

**Sensitivity of XENONnT to the $2\nu\text{EC}\beta^+$ Decay of
 ^{124}Xe and R&D Towards XLZD, the Next
Generation Liquid Xenon Observatory**

DISSERTATION

zur

Erlangung der naturwissenschaftlichen Doktorwürde
(Dr. sc. nat.)

vorgelegt der
Mathematisch-naturwissenschaftlichen Fakultät

der
UNIVERSITÄT ZÜRICH

von
ANDREA PALOMA CIMENTAL CHÁVEZ

aus
Mexiko

PROMOTIONSKOMMISSION
Prof. Dr. Laura Baudis (Vorsitz)
Prof. Dr. Florencia Canelli
Dr. Diego Ramírez García
Dr. Christian Wittweg

Zürich, 2026

A Ramiro.

*A quienes, sin dejar de pensar, trabajar y cuidar,
insisten en su autonomía.*

*No basta adaptarnos a una sociedad que
cambia en la superficie y permanece idéntica en la raíz.
No basta imitar los modelos que se nos proponen
y que son la respuesta a otras circunstancias que las nuestras.
No basta siquiera descubrir lo que somos.
Hay que inventarnos.*

— Rosario Castellanos, *El eterno femenino*.

ABSTRACT

Astrophysical and cosmological observations provide compelling evidence for an additional form of matter, known as dark matter, which constitutes approximately 84% of the total matter content of the Universe. Among the leading candidates to explain its fundamental nature are weakly interacting massive particles (WIMPs). Despite extensive experimental efforts over the past decades, including collider searches, as well as direct and indirect detection experiments, dark matter has yet to be observed in a laboratory experiment.

Located deep underground at the Laboratori Nazionali del Gran Sasso, XENONnT is a fourth-generation direct detection experiment employing a dual-phase xenon time projection chamber. The detector is instrumented with 494 Hamamatsu photomultiplier tubes to detect light signals produced by particle interactions within 5.9 tonne of liquid xenon.

The first part of this thesis presents several aspects of XENONnT, including a description of the facility and the performance of its photomultiplier tubes, which are continuously monitored and modeled using in-situ calibration data. Data from the first and second science runs are used to probe both low-energy electron recoil (ER) and nuclear recoil (NR) regions. No excess of low-energy ER events is observed, in contrast to previous results from XENON1T. The measured ER background rate, $15.8 \text{ events}/(\text{t} \cdot \text{y} \cdot \text{keV})$, is the lowest achieved to date in liquid xenon time projection chambers. Similarly, no significant excess is found in the WIMP search, leading to new upper limits on the WIMP–nucleon cross section, reaching a minimum of

$1.7 \times 10^{-47} \text{ cm}^2$ at a WIMP mass of $30 \text{ GeV}/c^2$. Additionally, the first measurement of solar neutrinos via coherent elastic neutrino–nucleus scattering in a liquid xenon detector is reported.

The thesis then addresses the characterization of the detector response in the MeV energy range, supported by a sophisticated simulation framework. Three key developments are introduced: data-driven yield models, a physics-motivated clustering algorithm for energy deposits, and a simulation-driven correction map for accurate reconstruction of multi-scatter signals. The resulting energy bias and resolution are studied, and a dedicated reconstruction algorithm for multi-scatter topologies is developed and validated, establishing the basis for MeV-scale searches in a detector traditionally optimized for low energies.

This characterization is applied to a sensitivity study of the decay of ^{124}Xe to the ground state of ^{124}Te via electron capture with positron emission and two neutrinos. With its full active mass, XENONnT is expected to yield approximately 137 ± 52 decay events per year. The feasibility of detecting this process is evaluated using full-chain simulations of signal and background. Dedicated background suppression techniques are developed to mitigate dominant gamma-ray backgrounds, including physics-driven selections, fiducial volume optimization, and a machine learning classifier. Despite these improvements, the search remains background-dominated, allowing only a lower limit on the half-life to be set. A binned profile likelihood fit yields a limit of 5.35×10^{20} years at 90 % confidence level.

The second part of this thesis focuses on the future liquid xenon observatory XLZD, in particular the full-scale vertical demonstrator developed at the University of Zurich. This setup aims to demonstrate electron drift over 2.6 m in liquid xenon. First results from Xenoscope on electron transport properties, measured with a purity monitor, are presented. The upgrade to a dual-phase detector is then described, including the design of a high-voltage feed-through capable of delivering 50 kV. Finally, results from the commissioning run are discussed.

Preface

This thesis documents my contributions to the XENON and XLZD dark matter programs, during my doctoral studies under the supervision of Prof. Dr. Laura Baudis. While the work reflects the collective effort of large collaborations, I explicitly highlight my individual contributions in each chapter below:

- Chapter 3: outlines a review of the XENONnT TPC photomultiplier tubes. My contributions include the continuous monitoring of the per-PMT gain stability and the development of a time-dependent gain model, together with its implementation as a correction in the XeDocs database, in collaboration with Luis Sanchez. This work corresponds to the second science run and builds on a modeling framework developed by Dr. Alexander Bismark. In addition, together with Mariana Rajado, I led a dedicated correction campaign following the identification of a failure in the calibration system. This effort involved the full reprocessing of all PMT calibration data from the second science run and the production of an updated gain model, as discussed in Section 3.5. These results were subsequently used in the WIMP and CE ν NS analyses presented in Chapter 4. Finally, the monitoring of the time evolution of the lone-hit area during the second science run is also part of my work.
- Chapter 5: presents the characterization of the detector response at the MeV energy scale. This chapter builds on the work of several collaborators. Dr. Diego Ramírez García and Henning Eißing Schultze refactored the previous simulation package into the modular **fuse** framework [1], used extensively through this thesis. Carlo Fuselli and Dr. Micha Weiss developed the data-driven yield models and the lineage clustering routines. My contributions focus on the validation and application of this framework. All simulation-data comparisons presented in this chapter are my work. I developed the simulation-driven correction map for the S1 signal, including the production of the underlying simulations, its implementation, and validation across multiple energies. The studies of energy resolution and energy bias are entirely my work. Finally, the development, validation and

implementation of the `multipeak` reconstruction algorithm were carried out in collaboration with Dr. Diego Ramírez García.

- Chapter 6: presents the sensitivity study of the ^{124}Xe decay via two-neutrino electron capture with positron emission. The majority of the work presented in this chapter is my own. I performed the full signal simulation chain, from DECAY0 generation to waveform level, and used these results to characterize the multi-scatter topology of the signal in the XENONnT TPC. I carried out the simulations of internal background components. I was also involved in the launch of a large-scale simulation campaign of detector materials, together with Saad El Morabit, Carlo Fuselli, Henning Eißing Schultze, and Dr. Diego Ramírez García. A significant part of my contribution consisted in enabling the use of these simulations for the analysis. I developed a dedicated algorithm to propagate primary-level information from Geant4 to event-level within the template builder framework developed by Dr. Maxime Pierre. I then carried out the large-scale transfer of the simulation dataset, comprising tens of hundreds of files, from the OSG infrastructure to the Midway cluster, and reprocessed the entire dataset with this algorithm. This step was essential for making the simulations usable in the analysis. I co-led the validation of these simulations together with Saad El Morabit and Ekaterina Ekoslova. Following this, I used the processed simulations to train a machine learning classifier for signal–background discrimination, to optimize the fiducial volume selection, and to construct the final background model. I then performed the statistical inference using the `alea` framework [2], deriving a lower limit on the half-life of the decay. Throughout this work, Dr. Diego Ramírez García was a close collaborator, who contributed to the development and refinement of the analysis strategy.
- Chapter 7: presents the Xenoscope facility for R&D. My contributions focus on the design, construction, and testing of the high-voltage feedthrough, developed in close collaboration with Dr. Frédéric Girard. I was responsible for its maintenance and for its integration during the 2.6 m TPC phase of Xenoscope. I also actively contributed to the upgrade from the purity monitor phase to the 2.6 m TPC commissioning phase. In this context, I produced the coated photocathode used in this phase and carried out its installation. The results of this commissioning run are reported in Ref. [3].

Table of Contents

List of Figures	VII
List of Tables	VIII
1 Dark Matter and Neutrinos: the Physics of Weakly Interacting Particles	1
1.1 Dark matter	2
1.1.1 Evidence	2
1.1.2 Dark matter candidates	4
1.1.3 Particle dark matter detection	6
1.1.4 Current status of WIMP direct detection	9
1.2 Neutrinos	10
1.2.1 Coherent elastic neutrino-nucleus scattering	11
1.2.2 Double weak decays	12
1.2.3 Detection of double weak decays: experimental considerations and status	14
2 The XENON Dark Matter Project	18
2.1 Detection principle	18
2.1.1 Liquid xenon as radiation detection medium	18
2.1.2 Signal generation in xenon	21
2.1.3 Signal detection and reconstruction	24
2.1.4 Background sources	28
2.2 The XENONnT experiment	31
2.2.1 From XENON10 to XENONnT	31
2.2.2 Dual-phase time projection chamber	32
2.2.3 Neutron veto	37
2.2.4 Muon veto	38
2.2.5 Calibration system	38
2.3 The next generation of detectors	40

3	Light Detection in the XENONnT TPC	41
3.1	The XENONnT photomultiplier tube	41
3.2	Photosensor calibration	45
3.3	PMT gain	46
3.3.1	Model-independent single photoelectron analysis	47
3.3.2	Time-dependent gain model	50
3.3.3	Gain implementation as a correction	52
3.4	Single photoelectron acceptance	53
3.5	Performance during science run one	55
3.5.1	Gain stability and LED timing effects	56
3.5.2	Lone hit evolution	59
3.6	Summary and outlook	62
4	Science Results from the XENONnT Experiment	63
4.1	Science data	63
4.2	Event reconstruction	65
4.3	Signal corrections	66
4.4	Event selection	70
4.5	Nuclear recoil dark matter results	72
4.6	^8B solar neutrino detection	76
4.7	Low energy ER searches	79
4.8	Summary and outlook	81
5	Modeling the XENONnT Response in the MeV Energy Range	82
5.1	The XENONnT simulations pipeline	83
5.2	Refined emission model for the MeV range	87
5.2.1	Lineage clustering	87
5.2.2	Data-driven electronic recoil yields	88
5.3	Detector response validation in the observable space	90
5.3.1	Data quality selection	91
5.3.2	S1 and S2 shape properties	94
5.4	Energy reconstruction for multi-scatter events	98
5.4.1	S1 relative light yield correction	99
5.4.2	Multi-scatter energy resolution and bias	103
5.5	Multiple S2-peak event reconstruction	107
5.6	Summary and outlook	109

6	Sensitivity Projection for the $2\nu\text{EC}\beta^+$ decay of ^{124}Xe	110
6.1	Signature of $2\nu\text{EC}\beta^+$ decay in a xenon dual-phase TPC	111
6.2	Analysis strategy	113
6.2.1	Signal simulation	114
6.2.2	Reconstructed energy selection	116
6.2.3	Signal distribution in the observable space	117
6.3	Internal backgrounds	118
6.4	Materials backgrounds	120
6.4.1	Rate conversion	121
6.4.2	SR0 materials simulation campaign	122
6.5	Background rejection	124
6.5.1	Fiducial volume selection	125
6.5.2	Machine learning discriminator	129
6.5.3	Signal and background estimates	137
6.6	Inference	138
6.6.1	Projected sensitivity	139
6.7	Summary and Outlook	141
7	Xenoscope: an R&D Platform for XLZD	144
7.1	The XLZD project	144
7.2	The Xenoscope facility	146
7.2.1	Cryogenics and xenon recirculation	149
7.2.2	Xenon recuperation and storage	150
7.2.3	Facility monitoring	152
7.3	Purity monitor	153
7.4	Dual-phase time projection chamber	156
7.4.1	High-voltage delivery system	158
7.4.2	Electron generation and light detection	164
7.5	Commissioning results	166
7.6	Next steps	171
A	XENONnT Material screening results	175
B	SR0 Material background simulation campaign	177
C	Xenoscope gas system	179

List of Figures

1.1	Galaxy cluster 1E 0657–56 also known as the “Bullet Cluster”.	2
1.2	Galaxy rotation curve and CMB angular power spectrum	3
1.3	Mass range of dark matter candidates.	5
1.4	WIMP–nucleon scattering rates and target dependence.	8
1.5	Current upper limits at 90 % C.L. on the SI WIMP-nucleon cross sections.	9
1.6	Mass parabolas for nuclear isobars.	13
2.1	Mass attenuation coefficient as a function of the photon energy in xenon.	19
2.2	Schematic of signal production in liquid xenon.	22
2.3	Light and charge yields for nuclear and electronic recoils in liquid xenon.	23
2.4	Detection principle of a xenon dual-phase TPC.	26
2.5	Simulated nuclear recoil and electronic recoil distributions.	27
2.6	The decay chains of ^{238}U and ^{232}Th	29
2.7	Evolution of the XENON dark matter detectors.	31
2.8	Photograph of the XENONnT experiment at LNGS.	33
2.9	CAD rendering of the XENONnT cryostat and TPC structure.	34
2.10	CAD rendering of the XENONnT experiment.	37
3.1	Photograph of a Hamamatsu R11410-21 photomultiplier tube and its working principle.	42
3.2	Schematic of the PMT calibration system used in the XENONnT TPC.	45
3.3	PMT waveforms recorded during a calibration and distribution of LED arrival times.	46
3.4	Integrated charge spectra and occupancy for an exemplary PMT.	47
3.5	Time evolution of the gain for four representative PMTs.	51
3.6	LED calibration spectra and SPE response of an exemplary PMT.	54
3.7	Example waveforms during the LED pulse arrival time shift.	54
3.8	Gain and occupancy as a function of the integration window end position.	56
3.9	Time evolution of relative gain differences.	57
3.10	Time evolution of the final median gains in SR1.	58

3.11	Relative standard deviation of measured and modeled PMT gains. . . .	59
3.12	Spatial distribution of the modeled PMT gain standard deviation. . . .	60
3.13	Time evolution of the median lone-hit area throughout SR1.	61
4.1	Accumulated livetime during SR0 and SR1.	64
4.2	Schematic of the XENONnT event reconstruction pipeline.	66
4.3	S2 single-scatter selection.	73
4.4	Detection efficiencies as a function of NR energy.	74
4.5	Dark matter search data in the (cS1, cS2) space.	75
4.6	Upper limits on the spin-independent WIMP–nucleon cross section. . .	75
4.7	Detection efficiency for low-energy NRs in XENONnT and the expected ^8B CE ν NS spectrum.	77
4.8	Measurement of the solar ^8B neutrino flux.	79
4.9	Comparison between the XENON1T and XENONnT low-energy ER spectra.	80
5.1	Schematic overview of the XENONnT simulation chain.	84
5.2	Illustration of the lineage clustering algorithm.	87
5.3	S1 width selections applied to ^{232}Th data.	91
5.4	S1 area fraction top selection applied to ^{232}Th calibration data. . . .	91
5.5	Schematic representation of an S2 signal.	92
5.6	Position distribution of ^{232}Th simulation and calibration data.	93
5.7	Comparison of ^{232}Th calibration data and simulations in the observable space.	94
5.8	S1 shape properties: data vs simulations.	95
5.9	S2 shape properties: data vs simulations.	96
5.10	S2 signal shape parameters vs z: data vs simulation.	97
5.11	S2 multiplicity: data vs simulation.	97
5.12	Spatial distribution of the energy bias in SR0 simulations.	99
5.13	TPC binning scheme and event distribution.	100
5.14	S1 area distribution in an exemplary spatial bin.	101
5.15	Difference between the simulation-driven and data-driven light yield maps. .	102
5.16	Energy bias vs R^2 after simulation-driven S1 correction.	103
5.17	Fit to ^{208}Tl energy spectrum with residuals.	104
5.18	Simulation-driven energy resolution curve.	105
5.19	Simulation-driven relative energy bias of multi-scatter events as a function of energy.	106
5.20	<i>Multipeak</i> S2 parameters: data vs simulations.	107
5.21	Ratio between $S2_0/S2_2$: data vs simulations.	109

6.1	Illustration of the two-neutrino electron capture with positron emission decay.	111
6.2	Simulation of the energy depositions produced by the $2\nu\text{EC}\beta^+$ decay in the XENONnT TPC.	112
6.3	DECAY0 simulation of the $2\nu\text{EC}\beta^+$ kinetic energy spectrum.	114
6.4	Signal waveforms for an example simulated event.	115
6.5	$2\nu\text{EC}\beta^+$ S2 peak multiplicity and S2 area distribution.	116
6.6	Simulated total reconstructed energy of $2\nu\text{EC}\beta^+$ decay.	117
6.7	Simulated signal distribution in the observable space.	118
6.8	Multi-scatter energy spectra of radiogenic background.	123
6.9	Relative contributions of individual detector materials to the total background.	124
6.10	Distributions of intrinsic and material background in the observable space.	125
6.11	Simulation-driven FoM.	127
6.12	FoM iso-contours and global FoM vs xenon mass.	128
6.13	Spatial distribution of the material background with the optimized FV.	129
6.14	Flowchart of the supervised learning pipeline.	130
6.15	BDT classifier training features.	132
6.16	Relative importance of the input features used in the BDT training.	134
6.17	Distribution of the BDT output score.	135
6.18	Asimov significance versus BDT score.	136
6.19	Final background model in the reconstructed energy space.	138
6.20	Expected lower limit on the $T_{1/2}$ of the ^{124}Xe decay via $2\nu\text{EC}\beta^+$	140
6.21	Profile likelihood scan as a function of the signal strength.	141
7.1	CAD rendering of the XLZD baseline design.	145
7.2	CAD rendering of the Xenoscope facility at UZH.	147
7.3	CAD rendering of the heat exchangers and cooling tower of Xenoscope.	149
7.4	The Xenoscope phases.	153
7.5	Electron drift lifetime measured in Xenoscope.	154
7.6	Drift velocity measured with the Xenoscope purity monitor.	155
7.7	Exploded view of the cathode high-voltage delivery system.	160
7.8	Air-to-vacuum feedthrough components.	161
7.9	SNeval components of the high-voltage delivery system during installation.	162
7.10	Cold-pressure test of the high-voltage feedthrough.	163
7.11	Photograph of the photocathode installed in the cathode plate and the SiPM array.	164
7.12	Operating conditions during the commissioning run.	167

7.13	Evolution of the drift and extraction fields during the commissioning run.	169
7.14	Mean gain of each SiPM tile during the commissioning run.	170
7.15	Workflow of the channel-trigger data acquisition mode.	171
7.16	Waveforms and hit pattern of a selected muon event.	172
C.1	Piping and Instrumentation Diagram of the Xenoscope facility.	180

List of Tables

1.1	Summary of double weak decay modes in xenon isotopes.	15
1.2	Decay modes of ^{124}Xe	17
2.1	Properties of xenon relevant for use in particle detectors.	20
3.1	PMTs excluded from the SR0 and SR1 analyses.	55
4.1	Overview of selection criteria for the WIMP, CE ν NS or low-energy ER searches.	71
4.2	Expected and best-fit values for the SR0 + SR1 WIMP search.	76
4.3	Expected and best-fit results of the CE ν NS search.	78
4.4	Expected and best-fit values for the low-energy ER search.	80
6.1	Preselection cumulative acceptances for signal and background.	125
6.2	Signal and background acceptances after applying the FV and the BDT cut selections.	137
6.3	Expected $T_{1/2}$ lower limits of the 2 ν EC β^+ decay of ^{124}Xe for different exposure scenarios.	140
7.1	Electrostatic field configuration of the Xenoscope electrodes.	168
A.1	Material contamination from screening campaign.	176
B.1	Summary of the SR0 material background simulation campaign.	178

Chapter 1

Dark Matter and Neutrinos: the Physics of Weakly Interacting Particles

In astroparticle physics, two central topics of current research are neutrinos and dark matter. Both are characterized by their elusive nature and are therefore referred to as weakly interacting particles. Neutrinos are fundamental particles within the Standard Model of particle physics, the theoretical framework describing fundamental particles and their interactions [4]. Over the past decades, substantial progress has been made in determining their properties and in understanding their role in astrophysics and cosmology. Dark matter, in contrast, is not accounted for within the Standard Model. Its existence is inferred from a wide range of astrophysical and cosmological observations [5–7], yet it has not been directly detected in laboratory experiments.

This chapter provides a brief introduction of dark matter and neutrinos, with emphasis on processes relevant to rare-event searches in xenon-based detectors. The evidence for dark matter is reviewed in Section 1.1.1, followed by a discussion of leading particle candidates in Section 1.1.2, focusing on weakly interacting massive particles (WIMPs). Detection approaches are introduced in Section 1.1.3, and the current experimental status is summarized in Section 1.1.4. The chapter then turns to neutrinos: Section 1.2 outlines the established framework before introducing two processes relevant for dark matter detectors, coherent elastic neutrino–nucleus scattering (CE ν NS) in Section 1.2.1 and double weak decays in Section 1.2.2. Finally, Section 1.2.3 summarizes the current experimental status of double weak decays in xenon.




Figure 1.1: Galaxy cluster 1E 0657–56 also known as the “Bullet Cluster”. Composite image showing the distribution of different components: galaxies observed in visible light (white), hot gas observed in X-rays (pink), and the total mass distribution inferred from gravitational lensing (blue). Figure from Ref. [9].

1.1 Dark matter

Several astrophysical and cosmological observations provide compelling evidence for the existence of an additional matter component in the Universe, which we refer to as dark matter. It constitutes a central component of the standard model of cosmology, Λ CDM model [8], which consists of a cosmological constant Λ that drives the accelerated expansion of the Universe, and cold dark matter (CDM), a non-relativistic matter component responsible for structure formation. In the following, the observational evidence for dark matter is briefly reviewed, starting from early astrophysical observations and progressing to measurements at larger cosmological scales.

1.1.1 Evidence

In the early 1930s, the Swiss astronomer Fritz Zwicky studied the Coma Cluster and measured the velocity dispersion of its galaxies [10]. Using the Virial theorem, he concluded that the observed velocities exceeded the values expected from the gravitational potential of the visible matter alone. To account for this discrepancy, he proposed the existence of an unseen mass component, which he referred to as “dark matter” [10].

At present, gravitational lensing provides a powerful tool for measuring the total mass of galaxy clusters by observing the deflection of light from distant luminous objects located behind the cluster along the line of sight. [11]. In contrast, the baryonic matter is traced through electromagnetic telescopes, most prominently via X-ray emission from the hot cluster gas [12]. Merging galaxy clusters provide an ideal laboratory for comparing baryonic and total mass distributions. A notable example is shown in Figure 1.1, where the total mass distribution of the 1E 0657-56 merging cluster [9], inferred from gravitational lensing, is shown in blue and overlaid with the hot gas distribution in pink. During the cluster collision, the gas component is slowed and becomes spatially separated from the collisionless components. In contrast, the lensing-derived mass




Figure 1.2: Left: Rotation velocity distribution of NGC 1705 galaxy as a function of the distance to the galactic center. The figure shows the contributions from the stellar disk (blue) and the gas (gray). Data from Ref. [18]. **Right:** Angular power spectrum of the CMB temperature anisotropies, showing the amplitude of fluctuations as a function of angular scale. The peak position encodes information on the baryonic, dark matter, and dark energy content of the Universe. Adapted from Ref. [19].

distribution follows the galaxies and is spatially offset from the gas. This indicates the presence of a dominant, effectively collisionless mass component not accounted for by the observed baryonic matter, consistent with the additional mass component inferred from earlier dynamical studies. Analyses of 72 similar merging clusters report a 7.6 σ significance for the presence of a collisionless mass component consistent with dark matter [13].

At galactic scales, evidence for dark matter arises from the radial velocity profiles of galaxies, as pioneered by Vera Rubin and Kent Ford in the 1970s [14, 15]. According to Newtonian dynamics, the rotation velocity $v(r)$ of stars at a distance r from the galactic center is given by:

$$v(r) = \sqrt{\frac{GM(r)}{r}}, \quad (1.1)$$

where $M(r)$ is the enclosed mass and G is the gravitational constant. Therefore, we expect the rotational velocity to decrease as $r^{-1/2}$ at large radii. Contrary to this expectation, observed rotation curves exhibit a flat behavior at large radii, as shown in Figure 1.2 left. This discrepancy can be explained by introducing the hypothesis that galaxies are embedded in extended, non-luminous matter distributions, known as dark matter halos [5]. Subsequent studies of galaxy rotation curves have consistently supported this conclusion [16, 17].

Further evidence for dark matter arises at cosmological scales through observations of the Cosmic Microwave Background (CMB). Discovered in 1965 [20], the CMB is relic electromagnetic radiation from the early Universe that originated approximately

380 000 years after the Big Bang. Specifically, it was produced during the recombination epoch, when the Universe cooled to about 3000 K, allowing free electrons and protons to combine into neutral hydrogen atoms. As a result, photons decoupled from matter and began to propagate freely. Since then, the expansion of the Universe has redshifted this radiation to a nearly perfect blackbody spectrum with a present temperature of (2.7250 ± 0.0006) K [21], corresponding to microwave wavelengths. While the CMB is remarkably uniform, it exhibits small temperature anisotropies, $\frac{\Delta T}{T}$, at the level of $\sim 10^{-5}$ ($\sim 30 \mu\text{K}$) [22], which provide a powerful probe of the matter density at the time of recombination and of the subsequent formation of large-scale structures in the Universe. The temperature anisotropies of the CMB are described using spherical harmonics, yielding an angular power spectrum that quantifies the amplitude of fluctuations as a function of their angular size on the sky. Figure 1.2 right shows the CMB temperature power spectrum. The characteristic peak heights and their spacing are sensitive to the relative energy densities of baryonic matter, dark matter, and dark energy [8, 23]. Measurements of these anisotropies by the WMAP and the Planck satellite [19, 24, 25] are interpreted within the six-parameter ΛCDM framework. By fitting the predicted angular power spectrum to the observed data, one obtains best-fit cosmological parameters of $\Omega_b = 0.049$ for baryonic matter, $\Omega_c = 0.26$ for cold dark matter, and $\Omega_\Lambda = 0.69$ for dark energy [24, 25], where the Ω_i denote the present-day energy density fractions relative to the critical density, $\rho_{\text{crit}} = 3H_0^2/(8\pi G)$ [8], corresponding to the density required for a spatially flat Universe.

The Standard Model of particle physics does not include a suitable dark matter candidate. Although neutrinos are electrically neutral and weakly interacting, they are too light and relativistic during structure formation, thus, they behave as hot dark matter and cannot account for the observed clustering of matter on small scales. Thus, it is assumed that dark matter must be composed of new particles that have not yet been directly detected in laboratory experiments. The next section provides an overview of selected candidates.

1.1.2 Dark matter candidates

The observations indicating the presence of a non-luminous mass component provide strong constraints on its physical properties. A viable dark matter candidate is expected to satisfy the following properties:

- **Cold:** Non-relativistic, or cold, at the time structures start to form in order to be consistent with the observed large-scale structure of the Universe. Relativistic particles suppress small-scale structure formation, as their large velocities allow them to free-stream out of dense regions.




Figure 1.3: Possible mass range of dark matter candidates spanning more than 90 orders of magnitude, together with the most researched dark matter particle candidates as distributed at their predicted mass range. Figure from Ref. [26]

- **Stable:** It must be stable, or at least long-lived on timescales comparable to the age of the Universe (~ 13.8 Gyr) [24], to account for its presence from the early Universe to the present day.
- **Neutral:** It must be electrically neutral, or have a sufficiently small electric charge (millicharge), to avoid energy loss through electromagnetic radiation, and thus be consistent with the CMB observations.
- **Massive:** It must possess non-zero mass to account for the observed gravitational effects attributed to dark matter.
- **Weakly interacting:** It must interact only very weakly with ordinary matter and exhibit negligible self-interactions, consistent with observations of merging galaxy clusters.

Dark matter candidates satisfying these criteria span an enormous mass range of 90 orders of magnitude, as shown in Figure 1.3, including fields, particles and macroscopic objects [26]. Some of the most prominent candidates include:

- **Primordial black holes (PBHs):** Macroscopic dark matter candidates consisting of hypothesized black holes formed in the early Universe from the collapse of large density fluctuations in the total energy density. Their abundance is tightly constrained by a variety of observations, including gravitational lensing and dynamical effects. PBHs with stellar or larger masses are strongly constrained and can only constitute a subdominant fraction, while very light PBHs would have evaporated via Hawking radiation by the present epoch. However, PBHs in the asteroid-mass range remain a viable candidate, as they are challenging to probe observationally [5, 27].
- **Sterile neutrinos:** Hypothetical fermions extending the Standard Model and accounting for neutrino masses. Unlike Standard Model neutrinos, which are

relativistic and too light to explain the observed dark matter density, keV-scale sterile neutrinos can be non-relativistic at the time of structure formation, making them viable dark matter candidates [6, 28].

- **Axions or Axion Like Particles (ALPS):** Axions are hypothetical light particles originally proposed to solve the strong CP problem in quantum chromodynamics [29]. ALPs are a broader class of light particles, not necessarily tied to the strong CP problem. Observations of stellar cooling place strong constraints on their interaction strengths, while cosmological and galactic observations constrain their masses. Ultra-light masses ($\leq 10^{-22}$ eV) are disfavored because they would suppress small-scale structure formation. Both axions and ALPs remain viable dark matter candidates [6, 30].
- **WIMPs:** A hypothesized generic class of electrically neutral particles interacting at the weak scale [5, 31], with masses in the range $1 \text{ GeV}/c^2 - 100 \text{ TeV}/c^2$ and interaction cross sections from 10^{-41} – 10^{-51} cm^2 [32]. These candidates are constrained by direct detection experiments, indirect searches, colliders, and the cosmological observations discussed in Section 1.1.1. Their relic abundance is assumed to be produced by the thermal freeze-out mechanism [33]. Being non-relativistic at the time of structure formation, they constitute a compelling candidate for cold dark matter.

Among these candidates, WIMPs are particularly well-motivated, as they arise naturally in several extensions of the Standard Model, such as supersymmetry [5] and theories with extra dimensions [5, 34]. Moreover, the annihilation cross section required to reproduce the relic abundance inferred from cosmological observations is of the order of the weak interaction scale a coincidence often referred to as the “WIMP miracle” [35]. This provides a framework in which the abundance at present date and interaction strength are correlated, implying that WIMPs may scatter off nuclei at rates accessible to current experiments. The following discussion focuses on WIMPs, including their detection strategies and current experimental status.

1.1.3 Particle dark matter detection

Three complementary dark matter search strategies exist with the aim of investigating the properties of dark matter particles, such as their mass and interaction cross section with baryonic matter. These are:

- **Collider searches:** These approaches aim to produce dark matter particles in high-energy collisions of Standard Model particles, such as at the Large Hadron Collider [36], where their presence is inferred from missing transverse energy and

momentum signatures.

- **Indirect detection:** This approach seeks the Standard Model particles produced by dark matter annihilation or decay in regions of high dark matter density, such as the galactic center or dwarf galaxies. Experiments include gamma-ray telescopes [37], neutrino observatories [38], and cosmic-ray detectors [39], which look for excesses over expected astrophysical backgrounds.
- **Direct detection:** These experiments aim to measure the energy deposited by dark matter particles scattering off nuclei (or electrons) of Standard Model matter in ultra-low background detectors.

In the following, the focus is on direct detection and key considerations relevant to detection technologies. The feasibility of directly detecting WIMPs in terrestrial detectors was first proposed in the 1980s [40]. If WIMPs couple, even weakly, to Standard Model particles, the interactions responsible for their annihilation in the early Universe also lead to scattering off nuclei. Given the non-relativistic velocities of galactic dark matter (220 km/s), these interactions are expected to occur predominantly as elastic scattering processes. The expected scattering rate off a target nucleus of mass m_N is given by [41]:

$$\frac{dR}{dE}(E, t) = \frac{\rho_\chi M}{m_N m_\chi} \int_{v_{\min}}^{v_{\text{esc}}} \frac{d\sigma}{dE}(v, E) v f_E(\vec{v}, t) d^3\vec{v}, \quad (1.2)$$

where M is the target mass of the detector, m_χ the dark matter mass, $f_E(\vec{v}, t)$ denotes the velocity distribution of dark matter particles in the frame of the Earth [42], and ρ_χ is the local dark matter density ($0.3 \text{ GeV}/\text{cm}^3$) [43], where the latter are modeled by the Standard Halo Model [44]. The quantity $v_{\min} = \sqrt{\frac{m_N E}{2\mu^2}}$ represents the minimum dark matter speed required to produce a nuclear recoil of energy E , where $\mu = \frac{m_\chi m_N}{m_\chi + m_N}$ is the reduced mass of the WIMP-nucleus system. The quantity $v_{\text{esc}} = 544 \text{ km/s}$ [45] represents the maximum velocity the dark matter can have and still be gravitationally bound to the Milky Way. The recoil energy spectrum is obtained assuming a Maxwell-Boltzmann velocity distribution for dark matter particles in the galactic halo. The resulting exponential dependence of the differential rate on the recoil energy is given by [41]:

$$\frac{dR}{dE} \propto \exp\left(-\frac{E}{E_0} \frac{4m_\chi m_N}{(m_\chi + m_N)^2}\right), \quad (1.3)$$

where E_0 is the most probable kinetic energy of the incident WIMP. For elastic WIMP-nucleus interactions, two main types of couplings are typically considered: spin-independent (SI) and spin-dependent (SD). The former corresponds to coherent scatter-




Figure 1.4: Left: SI WIMP–nucleon scattering rates for a WIMP mass of $m_\chi = 10 \text{ GeV}/c^2$ with cross-section $\sigma = 10^{-49} \text{ cm}^2$ (dark blue), and for $m_\chi = 200 \text{ GeV}/c^2$ with $\sigma = 10^{-45} \text{ cm}^2$ (light green). **Right:** SI WIMP–nucleon scattering rates in common target materials. Dashed lines correspond to $m_\chi = 10 \text{ GeV}/c^2$, while solid lines correspond to $m_\chi = 200 \text{ GeV}/c^2$. Figures generated using [46].

ing off the entire nucleus, where the WIMP interacts with all nucleons in phase, while the latter arises from the coupling to the nuclear spin, such that only unpaired nucleons contribute. The differential cross section for WIMP–nucleus scattering, $\frac{d\sigma}{dE}(v, E)$, is given by [41]:

$$\frac{d\sigma}{dE}(v, E) = \frac{m_N}{2\mu^2 v^2} (\sigma_{\text{SI}} F_{\text{SI}}^2(E) + \sigma_{\text{SD}} F_{\text{SD}}^2(E)), \quad (1.4)$$

where σ_{SI} and σ_{SD} denote the SI and SD cross sections, respectively, and F_{SI}^2 and F_{SD}^2 are the corresponding nuclear form factors. For SI interactions, a commonly used choice for the form factor is the Helm approximation [47], and the SI cross section is given by [41]:

$$\sigma_{\text{SI}} = \sigma_n \left(\frac{\mu}{\mu_n} \right)^2 A^2, \quad (1.5)$$

where μ_n is the reduced mass of the WIMP–nucleon system and A is the number of nucleons. The A^2 dependence implies that heavier target nuclei yield higher event rates. The expected SI recoil energy spectra for different WIMP masses and cross sections are shown in Figure 1.4 left, while the event rates for various target nuclei are shown in Figure 1.4 right, illustrating the dependence on A^2 . In particular, the spectra for heavier targets are shifted toward lower recoil energies compared to lighter targets. At the same time, heavier target nuclei lead to lower recoil energies, such that detectors using heavy targets must achieve sufficiently low energy thresholds to detect the resulting recoils.




Figure 1.5: Current upper limits at 90 % C.L. on the SI WIMP-nucleon cross sections. The experiments are direct detection searches, including CRESST-III [49], DarkSide-50 [50], Panda-X-4T [51], DEAP-3600 [52], SuperCDMS [53, 54], XENONnT [55], LZ [56], and the projected sensitivities of XLZD [57]. The xenon neutrino fog is also shown [58]. Figure generated using [59].

1.1.4 Current status of WIMP direct detection

To achieve sensitivity to such small SI cross sections, an optimal WIMP detector should combine a large target mass M , a high atomic mass number A , a low energy threshold and ultra-low background conditions. Several detector designs have been developed to exploit different target materials for dark matter detection. Among these, noble liquid detectors provide excellent performance [41, 48]. In particular, the xenon-based time projection chamber technology will be discussed in detail in Chapter 2.

For a given WIMP mass and interaction cross section, the expected number of events in a detector can be predicted by integrating the differential rate (Equation 1.2) above the detection threshold and multiplying by the exposure of the experiment, defined as the product of the active detector mass and the live time. If WIMPs with certain properties existed (those discussed in 1.1.2), they would produce a measurable number of events. The absence of such events therefore allows these regions of the cross-section-mass parameter space to be excluded.

The most stringent current bounds on SI WIMP–nucleus scattering cross sections for various target materials are shown in Figure 1.5, together with the irreducible neutrino background, often referred to as the “neutrino fog” [58], briefly discussed in Section 1.2.1. The latest XENONnT results are discussed in Chapter 4.

1.2 Neutrinos

The neutrino was first proposed by Wolfgang Pauli in 1930 [60] to account for missing energy, momentum, and angular momentum in nuclear beta decays. Enrico Fermi subsequently developed a theory of beta decay incorporating this particle and introduced the name “neutrino”. It was experimentally confirmed in 1956 through the detection of reactor antineutrinos via inverse beta decay, $\bar{\nu}_e + p \rightarrow e^+ + n$ [61]. A second neutrino flavor, the muon neutrino ν_μ , was discovered in the 1960s [62], while solar neutrino measurements in the late 1960s revealed a deficit with respect to theoretical predictions, a discrepancy known as the solar neutrino problem [63]. The third flavor, ν_τ , was experimentally observed in 2000 [64], by identifying τ particles produced in charged-current neutrino interactions.

Neutrino oscillations, i.e. the conversion between neutrino flavors, were proposed in the late 1950s [65, 66] and experimentally established in the late 1990s [67]. These observations demonstrated that neutrinos have nonzero masses and that flavor and mass eigenstates are not identical. Flavor states (ν_e, ν_μ, ν_τ) are superpositions of mass eigenstates (ν_1, ν_2, ν_3), which can be expressed as [68]:

$$|\nu_\alpha\rangle = \sum_{j=1}^3 U_{\alpha j} |\nu_j\rangle, \quad \alpha = e, \mu, \tau, \quad (1.6)$$

where $U_{\alpha j}$ are the elements of the Pontecorvo–Maki–Nakagawa–Sakata matrix [66], parametrized by three mixing angles and a CP-violating phase. Experimental observations of solar, atmospheric, reactor, and accelerator neutrinos have established this three-flavor mixing framework [69–74]. In particular, solar neutrino data are well described by the Mikheyev–Smirnov–Wolfenstein (MSW) effect [75, 76], a phenomenon arising from neutrino propagation in matter that enhances flavor conversion and explains the observed deficit of ν_e relative to standard model solar predictions.

Oscillation experiments measure two distinct mass-squared differences, $\Delta m_{ij}^2 = m_i^2 - m_j^2$, and are sensitive only to differences between neutrino masses, leaving the absolute neutrino mass scale undetermined. This defines two characteristic mass scales and allows for two possible mass orderings: normal and inverted [77, 78]. The measured mass differences exhibit a hierarchical pattern, with one significantly larger than the other, while the mixing structure is characterized by two large mixing angles and one smaller angle. The MSW effect further determines the sign of Δm_{21}^2 , implying $m_2 > m_1$ [75, 76].

Despite this progress, several fundamental properties of neutrinos remain unknown. In particular, the absolute neutrino mass scale has not yet been determined [78, 79],

and it is still unclear whether neutrinos are Dirac or Majorana particles [78, 80]. In addition, the origin of their small masses compared to other fermions is still not understood [81, 82]. The search for processes such as double weak decays provide a complementary experimental approach to address these open questions, as will be discussed in Section 1.2.2.

At the same time, neutrinos play an important role in rare-event searches, where their interactions with matter at low energies, in particular through coherent elastic neutrino–nucleus scattering, can give rise to measurable nuclear recoils $\mathcal{O}(\text{keV})$. In direct dark matter detection experiments, these interactions can produce signals indistinguishable from those expected from WIMPs, constituting an irreducible background [58].

1.2.1 Coherent elastic neutrino-nucleus scattering

Predicted in the 1970s, CEνNS is a neutral-current process in which a neutrino scatters elastically off an entire nucleus via Z^0 exchange [83]:

$$\nu_\alpha + (A, Z) \rightarrow \nu_\alpha + (A, Z), \quad (1.7)$$

where $\alpha = e, \mu, \tau$. The interaction becomes coherent when the momentum transfer is small compared to the inverse nuclear size, such that the neutrino effectively couples to the nucleus as a whole. In this regime, the differential cross section is given by [84]:

$$\frac{d\sigma}{dT} = \frac{G_F^2 M}{4\pi} Q_W^2 \left(1 - \frac{MT}{2E_\nu^2}\right) F^2(q^2), \quad (1.8)$$

where G_F is the Fermi constant, M is the nuclear mass, T is the nuclear recoil energy, E_ν is the incident neutrino energy, and $F(q^2)$ is the nuclear form factor, which accounts for the loss of coherence at large momentum transfer. The quantity:

$$Q_W = N - (1 - 4\sin^2 \theta_W)Z \quad (1.9)$$

is the weak nuclear charge, with N and Z the neutron and proton numbers, respectively, and θ_W the Weinberg angle. Since $\sin^2 \theta_W \approx 0.23$, the factor $(1 - 4\sin^2 \theta_W)$ strongly suppresses the proton contribution to Q_W . As a result, the weak charge is dominated by the neutron number, $Q_W \simeq N$, and the CEνNS cross section scales approximately proportional to N^2 . As a result, CEνNS has a relatively large cross section compared to other neutrino interaction channels at low energies [84, 85]. The recoil energies involved are very small, typically at the keV scale or below for neutrino energies of a few tens of MeV. This makes the experimental detection of CEνNS challenging and requires detectors with low energy thresholds and ultra-low backgrounds [85].

The process was first experimentally observed in 2017 using neutrinos from a spallation source [86], and has since been measured with different detector technologies and neutrino sources [87–90]. Neutrinos of astrophysical origin, in particular, neutrinos produced in the ^8B decay chain in the Sun have energies of up to ~ 15 MeV [91], making them particularly relevant for CE ν NS searches in dark matter detectors.

In liquid xenon detectors, the relatively large neutron number of xenon enhances the interaction cross section, leading to a non-negligible event rate. CE ν NS interactions produce low-energy recoils that are experimentally indistinguishable from those expected from WIMP scattering. For this reason, CE ν NS constitutes both a signal channel and an irreducible background for dark matter searches [92], defining a fundamental sensitivity limit for current and future experiments, known as the “neutrino fog” [58].

1.2.2 Double weak decays

Double weak decays, such as double β -decay first proposed by Maria Goeppert-Mayer in the 1930s [93], are nuclear isobaric transitions in which a parent nucleus (A, Z) transforms into a daughter nucleus $(A, Z \pm 2)$ via the simultaneous conversion of two nucleons [94]. These processes occur when no energetically allowed states can be reached through a single β -decay [95]. Their characteristic half-lives typically range from 10^{17} – 10^{25} years [78]. The process is illustrated in Figure 1.6, based on the semi-empirical Bethe–Weizsäcker approximation for nuclear masses [96]. In this framework, an additional contribution, the pairing term, accounts for the tendency of nucleons to form pairs with opposite spins [94, 96]. Consequently, for a fixed mass number A , the intermediate nucleus $(A, Z \pm 1)$ can have a larger mass than the initial and final nuclei. This makes single β -decay energetically forbidden, while allowing the double β transition. Depending on the nucleus, either two neutrons are converted into two protons or vice versa. The energy release, or Q -value, is given by the mass difference between the initial and final nuclei:

$$Q_{\beta\beta} = [M(A, Z) - M(A, Z \pm 2)] \cdot c^2. \quad (1.10)$$

Due to charge conservation, either two electrons ($\beta^-\beta^-$) or two positrons ($\beta^+\beta^+$) are emitted. Alternatively, one or both positron emissions can be replaced by the capture of atomic electrons by the nucleus, referred to as electron capture (EC) [97–99]. Four different modes of double β -decay are allowed within the Standard Model [99]:




Figure 1.6: Mass parabolas for nuclear isobars, following the semi-empirical Bethe–Weizsäcker formula. Due to the pairing term in the formula, even–even nuclei have lower masses than odd–odd nuclei. Thus, a single β -decay from (a) to (b) or from (e) to (d) is not energetically possible, while the direct transition from (a) to (c) or (e) to (c) is allowed. Adapted from Ref. [94].

$$\begin{aligned}
 2\nu\beta^-\beta^- : \quad & (A, Z) \rightarrow (A, Z+2) + 2e^- + 2\bar{\nu}_e \\
 2\nu\beta^+\beta^+ : \quad & (A, Z) \rightarrow (A, Z-2) + 2e^+ + 2\nu_e \\
 2\nu\text{EC}\beta^+ : \quad & (A, Z) + e^- \rightarrow (A, Z-2) + e^+ + 2\nu_e \\
 2\nu\text{ECEC} : \quad & (A, Z) + 2e^- \rightarrow (A, Z-2) + 2\nu_e.
 \end{aligned} \tag{1.11}$$

There is also the hypothesis that these transitions may occur in a neutrinoless mode. The absence of neutrinos in the final state fundamentally changes the physics interpretation, as in the Standard Model lepton number is conserved in all interactions. To date, no experiment has observed neutrinoless double beta decay. However, its discovery would provide insight into the fundamental nature of the neutrino: in particular, it would imply that neutrinos are their own antiparticles (so-called Majorana particles) [100].

The half-life of the two-neutrino double beta decay is given by [95]:

$$T_{1/2}^{2\nu} = (G_{2\nu} |M_{2\nu}|^2)^{-1}, \tag{1.12}$$

where $G_{2\nu}$ is the phase-space factor, and $M_{2\nu}$ is the nuclear matrix element (NME), which encodes the nuclear transition amplitude between the initial and final states. Precise measurements of two-neutrino double beta decay provide important tests of nuclear structure models and benchmarks for the calculation of NMEs [68, 95]. These are essential inputs for interpreting searches for neutrinoless double beta decay.

Similarly, for the neutrinoless mode, the half-life is given by [95]:

$$T_{1/2}^{0\nu} = (G_{0\nu} |M_{0\nu}|^2 \langle m_\nu \rangle^2)^{-1}, \quad (1.13)$$

where $G_{0\nu}$ is the phase-space factor, $M_{0\nu}$ is the NME, and $\langle m_\nu \rangle$ the effective Majorana neutrino mass. A detailed calculation of the phase-space factors can be found in [101].

The effective neutrino mass can be extracted from the measured half-life, provided that $G_{0\nu}$ and $M_{0\nu}$ are known. It is defined as [68]:

$$\langle m_\nu \rangle = \left| \sum_{i=1}^3 U_{ei}^2 m_i \right|, \quad (1.14)$$

and depends on the neutrino mass eigenvalues, mixing parameters, and Majorana phases. The dominant theoretical uncertainty arises from the NME, whose calculation is challenging [102]. Several nuclear structure approaches are employed, including the Interacting Shell Model (ISM) [103–105], the Quasiparticle Random Phase Approximation (QRPA) [106], and the Interacting Boson Model (IBM) [107]. A comprehensive review of these methods can be found in [102].

The observation of neutrinoless double beta decay would provide direct evidence of lepton number violation and establish the Majorana nature of neutrinos, with important implications for the origin of neutrino masses and for understanding the matter–antimatter asymmetry of the Universe [80] by a process known as leptogenesis [108].

