of Collection Wires	Standoff Distance
Risetime	N/A	0.02998	0.39303	0.02967
Energy	N/A	N/A	0.046703	-0.000163
Number of Collection Wires	N/A	N/A	N/A	0.02945
Standoff Distance	N/A	N/A	N/A	N/A

Table 5.10: Correlations between parameters for ^{137}Xe events.

A noteworthy observation from the correlation matrices is that energy and the number of triggered collection wires are negatively correlated for $\beta\beta$ event types while they are positively correlated for other event types. In the current EXO-200 analysis, events are discriminated as either single-site (2 or fewer wires) or multi-site (more than 2 wires). However, it is suggested from these correlation matrices that by adding the number of collection wires to the BDT instead of considering only single-site or multi-site, we can gain substantial discriminating power due to the anti-correlation. Indeed, the more we use energy for discrimination, the more we should use the number of collection wires – grouping into either SS or MS is essentially removing useful information. Furthermore, we see that the correlation between number of U-wires and energy is substantially different when comparing ^{238}U events to ^{232}Th events. This useful property can also be motivated by considering the distribution of the dominant interactions of photons with Xenon atoms at the energies that we consider in EXO-200 (refer to *Figure 5.1*), namely pair production, photoelectric absorption, and Compton scattering. The likelihood of Compton scattering for any given interaction between a photon and a Xenon atom depends on the energy of the photon. Thus, we expect to see a relationship between charge cluster multiplicity (i.e. the number of coincidentally triggered U-wires) and photon energy. In particular, as photon energies decrease towards 1 MeV, Compton scattering becomes a more likely interaction in Xenon. This is exemplified in the next paragraph when considering the distribution of the number of U-wires for ^{60}Co events in comparison to ^{232}Th or ^{238}U events.

We know that neutrinoless double beta decay in ^{136}Xe has a half-life that is likely longer than 10^{26} years (see *Table 3.1* and [24][29]). As a result, strict low background requirements must be imposed in EXO-200 and, eventually, in nEXO in order to obtain a good signal to background ratio in the ROI of our parameter space. All candidate detector components underwent radioassay screening prior to selection for use in the final EXO-200 detector. The EXO-200 detector also operates at a depth of 1500 m.w.e in order to shield the detection volume from background due to cosmic muon activation of detector components. Nevertheless, after an extensive background reduction/prevention effort, some backgrounds persist, contaminating the detection volume and, ultimately, the ROI. A detailed study on the types of events (signal and background) that occur in the EXO-200 detection volume can be found elsewhere [39]. The

result of the study is that, in addition to $0\nu\beta\beta$, five unique event types must be included in the signal & background fit model. The event types in the fit model are summarized in *Table 5.1*. The chosen backgrounds are those that were found to contribute a non-negligible number of counts to the data. The number of counts was determined through a combination of radioassay results and detailed detector simulations. Radioactive contamination is found to be mainly due to the daughter products of the ^{232}Th and ^{238}U decay chains. From the ^{232}Th chain the main concern is the gamma produced in the decay of ^{208}Tl , which, at 2614 keV (100 %), is relatively close to the ^{136}Xe Q-value. From ^{238}U comes the ^{214}Bi isotope, which has two β -decay paths with end-points of 2663 keV (1.7 %) and 3272 keV (18.2 %), and also a gamma at 2447.8 keV (1.5 %). The β -decay paths have energies close to the Q-value but because these events penetrate the detector from external origin, they can be distinguished from $0\nu\beta\beta$ using the standoff distance. The gamma path is distinguished by Compton scattering that produces a higher value for the number of U-wires and, again, by standoff distance. The larger correlation between energy and number of U-wires for ^{232}Th in comparison to ^{238}U is particularly interesting as it allows for discrimination between γ -producing events that are otherwise quite similar. The physical explanation for the difference between the correlations is explained by the decay chains. Clearly, at higher energies, ^{238}U tends to undergo β -decay instead of γ -decay while ^{232}Th only undergoes γ -decay. We have already explained the mechanisms that result in β -producing events triggering fewer U-wires in the detector than γ -producing events. Thus, at higher energies, ^{232}U and ^{232}Th become easier to distinguish based on counting the number of triggered U-wires in addition to the event energy. From cosmogenic activation of copper in the detector vessel and the TPC we get the ^{60}Co isotope which releases two prominent gammas, at 1173.2 keV and 1332.5 keV (99.88 %), in succession which causes a sum peak at 2505.7keV. This is close to the Q-value so we must look to the other parameters to discriminate this from $0\nu\beta\beta$. Fortunately Compton scattering of these two photons is highly likely and this produces many charge clusters. Indeed, Compton scattering at these energies is more likely than at the higher energies of the gammas from ^{232}Th and ^{238}U . Ultimately, this increases the number of triggered U-wires and makes ^{60}Co easily distinguishable from $0\nu\beta\beta$ (see *Figure 5.1* – bottom left). Through cosmogenic activation, ^{137}Xe is produced in an excited state capable of β -decay with a

Q-value of 4173 keV [38]. All of the events just mentioned are all fairly close to the $0\nu\beta\beta$ Q-value of 2457.83 keV and all are found to produce a non-negligible number of counts in the ROI.

Event Type	Origin
$0\nu\beta\beta$	^{136}Xe decay
$2\nu\beta\beta$	^{136}Xe decay
^{60}Co	Cosmogenic activation of Cu vessel
^{232}Th	Underground-inherent/far source
^{238}U	Underground-inherent/far source
^{137}Xe	Cosmogenic/neutron capture in ^{136}Xe

Table 5.11: Signal and background model for fitting in EXO-200 data analysis [38].

5.2 Event-by-Event Classification with EXO-200 Data using Ensemble Methods

In this section, the construction of various classification algorithms for event-by-event discrimination of the various event-types considered in the EXO-200 fit model are summarized. These classifiers all make use of the four measured parameters that were listed in section 5.1 of this chapter – no feature learning is employed.

An ensemble method in machine learning is an approach to learning in which multiple learning algorithms are combined to improve upon the result, with respect to some metric, that any one of the individual algorithms in the combination would produce on its own. These metrics are often related to error rate in classification as well as generalizability to new datasets. A simple example of an ensemble is a majority vote algorithm. In this scenario, several classifiers are trained to classify events as one of at least two classes. In the context of EXO-200, each classifier is trained to discriminate between events of types $0\nu\beta\beta$, $2\nu\beta\beta$, ^{60}Co , ^{232}Th , ^{238}U , ^{137}Xe . When considering a particular event, each classifier votes for which class it believes to be the true class of the event. The event is then classified as the class which received the most votes. In the scenario where all of the weak learners use the same base algorithm, the ensemble is called homogenous. Otherwise, the ensemble is called heterogeneous. In this work, we use homogenous ensembles when the boosting technique is employed and heterogeneous

techniques when stacking is employed. Each method is described with more detail in the rest of this chapter.

5.2.1 Boosting

In the context of classification algorithms the so-called *boosting* ensemble method is an approach to learning in which several weak learners are combined by a weighted average to produce a strong learner. A weak learner is often defined as a fast algorithm, such as a decision tree, that performs just better than random guessing (boosting will fail if the weak learner does worse than random guessing) while a strong learner tends to have a much higher accuracy. However, the reader should not dwell on the exact distinction between weak learners and strong learners in general. The distinction is more useful in a relative context – in the case of boosting the weighted average of the so-called weak learners should certainly be a stronger learner than any of the single trees. Therefore the weighted average is referred to as a strong learner while its contributors are referred to as weak learners.

Though the particular boosting method may vary in general, in this work we have considered adaptive boosting (AdaBoost). The reason for employing AdaBoost is that events that were misclassified in a particular weak learner within the strong learner should be given more weight in the subsequent weak learner, such that the latter weak learner is biased towards correcting the “mistakes” in the former weak learner. By contrast, and to enforce normality, correctly classified events have their weights reduced in the next weak learner. The pseudocode for a binary (i.e. two-class or two-event-type) AdaBoost algorithm is shown in *Figure 5.2*. The generalized multi-class version of this code is nearly identical to the binary case. It is derived from the statistical interpretation of AdaBoost here [66].

-
- Given training data: $(\vec{x}_1, y_1), \dots, (\vec{x}_m, y_m)$, where $\vec{x}_i \in X$, $y_i \in \{-1, 1\}$
 - Initialize uniform weights for weak learner D : $D_1(i) = \frac{1}{m}$, for $i = 1, \dots, m$
 - For $t = 1, \dots, T$:
 - Train a weak learner using distribution D_t
 - Obtain weak hypothesis, $h_t: X \rightarrow \{-1, 1\}$
 - Objective: select h_t with small error,

$$\varepsilon_t = P_i[h_t(\vec{x}_i) \neq y_i; D_t]$$

- Choose weight, $\alpha_t = \frac{1}{2} \ln \left(\frac{1-\varepsilon_t}{\varepsilon_t} \right)$
- Update:
For $i = 1, \dots, m$:

$$D_{t+1} = \frac{D_t(i) e^{-\alpha_t y_i h_t(\vec{x}_i)}}{Z_t}$$

, where Z_t ensures normalization for D_{t+1}

- Aggregate weak learners into strong learner,

$$H(\vec{x}) = \text{sign} \left(\sum_{t=1}^T \alpha_t h_t(\vec{x}) \right)$$

Figure 5.2: Pseudocode for AdaBoost. In the final step for the strong learner, the sign of the weighted average – the summation – is taken. This is to classify events as “-1-like” events or as “1-like” events. In practice the -1 and the 1 are placeholders for events of a particular nature such as beta-producing or gamma-producing events or, in the case of this analysis, as any of $0\nu\beta\beta$, $2\nu\beta\beta$, ^{60}Co , ^{232}Th , ^{238}U , ^{137}Xe . The above algorithm must be slightly modified for the latter 6 class case, however [66]. In this algorithm ε_t denotes the error of the t^{th} weak learner, α_t denotes the weight of the t^{th} weak learner in the weighted average, h_t denotes the prediction of the t^{th} weak learner, and D_t denotes the distribution of data points that is updated based on the rate of misclassification.

We can see that the algorithm works by aggregating several weak learners into a strong learner – a weighted average of the individual weak learners. Another important point is that weak learners with good performance (i.e. high classification accuracy) have a large weight in the weighted average while poorly performing weak learners have a low weight.

In *Figure 5.3* the weighted average procedure is illustrated for weak learners trained on two input features. We see that simple, linear cuts in each weak learner come together in the strong learner to produce a complex decision boundary.

Figure 5.3: Illustration of weighted average procedure in AdaBoost. Three weak learners acting as binary classifiers are combined to create a more complex decision boundary with superior classification accuracy to any of the individual learners.

5.2.2 Decision Trees

In this work, the weak learners used in AdaBoost have always been decision trees. When decision trees are used as the weak learners in AdaBoost, the strong learner is referred to as a *Boosted Decision Tree* (BDT) or a *forest*. A prerequisite to implementing a BDT is the decision tree itself. These can be understood by considering the binary decision tree (BiDT). A BiDT, graphically depicted in *Figure 5.4*, is a decision tree in which each node has two output options representing distinct event types in the context of a physics experiment.

Figure 5.4: Graphical representation of the logic of a BiDT. This particular tree would take two input variables, x & y , to make its cuts in classifying events as signal or background. The thresholds of the cuts are denoted by c, d, e which are learned thresholds on parameters, chosen to optimize discrimination.

The so-called *root node* is the first node in the tree and takes the original dataset $(X, \vec{y})$ as input, where an entry of X is some event/point, $\vec{x}_i$, and y_i is its corresponding event type. The algorithm splits the dataset into two subsets according to some learned threshold on one of the parameters that describes each $\vec{x}_i$. In our case, the threshold is set according to the *Gini Purity* metric and is optimized to separate distinct event types into distinct subsets. The parameter on which the threshold/cut is learned is chosen by iterating over the parameters that describe each point, $\vec{x}_i$ in the learned parameter space. After the cut on the root node, the tree increases in depth by adding nodes to each of the two original subsets/nodes. These new nodes are split by the same procedure as the root node – splitting the subsets into smaller subsets using some threshold on some parameter, learned by applying the *Gini Purity* metric. In *Figure 5.4*, the chosen parameters are denoted by x_i, x_j, x_k and the corresponding thresholds for cutting on these parameters are denoted by c, d, e respectively. The last set of subsets in the tree corresponds to signal-like and background-like events.

The multi-class version of decision trees is similar to the BiDT. In the multi-class case, the number of subsets after splitting a node will be equal to the number of classes. To achieve this,

the number of cuts at each node will be equal to the number of classes minus one. The process is otherwise the same as in the BiDT case.

5.2.3 Boosted Decision Trees Applied to EXO-200 Fit Model

Much of the effort in understanding the workings of a BDT is exerted in understanding the BiDT. The rest of the algorithm has already been summarized in the pseudocode of *Figure 5.2*. In the particular instance of this code where we use BiDTs as weak learners, the hypotheses, h_t , represent the cut (or combination of cuts) of each decision tree in the forest.

The forest that was developed for direct event classification of the EXO-200 fit model was trained and tested on Monte Carlo (MC) datasets simulated using *Geant 4*. The so-called *learning curves* of the forest are shown in *Figure 5.5*. The horizontal axis denotes the number of trees in the forest while the vertical axis denotes the error of the complete forest. The forest was simultaneously trained on 30,000 events from each of the six event types in the EXO-200 fit model and the events had equal weights. The choice of 30,000 events was determined via trial and error with different numbers of events. Using a lower number such as 10,000 produced a larger discrepancy between the learning curves of the training set and test set. Using more than 30,000 did not produce any significant improvement and so we used 30,000 to limit the computational complexity of the training, testing, and tuning steps of building the final forest. The test set consisted of 48,000 events of each of the six event types. This number was somewhat arbitrary for most of the event types. For ^{60}Co , only 48,000 events remained after using 30,000 for the training set. This number produced a small standard deviation for the estimates of the classification rates of ^{60}Co ($< 1\%$ of the classification rate in most cases – see *Figure 5.6*). Thus, to reduce computational complexity during the lengthy algorithm tuning campaign it was decided that the standard deviation was small enough with 48,000 events and so this number was used for each event type (i.e. the test set consists of 48,000 events of each of the six event types in the EXO-200 fit model). Equation (5.3.3) shows how the classification estimates are made and (5.3.4) shows how the variances of the estimates shrink with larger test sets.

The ultimate metric for performance of an analysis method in a rare event search such as EXO-200 is a sensitivity analysis (this is described in *Chapter 3*). However, prior to this work, there did not exist any method for doing statistical inference about parameters describing an entire dataset based on the output of direct event-by-event classification algorithms applied to the elements of the dataset. A formalism for statistical inference that can be used for such an analysis was therefore developed as part of this work and is outlined in section 5.3 of this chapter. Consequently, the objective of this work was not to conduct a sensitivity study, but to provide a proof of concept for the method outlined in section 5.3. As a result, sensitivity was not directly considered as a metric of performance when developing the algorithms that were developed for this work. The metric of performance was taken to be the average classification accuracy of all six event types in the fit model (e.g. the average of the diagonal entries of *Figure 5.6*).

As we add trees, the forest’s prediction on the training and test sets clearly improves. Eventually, though, we see that the curve begins to flatten, even as several thousand more trees are added to the forest. Thus, keeping all of the trees in the flat region is unnecessary, while removing these unnecessary trees from the forest reduces computational complexity in future steps such as Monte Carlo or bootstrap studies for statistical inference. To remove the unnecessary trees, we plot the learning curve and cut the forest somewhere along the lowest valley on the test set. Long valleys are convenient because they suggest further generalizability (imagine cutting at the approximate vertex of, say, a parabola versus a flat line ... missing the vertex results in large increase in error while you can’t really “miss” with a long valley). The length of these valleys can be manipulated by adjusting the so-called learning rate hyperparameter of the AdaBoost algorithm – smaller learning rates stretch the curve while larger rates compress the curve. In loose terms, the learning rate adjusts the stepsize in a gradient descent algorithm.

Figure 5.5: Learning curves for a BDT trained on EXO-200’s 6-class signal/background fit model. The dotted line denotes the performance of a single decision tree. Evidently the boosting procedure improves discriminating power. The confidence bands correspond to 68 % confidence intervals. The training (test) interval was obtained by bootstrapping the results on the training (testing) sets at each forest size (i.e. each number of trees on the horizontal axis). The error was bootstrapped 1000 times at each forest size.

A learning rate of 0.05 was found to be sufficiently low for our test set. The forest initially consisted of 25000 trees (here only 13000 are depicted since the entire region from 13000 to 25000 is flat) and was pruned to include only the first 8000 trees. The resulting performance of the forest after pruning to 8000 trees is summarized by its *classification matrix*, shown in *Figure 5.6*. The row index corresponds to the true class of an event while the column index labels the resulting classification. This entire matrix comes from training and testing on distinct Monte Carlo datasets.

On the row below the matrix, the performance on a calibration ^{228}Th dataset is shown. As explained in *Chapter 3*, at EXO-200 we often calibrate with a ^{228}Th source. Thus, we have tested our MC-trained BDT on a ^{228}Th calibration dataset to see if it generalizes well to real physics events. The reader may be concerned that the algorithm was trained on MC-produced ^{232}Th

events, not ^{228}Th events. However, these two atoms should not be explicitly considered as different event types in the context of our classification problem. First of all, both atoms are part of the same decay chain. The ^{232}Th isotope decays via α -decay followed by two β -decays to become a ^{228}Th isotope. Furthermore, the ^{232}Th events in MC simulations of the detector are so-called *far source* events that originate in regions well outside of the EXO-200 detector. Thus, many of these events do not deposit energy in the TPC via any of the decays mechanisms that precede the appearance of ^{228}Th in the decay chain. Therefore, many of the events in the MC dataset are effectively ^{228}Th events. We therefore expect that our algorithm will learn to classify ^{228}Th events as ^{232}Th events.

The performance of the algorithm on the MC and calibration sets was checked and we have bootstrapped the error rates on each set. We have found that the difference in performance between the sets is statistically significant before considering systematic errors between Monte Carlo and calibration studies (see *Figure 5.7*).

	$0\nu\beta\beta$	$2\nu\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
$0\nu\beta\beta$	0.811 ^{+0.002} _{-0.003}	0.055 ^{+0.001} _{-0.001}	0.011 ^{+0.0005} _{-0.001}	0.072 ^{+0.001} _{-0.002}	0.0365 ^{+0.0009} _{-0.001}	0.0148 ^{+0.0005} _{-0.001}
$2\nu\beta\beta$	0.057 ^{+0.001} _{-0.003}	0.801 ^{+0.002} _{-0.004}	0.0025 ^{+0.0002} _{-0.0002}	0.00023 ^{+0.00006} _{-0.0001}	0.139 ^{+0.002} _{-0.003}	0.0003 ^{+0.00008} _{-0.0001}
^{60}Co	0.042 ^{+0.0009} _{-0.002}	0.0046 ^{+0.0003} _{-0.001}	0.725 ^{+0.002} _{-0.004}	0.104 ^{+0.001} _{-0.003}	0.124 ^{+0.002} _{-0.004}	0.0003 ^{+0.00008} _{-0.0001}
^{232}Th	0.151 ^{+0.002} _{-0.004}	0.0440 ^{+0.0009} _{-0.002}	0.162 ^{+0.002} _{-0.004}	0.492 ^{+0.002} _{-0.005}	0.078 ^{+0.001} _{-0.002}	0.073 ^{+0.001} _{-0.003}
^{238}U	0.105 ^{+0.001} _{-0.003}	0.217 ^{+0.002} _{-0.004}	0.179 ^{+0.002} _{-0.003}	0.044 ^{+0.0009} _{-0.0022}	0.446 ^{+0.002} _{-0.005}	0.009 ^{+0.004} _{-0.001}
^{137}Xe	0.224 ^{+0.002} _{-0.004}	0.146 ^{+0.002} _{-0.004}	0.0205 ^{+0.0006} _{-0.002}	0.119 ^{+0.001} _{-0.004}	0.072 ^{+0.001} _{-0.003}	0.419 ^{+0.002} _{-0.005}
^{228}Th	0.227 ^{+0.002} _{-0.004}	0.00030 ^{+0.00008} _{-0.0002}	0.212 ^{+0.002} _{-0.004}	0.468 ^{+0.002} _{-0.002}	0.0320 ^{+0.0008} _{-0.002}	0.061 ^{+0.001} _{-0.002}
Calibration						

Figure 5.6: The classification matrix for the optimized BDT. Each entry denotes the probability that a true event, given by the row index, is classified as a particular event-type, given by the column index. The classifier was trained on MC events generated using Geant-4 and the result of this matrix are from testing on a MC dataset. Below the matrix is the result of testing the MC-trained BDT on events from a ^{228}Th calibration source located at position S5 in the detector (see *Figure 3.13*). Bold entries denote correct classification rates, red entries indicate the class to which each true event type is most often misclassified, and green entries indicate the class to which each true event type is most rarely misclassified. Errors correspond to 68 % intervals.