1.2.3 Detection of double weak decays: experimental considerations and status

In general, the expected number of detected decays in a given experiment is determined by the detector exposure and the decay half-life. It can be expressed as:

$$N = \ln(2) \cdot \frac{\eta \cdot \epsilon \cdot m \cdot t \cdot N_A}{m_u \cdot T_{1/2}}, \quad (1.15)$$

where η is the isotopic abundance of the target nuclei, ϵ the detection efficiency, m the detector mass, t the live time, N_A is Avogadro’s number, and m_u the molar mass of the target isotope. This expression highlights that the observation of processes with extremely long half-lives requires large exposures and sufficiently high isotopic abundances of the target nucleus.

In addition, the sensitivity of an experiment is determined by the level of background in

Isotope	Abundance [%] [113]	Q-value [keV]	Decay		Half-life	
			mode	Measured [yr]	Predicted [yr]	
^{124}Xe	0.095(5)	2856.73(12) [114]	2νECEC	$(1.1 \pm 0.2_{\text{st}} \pm 0.1_{\text{sy}}) \times 10^{22}$ [115]		
			0νECEC		2.1×10^{32} [116]	
			2νECβ ⁺		$(1.3 \pm 0.5) \times 10^{23}$ [116]	
			0νECβ ⁺		8.4×10^{27} [116]	
			2νβ ⁺ β ⁺		$(2.2 \pm 0.7) \times 10^{28}$ [116]	
			0νβ ⁺ β ⁺		1.5×10^{29} [116]	
^{126}Xe	0.089(3)	918.3(3.4) [117]	2νECEC	$> 1.9 \times 10^{22}$ [118]		
			0νECEC			
^{134}Xe	10.436(35)	824.1(0.3) [117]	2νβ ⁻ β ⁻	$> 8.7 \times 10^{20}$ [119]		
			0νβ ⁻ β ⁻		$> 1.1 \times 10^{23}$ [119]	
^{136}Xe	8.857(72)	2458.7(0.6) [120]	2νβ ⁻ β ⁻	$(2.16 \pm 0.01_{\text{st}} \pm 0.06_{\text{sy}}) \times 10^{21}$ [109]		
			0νβ ⁻ β ⁻		$> 3.8 \times 10^{26}$ [121]	

Table 1.1: Summary of double weak decay modes in xenon isotopes. The predicted lower limits on half-lives are given at 90% C.L. in the last column.

the energy region of interest. The scaling of the half-life sensitivity can be approximated as:

$$T_{1/2} \propto \begin{cases} \eta \cdot m \cdot \epsilon \cdot t & \text{background-free} \\ \eta \cdot \epsilon \sqrt{\frac{m \cdot t}{B \cdot \Delta E}} & \text{background-limited} \end{cases}, \quad (1.16)$$

where B is the background rate and ΔE defines the energy window used in the analysis. The definition of the energy window ΔE depends on the nature of the decay under study. In the neutrinoless case, the signal corresponds to a mono-energetic peak in the energy spectrum centered at the Q -value, and ΔE is directly determined by the detector energy resolution [109–111]. In contrast, for two-neutrino modes the signal is the sum of the emitted electron energies, which forms a continuous spectrum due to the presence of neutrinos. As a consequence, the region of interest is not uniquely defined by the detector resolution, but depends on the specific analysis strategy, including the choice of energy window and background modeling (see Chapter 6). Altogether, these considerations place stringent requirements on detector technologies, favoring large target masses, low backgrounds, and good energy resolution.

Among the 35 nuclei that can undergo double weak decays [112], xenon is particularly well suited to meet these requirements. In particular, liquid xenon detectors developed for dark matter searches provide large target masses, ultra-low backgrounds, and excellent energy resolution, making them powerful instruments for detecting double weak decays.

Double weak decay modes are energetically allowed in four xenon isotopes: ^{124}Xe and ^{126}Xe , which can undergo decay via electron capture and positron emission channels, and ^{134}Xe and ^{136}Xe , which primarily undergo double β^- decay. Two-neutrino decay modes to the ground state have been experimentally observed for ^{136}Xe [109, 110] and, more recently, the double electron capture mode in ^{124}Xe [115]. The latter was measured using a liquid xenon dark matter detector, demonstrating the potential of this technology for studying rare nuclear processes. For ^{126}Xe [118] and ^{134}Xe [119], no decay has been observed so far, and only lower limits on the corresponding half-lives have been established. The abundances of the four isotopes in natural xenon are shown in Table 1.1, together with their Q -values, decay modes, and corresponding measured or predicted half-lives. As shown in Table 1.1, the current leading experiments on measurements of ^{136}Xe double beta decay or the neutrinoless mode are past experiments: KamLAND-Zen [121] and EXO-200 [119], which are both enriched in ^{136}Xe .

A dedicated study on second-order weak processes in xenon-based TPCs [116] provides predictions for the half-lives of various decay modes of ^{124}Xe , including both two-neutrino and neutrinoless channels. Given its Q -value of $(2856.73 \pm 0.12) \text{ keV}$ [114], ^{124}Xe allows for six possible decay modes: $0\nu/2\nu$ double electron capture (ECEC), $0\nu/2\nu$ electron capture with positron emission ($\text{EC}\beta^+$), and $0\nu/2\nu$ double positron emission ($\beta^+\beta^+$). While the $2\nu\text{ECEC}$ decay has been observed by XENON1T [115] and recently by XENONnT [122], LZ [123] and PandaX [124], the other channels remain unobserved.

The half-lives of the ^{124}Xe decay modes shown in Table 1.1 are taken from Ref. [116], which were computed from Equation 1.12. In particular, the half-life of the $2\nu\text{EC}\beta^+$ channel is updated using the most recent measurement of the $2\nu\text{ECEC}$ half-life from Ref. [125]. Assuming that the NMEs of the $2\nu\text{ECEC}$ and $2\nu\text{EC}\beta^+$ channels are equal to first approximation, the measured $2\nu\text{ECEC}$ half-life can be used to constrain $M_{2\nu}$ and to infer the corresponding $2\nu\text{EC}\beta^+$ half-life through the ratio of phase-space factors [116]:

$$T_{1/2}^{2\nu\text{EC}\beta^+} = \frac{G_{2\nu\text{ECEC}}}{G_{2\nu\text{EC}\beta^+}} T_{1/2}^{2\nu\text{ECEC}}, \quad (1.17)$$

with $G_{2\nu\text{ECEC}} = (1.5 - 2.0) \cdot 10^{20} \text{ yr}$ and $G_{2\nu\text{EC}\beta^+} = (1.2 - 1.7) \cdot 10^{21} \text{ yr}$ [116]. This approximation is motivated by the fact that the expected difference between the NMEs of the two channels is small compared to the current experimental uncertainty on the measured $2\nu\text{ECEC}$ half-life [116].

Table 1.2 summarizes the emitted quanta in each decay mode together with the expected coincident signature. In general, decay modes involving electron capture are accompanied by atomic de-excitation products, such as X-rays and Auger electrons,

Decay mode	Emitted quanta	Coincidence
2νECEC	X-rays/ e_{Auger} , $2e^-$, 2ν	X-rays/ e_{Auger}
2νECβ ⁺	X-rays/ e_{Auger} , e^+ , 2ν	X-rays/ e_{Auger} , e^+ + 2γ
2νβ ⁺ β ⁺	$2e^+$, 2ν	$2e^+$ + 4γ
0νECβ ⁺	X-rays/ e_{Auger} , e^+	X-rays/ e_{Auger} , e^+ + 2γ
0νβ ⁺ β ⁺	$2e^+$	$2e^+$ + 4γ
0νECEC	X-rays/ e_{Auger} , 2 - 3γ	X-rays/ e_{Auger} and 3γ

Table 1.2: Different decay modes of ^{124}Xe with their corresponding coincidence signatures. The 2νECEC mode is the only one that has been already measured [115]. Adapted from Ref. [116].

resulting from the rearrangement of the atomic shell. When combined with β^+ emission, the produced positron subsequently annihilates, giving rise to two 511 keV gamma rays. As a result, these channels are characterized by distinctive coincidence signatures involving combinations of atomic de-excitation quanta and annihilation photons, which can be exploited for background suppression.

Among the available decay modes of ^{124}Xe , the 2νECβ⁺ channel is particularly promising for detection in xenon-based detectors. This is due to its relatively shorter predicted half-life compared to the other unobserved channels, as well as the natural abundance of ^{124}Xe , with approximately 1 kg of ^{124}Xe per tonne of natural xenon. In addition, xenon detectors offer $\mathcal{O}(1\%)$ relative energy resolution at its Q -value [126] and provide handles for background suppression through event topology and spatial information. The study of this decay mode also provides complementary information to the more commonly explored double β^- decay channels, contributing to the validation of NME calculations and improving the understanding of double weak processes. A detailed assessment of the sensitivity to the 2νECβ⁺ decay mode in a specific experimental configuration, namely in the XENONnT detector, is presented in Chapter 6, with the goal of quantifying the experimental reach of current-generation liquid xenon detectors for this process.

Chapter 2

The XENON Dark Matter Project

Dual-phase xenon time projection chambers (TPCs) have established themselves as the leading technology in direct dark matter searches in the mass range of $\sim 3 \text{ GeV}/c^2$ to $10 \text{ TeV}/c^2$ [127, 128] for more than a decade. In this context, the XENON dark matter project was conceived as a phased program, designed to achieve progressively higher sensitivity through a series of increasingly large dual-phase TPCs with improved background suppression, each building upon the expertise gained from its predecessors.

In this chapter, the fundamental detection principles of dual-phase xenon TPCs are introduced in Section 2.1 together with a review of the relevant microphysics properties of liquid xenon that enable this technology. Building on this foundation, Section 2.2 presents an overview of the XENON dark matter project, including a description of the current-generation detector and its main infrastructure.

2.1 Detection principle

Like other dual-phase liquid xenon time projection chambers, the XENON dark matter project exploits the intrinsic properties of LXe that make it a powerful medium for rare-event searches. This section discusses the properties relevant for the practical aspects of the detector and reviews the processes by which energy depositions in the LXe lead to measurable signals. The discussion follows the full chain from quanta production to their propagation, detection, and reconstruction.

2.1.1 Liquid xenon as radiation detection medium

Xenon is a noble gas that is naturally present in the atmosphere in trace amounts at $\sim 0.087 \text{ ppm}$ [129]. Industrially, it is obtained as a byproduct of air separation plants,




Figure 2.1: Mass attenuation coefficient as a function of the photon energy in xenon for photoelectric effect (solid purple), Compton scattering (solid orange) and pair production (solid violet). The total attenuation coefficient is given by the solid yellow line. Adapted from the National Institute of Standards and Technology (NIST) [132].

where oxygen and nitrogen are first extracted through cryogenic distillation and xenon is subsequently isolated from the residual gases.

The high atomic number of xenon ($Z = 54$) contributes to strong self-shielding against external radiation. This is reflected in the photon interaction cross sections, which scale approximately as $Z^5/E_\gamma^{7/2}$ for photoabsorption, Z/E_γ for Compton scattering, and $Z^2 \log(2E_\gamma)$ for pair production [130]. As a result, radiation is rapidly attenuated in liquid xenon, as illustrated in Figure 2.1, which shows the photon absorption coefficient over a wide energy range. For reference, the intensity of a 2.614 MeV gamma-ray is attenuated by a factor of e after traveling ~ 9 cm in liquid xenon. In addition, its large atomic mass and high liquid density (2.854 g/cm^3 at 177 K and 2 bar [131]) yield a high stopping power (dE/dx) for charged particles, which is exploited for particle identification based on the dependence of the signal size on the ionization density produced. Table 2.1 lists some of the atomic, thermodynamic, and optical properties of xenon relevant for detector design.

Liquid xenon is an efficient scintillator with a short scintillation timescale. For MeV electrons, the scintillation yield of liquid xenon is approximately 80 % of that of NaI(Tl), a widely used inorganic crystal scintillator [48]. Additionally, the mean energy to produce an electron-ion pair in liquid xenon is 15.6 eV [133], corresponding to an ionization yield of 64 pairs/keV, which is significantly higher than in liquid argon and gas xenon [134, 135].

Property	Value	Unit	Notes	Reference
Atomic				
Atomic number, Z	54			[138]
Atomic mass	131.293	u	Valid for natural xenon	[138]
Natural isotopes (abundance)	^{124}Xe (0.095)	%	$2\nu\text{ECEC}$, $T_{1/2} = 1.18(19) \cdot 10^{22}$ y	[139]
	^{126}Xe (0.089)	%	$2\nu\beta\beta$, unknown	[139]
	^{128}Xe : (1.910)	%	stable	[139]
	^{129}Xe : (26.401)	%	stable	[139]
	^{130}Xe : (4.071)	%	stable	[139]
	^{131}Xe : (21.232)	%	stable	[139]
	^{132}Xe : (26.909)	%	stable	[139]
	^{134}Xe : (10.436)	%	$2\nu\beta\beta$, $T_{1/2} > 2.8 \cdot 10^{22}$ y	[139]
	^{136}Xe : (8.857)	%	$2\nu\beta\beta$, $T_{1/2} = 2.17(6) \cdot 10^{21}$ y	[139]
Thermodynamic				
Density	2.854	g/cm^3	at 177 K, 2 bar	[131]
Boiling point	177.88	K	at 2 bar	[138]
Melting point	161.41	K	at 2 bar	[138]
Triple point	(161.38 ± 0.02)	K		[140]
	(0.8160 ± 0.0001)	bar		[140]
Critical point	289.7	K		[138]
	58.4	bar		[138]
Optical				
Scintillation wavelength	175	nm	at 168 K	[135, 141]
Refractive index	(1.69 ± 0.02)		at 589 nm	[138]
Attenuation length	(36.4 ± 1.8)	cm	at the triple point	[138]

Table 2.1: Atomic, thermodynamic, and optical properties of xenon relevant for use in particle detectors.

Xenon liquefies at a relatively high temperature of 177 K at 2 bar [136] compared to argon at 87.302 K [137], reducing the required cooling power and making cryogenic operation and handling more accessible. Dual-phase LXe TPCs are typically operated at temperatures of 177 K, which correspond to a xenon vapor pressure of 2 bar [136], where liquid and gas phases coexist in equilibrium. This thermodynamic regime enables the formation of a liquid–gas interface, which is an essential feature of dual-phase LXe TPCs, as outlined in Section 2.1.3.


Additionally, xenon contains several naturally occurring isotopes. Two isotopes, however, are of particular physics interest: ^{136}Xe , which enables searches for both two-neutrino and neutrinoless double beta decay ($2\nu\beta\beta$ and $0\nu\beta\beta$) [125], and ^{124}Xe , which undergoes a variety of double weak decay modes [116]. These processes are among the main rare event targets for current and next generation LXe TPCs [57, 142, 143], as

discussed in Section 1.2.2. In particular, the sensitivity of XENONnT to the decay of ^{124}Xe to the ground state of ^{124}Te via two-neutrino electron capture with positron emission ($2\nu\text{EC}\beta^+$) will be investigated in Chapter 6.

2.1.2 Signal generation in xenon

When an incoming particle interacts with liquid xenon, its energy is partitioned into three channels: excitation, ionization, and atomic motion (non-radiative heat). Only the first two channels give rise to measurable signals in a dual-phase TPC, the third channel is not directly observable. The relative contribution of each channel depends on the type of interaction. In LXe TPCs, the two main types of interactions are nuclear recoils (NR), when the incoming particle transfers momentum to the xenon nucleus as a whole, and electronic recoils (ER), which result from scattering off atomic electrons. ERs transfer most of their energy into excitation and ionization, while NRs dissipate a larger fraction into heat [144], reducing the detectable signal.

Upon initial recoil, excited atoms (or excitons, Xe^*) rapidly form excimers (Xe_2^*) which decay to the ground state through the emission of vacuum ultraviolet (VUV) photons with a characteristic wavelength of 175 nm [135, 141, 145]. Excimers are short-lived diatomic molecules formed when an excited xenon atom binds to a neighboring ground-state atom. They are stable only in the excited state and decay radiatively, producing the 175 nm scintillation photons characteristic of liquid xenon. The overall mechanism is described by [135]:



These excimers can de-excite through either a singlet or a triplet state, with lifetimes of 3 ns and 27 ns respectively [146]. Since the excimer molecule is only bound in the excited state and has no bound ground state, the emitted VUV photons cannot be reabsorbed. As a result, liquid xenon is effectively transparent to its own scintillation light. The occurrence ratio of singlet to triplet state varies between nuclear and electronic recoils, with nuclear recoils favoring the singlet state and electronic recoils favoring the triplet state [146]. In liquid xenon the ratio of lifetimes is too small to provide strong pulse-shape discrimination. On the other hand, in liquid argon, the singlet and triplet lifetimes are well separated, with values of 7 ns and 1.6 μs , respectively [135]. This enables efficient pulse-shape discrimination, typically achieving an ER rejection of 10^8 at 50 % NR acceptance [147].





Figure 2.2: Schematic of signal production in liquid xenon. An energy deposition produces ionization and excitation in proportions that differ for electronic and nuclear recoils. Excited xenon atoms form excimers, as described by Equation 2.1, which decay and emit VUV scintillation photons at 175 nm. Ionization produces electron–ion pairs, a fraction of the electrons recombine with ions, creating additional excitons and thus contributing further to scintillation, according to Equation 2.2. The remaining free electrons escape recombination [135].

In parallel, the ionization process creates electron-ion pairs. In the absence of an electric field, these electrons recombine with nearby xenon ions Xe^+ , producing additional excitons. These excitons form excimers, which in turn undergo the same process as the ones formed by direct excitation, as described in Equation 2.1. This process can be summarized by the following reactions [135]:



Since the drift field reduces the probability of recombination, the primary scintillation yield decreases while the number of free electrons available for extraction increases. This process gives rise to the anti-correlation between scintillation and ionization in liquid xenon. Higher fields enhance the electron yield at the expense of photon yield, whereas lower fields favor recombination and thus scintillation as shown in Figure 2.3. In practice, the observable quanta are the de-excitation photons and the electrons that escape recombination. The number of de-excitation photons and the number of electrons is




Figure 2.3: Light and charge yields for nuclear recoils (**left**) and electronic recoils (**right**) in liquid xenon as a function of recoil energy. Results are shown for several drift-field strengths, indicated by the color scale. Obtained using the Nobel Element Simulation Technique (NEST) v2.4.0 [148].

given by [149]:

$$N_\gamma = \left(\frac{N_{ex}}{N_i} + r \right) \cdot N_i, \quad (2.3a)$$

$$N_e = (1 - r) \cdot N_i \quad (2.3b)$$

where N_{ex} and N_i are the numbers of excitons and electron-ion pairs, respectively. Since the probability of recombination r fluctuates from event to event, the partition between N_γ and N_e also fluctuates. These recombination fluctuations introduce an irreducible variance in the scintillation and ionization channels when considered separately. However, the total number of observable quanta is conserved:

$$N_q = N_\gamma + N_e, \quad (2.4)$$

which depends only on the deposited energy and the mean energy required to produce a single observable quantum, known as the W -value. In liquid xenon, the W -value is commonly taken to be $W \simeq 13.7 \text{ eV}$ [141]. Recent measurements, however, indicate

smaller values [150, 151], leading to a discrepancy of up to $\sim 20\%$. Further independent studies are required to resolve this tension and establish a consistent reference value. This motivates the definition of a combined energy scale in terms of the total quanta,

$$E = W \cdot (N_\gamma + N_e), \quad (2.5)$$

where information from both scintillation and ionization channels is exploited. By summing the two contributions, recombination fluctuations largely cancel, yielding a precise estimate of the energy deposited. The expected yield of quanta varies with the recoil type and energy, and can be controlled by applying an external electric field, as shown in Figure 2.3. It is therefore convenient to define the photon yield P_y , as the number of scintillation photons produced per unit recoil energy, and the electron yield Q_y as the number of ionization electrons produced per unit recoil energy [135]:

$$P_y = \frac{\langle N_\gamma \rangle}{E}, \quad Q_y = \frac{\langle N_{e^-} \rangle}{E}. \quad (2.6)$$

The higher stopping power of nuclear recoils produces denser ionization tracks [135], which increases the probability of electron-ion recombination. As a consequence, NR interactions yield relatively more scintillation photons and fewer ionization electrons than ER interactions of the same energy. This leads to a systematically smaller charge-to-light ratio, Q_y/P_y , for NR interactions. The difference in this ratio, provides the primary handle for particle-type identification in liquid xenon TPCs [152, 153].

2.1.3 Signal detection and reconstruction

A dual-phase liquid xenon time projection chamber consists of a cylindrical active target volume monitored by photosensor arrays (see Chapter 3) located above and below the interaction region. The active volume is surrounded by highly reflective panels to maximize VUV light collection. Above the liquid, a thin gaseous xenon (GXe) layer with a thickness of $\mathcal{O}(\text{mm})$ is maintained. A cathode at the bottom of the TPC and two electrodes near the liquid-gas interface (gate and anode) establish the electric fields required to drift ionization electrons upward through the LXe and to extract them into the GXe, where they produce proportional scintillation light. A schematic representation of the detection principle is shown in Figure 2.4.

The photons and electrons described in the previous section give rise to the S1 and S2 signals, which constitute the core observables of LXe TPCs. The S1 signal is generated when VUV scintillation photons propagate through the liquid xenon. Their transport is governed by Rayleigh scattering, with a characteristic scattering length of $\sim 30\text{ cm}$ [154], before they are detected by the photosensor arrays. The intrinsic time profile of the

scintillation is determined by the decay of singlet and triplet excimer states, leading to a characteristic emission timescale of $\mathcal{O}(\text{ns})$ [135]. In practice, the detected S1 signal is further broadened by the photosensor response, as will be discussed in Chapter 3.

Electrons produced at the interaction site that do not recombine in the presence of the electric field, drift vertically toward the liquid–gas interface. The number of electrons reaching the interface is governed by the drift field, the extraction field, and the xenon purity via electron attachment during the drift. In LXe, the electron drift velocity, v_d , increases with field and saturates around $2\text{ mm}/\mu\text{s}$ for fields above $1\text{ kV}/\text{cm}$ [135, 155]. The drift velocity is also inversely proportional to the temperature and depends only weakly on the pressure conditions via the density of LXe [135]. During the drift, the electron cloud undergoes diffusion, characterized by a longitudinal component along the drift field direction (z), with increasing drift field suppressing this broadening, and transverse to the field, in the (x, y) plane [135]. The longitudinal $\sigma_{L,z}$ and transverse $\sigma_{T,R}$ spread of the electron cloud diffusion are correlated with the drift velocity in the following way [156]:

$$\sigma_{L,z} = \sqrt{\frac{2D_L \cdot t_d}{v_d^2}}, \quad (2.7)$$

$$\sigma_{T,R} = \sqrt{\frac{2D_T \cdot t_d}{v_d^2}}, \quad (2.8)$$

where t_d is the drift time, D_L is the longitudinal diffusion coefficient and D_T the transverse diffusion coefficient. In LXe, D_L is typically an order of magnitude smaller than D_T [135]. Reducing the longitudinal spread of the drifting electron cloud at higher drift fields improves the z -resolution and therefore the separation of nearby interaction sites along the drift direction [116].

The presence of electronegative impurities in the liquid xenon such as O_2 , H_2O and N_2O can strongly affect the survival probability of drifting electrons. These molecules attach electrons, reducing their number and consequently suppressing the S2 signal. The attachment probability depends mainly on the impurity concentration, the attachment rate constant, and the magnitude and uniformity of the applied drift field. These considerations highlight the requirement of maintaining high xenon purity ($\mathcal{O}(\text{ppb})$ [135]) and applying strong drift fields. Impurities may undergo photoionization if their ionization potential is below the energy of LXe scintillation photons [135]. This process can produce delayed S2 signals or charge tails that are not associated with the primary interaction, thereby complicating event reconstruction through additional pulses, pileup, and time-dependent variations in detector response [153, 156].




Figure 2.4: Detection principle of a xenon dual-phase TPC. When a particle recoils off a xenon atom, prompt scintillation light (S1) and ionization electrons are created. These electrons are drifted upward under the influence of an electric field established between the cathode and gate electrodes. At the interface, the electrons are extracted into the gas phase by an extraction field applied between the gate and anode electrodes. The extracted electrons produce a delayed proportional scintillation signal (S2) via electroluminescence. Both the S1 and S2 light signals are read out by the photosensor arrays on top and bottom of the LXe target. This enables transverse position reconstruction via the top PMT S2 response pattern, along with the computation of vertical interaction depth based on the drift time of the electron cloud. Illustration courtesy of L. Althüser.

Upon reaching the liquid surface, electrons must overcome a potential barrier arising from the polarization of the dense xenon medium, which stabilized them in the liquid phase. This makes the transition into the gas energetically disfavored. To achieve efficient extraction, an electric field of several kV/cm is typically applied between the anode and gate electrodes placed just below and above the liquid-gas interface [126, 157]. Following extraction into the gaseous xenon layer, they are accelerated by the extraction field, colliding and exciting xenon atoms that subsequently de-excite and emit additional scintillation photons at 175 nm via electroluminescence. The number of extracted electrons is proportional to the number of secondary photons produced, which are again detected by the PMTs, forming the S2 signal.

An interaction in a TPC is identified by the temporal coincidence of a prompt S1 signal and a delayed S2 signal, which together define a physical event. The (x, y) coordinates of the interaction vertex are reconstructed from the spatial distribution of the S2 light on the top PMT array, while the depth (z coordinate) is determined from the time difference between the S1 and S2, given the measured electron drift velocity.

The deposited energy is reconstructed using the combined S1 and S2 signals introduced




Figure 2.5: Simulated nuclear recoil and electronic recoil distributions at an electric drift field of 200 kV/cm. The shaded contours indicate the 1σ (solid) and 2σ (light) density regions. Simulations were performed with NEST v2.4.0 [148] and the Python interface `nestpy` 2.0.2 [158].

in Equation 2.5. To correlate the produced quanta with the measured signals, two detector parameters are introduced. The first is g_1 , that represents the average probability to detect a scintillation photon as a photoelectron in the PMTs, and the second is g_2 , which quantifies the average number of detected photoelectrons in the PMTs per ionization electron produced in the liquid xenon. The parameter g_1 accounts for the position-dependent light collection efficiency (LCE) and the PMT response. The latter includes the quantum efficiency (QE), defined as the probability that an incident photon produces a photoelectron in the photosensor, and the collection efficiency (CE), defined as the probability that a produced photoelectron reaches the amplification stage of the photosensor. More details on these parameters are provided in Section 3.1. The g_2 additionally incorporates the electron extraction efficiency ϵ_{ext} at the interface and the gas gain G , defined as the number of photoelectrons detected per extracted electron into the gaseous xenon. The latter, also referred to as single electron gain, accounts for the PMT efficiencies introduced earlier. Mathematically, g_1 and g_2 can be expressed as:

$$\begin{aligned} g_1 &= \frac{S1}{N_\gamma(x, y, z)} = \text{LCE}(x, y, z) \times \text{QE} \times \text{CE}, \\ g_2 &= \frac{S2}{N_{e^-}(x, y, z)} = \text{LCE}(x, y, z) \times \epsilon_{\text{ext}} \times G(x, y). \end{aligned} \tag{2.9}$$

Using Equation 2.9, the reconstructed energy as defined in Equation 2.5 can be expressed directly in terms of the measured S1 and S2 signals as follows:

$$E = W \cdot \left(\frac{S1}{g_1} + \frac{S2}{g_2} \right). \quad (2.10)$$

The g_1 and g_2 values are detector-dependent and are determined experimentally through dedicated calibration campaigns at various energies and with different radiation sources.

In dual-phase LXe TPC detectors, the most relevant interaction types are NRs and ERs. Discrimination between the two, is typically performed in the S1 versus S2 parameter space, where events populate two distinct bands, as shown in Figure 2.5. As discussed in Subsection 2.1.2, this separation originates from the different stopping power of electrons and recoiling xenon nuclei, which produce different ionization track densities and thus, distinct electron–ion recombination probabilities. This translates into a smaller S2/S1 ratio for NR interactions than ER interactions of the same energy. Current generation dual-phase LXe TPCs have achieved ER discrimination power greater than 99 % at 50 % NR acceptance [159–161]. In practice, the ER and NR bands are determined through dedicated calibrations. ER populations are measured using internal β -emitting sources, while NR populations are obtained with fast neutron sources.

2.1.4 Background sources

In any dark matter experiment, a variety of processes can deposit energy in the detector and mimic the signal of interest. Such events are collectively referred to as background. In LXe TPCs, the background sources can be broadly grouped into external, intrinsic and neutrino induced. In this context, this section presents an overview of the dominant background contributions.

As outlined in Section 1.1.3, the characteristic WIMP signature corresponds to a single-scatter, low-energy nuclear recoil uniformly distributed within the active target. While ERs, typically induced by gamma rays, electrons and neutrinos, dominate the total interaction rate, they can be strongly suppressed via the S2/S1 ratio. In contrast, NR backgrounds, particularly from neutrons and coherent elastic neutrino-nucleus scattering (CE ν NS), are far more difficult to distinguish from a WIMP signal and therefore represent the ultimate background for LXe dark matter searches.

High-energy cosmic muons with mean energies of $\mathcal{O}(100 \text{ GeV})$ can penetrate large amounts of rock overburden [162], where they induce secondary showers through a variety of processes, including deep-inelastic muon–nucleus scattering, nuclear spallation, photo-nuclear reactions, and hadronic interactions of secondary nucleons, pions,




Figure 2.6: The decay chains of ^{238}U (left) and ^{232}Th (right). Together with the ^{235}U decay chain, these are the dominant low-energy background sources in XENONnT. They produce neutrons through (α, n) reactions and spontaneous fission, while gamma rays from nuclear de-excitation generate Compton backgrounds. The beta decay of ^{214}Pb constitutes the dominant intrinsic low-energy ER background.

and kaons. These interactions produce fast neutrons with energies extending up to tens of GeV, which are partially moderated by the surrounding rock and detector shielding. Neutrons that reach the active target can generate NRs of a few keV through multiple elastic scatterings, giving rise to NR multi-site topologies. To suppress this background, LXe TPCs are located deep underground and surrounded by water Cherenkov muon veto (MV) systems [142, 143, 163] that enable tagging parent muons and their associated cascades.

Other relevant backgrounds arise from radiogenic sources, in particular from trace amounts of ^{232}Th and ^{238}U present in the detector materials and shielding. Their decay chains generate neutrons through spontaneous fission and (α, n) reactions (see Figure 2.6). Unlike the high-energy cosmogenic neutrons produced by muons, these radiogenic neutrons have energies of few MeV and can induce single-scatter NRs that are indistinguishable from a WIMP interaction. This motivates the use of extremely radiopure construction materials and a precise characterization of their intrinsic radioactivity, together with the integration of dedicated neutron veto (NV) systems [128, 163, 164],

which will be discussed in Section 2.2.3.

A small fraction of events populating the ER band (typically $\sim 0.5\%$) leaks into the NR band region. These events are referred to as ER background. They are mainly induced by β decays from daughter isotopes in the ^{232}Th and ^{238}U decay chains (see Figure 2.6). Additionally, ER background arises from gamma emission by ^{40}K , naturally occurring in some materials, as well as by ^{137}Cs and ^{60}Co . The latter if of anthropogenic origin in stainless steel and can also be cosmogenically produced in copper. The mitigation of this background begins with careful material selection during detector construction and is complemented by fiducialization. However, high-energy gamma rays ($\mathcal{O}(\text{MeV})$) have long attenuation lengths and can penetrate deep into the active target. They predominantly interact via Compton scattering (see Figure 2.1), producing multi-site events that are not always rejected by fiducialization. While the multi-scatter nature of these events provides a handle for rejection in WIMP dark matter searches, for multi-scatter signals at the MeV scale, the gamma background becomes dominant. In this case, additional rejection techniques based on machine-learning methods are therefore necessary for efficient signal-to-background discrimination, a topic that will be further discussed in Chapter 6.

ER backgrounds intrinsic to the xenon target are also present. These include the β -decay of ^{85}Kr ($T_{1/2} = 10.75$ years [165]), present as a trace contaminant in commercially produced xenon due to incomplete krypton removal during air separation, with the isotope itself being predominantly of anthropogenic origin [166]. Additional irreducible ER backgrounds arise from radioactive decays of xenon isotopes, in particular $2\nu\beta\beta$ decay of ^{136}Xe ($T_{1/2} = (2.165 \pm 0.016_{\text{stat}} \pm 0.059_{\text{sys}}) \times 10^{21}$ years [109]), as well as from the $2\nu\text{ECEC}$ decay of ^{124}Xe ($T_{1/2} = (1.13 \pm 0.18_{\text{stat}} \pm 0.05_{\text{sys}}) \times 10^{22}$ years [115]). These decays are uniformly distributed in the active volume and therefore cannot be suppressed by fiducialization. An important contribution comes from the isotope ^{222}Rn ($T_{1/2} = 3.8$ days) which is continuously emanated from the surfaces in the experiment that are in contact with xenon, which dissolves into the LXe becoming homogeneously distributed. In particular, the β transition of its daughter isotopes ^{214}Bi to ^{214}Po constitutes the dominant ER background source for low-energy searches. Furthermore, the β transition of ^{210}Bi to ^{210}Po followed by its subsequent α -decay can produce fast neutrons in (α , n) reactions.

Recently, the XENONnT [89] and PandaX-4T [90] collaborations reported the first observation of solar neutrinos from the ^8B decay via $\text{CE}\nu\text{NS}$. While the presence of this background has long been predicted, its direct measurement demonstrates that the dual phase LXe TPC technology has reached the sensitivity required to probe the so-called neutrino fog [58], where solar neutrinos become an irreducible background for WIMP




Figure 2.7: Evolution of the XENON dark matter detectors. For each experimental stage, XENON10 [167], XENON100 [168], XENON1T [169], and XENONnT [142], the mass of active LXe is indicated by the dark blue bars. The total LXe mass contained within the cryostat of each experiment is 25 kg (XENON10), 161 kg (XENON100), 3200 kg (XENON1T), and 8400 kg (XENONnT). The level of low-energy ER background ($\lesssim 20$ keV) for each detector is indicated by the red bars.

dark matter searches.

2.2 The XENONnT experiment

The XENONnT detector is the current phase of the XENON project, located at the Laboratori Nazionali del Gran Sasso (LNGS) at a depth of 3600 meters water equivalent (m.w.e.), where the cosmic muon flux is reduced by a factor of 10^6 compared to sea level [170]. It consists of a dual-phase LXe TPC designed to continue the search for WIMP dark matter, building upon the expertise and results of its predecessors.

Having outlined the main background sources relevant for dark matter searches, which drive the detector design, this section is aimed at presenting the detector architecture together with its veto systems, xenon storage and handling infrastructure, as well as the novel subsystems deployed in the current phase of the XENON project. A brief overview of the evolution of the XENON program, from XENON10 to XENONnT, is also provided.

2.2.1 From XENON10 to XENONnT

The XENON program has evolved through successive detector generations, each representing an increase in target mass and reduction in background levels. A comparison of different detector generations is shown in Figure 2.7, illustrating the relative target sizes and the corresponding improvements in background levels. XENON10 [167]

demonstrated the viability of dual-phase TPCs for dark matter searches and set competitive limits on the WIMP–nucleon cross section. Its successor, XENON100 [168], achieved an order of magnitude improvement in sensitivity through improved shielding, purification, and a larger target (see Figure 1.5). The XENON1T experiment [171] was the first tonne-scale liquid xenon dark matter detector, instrumenting 2 t of LXe in the active TPC volume and containing a total of 3.2 t inside the cryostat, achieving an overall improvement in sensitivity to spin-independent WIMP interactions of more than 3 orders of magnitude since the beginning of the XENON project. In addition, the $2\nu\text{ECEC}$ decay of ^{124}Xe was first observed with the XENON1T experiment [115].

XENONnT operates with an active LXe target of 5.9 t, nearly 400 times larger than that of the first generation. It was realized as a fast upgrade of XENON1T, decommissioned in 2020, and thus several subsystems have been inherited. Three such examples include the water Cherenkov muon veto, the GXe purification system, and the outer cryostat vessel. The experiment comprises three nested detectors, from inside outside: the dual-phase LXe TPC, the NV and the MV, and it is operated alongside a three-story service building that houses the supporting infrastructure, including systems for xenon cooling, purification, storage, and distillation, as well as the data acquisition and slow-control systems for monitoring the experiment.

To mitigate intrinsic backgrounds, several new subsystems were introduced in this generation, including the cryogenic distillation column dedicated to the removal of ^{222}Rn , the NV system, the LXe recirculation and purification system as well as a novel field shaping system for the TPC [172].

2.2.2 Dual-phase time projection chamber

The dual-phase XENONnT TPC, which features a cylindrical active region of 1613 mm in height and 1327 mm in diameter [142], is located at the core of the nested veto system. A CAD rendering of the XENONnT cryostat is shown in Figure 2.9. The cylindrical active target is surrounded by vertical panels of polytetrafluoroethylene (PTFE), forming a reflective side wall. These panels provide a highly reflective surface to maximize the light collection efficiency by the PMTs. Moreover, they shield VUV light produced in the non-instrumented LXe from entering the active region. PTFE also introduces background and operational challenges as it contains trace amounts of ^{238}U and ^{232}Th , making it a source of ^{222}Rn emanation. In addition, PTFE is a source of outgassing¹, which reduces xenon purity and limits charge collection efficiency. For these reasons, the use of PTFE in detector construction must be minimized. In XENONnT, the total

¹Release of impurities such as O_2 and N_2 .




Figure 2.8: Photograph of the XENONnT experiment at LNGS. The 10 m-high water tank on the left serves as a Cherenkov muon veto. It contains the neutron veto, which surrounds the cryostat housing the LXe TPC, together with the calibration structures shown in Figure 2.10. The three-story service building on the right hosts the cryogenics, purification and distillation systems, as well as the data acquisition and storage systems. Photograph by H. Schulze Eißing/ XENON Collaboration.

PTFE mass is restricted to about 128 kg. The PTFE panels are 3 mm thick, corresponding to the minimum thickness required for opacity to xenon scintillation light [173]. They are joined together in an interlocked structure using PTFE pillars, ensuring both mechanical stability and optical isolation of the active volume. Two arrays of PMTs are installed at the top and bottom of the cylindrical active volume, with 253 PMTs in the top array and 241 in the bottom, for a total of 494. The PMTs are held in place by a single solid 8 mm-thick PTFE support disk, which also fills the gaps between PMTs to provide additional light reflectivity. Each PMT is a 3-inch diameter Hamamatsu R11410-21, chosen for its low intrinsic radioactivity [169] and its high quantum efficiency of $\sim 30\%$ at the xenon scintillation wavelength at room temperature [174].

Five parallel grid electrodes (see Figure 2.9) are installed at different depths within the cylindrical TPC volume: (i) the bottom screening grid, located 1.5 mm above the bottom PMT array, (ii) the cathode, (iii) the gate, (iv) the anode, and (v) the top screening grid, located 28 mm below the top PMT array. The two screening grids are placed to minimize the field in front of the PMT photocathodes. The cathode and the gate define the drift field region, which extends 1486 mm, while the gate and anode define the electron extraction region through a gap of 8 mm. Two and four transverse support wires were installed below the gate and anode electrodes respectively, to counteract sagging of the electrode plane [142]. This deformation arises from the combined mechanical stress due to the size and weight of the electrodes and the electrostatic forces




Figure 2.9: **Left:** CAD rendering illustrating the XENONnT cryostat and TPC structure. **Right:** Close-up views of the top and bottom regions, highlighting the electrode configuration and PMT arrays [142]. Figure from Ref. [142].

applied during operation. The volume between the cathode and the bottom screening grid contains about 200 kg of LXe. In this region, only S1 signals are detected, as ionization electrons cannot be extracted in the absence of an upward drift field.

The cathode is negatively biased using a Heinzinger PNC 150000-1 NEG high-voltage supply connected through a custom-made feedthrough originally designed for XENON1T. The feedthrough consists of a conductive stainless-steel rod enclosed within an ultra-high-molecular-weight polyethylene insulator, cryofitted into a stainless-steel tube of 25.4 mm diameter to ensure vacuum tightness. The total length of the feedthrough is about 1.5 m [171].

The liquid level is regulated using a diving bell equipped with a gas inlet and an exhaust tube that opens into the LXe volume. The flow of GXe pressurizes the bell, and the balance between the gas pressure inside the bell and the hydrostatic pressure of the surrounding liquid determines the position of the liquid–gas interface. By adjusting the height of the lower opening of the exhaust tube, controlled remotely via a motion feedthrough, the equilibrium pressure, and therefore the liquid level inside the bell, can be set precisely. The liquid level is continuously monitored during operation using four parallel-plate capacitive short-level meters installed around the TPC circumference. These sensors provide a measurement precision of 0.2 mm over a dynamic range of 10 mm. During LXe filling, the liquid level is monitored by two capacitive level meters extending along the TPC height [142].

Two concentric sets of field-shaping electrodes are installed to ensure uniformity of the drift field, separated radially by 15 mm and arranged in a vertically interleaved pattern. The inner set consists of 71 oxygen-free high-conductivity (OFHC) copper wires in

direct contact with the outer edge of the PTFE wall to mitigate charge-up. The outer set consists of 64 OFHC copper bands, each 15 mm tall and 5 mm thick, forming 24-sided polygonal shape electrodes that surround the PTFE cylinder. Each stack is set at incremental potential via resistor chains from the cathode to the top field shaping ring, isolated from the gate [172].

The TPC is housed within a vacuum-insulated double-walled cryostat made of 5 mm thick low-radioactivity stainless steel [175]. The inner vessel contains the full 8.4 t inventory of LXe, while the outer vessel provides thermal insulation from the environment. The cryostat has a quasi-cylindrical shape and is sealed at the top and bottom by domed endcaps, which accommodate several access ports. A double-walled, 6 m cryogenic pipe is coupled to the upper dome, extending from the outer vessel to the cryogenics and purification systems located in the service building. This pipe accommodates all electrical, optical and sensor connections, as well as the gas purification system inlet/outlet lines. The bottom dome allows for an additional port used for connecting the cryostat to the inlet of the liquid purification system.

The cooling power required to maintain xenon in the liquid phase is provided by two redundant pulse tube refrigerators (PTRs) [176], each delivering 250 W of cooling power. The xenon storage system comprises two units, ReStoX1 [171] and ReStoX2 [142], which together can accommodate the full xenon inventory. ReStoX1 is directly connected to the bottom port of the cryostat through a vacuum-insulated line and is mainly used for rapid recovery of LXe in case of an emergency. ReStoX2 is designed for GXe storage. It is rated up to 60 bar can hold 10 tonnes of xenon at room temperature, more than the total xenon inventory at room temperature [142].

Unlike in XENON1T, xenon purification in XENONnT is performed in both the gas and liquid phases [142]. For the GXe purification, LXe is extracted from the inner vessel through a line running inside the cryogenic pipe, evaporated via a system of heat exchangers, and fed into the gaseous purification loop. The xenon gas is continuously circulated through high-temperature getters [177], operated in two parallel branches. Circulation is maintained by custom-developed magnetically coupled piston pumps [178], achieving a nominal flow rate of ~ 80 SLPM.

For LXe purification, additional cooling is provided by a LXe-LN₂ heat exchanger, which compensates the thermal load from the circulation system and ensures the thermal stability of the xenon during operation. The LXe is extracted from a vacuum-insulated port located at the bottom of the cryostat and circulated by a cryogenic liquid pump [179] capable of flow rates up to 4 LPM, corresponding to ~ 16 t/d. Impurities are removed by custom-designed vacuum-insulated purifiers containing SAES

St707 getter pills, which capture O_2 and other electronegative contaminants during standard operation of the system [180]. During detector commissioning, an Engelhard Q-5 purifier was also employed to remove the high impurity load originating from out-gassing of detector materials. However, because of its relatively large radon emanation rate, this purifier is used only during initial purification stages. Once the xenon purity reaches a sufficiently low O_2 -equivalent concentration, purification proceeds solely with the St707 getters. Under nominal conditions, the LXe purification system can reduce the impurity concentration to about 0.1 ppb (O_2 -equivalent) within one week. The purity of xenon is commonly quantified in terms of the electron lifetime, τ , which is defined as the time required for the number of electrons (N_e) at zero drift time (t_0) to decrease by a factor of e due to attachment to impurities. The electron lifetime is parametrized as:

$$N_e(t_d) = N_e(t_0) e^{-t_d/\tau}. \quad (2.11)$$

In XENONnT, electron lifetimes on the order of ~ 10 ms have been achieved [156].

The krypton impurities in the xenon are removed by cryogenic distillation in a dedicated column that takes advantage of the 10-times higher vapor pressure of krypton relative to the xenon at 178 K [166]. After a single pass through the column, the concentration of $^{\text{nat}}\text{Kr}$ is reduced by a factor of $(6.4^{+1.9}_{-1.4}) \times 10^5$.

In XENONnT, the achieved ^{222}Rn activity concentration within the 5.9 tonne liquid xenon volume is $(0.90 \pm 0.01_{\text{stat}} \pm 0.07_{\text{sys}}) \mu\text{Bq/kg}$ [181]. This reduction results mainly from the use of low-radon-emitting gas pumps, combined with extensive material selection and cleaning procedures, as well as the implementation of a novel online radon removal system integrated into the LXe recirculation loop [142]. The radon removal system operates based on cryogenic distillation, exploiting the 10-fold lower vapor pressure of radon relative to xenon [181]. Liquid xenon is extracted from the cryostat and fed into a distillation column, where it is heated by a reboiler, causing the xenon to evaporate. Because radon has a much lower vapor volatility, it remains predominantly in the liquid phase. Over time, radon accumulates at the bottom of the column and decays ($T_{1/2} = 3.8$ days). A fraction of the radon-depleted xenon gas is recondensed and returned to the detector, while the remainder exits from the top of the column and is reintroduced into the cryostat.




Figure 2.10: CAD rendering of the XENONnT detector. The setup reuses the XENON1T water tank, which contains 84 MV PMTs mounted on its inner walls. The white octagonal structure represents the reflector panels that enclose the NV region, which houses 120 PMTs. The stainless-steel outer vessel is similarly clad with reflector material to improve light collection in the NV. The external calibration systems are: beam pipe (pink), I-belt (blue), and U-tubes (green and red), which guide sources around the TPC. Figure from Ref. [142]

2.2.3 Neutron veto

The NV operates as a water Cherenkov detector. It was implemented to identify and reject TPC events that coincide with an interaction in the NV. Such coincidences are typically produced by radiogenic neutrons that manage to penetrate the TPC and subsequently escape further detection in LXe. The system was initially operated with pure demineralized water but was designed to ultimately use a 0.48 % (by mass) solution of gadolinium sulfate octahydrate ($\text{Gd}_2(\text{SO}_4)_3 \cdot 8\text{H}_2\text{O}$) to increase neutron capture efficiency [142]. Gadolinium has the highest thermal neutron capture cross section among naturally occurring elements, in particular the isotopes ^{155}Gd and ^{157}Gd with cross-sections of 5.2×10^4 barn and 2.15×10^5 barn [182], respectively, making it extremely efficient at capturing neutrons. The capture probability on Gd is $\sim 91\%$, compared to 9 % for hydrogen. Following neutron capture, the excited Gd nucleus de-excites within a few picoseconds through the emission of a cascade of high-energy gamma rays, releasing a total energy of about 8 MeV. In the case of hydrogen capture, a single 2.2 MeV

gamma ray is emitted during the de-excitation of the nucleus. These gamma rays generate Compton electrons in the surrounding water, which in turn produce Cherenkov light detected by the PMTs. The neutron capture occurs with a medium-characteristic delay after the initial interaction, with demineralized water being 200 μs [164] and, in gadolinium-loaded water, depending on the Gd concentration. For example, a capture time of about 110 μs has been measured for a Gd concentration of 0.1 % [183]. This distinctive delayed-coincidence signature used to tag neutron-induced events in the TPC. Neutrons that interact within the LXe active region and subsequently escape the TPC are moderated by water in the NV, typically traveling less than 20 cm before being thermalized and captured by a Gd nucleus (or hydrogen in the case of operation with pure water). To maximize the light collection efficiency, the walls and end caps of the NV are lined with 1.5 mm-thick expanded PTFE, a material with a reflectivity exceeding 99 %. This layer not only improves light collection but also provides optical isolation between the NV and the MV. It is instrumented with 120 low-radioactivity Hamamatsu R5912-100-10 8 inch PMTs [184], arranged in equidistant rows along the walls of the NV. These PMTs have a QE of $\sim 40\%$ at a wavelength of 350 nm. In the first science run (SR0) of XENONnT, using demineralized water, a 250 μs tagging window achieved a neutron-tagging efficiency of $(53 \pm 3)\%$ with a livetime loss of 1.6 % [164].

2.2.4 Muon veto

The muon veto serves both as an active water Cherenkov detector and as a passive shield against environmental gamma rays and neutrons [142]. It consists of ~ 700 t of demineralized water contained in a cylindrical tank with a height of 10 m and a diameter of 9.6 m. The inner surface of the tank is clad with a highly reflective foil to maximize LCE [185]. The detector is instrumented with 84 Hamamatsu R512ASSY 8 inch PMTs [186], arranged in five vertically spaced, equidistant rings along the tank walls. These PMTs have a quantum efficiency of about 30 % in the wavelength range of 300–600 nm [187]. The system is designed to actively tag cosmic muons, within a 300 ns coincidence window, and secondary hadronic showers produced through spallation reactions in the surrounding rock. MC simulations show that the detection efficiency for crossing muons is $\sim 100\%$ with this trigger [142]. The NV and the MV share the same water. However, the addition of Gd sulfate does not affect the MV performance. Both veto systems are shown in Figure 2.10.

2.2.5 Calibration system

A set of dedicated mechanical structures are installed in the XENONnT experiment to deploy external calibration sources, which are used for characterizing the TPC response

to well-known reference signals. These structures are shown in Figure 2.10. The beam pipe (in pink), the I-belt (in blue), which enables the placement of a calibration source in the orange box, and the U-tubes (in green and red), which allow the sources to be moved around the TPC.

In XENONnT, several types of external sources are used. A ^{228}Th source provides high-energy gamma rays for characterizing the detector response at the MeV scale, as will be extensively discussed in Chapter 5. A $^{241}\text{AmBe}$ source emits neutrons used to calibrate the NR response and to determine the neutron-tagging efficiency of the NV [164]. Finally, a novel ^{88}YBe photoneutron source produces quasi-monoenergetic 152 keV neutrons, enabling calibration of low-energy NR recoils down to sub-keV energies [188].

For internal calibration sources, radioactive isotopes are typically introduced through the gas system by extracting a small fraction of xenon, mixing it with the radioactive source, and reinjecting it into the TPC as part of the GXe flow. This method ensures a uniformly distributed activity within the active volume, which is particularly useful for studying spatially dependent detector effects.

In XENONnT, internal calibrations are performed using $^{83\text{m}}\text{Kr}$, ^{220}Rn , ^{222}Rn and ^{37}Ar :

- $^{83\text{m}}\text{Kr}$ is used to characterize the spatial uniformity of the detector response at low energies [189]. Is a metastable isomer with a half-life of 1.83 h, originating from the β -decay of ^{83}Rb . It de-excites in two steps via the emission of conversion electrons. The first transition releases a total electron energy of 32.2 keV, followed by a transition to the ground state with a half-life of 154 ns, accompanied by the emission of a 9.4 keV conversion electron.
- ^{220}Rn is used to characterize the low-energy ER response [190, 191]. The ^{220}Rn is supplied by ^{228}Th sources with an activity of 55 kBq, which exhibit a high ^{220}Rn emanation rate. Being a noble gas, ^{220}Rn can efficiently diffuse throughout the detector volume. The subsequent decays of ^{220}Rn ($T_{1/2} = 56\text{ s}$) and its daughter isotopes produce α , β , and γ radiation that can be used to calibrate a variety of detector features. In particular, after being introduced into the xenon target, the progeny isotope ^{212}Pb induces homogeneously distributed low-energy ($< 30\text{ keV}$) ER interactions and exhibits a β -spectrum with a Q-value of 560 keV. Because the decay chain is short-lived, the activity vanishes within a few days.
- ^{222}Rn is used to calibrate the low-energy ER background from its daughter isotope ^{214}Pb . It is introduced into the TPC via dedicated emanation sources connected to the xenon circulation system, allowing it to be uniformly distributed through-

out the LXe target. Its progeny include ^{214}Pb , which undergoes β -decay and constitutes the dominant low-energy ER background. By injecting an increased ^{222}Rn activity, the corresponding ^{214}Pb rate can be measured with high precision. After the calibration, the activity decreases over timescales of a few weeks due to radioactive decay and removal by the purification system.

- ^{37}Ar provides calibration near the energy threshold (~ 1 keV) [192]. It decays via electron capture with a half-life of 35.01 d and a Q-value of 813.9 keV. The atomic vacancies created are filled by outer-shell electrons, leading to the emission of Auger electrons and X-rays at 2.82 keV(90.2 %), 0.27 keV(8.7 %), 0.01 keV(1.1 %), where only the former two are used.

2.3 The next generation of detectors

The XENONnT detector represents the culmination of two decades of technological progress in LXe TPCs. Combining a multi-tonne LXe active mass with ultra-low background levels, its design and operation define the current state of the art in low-energy rare event searches. At present, XENONnT continues taking data and has demonstrated high sensitivity to a variety of rare event processes, including CE ν NS from solar neutrinos [89], light dark matter candidates [193], and ongoing WIMP searches [55], with potential for many additional channels [194–196]. In parallel with this physics program, the detector was further upgraded during summer 2025, including the implementation of new electrodes, as part of the continued effort to improve the stability and performance of the TPC electric-field configuration. After the upgrade, XENONnT entered recommissioning, and new science data are expected to be taken with improved electric-field conditions.

Motivated by this broad physics program, together with the relatively straightforward scalability of dual-phase LXe technology, the case for an even larger detector is compelling. Building on this foundation, the next generation detector aims to further push the WIMP–nucleus spin-independent sensitivity frontier while substantially expanding the physics reach beyond dark matter searches. In this context, the XLZD project will be discussed in Chapter 7, along with the path towards realizing the next stage in the development of the ultimate dual-phase liquid xenon observatory.

Chapter 3

Light Detection in the XENONnT TPC

The dual-phase xenon TPC technology relies on the detection of light produced following a particle interaction in LXe. In XENONnT, both scintillation and electroluminescence light are detected by arrays of 3-inch diameter R11410-21 PMTs manufactured by Hamamatsu. The PMTs are directly responsible for converting the light into measurable signals that are used for event reconstruction. Given the important role of the PMTs, their response and stability over time must be continuously monitored and modeled.

This chapter first introduces the XENONnT photomultiplier tube and summarizes its main characteristics in Section 3.1. Based on this description, Section 3.2 presents the calibration system used to monitor their response over time. Section 3.3 describes the approach adopted to extract the PMT gain, the modeling of its time evolution and its implementation in the event reconstruction chain. The photoelectron acceptance and its impact on signal reconstruction is described in Section 3.4. Finally, Section 3.5 reviews the PMT performance during the second science run (SR1), with emphasis on the gain stability and modeling.

3.1 The XENONnT photomultiplier tube

The Hamamatsu R11410-21 PMT used in XENONnT, is shown in Figure 3.1 left. This PMT model was specifically developed for operation in liquid xenon detectors. Its performance under cryogenic conditions, down to -110°C [174, 197], and in the presence of an electric field has been extensively characterized in multiple facilities, including the University of Zurich [169, 174, 198].




Figure 3.1: **Left:** Photograph of a Hamamatsu R11410-21 photomultiplier tube used in the XENONnT experiment. **Right:** Schematic of a Hamamatsu R11410-21 photomultiplier tube, showing its internal components and operating principle: an incident photon ejects a photoelectron from the photocathode, which is amplified through the dynode chain and collected at the anode as a current.

The working principle of a PMT is illustrated in Figure 3.1 right. When scintillation photons reach the PMT photocathode, charge carriers, commonly referred to as photoelectrons (PE), are emitted via the photoelectric effect. These photoelectrons are accelerated by an electric field established between the photocathode and the first dynode through a focusing electrode. Upon striking the first dynode, each photoelectron releases several secondary electrons. A series of dynodes, each held at a progressively higher electric potential, then accelerate these electrons successively from one stage to the next, repeating the process through all dynode stages. In this way, the initially small current is amplified to a measurable level at the final stage, in the anode, where the signal is collected. The output current is subsequently converted into a voltage signal by passing it through a resistor.

The R11410-21 PMT features a 3 inch flat-circular diameter window made of synthetic silica that is transparent to VUV light. It features a bi-alkali photocathode with a 64 mm diameter sensitive area that exhibits a mean QE of $(34.0 \pm 2.8) \%$ [174] at the xenon scintillation wavelength of 175 nm [135, 141] and a CE of approximately 95 % [197]. The R11410-21 PMT exhibits a non-negligible probability of double photoelectron emission of approximately 20 % for VUV photons incident on the photocathode [199]

The PMT body, made out of a Kovar alloy, is evacuated and constructed from materials selected for their extremely low intrinsic radioactivity [169]. In particular, the ceramic stem, which serves as an insulator within the dynode structure and constitutes the dominant source of radioactive background, has been the focus of extensive material optimization, resulting in total activities at the level of $\lesssim 10 \mu\text{Bq}$ per PMT for the ^{238}U chain [174].

The ratio of the anode output current to the photoelectric current emitted from the photocathode defines the PMT gain. The secondary electron emission ratio for one dynode stage, δ , is given by $\delta = a \cdot \tilde{V}^k$, where a is a constant and k depends on the dynode material and geometry, and $\tilde{V}$ is the inter-stage voltage. For a PMT with n identical dynode stages and an applied voltage between the anode and cathode, V , the total gain is expressed as [200]:

$$G = \delta^n = A \cdot V^{k \cdot n}, \quad (3.1)$$

where A is a proportionality constant that depends on the dynode structure and on the voltage distribution across the stages. For the XENONnT R11410-21 PMTs, with $n = 12$ and $k = 0.67 \pm 0.02$, the median A is $\sim 8 \times 10^4$, yielding a typical gain of $\sim 2 \times 10^6$ at 1500 V. Under normal operating conditions, the PMT response is linear, meaning that the output charge is proportional to the number of photoelectrons produced at the photocathode. Some photoelectrons emitted from the photocathode follow sub-optimal trajectories, for example by skipping a dynode or by being backscattered from a dynode surface. These photoelectrons produce under amplified output signals, which can account for up to approximately 20 % of the single photoelectron (SPE) charge spectrum [201].

Another intrinsic feature is the dark current, a small current present in the absence of physical interactions when the PMT is operated in complete darkness. It originates from processes such as thermionic emission of electrons from the photocathode and dynode surfaces, as well as scintillation in the PMT quartz window induced by cosmic rays. The thermionic emission rate depends strongly on the operating temperature, decreasing rapidly with cooling and becoming strongly suppressed at cryogenic temperatures. The measured dark current at the anode also depends on the applied bias voltage through the gain, since a higher electric field amplifies the emitted electrons more strongly. This current manifests as discrete output pulses, whose frequency is referred to as the dark count rate. The R11410 PMT features a dark count rate of about 50 Hz at -100°C [174].

PMTs contain residual gas molecules or atoms introduced during production or through small leaks, including H_2 , He, CH_4 , Ar, N_2 and Xe. To remove reactive gases from the vacuum body, the PMTs are equipped with a getter material. However, noble gases such as He, Ar, and Xe are not affected by the getter. Ionization of these residual gas molecules by accelerated electrons within the PMT produces positive ions. These ions drift back toward the photocathode or earlier dynodes under the applied electric field and, upon impact, release additional electrons. The amplification of these electrons results in delayed noise pulses known as afterpulses. These are typically observed with

a characteristic time delay following the primary signal. The gases present within the PMT vacuum can be identified in the time-delay afterpulse spectrum, as the delay time, derived from the ion drift in the electric field between the photocathode and focusing grid, scales with the square root of the ion mass-to-charge ratio [197]:

$$t = \left(1.13 \frac{\text{V}^{1/2}}{\text{cm}}\right) \sqrt{\frac{L^2 m}{V_0 q}}, \quad (3.2)$$

where m and q are the ion mass and charge, L is the distance between photocathode and focusing grid, and V_0 is the corresponding potential difference [197]. In the case of Xe and N₂, their presence is typically the result of tiny leaks either during operation in the TPC or during storage in air. As a result, their afterpulse rate can increase over time. This makes continuous monitoring essential in order to track potential leak development and to assess its impact on detector performance.

The transit time refers to the interval between the arrival of a photon at the photocathode and the moment when the corresponding anode pulse reaches its maximum amplitude. It is measured by illuminating the PMT at single-photon intensity and recording the delay between an external LED trigger and the resulting PMT signal [174]. The XENONnT PMT features a transit time of 46 ns [197]. The arrival of the electron avalanche at the anode is approximately 125 ns after the LED trigger. Some photoelectrons are emitted with different initial angles or kinetic energies, resulting in slightly different trajectories through the dynode chain. These variations cause fluctuations in the transit time among individual pulses. The full width at half maximum (FWHM) of the transit time distribution is referred to as the transit time spread (TTS). The XENONnT PMT features a mean TTS of (9.2 ± 0.5) ns [174].

For the XENONnT experiment, 153 PMTs were reused from XENON1T. Additionally, 341 PMTs were installed after undergoing a dedicated testing and selection campaign [174], chosen from 368 newly produced units. In total, 494 PMTs are instrumented in the TPC, with channels 0–252 assigned to the top array and channels 253–493 to the bottom array [142]. The PMTs are positioned in a hexagonal layout, with those exhibiting the highest QE assigned to the central region of the bottom array and then radially outward in order to maximize the S1 light collection efficiency. For the top array, PMTs with marginal performance were deliberately kept away from the outermost positions, where a failure would most strongly affect the S2 signal light distribution used for position reconstruction, which directly impacts the fiducial volume definition [142].

Understanding the PMT response under operating conditions is essential for accurate




Figure 3.2: Schematic of the PMT calibration system used in the XENONnT TPC. Pulsed blue LED light is generated outside the detector and delivered into the liquid xenon volume through a series of optical fibers for gain and occupancy calibration.

event reconstruction and long-term detector stability. Regular calibrations are therefore performed to monitor variations in performance parameters. These procedures are described in the following section.

3.2 Photosensor calibration

The PMT arrays are calibrated weekly using low-intensity blue (460 nm) LEDs coupled to a remotely controlled pulse generator (Berkeley Nucleonics Corporation model 505-4c). The light is transmitted into the TPC through 24 polymethyl methacrylate (PMMA) optical fibers inserted through the PTFE panels via 250 μm holes. The fibers are arranged in two concentric rings of 12 fibers each, positioned equidistantly at 25 cm and 20 cm below the gate electrode. The fibers run through the cryogenic pipe that connects the external air side, where the LEDs are housed in a black box, to the internal cryogenic side. A schematic of the calibration system is shown in Figure 3.2.

The LEDs are controlled by the pulse generator, which is externally triggered in synchronization with the digitizers in the data acquisition (DAQ) system. The trigger frequency is set to 2.4 kHz, allowing for the acquisition of approximately $\mathcal{O}(10^5)$ waveforms per PMT within 3 min. A NIM timing module is used to apply a delay of 0.8 μs to




Figure 3.3: Left: Example of PMT waveforms recorded during an LED calibration. The green shaded region indicates the region used to estimate the baseline level of the PMT signal. The blue shaded region defines the time window over which the waveform is integrated. **Right:** Distribution of LED pulse arrival times for all PMTs during the same calibration. The red horizontal lines denote the boundaries of the pulse integration window, and the black vertical line separates PMTs in the top and bottom arrays.

the pulse generator relative to the DAQ trigger so that the LED-induced signals appear centered in the $1.6\ \mu\text{s}$ acquisition window (160 samples of 10 ns width). An example of the LED arrival time distribution is shown in Figure 3.3. The first 40 samples, where no LED-induced signals are expected, are used to compute the waveform baseline. The signal area and amplitude are extracted using the [78, 116] sample window. This window was chosen during the commissioning run of XENONnT, and was defined in order to include the entire $\sim 150\ \text{ns}$ LED-induced pulse, including the spread in the arrival time.

The LED calibration data are used to monitor three main parameters that serve as indicators of stability and potential degradation: the gain, the SPE acceptance, and the afterpulse rate. These parameters are discussed in detail in the following sections.

3.3 PMT gain

The gain is a fundamental quantity in the calibration of the PMT response, as it represents the first step in the event reconstruction chain. It provides the conversion factor between the recorded signal in analog-to-digital converter (ADC) counts and the corresponding light intensity expressed in PE units. In the following subsections, we discuss the procedure used to calculate the gain, the concept and development of its time-dependent model, and the implementation of the resulting correction within the XENONnT event reconstruction chain.




Figure 3.4: **Left:** Integrated charge spectra for PMT 226 with the LED off (black) and on (teal). **Right:** Occupancy as a function of the ADC area threshold for the same PMT and calibration. The vertical blue line marks the chosen ADC threshold, selected as the value where the relative uncertainty of the occupancy falls below 2 % (red horizontal line in the lower panel) while the probability of observing zero LED-induced photoelectrons remains high.

3.3.1 Model-independent single photoelectron analysis

The first ingredient in the gain determination is the integrated charge spectrum of the detected photoelectrons. Since the production of photoelectrons is a stochastic process, the number of detected photoelectrons n follows a Poisson distribution:

$$P(n|\lambda) = \frac{\lambda^n e^{-\lambda}}{n!}, \quad (3.3)$$

where λ is called the occupancy, defined as the average number of photoelectrons produced by a PMT for a given light intensity and is given in units of photoelectrons per trigger [PE/trigger].

For a given occupancy, the distribution of the output signal area can be expressed as sum over the n -PE spectra, with each component weighted by its probability $P(n|\lambda)$. For scintillation signals that produce more than 5 PE, the n -PE charge distribution is well approximated by a Gaussian due to the central limit theorem [201]. However the overall shape of the distribution can be modified by sub-optimal PE trajectories in the amplification process, and thus, fitting a specific model can significantly bias the estimated mean of the SPE spectrum, which motivates the use of a model-independent method.

The PMT gain represents how much charge corresponds to one PE, namely the mean of the SPE distribution. In XENONnT, the PMT response is estimated without assuming any particular functional form of the SPE distribution, but using instead the statistical

approach described in Ref. [201]. To this end, two datasets are acquired during an LED calibration: (1) an externally triggered LED dataset taken at low light intensity to ensure operation in the SPE regime, and (2) a noise dataset recorded with the LED switched off. Figure 3.4 displays exemplary LED-on and LED-off datasets for PMT 226, recorded during SR1. It can be shown [201], that the mean E and variance V of the SPE response Ψ can be expressed as,

$$E[\Psi] = \frac{E[T] - E[B]}{\lambda}, \quad (3.4)$$

$$V[\Psi] = \frac{V[T] - V[B]}{\lambda} - E^2[\Psi], \quad (3.5)$$

where T is the total integrated charge of the LED-on dataset and B is the noise-only charge spectrum of the LED-off dataset. This method is reliable for occupancies up to $\mathcal{O}(10 \text{ PE})$. Using the Poisson probability for observing zero photoelectrons, $P(0|\lambda)$, in equation 3.3, the occupancy can be computed directly as:

$$\lambda = -\ln P(0|\lambda). \quad (3.6)$$

Experimentally, the probability of observing zero LED induced photoelectrons, is estimated as the fraction of LED triggers that produced no detected photoelectrons. If N_0 is the number of zero-PE triggers and N the total number of LED triggers, then:

$$P(0|\lambda) \sim \frac{N_0}{N}. \quad (3.7)$$

To estimate the ratio N_0/N , we define a charge threshold below which the contribution from LED-induced photoelectrons is negligible. In this low-charge region, the LED-off distribution contains only baseline noise and therefore characterizes the true zero PE response. Let A_B denote the number of LED-off entries falling below this threshold, and A_T the corresponding number of LED-on entries. Since the LED-on distribution consists of a mixture of noise and LED-induced signals, it will coincide with the LED-off distribution in this region only for triggers in which no photoelectron was produced. Because both datasets contain the same number of triggers N , the ratio A_T/A_B represents the fraction of LED-on triggers consistent with zero PE events. With the total number of LED-on triggers given by N , the occupancy and variance are [201]:

$$\lambda = -\ln\left(\frac{N_0}{N}\right) = -\ln\left(\frac{A_T}{A_B}\right), \quad (3.8)$$

$$V[\lambda] = \frac{e^\lambda + 1 - 2(A_B/N)}{A_B}. \quad (3.9)$$

The occupancy is then evaluated for several choices of the charge threshold (see Figure 3.4), provided the threshold lies within a region where photon-induced signals are not expected. This region corresponds to the negative side of the 0 PE peak, commonly referred to as the pedestal. From equation 3.9 one can see that $V[\lambda]$ decreases for higher A_B , which increases for higher threshold values. In contrast, if the threshold is set to high, we operate in the regime where the assumption of no LED-induced photoelectrons is no longer valid. Therefore, we choose the threshold value in which we still have enough statistics such that the relative error of the occupancy is below 2%. By propagating the uncertainties in equation 3.4, the total uncertainties of the single photoelectron mean are expressed as:

$$V[E[\Psi]]_{\text{stat}} = \frac{\lambda(E^2[\Psi] + V[\Psi]) + 2V[B]}{N\lambda^2} + \frac{E^2[\Psi](e^\lambda + 1 - 2(A_B/N))}{A_B\lambda^2}, \quad (3.10)$$

$$V[E[\Psi]]_{\text{sys}} = (E[T] - E[B]) \frac{V[\lambda]}{\lambda^2}, \quad (3.11)$$

where one observes that the uncertainty on the occupancy dominates the overall uncertainty in the gain estimate. As shown by Saldanha et al. [201], the uncertainty is minimized for occupancies of approximately $\lambda \sim 2$ PE/trigger. For this reason, each LED calibration is performed at three different illumination settings to ensure that the occupancies across all PMTs in the arrays remain close to this optimal value.

The PMT gain G is defined as the charge produced by a single photoelectron at the PMT output. Once the mean of the SPE charge distribution is determined, it can be converted from ADC units into an electron multiplication factor using:

$$G = \frac{E[\Psi] \cdot r}{f \cdot Z \cdot A \cdot q_e}. \quad (3.12)$$

In this expression, r denotes the ADC resolution in volts per ADC bin, which converts the analog input voltage into the corresponding digitized amplitude. The digitizer used in XENONnT is a 14-bit CAEN V1724, operating at a sampling frequency $f = 100$ MS/s

and providing a dynamic range of 2.25 V. The NIM timing module responsible for the trigger delay has an input impedance of $Z = 50\,\Omega$. The analog signal is further amplified by a factor $A = 10$ before digitization, and q_e denotes the elementary charge.

3.3.2 Time-dependent gain model

As introduced earlier, the PMTs are calibrated on a weekly basis, and the model-independent method outlined in the previous section is used to estimate their gains. Since PMT gains can vary over time due to changes in the detector operating conditions, a time-dependent gain model is required to accurately reconstruct signal amplitudes throughout the science run. In XENONnT, the estimated gains are modeled as a function of time using a custom gain modeling algorithm [202]. This approach was introduced to overcome several limitations of the procedure used in XENON1T, where PMT gains were treated as constant during a science run provided that their standard deviation remained below 3%. For PMTs that did not satisfy this stability criterion, the gain evolution was modeled manually using Fermi–Dirac functions, a procedure that required extensive tuning. Moreover, statistical fluctuations of the measured gains were not explicitly considered in the modeling. Variations in the PMT bias voltage also required manual corrections through empirical piecewise functions, further limiting the level of automation.

The approach adopted in XENONnT aims to smooth out statistical fluctuations in the measured gains while preserving the trend of its time evolution. The methodology is fully automated, enabling the modeling of all PMT gains within the time requirements of online analyses. The steps of the modeling procedure are illustrated in Figure 3.5.

In the first step, the gain model is produced using a custom algorithm that smooths the measured gain values with a Savitzky–Golay filter. For each data point, the algorithm fits the local gain evolution trend within a dynamic window that uses adjacent measurements (typically five points before and five after) and evaluates the quality of the fit. Based on this goodness-of-fit diagnostic, the algorithm can automatically adjust the window size or change the fit function (linear by default) to better capture the local trend in the gain evolution. The best-fit parameter values together with the smoothed gain value is then used to extrapolate the gain to both forward and backward timestamps, which is especially useful around bias voltage changes.

In a second step, the relative gain variation over the science run is evaluated by computing the difference between the 95th and 5th percentiles of the smoothed gain distribution. If this variation is below a 2.5% threshold, the gain is considered stable and the average measured value is used in the model. Otherwise, the smoothed, time-




Figure 3.5: Time evolution of the gain for four representative PMTs. In the first step, the measured gain values from the LED calibrations are smoothed. In the second step, if the stability criterion is satisfied, the gain is replaced by its average value, whereas PMTs failing the criterion maintain the smoothed gain values in the model. The top panels show examples from the top PMT array of XENONnT, and the bottom panels correspond to PMTs in the bottom array. The left-hand panels illustrate PMTs for which the stability criterion was met, while the right-hand panels show cases in which the criterion was not fulfilled.

dependent gain is used. The 2.5 % threshold was determined empirically from studies in the first science run of XENONnT [202], where it was found to provide the most effective criterion for identifying PMTs with stable gain behavior.

The PMT gain depends strongly on the applied bias voltage. During the commissioning of XENONnT, the voltages were tuned such that all PMTs operated at gains of approximately 2×10^6 . During science data-taking, however, the bias voltage may be intentionally reduced to mitigate intermittent light emission or PMT flashes, which leads to a sudden gain decrease. In the gain modeling workflow, the voltage settings for each PMT are retrieved from the slow control system and are used to track any voltage changes over time. When no gain calibration is available at a given timestamp, the gain is estimated using Eq. 3.1, allowing the model to interpolate or extrapolate the gain based on the actual operating voltage of the PMT.

During the ^{220}Rn calibration, all but eight PMTs are ramped down to zero high voltage in order to prevent damage caused by the high event rate. The eight PMTs are kept on to monitor the event rate while the valve used to inject the ^{220}Rn source into the xenon

circulation system is opened. After the injection period, all PMTs are ramped back to their nominal operating voltages, and LED calibrations are performed to verify the gain stability. Since the gain model follows the applied bias voltage through Equation 3.1, direct interpolation between these periods and nominal operating gains would therefore introduce artificial discontinuities in the time-dependent gain model. To avoid this, the measured gain is extrapolated to one second before and after each high-voltage transition, and the gain values during the ramping interval are obtained by interpolating between these two extrapolated points. Fluctuations are occasionally observed in the gains extracted from LED calibrations performed shortly after the ^{220}Rn calibration. In Figure 3.10, the short-term gain increase shortly before the 200th day after the start of SR1 corresponds in time with the ^{220}Rn calibration period. The origin of these fluctuations is not completely understood. However, they could be correlated to the repeated voltage changes.

3.3.3 Gain implementation as a correction

As of now, the gain is treated as the electron amplification factor of the PMT. In the next step, they are converted into the quantity `adc_to_pe`, defined as

$$\text{adc_to_pe} = \frac{\text{adc_to_e}}{G}, \quad (3.13)$$

where G is the electron amplification factor obtained from Equation 3.12, and `adc_to_e` is a constant, which depends on the technical specifications of the digitizer, that converts ADC counts to readout electrons corresponding to `adc_to_e` = 17 142.817 41 $\text{e}^-/\text{ADC count}$. Applying the `adc_to_pe` factor to a pulse area in ADC counts yields the corresponding area in PE units. This conversion is applied in the first stage of the event reconstruction chain of XENONnT.

As a time-dependent calibration constant, the `adc_to_e` values are stored and version-controlled within the **XeDocs** [203] metadata management tool together with other detector parameters. These include additional calibration constants and correction maps, collectively referred to as corrections, which account for time-dependent variations and other non-ideal detector effects, which will be discussed in Section 4.3. Each correction is associated with a specific detector configuration and timestamp, ensuring reproducibility among analyses and data releases. For the gain model, the time-dependent `adc_to_pe` values are stored as a set of discrete entries indexed by PMT number, timestamp, and detector subsystem (TPC, neutron veto, or muon veto). These calibrated values are uploaded to the **XeDocs** database in a standardized JSON format and versioned to allow for consistent use in data processing. PMTs that are disabled or excluded

are represented with `adc_to_pe` values set to zero. During event reconstruction, the `adc_to_pe` value for each PMT at the waveform timestamp are retrieved, which ensures that all pulse areas are converted to PE units using the correct, time-dependent gain value.

3.4 Single photoelectron acceptance

During science data taking, each PMT has its own self-trigger threshold, a channel-specific value defined relative to its baseline. This threshold prevents the detector from being continuously triggered by electronic noise and determines when signals from that channel are read out by the DAQ system. When a PMT signal exceeds its self-trigger threshold, it is recorded by the DAQ system and stored in a Mongo database [204]. The digitized PMT signals, referred to as *raw records*, are processed using dedicated algorithms implemented in the open-source `strax/straxen` [205, 206] framework developed for xenon-based detectors. The *raw records* acquired within a single DAQ trigger constitute a waveform.

The fraction of SPE signals that have amplitudes larger than the self-trigger threshold is the SPE acceptance. It quantifies the efficiency with which true single-photoelectron pulses are recorded above threshold and therefore directly impacts the low-energy threshold of the detector. It is obtained from the noise-subtracted amplitude spectrum, constructed by subtracting the LED-off from the LED-on spectrum. The calibration data used to extract the SPE acceptance are acquired with the same DAQ configuration as the gain calibration, but with reduced light level. This ensures operation in the low occupancy regime, where the majority of triggers correspond to the production of a single photoelectron. Figure 3.6 left shows these amplitude spectra for PMT 257, where the LED dataset is shown in purple and the noise dataset in orange. As in the datasets acquired for gain analysis, the LED-on spectrum contains contributions from LED-induced single photoelectrons as well as events from multiple photoelectrons, visible as the high-amplitude tail, dark counts, and residual noise. The noise spectrum, by contrast, includes only electronic noise and dark counts. The SPE acceptance is computed as:

$$\text{SPE}_{\text{acc}} = 1 - \frac{\sum_i^{\text{ADC}_{\text{th}}} n_i}{\sum_{\infty} n_i}, \quad (3.14)$$

where ADC_{th} is the self-trigger threshold of the PMT and n_i is the count in the i -th bin of the noise-subtracted amplitude spectrum. The LED-off spectrum is scaled in order to match the same noise level as the LED-on spectrum. The scaling factor is determined such that the LED and noise datasets agree below a chosen ADC threshold, typically




Figure 3.6: **Left:** Amplitude spectra of the LED calibration dataset (purple) and the corresponding noise dataset (orange); the inset highlights the pedestal region used for scaling. **Right:** Noise-subtracted SPE spectrum (pink) and SPE acceptance curve (black) for PMT 257, showing the acceptance at a 15 ADC threshold.




Figure 3.7: **Left:** Example waveforms from PMT 367 during a calibration in which the LED pulse arrival time shift was largest, causing pulses to extend beyond the original integration window used for gain calculation. **Right:** Distribution of LED arrival times for all PMTs in the same calibration, showing a systematic shift to later times beyond the fixed integration window (red), resulting in incomplete charge integration and underestimated gains.

in the region where only noise contributes. In Figure 3.6 left, a threshold of 5 ADC was used. The scaled noise spectrum is shown in light orange. The SPE spectrum is obtained by subtracting the scaled noise spectrum from the LED spectrum and is shown in the right panel of Figure 3.6. From the SPE spectrum, the survival function, shown in black, is computed as a function of the self-trigger threshold. This cumulative representation is the fraction of SPE signals with amplitudes greater than or equal to a given ADC value. The resulting acceptance varies depending on the PMT gain, noise level, and local threshold settings, typically ranging above 90 % for nominal operation. For the run and channel shown in Figure 3.6, a threshold of 15 ADC yields an SPE acceptance of ~ 94 %.

	Xenon afterpulses	Nitrogen afterpulses	Light emission	Unstable gain	High noise	Lost connection
SR0	177	121, 291, 313, 461	338	28	405	354
		156, 164, 350, 352	362			386
		393, 427				
SR1		442, 108	144			

Table 3.1: PMTs excluded from the SR0 and SR1 analyses, along with the reasons for their removal. In total, 20 PMTs were excluded, meaning that 95 % of the XENONnT PMTs remained operational during SR1.

3.5 Performance during science run one

Following the commissioning and science run zero (SR0) periods [156], XENONnT entered SR1, which acquired over 400 d of data alternating between regular detector and LED calibrations. During the commissioning run, the high voltage of each PMT was optimized through a dedicated scan aimed at keeping the high voltage as low as possible to avoid micro-light emission, while equalizing the gains at $\sim 2 \times 10^6$. This target gain was chosen because, at this operating point, most PMTs achieve an SPE acceptance above 90 % for a self-trigger threshold of 15 ADC samples, with the exception of a few channels affected by increased noise levels. The applied bias voltage was limited to a maximum of -1.5 kV as a strategy to mitigate premature aging of the PMT. During SR1, all PMTs operated at their commissioning bias voltages, with the exception of PMTs 108 and 144, whose voltages were adjusted to mitigate, respectively, a high afterpulsing rate and intermittent light emission. For physics analyses, signals from faulty PMTs are excluded from the event reconstruction chain. This is implemented by setting their gain to zero in the gain correction. During SR0, 17 out of the 494 PMTs were treated in this way due to issues such as elevated electronic noise, unstable gain behavior, light emission, increasing afterpulse rates, or damage to the cable connections. During SR1, three additional PMTs (108, 144, 442) were excluded from analysis. PMT 108 exhibited high afterpulse rate due to nitrogen outgassing. PMT 144 showed persistent light emission, which required a reduction of its bias voltage. This reduced its SPE acceptance to about 50 %, which is well below the 90 % threshold adopted in XENONnT. PMT 442 exhibited a growing afterpulsing rate due to nitrogen outgassing, projected to exceed 4 % by the end of SR1, which is well above the XENONnT threshold of 2 % for excluding a PMT from the analysis. Table 3.1 provides an overview of the PMTs excluded in each science run and the corresponding reasons.




Figure 3.8: Gain (left) and occupancy (right) as a function of the integration window end position. The scan was performed from sample 102 to 144 in steps of 4 ADC samples for five PMTs during a PMT calibration where the LED shift was at its maximum. The gray vertical line marks the end position used during SR0, which produced underestimated gains. The red vertical line indicates the updated endpoint at 132 ADC samples.

3.5.1 Gain stability and LED timing effects

A slow, seasonal time dependence was observed in the arrival time of the LED calibration pulses during SR1. The pulses shifted by several tens of nanoseconds with respect to their nominal timing, causing a fraction of them to fall partially outside the fixed integration window used in the gain extraction. The waveforms of the LED pulse during an exemplary calibration affected by this effect are shown in Figure 3.7 left, together with the distribution of the LED arrival time for all PMTs during the same calibration run. For comparison, Figure 3.3 illustrates the LED arrival time distribution under nominal conditions, where the pulses are contained within the expected integration window. As a result, the integrated LED pulse area and therefore the gains were underestimated during periods when the shift was largest. The median relative gain difference between the underestimated and the correct gains is shown in Figure 3.9, for each PMT array.

To make the gain extraction robust against these shifts, the integration window was enlarged to ensure that the full LED pulse area is always included. Because the LED pulse consistently drifted towards larger ADC samples, the left endpoint of the window was kept fixed at 78 ADC samples. To determine an appropriate right endpoint, we scanned the gain and occupancy for five representative PMTs while varying the right boundary in steps of 4 ADC samples, from 116 to 144, as shown in Figure 3.8. The endpoint must balance the full pulse integration against contamination from spurious signals. Since the gain becomes independent of the integration window boundary at around 130 ADC samples, we selected 132 ADC samples as the new right endpoint.




Figure 3.9: Time evolution of the mean relative gain difference between gains estimated with the original integration window and the corrected gains obtained with the enlarged window. Shaded bands denote the standard deviation of the mean among PMTs in each array. A clear seasonal dependence is observed, with the largest differences occurring in early spring. The broken axis indicates the XENONnT maintenance period with no available calibrations.

The PMT gains of SR1 were recalculated with this enlarged window defined from [78, 132] ADC samples. The largest effect occurs between January and June (see Figure 3.9), where the relative difference reaches up to about 8%. For LED calibrations taken in August, September, and October, the waveforms still exhibit a small timing shift, but the pulses remain fully contained within the original integration window, thus, the gains are not affected by this drift. In this period, a residual offset of roughly 2% is nevertheless present and is understood as a consequence of the larger integration window, which inevitably includes additional baseline and tail charge, slightly increasing the estimated gain.

The origin of the drift in the LED-pulse arrival time remains unknown. Since the time dependence appears to be seasonal, several temperature related hypotheses, both inside and outside the detector, have been considered. One possibility is a temperature gradient along the fibre that transports the LED light from the DAQ room to the XENONnT cryostat, however, no direct correlation has been found with the temperature of any component of the light system. Temperature-induced instabilities in the NIM timing module, which provides the trigger delay between the LED and the DAQ, are also under consideration, although no conclusive evidence has been found.

The final PMT gain evolution of SR1 is shown in Figure 3.10, plotted as the median for each PMT array. The various detector calibration periods are indicated by color bands, including ^{83m}Kr , ^{88}YBe , ^{220}Rn , $^{241}\text{AmBe}$, ^{232}Th and ^{222}Rn . The median gains are mostly




Figure 3.10: Time evolution of the final median gains in SR1 for the top (red) and bottom (blue) PMT arrays. The color vertical bands indicate periods when dedicated detector calibration campaigns were conducted for detector response modeling. The broken axis represents the maintenance period where no calibrations were taken.

stable throughout the science run, within 3 % for most PMTs, as shown in the black histogram of Figure 3.11. Short-term fluctuations during high-rate calibrations, such as the injection of ^{220}Rn shown in the blue vertical band, are visible throughout the science run. These variations are likely connected to voltage changes in the PMTs, which need a certain stabilization time after being biased to their operational voltage [207]. Since these fluctuations do not correspond to a physical change in the gain but arise from the temporary voltage manipulation, they are typically smoothed out by the gain modeling procedure.

The SR1 gain model was produced using the measured gains shown in Figure 3.10 as input, using the method described in Section 3.3.1. Figure 3.11 compares the standard deviation per-PMT of the measured gains with that obtained from the resulting model. Over the science run, the standard deviation of the PMT gain model remains below 3 % for all PMTs, which is in agreement with the 2 % standard deviation reported in SR0 [202]. The discrete peak at zero standard deviation corresponds to the 45 % of PMTs whose modeled gain met the 2.5 % stability requirement, and are therefore modeled with a constant gain. The position distribution of the gain standard deviation is shown in Figure 3.12, where one can see that the majority of the averaged PMTs are located in the bottom array. The PMTs in the top array show a larger standard deviation over the science run. Several of these PMTs are located near a region where sporadic S2 emission from single or few electrons occurs, suggesting that the observed gain variations are driven by the large light intensities observed during these periods




Figure 3.11: Distribution of the relative standard deviation of the measured and modeled PMT gains. Most PMTs exhibit gain stability better than 3% over the 400 d period. The spike at zero in the model distribution corresponds to the fraction of PMTs whose modeled gain satisfied the averaging criteria.

that spontaneously appeared during the data taking periods of SR1. However, as shown in Figure 3.10, after short-term fluctuations the gains typically return to their nominal values in the subsequent weeks, indicating that there is no long-term degradation.

The SPE acceptance during SR1 was similarly stable over time, with the median value per calibration remaining above 93% across all PMTs for a threshold of 15 ADC samples. No significant fluctuations were observed throughout the science run, nor when compared to the corresponding SR0 results [202].

3.5.2 Lone hit evolution

As described in Section 2.2.2, the basis of the data analysis are the S1 and S2 signals seen in the top and bottom photosensor arrays. To reconstruct these signals, the processing chain begins with the identification of *hits* (see Figure 4.2). A hit is defined as a time interval in which the waveform of an individual PMT exceeds its self-trigger threshold. This time interval extends from 30 ns before and 200 ns after the point where the threshold was surpassed. Hits that occur within 700 ns, regardless of the PMT in which they originate, are grouped into *peaklets*. In this first clustering stage, hits that are isolated from any other hit within this time window are stored and handled separately. These are called *lone hits* and are produced by single photoelectrons of diverse sources including dark counts and small contributions from scintillation light. Since their waveform corresponds to the response of a single detected photon, the area




Figure 3.12: Spatial distribution of the modeled PMT gain standard deviation across the PMT arrays, showing localized variations in gain stability.

of such hits is expected to follow the SPE distribution and therefore to be centered around 1 PE.

The time evolution of the lone-hit area is monitored as an independent proxy for the PMT gain stability. At this early stage of the reconstruction, the gain correction is the only one applied to the data. Therefore, any variation in the mean lone-hit area directly reflects changes in the effective PMT gain. To quantify this, the area distribution of lone hits is fitted with a Gaussian function over the interval 0.7–1.3 PE. The mean of the Gaussian fit is extracted for each PMT in each data-taking run, and the resulting values are used to construct the time-evolution curves.

During SR0, a time-dependent variation in the area distribution of lone hits was observed [202]. This variation exhibited both a long-term trend and several short-term fluctuations. The short-term behavior was found to correlate with specific detector conditions, including periods of high-rate calibration, data-taking in the absence of an electric field, and intervals shortly after turning PMTs on. The long-term trend was not fully understood. However, it showed a correlation with the temporal evolution of the light yield measured using α -particles at 5.5 MeV from the ^{222}Rn background.

The time evolution of the PMT gain did not reproduce the trends observed in the lone-hit population. This discrepancy is attributed to the fact that PMT gain calibrations are performed using LED light at wavelengths around 470 nm, whereas xenon scintillation light peaks in the vacuum ultraviolet range near 175 nm. As a consequence, variations affecting the PMT response to xenon scintillation may not be adequately captured by LED-based gain measurements. Given these observations, an empirical correction




Figure 3.13: Time evolution of the median lone-hit area throughout SR1. Short-term fluctuations coincide with calibration periods. The long-term trend, correlated with the light yield time evolution, was removed at a subsequent reconstruction stage via an effective light-yield correction and thus does not impact the XENONnT science data.

was derived directly from the lone-hit data and subsequently applied to the S1 signal correction. This procedure effectively removed the long-term time dependence observed in the light-yield evolution, resulting in a stable detector response throughout SR0.

In SR1, the monitoring of lone hits was carried out in the same manner as in SR0, the time evolution of the median per array is shown in Figure 3.13. Short-term variations were again observed, particularly during high-rate calibration periods such as the ^{220}Rn source injection. These variations are attributed either to the time required for the PMTs to stabilize after being ramped down to zero voltage and turned back on or to the possibility that PMTs switched off during the injection experience some degree of window charge-up. A long-term trend was also present, which was again correlated with the evolution of the light yield. At the time the lone-hit area evolution for SR1 was analyzed, a time-dependent light-yield correction for the S1 signal (see Section 4.3) had already been implemented to account for photoionization-induced contributions. As a result, the long-term trend in the light yield observed in SR1, was effectively removed. Consequently, no additional empirical correction based on lone hits was required for SR1.

Up to the end of SR1, the treatment of lone hits relied on modeling their area distribution with a simple Gaussian function. This approach has several limitations: it does not account for double-photoelectron emission, nor does it incorporate contributions from PMT noise, afterpulses, or light emission. This simplified model can bias the fitted parameters, and the inferred time-dependent behavior of the lone-hit population

may not fully reflect the true detector response. To address this, a more comprehensive model has been developed that includes all of these contributions and, in addition, provides improved monitoring of PMT noise, afterpulsing behavior, SPE acceptance, and self-trigger thresholds. This framework represents a step forward toward understanding the origin of the observed time-dependent trends in the lone-hits, and further studies will be required to fully understand the detector effects involved.

3.6 Summary and outlook

In this chapter, the calibration and monitoring of the XENONnT photosensors were presented, with emphasis on the determination of the PMT gain and its evolution over time. During SR1, the gains were extracted using the established model-independent method [201] and modeled with the custom time-dependent algorithm developed for XENONnT [202], providing an accurate description of the PMT response. The seasonal drift of the LED pulse timing observed during SR1 was identified and corrected.

This well-characterized light response forms the basis for the reconstruction of the S1 and S2 signals used in the physics analyses. The next chapter therefore presents the XENONnT SR1 science results, including the search for WIMP dark matter and the measurement of solar neutrinos via CE ν NS.

Chapter 4

Science Results from the XENONnT Experiment

XENONnT is primarily a dark matter detector. Its large target mass, low background level, and low energy threshold combined with ER/NR discrimination enable the study of other rare processes using a common dataset. These analyses share the event reconstruction framework [156], while differing in their event selection, fiducialization, and statistical interpretation [89, 122, 208]. In the previous chapter, the photosensors employed in XENONnT were presented, together with the procedure for converting raw signals into PE within the event reconstruction pipeline.

Building on this foundation, this chapter is organized as follows: Section 4.1 provides an overview of the data sets collected during the first science runs of XENONnT. Section 4.2 presents an overview of the event reconstruction pipeline and Section 4.3 outlines their correction procedures. The selection criteria are summarized in Section 4.4. Finally, the main scientific results are presented in Sections 4.5, 4.6, and 4.7. These correspond to NR searches for WIMP–nucleon interactions and ^8B CE ν NS, and to ER searches for new physics, all performed in a blind analysis approach.

4.1 Science data

XENONnT data acquisition is organized into periods corresponding to different detector conditions and configurations. The evolution of the accumulated livetime is shown in Figure 4.1, and the main features of each period are described below.

SR0: This science run spanned from May 1 to December 12, 2021 and achieved the lowest ER background level reported for xenon detectors [122]. During this period, a dedi-




Figure 4.1: Accumulated livetime during the SR0 and SR1 data-taking periods as a function of calendar time. Science and calibration periods are indicated by vertical color bands. The bottom panel shows the sub-division of each science run into their respective sub-periods and conditions. Figure courtesy of the XENON Collaboration.

cated data-taking mode, referred to as tritium-enhanced data (TED), was implemented. This mode was motivated by the low-energy excess observed in XENON1T [209] and involved bypassing the getter to test the hypothesis of an unaccounted ^3H -like background contribution. The analysis of TED data did not reveal a significant ^3H excess [122]. TED is excluded from the WIMP search analysis. A short circuit between the cathode and the bottom screening mesh limited the cathode voltage to -2.75 kV , instead of the design value of -30 kV . This resulted in electrode potentials of $V_{\text{anode}} = 4.9\text{ kV}$, $V_{\text{gate}} = 0.3\text{ kV}$, and $V_{\text{cathode}} = -2.75\text{ kV}$, corresponding to a drift field of 23 V/cm and an extraction field of 2.9 kV/cm [122, 210]. The initial measured ^{222}Rn activity concentration was $3.5\text{ }\mu\text{Bq/kg}$. After combined liquid and gas distillation (see Chapter 2), the concentration was reduced to below $1\text{ }\mu\text{Bq/kg}$, corresponding to a suppression by more than a factor of three [142]. After online distillation, the reached krypton concentration was $^{\text{nat}}\text{Kr/Xe} = (5.6 \pm 3.6) \times 10^{-14}\text{ mol/mol}$, well below the design goal of $1.0 \times 10^{-13}\text{ mol/mol}$ [142].