Figure 5.7: Histograms of the bootstrapped MLE for the correct classification rate of ²³²Th using MC data (left) and ²³²Th calibration data (right). The calibration data is obtained by placing the ²²⁸Th source at position S5 outside of the TPC volume (see Figure 3.13). The MC ²³²Th data is simulated as penetrating the detector approximately uniformly from all sides. The bootstrap procedure was repeated 5000 times for both MC and calibration data. A Gaussian fit is superimposed on each histogram.

Figure 5.8: Learning curves for performance of BDT on MC ²³²Th (left) and calibration ²²⁸Th (right).

While there is undoubtedly a difference in performance between MC and calibration data we find solace in the learning curves of the MC and calibration data, both depicted in Figure 5.8. The curves are identical, even in their fine structure, apart from a small pedestal. This striking resemblance in the fine structure suggests that it is always the same events which are being misclassified in the calibration dataset. This may be partly due to mislabelling of the calibration set. Indeed, we assume that all recorded events are from the Thorium calibration source in the calibration set but there is no reason for the various other event types to not contribute, at

least in part, the total event count; Nature does not stop to let us calibrate our detector! However, the calibration source certainly produces the overwhelming majority of counts during a calibration run. It is estimated that less than $1/1000^{\text{th}}$ of events are not actual Thorium events. Thus, mislabelling cannot account for the entire 3 % difference in error rate between MC and calibration Thorium (the 3 % difference in error rate should be read as the difference between the rates at the cut location along the learning curve only, and not as a constant ratio along the entire learning curve). In addition to this trivial explanation is the fact that the ^{228}Th calibration data does not penetrate the TPC uniformly as the MC ^{232}Th does. Rather, the source is placed just outside of the TPC at calibration position S5 (see *Figure 3.13*). There is therefore an imposed difference between how the two sources penetrate the detector and the ^{228}Th calibration distribution is not completely faithful to the true distribution of real ^{228}Th events of external origin that penetrate the detector (or vice versa). This difference could be enough to cause a small amount of ^{228}Th calibration events to jump across the decision boundary in the parameter space created by the forest, resulting in misclassification as some other event type. For example, ^{60}Co events actually do originate within detector components and could originate near the S5 calibration location [39]. Thus, it may be possible that the forest trained on MC ^{228}Th occasionally classifies a ^{228}Th calibration event as a ^{60}Co event more often than it would misclassify a real, non-calibration ^{228}Th event as a ^{60}Co event. If this were true, we would expect this to be reflected in the classification rate of ^{228}Th to ^{60}Co to be increased relative to the ^{232}Th to ^{60}Co rate. This effect can indeed be observed in *Figure 5.7*. This hypothesis remains to be verified in a future analysis.

In current EXO-200 analysis, there is a systematic error between MC ^{228}Th and calibration ^{228}Th . When considering the fraction of ^{228}Th events that are MS events, the analysis find a systematic difference of as much as $\approx 10\%$ at some energies of ^{228}Th events [24][29][93]. If such systematic errors were to be considered here, it is plausible that the distributions in *Figure 5.7* would overlap and that there would be no significant difference between the performance of the forest on ^{228}Th calibration data and ^{228}Th MC data. This hypothesis remains to be verified in a future analysis.

5.2.4 Diversity Pruning

In general, independence between classifiers in an ensemble is considered an asset in improving the error rate and the generalizability of the ensemble [67-69]. One measure of independence is diversity, where two classifiers are said to be diverse if they often classify identical events differently from each other. As such, we have attempted to improve our initial AdaBoost results by including a diversity parameter in the construction of the forest. The results of this endeavour are summarized in this section.

To incorporate diversity into our algorithm, the initial forest was pruned using a diversity metric. The diversity-based pruning procedure that was used, originally proposed by Diettrich and Margineantu, consists of four steps [69]. First, we train a large forest on our MC training set. We then compute the *Cohen Kappa Score*, k , between all pairs of trees in the forest. This score is defined by,

$$k = \frac{p_0 - p_e}{1 - p_e} \quad 5.2.1$$

where p_0 is the probability that two trees agree on the classification of a given event, and p_e is the probability of agreement by chance (i.e. 1/6 in our six-class case). Next, we remove all trees from the forest. Finally, we add trees back to the forest in pairs from the lowest kappa score (i.e. lowest rate of agreement) to the highest. If the pair of trees to be added includes a tree that has already been added back to the forest, we simply skip this pair and move on to the next candidate pair.

The learning curves of the original forest and the forest after diversity pruning are shown in *Figure 5.9*. The improvement due to pruning is an absolute error rate reduction of 0.04549 %. The 68% confidence band for the forest without diversity is depicted in *Figure 5.5*. At the point of the cut for the forest without diversity (i.e. 8000 trees) the error rate is $(38.6 \pm 0.9) \%$ (computed by averaging the errors along the diagonal of *Figure 5.5* and adding the corresponding variances in quadrature). Thus, an absolute error rate reduction of 0.0459 % produces a rate within the 68 % confidence bound at the cut location. However, *Figure 5.9*

illustrates a systematic error in rate reduction along the two learning curves. This suggests that the pruned forest has learned additional information about the parameter space and that it could be beneficial to include diversity considerations in the forest. This remains to be studied further in order to draw a final conclusion on the matter.

Figure 5.9: Learning curves for a MC test set with (red) and without (blue) diversity pruning. There is a modest improvement after pruning (green circle).

We have departed slightly from Diettrich's original algorithm in the diversity pruning process. In the original paper, all of the trees in the forest had similar error rates. In our forest, this is not at all the case, as shown in *Figure 5.10*. If two trees have a large difference in error rate, they will undoubtedly have a very low kappa score. In our case, this results in trees that were constructed early in the AdaBoost algorithm having very low kappa scores when paired with trees that were constructed late in the AdaBoost algorithm. We therefore end up repopulating our forest by essentially grabbing pairs at opposite ends of the error curve. In this sense, our diversity metric is being distorted; we are not actually achieving what we want to achieve, which is to select classifiers which perform well in different regions of the parameter space, complementing each other in a weighted average. If one tree *A* has a low error rate and tree *B* has a high error rate, then the region in which *B* does well likely has a large overlap with the region in which *A* does well. Outside of this overlap, *A* still has a large region in which it does

well while B almost always disagrees incorrectly. Thus, we gain essentially nothing by adding B to A and, in fact, probably reduce the performance due to the large region in which B incorrectly disagrees.

In constructing *Figure 5.9* we have modified the pairwise repopulation step as follows: We look for a region on the tree-by-tree error curve in which the tree-by-tree error rate is approximately constant or which varies only slightly relative to the average error rate. We then keep all trees that were built prior to this region and only consider the kappa pruning procedure in the trees built after these first 8000. Since all trees have approximately the same error rate in this region, the average error of the pairs will be essentially equal to the error of each individual tree. Thus, we have a more “pure” consideration of diversity during our pruning process than if we were to consider all trees, where the diversity metric is clearly affected by overall error itself. Instead, we want diversity based on *what* is being learned. For example, if the error rates of the trees are equal but the kappa score is low, then perhaps they perform well in different regions of the parameter space. Thus, by diversity-pruning with trees of comparable error rates, we remove redundancy, the presence of which can be shown to hinder performance in AdaBoost. We also avoid the result discussed earlier where adding tree B to tree A actually worsened performance due to the large difference in error rates of the two trees.

Figure 5.10: Error of each tree plotted in order of addition to the overall forest (left). Evidently the error rate of a given tree tends to increase the later it is constructed and added to the forest. Of course, the weights of each tree correspondingly decrease the later the tree is constructed (right).

5.2.5 Stacking BDT with Naïve Bayes and KNN

A second ensembling method referred to as *stacked generalization*, or simply *stacking*, was also implemented in the course of this work. This approach seeks to combine the output of several learning algorithms (base-learners) in order to build a final *stacked* meta-algorithm. The combined learner is called a meta-learner or meta-algorithm because rather than using the original input features of an event, it often takes the output (meta-features) of the base learners as input.

A good stacking algorithm will successfully identify the mutual information between base-learners in order to avoid confirmation bias/correlations. This can be interpreted to mean that the algorithm tries to remove mutual information in such a way as to make the output of algorithm *A* interpretable as a prior probability for algorithm *B* in a Bayesian updating of belief/probability. In general one expects high correlation between good classifiers, so using one to obtain a prior probability for the other is a flawed use of Bayes' formula. If mutual information is removed, it becomes appropriate to do this. A more complete description of stacking can be found here [70].

If we consider the interpretation of stacking that was just given, then the base learners of a stacked learner should be diverse or independent to achieve good performance. In this case, high diversity equates to reduced mutual information. In this work we have combined a BDT, a *K Nearest Neighbors* (KNN), and a *Naïve Bayes* (NB) to form a stacked ensemble. We have already mentioned the importance of diversity when combining classifiers and stacked ensembles to benefit from diversity of their component classifiers. In this case, we don't consider the rate of agreement between learners as our measure of diversity/independence. Rather, each of the base algorithms has a unique method for discriminating between event types. In a sense, they are each learning something different about the parameter space. The idea, then, is that because the algorithms are attempting to distinguish by different methods, they will hopefully complement each other and provide a better aggregate performance, particularly near the boundaries of overlapping PDFs in our multidimensional parameter space.

A NB classifier is fundamentally different to most classifiers in that it assumes that all input features are independent of each other. Therefore, unlike a boosted decision tree, for example, it does not map classification regions by finding the differences between classes with respect to how input features correlate differently between each other within a certain class. For example, the way that energy and standoff distance are correlated to each other for $2\nu\beta\beta$ events, is different to how the same two variables correlate to each other for Th-232 events. A boosted decision tree attempts to carve these correlations out of the parameter space while a NB does not. This diversity in the learning method compared to the BDT, as well as the simplicity and speed at which a NB classifier can be trained, are the reasons that it was included in the stacked ensemble. The classification matrix for a NB trained on the same data as our BDT is shown in *Figure 5.11*.

	$0\nu\beta\beta$	$2\nu\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
$0\nu\beta\beta$	$0.844^{+0.002}_{-0.002}$	$0.064^{+0.001}_{-0.002}$	$0.0170^{+0.0006}_{-0.001}$	$0.0252^{+0.0007}_{-0.002}$	$0.0278^{+0.0008}_{-0.002}$	$0.019^{+0.006}_{-0.005}$
$2\nu\beta\beta$	$0.059^{+0.001}_{-0.002}$	$0.882^{+0.001}_{-0.001}$	$0.0041^{+0.0003}_{-0.0002}$	$0.0004^{+0.00006}_{-0.00005}$	$0.052^{+0.002}_{-0.003}$	$0.00031^{+0.00008}_{-0.0001}$
^{60}Co	$0.055^{+0.001}_{-0.001}$	$0.0240^{+0.0002}_{-0.001}$	$0.727^{+0.002}_{-0.002}$	$0.066^{+0.001}_{-0.003}$	$0.124^{+0.002}_{-0.002}$	$0.001^{+0.00008}_{-0.0001}$
^{232}Th	$0.227^{+0.002}_{-0.004}$	$0.074^{+0.001}_{-0.001}$	$0.194^{+0.002}_{-0.004}$	$0.338^{+0.002}_{-0.002}$	$0.051^{+0.001}_{-0.002}$	$0.114^{+0.001}_{-0.003}$
^{238}U	$0.116^{+0.001}_{-0.003}$	$0.346^{+0.002}_{-0.004}$	$0.181^{+0.002}_{-0.003}$	$0.028^{+0.0009}_{-0.0022}$	$0.318^{+0.002}_{-0.002}$	$0.0091^{+0.0004}_{-0.0005}$
^{137}Xe	$0.253^{+0.002}_{-0.002}$	$0.169^{+0.002}_{-0.001}$	$0.0260^{+0.0078}_{-0.0008}$	$0.087^{+0.001}_{-0.004}$	$0.0450^{+0.0009}_{-0.001}$	$0.417^{+0.002}_{-0.002}$

Figure 5.11: Classification matrix of a Naïve Bayes' trained and tested on MC datasets (the same datasets as were used to train the BDT). Row indices correspond to the true type of an event while column indices correspond to the type to which an event was classified. Errors correspond to 68% intervals and were obtained via bootstrapping the dataset 100,000 times.

A KNN algorithm assigns an event to the class that is most common among its closest neighbouring points in the training set. Thus, the KNN is directly considering the density of points, essentially ignoring the values of the parameters describing each event in the parameter space. This is quite different to both the BDT and the NB learning approaches. Thus, the KNN may be able to perform better than the BDT or the NB in certain regions of the parameter space where the learning method of the BDT and the NB have trouble. This is simply a hypothesis and may not occur in practice. However, if it does, inclusion of the KNN for diversity in learning method is useful. This is the reason that the KNN was included in the stacked ensemble. The basic KNN is simple to implement as the user is only required to tune a single parameter,

namely the number of neighbouring points to consider during classification. Now, the KNN algorithm can be extended to operate as a weighted majority vote by weighting votes according to the distance between training points and new test points. Several distance metrics can be used in the weighting process (e.g. Euclidean, Manhattan, other Minkowski distances). It has been shown theoretically elsewhere that a fractional Minkowski distance is more robust to noise when compared to the Euclidean, Manhattan, or other integer Minkowski distances [71][72]. With this in mind and knowing that we certainly have a noisy dataset in any particle physics experiment, we have iterated over various fractional and integer values for the Minkowski metric and found that a p value of $1/10$ gives the best performance for our particular dataset. The general Minkowski distance between two points $\vec{x}, \vec{y}$ is given by,

$$D(\vec{x}, \vec{y}) = \left(\sum_{i=1}^d |x_i - y_i|^p \right)^{1/p} \quad 5.2.2$$

where d is the dimension of the parameter space in which a data point or event is plotted. The Euclidean and Manhattan distances are instances of the Minkowski distance with p equal to 2 and 1 respectively. Learning curves for a KNN using various values of p are shown in *Figure 5.12*, where we tested on 1200 MC events (200 of each event type) and trained on the same dataset as the BDT. We find that the KNN with a p value of $1/10$ and one which considers the 56 closest neighbours gives the best performance. The classification matrix for the KNN algorithm with a p value of $1/10$ and 56 neighbours, trained on the same MC dataset as our BDT, is shown in *Figure 5.12*.

Figure 5.12: Learning curves of the KNN for various values of p in the Minkowski distance metric. The horizontal axis denotes the number of neighbours considered by the algorithm in the voting step while the vertical axis denotes the overall error on 1200 test points. The test points are evenly distributed among the six event types (i.e. 200 points per event type).

	$0\nu\beta\beta$	$2\nu\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
$0\nu\beta\beta$	$0.802^{+0.002}_{-0.002}$	$0.055^{+0.001}_{-0.002}$	$0.0160^{+0.0006}_{-0.0005}$	$0.071^{+0.001}_{-0.002}$	$0.0394^{+0.0009}_{-0.0009}$	$0.0126^{+0.0005}_{-0.0005}$
$2\nu\beta\beta$	$0.058^{+0.001}_{-0.001}$	$0.794^{+0.002}_{-0.002}$	$0.0016^{+0.0002}_{-0.0001}$	$0.00041^{+0.00006}_{-0.00006}$	$0.143^{+0.002}_{-0.002}$	$0.00030^{+0.00008}_{-0.00009}$
^{60}Co	$0.045^{+0.001}_{-0.001}$	$0.012^{+0.0002}_{-0.001}$	$0.720^{+0.002}_{-0.002}$	$0.082^{+0.001}_{-0.003}$	$0.139^{+0.002}_{-0.002}$	$0.000^{+0.00008}_{-0.0001}$
^{232}Th	$0.154^{+0.002}_{-0.004}$	$0.049^{+0.001}_{-0.001}$	$0.209^{+0.002}_{-0.004}$	$0.453^{+0.002}_{-0.002}$	$0.075^{+0.001}_{-0.002}$	$0.057^{+0.001}_{-0.003}$
^{238}U	$0.101^{+0.001}_{-0.001}$	$0.225^{+0.002}_{-0.002}$	$0.179^{+0.002}_{-0.002}$	$0.0382^{+0.0009}_{-0.0009}$	$0.447^{+0.002}_{-0.002}$	$0.007^{+0.0004}_{-0.0005}$
^{137}Xe	$0.227^{+0.002}_{-0.002}$	$0.143^{+0.002}_{-0.001}$	$0.030^{+0.0007}_{-0.0007}$	$0.0137^{+0.002}_{-0.002}$	$0.045^{+0.0009}_{-0.0009}$	$0.389^{+0.002}_{-0.002}$

Figure 5.13: Classification matrix of a KNN trained and tested on MC datasets (the same datasets as were used to train the BDT). Row indices correspond to the true type of an event while column indices correspond to the type to which an event was classified. The number of neighbours used in the classification vote was 56. Errors correspond to 68 % intervals and were obtained via bootstrapping the dataset 100,000 times.

Once the KNN and the NB were optimized, they were stacked with the original BDT as follows: The KNN and NB were optimized to the test set. Each algorithm output a vector of predictions for the test set. These predictions were used as input variables along with the usual parameters of energy, U-wire risetime, number of U-wires, and standoff distance in a BDT that was trained on the test set that was used with the NB and the KNN. That is, we trained a BDT as before, except with 6 total parameters instead of 4, with the two new parameters being the predictions

made by the KNN and the NB – a hybrid of meta-features and measured parameters. The architecture is shown graphically in *Figure 5.14*. This stacked learner was optimized using a second test set (we require a second test set when stacking so as to avoid bias). The learning curves on the two datasets are depicted in *Figure 5.15* and the classification matrix is given in *Figure 5.16*. The learning curves appear to flatten after approximately 6000 trees have been added to the forest, perhaps even beginning to increase towards the end of the curve, and so the final forest was cut at 6000 trees. The classification matrix corresponds to a forest size of 6000 trees. Due to time constraints, this classification matrix does not consider the pairwise kappa pruning process that was outlined in the previous section. Nevertheless, the kappa pruning step remains to be confirmed to produce a statistically significant reduction in error rate before it is included in this classifier. It is also worth noting that the error rate on the test set is briefly lower than that on the training set when the forest contains approximately 100-200 trees (see *Figure 5.15*). While confidence bands analogous to those in *Figure 5.5* were not created for *Figure 5.15* due to time constraints (only the errors at the cut location were calculated, as shown in *Figure 5.15*), this observation is likely just a statistical fluctuation. Indeed, across classification problems, it is not an uncommon observation for a learning curve on a test set to briefly fluctuate below the learning curve of the training set in the early stages of training a classification algorithm. An explanation is the following: A classification algorithm creates boundaries in the parameter space to separate regions that are dominated by different types of events. The boundary itself moves or fluctuates around the optimal location in the parameter space as the classifier learns more about the space and various data points (i.e. events) jump from one side of the boundary to the other. In our case, and many others, the classifier adjusts its boundary more substantially in the early stages of learning (i.e. when much remains to be learned about the space). It is therefore not surprising that random fluctuations in the positioning of the boundary will happen to produce better performance on the test set at these stages. As the classifier learns more about the set on which it is trained, it will improve its performance on this set and it becomes increasingly unlikely for another random set to be classified more accurately than the classifier's training set. Of course, the difference between training and test performance cannot be allowed to grow to be too large as this suggests that

the classifier may have “overfit” to the training set and that its decisions boundaries in the parameter space describe only the points in that set and not necessarily the distribution from which they were drawn.