SR1: This science run took place between May 19, 2022 and August 8, 2023. Elevated ER background levels from ^{85}Kr and ^{37}Ar at the beginning of the SR led to its subdivision into SR1_a and SR1_b. The latter corresponds to the period following a cryogenic distillation campaign, which reduced these backgrounds to subdominant levels [55]. In parallel, the radon distillation system was operated at full capacity, resulting in a reduction of the ^{214}Pb background by a factor of 1.9 (SR1_a) and 1.7 (SR1_b) relative to SR0 [55]. Additionally, the liquid level was lowered by 0.2 mm and the anode voltage increased by 50 V compared to SR0. No changes in the drift field were made. These adjustments were implemented to mitigate the occurrence of so-called hot spots, i.e., localized bursts of single-electron (SE) emission originating from the top electrodes. In addition, a higher rate of small S2 signals after high-energy events was observed. This phenomena is attributed to photoionization of non-electronegative impurities, possibly introduced by the change of mode of the radon distillation column [55].

The calibration data collected during these science runs, indicated by the vertical color bands in Figure 4.1, were introduced in Section 2.2.5.

4.2 Event reconstruction

Event reconstruction is performed using the open-source framework **straxen** [206]. As outlined in Section 3.5.2, the process begins with the identification of PMT *hits*, defined as time intervals in which a PMT pulse exceeds its self-trigger threshold. Each hit window extends 30 ns before the threshold-crossing time and 200 ns after. Hits from any PMT are grouped together if they occur within 700 ns. These clustered hits are iteratively split into sub-clusters based on their timing information and the summed waveform of all hits in the cluster. The sub-clusters are called *peaklets*, which are classified as S1 or S2 *peaklets* according to their area, rise time, width and the fraction of the area observed by the top PMT array (AFT). An additional characteristic is evaluated: the *tight coincidence*, defined as the number of PMTs recording a hit within ± 50 ns of the peaklet's maximum amplitude. The required value depends on the analysis; for example, a tight coincidence of 3 is used in the WIMP search, while a value of 2 is adopted for the CEνNS analysis (see Section 4.6). A higher tight coincidence reduces the rate of spurious S1 signals but increases the energy threshold.

At the hit-cluster splitting stage, a physical S2 may be incorrectly split into multiple smaller peaklets, particularly for interactions at large drift distances where longitudinal diffusion broadens the signal. To account for this effect, peaklets that are sufficiently close in time are merged, under the assumption that they originate from the same physical S2. After this step, the merged S2 and the S1 peaklets are referred to as *peaks*.




Figure 4.2: Schematic of the XENONnT event reconstruction pipeline applied in **straxen**. Starting with PMT pulses, events are built through a sequence of intermediate objects (blue tiles). Peaklets that do not satisfy the S1 or S2 criteria are indicated with a question mark. The main processing steps are indicated in gray. Adapted from Ref. [156].

An event is defined as a time interval centered on an S2 peak and containing its corresponding S1 signal. The S2 is required to have an area ≥ 100 PE, and fewer than eight additional peaks must be present within a ± 10 ms time window. The largest S1 peak preceding the S2 within 2.45 ms is assigned as the main S1 of the event. The second-largest S1 and S2 signals are referred to as the alternative S1 and S2.

Once events are reconstructed, the three-dimensional interaction position is reconstructed. The vertical position (z_{obs}) of an event is obtained by multiplying the electron drift velocity, measured to be (0.675 ± 0.006) mm/ μ s in an electric field of 23 V/cm, by the time difference between the S1 and S2 signals. The horizontal position ($x_{\text{obs}}, y_{\text{obs}}$) is reconstructed using the S2 light pattern observed by the top PMT array, for which machine learning algorithms are employed.

The reconstruction described above provides the basis for all analyses presented in this chapter.

4.3 Signal corrections

Spatial and temporal detector effects impact the reconstruction of S1 and S2 signals and must be accurately characterized to ensure a uniform detector response across the full volume and over time. First, reliable position reconstruction is essential for applying corrections.

The non-uniform electric field introduces biases in position reconstruction, particularly near the TPC walls [172]. Field inhomogeneities have two main consequences: first, they lead to position-dependent photon and electron yields at fixed energy due to variations in the recombination process; second, they distort the drift paths of electrons, causing a mismatch between the reconstructed S2 position and the true interaction vertex. This effect manifests as a depth-dependent inward bias in the reconstructed radial coordinate. To correct for this, a three-dimensional data-driven field distortion correction (FDC) map is derived using $^{83\text{m}}\text{Kr}$ calibration data, which provides a monoenergetic and spatially uniform source of events. The correction is defined as a function of the observed position and is applied to recover a spatially uniform event distribution. In practice, a correction term is added to the observed radius to obtain the corrected coordinate, while the Cartesian coordinates are obtained through trigonometric calculations. A complete description of the FDC map is provided in Ref. [156, 211].

The position dependence of the S1 signal in the TPC arises from both geometric effects and electric field inhomogeneities. Geometric effects are primarily due to variations in photon path length and in the number of reflections that photons undergo before reaching a PMT. In practice, the position dependence of S1 is reflected in the light yield (LY) defined as the number of detected photoelectrons per unit deposited energy. Using Equations 2.6 and 2.9, the LY can be written as:

$$\text{LY}(\varepsilon(x, y, z), E) = \frac{\text{S1}(\varepsilon(x, y, z), E)}{E} = \text{LCE}(x, y, z) \cdot \text{QE} \cdot \text{CE} \cdot P_y(\varepsilon(x, y, z), E), \quad (4.1)$$

where LCE is the light collection efficiency, ε the drift field, E the energy deposited, QE is the quantum efficiency, CE is the collection efficiency and P_y is the photon yield, defined as the number of photons produced per unit energy deposited (ph/E). The spatial dependence of the light yield is therefore driven by two contributions: the LCE and the P_y . The LCE encodes purely geometrical effects and constitutes the dominant source of position dependence, while P_y depends on the local electric field and therefore reflects its spatial non-uniformities.

The goal of the S1 correction is to remove these position-dependent effects and recover a response that depends only on the deposited energy. The correction is formulated in terms of the LY and it is derived in two steps. First, the field-dependent component is removed using a detailed simulation of the electric field in the TPC [172, 211], combined with a model for the field dependence of the LY [212]. This step isolates the geometrical contribution. Subsequently, a relative light yield correction (LY_{rel}) is derived to equalize

the detector response across the fiducial volume. This correction is obtained using $^{83\text{m}}\text{Kr}$ calibration data. A detailed description of the correction methodology is given in Section 5.4.1, where a simulation-driven derivation is presented. The same procedure is applied for the data-driven correction.

In SR0, an additional time-dependent correction was applied to the S1 signal, namely the lone-hit correction discussed in Section 3.13. This correction is derived from the time evolution of the median lone-hit area, normalized to a reference value corresponding to science data-taking period, during which the variation of the lone-hit area was small compared to that seen in calibration periods. The combined effect of the position- and time-dependent corrections is to stabilize the LY over time and across the TPC.

In SR1, the time evolution of the S1 response was treated differently, mainly due to the addition of photoionization whose level varies in time. An anti-correlation between the light yield and the level of photoionization over time was observed. This effect was corrected using α decays from the ^{222}Rn chain present in the science data [55]. The correction is implemented as a normalization factor to the mean light yield over the full science run. Overall, the corrected S1 is:

$$\text{cS1} = \frac{\text{S1}(x, y, z, t)}{C_{x,y,z} \cdot C_t}, \quad (4.2)$$

where $C_{x,y,z}$ is the position-dependent correction and C_t is the time-dependent correction.

In analogy to the S1 LY, the position dependence of the S2 signal can be expressed in terms of the S2 charge yield (CY), defined as the S2 area per unit deposited energy:

$$\text{CY}(\varepsilon(x, y, z), E) = \frac{\text{S2}(\varepsilon(x, y, z), E)}{E}, \quad (4.3)$$

which depends on both the local electric field and the position-dependent detector response. First, the S2 signal is corrected for variations in the electric field. This correction rescales the signal to a reference field, resulting in a field-independent S2. Second, residual variations in the detector response across the (x, y) plane are corrected by equalizing the average S2 response across the TPC after removing drift-time effects. The position-independent $\widetilde{\text{S2}}$ signal is obtained as:

$$\widetilde{\text{S2}} = \frac{\text{S2}(x, y, z, t_d)}{C_{\text{CY}}(x, y) \cdot C_{xy}(x, y)}, \quad (4.4)$$

where C_{CY} accounts for field-dependent variations in the CY, and C_{xy} corrects for residual position-dependent variations in the S2 response across the (x, y) plane, arising from non-uniform light collection, extraction, and amplification in the gas phase. After these corrections, the S2 signal depends predominantly on the drift time through the electron survival probability, parametrized by Equation 2.11. The electron lifetime is extracted by fitting the $\widetilde{S2}$ signal as a function of drift time. The resulting lifetime $\tau(t)$ is then used to correct for electron losses during drift through a multiplicative factor. In SR0 and SR1, electron lifetimes of $\sim 10\text{--}30$ ms were achieved [156], which are much longer than the maximum drift time of 2.2 ms at the drift field 23 V/cm.

After accounting for the electrons that reach the liquid–gas interface, one must consider the extraction efficiency into the gas phase, as only extracted electrons contribute to the S2 signal via proportional scintillation. This introduces additional detector-dependent efficiencies in the reconstructed S2 signal. In particular, the correlation between the number of extracted electrons and the observed S2 signal is governed by the single electron gain (G), which characterizes the detector response to individual electrons. In addition, not all electrons reaching the liquid surface are successfully extracted into the gas phase. The extraction efficiency (ϵ_{ext}) can vary over time, particularly following changes in detector conditions such as anode voltage adjustments. These effects are monitored and quantified using calibration data. The G is determined from isolated single electron signals, whereas ϵ_{ext} is inferred from monoenergetic $^{83\text{m}}\text{Kr}$, for which the number of produced electrons is expected to be constant. Any observed variation in the reconstructed number of electrons over time is therefore attributed to changes in detector response and corrected accordingly. The corrected S2 is written as:

$$\text{cS2} = \widetilde{S2} \frac{\exp(t_d/\tau(t))}{\epsilon_{\text{ext}}(t, x, y) \cdot C_G(t, x, y)}, \quad (4.5)$$

where $C_G(t, x, y)$ is the correction factor obtained from G variations. Having obtained the corrected S1 and S2 signals, the energy is reconstructed using Equation 2.10. The detector parameters g_1 and g_2 are determined from calibration data by exploiting the anti-correlation between light and charge yields for several monoenergetic peaks. Here, g_1 denotes the photon detection gain expressed in photoelectrons per photon (PE/ph), while g_2 denotes the electron gain in photoelectrons per extracted electron (PE/ e^-) in SR0 $g_1 = (0.1515 \pm 0.0014)$ PE/ph and $g_2 = (16.5 \pm 0.6)$ PE/ e^- [156], in SR1 $g_1 = (0.13670 \pm 0.00001)$ PE/ph and $g_2 = (16.9 \pm 0.5)$ PE/ e^- [55].

4.4 Event selection

After applying the signal corrections, a set of selection criteria is introduced to isolate well-reconstructed and physically meaningful events, while suppressing background contributions. The specific selection depends on the physics channel under consideration. In this section, an overview of the selection criteria adopted for the WIMP and CE ν NS analyses is presented.

The MV and NV, which surround the TPC (see Sections 2.2.4 and 2.2.3), provide the means to tag interactions in the TPC that occur in coincidence with signals in these external veto systems. These vetoes are primarily used to reject cosmogenic muons and radiogenic neutrons that interact in the TPC and trigger in coincidence the NV.

Single-photon PMT hits and few-electron S2 signals can appear up to one second after high-energy interactions. These delayed signals have been observed to exhibit spatial and temporal correlations with preceding energetic events [153, 213, 214]. Although not physically meaningful, they can form events through random pairing of S1 and S2 signals, leading to so-called accidental coincidences (AC). These backgrounds are mitigated using dedicated selection criteria. In particular, the S2 time shadow, computed from the size of the preceding S2 signal divided by the time elapsed since that event, is used to suppress delayed signals following energetic interactions. In addition, local rates of small peaks (“ambience”) are reduced through an ambience selection, and single-electron signals induced by hot-spot activity are identified and rejected based on their reconstructed position.

The S1 and S2 quality cuts remove unphysical events based on their distribution in specific parameter spaces. For the S1 signal, several criteria are applied. Few-electron signals or afterpulses misidentified as S1s are rejected using the S1 width selection, as genuine S1 signals are characterized by a narrow time profile. Events in which an excessive fraction of the S1 light is recorded by a single PMT are removed, as these are typically associated with PMT flashes or afterpulses.

In addition, events with multiple photons detected in the same PMT are rejected, given the low probability for a true interaction to produce such a signature. Abnormal S1 light patterns, arising from reconstruction artifacts or unresolved multiple scatters, are suppressed. Finally, events with an unusually large fraction of S1 light observed in the top array are rejected, as they are typically associated with interactions occurring in the gas phase.

For the S2 signal, similar selections based on the hit pattern and the area fraction top (AFT) are applied. In addition, events with nonphysical drift times are rejected

Selection	Concept	WIMP	CE ν NS	Low ER
Livetime				
Muon veto	Rejects events in coincidence with a MV trigger	✓	✓	✓
Neutron veto	Rejects events in coincidence with a NV trigger	✓	✓	✓
Shadow of high-energy peaks	Rejects signals correlated in time and space with high-energy events	✓	✓	✓
Ambience	Rejects peaks nearby in time or space by many small peaks	✓	✓	✓
Hot spot veto	Rejects events correlated in time and space with hot spot location	✓	✓	
S1 - S2 quality criteria				
S1 width	Rejects signals with excessively large S1 width	✓		✓
S1 max. PMT area	Rejects S1 signals where a large fraction of the light is detected by a single PMT	✓	✓	✓
S1 per-PMT hits	Rejects S1 signals with multiple hits detected in the same PMT		✓	
S1 hit pattern	Rejects S1 signals with an abnormal hit pattern in the top PMT array given its reconstructed position	✓		✓
S1 AFT	Rejects S1 signals with a top-array light fraction inconsistent with the reconstructed position	✓		✓
S2 hit pattern	Rejects S2 signals with an abnormal hit pattern in the top PMT array	✓	✓	✓
S2 AFT	Rejects S2 signals with an abnormal fraction of light detected in the top PMT array	✓	✓	✓
S2 width	Rejects S2 signals with widths inconsistent with the expectation from longitudinal diffusion	✓	✓	✓
Single-scatter				
S1 single scatter	Rejects events without a unique S1–S2 pairing for the largest S2 signal	✓		✓
S2 single scatter	Rejects multi-scatter events	✓	✓	✓
Other				
S2 BDT	Rejects AC events using S2 shape parameters	✓	✓	
S1 BDT	Rejects AC events using S1 properties	✓	✓	
Fiducial volume	Rejects external radiation and poorly understood regions in the TPC	✓	✓	✓

Table 4.1: Overview of selection criteria, their concept and application to the WIMP, CE ν NS or low-energy ER analyses.

using the S2 width selection. The expected S2 width is modeled assuming Gaussian longitudinal diffusion, for which the variance increases linearly with drift time. Events are required to lie within the 1st and 99th percentiles of the normalized S2 width distribution as a function of S2 area. This selection efficiently removes events occurring in the gas phase, ACs, or multiple S2 signals merged into a single pulse. Additionally, Boosted Decision Tree (BDT) classifiers are trained on signal shape parameters in order to reduce the AC background.

WIMPs and neutrinos have a negligible probability of producing multiple scatters within the active volume, as expected for processes with very low interaction cross-sections. In contrast, many background sources can produce multiple interactions within the detector. For this reason, a single-scatter selection is applied. The selection evaluates whether the size of an alternative S2 is physically compatible with a single-scatter event by comparing it to the main S2 size in the parameter space of alternative S2 area versus main S2 area. The cut is parameterized as the sum of two Fermi–Dirac functions, as shown in Figure 4.3. These functions smoothly interpolate between a low-S2 and a high-S2 regime, with a transition around 3×10^4 PE. The resulting function defines the maximum allowed alternative S2 area as a function of the main S2 area. Events with an alternative S2 area below this threshold are classified as single scatter, while events above it are rejected as likely multiple-scatter events. The parameters of the model are derived from ^{220}Rn calibration data.

Finally, an inner volume of the TPC, referred to as the fiducial volume (FV), is defined to reduce external backgrounds by exploiting the self-shielding of LXe (see Table 2.1), whereby radiation is attenuated toward the detector interior. This selection suppresses gamma rays and neutrons entering the detector as well as surface backgrounds. The FV definition depends on the analysis: for the WIMP and low-energy ER searches it is optimized for signal-to-background ratio [55, 122], while for the CE ν NS analysis it is chosen such that surface background contributions can be neglected in the statistical inference [89].

4.5 Nuclear recoil dark matter results

The search for WIMP-nucleon interactions, both for spin-independent (SI) and spin-dependent (SD), is performed in the (cS1, cS2) parameter space. The region of interest (ROI) of both SR0 and SR1 is defined as: $\text{cS1} \in [0, 100]$ PE and $\text{cS2} \in [126, 12589]$ PE, as shown in Figure 4.5. A threshold of $\text{S2} > 200$ PE (SR0) and $\text{S2} > 320$ PE (SR1) is applied. All the events with a reconstructed energy below 20 keV and within the ER and NR contours are blinded. The signal detection efficiency at different stages (peak




Figure 4.3: Distribution of the ratio $S2_{alt}/S2$ as a function of S2 area for ^{220}Rn calibration data. The dashed line indicates the S2 single-scatter selection, separating single-site from multi-site events. Figure from Ref. [156].

reconstruction, event building and data selection) shown in Figure 4.4 as a function of recoil energy, are obtained using a data-driven approach validated with simulations [55].

The detector response to NR and ER is calibrated using dedicated datasets. The NR response model is obtained from $^{241}\text{AmBe}$ calibration data, using a clean sample of neutron events selected by tagging coincident gamma rays emitted alongside the neutrons [156]. The ER response model is derived from ^{220}Rn calibration data, primarily using the β decays of ^{212}Pb . In both cases, a LXe emission model, parametrized within the Noble Element Simulation Technique (NEST) [148] model, is fitted to the data to determine the detector response. These models provide a mapping between true deposited energy and the corresponding distribution in the observable (cS1, cS2) space for NR and ER interactions. The NR response is then folded with the WIMP recoil energy spectrum, yielding a predicted signal distribution in observable space. The WIMP spectrum is computed using the Helm form factor for SI interactions [47]. The same NR response model is also used to describe neutron-induced backgrounds, with their respective recoil spectra. The ER response model is folded with an ER spectrum to predict how the dominant background populate the observable space.

The background model includes five components, whose expected distribution in (cS1, cS2) space is shown in Figure 4.5. A summary of the expected rates for each background component is presented in Table 4.2 using a combined exposure of $3.1 \text{ t} \cdot \text{yr}$.

The most dangerous NR background are radiogenic neutrons, as they are largely indistinguishable from the expected WIMP signal in the observable parameter space. After selections, the total neutron expectation over the two SRs is (1.87 ± 0.47) events, which




Figure 4.4: Left: Detection efficiencies as a function of NR energy for SR1a (dashed) and SR1b (solid). The peak reconstruction efficiency (green) is determined by the S1 threefold tight coincidence requirement. Including event building and data selection (blue), and finally the ROI selection (black), reduces the overall efficiency. The efficiency exceeds 10 % in the energy ranges $[0.6, 13.5]$ keV_{ER} and $[3.8, 64.1]$ keV_{NR} . The expected recoil spectra for three WIMP masses, shown in orange, are displayed without efficiency corrections. The upper axis shows the corresponding conversion from NR to median ER energy. Figure from Ref. [55].

is estimated through full-chain simulations¹ and a custom framework based on the NR fitted model [215, 216]. Other NR backgrounds originate from solar and atmospheric neutrinos via CEvNS, which are subdominant. The dominant background is from ER interactions, with a total rate of 936 ± 35 events.

The AC background is another relevant contribution to the WIMP search due to its significant overlap with the signal ROI in the observable space, as shown in Figure 4.5. Its rate is estimated using a data-driven approach based on the measured rates of isolated S1 and S2 signals and the coincidence window used in event reconstruction. To model this background, isolated S1 and S2 waveforms are sampled from data, assigned random drift times, and combined into artificial events, which are processed through the standard reconstruction and selection chain. The resulting sample is used to construct templates in (cS1, cS2) space and to estimate the corresponding background rate. The surface background arises from β decays of ^{210}Pb originating on the TPC walls. It is modeled in the (cS1, cS2, R, z) parameter space using data from a sideband region outside the WIMP ROI. The spatial and signal distributions are derived from events reconstructed near or outside the TPC wall, and the resulting model is validated using events outside the fiducial volume, showing good agreement.

¹See Chapter 5 for an overview of the simulation framework.




Figure 4.5: Dark matter search data in the $(cS1, cS2)$ space. Each event is represented by a pie chart indicating the relative contribution of the best-fit model, including a $200 \text{ GeV}/c^2$ WIMP signal (orange contour). The size of each pie chart is proportional to the expected signal density at that location. The background distributions are shown for ERs (blue), ACs (purple), and surface events (green). The neutron background (yellow in the pie charts) has a distribution similar to that of the WIMP signal. Adapted from Refs. [55] and [210].




Figure 4.6: Upper limits on the spin-independent WIMP–nucleon cross section at 90 % confidence level as a function of WIMP mass (black line). The expected sensitivity is shown by the 68 % (green) and 95 % (yellow) bands, corresponding to the range of limits obtained under the background-only hypothesis, with the median expectation indicated by the dotted line. Results from XENONnT using SR0 data only [210], as well as from LZ [128] and PandaX-4T [51], are shown for comparison.

	SR0		SR1 _a		SR1 _b	
	Nominal	Best-fit	Nominal	Best-fit	Nominal	Best-fit
ER (flat)	134	(136 ± 12)	(430 ± 30)	(450 ± 20)	(151 ± 11)	(154 ± 10)
ER (³ H-like)	–	–	62	(40 ± 30)	101	80 ⁺¹⁸ ₋₁₇
ER (³⁷ Ar)	–	–	(58 ± 6)	(55 ± 5)	–	–
Neutron	(0.7 ± 0.3)	(0.6 ± 0.3)	(0.47 ± 0.19)	(0.45 ± 0.19)	(0.7 ± 0.3)	(0.7 ± 0.3)
CEvNS (solar)	(0.16 ± 0.05)	(0.16 ± 0.05)	(0.010 ± 0.003)	(0.010 ± 0.003)	(0.019 ± 0.006)	(0.019 ± 0.006)
CEvNS (atm.)	(0.04 ± 0.02)	(0.04 ± 0.02)	(0.024 ± 0.012)	(0.024 ± 0.012)	(0.05 ± 0.02)	(0.05 ± 0.02)
AC	(4.3 ± 0.9)	4.4 ^{+0.9} _{-0.8}	(2.12 ± 0.18)	(2.10 ± 0.18)	(3.8 ± 0.3)	(3.8 ± 0.3)
Surface	(13 ± 3)	(11 ± 2)	(0.43 ± 0.05)	(0.42 ± 0.05)	(0.77 ± 0.09)	(0.76 ± 0.09)
Total	152	(152 ± 12)	553	(550 ± 20)	257	(239 ± 15)
WIMP	–	1.8	–	1.1	–	2.1
Observed	152		560		245	

Table 4.2: Expected event yields for the nominal and best-fit models in SR0, SR1_a, and SR1_b, corresponding to a total exposure of 3.1 tonne-years, including an unconstrained WIMP signal with a mass of 200 GeV/ c^2 . The uncertainties in the nominal predictions correspond to the widths of the Gaussian constraints applied in the likelihood. Adapted from Ref. [55].

A blind analysis strategy was adopted. The search data are statistically interpreted using an unbinned extended likelihood constructed from three-dimensional models in cS1, cS2, and R, following the procedure described in Ref. [156]. These models are generated using toy Monte Carlo simulations based on the fitted yields and detection efficiencies.

In total, 152 (SR0) and 805 (SR1) events were observed within the ROI, of which 16 (SR0) and 27 (SR1) fall within the blinded region. No significant excess over the expected background is found. A profile log-likelihood ratio test is therefore used to set constraints on the WIMP–nucleon interaction cross section over the range of WIMP masses considered, for both SI and SD interactions, as shown in Figure 4.6. For SI interactions, the most stringent upper limit is $1.7 \times 10^{-47} \text{ cm}^2$ at a WIMP mass of 30 GeV/ c^2 [55].

4.6 ⁸B solar neutrino detection

Using the selection thresholds of the WIMP search, the expected CEvNS rate is only $\sim 0.17 \text{ t}^{-1} \text{ yr}^{-1}$ [89], corresponding to fewer than one event in the combined SR0 and




Figure 4.7: Detection efficiency for low-energy NRs in XENONnT and the expected ^8B CE ν NS spectrum. The light (dark) blue curve shows the efficiency for detecting S1 (S2) signals, and the black curve shows the combined efficiency. The CE ν NS spectrum is shown before (light red) and after (dark red) accounting for detector efficiency. The green curve shows a scaled ^{88}YBe calibration spectrum. Figure from Ref. [89].

SR1 exposure. To enhance the signal yield, the selection is relaxed by reducing the tight coincidence requirement from 3 to 2 PMT hits for the S1 signal, and by lowering the S2 threshold from 200 to 120 PEs. The resulting region of interest is defined by S1 signals with 2–3 PMT hits and $\text{cS2} \in [120, 500]$ PE. This lower threshold increases the expected CE ν NS event rate by a factor of ~ 17 compared to that obtained with the WIMP search selection [89]. The S1 and S2 acceptances together with the ^8B CE ν NS spectrum is shown in Figure 4.7.

Accessing this regime requires a precise understanding of the detector response to very low-energy NRs. For this purpose, a dedicated calibration was performed using an external ^{88}YBe neutron source [188], which emits quasi-monoenergetic neutrons of 152 keV, producing NRs up to about 3 keV. The observed data are well described by simulations, and the light and charge yields are extracted by fitting NEST [148] parametrization.

At such low thresholds, accidental coincidences constitute the dominant background. Prior to any suppression, the rates of isolated signals are approximately 15 Hz for S1 and 0.15 Hz for S2, resulting in ~ 400 events per day [89].

In addition to the quality selections summarized in Table 4.1, ACs are further suppressed using the S1 and S2 BDT selections and the S2 time shadow variable. The latter is introduced to identify spurious small S2 signals correlated with preceding large S2 pulses from earlier interactions. It quantifies the presence of such correlations within

	Nominal	Best-fit
AC (SR0)	7.5 ± 0.7	7.4 ± 0.7
AC (SR1)	17.8 ± 1.0	17.9 ± 1.0
ER	0.7 ± 0.7	$0.5^{+0.7}_{-0.6}$
Neutron	$0.5^{+0.2}_{-0.3}$	0.5 ± 0.3
Total background	$26.4^{+1.4}_{-1.3}$	26.3 ± 1.4
^8B	$11.9^{+4.5}_{-4.2}$	$10.7^{+3.7}_{-4.2}$
Observed	37	

Table 4.3: Expected and best-fit number of events for signal and background components in the ROI. The uncertainty on the expected signal includes contributions from the signal detection efficiency, as well as the LY and CY. The uncertainties on the background expectations reflect the widths of the Gaussian constraints used in the fit, while the ^8B signal is left unconstrained. Adapted from Ref. [89].

a defined time window, which allows the rejection of delayed, non-physical signals.

The AC model is validated using AC-rich datasets constructed from events passing all selection criteria except those targeting ACs and an independent calibration dataset using ^{37}Ar . The predicted and observed distributions show good agreement across all analysis dimensions, as confirmed by four-dimensional goodness-of-fit tests.

The subdominant background components, ER and radiogenic neutrons, together with the $\text{CE}\nu\text{NS}$ signal are modeled and included in a combined fit. The search is performed in a binned four-dimensional space: cS2, S2 time shadow, S1 BDT and S2 BDT. The total expected event rates for a $3.5 \text{ t} \cdot \text{yr}$ exposure are summarized in Table 4.3.

A total of 37 events were observed, compared to an expected background of $26.4^{+1.4}_{-1.3}$ events. The resulting excess of 10.6 events is consistent with the predicted ^8B $\text{CE}\nu\text{NS}$ signal of $11.9^{+4.5}_{-4.2}$ events.

From the statistical inference, the background-only hypothesis is disfavored at the 2.3σ level. The inferred ^8B solar neutrino flux is $4.7 \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$, in good agreement with previous measurements, such as that reported by Sudbury Neutrino Observatory (SNO) [217]. The profile likelihood as a function of the assumed ^8B flux shows a best-fit value consistent with earlier results from PandaX-4T [218] and XENON1T [219], and within the expected flux range. Interpreted as a measurement of the $\text{CE}\nu\text{NS}$ cross section, the result is fully consistent with the Standard Model prediction for neutrino–nucleus scattering in xenon [84, 85]. This measurement provides the first evidence of solar neutrino $\text{CE}\nu\text{NS}$ interactions in a LXe detector.




Figure 4.8: Measurement of the solar ^8B neutrino flux. **Top:** the 68% (90%) confidence interval obtained in this work is shown in black (gray). The 68% CL result from SNO [217], and the 90% CL upper limits from XENON1T [219] and PandaX-4T [218], are also shown for comparison. **Bottom:** the profile likelihood ratio test statistic q_μ as a function of the assumed ^8B neutrino flux (solid red line). The confidence intervals are derived using the Feldman–Cousins construction at 68% (black) and 90% (gray) confidence level. Figure from Ref. [89].

4.7 Low energy ER searches

Three channels are of interest in the low-energy ER analysis: solar axions, solar neutrinos with an enhanced magnetic moment, and bosonic dark matter. In addition, the presence of ^3H as a possible source of the XENON1T low-energy excess [209] could not be ruled out, as the TED showed no evidence of a tritium-like excess. The ROI for the low-energy analysis extends from 0 to 140 keV in the combined energy scale.

Nine components are included in the background model, summarized in Table 4.4. The dominant background source at low energies is the ^{214}Pb beta decay. Its rate is constrained using the decay rates of ^{218}Po and ^{214}Po , yielding an activity between (0.770 ± 0.038) and $(1.610 \pm 0.078) \mu\text{Bq/kg}$. The ^{85}Kr rate is constrained using measurements of the $^{\text{nat}}\text{Kr}$ concentration and the assumed abundance ratio $^{85}\text{Kr}/^{\text{nat}}\text{Kr} = 2 \times 10^{-11}$ [142].

The contribution from detector materials is estimated using the full-chain simulation framework together with results from the screening campaign [175, 220]. The resulting single-scatter gamma-ray background has a flat spectrum below 140 keV, with an expected rate of $(2.1 \pm 0.4) \text{ events}/(\text{t} \cdot \text{yr} \cdot \text{keV})$. The contribution from $2\nu\beta\beta$ decay of ^{136}Xe in the ROI is determined using the half-life reported in Ref. [109], together with the isotopic abundance of ^{136}Xe , using Equation 1.15. It becomes the




Figure 4.9: Comparison between the XENON1T and XENONnT low-energy ER spectra. **Right:** Fit to SR0 data using the background model B_0 . The fit result of B_0 is the red line. The subdominant AC background is not shown. Figures taken from Ref. [122].

Component	Constraint	Fit
^{214}Pb	(570, 1200)	960 ± 120
^{85}Kr	90 ± 60	90 ± 60
Materials	270 ± 50	270 ± 50
^{136}Xe	1560 ± 60	1550 ± 50
Solar neutrinos	300 ± 30	300 ± 30
^{124}Xe	—	250 ± 30
AC	0.70 ± 0.04	0.71 ± 0.03
^{133}Xe	—	150 ± 60
$^{83\text{m}}\text{Kr}$	—	80 ± 16

Table 4.4: Summary of background components with constraints and best-fit values between 1 and 140 keV using the SR0 exposure corresponding to 1.16 tonne-years. Adapted from Ref. [122].

dominant background above 40 keV. The 2ECEC process of ^{124}Xe is modeled as detailed in Ref. [125] and is left unconstrained in the fit. With an expected rate of (0.71 ± 0.04) events/(t · yr · keV), the AC background is subdominant but still included in the fit. The expected solar neutrino spectrum is obtained using the standard neutrino flux in the Large Mixing Angle MSW model and the Standard Model neutrino–electron scattering cross section [221, 222]. A 3% uncertainty is assigned to constrain the solar neutrino rate, based on the scattering rate computed in Ref. [223]. Trace amounts of ^{133}Xe , activated during $^{241}\text{AmBe}$ calibrations, are also included but left unconstrained in the fit due to their negligible rate and low impact in the ROI. $^{83\text{m}}\text{Kr}$ is included in the fit, as residual amounts remain after regular calibrations, and its rate is left unconstrained [122].

A fit to the background model was performed in reconstructed energy using an unbinned maximum likelihood approach. The data agree well with the background model, no

excess is observed, and a beyond the standard model interpretation of the XENON1T excess is excluded. The fitted background model is shown in Figure 4.9, and the best-fit values of the background components are reported in Table 4.4.

The average rate of events between 1 and 30 keV was (15.8 ± 1.3) events/(t · yr · keV), about five times lower than the (76 ± 2) events/(t · yr · keV) measured in XENON1T [209], as shown in Figure 4.9 left. This unprecedented background level is achieved through radon removal (see Chapter 2) and an improved material selection and cleaning. Treating tritium as a signal yields a best-fit rate of zero, with a 90% C.L. upper limit of 15 events/(t · yr). Based on this result, new parameter space is excluded via limits on the interactions of solar axions, bosonic dark matter, and solar neutrinos with an enhanced magnetic moment, as detailed in [122].

4.8 Summary and outlook

This chapter has presented the low-energy analysis framework of XENONnT and its application to the SR0 and SR1 datasets. The achieved sensitivity results from the combination of detector design, material screening, background reduction techniques, and a detailed reconstruction and correction framework. In particular, limitations induced by the low and non-uniform drift field are accounted for through reconstruction and corrections, together with dedicated calibrations of the detector response.

While the low-energy framework described here is shared by the WIMP, CEνNS and ER analyses, other channels require dedicated treatments beyond the standard reconstruction and correction scheme. This is the case for MeV-scale signatures, which motivates dedicated approaches useful for high-energy physics channels as will be discussed in Chapter 5.

Chapter 5

Modeling the XENONnT Response in the MeV Energy Range

Liquid xenon time projection chambers, widely used in the last two decades for low-energy ($\mathcal{O}(\text{keV})$), single-scatter searches [135], are increasingly being used to explore rare processes in the MeV energy range [194, 224–227]. The XENONnT experiment provides sensitivity to various MeV-scale processes, including neutrinoless double beta decay of ^{136}Xe and double weak decays of ^{124}Xe [125], and spectral shape measurements of radioactive decays [228]. This expanded physics program is enabled by the combination of ultra-low background conditions, good energy resolution, and sophisticated reconstruction and simulation frameworks [55, 89, 127, 193]. In particular, the simulation framework has played a central role in XENONnT analyses, providing essential input for detector response modeling and reconstruction in the presence of non-ideal detector operating conditions¹ that limit data-driven approaches. While its performance is well established in the low-energy regime, the extension to MeV-scale signals motivates a dedicated study of its validity and limitations in this energy range.

This chapter describes the modeling and characterization of the XENONnT detector response to ERs in the MeV energy range. A summary of the simulation framework and the emission model adopted in this work is presented in 5.1 and 5.2, while the agreement between data and simulation is discussed in 5.3. The multi-scatter energy reconstruction is described in 5.4, together with a simulation-driven correction for the position dependence of the S1 light collection. In 5.5, we present a processing algorithm developed for the reconstruction of multi-scatter event topologies, which will be used in Chapter 6 to study the XENONnT sensitivity to the $2\nu\text{EC}\beta^+$ decay of ^{124}Xe .

¹See Section 4.1 for details on the operating conditions of the first and second XENONnT science runs.

5.1 The XENONnT simulations pipeline

In XENONnT, the Monte Carlo (MC) framework refers to the software packages used to simulate the passage and detection of particles in the detector, together with the software infrastructure required to produce, manage, and process these simulations. The MC framework is used to model particle interactions in all detector subsystems, including the TPC, the neutron veto [164], and the muon veto [229]. In this chapter, the discussion is focused on the TPC response. An exhaustive description of the MC simulation framework is beyond the scope of this work, a detailed description can be found in Ref. [215, 216].

Large-scale MC productions are executed on high-throughput computing resources provided by the Open Science Grid (OSG) [230, 231], which supplies the computing and storage capacity required for the production of the large simulations datasets used in XENONnT. As illustrated in Figure 5.1, the MC framework is organized as a chain that begins with particle transport and culminates in the generation of digitized waveforms. The discussion in this section focuses on the two core components of this chain: the **Geant4** MC toolkit [232, 233], used for particle propagation and tracking, and the custom XENONnT package **fuse** (framework for unified simulation of events) [1, 55], which models the quanta emission in liquid xenon and detector response up to the waveform level. The resulting simulated waveforms are subsequently processed by the standard XENONnT data processor (**straxen**) [206] using the same algorithms applied to the measured data.

The simulation chain begins by defining the particle generators, which are the primary particles with their corresponding type, position, momentum, and direction. The **Geant4** implementation relies on a detailed description of the geometry and material properties of the XENONnT detector, as well as the physics lists used to describe the interaction processes that particles can undergo [215, 216]. **Geant4** provides step-by-step tracking information for each energy-depositing interaction in the xenon volume, including the underlying physical process and the type of particle that produced it.

In the subsequent step, referred to as *microphysics* simulation as illustrated in Figure 5.1, **fuse** processes the energy depositions by grouping them according to interaction type and origin, using our best knowledge of quanta emission in liquid xenon, as will be described in detail in Section 5.2. The number of light and charge quanta that the interaction yields is calculated based on the recoil energy and the electric drift field at the interaction site. In the next step, the generated quanta are propagated using their positions and emission times together with detector response parameters, which are described below, to model the light and charge transport in the detector. This is




Figure 5.1: Schematic overview of the XENONnT detector response simulation chain. The **Geant4 MC** simulates particle transport and provides detailed tracking information of energy deposits in the liquid xenon. These deposits are provided as input to **fuse**, which converts them into photons and electrons that are subsequently propagated through the detector while accounting for detector specific response. The PMT response is also simulated, producing output in the same format as real data, which is then processed by **straxen** in an identical manner to measured data.

the start of the *detector physics* (see Figure 5.1) simulation of the S1 signal. First, **fuse** incorporates NEST [148] to model the timing profile, which depends on the recoil type and drift field, of the prompt scintillation light with respect to the recoil time. The scintillation light is propagated from the interaction site by means of dedicated optical simulations performed with **Geant4** [156]. For the latter, the LXe optical properties (summarized in Table 2.1) are considered to obtain the per-PMT geometric light collection efficiency, defined as the position-dependent fraction of photons that reach the photocathode of a PMT. The photon propagation times as a function of interaction depth are estimated using these optical simulations, and individual photon arrival times are subsequently sampled from the corresponding depth-dependent time distributions. The photosensor detection efficiencies, including the quantum efficiency and collection efficiency described in 3.1, are incorporated to model the conversion of incident photons into the observed S1 signal in photoelectrons.

The simulation of the S2 signal begins with the propagation of the ionization electrons created at the interaction site, which is assumed to be point-like [141], toward the liquid

gas interface. The propagation of the electron cloud under the influence of the drift field is modeled using the Sorensen diffusion model [234]. As introduced in Section 2.1.3, the number of electrons reaching the interface depends on detector conditions such as the drift and extraction field strength, as well as on the xenon purity. In addition, electric field inhomogeneities can lead to electron losses due to absorption at the PTFE walls. These effects are simulated in **fuse** within the *detector physics* stage of the simulation chain (see Figure 5.1). During propagation, the electron cloud undergoes diffusion in the longitudinal (z) and transverse directions (x, y). The electric field inhomogeneities are taken into account by implementing an electric field map obtained from COMSOL simulations [211, 235] performed at the XENONnT average drift field, $E_d = 22.9 \text{ V/cm}$. The drift velocity of the electron cloud is derived from this map as the average value along the electron cloud path, and is computed using the PyCOMes software [236]. A simulation-driven field distortion correction accounts for the fact that electrons drift through a non-uniform electric field. A NEST drift velocity model is used together with the simulated electric field map, to compute the true drift time at a given depth. The depth is then reconstructed assuming a constant drift velocity, and the correction is defined as the difference between this reconstructed depth and the true interaction depth. For a given interaction vertex in the TPC, the electron cloud broadening in z is modeled using a data-driven longitudinal diffusion coefficient map $D_L(R, z)$ derived from $^{83\text{m}}\text{Kr}$ calibration data. This spatial broadening is converted into a temporal broadening using the averaged drift velocity according to Equation 2.7. The broadening in the transverse direction is modeled using a constant diffusion value across the TPC, $D_T = 4.841 \times 10^{-8} \text{ cm}^2/\text{ns}$, taken from Ref. [237]. Then, the number of electrons that reach the interface is corrected for the probability of electron loss due to attachment to electronegative impurities through Equation 2.11. In XENONnT, during SR0 we achieved an average $\tau = 10 \text{ ns}$ [156] which corresponds to a survival probability of $\sim 80 \%$.

The electrons are extracted into the gas xenon by means of a stronger field. Then, a uniform, data-driven electron extraction efficiency of 55 % is applied. Garfield++ [211, 238] is used to simulate electron trajectories in the gas phase. For every electron extracted, secondary scintillation photons are drawn from a Poisson distribution, considering the gas gain (discussed in 2.1.2) and the double photoelectron probability (discussed in 3.1). The photon arrival times of the S2 signal and its resulting PMT hit pattern are derived in the same fashion as for the S1 signal.

The final step of the **fuse** simulation, corresponding to the *PMTs and electronics* (see Figure 5.1), incorporates the PMT response and associated readout electronics. The photon arrival times entering this stage already include the mean photoelectron transit

time of 46 ns and a transit time spread of 9 ns [174, 197]. Each detected photon is converted into a photoelectron and assigned a channel-specific PMT gain corresponding to the detector operating conditions. The PMT and readout electronics response, including the voltage divider circuit, preamplifier, and digitizer bandwidth, is encoded through a single-photoelectron pulse shape template extracted from calibration data. Individual photoelectron responses are summed to form the analog PMT waveform. PMT afterpulses, introduced in Section 3.1, are simulated by adding delayed single-photoelectron pulses according to PMT and ion specific rates obtained from calibration data. The resulting waveform is then converted into digitized ADC counts, including the effects of the digitizer resolution. Finally, electronic noise measured from real data is added per-channel, yielding simulated S1 and S2 PMT waveforms in the same format as those recorded by the detector.

After simulating the *PMTs and electronics*, the waveforms are processed by **straxen** using the same reconstruction algorithms used for measured data. This includes applying the correction discussed in Sections 4.3. However, several aspects of the detector response are implemented in a simulation-driven, first-principles approach within **fuse**. These include per-PMT light-collection efficiency maps, electron drift and diffusion modeling, extraction efficiency at the liquid-gas interface, the secondary scintillation delay, as well as event-level corrections. For the latter, two such examples are the field-distortion correction and the S1 area position-dependence correction. The latter will be discussed in detail in Section 5.4.1. Taken together, these simulation-driven implementations allow **fuse** to provide realistic simulations of the detector response.

The performance of the **fuse** framework was originally optimized and extensively validated at the keV energy scale [55, 89, 188, 208]. When extending the simulation chain to the MeV energy scale, several assumptions that are valid in the keV regime cease to hold. This is primarily due to the significantly larger spatial extent and higher number of S2 peaks per event, or multiplicity, of energy depositions produced by MeV-scale interactions. In particular, the clustering algorithm based on spatial proximity and the photon and electron yields provided by NEST do not provide an accurate description of the detector response in this regime. To achieve agreement between data and simulation at high energies, three major upgrades were implemented in the framework: (i) a physics-motivated algorithm for grouping, or clustering, energy depositions, (ii) a data-driven yield model derived from ^{214}Bi and ^{212}Pb beta decays, and (iii) a simulation-driven, first-principles correction of the S1 area position dependence. These developments, together with their validation, are described in the following sections.




Figure 5.2: Illustration of the lineage clustering algorithm. Particles shown in black do not form clusters, only their interaction products. **Left:** A gamma-ray propagates and undergoes multiple interactions, each one of which initiates a separate lineage, whose products are clustered together. **Right:** An electron produced in a beta decay undergoes multiple interactions along its track, all of which are assigned to the same lineage. The dotted contours and the color indicate the resulting clusters to which the data-driven yields are applied.

5.2 Refined emission model for the MeV range

5.2.1 Lineage clustering

In XENONnT, the standard approach to cluster energy deposits at low energies relies on spatial and temporal proximity. The clustering ensures that a physically motivated interaction site can be defined for the emission model, which assumes localized energy depositions. At low-energies, the clustering is typically implemented using the DBSCAN² algorithm from the scikit-learn package [239]. The type of recoil of individual energy deposits is given by the Geant4 process information associated with each interaction step. Energy deposits occurring within a predefined time and distance thresholds are grouped together, with their energies summed to form a single effective interaction. The corresponding position and time are assigned as the energy-weighted average of the clustered deposits. Then, the interaction type assigned to each cluster is determined by the type of recoil that contributes the largest fraction of the total deposited energy among its constituent energy deposits. This approach has proven to describe the micro-

²Density-Based Spatial Clustering of Applications with Noise.

physics at the keV energy scale. However, it becomes unsatisfactory at MeV energies, where this clustering scheme can lead to nonphysical behavior. In particular, energy deposits originating from distinct particles, such as multiple electrons or gamma-rays in close spatial and temporal proximity, may be incorrectly merged into a single cluster. Conversely, extended electron tracks can lead to the artificial splitting of energy depositions into multiple clusters when the separation between successive deposits is smaller than the characteristic track length, which corresponds to the spatial extent of energy deposition by the electron in liquid xenon. These effects directly impact the assignment of scintillation and ionization yields (shown in Figure 2.3) by hiding the true topological structure of the energy deposition.

To address these limitations, a physics-driven approach was implemented in **fuse** [228]. In this approach, clusters are formed based on the origin and nature of the interactions, i.e., *lineage*, rather than just on spatial and temporal proximity. An illustrative example of the lineage clustering algorithm is shown in Figure 5.2. For electron interactions, all associated energy depositions (whether from ionization or scattering processes) are grouped within a single cluster. Cases involving Bremsstrahlung emission form separate clusters. For gamma-rays, each distinct interaction, such as pair production, Compton scattering or photo-absorption, initiates a separate cluster, with the option to merge clusters when successive interactions occur within a predefined spatial threshold set by default to zero. The distinction in the treatment of electron and gamma reflects the fundamentally different topologies of electron and gamma induced energy loss in liquid xenon, where electrons typically deposit energy continuously along long tracks, while gamma-rays interact through discrete, spatially localized processes.

This clustering defines the effective interaction site at which scintillation and ionization are produced, bridging the gap between the microscopic energy deposits simulated in Geant4 and the emission models used to generate quanta, which assume localized interactions. By preserving the underlying topology of the energy deposition, lineage clustering provides a basis for assigning electron and photon yields at the cluster level. This is particularly important at high energies, where extended tracks and multiple secondary particles would otherwise bias the effective quanta production. Building on this interaction definition, the next section introduces a data-driven model to describe the electronic recoil yields applied to these clusters.

5.2.2 Data-driven electronic recoil yields

The development of a data-driven yield model is motivated by the limitations of the NEST framework, which is commonly used to predict the scintillation and ionization yields per unit of deposited energy. NEST provides electronic recoil yields based on

a data-driven approach using input from several detectors under different electric field conditions [144, 148]. To model the electron yield, an empirical ten-parameter function is used, consisting of two smooth energy transitions that interpolate between low-energy and high-energy regimes [240]:

$$Q_y(E, \varepsilon, \rho) = m_1(\varepsilon, \rho) + \frac{m_2 - m_1}{\left[1 + \left(\frac{E}{m_3(\varepsilon)}\right)^{m_4(\varepsilon)}\right]^{m_9}} + m_5(\varepsilon) - \frac{m_5}{\left[1 + \left(\frac{E}{m_7(\varepsilon)}\right)^{m_8(\varepsilon)}\right]^{m_{10}}}, \quad (5.1)$$

where E is the energy deposited in the interaction cluster, ρ the liquid xenon density and ε is the electric field. The parameters m_1 - m_{10} control the overall normalization, asymptotic yields, transition energies, slopes, asymmetries, and electric-field dependence of the model, and are tuned to reproduce electron yield measurements over a wide range of energies and fields. At high energies, however, the interpolation performed by NEST across different detector configurations and electric field strengths does not provide satisfactory agreement with data.

NEST is fundamentally data-driven, the resulting parametrizations represent effective yields that include detector-specific effects such as energy resolution, signal reconstruction thresholds, and clustering performance. Consequently, the data used to constrain the model can differ not only in electric field and geometry but also in their ability to resolve spatially separated energy depositions. In detectors with limited resolution, multiple nearby interactions may be reconstructed as a single scatter, whereas higher resolution detectors can resolve them as multi-scatter. Because recombination and thus quanta production in liquid xenon are non-linear with deposited energy (see Figure 2.3), this topology dependence propagates into the inferred yields. The NEST parametrization therefore encodes, to some extent, these detector effects, which motivates a detector-specific tuning.

In this work, the data-driven yields are determined by fitting Equation 5.1 to the β -decays of ^{212}Pb and ^{214}Bi [228], from the ^{220}Rn and ^{222}Rn decay chains respectively, covering an energy range from 10 keV up to 3 MeV. The photon yield is calculated as the complementary fraction of quanta, using Equations 2.4 and 2.6. The yields are applied at the cluster level, replacing the standard NEST interpolation in the high-energy regime while maintaining the same low-energy ER and NR emission parametrization [208].

The combined use of lineage-based clustering and our data-driven β -yield model directly affects the number of photon and electron quanta produced per interaction. At the

event level, these changes manifest themselves in the reconstructed signal areas and, consequently, in the reconstructed energy. The impact of these improvements on the simulation framework is discussed in the following section.

5.3 Detector response validation in the observable space

At the MeV energy scale, interactions in the detector are predominantly multi-scatter, leading to complex event topologies that cannot be fully characterized by reconstructed energy alone. In this regime, the detector response must therefore be validated directly in the observable space, by comparing features that encode the spatial and temporal structure of the signals. To this end, comparisons between data and simulation are performed using uncorrected S1 and S2 signal shape observables, including their area, width, rise time, fraction of light detected by the top PMT array, and reconstructed position. These quantities probe the modeling of scintillation and ionization processes, as well as their propagation and detection.

The agreement between data and simulation is studied using ^{232}Th simulation and calibration data. For the latter, the source is deployed in the U-tubes (see Figure 2.10) at multiple positions relative to the TPC, including locations near the top and bottom of the detector (see Figure 5.6). For the comparison, the full energy range of the ^{232}Th decay chain is considered. The data quality selections described in Section 5.3.1 are applied to both data and simulations.

The calibration data were acquired on the 24th of March 2022 and consist of 30 minutes of data taking at each of the two source positions. This calibration was performed outside the official SR0 science run (see Section 4.1), during a period in which detector conditions differed from those of the science data. As a consequence, the data-driven corrections derived for SR0 are not optimal for this calibration dataset and may bias the reconstructed detector response. To account for this, a dedicated set of corrections was derived specifically for the ^{232}Th calibration period. In particular, the lone-hit and PMT gain corrections were obtained using the methods described in 3.3.3 and 3.5.2, while the electron lifetime correction was derived according to the procedure detailed in ???. The PMT gain correction was especially important, as the LED timing effects discussed in Section 3.5.1 were at their maximum during March 2022.




Figure 5.3: S1 width selections as a function of S1 area applied to ^{232}Th calibration data. The lower-bound cut on the S1 width at 90 % area (**left**) rejects narrow S1 signals, while an upper-bound cut on the S1 width at 50 % area (**right**) removes events with abnormally wide S1 signals. The solid pink lines indicate the empirically determined cut boundaries, and the shaded pink regions denote the excluded regions.




Figure 5.4: S1 area fraction top as a function of depth for ^{232}Th calibration data, shown as a color map. The selection boundaries are indicated by pink solid lines, with the rejected region is shaded.

5.3.1 Data quality selection

A set of selection criteria are applied to both simulated and measured events in order to ensure a reliable comparison between the two datasets. These selections aim to remove nonphysical or poorly reconstructed events [156], and will be described below.

S1 width: Two complementary selections based on the temporal spread of the S1 signal are applied. Events with wide S1 signals are identified by comparing the width containing the central 50 % of the S1 area to the total S1 size. An upper bound, defined as a function of S1 area using an empirical model, with separate parameterizations for small and large areas smoothly interpolated, rejects events that are too wide for their size (Figure 5.3 right). This removes pile-up, distorted waveforms, and gas-phase inter-




Figure 5.5: Schematic representation of an S2 signal showing the central 50 % region of the integrated S2 area. The rise time, also shown, is defined as the interval between the 10 % and 50 % quantiles of the cumulative area.

actions. Events with narrow S1 signals are rejected by comparing the width containing the central 90 % of the S1 area to the S1 size. A lower bound, given by a polynomial fit to calibration data (Figure 5.3 left), efficiently rejects α particles, which produce characteristically narrow pulses.

S1 area fraction top: The fraction of S1 light detected by the top PMTs depends on interaction depth. Events near the liquid–gas interface yield larger fractions, while deeper events yield smaller ones. Bulk events therefore form a narrow, approximately linear band in S1 area fraction top vs. reconstructed depth (z). Events outside this band are rejected using empirically determined linear upper and lower bounds, shown in Figure 5.4. This cut removes below-cathode events, which have low and nearly constant top fractions due to optical shielding and lack of S2, and gas-phase interactions, which show high top fractions and reconstructed positions near the interface ($z > 3.2$ cm).

S1 timing: The timing of the S1 pulse is constrained by requiring the time difference between the S1 start time and its weighted center time to be less than 200 ns. This requirement removes events with abnormally long S1 pulses, which are produced by pile up of events or PMT afterpulses.

Single S1: Events are selected in the parameter space defined by the area of the second-largest S1 pulse as a function of the main S1 area. Events with a second-largest S1 area greater than 200 PE are rejected. Such events may originate from S2 signals that are incorrectly classified as S1, or from multiple uncorrelated interactions within the same event window.

S2 width: A selection on the width of the S2 signal is applied in the parameter space defined by the S2 width at 50 % (illustrated in Figure 5.5) area and the drift time. Events are required to have an S2 width larger than 70 % of the expected value from the diffusion model at the corresponding drift time. This lower-bound selection removes events with nonphysical drift times, such as gas events, accidental coincidences, and merged S2 signals.




Figure 5.6: Position distribution of ^{232}Th calibration events from both top and bottom positions used for validation of the high-energy simulation framework. The simulated distribution (upper panel) is compared to calibration data (lower panel), shown after applying the S1 and S2 data quality selections described in this section. The FV boundaries are marked by pink dashed lines, the 95 % contour region is indicated in blue and the TPC limits by black solid lines.

S2 multiple scatter: The S2 single-scatter selection described in Section 4.4 is applied to both simulations and data in order to distinguish between single- and multi-scatter interactions. While the physics searches presented in Chapter 4 are restricted to single-scatter events, the high-energy regime considered here is dominated by multi-scatter interactions. Consequently, the simulation–data comparison in this section is performed using the multi-scatter event population.

Naive FV: An internal naive FV is defined with the following boundaries:

$$\begin{aligned} -133 \text{ cm} < z < -20 \text{ cm} \\ R < 63 \text{ cm}, \end{aligned} \tag{5.2}$$

where R is the radial coordinate of an event and z is the depth in the TPC. Figure 5.6 shows the spatial distribution of the selected events after applying all data quality criteria, for both data (lower panel) and simulations (upper panel). The TPC boundaries are indicated in black, the naive FV selection is indicated by pink dashed lines, and the




Figure 5.7: Comparison of ^{232}Th calibration data and simulations in the observable space. Calibration data are shown in dark blue with statistical uncertainties, while simulations are represented by their mean values (solid lines) and the associated 1σ spread (shaded bands). S2 signal area as a function of S1 signal area for simulations obtained with the refined emission model and data-driven yields (new) and for simulations without these (old).

95 % region is indicated by the blue contour.

The selection in Equation 5.2 is targeted to exclude events occurring close to the TPC boundaries. In these regions, multi-scatter interactions may be only partially contained within the active volume, leading to incomplete reconstruction, or may be affected by electric field non-uniformities near the electrodes, which can distort the reconstructed signals. Additionally, this selection also reduces contributions from radioactive decays in surrounding detector materials. A dedicated FV optimization targeting the $2\nu\text{EC}\beta^+$ decay channel is derived in Chapter 6.

5.3.2 S1 and S2 shape properties

At MeV energies, multiple S2 signals can be individually resolved provided sufficient spatial separation between interaction sites. In contrast, due to the fast propagation of scintillation light, S1 signals originating from multiple energy depositions are typically merged and reconstructed as a single S1. As a result, S2 observables play a central role in resolving multi-scatter topologies and therefore constitute a primary benchmark for validating the detector response model at high-energies. In this subsection, we compare data and simulation, with particular emphasis on S2 properties. The detector model is considered validated when simulations reproduce the overall shape and position dependence of these observables.

Figure 5.7 shows the S1 and S2 area parameter space for energies up to 2.6 MeV,




Figure 5.8: S1 shape properties comparison between data and simulations using ^{232}Th events. Distributions are shown separately for events reconstructed in the top and bottom TPC, corresponding to two different source deployment positions. The S1 rise time (**left**) and the central 50 % of the S1 width (**middle**) exhibit a discrepancy between data and simulation, attributed to waveform downsampling applied during peak reconstruction, as discussed in the text. The fraction of the S1 signal detected by the top PMT array (**right**) exhibits good agreement with the simulations.

comparing calibration data with simulation results. The gray curve corresponds to the previous simulation configuration, which does not include the refined emission described in Section 5.2. The blue curve represents the updated simulation, which exhibits a remarkable level of agreement with data across the full parameter space. This agreement demonstrates that the refined clustering algorithm and the data-driven yield model provide an accurate description of the conversion of energy deposits into observable quanta. As a result, the global response of the detector in terms of scintillation and ionization signals is consistently reproduced, establishing a solid foundation for the subsequent validation of peak-level observables and event topology.

The S1 shape parameters are shown in Figure 5.8. A discrepancy is observed in the rise time of simulated S1 signals when compared to data. The data distribution is shifted towards larger values relative to simulation, with a more pronounced effect in the top position events. This effect is attributed to the waveform downsampling applied during *peaklet* reconstruction. In **straxen**, S1 width and rise time are derived from the *peaklet* waveforms. To fit these waveforms into fixed-length arrays, downsampling is sometimes applied, resulting in S1 signals with different effective time resolutions. Since the S1 signals are intrinsically narrow, downsampling can significantly distort their temporal profile, leading to biased width and rise time estimates. The effect becomes more pronounced at high energies, where S1 signals extend over more samples and are therefore more likely to trigger downsampling during waveform reconstruction. Although this effect requires further investigation, it has a negligible impact on the multi-scatter search presented in Chapter 6, as this analysis relies on S2 observables.




Figure 5.9: S2 shape properties comparison between data and simulations using ^{232}Th events. The two distributions per panel correspond to different radioactive source deployment positions. **Left:** S2 rise time distribution. **Middle:** Central 50% width of the S2. **Right:** Fraction of the S2 signal measured by the top PMT array, which shows a small depth-dependent discrepancy discussed in the text.

Figure 5.9 shows the S2 rise time, S2 width at 50%, and the fraction of the S2 signal detected by the top PMT array. The light blue and dark blue histograms represent the top and bottom deployments of the ^{232}Th source, respectively. In general, the S2 time profile observed in simulations is in good agreement with the data, as seen in the rise time (left) and width (middle) distributions. Small offsets are observed in the S2 area fraction top shown in Figure 5.9 right, however the characteristic scale and shape of the distribution is also well reproduced for both source positions. Overall, these comparisons demonstrate that the simulation reproduces the main morphological features of the S2 signals, supporting the validity of the detector response model for the reconstruction of multi-scatter event topologies.

These observables are further studied as a function of z , as shown in Figure 5.10. The S2 rise time and the S2 width at 50%, shown in the left and middle panel respectively, are reproduced with excellent agreement over the entire depth of the TPC. The fraction of S2 light detected by the top PMT array, shown in Figure 5.10 left, exhibits small offsets at the bottom and top of the TPC, with this offset remaining within the 1σ band uncertainty. The observed 1–2% variation in the uncorrected S2 AFT arises from residual (x, y) -dependent optical response effects, as no S2 (x, y) correction is applied. This observable provides a validation of the optical response and light collection efficiency as a function of interaction depth. Since these observables are governed by the electron drift velocity, longitudinal diffusion in the liquid and the electroluminescence process in the gas phase, this agreement indicates that these effects are accurately modeled in the simulation.

The S2 multiplicity as a function of z is shown in Figure 5.11 for data and simulation. For this comparison, an area threshold of 200 PE was applied to data in order to




Figure 5.10: Comparison of simulation and data for S2 signal shape parameters as a function of z including the fraction of S2 light detected by the top PMT array (**left**), S2 rise time (**middle**), and the central 50% width (**right**). Solid lines show the dataset means, and shaded bands denote the 1σ intervals. The simulation and data agree within 1σ across all parameters over the full z range.




Figure 5.11: S2 multiplicity distribution for ^{208}Tl calibration data and simulation, shown for the top and bottom deployment positions of the external source. The shaded corresponds to the statistical uncertainty.

suppress small S2 signals. This threshold is chosen to be consistent with the characteristic area of few-electron background signals, which are not included in the simulation. These small S2s are associated with the delayed release of electrons following a high-energy interaction and are commonly referred to as electron trains. They originate from processes such as photoionization or the delayed emission of electrons trapped on impurities or detector surfaces [153, 213, 214]. After applying this selection, the comparison shows that data and simulation agree within the 1σ band.

A residual offset is nevertheless observed in both source deployment positions, with data exhibiting a slightly higher multiplicity than the simulation. This trend may be related to a small mismatch between the implemented and the actual source deployment geometry. In the simulation, the top and bottom source positions are placed at the same distance from the cryostat. The data, however, suggest a slight asymmetry (see Figure 5.6). Such a difference alters the spatial distribution of interaction sites and

changes the fraction of events occurring near detector boundaries. Events close to the boundaries are more susceptible to reconstruction effects that increase the number of identified S2 peaks. As a result, a different boundary event fraction between source deployments can produce an additional shift in the observed multiplicity.

Although this effect requires further study, the remaining mismatch does not limit the sensitivity of the high-energy searches presented in this work. This is due to the small area of the additional S2s, which makes them subdominant for background discrimination that primarily relies on the three largest S2 peaks, as discussed in Chapter 6. Additional observables also provide robust handles on the topology of multi-scatter events, as described in Section 5.5 and Section 6.5.2.

5.4 Energy reconstruction for multi-scatter events

The total deposited energy of a multi-scatter event is reconstructed by adding the contributions from all S1 and S2 peaks within the event window. In contrast to single-scatter reconstruction, no attempt is made to associate the S1 signal with individual interaction sites. Instead, an effective description of the total deposited energy is constructed following Equation 2.10, by combining the contributions from all valid S2 peaks. To this end, we define $\text{cS1}_{\text{multi}}$ and cS2_{sum} :

$$\text{cS1}_{\text{multi}} = \text{S1} \frac{\sum_i \text{LY}_{\text{rel}}(\vec{x}_i) \text{cS2}_i}{\sum_i \text{cS2}_i}, \quad \text{cS2}_{\text{sum}} = \sum_i \text{cS2}_i, \quad (5.3)$$

where the index i runs over all valid S2 peaks in the event and $\text{LY}_{\text{rel}}(\vec{x}_i)$ denotes the relative light yield correction at the S2 peak reconstructed position $\vec{x}_i$. The $\text{cS1}_{\text{multi}}$ represents an effective corrected S1 signal for multi-scatter events. Since typically only a single S1 pulse is observed, the position-dependent relative light yield correction is evaluated using a cS2-weighted average over all contributing S2 peaks. The cS2_{sum} is defined as the sum of the corrected S2 areas of all S2 peaks whose drift times lie between zero and the maximum allowed drift time of the TPC.

For each event, the reconstructed position is obtained as weighted averages over all its valid S2 peaks. The corrected S2 area, cS2, is used as a proxy for the energy deposited at each interaction site, such that interaction sites with larger ionization signals contribute more to both the reconstructed position and the effective light collection correction entering $\text{cS1}_{\text{multi}}$. This approach is particularly useful for applying position-dependent corrections and FV selections at the event level.

The validity of this effective energy reconstruction relies on the modeling of position-




Figure 5.12: Spatial distribution of the energy bias in SR0 simulations for a 511 keV gamma-ray, with a maximum bias of $\sim 3\%$. The energy is reconstructed using the SR0 data-driven S1 LY_{rel} map as input at the event-level stage, which leads to a clear position dependence in both the (x, y) and (R, z) planes.

dependent detector response effects. Since the definition of $cS1_{\text{multi}}$ accounts for the LY_{rel} evaluated at the reconstructed interaction positions, any position-dependent mis-modeling of the S1 response can propagate directly into the reconstructed energy, producing systematic biases. When applying the data-driven S1 LY_{rel} map to simulated events, the reconstructed energy systematically deviates from the nominal value in a position-dependent manner across the TPC. To quantify the effect, we evaluated the reconstructed energy as a function of position by binning events in the (x, y) plane and in cylindrical coordinates (R, z) . In each spatial bin, we compute the mean reconstructed energy using all events whose reconstructed positions fall within that bin. We then define the relative energy bias as the fractional deviation of this mean value from the nominal gamma-ray energy. The spatial distribution of the bias is shown in Figure 5.12 for a gamma-ray of 511 keV. This effect appears only in simulations across energies from a few keV to the MeV scale and predominantly impacts the corrected S1, motivating the derivation of a simulation-driven correction map that will be discussed in the next section.

5.4.1 S1 relative light yield correction

The S1 LY_{rel} map is an essential input to the energy reconstruction described in the previous section. The map is defined as position-dependent correction factors whose values are applied to the reconstructed S1 signal to compensate for spatial variations across the TPC volume. The goal of the S1 LY_{rel} map is to balance the light yield across the FV. Since the light yield depends not only on the optical response of the




Figure 5.13: Left: Binning scheme in cylindrical coordinates for a representative z slice in the bulk of the TPC. Each slice contains 192 bins, with approximately 2000 events per bin, corresponding to a statistical uncertainty of 2.2 %. **Right:** Distribution of the number of events per TPC bin. The red and blue histograms correspond to the two upper and two lower z -slices, respectively, which contain significantly fewer events than the central z -slices shown in gray.

detector but also on the scintillation photon yield (see Equation 4.1), which varies with the electric field strength, the field dependence of the S1 signal must be removed prior to deriving the correction map. To this end, a liquid xenon light-yield model [212] is used to predict the photon yield at a given electric field value, $P_y(\varepsilon(x, y, z))$. The field-independent S1 signal, $\widetilde{S1}$, is obtained by deconvolving the electric-field dependence of the light yield, using the ratio of the predicted photon yield at the local field to that at the average field in the FV used for the data-driven approach:

$$\widetilde{S1} = S1(x, y, z) \frac{P_y(\langle \varepsilon \rangle_{FV})}{P_y(\varepsilon(x, y, z))}. \quad (5.4)$$

Once the field dependence has been removed, the remaining variations in the light yield reflect purely geometric effects. Deriving a correction for these effects requires a calibration source that provides a uniform and monoenergetic energy deposition throughout the detector volume with high statistics. ^{83m}Kr calibration events satisfy these requirements and are therefore used for the derivation of the data-driven map [156]. To follow the same strategy in the simulation-driven approach, monoenergetic 41.5 keV energy depositions were simulated to mimic the ^{83m}Kr calibration events. In total, 2.0×10^7 primaries were simulated, yielding approximately 1.8×10^7 reconstructed events used for the map construction.

The map is derived by dividing the TPC into spatial bins defined in cylindrical coordinates (R, θ, z) . The binning is the same as in the data-driven approach and it is chosen




Figure 5.14: Example of the S1 area distribution in a spatial bin of the z -slice shown in Figure 5.13. The simulated events are shown in gray together with their statistical uncertainties. The red solid line represents the Gaussian fit to the distribution, and the corresponding fit residuals are shown in the bottom panel.

to balance the spatial resolution with statistical uncertainty. It consists of 48 slices in z and 8 radial rings per slice, a representative z -slice is shown in Figure 5.13 left. To maintain uniform bin volumes across the detector volume, the number of bins increases with radius. Specifically, the n th radial ring is divided into $2n - 1$ bins in θ , keeping the radial bin width constant while accounting for the increasing circumference at larger radii. To avoid elongated bins in the outer rings, the number of azimuthal bins is scaled by a factor of three, resulting in bins with comparable size in the R and θ directions. With this binning, each bin contains approximately 2000 events, corresponding to a statistical uncertainty of about 2.2%. The topmost and bottommost z -slices contain fewer events per bin due to reduced statistics in this region of the TPC, as shown in Figure 5.13 right.

A Gaussian function is fitted to the $\widetilde{S1}$ area distribution per bin to extract the mean value, $\langle \widetilde{S1} \rangle_i$, as shown in Figure 5.14. In the data-driven approach, the bin-wise mean values are normalized to the average over the FV. In the simulation-driven approach, this average is instead computed as a weighted mean of the bin-wise $\widetilde{S1}$ values, where the weights correspond to the number of events in each bin:

$$\langle \widetilde{S1} \rangle_{\text{FV}}^w = \frac{\sum_i \langle \widetilde{S1} \rangle_i \cdot N_i}{\sum_i N_i}, \quad (5.5)$$




Figure 5.15: Difference between the simulation-driven and data-driven light yield correction maps for SR0. The distribution is obtained by summing the bin-wise differences across all z -slices and projecting them onto the (x, y) plane, such that each bin represents the total contribution at a given (x, y) location integrated along the z -axis. The spatial dependence of the difference is consistent with the position-dependent energy bias observed in Figure 5.12, indicating that the simulation-driven correction addresses the spatial component of the bias.

where N_i denotes the number of events in the i -th bin. The resulting normalization factor is $\langle \widetilde{S1} \rangle_{FV}^w = 363.16$ PE. The correction is implemented in the reconstruction as a relative light-yield factor, LY_{rel} :

$$LY_{\text{rel}} = \frac{\langle \widetilde{S1} \rangle_i}{\langle \widetilde{S1} \rangle_{FV}^w}. \quad (5.6)$$

For each event, the corrected S1 signal is obtained as, $cS1 = \frac{S1}{LY_{\text{rel}}}$. The resulting correction factors are stored as a spatial map and interpolated to the reconstructed position of each event. To compare the simulation-driven and data-driven maps, the absolute difference between the two is computed and integrated over all z slices at each (x, y) position, yielding the distribution shown in Figure 5.15, which follows the position-dependent energy bias observed in Figure 5.12. This demonstrates that the simulation-driven correction captures the spatial component of the observed energy bias. After applying the simulation-driven LY_{rel} correction, the position-dependent component of the energy bias is strongly reduced. While in the bulk of the TPC the bias is approximately uniform around $\sim 1\%$ at 2.6 MeV, a residual radial dependence remains at large radii, reaching up to -3% as shown in Figure 5.16. This behavior is consistent with partial energy loss from multi-scatter interactions occurring near the detector boundary, where one or more scatters can escape the active volume. In addition, the average normalization of the data-driven and simulation-driven maps within the FV




Figure 5.16: Energy bias as a function of R^2 for a 2.6 MeV gamma-ray simulation, after applying the simulation-driven S1 correction map, using both multi and single scatter events. Each curve corresponds to a 6 cm slice in z . The vertical gray line marks the boundary of the outermost radial bin used in the LY_{rel} map derivation. In the bulk of the TPC, the bias is approximately uniform at the 1 % level across all z values, while at large radii the bias is increasingly negative, as discussed in the text.

differs (1.005 and 1.012, respectively), implying a net shift of the corrected S1 of $\sim 0.7\%$ when the simulation-driven map is used. In the following section, the relative energy resolution and residual energy bias are quantified as a function of energy using simulated monoenergetic gamma-ray lines.

5.4.2 Multi-scatter energy resolution and bias

The energy resolution quantifies the precision with which a detector can reconstruct the energy of an interaction. For a given monoenergetic energy deposition, the reconstructed energy is not a single value but a distribution with spread, that reflects intrinsic fluctuations in signal production and detector response. The energy resolution is commonly defined as the ratio of the standard deviation, σ , to the mean reconstructed energy E_r of a given peak. The mean reconstructed energy is obtained by fitting an exponentially modified Gaussian function to the reconstructed energy distributions of monoenergetic simulated sources in the keV to MeV range. It consists of a Gaussian core convolved with a one-sided exponential component toward higher energies. This model is chosen to account for a high-energy tail observed in most peaks. The tail is caused by afterpulses or electron trains produced after an energetic interaction that add charge to the true S2 signal, which broaden the distribution on the high-energy side. In the limit where the peaks exhibit no high-energy tail, the exponential component vanishes and the model reduces to a pure Gaussian. The fit model includes a linear background com-




Figure 5.17: Example fit (red) to the reconstructed energy, E_r , spectrum of the ^{208}Tl line expected at 2614.5 keV. The linear background and the Exponentially Modified Gaussian (EMG) components of the fit model are shown in orange and green, respectively. The bottom panel shows the residuals. The 1σ , 2σ , and 3σ bands are shown for reference.

ponent within an energy window centered on the expected peak position. The events are binned into a spectrum, and the exponentially modified Gaussian model is fitted to the binned counts in the selected energy window using a χ^2 minimization.

Eight sources were simulated: $^{83\text{m}}\text{Kr}$ (41.5 keV), $^{131\text{m}}\text{Xe}$ (163.9 keV), $^{129\text{m}}\text{Xe}$ (236.2 keV), the e^+e^- annihilation peak (511 keV), ^{214}Bi (609.3 keV, 1120 keV, 2204 keV), ^{60}Co (1173.2 keV and 1332.5 keV), ^{40}K (1461 keV) and ^{208}Tl (2614.5 keV). To characterize the energy resolution for multi-scatter topologies, only multi-scatter events were included in the analysis. A representative fit to the ^{208}Tl peak at 2614.5 keV is shown in Figure 5.17. From the fit parameters, the energy resolution is extracted as the ratio σ/E_r and fitted as a function of the reconstructed energy with an empirical function given by [241],

$$\frac{\sigma(E_r)}{E_r} = \frac{a}{\sqrt{E_r}} + b, \quad (5.7)$$

with a and b as fit parameters and E_r as the reconstructed energy. The dominant contribution to the resolution arises from the stochastic nature of quanta production in LXe, whose variance scales as $1/\sqrt{N_q}$ and therefore as $1/\sqrt{E}$, following Equation 2.5. At high energies, where N_q is large, these statistical fluctuations are expected to be




Figure 5.18: Energy resolution as a function of reconstructed energy obtained from monoenergetic sources simulations. The blue dashed line shows a fit to the empirical model in Equation 5.7, whose extrapolated values are $a = (0.347 \pm 0.005) \text{ keV}^{1/2}$, $b = (0.0025 \pm 0.0002)$. For comparison, the energy resolution measured from data using single-scatter events is also shown. The inset provides a zoom of the MeV range.

subdominant. In addition to the Poisson term, recombination fluctuations contribute to the variance as the fraction of electron–ion pairs that recombine is not fixed but fluctuates around a mean value. These fluctuations affect the relative partitioning between observable photons and electrons. In the combined energy scale, this effect largely cancels due to the intrinsic anti-correlation of S1 and S2, but residual fluctuations remain, particularly if light and charge collection are not perfectly corrected. In practice, recombination fluctuations are more relevant at low energies, while at MeV scales the statistical term dominates. Both contributions are covered by the parameter a in the resolution function. Even though custom corrections are applied to mitigate position-dependent variations, these are imperfect and result in residual fluctuations in the reconstructed S1 and S2. Electronic noise in the PMTs as well as statistical fluctuations from dynode multiplication provide further contributions [242], although these are generally subdominant at the MeV scale. All of these non-statistical contributions are absorbed in the constant term b of Equation 5.7. The resulting energy resolution curve is shown in Figure 5.18. The simulation-driven energy resolution curve is systematically lower than the data-driven reference. At 1 MeV it is lower by 14 %, increasing to 18 % at 2.6 MeV. This behavior is expected, as the simulation represents the intrinsic detector response under ideal conditions, while measured data include additional systematic effects that introduce a modest broadening of the energy distribution.

While the resolution characterizes the spread of reconstructed energies around the mean,




Figure 5.19: Relative energy bias of multi-scatter events as a function of energy derived from simulations of monoenergetic sources. The blue dashed line shows a fit to the empirical model given in Equation 5.8, with best fit parameters $A = (1.14 \pm 0.04)$, $k = (0.0090 \pm 0.0007) \text{ keV}^{-1}$, $c = (-0.9190 \pm 0.0008)$. The data-driven, single-scatter bias curve is shown in the black dotted line together with its 1σ uncertainty band [156].

the energy bias accounts for systematic shifts of the peak position from the true energy, reflecting detector effects not corrected by standard corrections. The bias of the reconstructed energy E_r with respect to its nominal value E_t is quantified as $(\mu(E_r) - E_t)/E_t$, where $\mu(E_r)$ denotes the mean peak position obtained from the fit. The relative energy bias measured using the monoenergetic peaks considered in this study is shown in Figure 5.19, together with an empirical model fitted to the data. The energy dependence of the bias is parametrized by [156]:

$$f(E_t) = A \arctan(k \cdot E_t) + c, \quad (5.8)$$

where A , k and c are free parameters. This functional form provides a smooth, monotonic description of the bias, which is linear at low energies and gradually saturates at higher energies. The similar energy dependence observed in data and simulations in Fig. 5.19 points to a common origin related to nonlinearities in the S1 and S2 signals, as discussed in the context of the energy resolution. The larger magnitude seen in data is attributed to additional instrumental and reconstruction effects, such as photoionization, afterpulses, and peak-finding imperfections, that are only partially modeled in the detector simulation.




Figure 5.20: Distributions of the *multipeak* S2 parameters for the five largest S2 peaks per event. ^{232}Th calibration data are compared with simulation using the ^{208}Tl 2.6 MeV gamma line. Each panel shows the distributions for source deployments at the top (blue) and bottom (pink) of the TPC. **Top:** S2 area. **Middle:** S2 R position. **Bottom:** S2 z position.

5.5 Multiple S2-peak event reconstruction

The standard `straxen` reconstruction chain identifies multiple S2 signals per event during peak finding. However, only the largest and second-largest S2 peaks and their properties are stored for analysis. As a result, the detailed peak structure of multi-scatter events is only partially preserved, limiting the ability to fully exploit their topology.

The validation of signal shape observables presented in Section 5.3.2, together with the improved energy reconstruction described in Section 5.4, demonstrates that the simulation and reconstruction chain provide an accurate description of the uncorrected observables. Building on this foundation and to address the limitations of the standard reconstruction chain, a dedicated peak-level reconstruction algorithm was developed in this work. This algorithm, referred to as *multipeak*, explicitly stores detailed information for multiple S2 peaks within each event. In the latter, at the peak level of the reconstruction chain and within each event, peaks are ranked by area and detailed information is stored for up to two S1 peaks and up to five S2 peaks. For every peak, the algorithm records the peak time relative to the event start, the waveform sampling resolution, the reconstructed peak area, and the fraction of the signals detected by the top PMT array. For S2 peaks, the PMT hit pattern is additionally stored through the per-channel area information. The S2 corrections derived at the event-level are applied

to each S2 peak. S2 positions are initially taken from the standard position reconstruction output and are subsequently corrected. First, a longitudinal correction is applied using a z -bias map parameterized in observed radius and depth. This correction accounts for the non-uniformity of the electric field along the drift direction, and yields a corrected depth coordinate for each peak.

Following the z -bias correction, field distortion corrections are applied using a three dimensional map evaluated at the intermediate corrected position. The map provides an effective radial displacement that accounts for distortions of the drift field, which bias reconstructed positions primarily in the radial direction, particularly near the TPC boundaries. This displacement is used to compute the corrected transverse coordinates by rescaling the original (x, y) vector while preserving its azimuthal direction. The corrected three-dimensional position and the corrected radius are stored for each S2 peak. The corrected S2 parameters stored by the *multipeak* plugin rely on the standard **straxen** correction maps. These are evaluated independently for each S2 peak at its corresponding reconstructed position.

The simulation–data comparison of several observables for the five largest S2 peaks, as extracted with the *multipeak* plugin, is shown in Figure 5.20. The comparison is performed using ^{208}Tl events from ^{232}Th calibration data. The reconstructed areas of the S2 peaks are well reproduced by the simulation, with the five S2 peaks exhibiting very similar distributions in both scale and shape.

For the S2₂, larger statistical fluctuations are observed, as expected from its lower rate, but the overall area shape remains consistent between data and simulation. Likewise, the z distributions show comparable trends across the five peaks, confirming that the detector response model captures the spatial separation of interaction sites within the detector. The radial distributions show very good agreement between data and simulation, indicating that the position and field corrections are valid at the peak level. The overall agreement demonstrates that the *multipeak* reconstruction reliably captures the dominant peak properties and topology of multi-scatter events, providing a robust basis for the analysis that follows. Additionally, the ratio between the corrected largest and third-largest S2 peaks ($S2_0/S2_2$) is shown in Figure 5.21. The agreement between data and simulation is strong, with most data points lying within the 1σ uncertainty band of the simulation. The agreement observed in this parameter space shows that we model not only the main S2 area accurately, but also the areas of the additional corrected S2 peaks. This constitutes an important step in validating our ability to reconstruct multi-scatter topologies reliably. This observable is particularly relevant for the analysis presented in Chapter 6, where it serves as one of the main handles for background rejection.




Figure 5.21: Distribution of the ratio between the largest and third-largest S2 peaks ($S2_0/S2_2$) per event, comparing the ^{208}Tl peak from ^{232}Th calibration data and simulation. Events from both the top and bottom source positions are included. The simulation reproduces the shape and normalization of this topology-sensitive observable. The lower panel shows the normalized residuals, with most bins within the 1σ uncertainty band.

5.6 Summary and outlook

The studies presented in this chapter show that the detector response is reproduced by the simulation at a level adequate for analyses in the MeV energy regime. The comparison of S1 and S2 observables demonstrates that the main signal properties and their dependence on detector position are described with sufficient accuracy to support studies driven by event topology. While some differences between data and simulation remain, these do not alter the qualitative description of event topology on which high-energy analyses rely. In addition to validating the detector response in this energy range, this work establishes the reconstruction tools and observables required to access peak-level topology, most notably through the *multipeak* algorithm. These ingredients enable the use of S2 topological information as the primary handle for distinguishing signal from background in MeV-scale searches. More broadly, this chapter provides a reference validation of the XENONnT TPC response at MeV energies, extending the characterization beyond the low-energy regime traditionally explored. The observables and reconstruction framework developed here are employed in Chapter 6, where they form the basis of the search for the ^{124}Xe decay via electron capture with positron emission in a regime dominated by multi-scatter radiogenic background.

Chapter 6

Sensitivity Projection for the $2\nu\text{EC}\beta^+$ decay of ^{124}Xe

In the previous chapter, we demonstrated that the response of the XENONnT TPC can be modeled with good accuracy at energies of a few MeV. This provides the basis for exploring rare-event signatures in this energy range and evaluating the discovery potential of the experiment beyond its primary dark matter program. In this chapter, we present the first sensitivity projection for the detection of $2\nu\text{EC}\beta^+$ decay of ^{124}Xe to the ground state of ^{124}Te . Among the possible decay modes of ^{124}Xe , this channel is expected to have the shortest half-life making it the most favorable candidate for observation [116, 243]. Its distinctive topology, combined with the characteristics of a xenon dual-phase TPC such as XENONnT, its large active target mass, isotopic composition, and good energy resolution, makes this technology particularly well suited for such a search. In contrast to other ER rare event searches [122], this decay is characterized by a multi-scatter signature, resulting in a background dominated by radiogenic sources and an energy spectrum that extends over a broad range. Consequently, traditional one-dimensional approaches based solely on reconstructed energy are of limited effectiveness. The analysis developed here extends standard strategies by incorporating additional information from event topology to improve discrimination between signal and background. The study is carried out in a simulation-driven approach built directly upon the detector modeling and multi-scatter reconstruction tools validated in Chapter 5.

This chapter is organized as follows: the experimental signature of the $2\nu\text{EC}\beta^+$ decay in a xenon dual-phase TPC is introduced in 6.1. The overall analysis strategy is then outlined in 6.2, where dedicated signal simulations are described and used to define the event preselections discussed in 6.2.2 and 6.2.3. The background sources are discussed




Figure 6.1: Illustration of the two-neutrino electron capture with positron emission decay to the ground state of ^{124}Te . A ^{124}Xe atom captures an electron and emits a positron together with two neutrinos, transforming into ^{124}Te . This transition is accompanied by the emission of X-rays and Auger electrons (X_k), which are produced when the atomic vacancy left by the captured electron is filled during the de-excitation process.

in 6.3 and 6.4, followed by the background-rejection techniques in 6.5, including the optimization of the fiducial volume in 6.5.1 and the development of a machine learning (ML) classifier in 6.5.2. Finally, the statistical framework and projected sensitivity are presented in 6.6, with the result reported as a 90 % confidence level lower limit on the half-life derived using a profiled likelihood-ratio test.

6.1 Signature of $2\nu\text{EC}\beta^+$ decay in a xenon dual-phase TPC

In the $2\nu\text{EC}\beta^+$ decay, a ^{124}Xe nucleus captures an orbital electron that coincides with the emission of a positron. In this work, we consider only the electron capture corresponding to the K-shell electrons, as they have the largest probability of being captured due to the proximity to the nucleus [99]. The electron capture creates a vacancy in the inner atomic shell, which is filled through atomic relaxation, resulting in the emission of X-rays and Auger electrons (X_k). In the transition to the ground state of ^{124}Te , two electron neutrinos are also emitted. The decay process is illustrated in Figure 6.1 and can be described by the following reaction:

$$^{124}\text{Xe} + e^- \rightarrow ^{124}\text{Te} + e^+ + 2\nu_e + X_k, \quad (6.1)$$

The decay features a characteristic Q -value of $(2856.73 \pm 0.12) \text{ keV}$ [114]. The K-shell atomic relaxation products have a fixed total energy of $E_k = (31.810 \pm 0.012) \text{ keV}$ [244], which is typically deposited locally near the decay vertex due to the $\mathcal{O}(\text{mm})$ mass




Figure 6.2: Simulation of the energy depositions produced by the $2\nu\text{EC}\beta^+$ decay in the XENONnT TPC. A positron (e^+) and atomic de-excitation quanta (X_k) are emitted simultaneously at the decay site, producing the earliest localized energy deposits (shown as the overlapped purple and red dots). The positron subsequently annihilates, yielding two back-to-back 511 keV gamma-rays (shown in teal and blue). These undergo multiple Compton scattering before photoabsorption indicated by the dotted lines connecting distinct scatters, generating spatially separated later energy deposits. Circles indicate individual energy depositions, with their size proportional to deposited energy and their color representing the time relative to the event start, except for the X_k which is indicated in red for clear visualization.

attenuation length of $\mathcal{O}(\text{keV})$ photons (see Figure 2.1). The maximum kinetic energy available to the final-state particles is $Q - 2m_e c^2 - E_k = 1802.92 \text{ keV}$, where the term $2m_e c^2$ accounts for the emission of the positron and the capture of an atomic electron. This energy is shared between the positron kinetic energy, $0 \leq T_{e^+} \leq 1802.92 \text{ keV}$, and the neutrinos, which carry away the remaining fraction and typically escape undetected.

A potential $2\nu\text{EC}\beta^+$ decay in a xenon dual-phase TPC would result in multiple ER scatterings. A schematic representation of the expected energy depositions following the decay is shown in Figure 6.2. Multiple distinct energy depositions are expected: (1) the atomic relaxation products are emitted in coincidence with the positron and are absorbed near the decay site; (2) the positron itself generates a short ionization track through a few mm before thermalization and subsequently annihilates nearly at rest; (3) upon annihilation, the positron produces two back-to-back gamma rays, each of which undergoes multiple Compton scatterings followed by photoabsorption. The number of energy depositions produced by this decay, together with the realistic signature in the XENONnT TPC will be characterized in detail in Section 6.2.1 by means of dedicated signal simulation.

The expected number of $2\nu\text{EC}\beta^+$ decay events is estimated from Equation 1.15, using $\eta = 9.5 \times 10^{-4}$ as the natural isotopic abundance of ^{124}Xe , $m_u = 131.293 \text{ g/mol}$, $m \cdot t$

is the exposure and $T_{1/2}$ as the predicted half-life of the decay. Based on the half-life prediction from [116], $T_{1/2} = (1.3 \pm 0.5) \cdot 10^{23}$ yr, and the total active LXe mass of 5.9 t, the expected number of decays is $\sim (137 \pm 52)$ events per year, where the dominant source of uncertainty comes from the uncertainty in the predicted half-life. From this estimate, the number of detectable signal events is smaller and is obtained by folding this expected rate with the analysis efficiency, which is determined from the full detector simulation and the selection chain described in the following sections.

In the next section we discuss a set of *preselection* criteria, designed to ensure good event reconstruction quality, full topology containment and background suppression. The resulting dataset will be used as input for developing the ML classifier, which provides further background discrimination.

6.2 Analysis strategy

The experimental challenge in searching for the $2\nu\text{EC}\beta^+$ decay lies in isolating rare multi-scatter signal ER events from a background dominated by high-energy gamma interactions. While both signal and background can populate similar energy ranges, their spatial and topological characteristics differ, providing a handle for discrimination. Signal events tend to produce multiple energy depositions with comparable S2 sizes, leading to a more even distribution of S2 area across the event. In contrast, the dominant radiogenic backgrounds, arising from single high-energy gamma rays, more often exhibit a topology with one dominant S2 peak accompanied by smaller secondary scatters. These differences motivate the use of observables that characterize the spatial extent, S2 multiplicity, and the distribution of S2 area among peaks for signal-background discrimination.

The sensitivity study presented in this work is based on a simulation-driven analysis in which both the signal and the relevant background components are modeled using the framework described in Chapter 5. The background simulations serve three purposes in the analysis. First, they provide the dataset used to train a multivariate classifier designed to exploit the characteristic multi-scatter topology of the $2\nu\text{EC}\beta^+$ signal. Second, they are used to optimize the fiducial volume by maximizing the signal-to-background ratio in spatial bins of the TPC. Third, they define the expected background distributions entering the statistical model. The event selection applied to data follows the same reconstruction and quality criteria discussed in Chapter 5.

In addition, a set of physics-motivated preselection requirements are imposed. These include a selection on the total reconstructed energy, aimed at retaining events in which




Figure 6.3: Kinetic energy distribution of 3×10^6 $2\nu\text{EC}\beta^+$ decay events generated with DECAY0. The dashed lines indicate the kinematic endpoint of the positron spectrum and the energy of the K-shell atomic relaxation.

all decay products are detected. Events are further required to fall within the signal-dominated region of the $(\text{cS1}_{\text{multi}}, \text{cS2}_{\text{sum}})$ parameter space. This region is defined by the 95th percentile contour of the simulated signal distribution.

The subsections below describe the signal simulation, the individual preselection criteria, and the optimization of the fiducial volume, together with their impact on the signal acceptance and background rejection. These ingredients define the dataset used for training a ML classifier and for building the statistical model employed in the sensitivity projection.

6.2.1 Signal simulation

Within the simulation framework of XENONnT, the generation of $2\nu\text{EC}\beta^+$ decay events is implemented at the stage of the primary particle generators (see Figure 5.1). Specifically, the decay products kinematics are generated with the DECAY0 event generator [245], which provides the momenta, emission times, and particle type. The kinetic energy distribution of 3×10^6 decays generated by DECAY0 is shown in Figure 6.3. The positron spectrum, shown in black, follows the expected continuous distribution, extending up to the endpoint at 1802.92 keV. The atomic relaxation kinetic energy, shown in pink, appears as a narrow peak at 31.81 keV, corresponding to the K-shell electron capture. The kinematics are passed as input to **Geant4**, where their passage through the detector geometry and interactions are simulated. The resulting energy




Figure 6.4: Signal waveforms for an example simulated event. The primary scintillation signal is shown in blue, while the ionization signals are shown in red, illustrating the multiple-scatter nature of the event, including one S2 (fifth panel) exhibiting a significantly smaller amplitude than the others.

depositions are subsequently processed by `fuse` and the standard reconstruction chain, following the sequence illustrated in Figure 5.1.

In this work, 3×10^6 signal events were simulated with the full XENONnT simulation chain. In an event, each interaction in the detector can, in principle, give rise to a distinct S2 signal, provided sufficient spatial and temporal separation. The atomic relaxation products are typically reconstructed as merged with the positron S2 due to their proximity, while each 511 keV annihilation gamma ray can generate multiple additional S2 signals through Compton scattering. Consistent with this picture, the signal topology is predominantly multi-scatter, with 85 % of simulated events reconstructed as multi-scatter and the remaining 15 % as single-scatter, arising from unresolved S2 signals. Given that the S1 light propagation occurs on $\mathcal{O}(\text{ns})$ timescales, the prompt scintillation from these multiple deposits is reconstructed as a single S1 pulse. The waveforms for a representative simulated signal event are shown in Figure 6.4. The S2 multiplicity distribution for multi-scatter events is shown in Figure 6.5 left. Although some events exhibit up to nine S2 peaks, the majority contain five or fewer. The corresponding S2 area distributions as a function of peak rank are shown in Figure 6.5 right. S2 peaks are ordered by decreasing area by construction. For this channel, the S2 area distribution shows a pronounced fall-off beyond the first few peaks, with peaks beyond the sixth typically below 10 PE. At these scales, the signals are comparable to contributions from few-electron emission and instrumental backgrounds such as afterpulses and photoionization induced delayed electrons, making it difficult to reliably distinguish physical S2 peaks from detector-induced signals. Consequently, higher-order peaks carry limited additional information, supporting the choice of using only the largest S2 peaks for topological discrimination.

The total deposited energy from the $2\nu\text{EC}\beta^+$ decay, shown in Figure 6.6, is reconstructed using Equations 5.3, where all S2 peaks contributing to the event are included. Because of the continuous kinetic energy available to the positron from the three-




Figure 6.5: **Left:** $2\nu\text{EC}\beta^+$ S2 peak multiplicity per event. Some events have up to nine S2s, although the majority of events have five or less. **Right:** Distribution of S2 areas as a function of peak index in the event. For each S2, the box shows the central 50% of events (interquartile range), the red line marks the median, and the whiskers indicate the spread of non-outlier values.

body decay, the reconstructed energy in the detector spans a wide range dominated by the positron contribution together with the annihilation gamma rays that contribute through subsequent interactions in the detector.

The resulting total reconstructed energy distribution overlaps with common ER backgrounds, reducing the effectiveness of a signal selection based purely on energy. Together, these considerations motivate an analysis focused on multi-scatter events, where peak-level observables can be used to exploit the distinctive topology of the signal. However, a selection in reconstructed energy remains useful, to retain events with fully contained topologies, as discussed in the next section.