Figure 5.14: Architecture of the stacked learner. The hidden layer includes a NB learner and a KNN learner. The final layer combines the outputs from these hidden layer algorithms with the parameters risetime, energy, number of triggered U-wires, and standoff distance as inputs for a BDT which makes the final prediction.

Figure 5.15: Learning curves of a stacked BDT taking predictions from a KNN and a NB as input along with the measured energy, U-wire risetime, number of triggered U-wires, and standoff distance.

	$0\nu\beta\beta$	$2\nu\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
$0\nu\beta\beta$	$0.818^{+0.002}_{-0.002}$	$0.0501^{+0.001}_{-0.0009}$	$0.0202^{+0.0006}_{-0.0006}$	$0.064^{+0.001}_{-0.001}$	$0.0360^{+0.0009}_{-0.0008}$	$0.0130^{+0.0005}_{-0.0005}$
$2\nu\beta\beta$	$0.059^{+0.001}_{-0.001}$	$0.799^{+0.002}_{-0.002}$	$0.0002^{+0.0002}_{-0.0001}$	$0.000^{+0.00006}_{-0.00006}$	$0.139^{+0.002}_{-0.002}$	$0.00030^{+0.00008}_{-0.00008}$
^{60}Co	$0.0370^{+0.0009}_{-0.0008}$	$0.005^{+0.0003}_{-0.0003}$	$0.740^{+0.002}_{-0.002}$	$0.094^{+0.001}_{-0.001}$	$0.125^{+0.002}_{-0.001}$	$0.00070^{+0.00008}_{-0.00009}$
^{232}Th	$0.149^{+0.002}_{-0.001}$	$0.044^{+0.0009}_{-0.0009}$	$0.177^{+0.002}_{-0.001}$	$0.492^{+0.002}_{-0.002}$	$0.072^{+0.001}_{-0.002}$	$0.066^{+0.001}_{-0.001}$
^{238}U	$0.101^{+0.001}_{-0.001}$	$0.214^{+0.002}_{-0.002}$	$0.175^{+0.002}_{-0.002}$	$0.0430^{+0.0009}_{-0.0009}$	$0.456^{+0.002}_{-0.002}$	$0.0094^{+0.0004}_{-0.0004}$
^{137}Xe	$0.232^{+0.002}_{-0.002}$	$0.140^{+0.002}_{-0.001}$	$0.0220^{+0.0006}_{-0.0007}$	$0.120^{+0.001}_{-0.001}$	$0.073^{+0.0009}_{-0.0009}$	$0.412^{+0.002}_{-0.002}$

Figure 5.16: Classification matrix of a stacked BDT taking predictions from a KNN and a NB as input along with the measured energy, U-wire risetime, number of triggered U-wires, and standoff distance. The BDT includes 6000 trees. Errors correspond to 68% intervals and were obtained via bootstrapping the dataset 100,000 times.

5.3 Confidence Intervals for General Event-by-Event Classification Algorithms

Creating confidence intervals for the number of counts coming from each event type in the EXO-200 fit model is a fundamentally different concept when considering event-by-event classification instead of the traditional approach. In the latter, events are plotted in the parameter space and a global fit is optimized to the distribution of points in the space. Difficulties concerning optimal bin sizes can occur with this method. This is especially true in regions of the space with low statistics. By contrast, the event-by-event approach will immediately classify a data point as what it believes to be its most likely true event type by determining the optimal binning on a 1D output variables that is a transformation of the combined input variables. In this case, the number of bins is equal to the number of classes (i.e. each bin represents a class or event type) and there is no need to optimize bin size manually. We then determine the most likely true composition of event types based on the binning made by the classification algorithm. In other words, we fit to a distribution of meta-variables – the classifications. Up until now, to my knowledge, a method for creating confidence intervals of this type for classification algorithms did not exist. In what follows, the method will be outlined and the results of its application to EXO-200 analysis will be summarized.

Suppose we want to train a classification algorithm to classify events as one of m classes, where each event is defined by d features. Suppose that the classification algorithm will attempt to break the d -dimensional feature space into regions in which different classes are highly separable. That is, into regions in which most of the events come from a single class. Each of

these regions will be labelled as class_i - like, where the subscript, i , denotes one of the m classes. A region is labelled according to which class is most abundant in that region.

If the above paragraph is true for a particular classification algorithm, then the distribution of events in the training set is important. For example, if the training set contains a far greater number of events from one class than it does from any of the other classes, it will immediately be biased to always predict that an event comes from that class. In the case of EXO-200 analysis, we have assumed no prior knowledge of the ratios of the abundances of events from the various classes. That is, the classification algorithm is trained on an equal number of events from each class. Of course, this doesn't accurately describe what is observed in our detector. For example, double-beta decay events are infrequent in comparison to ^{232}Th events. However, for the analysis developed here, this does not cause problems due to the following interpretation of what the classifier is learning: With any training dataset, to achieve the lowest error in any fixed region of the parameter space, a good classifier will learn to predict the highest frequency class in that region. It will adjust the regions themselves to optimize the overall classification accuracy. With uniform training data, if it is trying to label a region by which class is most abundant in that region, then it is essentially learning which class has the d -dimensional probability distribution function with the largest value in that region. A 1D example is shown in *Figure 5.17* below.

Figure 5.17: Overlapping probability distribution functions in 1D. The red, vertical line denotes the optimal classification cut. All points to the left of the line are labelled as class a and all points to the right of the line are labelled as class b .

In the figure, a vertical separation boundary is drawn at the intersection of the distributions of the different classes. Points to the right would be labelled as class b events while points to the left would be labelled as class a events. In the figure, frequency is plotted on the vertical axis, but if the sum of the frequencies of the classes are equal, then it is essentially the same thing as plotting and separating the PDFs of the two classes.

In practice, however, it may be true that the rate of a given event type is much larger than another. Therefore, although events of one class may have a larger PDF value in a given region than some other class, that other class may have a much higher rate and in effect will be observed in said region more often than the class with a higher same-class probability in that region. That is, the cut on the PDFs would not be the same as the cut on the frequencies. Thus, in practice you will achieve a lower classification error if you always predict the second class than you will if you always predict the first class. This is the complete opposite result of the truth! A remedy to this problem is to weight each region by the prior probability of observing each class. However, these weights are often unknown. This is especially true when attempting to observe never before observed signals as we are doing at EXO.

There is another way of dealing with the event rate problem that involves estimating the most probable true count of each class given a particular set of classified events of unknown identity. Suppose we observe a set of N events. A classification algorithm is used to classify each of these events as belonging to a subset of one of m classes (in our case a class corresponds to one of $0\nu\beta\beta$, $2\nu\beta\beta$, ^{60}Co , ^{232}Th , ^{238}U , ^{137}Xe). In this sense we are transforming an observation matrix of events with a hash function. Let T denote the classification algorithm, then,

$$\begin{aligned} T: D \in \mathbb{R}^{N \times d} &\rightarrow \vec{h} \in \mathbb{Z}^N \mid h_i \in \{1, 2, \dots, m\} \\ T: \vec{v} \in \mathbb{R}^d &\rightarrow h \in \mathbb{Z} \mid h_i \in \{1, 2, \dots, m\} \end{aligned} \tag{5.3.1}$$

where the classification algorithm is represented by a transformation, T , that transforms a $N \times d$ - dimensional matrix, D , into an N -dimensional row vector, $\vec{h}$, referred to as the *classification vector*. Each entry in $\vec{h}$ is some integer on $[1, m]$ corresponding to a class, where m is the total number of classes. The transformation is depicted graphically in *Figure 5.18*.

Figure 5.18: A classification algorithm acts as a transformation that takes a physics event parameterized by d parameters to an integer that represents an event type. The transformation is illustrated here as a hashing transformation.

Each time an event is classified, it is either correctly classified or it is misclassified as a member of one of the other $m-1$ classes. We interpret this to mean that members of a *fixed* class are classified according to an m -dimensional multinomial distribution. We can define a classification matrix,

$$P = \begin{pmatrix} p_{11} & \cdots & p_{1m} \\ \vdots & \ddots & \vdots \\ p_{m1} & \cdots & p_{mm} \end{pmatrix} \quad 5.3.2$$

, where entries common to a row correspond to the m classification probabilities of a given multinomial distribution (i.e. there are m multinomial distributions – one for each class – and each is defined by one of the m rows). Let's adopt the notation that the first subscript, i , of any p_{ij} corresponds to the true class while the second subscript, j , corresponds to the class to which an event is classified.

In practice, the entries of P are not known exactly because we don't know the closed forms of the PDFs. Instead, each entry is estimated by testing on a cross validation set after training with the estimates given by the MLEs of the probabilities of a multinomial distribution,

$$\hat{p}_{ij} = \frac{n_{ij}}{N_i} \quad 5.3.3$$

, where n_{ij} is the number of events of true types i classified as event type j , and N_i is the total number of events of true type i . Each entry in P has a width due to testing on a dataset of finite size. In principal, the width should become increasingly narrow as the cross validation set size approaches infinity (if we tested on every possible point we would know the classification probabilities exactly). It is straightforward to point out that this is the case using the properties

of a binomial distribution (this is sufficient to show the order on which variance gets smaller in the multinomial case). It is known that the Maximum Likelihood Estimate (MLE) for the probability of success, $\hat{p}$, in a binomial distribution is given by,

$$\hat{p} = \frac{n}{N} \quad 5.3.4$$

where n is the number successes (correctly classified events) given N events. The variance of the estimate is then given by,

$$\begin{aligned} V(\hat{p}) &= E(\hat{p}^2) - E(\hat{p})^2 \\ &= E\left(\frac{n^2}{N^2}\right) - E\left(\frac{n}{N}\right)^2 \\ &= \frac{1}{N^2} V(n) \\ &= \frac{1}{N^2} Np(1-p) \\ &\approx \frac{1}{N} \hat{p}(1-\hat{p}) \\ &= n \left(\frac{N-n}{N^3} \right), \quad \hat{p} = \frac{n}{N} \end{aligned} \quad 5.3.5$$

, where the second to last manipulation uses the estimates for the probability of success in place of the true probability. Clearly, then, as the size of the cross validation set increases, the standard deviation is proportional to $n^{1/2}/N$.

Aside:

Emphasis should be placed on the fact that the p_{ij} are determined by the performance on the cross validation set, not by performance on the training set. Thus, as the cross validation set size increases, the number of trees in, say, an AdaBoost algorithm doesn't need to increase in order to achieve the same performance (in general, if the algorithm is trained on more data, it will need to include more trees to reach its optimal performance). This allows us to compare widths consistently between cross validation set sizes as the underlying algorithm need not be retrained from dataset to dataset. Therefore, if we want to have narrower widths in P , we need only test on larger datasets and update P according to the results. Thus, if it is believed that the

classification algorithm has reached Bayes' limiting error and that it has used enough data so as not to overfit, it makes sense to use any new data in the test set as opposed to the training set.

Once a classifier has hashed/classified a set of observed physics events, we can take the vector of classified events and create a one-dimensional histogram where a bin represents an event type (depicted graphically in *Figure 5.19*). We refer to this as the *Observed Histogram*.

Figure 5.19: The observation vector can be binned into a 1D histogram referred to as the *Observed Histogram*.

The histogram can be represented as an m -dimensional vector, $\vec{n}_{obs}$, referred to as the *Observation Vector*,

$$\vec{n}_{obs} = (n_{obs,1}, \dots, n_{obs,m})^T \quad 5.3.6$$

where $n_{obs,i}$ denotes the number of events that were classified as a member of the i^{th} class. Each $n_{obs,i}$ will be a sum of correctly classified events of class i , and misclassified events of type $j \neq i$. Mathematically,

$$n_{obs,i} = \sum_{j=1}^m n_{ji} \quad 5.3.7$$

where n_{ji} is the number of events of true class, j , that are classified as events of class, i , in a particular set of events or experiment. This is unknown but the expectation value follows the multinomial prescription. Taking the expectation value,

$$E(n_{obs,i}) = \sum_{j=1}^m E(n_{true,j,i}) \quad 5.3.8$$

$$E(n_{obs,i}) = \sum_{j=1}^m p_{ji} n_{true,j} \quad 5.3.9$$

which implies the following system of linear equations relating observations to true values,

$$E(\vec{n}_{obs}) = P^T \vec{n}_{true} \quad 5.3.10$$

where $n_{true,j}$ denotes the number of events that are truly events from class j . These can be considered to be the entries of the *true histogram* and can be grouped into a vector of their own, the *true vector*,

$$\vec{n}_{true} = (n_{true,1}, \dots, n_{true,m})^T \quad 5.3.11$$

In practice, we are unaware of the true class of any given event. Thus, (5.3.6) can be constructed by applying the classifier, T , to an observed dataset, D , and subsequently binning the result, $\vec{h}$, while (5.3.11) must be estimated. In what follows, (5.3.11) is estimated by the method of binned maximum likelihood and the use of (5.3.9).

Once a set of observations, D , has been hashed and subsequently binned, and if the total number of observations, N , is taken to be fixed, then we can assume that the number of events in each bin follows a multinomial distribution. In such a case, the joint probability of a particular binning is given by,

$$f_{joint}(\vec{n}_{obs}; \vec{n}_{true}) = N! \prod_{i=1}^m \frac{1}{n_{obs,i}!} \left(\frac{v_{obs,i}}{n_{obs,i}} \right)^{n_{obs,i}} \quad 5.3.12$$

where $v_{obs,i}$ is the expected number of events in the i^{th} bin (which represents the i^{th} class).

Taking the logarithm produces the log-likelihood function,

$$\log L(\vec{\theta}) = \sum_{i=1}^m n_{obs,i} \log v_{obs,i}(\vec{\theta}) \quad 5.3.13$$

where $\vec{\theta}$ denotes a set of parameters that determines $v_{obs,i}$. The objective becomes to maximize the log-likelihood by varying the parameters, $\vec{\theta}$, to determine the MLEs, $\hat{\theta}$. Looking at (5.3.7) or (5.3.8), the $v_{obs,i}(\vec{\theta})$ can be expanded and the log-likelihood becomes,

$$\begin{aligned} \log L(\vec{\theta}) &= \sum_{i=1}^N n_{obs,i} \log(p_{1i} n_{true,1} + \dots + p_{mi} n_{true,m}) \\ &= \sum_{i=1}^N n_{obs,i} \log \left(\sum_{j=1}^m p_{ji} n_{true,j} \right) \end{aligned} \quad 5.3.14$$

where we see that $\vec{\theta} = \vec{n}_{true}$. Thus, by maximizing (5.3.12) with respect to $\vec{\theta}$, we determine the MLE of $\vec{n}_{true}$.

The log-likelihood must be maximized while taking into account 2 physical constraints. The first constraint is conservation of total number of events,

$$N = \sum_{i=1}^m \hat{n}_{true,i} \quad 5.3.15$$

and the second is that the estimated number of events of any class must be between zero and the total number of events,

$$0 \leq \hat{n}_{true,i} \leq N \quad 5.3.16$$

In this work, the optimization procedure is handled through *Python's scipy.optimize* library with its built-in *minimize* function [94]. We then look to building confidence intervals for the MLEs of the true number of events coming from each class. Since we now have a likelihood function, there are several methods for approaching this. In the next section we employ the bootstrap approach to build confidence intervals for EXO-200 Monte Carlo data.

5.4 Confidence Intervals for Classification Algorithms at EXO-200

In this section we summarize the application of the formalism from the previous section to the BDT trained and optimized on EXO-200 MC data. For the bootstrap step, the bootstrap percentile method was employed [73][74]. This method constructs, for example, a 68% confidence interval by choosing the central 68% of the re-sampled test statistics. The true coverage converges to nominal coverage as $O\left(\frac{1}{\sqrt{n}}\right)$, where n is the number of re-sampled test statistics, given that the following assumptions about the dataset hold: (1) The statistic is unbiased, (2) the statistic is homoscedastic, (3) The bootstrap statistic can be transformed into a standard normal distribution, (4) the untransformed and de-transformed distributions of the test statistic are the same, (5) the standard error of the bootstrap statistic and the sample statistics are the same.

Combined with the MLE formalism of the previous section, the bootstrap approach was used to construct 68% confidence intervals for the MLEs of three unique blind datasets. The classification algorithm that was applied to the datasets was the original BDT before applying kappa pruning or stacking with the KNN and NB. The results are summarized in the tables below. We refer to the blind datasets of *Table 5.12*, *Table 5.13*, *Table 5.14* as *Test Set 0*, *Test Set A*, *Test Set B*, respectively.

Event	True	Lower Limit	MLE	Upper Limit	Observed
$0\nu\beta\beta$	12	0.0	8.8	23.7	111
$2\nu\beta\beta$	329	295.5	320.0	346.2	349
^{60}Co	21	1.4	20.9	41.8	122
^{232}Th	269	238.5	266.9	291.2	154
^{238}U	339	305.5	346.6	380.9	226
^{137}Xe	30	20.6	36.7	52.2	38

Table 5.12: An example of bootstrap confidence intervals with nominal coverage of 68% applied to a dataset where the true event count of each event type (*True* column) is non-zero. We see that for this particular event distribution, we cannot distinguish 12 $0\nu\beta\beta$ from zero. The *Observed* column denotes the number of events from the dataset that were classified as a given event type by the BDT algorithm before applying the MLE correction.

Class	True	Lower Limit	MLE	Upper Limit	Observed
$0\nu\beta\beta$	0	0.0	0.0	15.2	0
$2\nu\beta\beta$	332	304.4	302.1	340.7	335
^{60}Co	21	8.3	13.9	23.1	46
^{232}Th	272	248.3	260.7	281.6	224
^{238}U	342	332.9	377.8	370.0	313
^{137}Xe	33	33.9	45.5	57.0	82

Table 5.13: An example of bootstrap confidence intervals with nominal coverage of 68% applied to a dataset where the true event count of $0\nu\beta\beta$ is zero. Coverage studies show that nominal and true coverage may not be equivalent for rare events (i.e. $0\nu\beta\beta$ events in this case).

Class	True	Lower Limit	MLE	Upper Limit	Observed
$0\nu\beta\beta$	0	0.0	0.0	15.2	2
$2\nu\beta\beta$	0	0.0	0.0	19.7	78
^{60}Co	1	0.0	4.8	27.2	34
^{232}Th	333	288.6	322.6	348.1	291
^{238}U	333	294.7	336.3	351.8	304
^{137}Xe	333	304.1	336.2	360.2	291

Table 5.14: An example of bootstrap confidence intervals with nominal coverage of 68% applied to a dataset where the true event count of multiple event types is zero.

The true coverage of various bootstrapping techniques converges to nominal coverage on different orders and under different assumptions, so coverage studies are customary in the field. In *Table 5.15* we show the true coverage of three different bootstrap methods given a nominal coverage of 68%. We tested 438 unique datasets respecting the same event distribution as *Test Set 0*. We resampled each set 10,000 times – that is we computed 10,000 re-sample MLEs for each distribution. For the percentile method we’re approaching the true coverage nicely after 10,000 resamples for most event types. Some types may be showing overcoverage but this, at least, is preferred to undercoverage. Also, we can easily resample much more than 10,000 times which will further reduce coverage error. Given our resources, a coverage study using more than 10,000 samples would have taken on the order of two weeks to complete for 438 datasets, and so we settled for 10,000. Now that the percentile method appears to be promising, future studies using 100,000 resamples should be conducted to look for improved convergence.