6.2.2 Reconstructed energy selection

The total reconstructed energy spectrum for simulated $2\nu\text{EC}\beta^+$ events reconstructed as multi-scatter is shown in Figure 6.6. The measurable energy arises from the positron kinetic energy ($0 \leq T_{e^+} \leq 1802.92 \text{ keV}$), the two 511 keV annihilation gamma rays, and the 31.8 keV atomic de-excitation products following the electron capture process. Since the neutrinos escape undetected, the reconstructed signal spectrum is continuous up to the decay Q -value of 2856 keV.

The escape of one or more annihilation gamma rays reduces the reconstructed energy by up to 511 keV per escaping gamma, thereby shifting partially contained events toward lower reconstructed energies. Consequently, events with reconstructed energies below 1053 keV cannot originate from fully contained signal topologies. The threshold at




Figure 6.6: Total reconstructed energy distribution of 3×10^6 $2\nu\text{EC}\beta^+$ decay events, shown for events reconstructed as multi-scatter. The signal spectrum shape is correlated with distinct event topologies, as discussed in the text. The dotted lines indicate the energy preselection boundaries applied in this analysis.

1053 keV corresponds to the combined energy of the two annihilation gamma rays and the atomic relaxation products. Events below this threshold exhibit signatures more similar to radiogenic gamma-ray background events. Motivated by this consideration, an energy preselection is applied:

$$1053 < E < 2856 \text{ keV}.$$

This selection defines the region in which fully contained signal events are expected to be reconstructed. Partially contained topologies may still contribute within the accepted energy range due to the continuous sharing of energy between the positron and the emitted neutrinos. The requirement preserves $\sim 93\%$ of simulated signal events and defines the starting point of the topology-based analysis. While not intended as a dedicated background rejection criterion, the selection also removes background events outside the accepted energy window, summarized in Table 6.1.

6.2.3 Signal distribution in the observable space

A selection in the observable space ($cS1_{\text{multi}}, cS2_{\text{sum}}$) is defined with the purpose of suppressing additional background events passing the previous preselections but lying outside the 95% signal distribution in this space. Additionally, the region outside the contour naturally defines a sideband, which can be useful for validating the classifier on measured data and for potential unblinding strategies. The latter is, however, beyond




Figure 6.7: Simulated signal distribution in the observable space. The color map shows the signal event density with only data quality selections applied. Solid and dashed black lines correspond to the 95 % and 68 % contour regions, respectively. Gray dashed lines indicate the constant energy curves defining the energy cut in Equation 6.2.2.

the scope of this work.

The simulated $2\nu\text{EC}\beta^+$ signal distribution in the observable space is shown in Figure 6.7. The selection in this parameter space is defined by a closed, smoothed contour (shown in black) that encloses 95 % of the signal events that survive the previous preselections, and retains approximately 95 % of the preselected signal sample. The corresponding ROI in the reconstructed energy space is given by the projection of this selection onto the energy axis, and is used in the statistical analysis presented in Section 6.6.

In the following, *preselections* refer to the data quality cuts described in Chapter 5 together with the reconstructed energy and $(\text{cS1}_{\text{multi}}, \text{cS2}_{\text{sum}})$ selections. The signal and background acceptance after each preselection step, are summarized in Table 6.1. Before introducing the techniques used for background mitigation, the next sections provide an overview of the main background sources relevant to the $2\nu\text{EC}\beta^+$ decay search.

6.3 Internal backgrounds

Intrinsic backgrounds originate from radioactive isotopes dissolved in the LXe target itself, which are often emanated from detector materials and cannot be mitigated through fiducial volume selections. These backgrounds are particularly relevant for rare event searches in the high-energy regime and can contribute significantly to the expected event rate in the region of interest.

The main intrinsic background source relevant to the $2\nu\text{EC}\beta^+$ decay analysis is ^{222}Rn , produced as a daughter of ^{226}Ra , which itself is part of the ^{238}U decay chain (see Figure 2.6). Trace amounts of ^{238}U and its progeny are naturally present in most detector materials. Over time, ^{222}Rn is released from these materials, either through surface emanation or diffusion, and can enter the xenon purification and recirculation system. Being chemically inert and gaseous, radon readily mixes with the xenon gas and subsequently dissolves into the LXe bulk, where it becomes uniformly distributed and contributes to the intrinsic background.

All transitions in the ^{222}Rn decay chain proceed via either α or β -decay. Due to their high mass, α particles generate dense and localized ionization tracks in LXe, which leads to increased recombination. This produces a large, narrow S1 signal and a suppressed S2 compared to typical ER interactions. As a result, these events exhibit a characteristically low S2/S1 ratio. Due to this, α events are eliminated by the S1 width cut and further suppressed by the (cS1_{multi}, cS2_{sum}) selection.

The ^{214}Bi β^- decay to ^{214}Po has a half-life of 19.9 minutes and a Q-value of 3.27 MeV. It proceeds through multiple branches, including several that emit high-energy gamma rays. The most prominent lines are: 609.3 keV (46 %), 1121 keV (16%), 1764.5 keV (15%), 2204.1 keV (5%) [246]. The first line falls below the energy preselection threshold of 1053 keV and is effectively rejected. However, the higher-energy gamma rays can undergo multiple Compton scattering in the LXe and produce MS topologies. The contribution of this background is evaluated using full-chain simulations, in which 1.2×10^6 ^{214}Bi decays are generated. From these simulations, the fraction of MS events passing the preselections is 28.82 %.

Following the β^- decay of ^{214}Bi to ^{214}Po , the latter undergoes an α -decay to ^{210}Pb with a characteristic half-life of 164 μs . This short-lived sequential decay results in a distinctive topology, commonly referred to as a BiPo event, which can be used for tagging and identifying ^{214}Bi contamination. BiPo events are typically characterized by two S1 and two S2 signals, one pair from the β -decay and another from the subsequent α -decay. Due to the short half-life of ^{214}Po , both decays often fall within the same acquisition window. In cases where the β transition proceeds via excited states, additional S2 peaks from de-excitation gamma-rays may also be present. The combination of low S2 and high S1 from the α -decay, along with the overall multi-S1 and multi-S2 structure, results in a recognizable coincident topology that can be used to tag and reject those events. In XENONnT, the data-driven rejection efficiency is 99.97 %.

The ^{214}Bi decays occurring in the LXe shell, the region between the TPC boundaries and the cryostat, cannot be identified by this coincidence. In such cases, the α particle

is not detected because it deposits its energy outside the sensitive LXe volume, while the associated gamma-rays can enter the active region and produce detectable signals that lack the time coincident α , making them more difficult to tag and reject. Since these events are located close to the edges of the TPC, they are effectively rejected by the fiducial volume selection. Their contribution is evaluated using full-chain simulations, in which 2.2×10^6 ^{214}Bi decays in the LXe shell are generated, from which the preselection acceptance is estimated to be 28.6 %.

The rate of ^{214}Bi decays in the active LXe volume and in the LXe shell is constrained using α -decays from the ^{222}Rn chain. Since some positively charged daughter ions can drift to detector surfaces and decay outside the active volume, the chain is not necessarily in equilibrium. The activity of ^{214}Bi is constrained between the measured rates of ^{218}Po and ^{214}Po α -decays, corresponding to 0.809–1.763 $\mu\text{Bq/kg}$ during SR0. The lower bound corresponds to the scenario of maximal daughter plate-out and is used in this work to scale the simulations to the expected activity, calculated using Equation 6.4.

6.4 Materials backgrounds

The materials used for the detector construction underwent several screening stages that determined whether their radiopurity level is sufficiently low for the XENONnT background requirements in the WIMP ROI. The trace radioimpurities present in the detector components, however, constitute the dominant background in the energy range relevant for this search. The most concerning transitions are those produced by the decay chains of ^{238}U and ^{232}Th (see Figure 2.6). For the former, the beta transitions to excited states and their accompanying relaxation gamma-rays produce multi-scatter ERs within the ROI. A representative example is the ^{214}Bi β^- decay.

A dedicated screening measurement campaign reported in Ref. [175, 247], was performed in order to determine the activity of the isotopes of the materials. The isotopes considered are ^{232}Th , ^{238}U , ^{235}U , ^{40}K and ^{60}Co . A break in secular equilibrium has been observed in several materials used in the experiment, as reported in Ref. [175, 247]. In particular, there is disequilibrium between ^{238}U and its daughter ^{226}Ra , as well as between ^{232}Th and ^{228}Th . Under secular equilibrium, each isotope in a decay chain has the same activity as its parent, meaning the daughter is produced and decays at the same rate. This condition is typically established when the half-life of the parent is much longer than that of the daughter. However, equilibrium can be disrupted when chemical or physical processes separate parent and daughter isotopes, or when screening measurements reveal different concentrations of isotopes within the same chain. To

reflect this in the MC simulations, the decay chains of ^{238}U and ^{232}Th are split into two parts, an early and a late part, corresponding to isotopes before and after the observed break. Although material screening results indicate disequilibrium between ^{232}Th and ^{228}Ra , the decay chain is split at ^{228}Th in the simulations for practical reasons. This choice does not impact the physics of the background model, as ^{228}Ra , ^{228}Ac , and ^{228}Th rapidly reach equilibrium with each other due to their relatively short half-lives.

$$\begin{array}{ll} ^{238}\text{U}_{\text{early}} : & ^{238}\text{U} \rightarrow ^{230}\text{Th} \\ ^{238}\text{U}_{\text{late}} : & ^{226}\text{Ra} \rightarrow ^{206}\text{Pb} \end{array} \qquad \begin{array}{ll} ^{232}\text{Th}_{\text{early}} : & ^{232}\text{Th} \rightarrow ^{228}\text{Ac} \\ ^{232}\text{Th}_{\text{late}} : & ^{228}\text{Th} \rightarrow ^{208}\text{Pb}. \end{array}$$

Combined with the measured activities from material screening, these simulations enable a realistic estimation of the expected number of background events for the XENONnT exposure considered in this analysis.

Since the screening measurements were performed prior to the SR0 data taking period (see Section 4.1), the activities of short-lived isotopes, such as ^{60}Co ($T_{1/2} = 5.27\text{ y}$) and ^{228}Th ($T_{1/2} = 1.91\text{ y}$), are scaled to reflect their decay over time. For example, the activity of ^{60}Co at the end of SR0 (12 December 2021) is calculated using the radioactive decay law:

$$A_{\text{present}}(t) = A_{\text{screening}} \cdot e^{\frac{-\ln 2 \cdot t}{T_{1/2}}}, \quad (6.2)$$

where $A_{\text{screening}}$ is the activity measured during the screening campaign, t is the elapsed time between the screening and 12 December 2021, and $T_{1/2}$ the half-life.

Since detector materials surrounding the active volume constitute the main source of high-energy gamma rays, and thus are the dominant background in the $2\nu\text{EC}\beta^+$ sensitivity study, we performed a full-chain simulation campaign using the framework described in Section 5.1. This work represents the first simulation of the XENONnT materials background carried out with the **fuse** framework.

6.4.1 Rate conversion

To obtain a physically meaningful interpretation of the simulations, the number of simulated events passing the selection criteria must be translated into an expected event rate in the detector. The ER background rate, $R_{i,j}$, for a radioactive isotope i in detector component j is given by

$$R_{i,j} = \frac{N_{i,j}^{\text{sel}}}{T_{i,j}^{\text{eff}} \times M_{\text{FV}}}, \quad (6.3)$$

where $N_{i,j}^{\text{sel}}$ is the number of simulated events that satisfy the selection criteria, M_{FV} is the fiducial mass corresponding to the fiducial volume discussed in Section 6.5.1, and $T_{i,j}^{\text{eff}}$ is the effective livetime. The latter represents the equivalent livetime required to accumulate the same statistics in data and is defined as

$$T_{i,j}^{\text{eff}} = \frac{N_{i,j}^{\text{G}}}{W_i \times M_j \times A_{i,j}}, \quad (6.4)$$

where $N_{i,j}^{\text{G}}$ is the number of events generated in **Geant4** for isotope i in component j , M_j is the mass of component j , and $A_{i,j}$ is the activity of isotope i in that component. Since **Geant4** treats individual steps in a decay chain as separate events, a single simulated event does not correspond to a full decay sequence. To account for this, a weighting factor W_i is applied, taken as the mean decay-chain length for isotope i .

To use the screening results from Ref. [175, 247] to obtain the activities of the components simulated in the SR0 materials simulation campaign (Appendix B.1), an additional reconstruction algorithm was implemented in **fuse** for this work, referred to as *event primary info* throughout this work. For each reconstructed event, the algorithm associates microphysics clusters whose time intervals are fully contained within the event window, and propagates their primary **Geant4** coordinates, $(x_{\text{pri}}, y_{\text{pri}}, z_{\text{pri}})$, to the event level. It also records the total deposited energy associated with that primary interaction. Screening activities are reported separately for individual detector sub-parts, while the simulations are performed at the parent component level. Simulated events are therefore selected according to their primary interaction coordinates and weighted using the activity of the corresponding sub-part. This procedure ensured a consistent normalization of the simulated rates to the screening measurements, which is a necessary input for the background rejection techniques discussed in the next section.

6.4.2 SR0 materials simulation campaign

A dedicated detector material full simulation campaign was carried out to model the backgrounds relevant for the $2\nu\text{EC}\beta^+$ sensitivity study. Each detector component was simulated separately considering the full decay chain of the isotopes listed above, using detector conditions corresponding to SR0. An overview of the number of generated primaries per component is provided in Appendix B, which was a large-scale production




Figure 6.8: Multi-scatter energy spectra of radiogenic background components within the ROI. The distributions are obtained from simulations that include all detector components, with the event rates normalized to the full active volume and expressed in units of events per year with a 10 keV binning. All preselection criteria have been applied prior to constructing the spectra.

carried out in OSG. Within this campaign, together with the event-level signal features accessible via standard low-energy reconstruction algorithms, the *multipeak* reconstruction algorithm was employed to access peak-level information in multi-scatter events, enabling the reconstruction of multiple S2 signals produced by high-energy gamma interactions from detector materials. The resulting topology information is later used as input to the ML discriminator.

The simulation output produced on OSG is generated per job in compressed tarball files. For each component–isotope combination, the datasets are transferred to the Midway remote server for analysis, extracted, and merged into a single sample for processing. The resulting merged datasets are used throughout this work for subsequent studies. This includes the energy bias and resolution evaluation presented in Chapter 5, and for the $2\nu\text{EC}\beta^+$ sensitivity projection, including the processing with *event primary info* algorithm, the application of the preselections described above, and the construction of the material background model. Basic validation checks were performed to ensure the consistency of the simulations, including the spatial distributions and reconstructed energy spectra, shown in Figure 6.8. Figure 6.9 shows the relative contributions of individual detector materials, indicating that the dominant contributions originate from the inner cryostat, PMTs, stainless steel bell, and outer cryostat.




Figure 6.9: Relative contributions of individual detector materials to the total background from ^{60}Co , ^{40}K , ^{232}Th , ^{226}Ra , ^{235}U , ^{238}U , and ^{228}Th within the ROI.

6.5 Background rejection

Having discussed the main background sources, we now quantify the efficiency of the preselection cuts within the ROI defined in the previous sections. The acceptance of each background component, as well as of the signal, after successive preselection steps is summarized in Table 6.1. Figure 6.10 shows the distributions of background events in the observable space after applying data quality and energy selections. As illustrated, these events populate regions that significantly overlap with the signal, leading to a non-negligible fraction surviving the preselection, as quantified in Table 6.1.

While the preselection cuts retain a high fraction of the signal, they are not sufficient to suppress background contributions to the required level. In particular, external gamma rays emitted by detector materials dominate after preselection, yielding $\mathcal{O}(10^6)$ events per year in the active volume. This exceeds the expected signal rate of ~ 113 events per year by several orders of magnitude and therefore constitutes the primary limitation of the analysis. Intrinsic backgrounds from ^{214}Bi , although less abundant, remain non-negligible after preselection, contributing at the level of $\mathcal{O}(10^3)$ events per year from both the LXe bulk and the LXe shell. These results demonstrate that the preselection alone does not provide sufficient discrimination between signal and background, due to both the large initial activity of detector materials and the similarity in event topology between signal and multi-scatter gamma interactions.




Figure 6.10: Distributions of intrinsic (**left**) and two representative material background sources (**right**) in the observable space after applying data quality and energy selection cuts (indicated in gray dashed lines). The shaded regions indicate the 68 % and 95 % density intervals. The black contour encloses the 95 % signal region. Events within this contour survive the preselections and are quantified in Table 6.1.

Source	Data quality [%]	Energy window [%]	cS1–cS2 [%]	Total rate [evt/yr]
Materials	88.95	35.62	10.23	4.18×10^6
LXe ^{214}Bi	95.05	75.14	28.82	610.22
LXe shell ^{214}Bi	98.53	42.35	28.66	1559.62
Signal	97.36	86.55	82.53	113.06

Table 6.1: Preselection cumulative acceptances in percentage for the different background sources and for the signal within the ROI. The corresponding total rate of each source is included in the last column, where the rate for ^{214}Bi has the BiPo selection efficiency applied.

To achieve the level of background suppression required for a sensitive search, additional discrimination strategies are necessary. This section presents two complementary approaches. The first defines an inner fiducial volume within the active region of the detector, exploiting the self-shielding of liquid xenon to reduce contributions from external gamma rays. In this work, the FV is determined through a simulation-driven optimization described in 6.5.1. The second approach employs a binary machine learning classifier to exploit subtle differences in event topology between signal and background. The development and performance of this classifier are presented in 6.5.2.

6.5.1 Fiducial volume selection

An optimized definition of the fiducial volume is required to suppress the material-induced background while preserving the maximal target mass. The FV is therefore determined through a simulation-driven optimization based on a figure of merit (FoM) proportional to the expected sensitivity of the analysis, defined using the standard

significance estimator $S/\sqrt{B}$ [248], where S is the signal and B the background rate.

The optimization is derived using multi-scatter events from the material background simulation campaign described in 6.4. To characterize the spatial dependence of the material background, the detector volume is partitioned into bins in the (R, z^2) plane, where the event positions, considering the corrected position of the primary S2 signal ($S2_0$), are used to assign events to spatial bins. The (R, z^2) plane is divided into a grid of 25×25 equally sized bins. For each bin, the background rate B_{ij} is evaluated, yielding a two-dimensional background density map.

Since the expected $2\nu\text{EC}\beta^+$ signal rate is uniform throughout the TPC, the signal yield in each spatial bin is proportional to the xenon mass contained in that bin. Because the binning is uniform in volume, this mass is approximately constant across bins. Under these conditions, the local optimization reduces to minimizing the background contribution, and a convenient proxy for the signal sensitivity is given by:

$$\text{FoM}_{ij} = \frac{1}{\sqrt{B_{ij}}}, \quad (6.5)$$

which preserves the spatial structure of the optimal sensitivity contours and serves as the basis for defining candidate FV boundaries. Iso-contours of the FoM_{ij} are constructed in the (R, z^2) plane. The result is shown in Figure 6.11, where the color map indicates the FoM_{ij} distribution in the TPC, and three contours at constant FoM_{ij} values are highlighted for reference. The observed FoM_{ij} distribution follows the expected spatial trend. It is largest in the central region of the TPC, where LXe self-shielding suppresses material induced backgrounds, and decreases toward larger radii and high z as the external radiogenic contribution increases.

To evaluate the performance of different FV definitions, a set of candidate contours is selected and described by a radius profile $R(z)$. The profile is modeled using two semi-superellipses describing the upper and lower regions of the contour, smoothly connected by a sigmoid blending function. The upper and lower regions are independently parametrized as:

$$R_{\text{top}}(z) = R_0^{\text{top}} \left[1 - \left(\frac{z - z_0^{\text{top}}}{z_{\text{top}} - z_0^{\text{top}}} \right)^{p_{\text{top}}} \right]^{1/q_{\text{top}}}, \quad (6.6)$$

$$R_{\text{bot}}(z) = R_0^{\text{bot}} \left[1 - \left(\frac{z_0^{\text{bot}} - z}{z_0^{\text{bot}} - z_{\text{bot}}} \right)^{p_{\text{bot}}} \right]^{1/q_{\text{bot}}}, \quad (6.7)$$




Figure 6.11: Simulation-driven FoM, represented by the color scale, in the R, z^2 parameter space. The three white lines indicate iso-contours of constant FoM at selected levels (0.13, 0.08, 0.045).

where R_0 is the maximum radius, z_0 denotes the height of the widest section, z_{top} and z_{bot} define the endpoints of the contour and p and q control the curvature of the profile. The transition height z_{mid} is defined as the location of the maximum radius of the extracted profile. The full parametrized contour is then obtained via:

$$R(z) = w(z) R_{\text{top}}(z) + [1 - w(z)] R_{\text{bot}}(z), \quad (6.8)$$

where $w(z)$ is a sigmoid function centered at z_{mid} , ensuring a smooth transition between both regions. The parameters of the upper and lower semi-superellipses are determined independently through a least squares fit. In Figure 6.12 left three representative contours are shown with their respective fitted functions.

The enclosed mass is then obtained by integrating the parametrized boundary. For each fitted contour, we compute a global figure of merit FoM_{FV} using:

$$\text{FoM}_{\text{FV}} = \frac{m_{\text{FV}}}{\sqrt{R_{\text{FV}}}}, \quad (6.9)$$

where m_{FV} is the xenon mass inside the volume, also known as fiducial mass, delimited by the respective candidate FV. The FoM dependence on the xenon mass, shown in Figure 6.12 right, exhibits the expected trade-off between background rejection and signal acceptance: small fiducial mass reduce background but also significantly decrease signal, while large fiducial mass increase the background rate. The optimal FV, shown in Figure 6.13 is therefore defined by the contour that maximizes the global FoM. In




Figure 6.12: Left: FoM_{ij} iso-contours overlaid with their corresponding fitted functions. **Right:** Global FoM as a function of the xenon mass enclosed by each candidate contour. The parabolic trend reflects the trade-off between background rejection and signal acceptance as the fiducial volume is varied. The optimal FV is defined by the contour that maximizes the fiducial mass.

this case, it corresponds to the iso-contour 0.08, which encloses a total xenon mass of 905 kg. This selection yields a signal acceptance of 11 % and a background rejection of 99.9 %.

A natural question is whether the BDT classifier should be applied prior to the FV optimization, such that the contour scan is performed on the post-BDT event sample. However, the method presented in this section exploits the fact that the dominant background exhibits strong position dependence in the detector, such that enlarging or shrinking the candidate volume corresponds to a well-defined tradeoff between signal acceptance and background contamination. For the optimization to be meaningful, the residual background map must retain a spatial gradient so that the contour scan traces a physically motivated boundary, as shown in Figure 6.12 left. Applying a BDT cut prior to the FV optimization was found to significantly alter the spatial distributions of signal and background in a strong, non-uniform way. As a consequence, the post-BDT background map exhibits weaker spatial coherence and less smooth structure. The resulting local FoM map becomes increasingly irregular, and the FV contours begin to follow small-scale spatial features rather than a smooth and physically interpretable trend. In this regime, the FV optimization is no longer driven by the detector geometry and background distribution, but instead becomes sensitive to where the classifier happens to perform better or worse across the detector volume. For this reason, the BDT is not applied prior to the FV optimization. Instead, the FV is defined using the pre-BDT background distribution, which preserves the spatial structure required for a robust geometric optimization. The BDT selection is then applied within the chosen fiducial region when evaluating the final sensitivity.




Figure 6.13: Spatial distribution of the material background, with the optimized FV derived in this section shown in pink. The selected FV encloses a xenon mass of 905 kg and achieves a material background rejection efficiency of approximately 99.9 %.

6.5.2 Machine learning discriminator

As discussed previously, the main challenge of this analysis lies in effectively isolating the $2\nu\text{EC}\beta^+$ decay signal from background, dominated by high-energy gamma-rays emitted by detector materials. These gamma-rays can undergo multiple Compton scatterings, producing multi-site energy depositions that partially mimic the signal topology, although they generally arise from a single gamma cascade rather than from the coincidence of several particles. Adding to this, the $2\nu\text{EC}\beta^+$ signal exhibits a broad and continuous energy spectrum, which prevents the use of a narrow energy window for event selection. Given the exposure collected during SR0, the expected number of signal events after preselection is on the order of a few tens. This imposes the requirement of maximizing the signal acceptance while achieving strong background suppression. To address this, we employ ML techniques to exploit complex multi-dimensional correlations among event observables. This approach increases our ability to discriminate signal from background beyond what is achievable through traditional cut-based methods.

The classifier used in this analysis follows a supervised learning approach [249], in which simulated signal and background events are provided as labeled examples during training. The algorithm learns to map reconstructed observables onto a single discriminant variable that quantifies how signal like or background like a given event is.

The model is based on boosted decision trees [250, 251], where each individual tree




Figure 6.14: Flowchart illustrating the training and evaluation pipeline of the supervised learning model. The simulated dataset is split into training, validation, and test samples. The training and validation sets are used for model fitting and hyperparameter optimization, while the final performance is evaluated on the test sample.

performs a sequence of binary splits on input observables to separate signal from background classes. At each split, the algorithm selects the variable and threshold that best improves the separation according to an impurity metric, such as the Gini index¹, producing regions that are increasingly enriched in one class. Each terminal node, or leaf, assigns a prediction based on the class composition of events falling into that region. While a single decision tree provides an interpretable classifier, it is typically limited in performance and sensitive to statistical fluctuations.

To improve robustness and discrimination power, multiple trees are combined into an ensemble using boosting techniques. In this analysis, gradient boosting [251, 252] is employed, in which trees are added sequentially to minimize a loss function. For binary classification, the loss corresponds to the binomial deviance, which measures the disagreement between the predicted and true labels. Each successive tree is trained on the residuals of the current ensemble prediction, effectively focusing on events that are most difficult to classify.

The final classifier output is constructed as the sum of the contributions from all trees in the ensemble, yielding a score between 0 and 1. This continuous score is interpreted as a measure of signal compatibility and is later used as a selection variable in the analysis. In this work, the classifier output is used as a cut to reject events below an optimum threshold determined through an Asimov significance study.

¹The Gini index is a metric used in decision trees to measure how mixed the class labels are within a node.

Model description

The classifier used in this analysis is a Gradient Boosted Decision Tree (GBDT) implemented with `scikit-learn` [253], a widely used multivariate technique in particle-physics applications [254, 255]. GBDTs are well suited for this problem as they capture nonlinear correlations among observables while remaining robust in regimes with limited signal statistics.

The model is trained in a supervised learning framework using simulated datasets, with signal events from the $2\nu\text{EC}\beta^+$ decay simulation and background events from the material background campaign. After feature extraction, the samples are combined into a single dataset with binary labels (1 for signal, 0 for background), and split into statistically independent training and test sets using a stratified split (70% / 30%) to preserve class proportions. The full workflow is illustrated in Figure 6.14.

Input variables are standardized using parameters derived from the training sample and applied consistently to both datasets. To account for class imbalance, inverse-frequency event weights are assigned during training.

The hyperparameters of the BDT are optimized using k-fold cross-validation on the training dataset. The optimal configuration is selected by maximizing the mean validation Receiver Operating Characteristic (ROC). The best and most stable performance is achieved with 200 estimators and a learning rate of 0.1, which are adopted for the final model.

Discrimination space

The objective of the classifier is to use correlations between multiple physics-driven features that summarize the event topology. Instead of overloading the model with raw peak information already constrained by preselection cuts, the BDT is trained on event-level features that characterize the global event structure. It does not attempt to identify individual quanta of the decay, but rather to distinguish signal events as a whole based on their characteristic topology.

For this purpose, events passing the preselection are evaluated using five per-event features, shown in Figure 6.15 where the distributions are displayed for both signal and background events, described below:

Maximum transverse and longitudinal separation between S2 peaks: The multi-scatter topology is characterized by an extended cascade of energy deposits. To capture this, we define $\Delta r_T^{\max}$ and $\Delta z^{\max}$ as the maximum transverse and longitudinal




Figure 6.15: Features used for training the BDT classifier. (a) Maximum transverse separation between S2 peaks within the same event. (b) Maximum longitudinal separation between S2 peaks within the same event. (c) Ratio between the largest and third largest S2 peaks in the event. (d) Number of reconstructed S2 peaks per event. (e) Fraction of the total S2 signal not contained in the dominant peak. The background shown corresponds to radiogenic sources.

separations among the five largest S2 peak pairs (i, j) as:

$$\Delta r_T^{\max} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}, \Delta z^{\max} = |z_i - z_j|, \quad (6.10)$$

Distances built from a fixed peak pair are not uniquely interpretable, since reconstructed S2 peaks cannot be unambiguously associated with a specific particle or interaction within the event. In contrast, the maximum separation provides an observable that reflects the global spatial extent of the cascade without relying on an event-by-event labeling of peaks. Since detector spatial resolution differs between transverse and longitudinal directions, the two distances are treated separately in the analysis.

Ratio of the main and third largest S2 peak areas: To encode information about the relative sizes of the S2 signals, we define a training feature based on the ratio between the largest corrected S2 peak in an event (cS2_0) and the third largest (cS2_2). In signal events, cS2_0 is typically produced by the energy deposition of the e^+ , which is expected to yield the largest S2. However, if e^+ carries less than 511 keV, its S2 size may become comparable to that of the peaks induced by gamma-ray. The second largest peaks (cS2_1) and third generally correspond to one of the gamma-ray or, in some cases, to the e^+ . The ratio $\frac{\text{cS2}_0}{\text{cS2}_2}$ therefore provides a measure of the hierarchy between dominant and subdominant S2 signals within the same event.

Energy fraction not contained in the main S2 peak: While $\frac{\text{cS2}_0}{\text{cS2}_2}$ captures correlations between specific peaks, it remains sensitive to the presence and reconstruction of individual subdominant S2s. To provide a more global characterization of the event topology, we also consider the fraction of the total S2 signal not contained in the dominant peak, defined as $1 - \frac{\text{cS2}_0}{\text{cS2}_{\text{tot}}}$, where cS2_{tot} is the sum of the corrected areas of all valid reconstructed S2 peaks in the event. This variable quantifies how distributed the S2 signal is among multiple S2 peaks. Events with one dominant interaction are characterized by low values, whereas events with multiple comparable energy deposits exhibit higher values. As such, this observable provides a robust measure of the degree of multi-site structure in the event, complementing the more local hierarchy captured by $\frac{\text{cS2}_0}{\text{cS2}_2}$.

Number of S2 peaks per event: The S2 multiplicity, N_{S2} , is included as a training feature, as the signal of interest is characterized by multiple spatially separated energy deposits. Both signal and background events can exhibit multi-site topologies, however, signal events tend to produce a larger number of reconstructed S2 peaks, reflecting a more fragmented pattern of energy deposition. In contrast, background events more frequently exhibit fewer resolved S2 peaks, although higher multiplicities are also possi-




Figure 6.16: Relative importance of the input features used in the BDT training. The importance score reflects the contribution of each variable to the decision making process, quantified by the fraction of total splits in which the feature is used across the ensemble.

ble, particularly for backgrounds involving multiple Compton scatterings. As a result, N_{S2} provides useful discrimination power, especially when combined with observables that characterize the spatial separation and relative size of the S2 peaks.

The relevance of this set of training features was evaluated via a k -fold cross-validation (with $k = 5$). In this procedure, the dataset is split into k parts, in each iteration four parts are used for training and the remaining part is used for testing. This way, the model is trained and evaluated five times, each time with a different test set. If the model performance is consistent across all folds, it indicates that the results are stable and not strongly dependent on a particular split of the data. From this process, we obtain a distribution of the area under the curve (AUC) values together with its standard deviation σ . This information is important because it provides a measure of the statistical fluctuations in the performance of the model. Later, when testing what happens if we exclude a given feature, we can compare the resulting AUC against this distribution and check whether the difference is statistically significant.

Performance metrics and application

Once the model is trained, the relevance of the input features is evaluated using the feature importance score. This metric reflects how frequently each variable is used to perform splits across the ensemble of decision trees, and therefore provides insight into which observables contribute most to the classification.

The resulting feature importance is shown in Figure 6.16. The fraction of S2 signal not contained in the dominant peak is the most relevant variable, accounting for ap-




Figure 6.17: **Left:** Distribution of the BDT classifier output score for signal (red) and background (blue) events. **Right:** ROC curve corresponding to the classifier performance, with an area under the curve (AUC) of 0.857. The gray dashed line indicates the performance of a random classifier for reference.

proximately 38 % of the total importance. The second most important feature is the maximum longitudinal separation between S2 peaks, contributing about 32 %, highlighting the relevance of the spatial extent of the event along the drift direction. This is followed by the ratio between the largest and third largest S2 peaks, which contributes around 20 %, capturing the hierarchy between dominant and subdominant energy deposits. The remaining features, namely the transverse separation between S2 peaks and the S2 multiplicity, contribute only marginally to the overall performance of the classifier.

As mentioned previously, the classifier assigns a score between 0 and 1 to each event in the test dataset, reflecting its probability of being signal like. The distribution of this score for signal and background events is shown in Figure 6.17 left. Ideally, signal and background would populate distinct regions of the score, with signal events clustering near 1 and background near 0. In practice, a partial overlap between the two distributions is observed, reflecting the intrinsic similarity between signal and background topologies in this channel. Nevertheless, a clear separation is visible, with signal events preferentially populating higher scores.

To classify events, a threshold can be applied to the BDT score. Varying this threshold defines a trade-off between signal efficiency and background rejection. This correlation is described by the ROC curve, shown in Figure 6.17 right, where each point corresponds to a different threshold choice. The ROC curve displays the true positive rate (TPR) as a function of the false positive rate (FPR). The TPR, also referred to as signal efficiency, represents the fraction of signal events correctly identified as signal. The FPR corresponds to the fraction of background events incorrectly classified as signal.




Figure 6.18: Asimov significance versus BDT classifier score for the $2\nu\text{EC}\beta^+$ decay of ^{124}Xe . The best significance is indicated by the vertical gray dotted line at 0.692.

The overall performance of the classifier is quantified by the area under the ROC curve (AUC), also shown as in Figure 6.17 right. The value 0.857 indicates a strong discrimination power, although not a perfect separation. This is consistent with the intrinsic complexity of the problem, as both signal and the dominant background exhibit multi-site topologies with overlapping characteristics.

The optimal threshold on the classifier score is determined by maximizing the expected sensitivity of the analysis. For this purpose, the Asimov significance [248], $Z_A(th)$, is used as the figure of merit:

$$Z_A(th) = \sqrt{2 \left[(s + b) \ln \left(1 + \frac{s}{b} \right) - s \right]}, \quad (6.11)$$

where $s = s(th)$ and $b = b(th)$ are the expected rate of signal and background, respectively, at a given classifier threshold th . The signal and background acceptances at each threshold are obtained directly from the classifier response, corresponding to the true positive rate and false positive rate discussed in the previous section. Increasing the threshold leads to stronger background rejection at the cost of reduced signal efficiency, while lowering the threshold increases signal acceptance but admits a larger fraction of background events. The optimal working point is therefore determined by balancing these effects.

The signal acceptance is defined as:

$$\epsilon_{sig}^{BDT}(th) = \frac{\text{True Positive}}{\text{True Positive} + \text{False Negative}}, \quad (6.12)$$

Source	Fiducial Volume [%]	BDT [%]	Total rate [events/yr]
Materials	0.14	39	(2038 ± 55)
LXe ^{214}Bi	9.7	69	(41.0 ± 0.1)
LXe shell ^{214}Bi	0.41	55	(3.7 ± 0.009)
Signal	11	79	(10 ± 4)

Table 6.2: Signal and background acceptances after applying the FV and the BDT cut selections. Percentages are reported per step with respect to the initial dataset. The FV acceptance for material backgrounds is computed as a weighted average over all components. The total event rate is reported together with its associated systematic uncertainty. For the material backgrounds, this uncertainty arises from the screening uncertainties reported in Ref. [175, 247], while for the signal it is dominated by the uncertainty in the predicted half-life [116].

while the background acceptance is given by:

$$\epsilon_{bkg}^{BDT}(th) = \frac{\text{False Positive}}{\text{False Positive} + \text{True Negative}}. \quad (6.13)$$

The expected background rate is computed from Equation 6.3 and scaled by $\epsilon_{bkg}^{BDT}(th)$, while the signal rate is computed from Equation 1.15 and scaled by $\epsilon_{sig}^{BDT}(th)$.

The Asimov estimate of the significance as a function of the classifier threshold is shown in Figure 6.18. The maximum significance is obtained at a threshold of 0.692. At this operating point, the signal acceptance within the fiducial volume is 79%, while the material background acceptance is 39%. This corresponds to a substantial suppression of the background while retaining a large fraction of the signal.

6.5.3 Signal and background estimates

The signal and background templates are normalized to the expected number of events per year after selection in a fiducial volume containing 0.905 t of LXe. Therefore, a livetime of 1 year corresponds to an exposure of $0.905 \text{ t} \times \text{yr}$.

As shown in Table 6.2, the combined effect of the fiducial volume and BDT selections significantly reduces the background contributions, in particular for external gamma rays from detector materials. Compared to the preselection stage, this corresponds to a reduction of several orders of magnitude in the dominant background component.

Despite this suppression, material induced backgrounds remain the dominant contribution, yielding ~ 2038 events per year, while the expected signal rate is ~ 10 events per year. This corresponds to a signal-to-background ratio of $\mathcal{O}(10^{-3})$, indicating that the analysis remains background limited. Intrinsic backgrounds from ^{214}Bi are subdominant




Figure 6.19: Final background model in the reconstructed energy space, showing the contribution of each background component and the signal model after all selection cuts.

after all selections, contributing at the level of a few tens of events per year.

The resulting signal and background templates, shown in Figure 6.19, define the final model in reconstructed energy space, which is used in the statistical analysis described in the following section.

6.6 Inference

The expected sensitivity of the $2\nu\text{EC}\beta^+$ decay of ^{124}Xe is evaluated using a binned Poisson log-likelihood in reconstructed energy space, using a 10 keV binning. This choice ensures sufficient counts per bin across the ROI. For this study, the `alea` statistical inference framework [2] was used. Following the notation of Ref. [248], the number of observed events in bin i is denoted by n_i , while the expected number of events is given by $\lambda_i(\mu, \boldsymbol{\theta})$, where μ is the signal strength parameter and $\boldsymbol{\theta}$ represents nuisance parameters associated with the background normalizations.

The likelihood function is then defined as the product of Poisson probabilities over all bins:

$$\mathcal{L}(\mu, \boldsymbol{\theta}) = \prod_i \text{Pois}(n_i | \lambda_i(\mu, \boldsymbol{\theta})). \quad (6.14)$$

In this work, no nuisance parameter constraints are included, and the inference there-

fore corresponds to a statistics only sensitivity estimate. Following the formalism of Ref. [248], the statistical inference is based on the profile likelihood ratio:

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})}, \quad (6.15)$$

where $\hat{\mu}$ and $\hat{\boldsymbol{\theta}}$ are the values that globally maximize the likelihood, and $\hat{\boldsymbol{\theta}}(\mu)$ denotes the values of the nuisance parameters that maximize the likelihood for the fixed μ . The test statistic used to set upper limits is then defined as:

$$q_\mu = -2 \ln \lambda(\mu), \quad (6.16)$$

which applies in the regime $\hat{\mu} \leq \mu$, relevant for upper-limit calculations [248]. The upper limit μ_{up} at a given confidence level is obtained from the condition $q_\mu = q_{\text{crit}}$, where $q_{\text{crit}} = 2.71$ for a 90% confidence level.

In this work, the expected sensitivity is evaluated by simulating 5000 pseudo-experiments generated under the background-only hypothesis ($\mu = 0$). For each pseudo-experiment, the upper limit μ_{up} is determined from a profile likelihood scan. The distribution of μ_{up} values is then used to extract the expected median limit, as well as the corresponding one and two standard deviation bands.

The limit on the signal strength is translated into a lower limit on the half-life of the process, using the inverse proportionality between the decay rate and the half-life as described in Equation 1.15.

6.6.1 Projected sensitivity

The resulting sensitivity, shown in Figure 6.20, is evaluated for three different exposure scenarios, corresponding to the accumulated livetime of the three science runs of XENONnT. The first data point corresponds to SR0 conditions, for which the detector response and background model were characterized in this work.

Considering the SR0 livetime of 97 days and a fiducial mass of 0.905 t, corresponding to an exposure of $0.24 \text{ t} \cdot \text{yr}$, the median expected lower limit on the half-life is:

$$T_{1/2} > 5.35 \times 10^{20} \text{ yr at 90\% C.L.} \quad (6.17)$$

Two additional data points are included, corresponding to the cumulative exposures




Figure 6.20: Expected median lower limit on the $T_{1/2}$ of the ^{124}Xe decay via $2\nu\text{EC}\beta^+$ as a function of exposure, assuming a fiducial mass of 0.905 t. The first point corresponds to SR0 only, while the second and third points represent the cumulative exposures of SR0+SR1 and SR0+SR1+SR2, respectively.

Science Run	Livetime [yr]	Exposure [t·yr]	$T_{1/2}$ lower limit [yr]
SR0	0.266	0.241	5.35×10^{20}
SR0 + SR1	0.956	0.865	1.0×10^{21}
SR0 + SR1 + SR2	2.408	2.179	1.6×10^{21}

Table 6.3: Expected median $T_{1/2}$ lower limits of the $2\nu\text{EC}\beta^+$ decay of ^{124}Xe for different exposure scenarios, corresponding to the livetime of three science runs of XENONnT and considering a fiducial mass of 0.905 t.

obtained by combining SR0 with SR1, and SR0 with SR1 and SR2, respectively. These points provide an estimate of the sensitivity achievable with the currently available data of XENONnT, and are summarized in Table 6.3. It should be noted, however, that these projections are obtained by scaling the exposure while keeping the detector response and background model fixed to those of SR0. A proper evaluation would require dedicated simulations for SR1 and SR2, accounting for possible variations in detector performance, reconstruction, and background conditions.

Potential sources of systematic uncertainty include the signal and background selection efficiencies as well as the background normalization. The treatment of systematic uncertainties would require the introduction of additional nuisance parameters and constraints in the likelihood, which would generally lead to a degradation of the expected sensitivity. This is, however, beyond the scope of the present study. In addition, finite simulation statistics may introduce uncertainties in the templates, although this effect is expected to be subdominant given the large number of simulated events available for




Figure 6.21: Profile likelihood scan as a function of the signal strength μ for a representative background-only dataset corresponding to the SR0 exposure. The test statistic $q_\mu = -2\Delta \ln \mathcal{L}$ is shown relative to its minimum at $\hat{\mu} = 0$, indicating consistency with the background-only hypothesis. The horizontal dashed lines correspond to the 90% and 95% confidence level thresholds.

each background source.

An example of the profile likelihood scan as a function of the signal strength μ is shown in Fig. 6.21. The best-fit value is found at $\hat{\mu} = 0$, indicating that the data are consistent with the background-only hypothesis. The 90 % confidence level upper limit is obtained from the intersection with the threshold q_{crit} .

6.7 Summary and Outlook

In this chapter, we presented the first sensitivity projection for the search for the $2\nu\text{EC}\beta^+$ decay of ^{124}Xe to the ground state of ^{124}Te in XENONnT. Building on the detector response validation of Chapter 5, in this study dedicated full-chain signal and background simulations were performed, and the analysis strategy was constructed around three main ingredients: physics motivated preselections, a simulation-driven FV optimization, and a ML discriminator.

The study shows that the search is strongly background limited. After all selections, the dominant background remains the radiogenic contribution from detector materials, which exceeds the expected signal rate by several orders of magnitude in the final selected sample. As a result, the projected sensitivity remains well below the half-life range predicted for this decay, indicating that an observation is not yet achievable with the present analysis configuration and currently available exposure.

At the same time, this work represents an important methodological step toward a fu-

ture search with data. It establishes a high-energy analysis framework for XENONnT capable of studying multi-scatter MeV scale events beyond standard low-energy analyses, and provides the first full-chain material background simulation campaign processed within the `fuse` framework. In addition, it demonstrates that topology discrimination offers an improvement over traditional cut-based selections.

Several improvement opportunities identified in this work point directly to the next steps. A first limitation is that the inference is based on a one-dimensional fit in reconstructed energy. This choice represents a controlled baseline for a first sensitivity study, but it cannot fully exploit the information carried by the event topology. Extending the inference to a multivariate analysis is, however, non-trivial, as it would require sufficient simulation statistics and a robust modeling of the signal and background distributions across a wide energy range for reliable multidimensional template building.

A second area with improvement opportunity concerns the treatment of detector and reconstruction effects specific to multi-scatter events. Phenomena such as limited S2 separation, photoionization, afterpulses, and delayed-electron backgrounds can distort the reconstructed topology and reduce the discriminating power of the observables used in the classifier. A more systematic characterization of multi-scatter reconstruction efficiencies in data, together with improved modeling and mitigation of instrumental high-energy effects, will be important for future analyzes.

A third area with an opportunity for improvement is that the projected sensitivity is derived within a statistically only framework, and a future search with data will require a full treatment of systematic uncertainties. The dominant sources of are expected to arise from the normalization of the material background model, driven by the material screening uncertainties, and from the total selection efficiency. The latter includes contributions from the preselections, fiducial volume, and BDT acceptances, as well as potential detector mismodeling of topology observables. In this context, the residual energy bias uncertainty derived in Chapter 5 is small compared with the 10 keV binning used in the inference.

The projections beyond SR0 were obtained by scaling the SR0 model in exposure while keeping the detector response and background conditions fixed. A realistic search using all available data will require dedicated simulations for each science run. In this context, potential detector changes, such as variations in the electric field, may also affect the reconstruction of multi-site topologies. For example, a stronger field could improve the separation of nearby S2 signals and therefore improve some of the topology observables used here, although the net impact on signal-to-background discrimination would need to be quantified with dedicated simulations because similar effects may also modify the

background response.

In summary, this work does not yet demonstrate discovery potential for the $2\nu\text{EC}\beta^+$ decay in XENONnT, but it does establish the essential analysis framework for such a search. The study identifies the dominant limitations clearly and introduces several tools that will be central for future high-energy rare-event analyses in XENONnT. In that sense, the main outcome of this chapter is both a first benchmark sensitivity and a substantial step toward a data search for $2\nu\text{EC}\beta^+$ and other rare ^{124}Xe decay channels.

Chapter 7

Xenoscope: an R&D Platform for XLZD

The XENON-LUX-ZEPLIN-DARWIN (XLZD) project [57] arises as a joint initiative between the established XENON [142] and LZ [143] collaborations, and the DARWIN R&D collaboration [256]. The XLZD detector will be the ultimate dual-phase xenon TPC for direct dark matter detection, while also serving as a multipurpose observatory for other rare event searches [57]. To realize XLZD, the active LXe target will be scaled to several tens of tons, far exceeding the target masses of current-generation TPCs [142, 143, 257]. Building a detector of this scale presents major technological challenges including long electron-drift lengths, ultra-high xenon purity, high-voltage (HV) delivery and efficient light detection.

This chapter focuses on the R&D activities performed within the Xenoscope project [258], a vertical LXe TPC demonstrator built at the University of Zurich to address the electron drift challenge over meter-scale distances. The XLZD baseline detector concept and physics case are presented in 7.1. The Xenoscope facility and its supporting infrastructure are introduced in 7.2. The results of the purity monitor phase are summarized in 7.3. The upgrades performed to commission Xenoscope as a dual-phase TPC are discussed in 7.4, including the implementation of optical readout, liquid-level control, and high-voltage delivery. Finally, the first results from the dual-phase TPC commissioning run of Xenoscope are presented in 7.5.