Event	0v $\beta\beta$	2v $\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
Basic Bootstrap	0.506	0.723	0.486	0.687	0.668	0.666
Percentile Bootstrap	0.683	0.737	0.678	0.748	0.698	0.664
Normal Bootstrap	0.566	0.728	0.584	0.730	0.687	0.639

Table 5.15: True coverage of various bootstrap confidence interval techniques applied to the event distribution of *Test set 0*. These numbers were obtained by creating confidence intervals for 438 unique datasets following the event distribution of *Test Set 0*. Each set used 10000 resamples during the bootstrap procedure.

A similar coverage study was performed with datasets following the event distributions of *Test Set A* to check behaviour of our intervals near boundaries (i.e. 0 and 1000). Since the percentile method performed best on *Test Set 0*, the other two methods were not considered in this second coverage study.

Event	0v $\beta\beta$	2v $\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
Coverage	0.522	0.752	0.694	0.760	0.706	0.716

Table 5.16: True coverage, using percentile bootstrap, for nominal coverage of 68% with datasets respecting the event distribution of *Test Set A*.

As can be seen for 0v $\beta\beta$ in *Table 5.16* we deviate from nominal coverage drastically for an event count that approaches zero. However, in the case that the bootstrap statistics are MLEs, this is remedied in a straightforward manner. It has been shown elsewhere [75] that the MLE for the mean of a normal distribution, with a constraint of non-negativity for an estimate of the mean (i.e. $y^* = \max(y, 0)$, where y is the estimate and y^* is the estimate after imposing the non-negativity constraint), the exact coverage of a percentile bootstrap confidence interval is given by,

$$p(\theta_0 \in [q_{\alpha_1}^*, q_{\alpha_2}^*]) = \begin{cases} 1 - \alpha_1 - \alpha_2 & ; \frac{1}{n^2} \theta_0 > \Phi^{-1}(1 - \alpha_2) \\ 1 - \alpha_1 & ; \frac{1}{n^2} \theta_0 \leq \Phi^{-1}(1 - \alpha_2) \end{cases} \quad 5.4.1$$

where θ_0 is the true value of the mean, α_1 & α_2 , are the quantiles of the tails of the confidence interval, $q_{\alpha_1}^*$ & $q_{\alpha_2}^*$ are the values of the estimate corresponding to those quantiles, n is the number of samples in the distribution, and Φ is the cumulative distribution function of a standard normal random variable.

The distribution of a MLE approaches a normal distribution in the large sample limit. In the bootstrap procedure we can impose the large sample limit by resampling until we have a large number of sample statistics. Thus, the above equation can be used to calculate our true coverage given a chosen nominal coverage for datasets with very low event counts. For example, in *Table 5.16* the second line of (5.4.1) certainly holds for $0\nu\beta\beta$ events as the true count is zero (just substitute zero for θ_0). If we add back 0.16 to the true coverage of $0\nu\beta\beta$ in *Table 5.16* we see that we regain the nominal coverage. Thus, for rare event searches, we can take a conservative approach and assume that a 68% nominal coverage will give a true coverage of 52%. What's more, we can actually calculate the value of the true parameter for which the change in coverage occurs. For example, if θ_0 is taken to represent the true count $0\nu\beta\beta$ events, then a true count of zero means the true coverage for a nominal bootstrap coverage of $1 - \alpha_1 - \alpha_2$ will always be $1 - \alpha_2$. For a 68% interval, the second line will apply if $\frac{1}{n^2}\theta_0 < \approx 1$. Therefore, since n is very large in practice, the first line will always apply unless $\theta_0 = 0$ (i.e. unless the event count of $0\nu\beta\beta$ is truly zero).

Another observation from *Table 5.16* worth discussing is that all event types apart from $0\nu\beta\beta$ are overcovered. This is probably due to the constraint which limits the minimum event count of any event to be no less than zero. Because of this constraint, we must correspondingly add counts to the count distribution of another event type when we cut off $0\nu\beta\beta$'s count distribution at zero. It is thus interesting to note that the most overcovered (undercovered) events are also those event types as which $0\nu\beta\beta$ events are most (least) often misclassified. At EXO, since we are mostly interested in the $0\nu\beta\beta$ count, this systematic overcoverage in other event types is perhaps not so problematic. However, if many event counts fall close to the boundaries, problems may arise.

Finally, for completeness, a Monte Carlo study was done to determine the mean and median bias of the MLEs. This was done by creating 10,000 unique datasets defined by the event distribution of *Test Set 0*, computing the MLEs for each set, and then taking the average(median) of these MLEs to determine if the mean(median) approaches the true event distribution in the large sample limit. This was then repeated for datasets following the event

distributions of *Test Set 0*, *Test Set A*, and finally *Test Set B*. The results of the bias calculations are summarized in *Table 5.17*. The medians tend to have a smaller bias than the means. In other words, the MLE is overestimating and underestimating at nearly equal rates. That the median performs better suggests that the distribution of the estimate is skewed, which is not unexpected considering the boundary constraints on the estimates.

Value	$0\nu\beta\beta$	$2\nu\beta\beta$	^{60}Co	^{232}Th	^{238}U	^{137}Xe
Test Set 0 MLE - mean	13.7	329.6	22.8	267.4	336.4	30.0
Test Set 0 MLE - median	12.1	329.6	21.3	267.6	336.9	29.6
Test Set 0 - True	12	329	21	269	339	30
Test Set A MLE - mean	5.8	332.8	23.7	268.4	337.0	32.3
Test Set A MLE - median	0.0	332.7	22.4	268.5	337.5	32.0
Test Set A - True	0	332	21	272	342	33
Test Set B MLE - mean	8.1	8.9	10.2	327.4	317.3	328.0
Test Set B MLE - median	1.5	2.7	4.3	327.5	318.4	328.2
Test Set B - True	0	0	1	333	333	333

Table 5.17: Mean and median biases of the binned MLE method for three unique event distributions. The median bias tends to be smaller than the mean bias for all datasets.

In general, the bootstrap approach may or may not work for a given dataset. The coverage studies summarized in this section have shown that the percentile method is quite applicable to EXO-like data. However, there are issues near the zero boundary and, while this can be remedied with (5.4.1), in practice we don't know the true count and therefore we don't know which expression in the equation holds. By always adopting the conservative approach, we lose power in our analysis. Nevertheless, there now exists a likelihood function for obtaining the MLEs of the number of events of each type – something that was not achievable prior to this work. Thus, we can look to estimation approaches other than bootstrapping to construct our intervals. As of January 2019, the EXO-Deep team is working on applying Wilks' Theorem, using the same likelihood method formalism that was developed earlier in the chapter. Wilks' Theorem essentially replaces the bootstrap approach. Because of the ordering principle in

Wilks' Theorem, we do not have to worry about coverage error. However, we lose the benefit of the speed at which a bootstrap study can be completed.

5.5 Discussion Regarding the Choice of Metric for Classifiers

In section 5.2 it was stated that the ultimate metric for the performance of an analysis technique in a rare event search is a sensitivity study. However, we saw that such a study was not the objective of this work. Rather, the objective was to develop a formalism to be able to do this kind of study at all and to provide a proof of concept for the method. The method was developed in section 5.3 and the proof of concept was illustrated via coverage studies of confidence intervals in which we saw that nominal coverage could be obtained (apart from difficult cases near estimate boundaries). We thus adopted a different metric of performance, namely the average classification accuracy across all event types in the fit model.

In section 5.3 it is shown that, for a dataset of unknown event distribution, the sum of the estimates of the true number of events originating from each of the six classes must sum to the total number of events (i.e. conservation of total number of events is imposed). Therefore, if one event type is particularly poorly understood by the classifier (i.e. has a low correct classification rate – corresponding to a diagonal entry for this event type in the classification matrix that is close to random guessing), then we can interpret this to mean that we have little information about this class and that the estimate for the true number of events from this class will fluctuate widely. To conserve the total number of events in the estimates, the other estimates must jointly fluctuate in the opposite direction. This is equivalent to saying that, if one diagonal element of the classification matrix is low, the off-diagonal elements of the same row must correspondingly increase. Therefore, there is reason to believe that optimizing based on the average accuracy across event types will produce narrow confidence intervals across all events – ultimately this improves sensitivity across event types. On the other hand, EXO-200 analysis is particularly concerned with $0\nu\beta\beta$ events. Therefore, we intuitively do not want events to be frequently misclassified as $0\nu\beta\beta$, as this contaminates the $0\nu\beta\beta$ bin and may produce poor estimates for the true $0\nu\beta\beta$ counts in the statistical inference step. Furthermore it is not necessarily the case that a classifier will frequently misclassify events of a poorly

understood event type as $0\nu\beta\beta$ events. Rather, it may classify these misclassified events as one of the other remaining four event types. In such a case, we are not particularly concerned with the poor accuracy in this class. That is, even if off-diagonal elements of a row in the classification matrix increase due to poor performance on the diagonal entry, it is not necessarily true that the entry in the column corresponding to $0\nu\beta\beta$ will increase.

If we define the purity of the $0\nu\beta\beta$ column in the classification matrix to be equal to the diagonal entry in the $0\nu\beta\beta$ column (i.e. the correct classification probability of $0\nu\beta\beta$) divided by the sum of all of the entries in the column, then we can reason that a high purity would produce narrow confidence intervals and, by extension, improved sensitivity. We do this by recalling that, for fixed decision boundaries in our parameter space (i.e. a trained, unchanging classifier), we can interpret rows of the classification matrix as vectors of probabilities that define multinomial distributions for each of the true event types in our fit model. Combining this with (5.3.7) allows us to conclude that the number of counts in the $0\nu\beta\beta$ bin will be the sum of six multinomial random variables (one for each event type) given a fixed number of events originating from each true event type in an unknown dataset. Since the rates at which each event type interacts in our detector are independent of each other, we can sum the variances of each of the contributing multinomial random variables to determine the variance of the counts in the $0\nu\beta\beta$ bin for a fixed true distribution of event types in an unknown dataset. Thus, for a fixed true distribution in the unknown set, the variance of $n_{obs,1}$, as defined in (5.3.7), is given by

$$V(n_{obs,i}) = \sum_{j=1}^m N_j p_{j1}(1 - p_{j1}) \quad 5.5.1$$

where $i = 1$ corresponds to the index for $0\nu\beta\beta$, and all variables carry forward their definitions from section 5.3. From (5.5.5) it is quite clear that the variance for the number of events classified as $0\nu\beta\beta$ is reduced for an improved purity given a fixed correct classification rate for $0\nu\beta\beta$ (i.e. given a fixed p_{11}) and a fixed true distribution of event types in the unknown set. Therefore, the purity metric may be a superior metric in a rare event search for $0\nu\beta\beta$ as it will produce narrow narrower distributions in the $0\nu\beta\beta$ bin of classified events.

However, this is not the full story as we apply the MLE in (5.3.14) to estimate the true number of $0\nu\beta\beta$ events and this does not depend only on the counts in the $0\nu\beta\beta$ bin, but of the counts in the bins for all event types. Therefore, if we sacrifice the correct classification of other events too much to ensure favourable $0\nu\beta\beta$ purity, the likelihood function may become unstable due to the algorithms poor performance (i.e. large variances) in the bins of all of the other event types. Ultimately, this could hinder sensitivity.

We therefore conclude that there is likely a trade-off between good purity results and good average correct classification results when it comes to moving on to a sensitivity study. Supposing that we choose the learning metric of our classifier to be the purity of the $0\nu\beta\beta$ bin, we run into difficulties related to the loss functions of the learning algorithms. Indeed, the loss functions of these algorithms are built to iteratively search for optimal classification accuracy of all training events. Therefore, if our desired metric is the purity of a particular bin, we must somehow increase the weight of this bin to, in a sense, distort the loss function. However, in a search for $0\nu\beta\beta$, we do not know the rate of occurrence of this event type in our detector in comparison to the rates of other event types. In fact, $0\nu\beta\beta$ may not exist at all. It thus becomes a difficult problem to determine how to weight the data to favour purity of the $0\nu\beta\beta$ bin. Without knowledge of the event rates, we must choose the weight of event types arbitrarily (e.g. is a $0\nu\beta\beta$ event twice as important as a background event? Three times? etc... the answer is not obvious).

At this point, it should be noted that, while the average accuracy was used as the metric for constructing the classifiers of section 5.2, the purity in the stacked classifier is superior to the purity in the forest corresponding to *Figure 5.6*. This is in addition to the stacked classifier also having a superior average accuracy. This can partly be attributed to the fact that purity tends to increase as the correct classification (the numerator of the purity calculation is increasing). However, there is a point at which purity is hindered by increasing correct classification as the distributions of event types can overlap significantly in certain regions of our parameter space. If, for example, we insist on increasing the correct classification probability of true $0\nu\beta\beta$ events in a region that is dominated by another event type, then, by definition, we must accept an increase in the probability that the event type which dominates this region will be misclassified

as $0\nu\beta\beta$. This is just an example of the trade-off between making errors of the first or second kind when making cuts in a parameter space for particle selection. The concept is explained further elsewhere [95].

Optimizing the trade-off between average accuracy and purity of the $0\nu\beta\beta$ bin therefore remains an interesting research avenue and we expect that the best answer to the trade-off between purity and average accuracy can only be provided by in depth sensitivity studies. Again, such studies were not the objective of this work but are being pursued by the EXO-Deep group and a small group at Carleton University in particular. Until such studies are completed, one alternative that makes use of simple learning curves for incorporating purity considerations into the learning procedure is as follows: There is a large flat region in many of the learning curves of section 5.2. While the average accuracy across classes is unchanged in this flat region, it is possible that the purity of the $0\nu\beta\beta$ bin varies and is optimal at particular location along the region. We could thus create a second learning curve within this flat region in which the purity of the $0\nu\beta\beta$ bin is plotted against the number of trees in the flat region. Optimizing the purity for a fixed average accuracy across classes would, presumably, improve the sensitivity versus the case without purity optimization. We have calculated the purity of $0\nu\beta\beta$ (again, where purity is defined as the correct classification probability of $0\nu\beta\beta$ divided by the sum of all of the entries in the $0\nu\beta\beta$ column of the classification matrix) at various stages of the learning process for the forest summarized by *Figure 5.6*. The evolution of the purity is depicted in the figure below as a learning curve.

Figure 5.20: *Left:* Purity of $0\nu\beta\beta$ column in the classification matrix corresponding to *Figure 5.5* as a function of the number of trees in the forest. *Right:* Zoomed view of the left plot corresponding to the flat region of the learning curve in *Figure 5.5*.

We see that the purity tends to increase with the size of the forest. On the right of the figure is a zoomed view of the curve from forest sizes of approximately 6000 to 13000 trees. Upon considering *Figure 5.5* we see that this region corresponds to the flat region of the learning curve based on average accuracy across classes. It also appears that the purity of $0v\beta\beta$ in this region could be decreasing as more trees are added. A statistical analysis remains to be completed to determine if the purity curve is increasing significantly in this region. If it is a significant increase, it would be interesting to consider cutting somewhere between 6000 and 8000 trees, where the purity seems to be lowest, as opposed to at 8000 where the forest of *Figure 5.5* was cut. In this approach we would keep the benefit of optimal overall accuracy (since it is unchanged in this region) while adding the benefit of better $0v\beta\beta$ purity. This strategy should be considered in future studies that incorporate sensitivity studies.

Chapter 6

Barium-Tagging Simulations

The third chapter of this thesis ended with a discussion of background limitations of nEXO. This chapter summarizes advancements that have been made for reducing these limitations through a measurement called Barium-tagging. If such a measurement is successfully made in nEXO, 100% of non- $\beta\beta$ backgrounds can be rejected. As will be seen in the first section of this chapter, this will greatly increase the sensitivity of nEXO, allowing deeper penetration into the parameter space of the normal mass hierarchy.

6.1 – Sensitivity with Barium-Tagging for nEXO

In EXO-200 and nEXO, background processes near the $0\nu\beta\beta$ Q-value in ^{136}Xe obscure the $0\nu\beta\beta$ signal. In *Chapter 3* we saw the degree to which backgrounds can decrease sensitivity and discovery potential. We saw that experiments that have excellent energy resolution, such as bolometer and Ge-diode experiments, are able to reject almost all backgrounds based solely on energy cuts. We saw that this increases the maximum attainable sensitivity for these experiments (see *Figure 3.16*).

Because nEXO in its current design will not achieve the excellent energy resolutions of some other, solid-state, $0\nu\beta\beta$ experiments (although it will do far better than, for example, liquid scintillator experiments), we look elsewhere to discriminate backgrounds near the Q-value. We have already discussed the possibility of increasing voltage to improve resolution of measured event parameters as well as the possibility of employing complex learning algorithms to better discriminate within the parameter space of events. Other methods that have not yet been discussed are, for example, operating deeper underground to reduce cosmic backgrounds, using low-activity materials for detector construction and using a larger LXe mass for self-shielding purposes (suppresses Thorium, Uranium, Radon backgrounds for example). Nevertheless, none of these methods can completely eliminate backgrounds in the ROI over a period of several years. Since nEXO will, ideally, be taking data for 10 years, this problem

persists. In this chapter we discuss advancements in an approach, called *Barium-tagging*, for addressing this problem.

Barium-tagging has been proposed as a method of rejecting all non- $\beta\beta$ backgrounds (i.e. ^{60}Co , ^{232}Th , ^{238}U , ^{137}Xe). From the description of $0\nu\beta\beta$ and $2\nu\beta\beta$ given in Chapter 2, we see that the daughter nucleus of the two decays is ^{136}Ba ,

6.1.1

where $2\bar{\nu}_e$ is present in $2\nu\beta\beta$ and not in $0\nu\beta\beta$.

The Barium-Tagging method, originally proposed by Moe [76], proposes the extraction of the Barium daughter from the TPC followed by its identification through laser induced fluorescence (LIF).

The following graph demonstrates the increase in sensitivity and discovery potential that is obtained by employing Ba-tagging. This dramatic improvement in mass sensitivity is the motivation for employing Ba-tagging. The calculation leading to the plot can be found elsewhere [22]. The improvement in mass sensitivity increase occurs because Ba-tagging permits complete rejection of non- $\beta\beta$ background processes near the Q-value of interest. Indeed, for a background that scales linearly with Nt , where N is the number of ^{136}Xe atoms in the detector and t is the running time of the experiment, the mass sensitivity scales like $(\sqrt[4]{Nt})^{-1}$. With complete rejection of such backgrounds, sensitivity would scale like $(\sqrt{Nt})^{-1}$ [76]. Consequently, for the case where nEXO collects data for five years, the inclusion of Ba-tagging allows us to claim a discovery anywhere in the inverted hierarchy and also allows us to probe parts of the normal hierarchy.

Figure 6.1: Sensitivity of nEXO with and without Ba-tagging, compared to final EXO-200 sensitivity. Sensitivity is computed with a 90% confidence interval [22]. The red ribbon denotes the space occupied by the inverted hierarchy, and the blue ribbon denotes the normal hierarchy.

The proposed LIF concept makes use of the convenient electronic ground state of the singly ionized Barium daughter from (6.1.1). Singly ionized barium has a single valence electron in its $6S_{1/2}$ orbital and can thus be treated mathematically as a single electron atom. It is found that barium has a strong allowed transition between the $6^2S_{1/2}$ and $6^2P_{1/2}$ states with a characteristic wavelength of $\approx 493\text{nm}$. A transition of less prominence occurs between the $6^2S_{1/2}$ and meta-stable $5^2D_{3/2}$ states with a characteristic wavelength of $\approx 650\text{nm}$ (see *Figure 6.2 & Table 6.1*) [77].

Figure 6.2: The single electron system of $^{136}_{56}\text{Ba}^+$ with branching ratios of the emission/absorption spectrum.

Transition	Wavelength (nm)	Transition Rate (MHz)	Branching Ratio
$6^2P_{1/2} \rightarrow 6^2S_{1/2}$	493.54544	92.1 ± 0.9	0.7293 ± 0.0002
$6^2P_{1/2} \rightarrow 5^2D_{3/2}$	649.86934	34.2 ± 0.3	0.2707 ± 0.0002

Table 6.1: Properties of transitions of interest in $^{136}_{56}\text{Ba}^+$ tagging system. The wavelengths shown correspond to vacuum conditions [21][77] .

It is clear from the branching ratio column of *Table 6.1* that the $6^2P_{1/2}$ will not always undergo spontaneous decay into the $6^2S_{1/2}$ ground state. Rather, there is a large branching ratio corresponding to the decay into the $5^2D_{3/2}$ state. These two decay paths motivate a two-laser system for laser induced fluorescence of $^{136}_{56}\text{Ba}^+$. A first laser tuned to the $6^2S_{1/2} \rightarrow 6^2P_{1/2}$ transition wavelength is employed to pump ground state valence electrons into the $6^2P_{1/2}$ excited state. Fluorescence of the same wavelength is then observed when the excited state decays back to the ground state. However, because the $6^2P_{1/2}$ state also decays to the meta-stable $5^2D_{3/2}$ state in $\approx 27\%$ of its decays, if this were the only irradiating wavelength, electrons would eventually be lost to the $5^2D_{3/2}$ state which would hinder the potential for identification of singly ionized barium atoms. The much longer lifetime of the meta-stable state increases the possibility of losing a trapped ion before it is detected. To avoid this, a second laser of wavelength corresponding to the $5^2D_{3/2} \rightarrow 6^2P_{1/2}$ transition is introduced in order to pump electrons back into the excited P-state from which they may again decay to the ground state.

Work has been done at Carleton University in which single $^{136}_{56}\text{Ba}^+$ fluorescence has been observed in an ion trap [21][65], demonstrating that Ba-tagging of the Barium daughter in $0\nu\beta\beta$ experiments is feasible. The remaining difficulty, in the case of EXO, is to extract the Barium ion from the TPC itself and deposit it in an ion trap. The results summarized in this chapter address this problem via two unique extraction approaches.

It should be noted at this point that the Barium daughter in $0\nu\beta\beta$ of ^{136}Xe will be singly ionized in LXe and doubly ionized in GXe, while the proposed laser tagging system requires a single valence electron. For the GXe case Ba^{++} must be reduced. This is discussed briefly in section 6.3.1.

6.2 The Electrostatic Probe

One of the earliest ideas in the Ba-Tagging effort was to construct a probe that could be inserted into a LXe TPC and that could attract and retain a Ba^+ atom to its tip. The probe would then be extracted and transported to an ion trap where the Ba^+ particle would be separated from the probe and trapped for laser induced fluorescence with the two-laser system described above. Until recently, all proposed schemes in which a LXe was the source medium (as opposed to GXe) were to extract the Ba^+ ion using an electrostatic grabbing probe that could be inserted and extracted from the LXe volume as needed. At least four schemes have been pursued – the cold probe, the cold probe with an optical fibre for in-situ tagging, a hot probe, and the Resonant Ionization Spectroscopy (RIS) probe [78-81]. The cold probe freezes a layer of Xenon to its surface and then attracts a Ba^+ ion to its tip. The Ba^+ ion is captured in the solid Xenon layer and is subsequently released from the probe into an ion trap by thawing the Xenon. The cold probe with a built-in optical fibre for in-situ tagging is similar but does not require the ion to be released into an ion trap. Rather, the ion is tagged on the tip by a fibre laser built into the core of the probe. The hot probe requires no freezing of Xenon. Instead, the Ba^+ is attracted to the surface of the probe itself. It is subsequently ejected from the probe into an ion trap by heating the probe to $\sim 1000\text{K}$. In the the final method, the RIS probe, the Ba^+ ion is attracted to the surface of a semiconductor. The ion is released from the semiconductor into the ion trap with an ablation laser and RIS lasers to ionize only the Barium atom.

Each of these methods requires electrostatic attraction of a Ba^+ ion to the probe's tip. Prior to this work, the hot probe and the cold probe methods were operated in the laboratory using clusters of ions [80][81]. However, there has been no work done to understand the attraction between single ions and the electrostatic probe (*electroprobe*). Of course, understanding the dynamics of single ion capture is important for a collaboration like EXO where we are looking for the single daughter of a double-beta decay. In this section, the results of simulations that consider the question of single ion attraction to an electrostatic probe are summarized. The particular method by which the ion will be separated from the probe is left for future work. That is, the results summarized here are useful in the attraction step of all electroprobe

approaches. Which method is superior will thus depend on the degree of success that is achieved in the separation step of each method.

6.2.1 The Barium-Tagging Line with the Electrostatic Probe

The design of a Ba-tagging line that employs any of the electroprobe methods can be broken down into at least five detection steps that can be optimized somewhat separately – the TPC, the electroprobe extraction step, separation of an ion from the electroprobe, the ion trap, and the two-laser LIF system. In *Figure 6.3* the tagging line is shown chronologically with respect to the order in which detector components will be used to tag a hypothetical Ba^+ ion. First, the TPC will measure the parameters of an event to determine whether it is or is not within the ROI. This analysis must also provide an accurate measurement of the event location. This location is used in the extraction/attraction step where, if an event falls within the ROI, the electroprobe is inserted into the TPC and moved towards the event's location. Once the ion has been collected, the electroprobe is removed from the TPC and is subsequently moved into the vicinity of the ion trap. At this stage, the separation step, denoted by a black box in *Figure 6.3*, takes place. This method will vary with the attraction mechanism of the electroprobe. Once the ion is separated from the probe and ejected into the ion trap, the two-laser system is used to fluoresce the ion to confirm its presence (or absence, as the case may be).

Figure 6.3: Ba-tagging line for a LXe TPC scheme with an electrostatic probe as the extraction device. There are five steps in the line. From left to right they are: the TPC, attraction to the electroprobe, removal from the electroprobe, the ion trap, the two-laser system.

6.2.2 COMSOL Model of Electroprobe

In this section we summarize the model that was built for simulating the attraction between singly ionized Barium atoms and an electrostatic probe. Appropriate boundary conditions and simplifications to the geometry are discussed.

In the extraction step, we must consider three forces that affect the motion of Ba^+ . Two are created by the probe itself. These are the electrostatic attraction between Ba^+ & the electroprobe, and the drag force exerted on the Barium ion due to the flow field created by the displacement of LXe by the probe. Apart from these two forces, there is also the constant electric field created by the TPC itself. A model of the extraction procedure, considering each of these forces simultaneously, was built and simulated using *COMSOL Multiphysics*, a finite element method (FEM) simulation software. The CAD model includes a simplified TPC, and a moving electroprobe with a spherical tip of diameter equal to the diameter of the electroprobe's shaft (see *Figure 6.4*). Other tip shapes have not yet been simulated but this is a topic of interest for future simulations. The dimensions of the model closely resemble the dimensions of a TPC currently stored at Carleton University that was built for the purposes of developing an extraction device for Ba-Tagging (i.e. this simulation is designed to apply nicely to a future experiment). We will see, however, that the results are mostly applicable to any TPC geometry as long as we remain *far enough* from the TPC components (i.e. field shaping rings, cathode/anode planes). This is essentially because fluid flow and electric fields are not disturbed by the TPC itself when we are far from its edges (under the assumption that the probe is moving slowly *enough* to satisfy the laminar flow assumption). The COMSOL CAD model of the TPC and the electroprobe is depicted in *Figure 6.4*, and its dimensions are summarized in the *Table 6.2*. The electrostatic boundary conditions of the simulation are also summarized in this table. The LXe conditions are set to 1.5 bar and 165 K. The saturation curve of Xenon is depicted in *Figure 6.5* to put these values in context.

Figure 6.4: CAD model of the TPC and probe simulated in COMSOL. Inner circular planes are for the purpose of a FEM meshing trick and do not affect physics results.

Figure 6.5: Saturation curve for Xenon. The leftmost orange point denotes the triple point. The red point denotes the conditions of LXe in the simulated TPC.

The TPC is simplified in simulations by removing its fine structure. For example, the cathode and the anode are simulated as solid, circular boundaries of a fixed potential, while they are crossed wires planes in practice. The field shaping rings are simplified as a continuous

cylindrical shell with a potential that ramps linearly from the anode to the cathode. At points far from the edges of the TPC, these simplifications don't affect the flow field or the electric fields. The resulting potential throughout the simulated volume is depicted in *Figure 6.6*.

Geometry Component/Boundary Condition	Value
TPC Diameter	12 cm
TPC Height	15 cm
Probe Diameter	3 mm
Cathode Potential	-8 kV
Anode Potential	0 kV
Probe Potential	Float*
Initial Probe Depth**	1cm
LXe Pressure	1.5 bar
LXe Temperature	165K

Table 6.2: Parameters for COMSOL simulation of electroprobe extraction of Ba^+ in a TPC.

*The surface floating probe surface takes on the potential of the cathode in the simulation.

** To avoid the difficulty of simulating the electroprobe's mesh passing through the cathode's mesh, the electroprobe begins the simulation partially submerged in the LXe.

Figure 6.6: Potential of surfaces in electroscope simulation. There is an imposed, linear change in potential along the walls of the TPC (walls are continuous – a simplification of the field shaping rings) from the anode to the cathode which are held at 0kV and -8kV respectively. The electroscope is set to a floating potential condition. It can be seen here that the surface of the probe takes on the potential of the cathode through which it passes during its descent.

In our *COMSOL* simulations, the drag force on Ba^+ is calculated using Stokes' drag law which applies to small spherical particles moving in a fluid. The force is defined by,

$$\vec{F}_d = 6\pi\eta R\vec{v} \quad 6.2.1$$

where η is the dynamic viscosity of LXe, $\vec{v}$ is the relative velocity of the Ba^+ particle with respect to the LXe that surrounds it, and R is the effective radius of the Ba^+ ion. Because the Barium particle is ionized, it will have an effective radius that is much larger than, say, its Van Der Waals radius of $5.4 \text{ \AA}$. It is believed that this is due, in large part, to the so-called *Snowball Effect* where a small layer of LXe forms around the Barium particle as it travels through the LXe, mostly caused by electrostrictive pressure. A study conducted by members of the EXO collaboration and other researchers were able to measure the effective radii of Ba^+ under various conditions [83]. We use those results in our simulations with an effective Ba^+ radius of approximately $8.5 \text{ \AA}$. The uncertainty on this value was not considered in any simulations.

The motion of the probe is completely vertical and defined in piecewise fashion as follows,

$$v = \begin{cases} -1 \frac{cm}{s} & , \quad t < 5.2s \\ -1 \frac{cm}{s} + 0.25 \frac{cm}{s^2} & , \quad 5.2s \leq t < 7.2s \\ 0 & , \quad 7.2s \leq t < 10.2s \end{cases} \quad 6.2.2$$

where v is the speed of the probe (positive direction being from cathode to anode, and t is time. In other words, the probe descends into the TPC at a speed of -1cm/s for 5.2s, then accelerates constantly for 2s down to a speed of 0cm/s. The simulation then carries on for 3s while the probe is at rest. It can be seen in *Figure 6.7* of the next section, which is a snapshot of the simulations at $t = 0$, that the probe begins slightly submerged in the LXe. This may not be the case in practice but it allows us to conveniently ignore the extremely difficult problem of simulating slip-flow at the boundary of LXe and GXe.

6.2.3 COMSOL Results for Single Ion Attraction with Electroprobe

Three scenarios – A, B, C (see *Figure 6.7*) – of the electroprobe were simulated. Each scenario was defined by the position of lowest descent, or “stopping point”, of the probe in the TPC. In all three scenarios, the stopping point of the electroprobe was at a height of 7.65cm, with the anode being at a height of zero and the cathode at 15cm. The radial position of the probe in *Scenario A* was zero (i.e. the probe descended along the longitudinal axis of the TPC). In *Scenario B* and *Scenario C*, it was shifted off the longitudinal axis by 3cm and 5cm respectively.

Figure 6.7: Electroprobe simulation scenarios. A scenario is defined by the position at which the electroprobe comes to rest in the TPC. The “height” parameter refers to the height at which the probe come to rest, with the anode being at a height of zero. The “distance from rings” parameter is the shortest distance from the probe’s stopping point to the rings. The TPC coordinates are given in centimetres.

For each scenario, the velocity map of a Ba^+ particle can be plotted according to,

$$\vec{v}(\vec{r}) = \frac{q}{6\pi R\eta} \vec{E}(\vec{r}) + \vec{v}_{\text{LXe}}(\vec{r}) = \mu \vec{E}(\vec{r}) + \vec{v}_{\text{LXe}}(\vec{r}) \quad 6.2.2$$

where μ is the mobility of Ba^+ in LXe, $\vec{E}$ is the electric field, and $\vec{v}_{\text{LXe}}(\vec{r})$ is the LXe velocity. The equation tells us that the velocity of Ba^+ is equal to its mobility multiplied by the electric field *plus* the velocity of the LXe around it – a boost into the frame of reference frame of the LXe itself.

This vector field describing the velocity of a Ba^+ was plotted for the three simulation cases. The results are shown in *Figure 6.8* and *Figure 6.9*. Scenarios *A* & *B* were found to be negligibly different within the capture regions and so only the results from the *Scenario A* and *Scenario C* are included in the figures.

Figure 6.8: Velocity maps for Ba^+ at various times during the motion of the probe. Each image corresponds to the map at a different cross-sectional time during the motion of electroprobe as defined by (6.2.2). *Top left:* probe moving at -1cm/s (time = 1s), *Top Right:* probe moving at speed of -1cm/s (time = 5.2s), *Bottom Left:* instant that probe comes to rest (time = 7.2s), *Bottom right:* probe has been at rest for 3s (time = 10.2s).

Figure 6.9: Velocity map for Ba^+ at instant that probe comes to rest (see equation 6.2.2) in *Scenario C*. **Top:** Velocity maps showing cut planes for two bottoms plots. **Bottom Left:** Velocity for *Cut Map A*. **Bottom Right:** Velocity for *Cut Map B*.

A particularly useful result to note in *Figure 6.8* is the maximum speed of Ba^+ as it converges towards the probe, of 2.51 cm/s (see the velocity map during the descent of the probe in *Figure 6.9 – bottom left*). If this value is substituted into (6.2.2) along with the maximum LXe flow speed just beneath the probe of 1 cm/s, we can calculate that we can add, at most, another 2.51 cm/s to the probe's speed (total probe speed of 3.51 cm/s) before all of the velocity map points away from the probe. Once the probe comes to rest, the drag force is reduced and the velocity map will slowly begin to converge towards the tip of the probe. If we let the probe sit at rest in the TPC for a long enough period, ions will converge to its tip. However, as the drift speed of Ba^+ for a field of, say, 1kV/cm is only approximately 2mm/s [83], this process can become impractically lengthy. One difficulty with this would be that, while the electroprobe is submerged, the field in the TPC is altered and data collection for other events will be perturbed. Understanding this perturbation would require a completely new R&D campaign. It is much more useful to choose some speed below 3.51 cm/s. With such a choice, ions will converge towards the probe during its descent and the amount of time that the probe must remain at rest in the TPC can be significantly reduced.

In *Figure 6.8* it can also be seen that the maximum Ba^+ speed when the probe's tip is near the anode is lower than its maximum speed when the probe is far from the anode, even when the probe is moving with the same velocity. This is explained as follows: when the probe's tip is less submerged, a smaller fraction of the probe is within the TPC to create an electrostatic attraction between the probe and the ion in the vertical direction. The field lines emanating from the part of the probe not contained within the detection volume do not penetrate the detection volume but instead end primarily on the field cage. As the probe descends, more of the probe is contained within the boundary and a greater attraction is created between the probe and the ion in the vertical direction due to these previously excluded field lines now being able to end on the ion itself. This observation means that there may be a z-dependence for the maximum acceptable probe speed. While *Figure 6.8* clearly shows that the velocity map still converges toward the tip of the probe when it is near the anode, the total volume for which convergence is observed appears to be reduced when compared to when the probe is deeper

in the TPC. The amount of reduction has not yet been quantified in simulations but may be of interest in the future.

In addition to mapping the velocity maps in each scenario, we also studied the capture probability of Ba^+ as a function of its position in the TPC. The origin was defined to be at the probe's point of lowest descent (i.e. the origin varied in each of the three scenarios). In each simulation, a uniform distribution of 4000 ions was generated within a cylindrical region with height equal to 2.65 cm, and radius of 7 mm. This region was chosen by combining the velocity map results with some quick back of the envelope calculations (not included here) – it was expected that anything outside of this region would be next to impossible to capture on the probe tip.

Our COMSOL simulations did not incorporate any kind of statistical fluctuations. For example, the electroprobe did not vibrate and the initial velocity of Ba^+ ions was always zero. In other words, the simulations were deterministic – an ion with some initial position in the TPC was either always be captured, or never captured. With the deterministic model, a region of 100 % capture probability was able to be defined for each simulated scenario. Since regions are either 100 % capture (a.k.a *capture regions*) or 0 % capture (a.k.a *loss regions*). The capture region for *Scenario A* is presented in 6.11. The figure was created using cylindrical coordinates as follows: For a fixed z-coordinate (representing distance from the tip of the probe), radial steps of 0.5 mm were taken until the point at which ions could no longer be captured was reached (where the probe tip is at a radial coordinate of zero). The value of the radius halfway between this step and the step just before this step was then labelled as the maximum radial capture coordinate for the fixed z-coordinate with an uncertainty equal to $\frac{1}{2}$. The z-coordinate was then incremented by 1 mm and the process was repeated. This process defined a series of circular planes with 100 % capture probability. The planes were connected by linear interpolation to create the full capture region. In the case of *Figure 6.10*, the probe remains at rest for 3s while ions continue to drift in the LXe.

Figure 6.10: Capture region for *Scenario A*. Radial and longitudinal dimensions are not to scale.

The shape of the capture region in *Figure 6.10* consists of distinct sub-regions of different radial sizes. In the first region, ions beginning at a radial position outside of 5.5mm move too far vertically before converging towards the probe's tip and end up colliding with the side of the probe or simply do not converge enough radially to reach the probe at all. This case is shown in *Figure 6.11* below.

Figure 6.11: Convergence of ions towards the electroprobe. This snapshot shows the trajectories of each ion from time zero to the point at which the probe comes to rest (i.e. from 0 s to 7.2 s). There are three cases: 1) Some ions move too far vertically before converging towards the probe resulting in collisions with the sides of the probe (cyan/light blue-green points) 2) Others simply do not converge enough radially in the duration of the probe's descent (dark blue points). 3) Finally, many converge nicely to the tip of the probe (red and orange points).

Once we get to the second region, some of the ions beginning outside of 5.5 mm are able to converge sufficiently before moving too much vertically to be captured. Thus, the capture radius increases to 6.0 mm. Something similar happens in the third region where ions can now begin with radial coordinates as large as 7.0 mm and still be captured. The capture region becomes more complicated after this last sub-region. First, the capture radius shrinks to 5 mm, then to 0.5 mm. Presumably the added radial distance outside of 5 mm results in too great a distance to travel to get to the probe tip before we begin removing it from the TPC (i.e. before the simulation ends), and eventually a radial coordinate greater than 0.5 mm becomes too large. But this cannot be the case because the capture radius then increases suddenly and significantly to 7.0 mm before eventually going back down to 0.5 mm. This seemingly strange

result can be explained by the flow field of LXe after the electroprobe has come to rest at its point of lowest descent. This field is shown in *Figure 6.12* at times of 1, 2, and 3 seconds after the probe has come to rest.

Figure 6.12: Flow field after various rest times for *Scenario A*.