7.1 The XLZD project

The main physics channel explored in dual-phase xenon TPCs is spin-independent WIMP–nucleus scattering, for which LZ, XENONnT, and PandaX-4T have reported




Figure 7.1: CAD rendering of the XLZD baseline design. The double walled stainless-steel cryostat (inner and outer vessels) encloses the cylindrical $3\text{ m} \times 3\text{ m}$ dual-phase TPC, whose main components are indicated: the photosensor arrays at the top and bottom of the active volume, the field cage, and the PTFE reflector panels. Figura taken from [57]

the most stringent exclusion limits to date [51, 55, 56]. Achieving sensitivity across a broader region of the WIMP mass-cross-section parameter space requires detectors with substantially larger target masses and stable operation over long periods of time. XLZD is envisioned to meet these requirements, reaching projected sensitivities down to spin-independent WIMP-nucleon cross-sections of order $1 \times 10^{-49} \text{ cm}^2$ (see Figure 1.5), about an order of magnitude beyond the current state of the art.

Like its predecessors, the physics reach of XLZD extends well beyond WIMP–nucleus interactions. The rich isotopic composition of xenon (summarized in Table 2.1) enables the investigation of $2\nu/0\nu\beta\beta$ of ^{136}Xe [194], together with double-weak decays of ^{124}Xe [116]. For the latter, the $2\nu\text{ECEC}$ decay mode, observed for the first time in the XENON1T detector [115], has already become a relevant background for WIMP searches in current generation detectors [55, 56]. In addition, as discussed in Chapter 6, the $2\nu\text{EC}\beta^+$ mode is already experimentally accessible for dual-phase xenon TPCs. XLZD will enable more precise half-life measurements for these processes and improve sensitivity to the corresponding neutrinoless modes. Furthermore, its large target mass and ultra-low background conditions will enable sensitivity to astrophysical neutrinos, including solar neutrinos from the ^8B cycle through $\text{CE}\nu\text{NS}$ [89, 90], pp -neutrinos [195] and neutrino bursts from galactic supernovae, via the inter-experiment Supernova Early Warning System (SNEWS2.0) network [196, 259].

The XLZD detector is envisioned as a cylindrical dual-phase xenon TPC (see Figure 7.1) with an active volume of approximately $3\text{ m} \times 3\text{ m}$, corresponding to a fiducial target mass of 60 t [57, 260]. The TPC will be instrumented with 2364 low-background photosensors arranged in two arrays to detect the prompt S1 and delayed proportional S2 scintillation signals characteristic of dual-phase TPCs (see Section 2.2.2). In the nominal design, 3 inch diameter Hamamatsu PMTs are foreseen, however, extensive characterization campaigns have also been conducted to evaluate alternative photosensor technologies, including 2 inch Hamamatsu PMTs [202] and silicon photomultipliers (SiPMs) [261]. The electric drift field must be between 240 V/cm and 290 V/cm in order to achieve optimal NR and ER discrimination [57]. The TPC will be housed in a low-background, double walled cryostat and will be surrounded by active neutron and muon veto systems.

Scaling the dual-phase LXe TPC technology to such dimensions introduces significant challenges in background control, signal readout, xenon handling, purity maintenance among others. In addition, establishing a uniform drift field across a 3 m distance requires optically transparent electrodes (see Section 2.2.2) and constitutes a major engineering challenge. At this scale, mechanical robustness must be ensured to prevent sagging while preserving high optical transparency. Drifting electrons over several meters imposes demanding design goals on the xenon purification and high-voltage systems, which must deliver voltages up to $\mathcal{O}(100\text{ kV})$ reliably. High drift fields are not only essential for optimal NR/ER discrimination but also help mitigate background sources related to longer drift times at lower fields, such as the temporal overlap of S1 and S2 signals from independent interactions within a single acquisition window and randomly paired isolated signals that do not originate from a physical interaction.

To address these issues, two large scale demonstrators have been developed: Xenoscope [258] at the University of Zurich and PANCAKE [262] at the University of Freiburg. Both demonstrators were originally conceived within the DARWIN [256] program, whose baseline design featured a dual-phase TPC with a height and diameter of 2.6 m. Consequently, the demonstrators were built to DARWIN dimensions rather than to the larger baseline foreseen for XLZD, and serve as testing platforms for validating the technologies required for the next-generation detector.

7.2 The Xenoscope facility

Current dual-phase LXe TPCs operate with drift lengths of up to about 1.5 m [142, 143]. Xenoscope, the XLZD vertical demonstrator, aims to characterize electron transport over the 2.6 m drift distance foreseen for XLZD. Its primary physics goal is to demon-




Figure 7.2: CAD rendering of the Xenoscope facility at UZH. The figure highlights the main subsystems of the demonstrator, as well as the supporting infrastructure such as the xenon gas storage, recovery, and recirculation systems. The 2.6 m dual-phase TPC is shown enclosed within the double-walled cryostat, together with main elements of the cooling system, including the pre-cooler, heat exchanger, and cooling tower. Figure taken from [258].

strate electron drift under standard operating conditions, as well as to enable systematic studies of electron diffusion in both the longitudinal and transverse directions. In addition, Xenoscope provides an opportunity to investigate xenon optical properties, such as Rayleigh scattering and absorption lengths, which are essential for the design and optimization of next-generation dual-phase TPCs. A render of the Xenoscope facility is shown in Figure 7.2, with the main subsystems highlighted.

Xenoscope was commissioned in two successive phases (see Figure 7.4). The first phase, operated in spring 2022, consisted of a single-phase LXe, 0.53 m-tall purity monitor designed to estimate the concentration of electronegative impurities over the drift length [263]. In this configuration, signal readout was based on charge collection, as will be described in Section 7.3. Following the operation of the purity monitor phase, several subsystems were added for the full-scale 2.6 m dual-phase TPC. In particular, the signal readout was redesigned to transition from charge to light detection, and both the liquid-level control system and the high-voltage delivery system were newly implemented.

Xenoscope operates a xenon inventory of approximately 350 kg, stored in aluminum bottles placed adjacent to the detector frame. During operation, the LXe is contained within a double-walled, vacuum-insulated stainless-steel cryostat. The cryostat is sus-

pendent from a precision leveling system [264] mounted on the upper platform of the support structure. The latter is constructed from 50 mm $\times$ 50 mm extruded aluminum profiles and comprises three main sections: the inner frame, the outer frame, and the staircase assembly. With an overall cubic footprint of 4 $\times$ 4 $\times$ 4 m³, the outer frame is enclosed with acrylic panels to minimize dust accumulation during the assembly of internal components. Next to the power distribution cabinet, the staircase section and the area behind it are secured with interlocks. Behind the staircase are the HV power supply, the DAQ system, the slow-control server, and an instrumentation rack. The interlock system ensures that, during operation, the perimeter containing these subsystems remains fully protected.

The outer cryostat vessel is designed in a modular configuration to facilitate assembly of the full height setup. It consists of six ISO-K-400 sections, the uppermost section features a conical shape to accommodate the routing of cables and other services. The outer vessel is anchored to the inner frame via four support arms to prevent mechanical motion during operation. The inner cryostat vessel follows a similar modular design, comprising six electropolished stainless-steel tube sections terminated with DN250CF flanges, with the uppermost section sharing the same conical geometry as the outer vessel.

Both the inner and outer vessels are fitted together in a flange assembly composed of an outer ISO-K-500 flange and an inner DN350CF flange. The two flanges are connected by six concentrically oriented, stress-relieving swivel rods of 20 mm diameter, suspended by 10 mm dowel pins. The top flange assembly accommodates six access ports, five smaller ones used to route cabling and auxiliary systems from the LXe side to the air side, and a sixth, larger port that houses the cooling tower. All ports are connected by axial-displacement bellows to allow for differential thermal contraction and mechanical flexibility.

Xenoscope employs a custom-designed leveling system developed to minimize tilting of the liquid–gas interface. It consists of a 40 mm-thick aluminum–magnesium alloy triangular plate supported symmetrically by three adjustable leveling legs [264]. The plate and legs are mechanically decoupled by polyurethane vibration-damping pads, which reduce the transmission of mechanical noise from the surrounding instrumentation to the detector. The top flange assembly is attached to the underside of the triangular plate, providing stable support and alignment of the cryostat during operation [264].




Figure 7.3: **Left:** CAD rendering of the two heat exchangers installed inside the vacuum vessel. **Right:** CAD rendering of the cooling tower and the top flange assembly. During recirculation, LXe is extracted from the top of the inner vessel and directed through the heat exchanger system, where it evaporates while transferring its cold to the incoming purified gas. The xenon vapor then passes through the vaporizer and post-heater before entering the purification system. The color gradient illustrates the vaporization process of the outgoing LXe. After purification, the xenon gas is pre-cooled in the smaller heat exchanger and fully condensed in the larger one. The condensed xenon then flows through the cooling tower and is reintroduced at the bottom of the inner vessel. Figures taken from [258].

7.2.1 Cryogenics and xenon recirculation

The goals of the cryogenics and recirculation systems are to maintain the xenon at stable thermodynamic conditions and achieve high purity. Before filling the detector with LXe, the inner vessel undergoes an initial cooling stage to shorten the overall cooldown time. A liquid-nitrogen-based precooling system (see Figure 7.2) is installed in the upper cylindrical section and serves as an auxiliary system during the cooldown and filling stages, to help bring the temperature from the room level to values close to the xenon liquefaction point. The pre-cooler consists of four stainless-steel segments connected in series and mounted around the top section, in direct contact with the inner vessel wall. Their curved geometry matches the contour of the upper region, ensuring efficient and uniform thermal contact during precooling. Liquid nitrogen (LN_2) is supplied from an external storage Dewar via two flexible metal hoses connected to a liquid feedthrough at the bottom of the outer cryostat, enabling efficient cooling of the vessel. The addition of the pre-cooler reduces the filling time by a factor of ~ 4 [261].

After the inner vessel is pre-cooled, the standard recirculation procedure is initiated. Recirculation is the process by which LXe is extracted from the top of the inner cryostat

vessel, evaporated, purified and subsequently recondensed and returned to the inner vessel in liquid form. The purification is performed through a hot zirconium getter manufactured by SAES Pure Gas [265], which is designed to remove electronegative impurities such as O_2 , H_2O , N_2 or CO . Two heat exchangers (see Figure 7.3) manufactured by Kelvion Holding GmbH [266] enable efficient thermal coupling between the outgoing cold LXe and the incoming warm purified GXe. The purified GXe leaving the getter transfers its heat to evaporate the returning LXe, while the latter cools the GXe before it reaches the cooling tower. This process efficiently recovers the latent heat associated with the liquid-gas phase change, thereby minimizing the external cooling power required.

The cooling power required to liquefy and maintain xenon in its liquid phase is provided by an Iwatani Cooperation PTR [176], model PC-150, rated to ~ 200 W at 165 K. The PTR cools an oxygen-free copper cold head, which serves as the xenon condenser. Here, the cooled GXe is fully liquefied upon contact with the cold surface and the resulting LXe is fed through a hose back into the bottom of the inner cryostat vessel, which assures that all the xenon in the inner vessel is circulated and purified. Both the heat exchangers and the cold head are instrumented with several PT100 Resistance Temperature Detectors (RTD) to monitor the temperature in the different stages. In addition, a 180 W heat band driven by a Cryo-con Model 24C Proportional-Integral-Derivative (PID) controller [267], provides active stabilization of the temperature and pressure of the xenon in the cold-head. The heat band is mounted in a cylindrical copper cup located between the cold head and the PTR, as shown in Figure 7.3.

To maintain the cryogenic conditions inside the inner vessel and minimize heat losses, thermal insulation is implemented using two complementary methods: a vacuum gap between the inner and outer vessels, which suppresses convective and conductive heat transfer, and alternating layers of highly reflective foil or multi-layer insulation (MLI) wrapped around the cryostat and cryogenic system to limit radiative heat exchange. The vacuum is achieved by means of two HiPace 80 turbo molecular pumps from Pfeiffer Vacuum GmbH [268] deployed near the heat exchanger and at the ground floor of the facility, near the bottom of the cryostat. The HiPace pumps are backed up by a common HiScroll 12 roughing pump from Pfeiffer Vacuum GmbH [268].

7.2.2 Xenon recuperation and storage

When not in use, the GXe is stored in nine aluminum bottles arranged in two rows within the storage-cylinder array located next to the outer frame. The storage array has sufficient volumetric capacity to accommodate the entire xenon inventory, which would reach a pressure of 58.4 bar at room temperature, placing it in the supercritical

regime.

Each bottle has a capacity of 40 L, and both the bottles and their valves are rated for pressures up to 200 bar. Four of these bottles are housed inside Dewar flasks for cooling and are interconnected to form the GXe recuperation system based on cryopumping. The latter is the process of transferring xenon by creating a pressure difference between two connected regions of the system. In Xenoscope, the xenon is transferred from the inner vessel into the storage cylinders by creating a low pressure region in the cold bottles, allowing xenon to flow towards them. To enable this, the Dewar flasks are cooled using an automated LN₂ filling system regulated by PT100 temperature sensors installed at different heights within each Dewar. Each row of bottles is suspended from a weighing system with a 250 kg capacity, enabling continuous monitoring of the xenon mass during operation. Additionally, four 1400 W band heaters are mounted around the upper section of each bottle to prevent ice formation on the valves, which could otherwise lead to xenon leaks.

Cryopumping begins by cooling one to four bottles, depending on the amount of xenon to be recovered, to LN₂ temperature, creating the required low-pressure region. The resulting pressure difference between the cold bottles and the inner vessel drives the xenon flow toward the bottles. As xenon condenses and freezes within the bottles, the pressure drops further, enhancing the cryopumping effect. Meanwhile, the evacuation of the inner vessel causes its pressure to decrease, reducing the pressure difference between the vessel and the cold bottles. Consequently, the cryopumping rate slows down asymptotically as the system approaches equilibrium. When only residual sub-mbar amounts of xenon remain in the inner vessel, the final recovery stage is performed using a parallel buffer volume located in the main gas panel. The recuperation process is considered complete once all xenon from the inner vessel has evaporated and been cryopumped into the cold bottles. After recuperation, the xenon inventory is equalized while warming up among all bottles for long-term storage.

The recovery of 350 kg of xenon typically consumes four 160 L LN₂ Dewars over the course of one week, corresponding to an average consumption rate of approximately 72 kg/d. The duration of the recovery process is determined by the evaporation rate of LXe inside the inner vessel as well as by the gas flow rate. The experience gathered during the facility commissioning in early 2021, when the inner vessel was first filled with (80.66 ± 0.10) kg of xenon, showed that recovery through cryopumping required about 2 d at an average flow of 5 slpm, corresponding to roughly 1.75 kg/h. Extrapolating this rate to the full xenon inventory yields total recovery times of approximately 200 h, which translates into several weeks of operation.

Because Xenoscope was conceived as a testing facility, frequent xenon filling and recovery cycles are unavoidable. To address the long recovery times and reduce LN₂ consumption, a liquid recovery system was implemented. The system is based on a custom-built spherical pressure vessel, Box of Xenon (BoX), into which LXe is drained by gravity through a vacuum-insulated cryogenic line connected to the bottom outlet of the cryostat. BoX is rated for pressures up to 90 bar and temperatures down to 77 K. It serves as a storage vessel for xenon in both the liquid and gas phases, with a total capacity of approximately 350 kg at room temperature, corresponding to a pressure of approximately 63 bar. A warm outlet line connects BoX to the high-pressure line, allowing the xenon to return in the gaseous phase to the storage array once it has reached ambient temperature. For LXe recovery, BoX is initially cooled down by flushing LN₂ through a copper block located beneath it. During LXe recovery, the pressure in BoX rises to about 3 bar due to the hydrostatic pressure of the LXe column. The warm outlet line must be opened during this process to allow GXe to flow back into the cryostat, maintaining pressure equilibrium and ensuring a continuous flow of LXe into BoX. The LN₂ consumption during liquid recovery is approximately 3.6 kg/h. At the end of the liquid recovery process the remaining GXe can be cryopumped into the storage array or into BoX by freezing the xenon. As the recuperated xenon warms isochorically inside BoX, its temperature and pressure rise beyond the critical point, causing it to transition into the supercritical state.

The main gas panel contains the fittings and valves required to route GXe between the storage bottles, the cryogenic system, and the purification loop. The gas system includes 0.5 inch pipes, manual valves rated for up to 68.9 bar and a high-pressure line equipped with valves rated for up to 344 bar. Several components of the purification system are integrated into the main gas panel, including the hot getter and a double-diaphragm xenon compressor manufactured by KNF Neuberger [269], capable of delivering outlet pressures up to 3 bar and flow rates up to 250 slpm in air. Xenon flow is regulated by a Teledyne Hastings [270] mass flow controller and is continuously monitored with an HFC-201 flow meter. A detailed Piping and Instrumentation Diagram (P&ID) of the Xenoscope facility is shown in Figure C.1 of Appendix A.

7.2.3 Facility monitoring

To guarantee safety and provide data for robust analysis, reliable monitoring is essential. The Xenoscope facility is monitored and controlled in real time through a modular, automated slow control system based on a microservice architecture [271].

A Programmable Logic Controller (PLC) and a Raspberry Pi collect data from sensors and control hardware components such as valves, pumps, and flow controllers. Com-




Figure 7.4: CAD renders of the two Xenoscope phases. **Left:** Charge-readout purity-monitor setup, showing the fiber, photocathode and the associated components. **Right:** Dual-phase 2.6m TPC with optical readout, highlighting the top-stack electrode system and the SiPM array. Figures taken from [258]

munication with the rest of the system is achieved through standard network protocols (OPC and HTTP).

The data collected are continuously stored in Prometheus, a time-series database optimized for monitoring and alarming. The data are visualized in Grafana, where users can access live dashboards, define alarm thresholds, and identify unusual behaviour. Users can send commands and adjust settings directly from their browser interface by means of a proxy service and a web-based front-end framework which allow bidirectional communication. The alarm rules configured in Grafana trigger notifications via Prometheus, which can alert users by email or SMS in the case of anomalies. Finally, all components, except the microcontroller software, run inside isolated Linux containers managed by Kubernetes, which automatically handles deployment and scaling and ensures that all services are always running reliably.

7.3 Purity monitor

As mentioned in Section 7.2, the proof of concept for electron drift in Xenoscope was planned in two stages. The first stage was designed to benchmark the purification system and demonstrate the achievable electron lifetimes under stable operating conditions.

The purity monitor operates by measuring the attenuation of a known population of electrons, $N(t_0)$, after drifting for a time t_d in LXe under an electric field. As discussed




Figure 7.5: Electron drift lifetime measured in Xenoscope. Three xenon recirculation flow rates were used, as indicated by the purple, blue and green markers. Dashed vertical lines denote changes in recirculation flow, while dash-dotted lines indicate irregularities in pressure or flow conditions. The red line shows the best-fit model, and the bottom panel displays the corresponding residuals. Figure taken from [263].

in Section 2.2.2, the xenon purity is quantified by the electron drift lifetime, τ (see Equation 2.11) with τ defined as:

$$\tau = \frac{1}{\sum_i k_i n_i}, \quad (7.1)$$

where k_i is the attachment rate for each impurity type and n_i its concentration. The survival probability of drifting electrons therefore depends directly on the concentration of electronegative impurities in the LXe. Commercial xenon typically contains impurities at the ppm level, mainly O_2 , H_2O , N_2 , CO , and traces of organic compounds originating from detector materials and their outgassing.

Both phases share the same optical components and electron production principles. For a detailed description of these systems, see Section 7.4.2. Here, we briefly describe the charge collection principle of the PM.

Electrons are generated via the photoelectric effect by illuminating a gold-coated photocathode with VUV photons from a xenon discharge lamp. The light is delivered through an optical fiber running along the height of the TPC down to the photocathode in the cathode plate. They drift through the extraction, drift, and collection regions, defined by four vertically arranged parallel electrodes: the cathode plate, cathode grid, anode grid, and anode plate. The cathode and anode are solid stainless-steel disks, while the grids consist of hexagonally etched stainless-steel meshes. The field cage extends




Figure 7.6: Left: Drift velocity measured with the Xenoscope purity monitor, varying the electric field from 25–75 V/cm in 5 V/cm increments. The results are compared to literature values from Gushchin (165 K) [272], EXO-200 (167 K) [237], Xurich II (2018, 184 K) [273], Xurich II (2021, 177 K) [264], HeXe (174 K) [212], XeBRA (173 K) [274], NEST v2.3.8 [148]. **Right:** Longitudinal diffusion coefficient measured in Xenoscope for an electron lifetime of $\tau = (649 \pm 23) \mu\text{s}$, compared with results from the purity monitor of Njoya et al. ($\tau \sim 35 \mu\text{s}$) [275], the TPC of Hogenbirk ($\tau \sim 430 \mu\text{s}$) [276], and NEST [148]. The NEST v2.3.7 model (solid grey curve) predicted vanishing diffusion at low drift fields, a behavior corrected in v2.3.8 (solid blue curve). Figures taken from [263].

from the cathode grid to the anode grid, providing a total drift length of 0.53 m. It is composed of 34 oxygen-free high-conductivity copper rings with an inner diameter of 15 cm, held together by interlocking PTFE spacers and connected through a parallel resistor chain of $5 \text{ G}\Omega$ each.

The motion of charge carriers in the extraction region induces a current on the anode. The electrons then drift through the uniform electric field between the cathode and anode grids, where some may attach to residual impurities in the LXe. Upon reaching the anode grid and entering the collection region, they induce a signal that grows until all surviving electrons are collected. This current induced at the anode is converted into a voltage signal by a transimpedance amplifier, then digitized and integrated to determine the charge, i.e., the number of extracted and surviving electrons.

The purity monitor run was performed with the inner vessel filled with 343 kg of xenon at $(177.6 \pm 0.1) \text{ K}$ and $(2.050 \pm 0.001) \text{ bar}$, with the purity monitor fully immersed in the liquid. During data taking, the cathode was biased at -2710 V , the cathode grid at -2650 V , the anode grid was grounded, and the anode held at 500 V . These settings defined three field regions: an extraction gap of 18 mm with a field of about 33 V/cm, a main drift region of 503 mm at 53 V/cm, and a collection gap of 10 mm with a field of approximately 500 V/cm. The full run spanned 89 d, from March to June 2021, with the results of published in Ref. [263].

Two measurement campaigns were then performed. The first, focused on determining

the electron lifetime at three xenon recirculation flow rates (30, 35, and 40 slpm) at constant electric field. The time evolution of the measured lifetime, together with the corresponding fit model, is shown in Figure 7.5, from which an electron lifetime of $(660 \pm 60) \mu\text{s}$ was obtained. This value is comparable to that achieved in the XENON1T experiment [156].

The second measurement campaign investigated electron transport properties by varying the drift field between 25–75 V/cm in steps of 5 V/cm. After the electron lifetime reached a stable plateau, data for drift-field-dependent studies were collected on day 89 of the run. The extraction field was adjusted to preserve the same extraction-to-drift field ratio used in the lifetime measurements, while the collection field was kept constant at 500 V/cm. The measured electron drift velocity is shown in Figure 7.6, together with results from other LXe setups operating at different field strengths. The measurements are consistent with previous results in LXe and with the NEST prediction, which is derived from data-driven empirical models.

The longitudinal diffusion coefficient (see Section 2.1.3) of the electron cloud was also derived from the drift-velocity dataset. It was obtained from the width of the charge signal at the anode, determined by a Gaussian fit. The resulting values are shown in Figure 7.6 and are consistent with previous measurements in LXe. The NEST version 2.3.7 does not include values at low drift fields (below 100 V/cm) and yields longitudinal diffusion coefficients smaller than those obtained in our purity monitor measurements and in other low-field measurements, as illustrated in Figure 7.6. Following discussions with NEST developers, our results were incorporated into an updated version, NEST v2.3.8, whose prediction, shown in Figure 7.6 as the solid blue line, provides an improved description of the longitudinal diffusion observed at low drift fields.

7.4 Dual-phase time projection chamber

Following the successful operation of the purity monitor phase, Xenoscope was upgraded to its full-height dual-phase configuration. This section describes the principal upgrades performed to commission Xenoscope as a dual-phase TPC, including the installation of the field cage and electrode stack, the photon detection system, the liquid-level monitoring and control system, and the high-voltage delivery system.

The Xenoscope field cage defines the uniform drift field within the active volume and ensures stable electron transport from the photocathode to the extraction region. It consists of 173 oxygen-free high conductivity copper field-shaping rings arranged in five modular sections, each approximately 50 cm in length. The rings are connected through

two parallel resistor chains, each composed of $1\text{ G}\Omega$ resistors, providing electrical redundancy and ensuring uniform potential grading along the drift axis. The mechanical support is provided by six polyamide imide (PAI) pillars coupled with PTFE holders, chosen for their high dielectric strength, mechanical stability, and low outgassing rates compared to other polymers.

Prior to assembly, all copper components were chemically cleaned in the local laboratories and subsequently electropolished to reduce microscopic surface roughness and suppress potential sites for field emission or contamination. The cleaning procedure included a gentle mechanical pre-cleaning, immersion in high-purity acetic acid to remove surface oxides and corrosion, ethanol rinsing followed by rapid drying to avoid re-oxidation, and ultrasonic cleaning using a 10 % solution of acidic ELMA detergent. Finally, the components were briefly passivated in dilute ($<1\%$) citric acid prior to storage to prevent oxidation. The lowest field-shaping ring is mechanically and electrically coupled to the stainless-steel cathode plate, which houses the photocathode described in Section 7.4.2. At the upper end, two additional electrodes, the gate and the anode (see Figure 7.4), complete the configuration required to establish the drift and extraction fields. These electrodes consist of fine hexagonal meshes framed by stainless steel rings and are separated by 12 mm to establish the extraction and electroluminescence regions.

During inspection, a slight misalignment between the gate and anode meshes was identified, resulting from a manufacturing imperfection. This geometric offset limits the maximum achievable bias between the two electrodes, as higher voltages would distort the extraction field uniformity. For this reason, once assembled in the inner vessel, the anode and gate were tested up to 5 kV and -1 kV , respectively, corresponding to a potential difference of 6 kV. No discharge activity was observed during the test, although a small increase in the anode current was detected as the voltage was raised. This current remained well below the power-supply limit of $300\text{ }\mu\text{A}$ and therefore does not represent a risk for operating the top stack at these voltages.

In the 2.6 m phase, photon detection is carried out with a custom-designed SiPM array using VUV-sensitive MPPCs from Hamamatsu (model S13371-6050CQ-02). The array consists of 192 $6 \times 6\text{ mm}^2$ sensors, grouped in sets of four (2×2) to form quads (see Figure 7.11 right). Four quads make up one tile, which corresponds to a single PCB readout channel, whose amplified output constitutes the analog summed signal. In total, 12 tiles cover an effective area of about $24 \times 24\text{ mm}^2$ each. The tiles are mounted on a 160 mm diameter stainless-steel plate, which is fixed to PAI pillars using M2 screws. A PTFE cover, fastened to the plate with eight screws, secures the array when installed upside down and prevents the tiles from dropping toward the anode located

15 mm below. The SiPM array is powered and read out through a custom potted electrical feedthrough (VACOM) mounted on a DN40CF flange at the top of the setup. It routes bias and signal lines via Kapton-insulated and coaxial PTFE cables to the external patch panel connected to the digitizers. The cables used to power the array are grouped in a breakout box equipped with current limiters to ensure safe operation.

The liquid level must be monitored during filling, recovery, and operation, with particular precision at the liquid–gas interface. A two-segment capacitive level meter installed along the field cage measures the capacitance difference between liquid and GXe, allowing continuous tracking of the LXe height. Three shorter capacitive sensors mounted on PAI pillars provide fine measurements between the gate and anode and enable tilt monitoring. The liquid level is adjusted with a weir system above the field cage, consisting of two concentric cylinders with a central bellows that sets the overflow height.

7.4.1 High-voltage delivery system

Current generation dual-phase TPCs typically operate with electric drift fields in the range of 20 to 200 V/cm [142, 143]. In Xenoscope, the target field of 190 V/cm requires a cathode bias of up to -50 kV, provided by a dedicated high-voltage supply, transmission cable, and vacuum feedthrough. The feedthrough transfers the high voltage from the air side to the cathode plate inside the inner vessel while preserving the vacuum between the outer and inner cryostat volumes. The objective of the HV delivery system is to deliver this potential stably and reliably throughout long-term operation. This section describes the design considerations, challenges, and implementation of the HV delivery system in Xenoscope.

In detectors for low-energy rare event searches, the HV system must minimize spurious electron emission from metallic surfaces. Such emission may arise from photoelectric, thermal, or field-emission processes and can be further enhanced by positive-ion accumulation on electrodes [277]. Spurious electrons can initiate dielectric breakdowns or contribute to few-electron backgrounds relevant for low-energy searches [153, 214]. To suppress emission sites and reduce local field enhancements, fine-wire or mesh electrodes are typically employed instead of solid plates. In dual-phase xenon TPCs, this geometry also ensures high optical transparency, enabling prompt scintillation light to reach both top and bottom photosensor arrays.

Beyond spurious emission, the stability of the applied field also depends on the dielectric properties of the liquid xenon itself. The dielectric breakdown in noble liquids is influenced by many parameters and has been widely studied [274, 277, 278], although many aspects remain inconclusive. The breakdown strength of LXe depends on electrode ge-

ometry, surface finish, stressed area, and liquid purity. Although intrinsic breakdown fields of order MV/cm have been measured under ideal laboratory conditions, practical systems typically reach instability at much lower fields due to microscopic imperfections or gas microbubbles near electrode surfaces [278, 279]. Particulate or gaseous inclusions enhance local field concentrations and promote breakdown, whereas electronegative impurities can mitigate discharge onset by capturing free electrons. Photoelectric and field-emission processes can also generate seed electrons at electrode surfaces, especially under VUV illumination from xenon scintillation light, whose 7 eV photons exceed the stainless-steel work function. Conditioning the electrodes through repeated voltage ramping helps suppress such emission sites by locally smoothing asperities and removing surface contaminants [279–281].

In Xenoscope, the drift field is established by applying a potential difference between the stainless steel cathode plate and the gate electrode. The cathode is biased through a custom HV delivery system inserted upward from the lower section of the outer vessel. An exploded view of the high-voltage delivery system is shown in Figure 7.7, highlighting the interfaces between air, vacuum, and liquid xenon. In the following paragraphs, individual components are referenced by the numbers in parentheses, corresponding to the labels shown in Figure 7.7, unless otherwise specified. The HV power supply is a Heinzinger PNC 100000 negative polarity unit rated for 100 kV and a maximum output current of 1 mA. The voltage is transmitted from the power supply to the cryostat using a Heinzinger cable model HVC100 [282] rated for 100 kV and with a maximum bending radius of 280 mm. The cable has four layers, from its center: a copper inner conductor (6), a white high-molecular weight polyethylene (HMWPE) insulator (5), a copper shielding braid (4), and a polyvinyl chloride (PVC) sheath insulator (3). On the air side, which extends from the power supply to the TPC, the cable is further protected from environmental exposure by a flexible plated steel conduit (1) [283].

The air-to-vacuum feedthrough consists of a high-voltage cable cryofitted in-house into a custom DN40CF flange (9), which is mounted at the bottom section of the outer cryostat vessel, represented by the blue rectangle in Figure 7.7. To create the vacuum-tight seal, the two outer insulating layers of the cable are removed, leaving the HMWPE core exposed. The cable is then cooled to 77 K by immersion in liquid nitrogen, resulting in a thermal contraction of approximately 300 μm . The DN40CF flange contains a central bore slightly smaller than the nominal cable diameter (shown in Figure 7.8 (a)), which is gradually enlarged until the cold, contracted cable can be inserted. Upon warming to room temperature, the cable expands and forms a vacuum-tight seal within the flange. Electrical grounding of the DN40CF flange is achieved by clamping the copper shielding braid of the HV cable using an M24 screw, as illustrated in Figure 7.8 (c), along with




Figure 7.7: Exploded view of the cathode high-voltage delivery system across the air, vacuum and LXe regions: (1) plated steel outer shield, (2) outer shield connector, (3) PVC cable outer layer, (4) copper ground shield, (5) HMWPE cable insulator, (6) copper inner conductor, (7) cable terminal, (8) grounding M24 screw, (9) custom DN40CF flange, (10) DN40CF gasket, (11) HMWPE insulator, (12) stainless steel connector, (13) stainless steel spring, (14) DN125CF flange, (15) ceramic feedthrough, (16) round-edged high-voltage terminal, (17) cup-spring mechanism, (18) cathode plate.

several additional photographs illustrating the fabrication process.

After conducting the first conditioning campaign of the HV feedthrough (see below), a leak was identified in the cryofitted region. The junction between the cable and flange was subsequently sealed with Torr Seal vacuum epoxy [284] to restore vacuum tightness, as shown in Figure 7.8 (f). The epoxy was applied and cured *in situ* after the full assembly was installed.

The vacuum-to-LXe feedthrough extends the high-voltage connection from the vacuum region into the inner vessel. Its components are enumerated in Figure 7.7. A commercial ceramic feedthrough (15) from CeramTec [285], rated for 100 kV, 4 K, and 8.6 bar is coupled to the inner conductor of the HV cable (6) through a spring-loaded stainless-steel connector (12) housed within a custom-made HMWPE insulator (11). The insulator features multiple rounded ridges along its surface to increase the creepage path¹ and thereby reduce surface leakage currents. The stainless-steel connector shown (12), includes a receptacle designed to securely hold the ceramic feedthrough

¹Shortest distance along the surface of an insulating material between two conductive parts at different electrical potentials.




Figure 7.8: The air-to-vacuum feedthrough components during fabrication. **a)**: Grounding M24 screw and custom DN40CF flange. **b)**: Clamping the copper shielding braid to the custom flange. **c)**: Cryofitted high-voltage cable inserted through the flange. **d)**: Same as **c)**, with the M24 screw threaded. **e)**: View of the cryofitted HMWPE section of the HV cable. **f)**: Junction between the cable and flange sealed with vacuum epoxy to ensure leak tightness.

pin, which ensures reliable electrical contact and stable high-voltage transmission toward the LXe side. The ceramic feedthrough itself is welded to a DN125CF flange (14), which is directly attached to the bottom nipple of the inner vessel. At its upper end, the ceramic feedthrough terminates in a threaded pin, to which a rounded stainless-steel contact pin (16) is attached. This contact pin interfaces mechanically and electrically with the cup-spring mechanism (17) that delivers the high voltage to the cathode plate (18). The cup-spring mechanism comprises a round-edged stainless steel cylinder that houses a concave surface and a coil spring. At one end, the cylinder is threaded to a round-edged cup that is attached to the cathode plate (see Figure 7.9 (c)). The opposite end of the cylinder makes contact with the pin from the ceramic feedthrough, establishing reliable mechanical and electrical connectivity. The design of the cup-spring mechanism ensures robust but flexible electrical contact with the cathode, accommodating potential thermal contractions at LXe temperatures. The cathode plate and the cup-spring mechanism feature rounded edges in order to reduce high field regions.

In Figure 7.9 different photos of the high-voltage delivery system during installation and before closing the inner and outer vessel are shown.

Prior to assembly in the inner vessel, the air-to-vacuum feedthrough was tested in a small-scale vacuum setup consisting of a chamber sealed with a DN125CF flange and equipped with a DN40CF port through which the cryofitted HV cable assembly was installed. This helium leak test demonstrated a leak rate of 1×10^{-9} mbar L/s, confirming the vacuum tightness of the cryofitted feedthrough. Afterward, the air-to-vacuum and vacuum-to-LXe feedthroughs were installed in the outer vessel, and a subsequent leak test of the assembled system yielded a leak rate of 1×10^{-8} mbar L/s, meeting the required vacuum specifications for operation. The feedthrough was then installed in the inner vessel and electrically conditioned to improve its resistance to electrical breakdown. This process involved gradually increasing the applied voltage while in-




Figure 7.9: **a):** View of the outer vessel during installation of the air-to-vacuum feedthrough. The HV cable is routed inside DN40CF tubes to protect the cryofitted section and the exposed HMWPE, which houses the HV tip that connects directly to the pin of the ceramic feedthrough (FT) in the vacuum region. **b):** View of the inner vessel showing the bottom section with the ceramic FT already installed. **c):** Cup-and-spring system attached to the cathode plate, used to transmit the high voltage from the ceramic FT to the cathode.

tentionally allowing controlled discharges to occur, thereby removing microscopic field emitters that can initiate breakdown events [279, 281]. The conditioning campaign was carried out in the inner vessel, prior to the installation of other detector subsystems (field cage, liquid-level control, and photosensors), to ensure that no sensitive components were at risk of damage during the procedure.

The voltage was increased stepwise while continuously monitoring the HV power supply for fluctuations in voltage and current, as well as the outer cryostat pressure for monitoring electrical discharges (see Figure 7.10 left). When a breakdown was detected, the voltage was reduced and held at a stable value for 10–15 min before resuming the ramp up. The process was repeated over several days, progressively extending the stable operating range. After conditioning, a maximum cathode bias of -50 kV was




Figure 7.10: Left: Correlation between the pressure in the outer vessel and the voltage during the high-voltage feedthrough conditioning campaign, used to monitor electrical discharges. The data corresponds to day 12 out of 25 of the campaign, when 50 kV was reached for the first time. Discharge events were identified as pressure spikes coincident with voltage drops. **Right:** Pressure and temperature evolution inside the inner vessel during the cold-pressure test. The vacuum-to-LXe FT was cooled to 120 K by filling the inner vessel with LN₂ until the ceramic FT was covered. After several hours, overpressure was generated by injecting N₂. Once 4.05 bar were reached, the system stability was monitored for approximately 5 minutes.

sustained without discharge for over nine consecutive days, corresponding to a drift field of approximately 190 V/cm across the full Xenoscope TPC height.

In addition to the electrical conditioning campaign, temperature and pressure tests were carried out to assess the mechanical integrity and leak tightness of the ceramic feedthrough and the welded DN125CF flange under cryogenic conditions. The ceramic feedthrough assembly was immersed in LN₂ to reproduce the thermal conditions expected during xenon filling cycles. Overpressure conditions were established by injecting N₂ inside the inner vessel. During the test, the pressure, temperature, and vacuum level were continuously monitored to detect any signs of leakage or mechanical deformation. The system exhibited stable behavior at an overpressure of 4 bar and a temperature of approximately 120 K, with no measurable pressure loss or vacuum degradation over a five-minute observation period, as shown in Figure 7.10. These results confirmed the ability of the feedthrough to withstand the combined thermal and mechanical stresses expected during operation and cryogenic transitions.

Following successful validation of the mechanical and electrical stability of the feedthrough assembly, the system was integrated into the full Xenoscope setup for the TPC commissioning, as discussed in Section 7.5.




Figure 7.11: Left: Photograph of the gold-coated photocathode installed in the cathode plate. An optical fibre delivers VUV light from a xenon discharge lamp, generating electrons via the photoelectric effect, which are then drifted through the xenon bulk. **Right:** Photograph of the silicon photomultiplier array, showing the custom PTFE cover, a tile composed of four quads, and a single quad, which covers an effective area of 24 mm^2 .

7.4.2 Electron generation and light detection

As introduced in Section 7.3, electrons are generated via the photoelectric effect by illuminating a photocathode integrated into the stainless-steel cathode plate at the bottom of the field cage (see Figure 7.11 left). The photons are produced by a 60 W xenon gas discharge lamp (model L7685, Hamamatsu [286]), which emits in the 190–2000 nm spectral range, corresponding to photon energies between approximately 0.6 eV and 6.5 eV. The light is collimated by a UV-grade sapphire lens placed at the lamp output and then coupled into an optical fiber that guides the photons toward the photocathode region. The lamp is externally triggered by a pulse generator operated at a maximum frequency of 40 Hz, with a pulse width of $2.9 \mu\text{s}$ at its maximum allowed voltage of 1 kV. The optical fiber is a 3 m multimode, UHV-compatible 600 μm polyimide-coated fiber from LewVac [287], selected for its resistance to VUV-induced degradation and suitability for high-vacuum environments.

The optical fiber maintains a high transmittance of approximately 80 % for several hours of continuous operation under exposure to high-energy UV photons [288]. The fiber is routed through a vacuum feedthrough and terminates 1–2 mm from the photocathode, with its end directly aligned to illuminate the surface. A PTFE holder secures its position during operation and ensures mechanical stability (see Figure 7.11 left). The photocathode comprises a 2 mm-thick, 30 mm-diameter quartz substrate coated with a 5 nm titanium adhesion layer and a 40 nm gold film. Quartz, being transparent and

having low reflectance, minimizes photon losses and enhances light transmission to the gold layer, thereby increasing the photoelectron yield compared to metallic substrates. The thin titanium layer serves solely as an adhesion booster for the gold coating and has a negligible effect on the optical or photoemissive performance of the photocathode. Gold was selected as the photoemissive material for its chemical inertness and stability, it does not oxidize and offers a work function of approximately 5.1 eV, which is within the peak photon energy range of the xenon discharge lamp [286].

The photocathode coatings were deposited using an automatic sputter coater [289], model Q150TS from Quorum, which produced a uniform 40 nm gold film by vapor deposition. The target thickness was chosen based on the effective photoemission depth of gold [290], corresponding to the region within the metal layer where incident photons generate photoelectrons with sufficient energy to escape the surface. Prior to deposition, the quartz substrates were cleaned in a 1 % aqueous solution of ELMA concentrate (1 part ELMA to 99 parts distilled water), followed by ultrasonic cleaning in distilled water for 15 minutes to remove surface contaminants. The coating process was performed under high vacuum conditions ($\mathcal{O}(1 \times 10^{-3} \text{ mbar})$), ensuring good adhesion and uniform thickness on the photocathode surface. Before deployment in the TPC, the photocathode was stored over more than 40 d in a vacuum desiccator for outgassing purposes.

Based on the lamp power, pulse duration, and estimated optical coupling efficiency, the expected number of photoelectrons emitted per pulse can be approximated as

$$N_e^- \sim \Phi_\gamma \times \Omega \times T \times \text{QE}, \quad (7.2)$$

where Φ_γ is the photon flux from the 60 W xenon lamp, $\Omega \approx 0.5^2$ is the effective solid-angle fraction reaching the photocathode, T accounts for transmission losses in the optical path (assuming an attenuation of 3 %/m), and $\text{QE} = 10^{-7}$ is the quantum efficiency of the gold film [291]. Under these conditions, the expected number of emitted photoelectrons per pulse is approximately 6×10^5 .

While the photocathode enables controlled electron generation at the bottom of the TPC, the detection of electroluminescence light at the top is performed by a SiPM array, as described in 7.2. The calibration of the SiPMs is essential in order to monitor its performance over time and to estimate the gain, that is fundamental in the reconstruction of the signals. The calibration system consists of a blue light emitting

²The solid angle fraction was computed from the geometry between the fiber and the photocathode, assuming isotropic emission from the fiber and using the analytical expression for a circular aperture.

diode (~ 460 nm) placed in a black box on the air side. The light is transmitted via an optical silica fibre (600 μ m diameter) covered by a black protection to optically isolate from the environment light. The fibre has an SMA connector on both ends, on one end connected to the black box and on the other connected to one of the vacuum feed-troughs instrumented at the top flange of Xenoscope. On the cryogenic side, the light is split into three naked silica fibers, with transmission band between 250–1200 nm. The three fibers, were polished and inspected for light transmission before installation. Once installed in the cryogenic side, they were routed towards the top stack of the field cage passing along the PAI pillars and pointing up to the photosensor array. The three fibers were distributed with a 120° separation between each other to achieve an homogeneous light coverage. During the installation operations, one of the three fibres was damaged which leaves the calibration system with only two functional fibres. All the tiles from the SiPM array are illuminated with only two fibres and the potential light non-uniformity will be compensated for by using different light levels for different channels.

7.5 Commissioning results

The commissioning campaign of the dual-phase 2.6 m TPC was conducted over a four-month period, from April to August 2024. During this phase, the xenon gas purifier was offline for maintenance, resulting in poor purity levels in the LXe. The xenon used had been purified in the purity monitor run, but its impurity concentration was no longer actively controlled. While moderate levels of electronegative impurities can, in some cases, quench dielectric breakdown by capturing charge carriers and suppressing feedback processes, excessive or volatile impurities may instead promote bubble-assisted discharges. The absence of purification therefore prevented stable and predictable high-voltage conditions, as the impurity content could either mitigate or exacerbate discharges depending on its composition and level [278, 279]. In addition, the electron attenuation length was reduced, preventing the observation of S2 signals that originated significantly below the gate electrode. Consequently, the primary objective of this run was not physics data acquisition, but rather the commissioning and performance validation of the newly implemented subsystems: the SiPM readout array, the liquid-level control system, and the high-voltage delivery system.

Xenon filling and level setting

The filling campaign began in April 2024, following a complete leak test and vacuum outgassing of the xenon handling system. Between late April and May, operations pro-




Figure 7.12: Operating conditions during the 20d filling campaign of the commissioning run of the 2.6m-tall dual-phase TPC. **Top:** Pressure and temperature evolution inside the inner cryostat vessel. **Bottom:** Xenon mass increase throughout the filling period together with the corresponding liquid-xenon level in the TPC.

gressed from the initial cooldown and first xenon transfer to active gas recirculation and system stabilization. The pre-cooler was first employed on April 24 to assist the cool down of the inner vessel and was later used during the final large-scale fillings to regulate flow and pressure. Throughout the process, the cold head temperature (typically maintained between 168–172 K) and pump flow were finely tuned to keep the inner vessel pressure stable at approximately 2 bar, avoiding both xenon freezing and overpressure conditions. The full filling campaign extended over about two weeks, as shown in Figure 7.12, with typical flow rates around 50 slpm. The xenon mass transferred to the detector was tracked by comparing the inflow measured by the mass flow controller to the recirculation flow through the purification loop, indicated in Figure C.1 by FC01 and FM01, respectively. After transferring roughly 40 kg of xenon, the liquid level became visible in the long level meters. In total, ~ 360 kg were filled into the inner vessel until the liquid level reached the anode.

The short level meters were calibrated by gradually raising the height of the LXe until each sensor was fully submerged, allowing for the determination of their maximum capacitance $C_{\max}$. The point at which the liquid first made contact with the lower edge of the sensor produced a sudden increase in capacity due to capillary effects, defining C_{cap} , while the baseline reading of the gas phase corresponded to $C_{\min}$. These three reference values enabled a direct conversion from capacitance to liquid height using the linear calibration relation $z(C) = 30\text{mm} \cdot (C - C_{\text{cap}})/(C_{\max} - C_{\min})$. Based on

Region	Distance [cm]	Electric Field [V/cm]	
		LXe	GXe
Anode - gate	1.2	3.5×10^3	6.7×10^3
Gate - cathode	260	100	—

Table 7.1: Electrostatic field configuration of the Xenoscope electrodes and the regions between them. For these settings, the distance between the liquid–gas interface was 6 mm. The electric drift field values were maintained for five consecutive days, while the extraction electric field was maintained for ten days.

the calibrated readings and the fine adjustment provided by the height-tunable weir system, the liquid level was then set approximately 6 mm above the gate electrode, centered between the gate and anode.

High-voltage setting

Once the liquid level was properly established, the electrode biasing procedure began, starting with the anode and gate electrodes. The voltages were increased gradually in 50 V steps to monitor stability and prevent discharges, while the current of the power supply was continuously monitored. During the initial ramp-up, the anode power supply tripped when the anode reached 2.6 kV and the gate was at -1.6 kV. Following this event, the system was successfully reconditioned and ramped again without further interruptions, reaching a stable potential difference of 7.2 kV, distributed as 3.6 kV on the anode and -3.6 kV on the gate. This voltage difference established the extraction region across the liquid–gas interface, corresponding to electric fields of 3.5 kV/cm in LXe and 6.7 kV/cm in GXe, as determined by their dielectric permittivities, $\epsilon_{\text{LXe}} = 1.85$ and $\epsilon_{\text{GXe}} = 1.00127$ [138]. Under these conditions, the electron extraction efficiency from liquid to gas is approximately 75 %, as inferred from the data presented in Ref. [211].