It can be seen that the flow field directly below the electroprobe's tip goes to zero more rapidly than the field near the outer radius of the probe. This suggests the following explanation for the seemingly odd behaviour: The ions that begin in the regions with capture radii of 5.5 mm and 0.5 mm run into the flow field at the edge of the probe. This field inhibits their continued radial convergence and they end up passing the probe tip before capture can occur. The ions that begin deeper in the LXe are able to converge further radially before being in close proximity to the probe, allowing them to avoid this flow field. Instead, they approach from directly beneath the probe where there is little to no flow. This is why the capture radius is able to increase back out to 7 mm. Finally, at some initial depth, the ions are simply too deep to travel the distance to the probe unless they begin directly beneath it, and the capture radius drops back down to 0.5 mm.

It should be noted that the size of the capture region changes as a function of the amount of time, or *rest time*, that the probe is allowed to sit in the TPC at rest before it is withdrawn. A summary of the capture region for various rest times in *Scenario A* is given in *Table 6.3*.

Rest Time = 1s		Rest Time = 2s		Rest Time = 3s	
Capture Radius ($\pm 0.25\text{mm}$)	Distance from Probe Tip (mm)	Capture Radius ($\pm 0.25\text{mm}$)	Distance from Probe Tip (mm)	Capture Radius ($\pm 0.25\text{mm}$)	Distance from Probe Tip (mm)
5.75	1.0 - 6.0	5.75	1.0 - 6.0	5.75	1.0 - 6.0
6.25	7.0 - 8.0	6.25	7.0 - 8.0	6.25	7.0 - 8.0
7.0*	9.0 - 11.0	7.0*	9.0 - 11.0	7.0*	9.0 - 11.0
5.25	12.0 - 13.0	5.25	12.0 - 13.0	5.25	12.0 - 13.0
0.75	14.0 - 15.0	0.75	14.0 - 15.0	0.75	14.0 - 15.0
3.25	16.0	7.0*	16.0	7.0*	16.0
0.75	17.0	6.75	17.0	7.0*	17.0
0.75	18.0	4.75	18.0	7.0*	18.0
0.75	19.0 - 25.0	0.75	19.0 - 25.0	0.75	19.0 - 25.0

Table 6.3: Description of capture regions for various rest times in *Scenario A*. A rest time of 3s corresponds to the capture region in *Figure 6.10*. The bolded region in the table marks the sub-regions the capture region for which a time-dependence on the capture radius has been observed. Ranges of z-coordinate are used for some entries to conserve space. *Because the maximum initial radial coordinate of Ba^+ ions was 7.0 mm, any capture radii labelled as 7.0 mm may actually be greater than 7.0 mm.

Each row in *Table 6.3* corresponds to the z-coordinate and maximum radius defining a capture plane analogous to the planes depicted in *Figure 6.10*. The case for a rest time of 3 s is shown in *Figure 6.10*. The bolded region in the table marks the sub-regions of the capture region in which a time-dependence has been observed for the capture radius. The differences in values for different rest times show that the capture radius can be increased significantly by allowing the electroprobe to linger at rest in the TPC for only an additional one or two seconds.

So far we have only considered the capture region for *Scenario A* (and we have stated that it does not differ from *Scenario B* – at least not according to our chosen radial and vertical step sizes). The capture region for *Scenario C* is different to that of *Scenario A* in at least two ways. First, there is a trivial difference in the capture radii themselves. Second, we have found that the capture probability develops an angular dependence (where we continue to use cylindrical

coordinates with the z-axis along the longitudinal axis of the electroprobe). Both differences can be attributed to the broken cylindrical symmetry as the probe moves away from the center of the TPC. The results of *Table 6.4* describe the capture region for *Scenario C*, just as *Table 6.3* describes *Scenario A*, and show the trivial difference in capture radii for the two scenarios. Furthermore, though the angular dependence has not been considered in depth, we have shown that there is certainly an angular dependence as follows: We defined a vertical cut plane along the longitudinal axis of the electroprobe that split the ions into two groups. The cut plane is chosen such that the electric field and flow fields are different in each region (see *Figure 6.13*). We then determined that capture regions using exactly the same method as before. Now, because the distribution of ions is uniform (a cylindrical region with height equal to 2.65 cm, and radius of 7 mm beginning at the point of lowest descent of the probe),

Figure 6.13: Cut plane for splitting ion distribution into two regions in *Scenario C*. The image on the right is an overhead view of the image on the left. The overhead image is included to make it clear that the cut plane passes through the center of the probe. The cut plane breaks geometric symmetry producing different electric fields and flow fields on either side of the plane.

any difference in the capture regions for the two groups of ions suggests that there must be an angular dependence. As it turns out, we can see a difference for the two groups. The results are summarized in *Table 6.5* for a rest time of 1s. Further analysis of this angular dependence would be particularly useful for cases where an ion's initial position is close to a field shaping ring. For example, accurately quantifying the capture region near the field shaping rings may

suggest that it is feasible to keep the probe fairly far away from the rings while still attracting an ion that is close to the rings. One could then deliberately “miss” an ion in the radial direction, being sure that it would drift sufficiently radial to still be captured, so as to keep the probe far away from the field shaping rings and to prevent possible sparking between the two electrodes.

Rest Time = 1s		Rest Time = 2s		Rest Time = 3s	
Capture Radius (± 0.25 mm)	Distance from Probe Tip (mm)	Capture Radius (± 0.25 mm)	Distance from Probe Tip (mm)	Capture Radius (± 0.25 mm)	Distance from Probe Tip (mm)
3.25	1.0 - 2.0	3.25	1.0 - 2.0	3.25	1.0 - 2.0
2.75	3.0 - 4.0	2.75	3.0 - 4.0	2.75	3.0 - 4.0
3.25	5.0 - 6.0	3.25	5.0 - 6.0	3.25	5.0 - 6.0
0.25	7.0 - 8.0	0.25	7.0 - 8.0	0.25	7.0 - 8.0
0.75	9.0	0.75	9.0	0.75	9.0
7.0*	10.0 - 12.0	7.0*	10.0 - 12.0	7.0*	10.0 - 12.0
4.75	13.0	4.75	13.0	4.75	13.0
0.75	14.0	0.75	14.0	0.75	14.0
0.25	15.0	0.25	15.0	0.25	15.0
1.25	16.0	1.25	16.0	1.25	16.0
2.75	17.0	7.0*	17.0	7.0*	17.0
0.25	18.0	5.25	18.0	7.0*	18.0
0.25	19.0	3.25	19.0	3.25	19.0
0.25	20.0	0.25	20.0	2.25	20.0
0.25	21.0	0.25	21.0	1.75	21.0
0.25	22.0	0.25	22.0	0.25	22.0
1.25	23.0	1.25	23.0	1.25	23.0
0.25	24.0 - 25.0	0.25	24.0 - 25.0	0.25	24.0 - 25.0

Table 6.4: Description of capture regions for various rest times in *Scenario C*. The bolded region in the table marks the sub-regions the capture region for which a time-dependence on the capture radius has been observed. The bolded region in the table marks the sub-regions the capture region for which a time-dependence on the capture radius has been observed.

*Because the maximum initial radial coordinate of Ba^+ ions was 7.0 mm, any capture radii labelled as 7.0 mm may actually be greater than 7.0 mm.

Inner Region [*]		Outer Region [*]	
Capture Radius (± 0.25 mm)	Distance from Probe Tip (mm)	Capture Radius (± 0.25 mm)	Distance from Probe Tip (mm)
5.25	1.0 - 2.0	3.25	1.0 - 2.0
2.75	3.0	6.75	3.0
2.75	4.0	3.25	4.0
5.75	5.0 - 6.0	3.25	5.0 - 6.0
0.5	7.0 - 8.0	0.25	7.0 - 8.0
7.0 ^{**}	9.0	0.75	9.0
7.0 ^{**}	10.0 - 12.0	7.0 ^{**}	10.0 - 12.0
4.75	13.0	6.25	13.0
7.0 ^{**}	14.0	0.75	14.0
0.75	15.0	0.25	15.0
1.75	16.0	1.25	16.0
2.75	17.0	2.75	17.0
0.75	18.0	0.75	18.0
0.75	19.0 - 21.0	0.25	19.0 - 21.0
0.75	22.0	0.75	22.0
1.25	23.0	1.25	23.0
0.25	24.0	0.75	24.0
0.75	25.0	0.25	25.0

Table 6.5: Capture regions for the sub-regions of *Scenario C* with sub-regions defined in *Figure 6.13*. The different capture radii for equal z-coordinates demonstrates an angular dependence in capture probability for this scenario. Ranges of z-coordinate are used for some rows to conserve space.

^{*}Refer to *Figure 6.13* for definitions of *Inner Region* and *Outer Region* headings.

^{**}Because the maximum initial radial coordinate of Ba^+ ions was 7.0 mm, any capture radii labelled as 7.0 mm may actually be greater than 7.0 mm.

The surface plots and capture regions presented in this section demonstrate the convergence of ion trajectories towards the electroprobe and highlight regions in which capture is most likely in practice. We see that there are large volumes where the capture radius is as large as 7.0 mm. Considering that an experiment like nEXO is expected to have a positional

reconstruction resolution of $\sim 1\text{-}2$ mm, this is a promising result. That there is a large amount of room for error in both the radial and vertical directions suggests that there should be no great difficulty in capturing ions even after realistic probe vibrations and initial Ba^+ velocities are added to the model. Outside of adding realism with vibrations and initial velocities, there are clearly practically interesting ideas to be pursued in studying the angular dependence of the capture region in more detail than has been presented here. The reduction in attraction towards the electroprobe as it nears the anode may also provide useful results after further study.

To end this section we emphasize that most geometrical deviations from the simulated TPC should not have an effect on the conclusions that have been drawn here. The drag force exerted by the flow field will certainly not be affected by a change in TPC dimensions unless the Ba^+ is very close to the field shaping rings – a situation which is unlikely to occur in practice. Similarly, the attraction between Ba^+ and the electroprobe shouldn't be affected by TPC dimensions until Ba^+ is close to the field shaping rings (assuming the probe is at the same potential as in the current simulated case of -8kV). However, if the electrostatic boundary conditions are changed, the current results for capture probabilities will not generalize. Simulations will have to be repeated. Nevertheless, the *COMSOL* model that was developed here is easily adaptable and future models with different boundary conditions need only modify the existing file slightly. The simulation will take approximately one day to repeat for a new geometry. Thus, the results of this section can be directly applied to the TPC to be used in Ba-Tagging studies at Carleton University, while the *COMSOL* file can quickly be adapted to future, larger-scale TPCs such as nEXO. In any case, the results here enable the Ba-tagging R&D group to advance from capturing clusters of ions, which has been achieved in the past [79][80], to single ion capture. Should this endeavour prove successful, we can proceed to separation of ions from the electroprobe and ultimately towards the trapping and fluorescence steps of the complete Ba-tagging scheme.

6.3 Capillary Method of Ba^+ Extraction from a LXe TPC

I am an old man now and when I die and go to Heaven there are two matters on which I hope for enlightenment. One is quantum electrodynamics and the other is the turbulent motion of fluids. About the former I am really rather optimistic.

– Sir Horace Lamb

6.3.1 The Barium-Tagging Line with the Capillary Method

The motivation for using a capillary extraction technique with an induced phase change is that it allows combines a LXe TPC with the simplicity of a GXe Ba-tagging line. For example, when considering a cold probe where Ba^+ is captured in xenon ice of the probe's tip, it is unknown whether Ba^+ will reduce to Ba or remain ionized. The GXe tagging line has no such issues with charge reduction. Furthermore, while LXe TPC methods have been able to extract clusters of ions, attempts to separate single ions from the extraction device have partially failed. For a GXe TPC, the typically proposed tagging line does not require separation of barium from an extraction device [84]. The implementation of the capillary method calls for a minor modification of the typical gas-phase TPC Ba-tagging line which is depicted in *Figure 6.14*.

Figure 6.14: Concept of a Ba-tagging line for a GXe TPC operating at 10 bar. The above diagram is the general concept while the lower diagram depicts a setup that has already been tested at McGill University [85].

In *Figure 6.14* the tagging line commences with a GXe TPC. In this tagging scheme, when an event near the $0\nu\beta\beta$ Q-value occurs, the electric field in the TPC is modified to drift ions towards a hole in the TPC wall. An ion passes through this hole into a converging-diverging nozzle that takes in high pressure, subsonic flow and outputs low pressure supersonic flow. This flow is passed through an RF-funnel which focuses Ba^{++} and separates it from GXe and

background impurities, and then passes it to an ion guide for further filtration [85][86]. Next is a collision cell in which Ba^{++} is converted into Ba^+ by the following reaction,

where *TEA* is a commercially available gas called triethylamine. In the past, the Ba-tagging group at Carleton has shown that (6.3.1) occurs with 100% transmission efficiency in a collision cell at the appropriate settings [87]. The now singly charged barium ion is passed to a linear Paul trap (LPT) for identification by LIF. The purpose of (6.3.1) is to enable this induced fluorescence by excitation of the 4 s valence electron in Ba^+ .

The filtration steps that make use of the RF-funnel and an ion guide to separate Ba^{++} from GXe and any other contaminants are where the GXe TPC scheme has been superior to the LXe TPC scheme so far. Because of the characteristic charge to mass ratio of particles, GXe particles and other charged particles are easily filtered with the RF-funnel and the ion guide. Indeed, this is an old and well-understood technique for separating particles of different charge to mass ratios. By contrast, separation of single ions from electrostatic probes in a controlled manner requires further engineering and research. For example, in the past, EXO's Ba-tagging group has failed to separate Ba^+ particles from a cold probe, supposedly due to the strong image charge created by the Ba^+ on the metallic probe.

While the GXe TPC scheme has its advantages in the context of Ba-tagging, nEXO will operate with a LXe TPC due to its advantages outside of the context of Ba-tagging [22]. This is why the capillary method has become interesting. In the capillary scheme, the GXe TPC is replaced by a LXe TPC, the collision cell is removed (Ba^{++} readily reduces to Ba^+ in LXe eliminating the need for a collision cell), and the capillary is placed between the TPC and the converging-diverging nozzle. The capillary then allows us to combine the LXe TPC with the GXe tagging line as follows: The capillary is inserted into the LXe TPC at a pressure of, say, 1.5 bar near a hypothetical $\beta\beta$ event. The other end of the capillary is housed in a chamber of lower pressure. The pressure difference across the capillary causes LXe to flow from the TPC to the low-pressure chamber,

dragging the Ba^+ with it. The chamber is presumably able to be kept at some fixed low pressure via continuous evacuation with something similar to McGill's nozzle/funnel system. The capillary itself is heated to some temperature above the boiling point of LXe, inducing a phase change through nucleate boiling. The GXe is then ejected into the low-pressure chamber along with the Ba^+ ion. The rest of the tagging line is then the typical GXe line with the collision cell removed. The complete capillary line is depicted in *Figure 6.15*.

Figure 6.15: Ba-tagging line for the capillary extraction scheme.

6.3.2 – Design of Flow Conditions in Capillary

As xenon flows from the TPC into the low-pressure chamber via the capillary, the capillary itself is heated to induce nucleate boiling of the LXe. As the phase change from LXe to GXe progresses along the capillary, various flow regimes can be defined. These regimes are illustrated in *Figure 6.16*.

Figure 6.16: Flow regimes in a fluid undergoing nucleate boiling in a long tube. The temperature of the tube versus the fluid temperature are depicted on the left of the figure for a typical fluid. Flow regimes and the type of heat transfer that dominates each regime are given on the right of the figure. Red arrows indicate the convergence of the gas phase as bubbles are ejected from the wall and the divergence of a small annular liquid region that replaces the volume of escaping bubbles.

It can be seen in *Figure 6.16* that the inlet (TPC end of the capillary) is a region of single-phase LXe flow, while the outlet (low-pressure chamber end of the capillary) is a region of single-phase GXe flow. Between these flow regimes is a region of LXe/GXe mixture. To study the behaviour of Ba^+ as it is dragged along the capillary by the flowing xenon three simulations

cases have been considered – the single-phase LXe flow regime, the single-phase GXe flow regime, and the mixture regime.

In general, the flow in any of these regimes can be laminar – the smooth flow of liquid in immiscible layers – or turbulent – characterized by eddies. In our case, the objective is to design flow conditions in which a Ba^+ ion can be transported from one end of the capillary to the other without loss – the ion is presumed lost if it makes contact with the capillary wall. For this reason, it is desirable to design the flow to be laminar, where the streamlines (i.e. the trajectories that a massless or very light particle follows due to the fluid drag force) are fairly straight from the inlet to the outlet of the capillary. Now, in the mixture regions of the flow, we cannot expect laminar flow. However, in general, the diverging behaviour of bubbles/gas from surfaces during nucleate boiling actually results in the convergence of streamlines for enclosed flow systems like a tube (see *Figure 6.17*). This phenomenon results from the buoyancy of the gas phase. Thus, we naively expect that ions will not be lost in great numbers in such a region of flow. In the pure LXe and GXe regions, by contrast, a turbulent flow has no obvious reason to converge towards the center of a capillary. Indeed, the eddies within the flow may end up launching particles towards the capillary walls resulting in heavy losses. Thus, the imposed boundary conditions of our capillary system must be designed to ensure a laminar flow in these regions.

Figure 6.17: Illustration of nucleate boiling at a rectangular interface between a liquid and a heated surface.

In order to simplify the characterization of flow, a dimensionless number, the Reynolds number, has become a standard value for choosing which set of equations, laminar or

turbulent, to use when modelling flow. A more detailed discussion of this quantity can be found in *Appendix A*. The range of values in which the Reynolds number falls suggests which set of equations should be used for modelling flow. This number, denoted by R_e , is defined by,

$$R_e = \frac{\rho D v}{\mu} \quad (6.3.2)$$

where ρ and μ are the fluid's density and dynamic viscosity respectively, D is a characteristic length scale, and v is a characteristic speed. As R_e increases from zero we transition from creeping to laminar to turbulent flows. As a rule of thumb, unobstructed flows can be modelled as laminar for $R_e < 2300$ and must be modelled as turbulent for $R_e > 4000$. Between these values is a complex mixed region.

In practice, there are some obscure points in the definition of R_e . For one, the density and dynamic viscosity of a fluid can depend on temperature and pressure. In our case, since the objective is to induce a phase change along a capillary, the fluid temperature and pressure will change throughout the motion. Thus, in order to ensure laminar flow conditions throughout the flow volume, the values for temperature and that produce the largest R_e are chosen. If this R_e is within the laminar regime, then the entire flow volume will be within the laminar regime. Outside of changing temperature and pressure conditions, there is also the issue of defining how the “characteristic” values are chosen. For long, cylindrical pipes in which the length is much greater than the diameter, the characteristic length scale of the system is taken to be the pipe diameter. This is certainly the case for our model in which the capillary has a diameter of $\sim 100 \mu\text{m}$ and a length of $\sim 2 \text{ m}$. The characteristic speed is usually taken to be the average speed of the flow (half the maximum speed for laminar flow).

In addition to the Reynolds number, the Mach number (i.e. the ratio of a fluid's speed to the speed of sound in the fluid) must also be considered to ensure laminar flow conditions. Large Mach numbers produce compressible flows in which the concept of laminar flow disappears. Typical regions for values of Mach number and Reynolds number that produce incompressible, laminar flows are shown in *Figure 6.18*. These regions are rules of thumb and, generally, the regions should be verified by experiments for a given fluid. In our simulations, since this data is

unavailable for LXe/GXe, we simply attempt to stay safely within the typical incompressible and laminar regions to account for the possibility that the regions are smaller for Xenon in practice.

Figure 6.18: Typical, experimentally determined flow types (compressible or incompressible) for various Mach numbers (left) and Reynolds number regions producing laminar or turbulent flow (right).

The *head loss* (a.k.a pressure drop) of a fluid flowing through a circular pipe is described by the Darcy-Weisbach equation,

$$\Delta P = \rho g h_f \quad (6.2.4)$$

where g is the gravitational constant and h_f is a collection of variables grouped together such that (6.2.4) resembles the equation for hydrostatic pressure. It is defined by,

$$h_f = f_D \frac{L}{D} \frac{\overline{||\vec{v}||}^2}{2g} \quad (6.2.5)$$

where L is the pipe's length and D is its diameter. For the case of laminar flow, Poiseuille was able to show that f_D is equal to $64/R_e$ in his studies of pressure induced flow.

The head loss, defined by (6.2.4), paired with the Reynolds number are the two variables that are used in this work to design a capillary extraction system in which the flow may be modelled as laminar. If the LXe is to be extracted into a lower pressure environment, as per the Ba-tagging scheme described in section 6.3.1, then the pumping speed of the vacuum in the low pressure region of the apparatus can be used to define the maximum acceptable speed of GXe and LXe (while still enforcing the laminar condition). The plots in *Figure 6.19* show the Reynolds

number, for both GXe and LXe, at the maximum pumpable flow rate that was reported by Thomas Brunner's team at McGill University with their RF-Funnel/nozzle extraction system [85]. To obtain the plots, a capillary diameter of 180 μm was used (a commercially available diameter).

Figure 6.19: Plots of the Reynolds number for GXe (left) and for LXe (right) when the capillary is designed to displace the maximum mass of fluid that can be pumped at the McGill University RF-Funnel [86].

In *Figure 6.19*, it can be seen that the Reynolds number in both GXe and LXe is safely under 1000 for nearly any temperature-pressure pairing where the temperature is under 200 K. If we consider the cross section of xenon's phase diagram shown in *Figure 6.5*, we see that a phase change from LXe to GXe can be induced at a temperature of around 170 K. Furthermore, we even see that as the pressure decreases for GXe, the Reynolds number falls ever lower into the safe, laminar regime. Since the Darcy-Weisbach equation tells us that pressure will decrease along the pipe, this enables us to superheat the GXe to temperatures much higher than the phase change temperature of approximately 170 K.

In order to scale these results to various other capillary diameters, the following version of the average fluid speed (used to get the Mach number) and the Reynolds number can be used,

$$\langle v \rangle = \frac{1}{32} \frac{\Delta P D^2}{L \mu} \quad (6.3.8)$$

$$R_e = \frac{1}{32} \frac{\rho}{\mu^2} \frac{\Delta P}{L} D^3 \quad (6.3.9)$$

where we only need to change the diameter to update the result for an unchanged ratio of head loss to capillary length. The sharp dependence on diameter suggests that the pressure gradient along the capillary will have to be changed significantly to accommodate a tube of, say, 1mm diameter instead of a capillary. As the expected positional resolution of physics events in a TPC is ~ 1 mm, such a change may have to be made in order to suck Ba^+ into the capillary successfully if other methods for doing this are not devised.

6.3.3 Capillary Simulations in *SIMION*

In the laminar flow case of single-phase flow, we expect that a neutral, massless particle will follow the streamlines of the flow field. Since streamlines in laminar pipe flow are fairly straight from inlet to outlet, we reasonably expect 100% transmission. Since the mass of a single Barium atom is very small, it is reasonable to expect the same result for a neutral Barium atom. In the following plots, we show results from COMSOL that confirm this expectation when a Stokes' drag force is applied to the ions in the same way as in the electroprobe simulations of *Chapter 6 – 6.2*. The capillary in the simulation was designed with a diameter of $180\text{ }\mu\text{m}$ and a length of 30 cm. This length is not expected to model the length of a real capillary, which is expected to be on the order of a few metres in length. However, using a length any longer than this causes meshing difficulties due to the aspect ratio of the system. To remedy this, we proceeded as follows. We assumed a real capillary length of 2 m, a length greater than the depth of the proposed nEXO TPC, therefore allowing the capillary to reach any event location in the LXe volume. We then assumed that the capillary would eject xenon into a perfect vacuum at its outlet. If we consider (6.3.8) and (6.3.9), this is essentially a worst case scenario when considering the Reynolds number as it maximizes the head-loss, ΔP . That is, if this pressure condition produces laminar flow, so too will a higher pressure at the outlet. For a TPC pressure of 1.5 bar, this produces a pressure gradient of 0.75 bar/m. This pressure gradient was used to calculate the equivalent pressure at 30 cm along the 2 m capillary. This pressure was then set as the outlet pressure in our COMSOL simulations where the capillary was 30 cm in length. Two simulations were run with these conditions – one for LXe and one for GXe. Select results of

these simulations are shown below in *Figure 6.20* where we see the streamlines of the flow near the outlet of the capillary. We see that the streamlines are straight and non-overlapping, suggesting that laminar flow holds under the applied boundary conditions.

Figure 6.20: *Left:* Streamlines for the single-phase LXe flow with applied 0.75bar/m pressure gradient in 1 180um diameter capillary. Results are obtained from COMSOL FEM software with the *laminar flow* module. The analogous single-phase GXe version of this plot is identical apart from the fact that the flow speed is greater for GXe. *Right:* Streamlines near the outlet in single-phase GXe flow with the same applied BCs. The red arrows indicate the direction of GXe flow at the cross-section of the capillary's outlet.

Near the outlet, we see a divergence away from the capillary's center. This is a consequence of the constant pressure boundary condition applied to the outlet. Indeed, if the flow were to be ejected into some chamber being pumped to hold some fixed pressure, we would expect a diverging jet as the flow transitioned from the capillary to the chamber. In other words, we see nothing unexpected here and we can conclude that the chosen boundary conditions produce laminar flow in both LXe and GXe. Furthermore, we have calculated mass fluxes (average speeds) of 0.01 g/s (16 cm/s) and 0.00036 g/s (300 cm/s) for the LXe and GXe simulations respectively. First of all, with Mach numbers of ~ 0.0002 in LXe and ~ 0.02 in GXe, both flow speeds produce Mach numbers within the incompressible flow regime (see *Figure 6.18*). Secondly, upon comparing the mass fluxes to *Figure 6.19* we can conclude that the applied boundary conditions of the flow produce a mass flux that is easily handled by the nozzle/RF-funnel apparatus at McGill University that can pump up to 0.0575 g/s [85].

In practice, we expect to have single ionized Barium moving through the capillary. Thus, it becomes important to understand the behaviour of Ba^+ when electrostatics is considered (e.g.

the image charge on the capillary). This has been studied using *SIMION 8.0*. In the *SIMION* simulation, it is assumed that the capillary's voltage is floating and that it passes through the cathode to enter the TPC. Thus, much like in the TPC case, we assume that the capillary surface takes on the potential of the cathode and we set its potential to -8 kV for the simulation. This allows us to avoid the difficulties of simulating a floating capillary and having to design a TPC in which to house it in *SIMION*. We use the maximum flow speed from *COMSOL* simulations to program a perfect, quadratic laminar flow Stokes' drag force (see *equation 6.2.1*) into the simulation. To compute the drag force, the effective radius of Ba^+ is set to 8.5 Å just as in the electroprobe simulations. Finally, Barium ions are defined to enter through the inlet of the capillary with the same velocity as the flow field. The results for a simulation of 5000 ions are summarized in *Figure 6.21*.

Figure 6.21: *Left:* Ba^+ trajectories along the capillary in *SIMION* simulations. *Left:* Side view of ions trajectories in the capillary. Green lines mark trajectories. *Right:* Red dots mark positions of Ba^+ ions at the inlet and outlet of the capillary. Red dots are used to mark the beginning and ending points (notice that there aren't any in the left image, signifying no losses along the capillary).

The results show that Ba^+ ions pass essentially straight through the capillary along the LXe streamlines. In the case of perfect, quadratic laminar flow, 100% transmission is achieved for a test of 5000 Ba^+ ions. To determine whether this is a result of the image charge causing the ions to converge or whether the fluid drag is preventing the ions from diverging, we have removed the drag force in a second simulation. The result is depicted in the figure below. We see that the ions move along only a small fraction of the length of the capillary before diverging. We can therefore conclude that the force created by the image charge is divergent, but that the fluid drag is sufficient to prevent such divergence.

Figure 6.22: Trajectories of Ba^+ ions in a capillary. Red dots mark the beginning and ending points of an ion's trajectory through the capillary. The farthest that any of them travel before colliding with the wall of the capillary is 14 cm.

The results of this section have summarized advancements in understanding the feasibility of the capillary extraction method. A previously unanswered question about the nature of the force on Ba^+ ions created by its image on the capillary is now answered. We can conclude that while the image charge force is divergent, the fluid drag is sufficient to prevent losses along a capillary of 2 m in length. There is also room for further R&D related to the capillary before beginning construction of a physical apparatus. In particular, the nucleate boiling regime remains to be fully simulated and understood. We expect to see an acceleration in the flow speed in order to conserve the flow of mass. The results here show that we can expect to preserve laminar flow as we transition from LXe to GXe flow regimes up to a GXe speed of thousands of centimetres per second for a capillary of 180 μm diameter. This is an extraordinarily wide range of parameter space and certainly adds to the feasibility of employing the capillary method in experiment. Finally, the nature of the jet at the outlet of the capillary provides an interesting R&D path for combining work already done on the nozzle/RF-Funnel system at McGill University with capillary studies that have been done at Carleton University. Such a combination would be a great step in advancing towards completion of the Barium-tagging line.

Chapter 7

Closing Remarks

The objective of this work has been to produce advancements in three methods for reducing backgrounds in EXO. The three methods that have been explored are high voltage ramping in LXe TPCs, machine learning analysis, and Barium-tagging.

In Chapter 4 the results of the early stages of an approach to apply pattern recognition to high voltage breakdown in liquid xenon were presented. For the first time, the results from various high voltage R&D groups throughout the EXO collaboration were pooled in an attempt to predict the threshold for high voltage breakdown in a liquid xenon time projection chamber. While the result of the pattern recognition approach greatly overestimated the breakdown voltage in comparison to experimental results, there remain many R&D paths for improving upon the base method. Of particular interest for future work are at least two points. The first is the optimization of the pattern recognition algorithm. For example, feature selection can be considered. Indeed, in the papers that inspired the use of the pattern recognition method, it was found that performance could be improved if only a subset of the full feature list was used [61][62]. There is also the question of which values of the SVM hyperparameters will produce the best performance. The second point is to apply the proposed method for incorporating the statistical nature of breakdown in liquid dielectrics into the SVM. This endeavour calls for EXO's high voltage groups to return to the laboratory to apply the proposed systematized method for recording breakdown measurements across all of EXO high voltage R&D. This method provides an exciting opportunity to unify the research of various groups of authors into a framework for quantifying the expected breakdown voltage in liquid dielectrics [53][61]. Should the result prove to be successful, it will provide a practical approach for safely ramping to high voltages in liquid dielectrics in general. This is immediately practical in experiments like nEXO, as we have seen, but also in experiments like *DUNE*, the long-baseline neutrino oscillation experiment which should be able to determine the mass hierarchy of neutrino. This experiment is intended to operate with tonne-scale liquid drift chambers much like nEXO. Finally, because this method

allows different apparatus geometries to inform each other, we can imagine some point in the future when various high voltage R&D groups around the world pool their data into a single database. Here, authors can use the each other's results to design new experiments in which liquid dielectrics are used.

In the machine learning analysis of Chapter 5 we were able to employ complex algorithms such as *stacking* and *diversity pruning* to build upon the past work of Warren Cree. Perhaps the most important result in this chapter was the derivation of a likelihood function for estimating the true distribution of event types given a vector of classified events. In the past, particle physics groups have been unable to do this directly and work-around techniques had to be employed to construct confidence intervals for the true distribution of events. Work continues to be done by myself, and others within EXO's machine learning group, in advancing the rigorous mathematical formalism of this technique. These new developments should prove to have general applicability in the context of classification algorithms applied to particle physics experiments and beyond.

Finally, results of Barium-tagging simulations summarized in Chapter 6 have provided insight into the nature of the interaction of single Ba^+ atoms with various extraction apparatuses. In the past, even with the cold probe approach, which has undoubtedly advanced more than other proposed electroprobe methods [79][80][88-90], we have only been able to capture ion clusters [79][80][88]. Because we have been able to determine the capture radius for a simple model of an electrostatic probe, these results will allow the various groups involved in the barium-tagging effort across nEXO to move away from capturing clusters of ions to single ion capture. Results from simulations of the capillary method are perhaps less complete than those of the electroprobe. Nevertheless, the results appear to suggest there is a high transmission probability for Ba^+ . Indeed, we have been able to conclude that while the image charge force causes the Ba^+ to diverge, the fluid drag force is sufficient to prevent this from happening. The next logical step for research into the capillary method is no doubt to complete simulations of the nucleate boiling flow regime. This is the only remaining region to be simulated before it can be concluded that the capillary is, or is not, worth pursuing in practice. Should these results continue to show high transmission probability, it is possible that the Barium-Tagging groups at

McGill University and Carleton University may begin developing an experimental line for connecting the capillary to McGill's RF-funnel. This would be a major advancement in the overall Ba-tagging endeavour.

To summarize, this work has been able to advance research into three methods of background reduction capabilities for EXO – high voltage for improved measurement resolution, machine learning for event discrimination, and Barium-tagging for 100% rejection of non- $\beta\beta$ events. While progress has been made, we have seen that each of these campaigns presents further opportunities for the EXO collaboration to continue in an effort further improve its background reduction capabilities. Such improvements will allow EXO to constrain limits for a theory of Majorana neutrinos ever further. We can therefore expect exciting new developments within the EXO collaboration and the field of neutrino physics in the years to come.

Appendix A – The Reynolds Number

This section explains the physical significance of the Reynolds number by invoking the Buckingham Pi Theorem and the dimensionless form of the Navier-Stokes equations – a set of equations that describes fluid flow.

The Buckingham Pi Theorem is an extraordinarily useful tool in dimensional analysis. In physical problems described by unknown equations, it can be used to derive dimensionless variables that characterize the physics of a given system. It does not, however, produce the equations in which these variables are used – a task left to the researcher. Once the correct equation is found, one should always be able to render it dimensionless with a form in which the variables derived by use of the Buckingham Pi Theorem are present. In other words, all problems can be reduced to dimensionless problems using the appropriate Buckingham-Pi variables for the problem. This discovery is generally attributed to Lord Rayleigh [92].

The Buckingham Pi theorem states the following: Suppose a physical system can be described by some equation,

$$0 = f(q_1, q_2, \dots, q_n) \quad (\text{A.1})$$

, where q_i is the i^{th} physical variable, n is the number of variables required to describe the system, and f is some function of the n variables. Then it is possible to restate (A.1) as a dimensionless equation of dimensionless variables,

$$0 = F(\pi_1, \pi_2, \dots, \pi_m) \quad (\text{A.2})$$

, where π_i is the i^{th} dimensionless variable describing the system, m is the number of dimensionless variables required to describe the system, and F is some function of said variables. The total number of dimensionless variables is given by,

$$m = n - k \quad (\text{A.3})$$

, where k is the total number of fundamental units (e.g. mass, time, length; while, for example, the Newton is not fundamental if mass, time, length are taken as fundamental as it can be decomposed into these three units) in (A.1).

For the case of drag on a spherical particle in a fluid, the relevant variables for the corresponding version of (A.1) are the diameter of the particle, D , speed of the fluid, v , density of the fluid, ρ , viscosity of the fluid, μ , and the applied force (e.g. pressure gradient), F . Here we have ignored variables related to compressibility of the fluid. In order for this to be valid, the Mach number of the fluid must be in the incompressible range (another laminar condition). Since this condition is addressed in *Chapter 6*, we can safely ignore such variables in what follows. To determine the π_i corresponding to (A.2), the number of fundamental units must be determined,

$$[D] = L$$

$$[v] = \frac{L}{T}$$

$$[\rho] = \frac{M}{L^3}$$

$$[\mu] = \frac{M}{LT}$$

$$[F] = \frac{ML}{T^2}$$

, where square brackets are an operation that converts a variable to its units, and where M , L , and T represent the units of mass, length, and time respectively.

There are clearly five variables while there are only three fundamental units describing them. Thus, by (A.3), there must be two dimensionless variables describing the system in a dimensionless form. It can be shown that these two dimensionless variables are the Reynolds number and a variable proportional to head loss [92][93]. The latter is essentially the Darcy-Weisbach equation that was seen in *Chapter 6* for describing the average speed of pipe flow across which is a fixed pressure gradient.

The choice of fundamental units in the above discussion only applies if the Mach number is sufficiently low to have a non-compressible fluid. This is why we also need to consider the Mach number in *Chapter 6* when we consider laminar flow.

In *Chapter 6* we considered various statements about the so-called Reynolds number, defined by,

$$R_e = \frac{\rho D v}{\mu} \quad (\text{A.4})$$

Here we will “derive” the number. We begin with the dimensionless form of the incompressible Navier-Stokes equation (compressible fluids are far more complex and are not dealt with in this work – indeed, we avoid it by design of the apparatus in practice),

$$\rho \frac{D\vec{v}}{Dt} = -\vec{\nabla} p + \mu \nabla^2 \vec{v} + \rho \vec{f} \quad (\text{A.5})$$

, where (A.5) is the Lagrangian derivative.

We can multiply (A.5) by $\frac{L}{\rho \|\vec{v}\|^2}$, where L is a characteristic length and $\|\vec{v}\|$ is the average velocity, to get,

$$\frac{L}{\|\vec{v}\|^2} \frac{D\vec{v}}{Dt} = -\frac{L}{\rho \|\vec{v}\|^2} \vec{\nabla} p + \mu \frac{L}{\rho \|\vec{v}\|^2} \nabla^2 \vec{v} + \frac{L}{\|\vec{v}\|^2} \vec{f} \quad (\text{A.6})$$

We then group variables to simplify,

$$\vec{v}' = \frac{\vec{v}}{\|\vec{v}\|}, p' = \frac{p}{\rho \|\vec{v}\|^2}, \vec{f}' = \vec{f} \frac{L}{\|\vec{v}\|^2}, \frac{\partial}{\partial t'} = \frac{L}{\|\vec{v}\|} \frac{\partial}{\partial t}, \vec{\nabla}' = L \vec{\nabla}$$

, which leaves,

$$\frac{D\vec{v}'}{Dt'} = -\vec{\nabla}' p' + \frac{\mu}{\rho L \|\vec{v}\|} \nabla'^2 \vec{v}' + \vec{f}' = -\vec{\nabla}' p' + \frac{1}{R_e} \nabla'^2 \vec{v}' + \vec{f}' \quad (\text{A.7})$$

The rightmost form shown in (A.7) provides a clear description of the relevance of the Reynolds number – as it grows, the viscous term (i.e. the second term in the rightmost form of (A.7)) becomes ever less relevant. This is essentially the laminar flow condition. Thus, very small Reynolds numbers, $R_e \ll 1$, signify that the inertial force term in the Navier-Stokes equation is negligible and we get *creeping flow*. In this flow regime the viscous force is equal to the pressure force minus the external force. As R_e increases we transition from creeping to laminar to turbulent flows: In unobstructed flows, flow can typically be modelled as laminar for $R_e < 2300$. It must be modelled as turbulent for $R_e > 4000$. Between these values is a complex mixed region.

The Reynolds number can be thought of as a scale factor for the size of periodic motion in the fluid, usually called *eddies*. As previously stated, the Reynolds number is a ratio of inertial forces to viscous forces. Thus, as this number grows, the momentum of the fluid is not significantly restricted by the viscosity of the fluid. Because of this, one must focus on ever smaller length scales in the flow before viscous forces become important again. However, before they become important, we expect to see turbulence. In this sense, the Reynolds number describes turbulence at the largest scale of flow – if there isn't any turbulence at the largest scale then there can't be any at the smallest.

The reader may be concerned that the initial multiplicative factor used to render the Navier-Stokes equation dimensionless is arbitrary – surely there are other combinations of variables that could achieve the same result. However, as we have seen, the Buckingham Pi theorem produces the Reynolds number as one of the two dimensionless numbers describing fluid motion. This guides the choice of multiplicative factor – it is chosen to render the equation dimensionless while deliberately causing the Reynolds number to appear. In this sense, the Buckingham pi theorem and the above manipulation are complementary. The Buckingham Pi Theorem only produces dimensionless variables and not the equations governing their relationships to each other. The equations must be found and so they have been in the derivation above.

B Increased Electric Field Strength due to Acrylic Struts

In *Chapter 4* it was stated that acrylic struts will increase the electric field in LXe for a gap-like geometry if the dielectric constant of the acrylic is greater than that of LXe. The figure below illustrates this phenomenon by simplifying the geometry into a rectangular case where a Gaussian pillbox can be used to determine the field. This figure was taken from a slideshow made by Sebastien Delaquis in his own thesis work. In the figure, the dielectric-filled gap is defined by width, $D-h$, where the acrylic strut has width, h . The strut has dielectric constant, ϵ_p , while the liquid dielectric has dielectric constant, ϵ_0 . Capacitance is denoted by C , charge by Q , surface area by A , field without any dielectric by E_0 , and field with the dielectric by $\hat{E}_0$. For small values of $D-h$ (i.e. for small gaps), the field becomes large.

The Electric Field in the gap

$$\oint \epsilon_0 \hat{E}_0 \cdot d\mathbf{s} = Q \quad \therefore \hat{E}_0 = Q / \epsilon_0 A$$

Imagine two capacitors in series

$$C V = Q \quad C = C_0 C_p / (C_0 + C_p)$$

$$C = \frac{\epsilon_0 \epsilon_p A}{\epsilon_0 h + \epsilon_p (D-h)}$$

$$\hat{E}_0 = \frac{\epsilon_0 \epsilon_p A V / \epsilon_0 A}{\epsilon_0 h + \epsilon_p (D-h)}$$

Identify the result without any dielectric :

$$E_0 = V/D,$$

$$\hat{E}_0 = \frac{E_0}{\frac{h}{D} \left(\frac{\epsilon_0}{\epsilon_p} - 1 \right) + 1}$$

Hence, one sees that for $\epsilon_p \gg 1$, the ratio $R = \hat{E}_0 / E_0 \rightarrow D / (D-h)$, which for a small gap can be large.

For a fixed dielectric constant, the max possible is ϵ_p / ϵ_0 . This is illustrated by the plot to the left, where we have taken the EXO-200 values $D = 11$ mm and $h = 3$ mm, (cathode to vessel gap, and acrylic "comb" to vessel gap)

Figure B.1: Physical origin of the increased electric field in the gap between an acrylic strut and a field shaping ring in a TPC.

References

- [1] SNO Collaboration, *Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by 8B solar neutrinos at the Sudbury Neutrino Observatory*, Phys. Rev. Lett. 87 (2001)
- [2] Fukuda, Y. *et al.* *Evidence for oscillation of atmospheric neutrinos*. Phys. Rev. Lett. 81, 1562–1567 (1998).
- [3] Furry, W. H. *On transition probabilities in double beta-disintegration*. Phys. Rev. 56, 1184–1193 (1939).
- [4] Chadwick, J. *Possible Existence of a Neutron*. Nature 129, 402–402 (1932).
- [5] Cowan, C. L., *et al.* *Detection of the Free Neutrino: A Confirmation*. Science, vol. 124, no. 3212, 1956, pp. 103–104. JSTOR, JSTOR, www.jstor.org/stable/1751492.
- [6] Bethe, H. A. & Critchfield, C. L. *The formation of deuterons by proton combination*. Phys. Rev. **54**, 248–254 (1938).
- [7] Griffiths, D. *Introduction to Elementary Particles*. (Wiley, 2008).
- [8] Schwetz, T., Tórtola, M. & Valle, J. W. F. *Three-flavour neutrino oscillation update*. New J. Phys. **10**, (2008).
- [9] Bellerive, A., Klein, J. R., McDonald, A. B., Noble, A. J. & Poon, A. W. P. *The Sudbury Neutrino Observatory*. Nucl. Phys. **B 908**, 30–51 (2016).
- [10] G. Danby; J.-M. Gaillard; K. Goulianos; L. M. Lederman; N. B. Mistry; M. Schwartz; J. Steinberger (1962). *Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos*. Phys. Rev. Lett. **9**, 36-44 (1962).
- [11] K. Kodama *et al.*, *Observation of Tau Neutrino Interactions*, Phys. Rev. Lett. **B 504**, 218–224, (2001).
- [12] F. Capozzi, G.L. Fogli, E. Lisi, A. Marrone, D. Montanino and A. Palazzo, *Status of three-neutrino oscillation parameters, circa 2013*, Phys. Rev. D 89 (2014)
- [13] Nosek, T. *Effects of Matter in Neutrino Oscillations and Determination of Neutrino Mass Hierarchy at Long-baseline Experiments*. (2016). arXiv:1612.09132

- [14] Goeppert-Mayer M. *Double Beta-Disintegration*. Phys. Rev. 48, 512 (1935).
- [15] Majorana, E. *A symmetrical theory of electrons and positrons*. Nuovo Cim. **14**, 171 (1937).
- [16] Felix Boehm and Peter Vogel. *Physics of Massive Neutrinos*, 2nd ed. (Cambridge University Press, (1992).
- [17] Kayser, B. *Are neutrinos their own antiparticles?* J. Phys. Conf. Ser. **173**, (2009).
- [18] Millea, M. *Planck 2018 results. VI. Cosmological parameters*. (2018). arXiv:1807.06209 [astro-ph.CO].
- [19] Peltoniemi, J. *Solar Neutrinos: Experiments*. The Ultimate Neutrino Page, 10 Apr. 2005, cupp.oulu.fi/neutrino/nd-sol2.html
- [20] Cousins, R. D. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D - Part. Fields, Gravit. Cosmol.* **57**, 3873–3889 (1998).
- [21] Killick, R. S. *Observation of Singly Charged Barium Ions in a Buffer Gas: Towards a Functional Barium-Tagging System for Use in the Enriched Xenon Observatory*. Master's Thesis, Carleton University (2015).
- [22] Albert, J. B. *et al. Sensitivity and discovery potential of the proposed nEXO experiment to neutrinoless double- β decay*. 1–23 (2018). arXiv:1710.05075v2 [nucl-ex].
- [23] M. Tanabashi et al. (Particle Data Group), Phys. Rev. D 98, 030001 (2018).
- [24] Gando, A. *et al. Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen*. Phys. Rev. Lett. **117**, 1–6 (2016).
- [25] Alfonso, K. *et al. Search for Neutrinoless Double-Beta Decay of ^{130}Te with CUORE-0*. Phys. Rev. Lett. **115**, 102502 (2015).
- [26] Agostini, M. *et al. Improved Limit on Neutrinoless Double- β Decay of Ge 76 from GERDA Phase II*. Phys. Rev. Lett. **120**, 1–5 (2018).

- [27] Arnold, R. *et al.* Search for neutrinoless double-beta decay of ^{100}Mo with the NEMO-3 detector. Phys. Rev. D **89**, 111101 (2014).
- [28] Arnold, R. *et al.* First Results of the Search for Neutrinoless Double-Beta Decay with the NEMO 3 Detector. Phys. Rev. Lett. **95**, 182302 (2005).
- [29] Albert, J. B. *et al.* Search for Neutrinoless Double-Beta Decay with the Upgraded EXO-200 Detector. Phys. Rev. Lett. **120**, 72701 (2018).
- [30] J. Kotila and F. Iachello. Phase-space factors for double- β decay. Phys. Rev. C **85**, 034316 (2012).
- [31] W.M. Haynes, *Handbook of Chemistry and Physics*, 91 ed. (CRC Press, 2011).
- [32] Maneschg, W. Present status of neutrinoless double beta decay searches. Talk at NuPhys2016 (2017). arXiv:1704.08537v1 [physics.ins-det]
- [33] 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. Phys. Rev. C **91**, 034304 (2015).
- [34] Fedor .imkovic, Vadim Rodin, Amand Faessler, and Petr Vogel. $0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements, quasiparticle random-phase approximation, and isospin symmetry restoration . Phys. Rev.C **87**, 045501 (2013).
- [35] Dobi, A. *et al.* A Xenon Gas Purity Monitor for EXO. (2011). arXiv:1106.1812v1 [physics.ins-det]
- [36] Leport, F. *et al.* A magnetically driven piston pump for ultra- clean applications. Rev. Sci. Instrum. **82**, 105114 (2011).
- [37] M. Auger *et al.*, EXO Collaboration. The EXO-200 detector, part I: detector design and construction. JINST **7**, P05010 (2012).
- [38] Albert, J. B. *et al.* Cosmogenic backgrounds to $0\nu\beta\beta$ and $2\nu\beta\beta$ in EXO-200. J. Cosmol. Astropart. Phys. **4** (2016).

- [39] Albert, J. B. *et al.* *Investigation of radioactivity-induced backgrounds in EXO-200*. Phys. Rev. C **92**, 15503 (2015).
- [40] Albert, J. B. *et al.* *Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector*. Phys. Rev. C **89**, 15502 (2014).
- [41] Conti, E. *et al.* *Correlated fluctuations between luminescence and ionization in liquid xenon*. Phys. Rev. B **68**, 54201 (2003).
- [42] Collaboration, T. E.-200 *et al.* *Search for Majorana neutrinos with the first two years of EXO-200 data*. Nature **510**, 229-234 (2014).
- [43] Delaquis, S. *Construction and operation of the EXO-100 cryogenic facility for R & D in liquid xenon : advances in barium ion tagging*. Doctoral Thesis, Universität Bern (2015).
- [44] Albert, J. B. *et al.* *Measurement of the drift velocity and transverse diffusion of electrons in liquid xenon with the EXO-200 detector*. Phys. Rev. C **95**, 25502 (2017).
- [45] U.S. Department of Energy Office of Environmental Management, *Accident Investigation of the February 5, 2014, Underground Salt Haul Truck Fire at the Waste Isolation Pilot Plant, Carlsbad NM*, [http://energy.gov/ehss/downloads/accident-investigation-february-5-2014-underground-salt-haul-truck-fire-waste \(March 2014\)](http://energy.gov/ehss/downloads/accident-investigation-february-5-2014-underground-salt-haul-truck-fire-waste-(March-2014)), accessed 9th of June 2015.
- [46] U.S. Department of Energy Office of Environmental Management, *Radiological Release Accident Investigation Report - Phase 1 Radiation Report*, [http://energy.gov/em/downloads/radiological-release-accident-investigation-report-phase-1-radiation-report, \(April 2014\)](http://energy.gov/em/downloads/radiological-release-accident-investigation-report-phase-1-radiation-report,(April-2014)), accessed 9th of June 2015.
- [47] SNOLAB, *User's Handbook*, 2nd revision, 26 June 2006
- [48] Aprile, E. & Doke, T. *Liquid xenon detectors for particle physics and astrophysics*. Rev. Mod. Phys. **82**, 2053–2097 (2010).
- [49] W. T. Chen *et al.*, *Measurement of the Transverse Diffusion Coefficient of Charge in Liquid Xenon*, Defect and Diffusion Forum, Vols. 326-328, pp. 567-572, (2012)

- [50] Lockwitz, S. & Jostlein, H. *A study of dielectric breakdown along insulators surrounding conductors in liquid argon*. J. Instrum. **11**, P03026--P03026 (2016).
- [51] Rebel, B. *et al.* *High voltage in noble liquids for high energy physics*. J. Instrum. **9**, T08004--T08004 (2014).
- [52] Baylin, I. E. *Electrical Breakdown of Liquid Dielectrics (A Review)*. Foreign Technology Division, Wright-Paterson Air Force Base, Ohio (1972).
- [53] L. A. Dissado, J. C. Fothergill, S. V. Wolfe and R. M. Hill, *Weibull Statistics in Dielectric Breakdown; Theoretical Basis, Applications and Implications*, IEEE Transactions on Electrical Insulation, Vol. EI-19, **no. 3**, pp. 227-233 (1984).
- [54] T. J. Lewis, *The basic processes of conduction in dielectric liquids*, Proceedings of 1993 IEEE 11th International Conference on Conduction and Breakdown in Dielectric Liquids (ICDL '93), Baden-Dattwil, Switzerland, pp. 32-41 (1993).
- [55] T. J. Lewis, *Basic electrical processes in dielectric liquids*, IEEE Transactions on Dielectrics and Electrical Insulation, vol. 1, **no. 4**, pp. 630-643 (1994).
- [56] Clark, F. M. *The Electrical Breakdown of Dielectric Liquids*. Journal of the Franklin Institute, vol. 216, **no. 4**, 429–458 (1933).
- [57] Dissado, L. A. & Hill, R. M. *Non-exponential decay in dielectrics and dynamics of correlated systems*. Nature **279**, 685–689 (1979).
- [58] Dissado, L. A. *Dielectric Relaxation and the Structure of Condensed Matter*. Phys. Scr. **T1**, 110–114 (1982).
- [59] Dissado, L. A. *Theoretical basis for the statistics of dielectric breakdown*. J. Phys. D. Appl. Phys. **23**, 1582–1591 (1990).
- [60] Hill, R. M. & Dissado, L. A. *Examination of the statistics of dielectric breakdown*. J. Phys. C Solid State Phys. **16**, 4447–4468 (1983).

- [61] Qiu, Z., Ruan, J., Huang, D., Pu, Z. & Shu, S. *A prediction method for breakdown voltage of typical air gaps based on electric field features and support vector machine*. IEEE Trans. Dielectr. Electr. Insul. **22**, 2125–2135 (2015).
- [62] Huang, C. & Ruan, J. *Prediction study on power frequency breakdown voltage of ring structure short air gaps*. ICHVE 2016 - 2016 IEEE Int. Conf. High Volt. Eng. Appl. 6–9 (2016).
- [63] *The Weibull Distribution*. (n.d.). Retrieved March 10, 2018, from http://reliawiki.org/index.php/The_Weibull_Distribution
- [64] Delaquis, S. et al. *Deep neural networks for energy and position reconstruction in EXO-200*. J. Instrum. **13**, P08023–P08023 (2018).
- [65] Cree, W. *Measurement of a new half-lifetime lower limit for neutrinoless double-beta decay in ^{136}Xe using the full EXO-200 phase 1 data set and improvements to barium tagging for nEXO*. Master's Thesis, Carleton University (2017).
- [66] Hastie, T., Rosset, S., Zhu, J. & Zou, H. *Multi-class AdaBoost*. Stat. Interface **2**, 349–360 (2013).
- [67] L. I. Kuncheva, C. J. Whitaker, C. A. Shipp and R. P. W. Duin, *Is independence good for combining classifiers?*. Proceedings 15th International Conference on Pattern Recognition. ICPR-2000, Barcelona, Spain, 2000, pp. 168-171 vol.2.
- [68] Bonab, H. & Can, F. *Less Is More: A Comprehensive Framework for the Number of Components of Ensemble Classifiers*. (2017). arXiv:1709.02925v2 [cs.LG]
- [69] Margineantu, D. D. & Dietterich, T. G. *Pruning Adaptive Boosting*. Proceedings of the Fourteenth International Conference on Machine Learning, Nashville, Tennessee, USA, July 8-12, 1997 (ed. Fisher, D. H.) 211–218 (Morgan Kaufmann, 1997).
- [70] Wolpert, D. H. *Stacked generalization*. Neural Networks **5**, 241–259 (1992).
- [71] Francois, D., Wertz, V. & Verleysen, M. *Non-Euclidean metrics for similarity search in noisy datasets*. Eur. Symp. Artif. Neural Networks, 339–344 (2005).

- [72] Aggarwal C.C., Hinneburg A., Keim D.A. (2001) *On the Surprising Behavior of Distance Metrics in High Dimensional Space*. Van den Bussche J., Vianu V. (eds) Database Theory — ICDT 2001. ICDT 2001. Lecture Notes in Computer Science, vol 1973. Springer, Berlin, Heidelberg
- [73] Efron, B., & Tibshirani, R. (1994). *An introduction to the bootstrap*. New York: Chapman & Hall.
- [74] DiCiccio, T., & Efron, B. (1996). *Bootstrap Confidence Intervals*. Statistical Science, 11(3), 189-212.
- [75] Wang, C., Marriott, P., Pengfei, L., & Li, P. (2018). *Asymptotic coverage probabilities of bootstrap percentile confidence intervals for constrained parameters*. arXiv:1712.02469 [math.ST]
- [76] M. K. Moe, *Detection of neutrinoless double-beta decay*, Phys. Rev. C **44** (1991) R931–R934.
- [77] Danilov, M. et al. (the EXO Collaboration), *Detection of very small neutrino masses in double-beta decay using laser tagging*, Phys. Lett. B 480, 12–18 (2000).
- [78] Dolinski, M. J. *The Enriched Xenon Observatory: EXO-200 and Ba⁺ tagging*. Nucl. Phys. B - Proc. Suppl. **229–232**, 124–127 (2012).
- [79] Mong, B, et al. (the EXO Collaboration), *Spectroscopy of Ba and Ba⁺ in solide xenon for barium tagging in nEXO*, Phys. Rev. A 91, 022505 (2015).
- [80] Twelker, K. et al (the EXO Collaboration), *An apparatus to manipulate and identify individual Ba ion from bulk liquid Xe*, Rev. Sci. Instrum. **85**, 095114 (2014); arXiv:1407.0618
- [81] Rowson, P.C. SLAC. *Stanford Linear Accelerator Center, Cryogenic Ion Capture Probe R&D at / SLACPUB-17110*, (2017).
- [83] Jeng, S.-C., Fairbank, W. M. & Miyajima, M. *Measurements of the mobility of alkaline earth ions in liquid xenon*. J. Phys. D. Appl. Phys. **42**, 35302 (2008).
- [84] Sinclair, D. et al. (the EXO Collaboration), *Prospects for Barium Tagging in Gaseous Xenon*, J. Phys: Conf. Ser. 309, 012005 (2011).

- [85] Brunner, T. *et al.* *An RF-only ion-funnel for extraction from high-pressure gases*. *Int. J. Mass Spectrom.* **379**, 110–120 (2015).
- [86] Querci, L., Varentsov, V., Günther, D. & Hattendorf, B. *An Rf-only ion funnel interface for ion cooling in laser ablation time of flight mass spectrometry*. *Spectrochim. Acta Part B At. Spectrosc.* **146**, 57–68 (2018).
- [87] Rollin, E. *Barium Ion Extraction and Identification from Laser Induced Florescence in Gas for the Enriched Xenon Observatory*. Doctoral Thesis, Carleton University (2015).
- [88] Chambers, C. *et al.* *Imaging individual Ba atoms in solid xenon for barium tagging in nEXO*. (2018). arXiv:1806.10694v2 [physics.ins-det]
- [89] B.M. Davis, J.G. McCaffrey, *Imaging individual barium atoms in solid xenon for barium tagging in nEXO*. *J. Chem. Phys.* **144**, 1044308 (2016).
- [90] B.M. Davis, B. Gervais, J.G. McCaffrey, *An investigation of the sites occupied by atomic barium in solid xenon – A 2D-EE luminescence spectroscopy and molecular dynamics study*. *J. Chem. Phys.* **148**, 124308 (2018).
- [91] Ellis, C. D. & Wooster, W. A., *Continuous Spectra of β -Rays*, *Nature (London)* **119**, 563-564 (1927).
- [92] scikit-learn developers. (2018). *sklearn.svm.SVC*. Retrieved May 5, 2019 from, <https://scikit-learn.org/stable/modules/generated/sklearn.svm.SVC.html#sklearn.svm.SVC>
- [93] Collaboration, T. E.-200 *et al.* *Search for Majorana neutrinos with the first two years of EXO-200 data*. *Nature* **510**, 229 (2014).
- [94] The SciPy community. (February 10, 2019). *scipy.optimize.minimize*. Retrieved May 5, 2019 from <https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html>
- [95] Cowan, G. *Statistical Data Analysis*. (Clarendon Press, 1998).