The cathode electrode was manually ramped up in parallel with the top stack to prevent a downward flux of electrons and thus potential damage to the thin gold film of the coated photocathode. The voltage was increased in steps of approximately 200 V every 2 min, allowing sufficient time for electrical stabilization between steps. Once stable operation was confirmed, the cathode voltage was gradually raised to -26 kV after about 7 d of conditioning, corresponding to a drift field of approximately 86 V/cm across the 2.6 m-long TPC. After four days of continuous data acquisition, the cathode bias was further increased to -29 kV, achieving a drift field of about 100 V/cm, as shown in Figure 7.13. The current was continuously monitored, with no concerning increase or signs of discharge observed. A summary of the electrostatic field configuration is provided in Table 7.1.




Figure 7.13: Evolution of the drift and extraction fields as a function of time since July 5, 2024. The drift field (red) and extraction field (blue) are derived from the monitored high-voltage settings of the cathode, gate, and anode electrodes. The gray vertical dashed line indicates the moment of the top stack trip.

SiPM calibration and muon data acquisition

The SiPM gain was monitored through regular LED calibrations. The signals were recorded by two CAEN v1724 digitizers [292], connected in a daisy-chain configuration and interfaced with a computer. For the gain calibration, the SiPMs were illuminated using five different LED light intensities, which ensures that each sensor receives sufficient photons to fully populate the SPE spectrum. The charge was reconstructed by integrating a 400 ns window that encompasses the entire LED signal. The resulting SPE charge spectra were fitted with a Gaussian function, and the mean values were used to compute the gain of each SiPM tile. Data handling and processing were performed using PyLArS [293], a custom Python framework developed for SiPM characterization and peak processing.

Figure 7.14 left shows the gain evolution of all tiles during the period from 2024-07-10 to 2024-07-24, when the drift and extraction fields were established, as well as the liquid level. During this period, most of the tiles exhibited a standard deviation below 1 %, as shown in Figure 7.14 right, with gains consistent with standard photosensor operation in current-generation liquid xenon detectors. Tiles B, G, and M consistently showed lower gains, likely due to higher local board temperatures caused by the presence of PT100 sensors mounted for temperature monitoring.

Data acquisition in dual-phase mode was performed using both muon-coincidence and channel-trigger modes. In channel-trigger mode, illustrated in Figure 7.15, signals are recorded independently of the muon-coincidence trigger, allowing for detector self-triggering at a manageable rate (<1 MHz). Signals from the three central detector tiles




Figure 7.14: **Left:** Mean gain of each SiPM tile measured from LED calibrations performed almost daily, with 16 calibration runs acquired over a period of approximately 12 days (2024-07-10 to 2024-07-24) at a bias voltage of 48.6 V. **Right:** Distribution of the standard deviations of the tile gains. The maximum relative standard deviation of the gain per tile remains below 3.5 % during stable operation.

(C, F, and G) were used (see the tile positions in Figure 7.16). Each signal was split using a CAEN N625 linear Fan-in/Fan-out module into three outputs: one directed to the digitizer for recording, one to an oscilloscope for live monitoring and debugging, and one to a discriminator with a fixed threshold of 15 mV. The discriminator outputs were then sent to the coincidence logic unit, configured for a triple-coincidence scheme. A trigger was generated only when all three tiles registered signals above threshold within the coincidence window.

In the muon-coincidence configuration, two BC-412 plastic scintillator panels were installed on opposite sides of the cryostat, with a vertical separation of approximately 30 cm, with the upper panel just above the height of the liquid–gas interface. With this geometry, only muons with an incident angle of about 10° relative to the vertical can traverse both scintillation panels and produce a coincidence signal. When a muon passed through both scintillators, the emitted light was detected by two PMTs. Each PMT’s analog signal was split: one copy was sent to the CAEN v1724 digitizer for waveform recording, and the other to a CAEN N840 leading-edge discriminator, which generated a logic pulse when the signal exceeded a noise-level threshold. The discriminator outputs were fed into a CAEN N455 coincidence logic unit operating in AND mode. A coincidence between the two scintillators produced a trigger pulse, initiating data acquisition in the digitizer. The same trigger was also sent to an ORTEC 997 analog counter for real-time monitoring of muon event rates.

Figure 7.16 shows a representative waveform of a correlated S1–S2 event from a muon traversing the TPC in the LXe region near the gate. The trigger signal from the scintillator PMTs is coincident with the S1 pulse, visible as a sharp feature at approximately




Figure 7.15: Workflow of the channel-trigger data acquisition mode. The analog signal from each tile is routed to a fan-in/fan-out module, which distributes the signal to three parallel paths: an oscilloscope for monitoring, a digitizer for waveform recording, and a discriminator. The discriminator applies a fixed threshold of 15 mV above the noise level and converts the analog signal into a logic pulse. The logic pulses from the three selected tiles are then sent to a coincidence unit, which generates a trigger only when a three-fold coincidence condition is satisfied within the predefined time window.

2 μs with a rise time of about 0.1 μs . The subsequent S2 signal appears as a broader pulse, lasting between 1–6 μs , with a characteristic rise time of roughly 0.2 μs . The S1 and S2 areas are 4.2×10^3 PE and 5.96×10^4 PE, respectively. Three characteristic event topologies were reconstructed: (i) S1-like pulses without a correlated S2, (ii) S2-like wide pulses without a preceding S1, and (iii) correlated S1–S2 pairs with time separations between 0 μs and 10 μs . Events with larger time separations, corresponding to interactions deeper in the TPC, were not observed. This absence is consistent with the lack of xenon purification during the measurement, which caused drifting electrons to be absorbed by impurities before reaching the extraction region. The detection of correlated S1–S2 signals from cosmic muons confirms the correct operation of the dual-phase TPC and all associated subsystems.

7.6 Next steps

The successful commissioning of the 2.6 m Xenoscope TPC marks a major milestone in demonstrating the stable operation of large-scale dual-phase xenon detectors. The simultaneous validation of the high-voltage delivery, liquid-level control, and photosensor systems confirms the readiness of the facility for extended data-taking. A subsequent run in 2025 further explored subsystem performance, providing valuable operational lessons on the xenon compressor stability, the high-voltage robustness and electron




Figure 7.16: Left: Waveform of an S1–S2 event from a muon crossing the LXe region near the gate. The top panel shows the summed waveform and muon trigger; the bottom shows individual channels. **Right:** Corresponding hit pattern with reconstructed interaction vertex (cross) from the S2 light distribution; the red circle marks the 16 cm inner diameter of the field cage. Figures taken from [3].

production, but not yet demonstrating electron drift across the full vertical length.

Further runs are scheduled for 2026 to demonstrate sustained electron drift across the full 2.6 m drift length and to perform quantitative studies of electron transport and diffusion in LXe. These measurements will provide critical input for the design of the next-generation XLZD detector. In parallel, future upgrades are planned to enhance the optical characterization capabilities of Xenoscope, such as large-scale measurements of light attenuation and scattering, and to increase the spatial granularity of the SiPM readout plane, aiming for improved topological event reconstruction within a few years.

Summary and outlook

The nature of dark matter remains one of the most compelling open questions in modern physics. A wide range of astrophysical and cosmological observations indicate that a significant fraction of the matter content of the Universe consists of a non-luminous and non-baryonic component, whose fundamental nature is still unknown. Among the most well-motivated particle candidates are weakly interacting massive particles (WIMPs), which arise naturally in several extensions of the Standard Model and are expected to interact only rarely with ordinary matter.

Over the past decades, dual-phase liquid xenon time projection chambers have emerged as the leading technology, combining large target masses with excellent background suppression and event reconstruction capabilities. The XENON program has played a central role in this effort, improving sensitivity over successive generations of detectors. XENONnT represents the latest stage of this development, operating a multi-ton liquid xenon target with unprecedented radiopurity and detector performance. In this context, this thesis addresses detector performance, physics analyses, and the extension of its characterization to MeV-scale energies.

The response of the photomultiplier tubes was continuously monitored and modeled, resulting in a gain stability better than 3% over a 400 d period. This gain model was used in the second science run and contributed to the physics results, including leading constraints on the spin-independent WIMP–nucleon cross section, with an upper limit of $1.7 \times 10^{-47} \text{ cm}^2$ at a WIMP mass of 30 GeV/ c^2 [55], and the first observation of solar neutrinos via coherent elastic neutrino–nucleus scattering in a liquid xenon detector, corresponding to a measured ^8B solar neutrino flux of $4.7 \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$ [89].

The detector response was characterized in the MeV energy range through a dedicated simulation framework, including data-driven yield models, a clustering algorithm for energy deposits, and a simulation-driven correction map for S1 signals. Together with a dedicated reconstruction algorithm for multi-scatter topologies, these developments establish the detector-modeling tools required to assess the feasibility of MeV-scale searches in XENONnT.

This framework was applied to a sensitivity study of the ^{124}Xe decay via two-neutrino electron capture with positron emission. Full-chain simulations of signal and background were performed, including large-scale material background modeling and rejection techniques based on fiducialization and machine learning. After all selections, the analysis remains background-limited. A projected lower limit on the half-life of 5.35×10^{20} years at 90% confidence level is obtained, providing a benchmark for future searches and identifying radiogenic backgrounds as the main limitation in this energy regime.

Developments toward next-generation liquid xenon detectors were carried out within the XLZD program. The Xenoscope facility at the University of Zurich serves as a full-scale vertical demonstrator for key detector technologies relevant to XLZD. In this context, a high-voltage feedthrough capable of delivering up to 50 kV was designed, constructed, and commissioned for operation in Xenoscope. The upgrade to a 2.6 m dual-phase time projection chamber enabled the first operation of the detector in this configuration and the observation of S1 and S2 signals. While electron drift over the full length has not yet been achieved, these results constitute an important step toward demonstrating this capability in future runs.

These results reflect the current status of liquid xenon detectors as mature instruments for low-energy rare-event searches. At the same time, the characterization of the XENONnT response at the MeV scale broadens the scope of the detector beyond its traditional low-energy program and advances the feasibility of rare-event searches in this regime. Next-generation experiments such as XLZD aim to extend this program by scaling detector mass to the tens-of-ton level, pushing sensitivity into new regions of parameter space. The results obtained with Xenoscope show that key technological challenges, such as electron drift over distances of up to 2.6 m and high-voltage delivery, are being progressively addressed, with the studies presented here representing initial steps toward their realization. Liquid xenon detectors will therefore continue to play a central role in the search for dark matter and rare processes, with future progress driven by detector scale, background control, and improved modeling of detector response.

Appendix A

XENONnT Material screening results

Component	^{238}U	^{60}Co	^{40}K	^{137}Cs	^{228}Th	^{235}U	^{232}Th	^{226}Ra
Lower / Top copper ring	< 1.20	< 0.03	0.28 ± 0.14	< 0.05	0.03 ± 0.05	< 0.55	0.04 ± 0.05	< 0.03
Field shaping wires	0.03 ± 0.01	–	< 0.28	–	< 0.01	–	< 0.01	–
Field shaper rings	< 3.31	0.43 ± 0.04	0.56 ± 0.23	< 0.03	< 0.04	< 0.17	< 0.04	< 0.09
Bottom PMT plate	< 1.20	0.18 ± 0.01	< 0.28	< 0.05	< 0.03	< 0.55	< 0.04	< 0.03
Out. cryostat elongation	4.00 ± 2.00	0.61 ± 0.05	1.40 ± 0.20	0.03 ± 0.02	0.57 ± 0.06	< 0.31	1.20 ± 0.2	1.34 ± 0.09
Out. cryostat shell	2.40 ± 0.70	9.70 ± 0.80	< 2.70	< 0.64	< 0.36	0.11 ± 0.03	0.21 ± 0.06	< 0.64
Out. cryostat flanges	1.30 ± 0.40	37.8 ± 2.0	< 2.70	< 0.21	4.50 ± 0.60	0.80 ± 0.23	0.14 ± 0.06	0.41 ± 0.10
Inn. cryostat shells	3.70 ± 0.60	28.3 ± 2.1	1.60 ± 0.60	< 0.21	0.50 ± 0.10	0.70 ± 0.30	0.10 ± 0.08	0.30 ± 0.10
Inn. cryostat flanges	1.40 ± 0.40	37.3 ± 0.90	< 5.60	< 1.50	4.50 ± 0.60	0.06 ± 0.02	0.21 ± 0.06	< 4.00
Inn. cryostat bottom dome	4.00 ± 2.00	0.61 ± 0.05	1.40 ± 0.20	0.03 ± 0.02	0.57 ± 0.06	0.70 ± 0.31	1.20 ± 0.20	1.34 ± 0.09
Bell side wall	3.70 ± 0.60	0.28 ± 0.21	1.60 ± 0.60	< 0.21	0.50 ± 0.10	0.80 ± 0.23	0.14 ± 0.06	0.41 ± 0.10
Bell side wall bottom lip	4.00 ± 2.00	0.61 ± 0.05	1.40 ± 0.20	0.03 ± 0.02	0.57 ± 0.06	< 0.31	1.20 ± 0.20	1.34 ± 0.09
PTFE sliding reflectors	2.50 ± 0.30	0.40 ± 0.10	< 1.70	< 0.16	0.40 ± 0.10	< 0.56	0.40 ± 0.10	0.60 ± 0.10
PTFE blocking reflector	0.11 ± 0.07	–	0.08 ± 0.03	< 0.01	< 0.04	< 0.05	0.04 ± 0.02	0.15 ± 0.03
PTFE bottom TPC	0.51 ± 0.07	–	8.00 ± 1.00	< 0.07	< 0.09	< 0.08	0.12 ± 0.03	< 0.03
TPC PMT bases	1.46 ± 0.14	–	0.29 ± 0.04	< 0.02	< 0.05	0.06 ± 0.01	0.40 ± 0.05	0.65 ± 0.03
Neutron veto PMTs	10.20 ± 3.26	0.68 ± 0.07	13.47 ± 2.89	0.16 ± 0.09	0.24 ± 0.04	0.41 ± 0.13	0.46 ± 0.19	0.49 ± 0.15

Table A.1: Material contamination from screening campaign. Values in $\mu\text{Bq/kg}$ [175]. For PMTs, values represent the average of reused XENON1T units and newly produced ones. Upper limits are quoted where no activity above background was observed and are conservatively treated as central values without uncertainty. Components or isotopes without available measurements are excluded from the analysis.

Appendix B

SR0 Material background simulation campaign

Component	Co60	K40	Ra226	Th228	Th232	U235	U238
CopperFieldGuard	1.93×10^7	1.00×10^8	3.82×10^7	8.03×10^7	9.57×10^7	3.66×10^7	5.00×10^8
CopperFieldShaperRing	1.66×10^6	5.00×10^7	9.30×10^7	4.93×10^7	5.05×10^7	3.63×10^7	4.45×10^7
CopperPMTPlate	2.21×10^6	1.42×10^7	1.79×10^7	1.42×10^7	1.13×10^7	3.66×10^8	3.66×10^8
CopperTopRing	1.86×10^7	3.86×10^7	3.77×10^7	3.77×10^7	4.02×10^7	3.07×10^7	2.88×10^7
InnerCryostat	7.03×10^7	7.98×10^7	6.95×10^8	1.83×10^8	1.60×10^8	1.83×10^8	4.75×10^8
InnerCryostatFlange	1.31×10^8	1.89×10^9	7.08×10^8	2.86×10^8	2.87×10^8	1.00×10^{10}	1.54×10^8
InnerCryostatRings	7.11×10^6	2.00×10^8	4.03×10^7	2.00×10^8	7.03×10^7	6.29×10^9	1.42×10^7
NvPMT	6.29×10^9	6.29×10^9	6.29×10^9	6.29×10^9	6.29×10^9	6.29×10^9	6.29×10^9
OuterCryostat	1.15×10^8	5.73×10^7	7.44×10^8	1.54×10^8	1.79×10^8	1.44×10^8	1.37×10^9
OuterCryostatFlange	4.34×10^8	3.50×10^9	1.42×10^9	2.96×10^9	2.00×10^9	1.40×10^{10}	1.62×10^8
OuterCryostatRings	3.93×10^7	3.50×10^8	1.30×10^8	2.86×10^8	2.00×10^8	7.32×10^9	3.38×10^7
PMT	3.83×10^7	2.31×10^8	9.12×10^7	3.61×10^7	3.60×10^7	9.14×10^7	9.33×10^8
SSBell	7.50×10^7	4.63×10^7	4.66×10^7	4.54×10^7	4.68×10^7	4.64×10^7	9.14×10^7
SSElectrodeRing	1.26×10^7	4.69×10^7	4.69×10^7	9.87×10^7	9.89×10^7	2.98×10^7	5.03×10^7
TeflonBottomTPC	3.21×10^7	5.00×10^7	3.17×10^7	2.74×10^7	9.83×10^7	3.43×10^7	9.78×10^7
TeflonPillar	2.08×10^7	5.00×10^7	9.74×10^7	8.90×10^7	4.97×10^7	3.30×10^7	3.31×10^7
TeflonTPC	2.60×10^7	5.00×10^7	4.82×10^7	4.78×10^7	4.83×10^7	4.89×10^7	3.12×10^7

Table B.1: Summary of the SR0 material background simulation campaign. Given the large-scale, the simulations were carried out in the OSG server during around three months. For each parent component, the number of simulated primaries $N_{\text{primaries}}$ per isotope are indicated.

Appendix C

Xenoscope gas system




Figure C.1: Piping and Instrumentation Diagram (P&ID) of the Xenoscope facility. The main subsystems are indicated by solid colored outlines: purification gas panel (blue), storage array (orange), heat-exchanger gas panel (pink), and cooling tower with the double-walled cryostat (purple). High-pressure, low-pressure, and vacuum lines are indicated by line thickness. The component naming scheme is given in the table on the right-hand side.

Bibliography

- [1] H. S. Eißing, D. Xu, D. R. García, C. Fuselli, Z. Xu, L. Principe et al., *XenonnT/fuse: 1.4.4*, Zenodo, Jan., 2025, 10.5281/zenodo.14740249.
- [2] D. Xu, R. Hammann, K. D. Morå, L. Yuan, Y. Ma, L. Hoetsch et al., *XenonnT/alea: v0.3.5*, Zenodo, Mar., 2026, 10.5281/zenodo.18840642.
- [3] M. Adrover, M. Babicz, L. Baudis, Y. Biondi, A. Bismark, C. Capelli et al., *Commissioning of the 2.6 m tall two-phase xenon time projection chamber of xenoscope*, [Journal of Instrumentation](#) **20** (2025) P04013.
- [4] M. K. Gaillard, P. D. Grannis and F. J. Sciulli, *The Standard model of particle physics*, [Rev. Mod. Phys.](#) **71** (1999) S96 [[arXiv:hep-ph/9812285](#)].
- [5] G. Bertone and D. Hooper, *History of dark matter*, [Rev. Mod. Phys.](#) **90** (2018) 045002 [[arXiv:1605.04909](#)].
- [6] C. Balazs, T. Bringmann, F. Kahlhoefer and M. White, *A Primer on Dark Matter*, [Astrophysics](#) **5** (2026) 17 [[arXiv:2411.05062](#)].
- [7] D. Clowe, M. Bradac, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones et al., *A direct empirical proof of the existence of dark matter*, [Astrophys. J. Lett.](#) **648** (2006) L109 [[arXiv:astro-ph/0608407](#)].
- [8] A. R. Liddle, *An introduction to modern cosmology*. 1998.
- [9] NASA/CXC/CfA/M. Markevitch et al., NASA/STScI, Magellan/U. Arizona/D. Clowe et al. and ESO WFI, *Bullet cluster composite image*, <https://chandra.harvard.edu/photo/2006/1e0657/>, 2006.
- [10] F. Zwicky, *Die Rotverschiebung von extragalaktischen Nebeln*, [Helv. Phys. Acta](#) **6** (1933) 110.
- [11] R. Massey, T. Kitching and J. Richard, *The dark matter of gravitational lensing*, [Rept. Prog. Phys.](#) **73** (2010) 086901 [[arXiv:1001.1739](#)].

- [12] H. Katayama and K. Hayashida, *X-ray study of the dark matter distribution in clusters of galaxies with chandra*, [Adv. Space Res. **34** \(2004\) 2519](#) [[arXiv:astro-ph/0405363](#)].
- [13] D. Harvey, R. Massey, T. Kitching, A. Taylor and E. Tittley, *The nongravitational interactions of dark matter in colliding galaxy clusters*, [Science **347** \(2015\) 1462](#).
- [14] V. C. Rubin and W. K. Ford, Jr, *Rotation of the andromeda nebula from a spectroscopic survey of emission regions.*, [Astrophys. J. **159**: 379-403\(Feb 1970\). \(1970\)](#) .
- [15] K. C. Freeman, *On the disks of spiral and SO Galaxies*, [Astrophys. J. **160** \(1970\) 811](#).
- [16] V. C. Rubin, W. K. Ford, Jr. and N. Thonnard, *Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa through Sc*, [Astrophys. J. Lett. **225** \(1978\) L107](#).
- [17] Y. Sofue and V. Rubin, *Rotation curves of spiral galaxies*, [Ann. Rev. Astron. Astrophys. **39** \(2001\) 137](#) [[arXiv:astro-ph/0010594](#)].
- [18] F. Lelli, S. S. McGaugh and J. M. Schombert, *SPARC: Mass Models for 175 Disk Galaxies with Spitzer Photometry and Accurate Rotation Curves*, [Astron. J. **152** \(2016\) 157](#) [[arXiv:1606.09251](#)].
- [19] Planck collaboration, *Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters*, [Astron. Astrophys. **594** \(2016\) A11](#) [[arXiv:1507.02704](#)].
- [20] D. T. Wilkinson, *Measurement of the Cosmic Microwave Background at 8.56-mm Wavelength*, [Phys. Rev. Lett. **19** \(1967\) 1195](#).
- [21] D. J. Fixsen, *The Temperature of the Cosmic Microwave Background*, [Astrophys. J. **707** \(2009\) 916](#) [[arXiv:0911.1955](#)].
- [22] K. M. Górski, *Cosmic microwave background anisotropy in the COBE DMR 4-year sky maps*, in *31st Rencontres de Moriond: Microwave Background Anisotropies*, pp. 77–84, 1996, [arXiv:astro-ph/9701191](#).
- [23] W. Hu and S. Dodelson, *Cosmic Microwave Background Anisotropies*, [Ann. Rev. Astron. Astrophys. **40** \(2002\) 171](#) [[arXiv:astro-ph/0110414](#)].

- [24] Planck collaboration, *Planck 2018 results. VI. Cosmological parameters*, [Astron. Astrophys. **641** \(2020\) A6](#) [[arXiv:1807.06209](#)].
- [25] WMAP collaboration, *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results*, [Astrophys. J. Suppl. **208** \(2013\) 20](#) [[arXiv:1212.5225](#)].
- [26] M. Cirelli, A. Strumia and J. Zupan, *Dark Matter*, [arXiv:2406.01705](#).
- [27] A. M. Green and B. J. Kavanagh, *Primordial Black Holes as a dark matter candidate*, [J. Phys. G **48** \(2021\) 043001](#) [[arXiv:2007.10722](#)].
- [28] K. N. Abazajian, *Sterile neutrinos in cosmology*, [Phys. Rept. **711-712** \(2017\) 1](#) [[arXiv:1705.01837](#)].
- [29] T. Mannel, *Theory and phenomenology of CP violation*, [Nucl. Phys. B Proc. Suppl. **167** \(2007\) 115](#).
- [30] D. J. E. Marsh, *Axion Cosmology*, [Phys. Rept. **643** \(2016\) 1](#) [[arXiv:1510.07633](#)].
- [31] G. Steigman and M. S. Turner, *Cosmological Constraints on the Properties of Weakly Interacting Massive Particles*, [Nucl. Phys. B **253** \(1985\) 375](#).
- [32] L. Roszkowski, *Particle dark matter: A Theorist's perspective*, [Pramana **62** \(2004\) 389](#) [[arXiv:hep-ph/0404052](#)].
- [33] H. Baer, K.-Y. Choi, J. E. Kim and L. Roszkowski, *Dark matter production in the early Universe: beyond the thermal WIMP paradigm*, [Phys. Rept. **555** \(2015\) 1](#) [[arXiv:1407.0017](#)].
- [34] D. Hooper and S. Profumo, *Dark Matter and Collider Phenomenology of Universal Extra Dimensions*, [Phys. Rept. **453** \(2007\) 29](#) [[arXiv:hep-ph/0701197](#)].
- [35] J. L. Feng, *The WIMP paradigm: Theme and variations*, [SciPost Phys. Lect. Notes **71** \(2023\) 1](#) [[arXiv:2212.02479](#)].
- [36] T. M. Hong, *Dark matter searches at the LHC*, in *5th Large Hadron Collider Physics Conference*, 9, 2017, [arXiv:1709.02304](#).
- [37] Fermi-LAT collaboration, *Searching for Dark Matter Annihilation from Milky Way Dwarf Spheroidal Galaxies with Six Years of Fermi Large Area Telescope Data*, [Phys. Rev. Lett. **115** \(2015\) 231301](#) [[arXiv:1503.02641](#)].
- [38] D. Hooper and F. Vassallo, *Searching for dark matter annihilation in the Sun with the IceCube Upgrade*, [Phys. Rev. D **112** \(2025\) 035014](#) [[arXiv:2505.06734](#)].

- [39] J. Kopp, *Constraints on dark matter annihilation from AMS-02 results*, [Phys. Rev. D **88** \(2013\) 076013 \[arXiv:1304.1184\]](#).
- [40] M. W. Goodman and E. Witten, *Detectability of Certain Dark Matter Candidates*, [Phys. Rev. D **31** \(1985\) 3059](#).
- [41] M. Schumann, *Direct Detection of WIMP Dark Matter: Concepts and Status*, [J. Phys. G **46** \(2019\) 103003 \[arXiv:1903.03026\]](#).
- [42] N. W. Evans and J. H. An, *Distribution function of the dark matter*, [Phys. Rev. D **73** \(2006\) 023524 \[arXiv:astro-ph/0511687\]](#).
- [43] J. Bovy and S. Tremaine, *On the local dark matter density*, [Astrophys. J. **756** \(2012\) 89 \[arXiv:1205.4033\]](#).
- [44] A. Cooray and R. K. Sheth, *Halo Models of Large Scale Structure*, [Phys. Rept. **372** \(2002\) 1 \[arXiv:astro-ph/0206508\]](#).
- [45] M. C. Smith et al., *The RAVE Survey: Constraining the Local Galactic Escape Speed*, [Mon. Not. Roy. Astron. Soc. **379** \(2007\) 755 \[arXiv:astro-ph/0611671\]](#).
- [46] J. Aalbers, B. Pelssers, J. R. Angevaare and K. D. Morå, *Jelleaalbers/wimprates: v0.5.0*, Zenodo, 2022, 10.5281/zenodo.2604222.
- [47] R. H. Helm, *Inelastic and Elastic Scattering of 187-Mev Electrons from Selected Even-Even Nuclei*, [Phys. Rev. **104** \(1956\) 1466](#).
- [48] E. Aprile, A. E. Bolotnikov, A. L. Bolozdynya and T. Doke, *Noble Gas Detectors*. Wiley, 2008, [10.1002/9783527610020](#).
- [49] CRESST collaboration, *First results from the CRESST-III low-mass dark matter program*, [Phys. Rev. D **100** \(2019\) 102002 \[arXiv:1904.00498\]](#).
- [50] DarkSide collaboration, *Low-Mass Dark Matter Search with the DarkSide-50 Experiment*, [Phys. Rev. Lett. **121** \(2018\) 081307 \[arXiv:1802.06994\]](#).
- [51] PandaX collaboration, *Dark matter search results from 1.54 Tonne · Year exposure of pandax-4t*, [Phys. Rev. Lett. **134** \(2025\) 011805](#).
- [52] DEAP collaboration, *Search for dark matter with a 231-day exposure of liquid argon using DEAP-3600 at SNOLAB*, [Phys. Rev. D **100** \(2019\) 022004 \[arXiv:1902.04048\]](#).
- [53] SuperCDMS collaboration, *Results from the Super Cryogenic Dark Matter*

- Search Experiment at Soudan*, *Phys. Rev. Lett.* **120** (2018) 061802 [[arXiv:1708.08869](#)].
- [54] SuperCDMS collaboration, *Low-mass dark matter search with CDMSSlite*, *Phys. Rev. D* **97** (2018) 022002 [[arXiv:1707.01632](#)].
- [55] XENON collaboration, *WIMP Dark Matter Search using a 3.1 tonne $\times$ year Exposure of the XENONnT Experiment*, [arXiv:2502.18005](#).
- [56] LZ collaboration, *Dark matter search results from 4.2 Tonne–Years of exposure of the lux–zeplin (lz) experiment*, *Phys. Rev. Lett.* **135** (2025) 011802.
- [57] XLZD collaboration, *The XLZD Design Book: Towards the Next-Generation Liquid Xenon Observatory for Dark Matter and Neutrino Physics*, [arXiv:2410.17137](#).
- [58] C. A. J. O’Hare, *New Definition of the Neutrino Floor for Direct Dark Matter Searches*, *Phys. Rev. Lett.* **127** (2021) 251802 [[arXiv:2109.03116](#)].
- [59] K. Dundas Morå, I. Olcina, X. Chen and Y. Ma, *ddimit_plot*, 2026.
- [60] W. Pauli, *Dear radioactive ladies and gentlemen*, *Phys. Today* **31N9** (1978) 27.
- [61] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse and A. D. McGuire, *Detection of the free neutrino: A Confirmation*, *Science* **124** (1956) 103.
- [62] G. Danby, J.-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz et al., *Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos*, **9** (1962) 36.
- [63] R. Davis, Jr., D. S. Harmer and K. C. Hoffman, *Search for neutrinos from the sun*, *Phys. Rev. Lett.* **20** (1968) 1205.
- [64] DONUT collaboration, *Observation of tau neutrino interactions*, *Phys. Lett. B* **504** (2001) 218 [[arXiv:hep-ex/0012035](#)].
- [65] B. Pontecorvo, *Inverse Beta Processes and Nonconservation of Lepton Charge*, *Sov. Phys. JETP* **7** (1958) 172.
- [66] Z. Maki, M. Nakagawa and S. Sakata, *Remarks on the unified model of elementary particles*, *Prog. Theor. Phys.* **28** (1962) 870.
- [67] Super-Kamiokande collaboration, *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (1998) 1562 [[arXiv:hep-ex/9807003](#)].

- [68] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*. 2007, [10.1093/acprof:oso/9780198508717.001.0001](#).
- [69] 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** \(2001\) 071301](#) [[arXiv:nucl-ex/0106015](#)].
- [70] SNO collaboration, *Direct evidence for neutrino flavor transformation from neutral-current interactions in SNO*, [AIP Conf. Proc. **646** \(2002\) 43](#).
- [71] KamLAND collaboration, *First results from KamLAND: Evidence for reactor anti-neutrino disappearance*, [Phys. Rev. Lett. **90** \(2003\) 021802](#) [[arXiv:hep-ex/0212021](#)].
- [72] Daya Bay collaboration, *Observation of electron-antineutrino disappearance at Daya Bay*, [Phys. Rev. Lett. **108** \(2012\) 171803](#) [[arXiv:1203.1669](#)].
- [73] K2K collaboration, *Indications of neutrino oscillation in a 250 km long baseline experiment*, [Phys. Rev. Lett. **90** \(2003\) 041801](#) [[arXiv:hep-ex/0212007](#)].
- [74] T2K collaboration, *Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam*, [Phys. Rev. Lett. **107** \(2011\) 041801](#) [[arXiv:1106.2822](#)].
- [75] L. Wolfenstein, *Neutrino oscillations in matter*, [Phys. Rev. D **17** \(1978\) 2369](#).
- [76] S. P. Mikheyev and A. Y. Smirnov, *Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos*, [Sov. J. Nucl. Phys. **42** \(1985\) 913](#).
- [77] F. Capozzi, E. Di Valentino, E. Lisi, A. Marrone, A. Melchiorri and A. Palazzo, *Global constraints on absolute neutrino masses and their ordering*, [Phys. Rev. D **95** \(2017\) 096014](#) [[arXiv:1703.04471](#)].
- [78] Particle Data Group collaboration, *Review of particle physics*, [Phys. Rev. D **110** \(2024\) 030001](#).
- [79] KATRIN collaboration, *Beta decay and neutrino mass: KATRIN and beyond*, [PoS NOW2024 \(2025\) 049](#).
- [80] S. M. Bilenky and C. Giunti, *Neutrinoless double-beta decay: A brief review*, [Mod. Phys. Lett. A **27** \(2012\) 1230015](#) [[arXiv:1203.5250](#)].
- [81] P. Minkowski, *$\mu \rightarrow e\gamma$ at a Rate of One Out of 10^9 Muon Decays?*, [Phys. Lett. B **67** \(1977\) 421](#).

- [82] T. Yanagida, *Horizontal gauge symmetry and masses of neutrinos*, Conf. Proc. C **7902131** (1979) 95.
- [83] D. Z. Freedman, *Coherent effects of a weak neutral current*, **9** (1974) 1389.
- [84] D. Z. Freedman, *Coherent Neutrino Nucleus Scattering as a Probe of the Weak Neutral Current*, *Phys. Rev. D* **9** (1974) 1389.
- [85] J. Barranco, O. G. Miranda and T. I. Rashba, *Probing new physics with coherent neutrino scattering off nuclei*, *JHEP* **12** (2005) 021 [[arXiv:hep-ph/0508299](#)].
- [86] COHERENT collaboration, *Observation of Coherent Elastic Neutrino-Nucleus Scattering*, *Science* **357** (2017) 1123 [[arXiv:1708.01294](#)].
- [87] J. Colaresi, J. I. Collar, T. W. Hossbach, C. M. Lewis and K. M. Yocum, *Measurement of Coherent Elastic Neutrino-Nucleus Scattering from Reactor Antineutrinos*, *Phys. Rev. Lett.* **129** (2022) 211802 [[arXiv:2202.09672](#)].
- [88] N. Ackermann et al., *Direct observation of coherent elastic antineutrino–nucleus scattering*, *Nature* **643** (2025) 1229 [[arXiv:2501.05206](#)].
- [89] XENON collaboration, *First Indication of Solar $B8$ Neutrinos via Coherent Elastic Neutrino-Nucleus Scattering with XENONnT*, *Phys. Rev. Lett.* **133** (2024) 191002 [[arXiv:2408.02877](#)].
- [90] PandaX collaboration, *First Indication of Solar $B8$ Neutrinos through Coherent Elastic Neutrino-Nucleus Scattering in PandaX-4T*, *Phys. Rev. Lett.* **133** (2024) 191001 [[arXiv:2407.10892](#)].
- [91] J. N. Bahcall, E. Lisi, D. E. Alburger, L. De Braekeleer, S. J. Freedman and J. Napolitano, *Standard neutrino spectrum from B -8 decay*, *Phys. Rev. C* **54** (1996) 411 [[arXiv:nucl-th/9601044](#)].
- [92] J. Billard, L. Strigari and E. Figueroa-Feliciano, *Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments*, *Phys. Rev. D* **89** (2014) 023524 [[arXiv:1307.5458](#)].
- [93] M. Goeppert-Mayer, *Double beta-disintegration*, *Phys. Rev.* **48** (1935) 512.
- [94] R. Saakyan, *Two-Neutrino Double-Beta Decay*, *Ann. Rev. Nucl. Part. Sci.* **63** (2013) 503.
- [95] J. D. Vergados, H. Ejiri and F. Simkovic, *Theory of Neutrinoless Double Beta Decay*, *Rept. Prog. Phys.* **75** (2012) 106301 [[arXiv:1205.0649](#)].

- [96] K. S. Krane, *Introductory nuclear physics*. Wiley, New York, NY, 1988.
- [97] M. Doi and T. Kotani, *Neutrino emitting modes of double beta decay*, [Progress of Theoretical Physics](#) **87** (1992) 1207
[arXiv:<https://academic.oup.com/ptp/article-pdf/87/5/1207/5272207/87-5-1207.pdf>].
- [98] R. G. Winter, *Double k capture and single k capture with positron emission*, [Phys. Rev.](#) **100** (1955) 142.
- [99] M. Doi and T. Kotani, *Neutrino emitting modes of double beta decay*, [Prog. Theor. Phys.](#) **87** (1992) 1207.
- [100] E. Majorana, *Teoria simmetrica dell'elettrone e del positrone*, [Nuovo Cim.](#) **14** (1937) 171.
- [101] J. Kotila and F. Iachello, *Phase space factors for double- β decay*, [Phys. Rev. C](#) **85** (2012) 034316 [arXiv:1209.5722].
- [102] J. Engel and J. Menéndez, *Status and Future of Nuclear Matrix Elements for Neutrinoless Double-Beta Decay: A Review*, [Rept. Prog. Phys.](#) **80** (2017) 046301 [arXiv:1610.06548].
- [103] J. Menendez, A. Poves, E. Caurier and F. Nowacki, *Disassembling the Nuclear Matrix Elements of the Neutrinoless beta beta Decay*, [Nucl. Phys. A](#) **818** (2009) 139 [arXiv:0801.3760].
- [104] E. Caurier, F. Nowacki and A. Poves, *Nuclear Structure Aspects of the Neutrinoless Double Beta Decay*, [Eur. Phys. J. A](#) **36** (2008) 195 [arXiv:0709.0277].
- [105] M. Horoi, S. Stoica and B. A. Brown, *Shell-model calculations of two-neutrino double- β decay rates of ^{48}Ca with the *gpxf1a* interaction*, [Phys. Rev. C](#) **75** (2007) 034303.
- [106] P. Vogel and M. R. Zirnbauer, *Suppression of the Two Neutrino Double beta Decay by Nuclear Structure Effects*, [Phys. Rev. Lett.](#) **57** (1986) 3148.
- [107] J. Barea and F. Iachello, *Neutrinoless double- β decay in the microscopic interacting boson model*, [Phys. Rev. C](#) **79** (2009) 044301.
- [108] W. Buchmuller, R. D. Peccei and T. Yanagida, *Leptogenesis as the origin of matter*, [Ann. Rev. Nucl. Part. Sci.](#) **55** (2005) 311 [arXiv:hep-ph/0502169].

- [109] EXO-200 collaboration, *Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector*, *Phys. Rev. C* **89** (2014) 015502 [[arXiv:1306.6106](#)].
- [110] KamLAND-Zen collaboration, *Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen*, *Phys. Rev. Lett.* **117** (2016) 082503 [[arXiv:1605.02889](#)].
- [111] LEGEND collaboration, *Neutrinoless double beta decay search with LEGEND*, *PoS NOW2022* (2023) 075 [[arXiv:2501.10046](#)].
- [112] V. I. Tretyak and Y. G. Zdesenko, *Tables of Double Beta Decay DATA—AN Update*, *Atomic Data and Nuclear Data Tables* **80** (2002) 83.
- [113] J. Meija, T. B. Coplen, M. Berglund, W. A. Brand, P. D. Bièvre, M. Gröning et al., *Isotopic compositions of the elements 2013 (iupac technical report)*, *Pure and Applied Chemistry* **88** (2016) 293.
- [114] D. A. Nesterenko, K. Blaum, M. Block, C. Droese, S. Eliseev, F. Herfurth et al., *Double- β transformations in isobaric triplets with mass numbers $A = 124, 130$, and 136*, *Phys. Rev. C* **86** (2012) 044313.
- [115] XENON collaboration, *Observation of two-neutrino double electron capture in ^{124}Xe with XENON1T*, *Nature* **568** (2019) 532 [[arXiv:1904.11002](#)].
- [116] C. Wittweg, B. Lenardo, A. Fieguth and C. Weinheimer, *Detection prospects for the second-order weak decays of ^{124}Xe in multi-tonne xenon time projection chambers*, *Eur. Phys. J. C* **80** (2020) 1161 [[arXiv:2002.04239](#)].
- [117] F. G. and Kondev, , and S. Naimi, *The nubase2016 evaluation of nuclear properties*, *Chinese Physics C* **41** (2017) 030001.
- [118] XMASS collaboration, *Improved search for two-neutrino double electron capture on ^{124}Xe and ^{126}Xe using particle identification in XMASS-I*, *PTEP* **2018** (2018) 053D03 [[arXiv:1801.03251](#)].
- [119] EXO-200 collaboration, *Searches for double beta decay of ^{134}Xe with EXO-200*, *Phys. Rev. D* **96** (2017) 092001 [[arXiv:1704.05042](#)].
- [120] P. M. McCowan and R. C. Barber, *Q value for the double-beta decay of Xe-136* , *Phys. Rev. C* **82** (2010) 024603.
- [121] KamLAND-Zen collaboration, *Search for Majorana Neutrinos with the Complete KamLAND-Zen Dataset*, *Phys. Rev. Lett.* **135** (2025) 262501 [[arXiv:2406.11438](#)].

- [122] E. Aprile, K. Abe, F. Agostini, S. A. Maouloud, L. Althueser, B. Andrieu et al., *Search for new physics in electronic recoil data from XENONnT*, [Physical Review Letters](#) **129** (2022) .
- [123] LZ collaboration, *Two-neutrino double electron capture of ^{124}Xe in the first LUX-ZEPLIN exposure*, [J. Phys. G](#) **52** (2025) 015103 [[arXiv:2408.17391](#)].
- [124] PandaX-4T, PandaX collaboration, *Measurement of two-neutrino double electron capture half-life of ^{124}Xe with PandaX-4T*, [JHEP](#) **05** (2025) 119 [[arXiv:2411.14355](#)].
- [125] XENON collaboration, *Double-Weak Decays of ^{124}Xe and ^{136}Xe in the XENON1T and XENONnT Experiments*, [Phys. Rev. C](#) **106** (2022) 024328 [[arXiv:2205.04158](#)].
- [126] XENON100 collaboration, *Observation and applications of single-electron charge signals in the XENON100 experiment*, [J. Phys. G](#) **41** (2014) 035201 [[arXiv:1311.1088](#)].
- [127] XENON collaboration, *First Search for Light Dark Matter in the Neutrino Fog with XENONnT*, [Phys. Rev. Lett.](#) **134** (2025) 111802 [[arXiv:2409.17868](#)].
- [128] LZ collaboration, *Dark Matter Search Results from 4.2 Tonne-Years of Exposure of the LUX-ZEPLIN (LZ) Experiment*, [Phys. Rev. Lett.](#) **135** (2025) 011802 [[arXiv:2410.17036](#)].
- [129] S.-C. Hwang et al., *Noble gases*, in *Kirk-Othmer Encyclopedia of Chemical Technology*, John Wiley & Sons, Ltd, (2005), [DOI](#).
- [130] W. R. Leo, *Techniques for Nuclear and Particle Physics Experiments: A How to Approach*. Springer Berlin, Heidelberg, 1987.
- [131] P. J. Linstrom and W. G. Mallard, *Nist chemistry webbook*, NIST Standard Reference Database Number 69, 2023.
- [132] National Institute of Standards and Technology, *Xcom: Photon cross sections database*, <https://physics.nist.gov/cgi-bin/Xcom/xcom2?Method=Elem&Output2=Hand>, 2025.
- [133] M. Miyajima, T. Takahashi, S. Konno, T. Hamada, S. Kubota, H. Shibamura et al., *Average energy expended per ion pair in liquid argon*, [Phys. Rev. A](#) **9** (1974) 1438.

- [134] E. Shibamura, A. Hitachi, T. Doke, T. Takahashi, S. Kubota and M. Miyajima, *Drift velocities of electrons, saturation characteristics of ionization and w-values for conversion electrons in liquid argon, liquid argon-gas mixtures and liquid xenon*, [Nuclear Instruments and Methods](#) **131** (1975) 249.
- [135] E. Aprile and T. Doke, *Liquid xenon detectors for particle physics and astrophysics*, [Rev. Mod. Phys.](#) **82** (2010) 2053.
- [136] F. Theeuwes and R. J. Bearman, *The p , v , t behavior of dense fluids v. the vapor pressure and saturated liquid density of xenon*, [The Journal of Chemical Thermodynamics](#) **2** (1970) 507.
- [137] E. W. Lemmon, R. T. Jacobsen, S. G. Penoncello and D. G. Friend, *Thermodynamic properties of air and mixtures of nitrogen, argon, and oxygen from 60 to 2000 K at pressures to 2000 MPa*, [Journal of Physical and Chemical Reference Data](#) **29** (2000) 331 [[arXiv:https://pubs.aip.org/aip/jpr/article-pdf/29/3/331/12202221/331_1online.pdf](https://pubs.aip.org/aip/jpr/article-pdf/29/3/331/12202221/331_1online.pdf)].
- [138] D. R. Lide, *CRC Handbook of Chemistry and Physics, 88th Edition*, vol. 130. CRC Press, 2008, [10.1021/ja077011d](#).
- [139] R. Firestone, V. Shirley, C. Baglin, S. Chu and J. Zipkin, *The 8th edition of the Table of Isotopes*. Springer, 1997.
- [140] R. C. Kemp, W. R. G. Kemp and P. W. Smart, *The triple point of xenon as a possible defining point on an international temperature scale*, [Metrologia](#) **21** (1985) 43.
- [141] C. E. Dahl, *The physics of background discrimination in liquid xenon, and first results from Xenon10 in the hunt for WIMP dark matter*, Ph.D. thesis, Princeton U. (2009).
- [142] XENON collaboration, *The XENONnT dark matter experiment*, [Eur. Phys. J. C](#) **84** (2024) 784 [[arXiv:2402.10446](#)].
- [143] LZ collaboration, *The LUX-ZEPLIN (LZ) Experiment*, [Nucl. Instrum. Meth. A](#) **953** (2020) 163047 [[arXiv:1910.09124](#)].
- [144] B. Lenardo, K. Kazkaz, A. Manalaysay, J. Mock, M. Szydagis and M. Tripathi, *A Global Analysis of Light and Charge Yields in Liquid Xenon*, [IEEE Trans. Nucl. Sci.](#) **62** (2015) 3387 [[arXiv:1412.4417](#)].
- [145] K. Fujii, Y. Endo, Y. Torigoe, S. Nakamura, T. Haruyama, K. Kasami et al., *High-accuracy measurement of the emission spectrum of liquid xenon in the*

- vacuum ultraviolet region*, [Nuclear Instruments and Methods in Physics Research A](#) **795** (2015) 293.
- [146] A. Hitachi, T. Takahashi, N. Funayama, K. Masuda, J. Kikuchi and T. Doke, *Effect of ionization density on the time dependence of luminescence from liquid argon and xenon*, [Phys. Rev. B](#) **27** (1983) 5279.
- [147] ArDM collaboration, *Backgrounds and pulse shape discrimination in the ArDM liquid argon TPC*, [JCAP](#) **12** (2018) 011 [[arXiv:1712.01932](#)].
- [148] M. Szydagis, *Noble element simulation technique*, Zenodo, Jan., 2023, 10.5281/zenodo.7577399.
- [149] M. Szydagis et al., *A review of NEST models for liquid xenon and an exhaustive comparison with other approaches*, [Front. Detect. Sci. Tech.](#) **2** (2024) 1480975 [[arXiv:2211.10726](#)].
- [150] L. Baudis, P. Sanchez-Lucas and K. Thieme, *A measurement of the mean electronic excitation energy of liquid xenon*, [Eur. Phys. J. C](#) **81** (2021) 1060 [[arXiv:2109.07151](#)].
- [151] EXO-200 collaboration, *Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment*, [Phys. Rev. C](#) **101** (2020) 065501 [[arXiv:1908.04128](#)].
- [152] E. Aprile, C. E. Dahl, L. DeViveiros, R. Gaitskell, K. L. Giboni, J. Kwong et al., *Simultaneous measurement of ionization and scintillation from nuclear recoils in liquid xenon as target for a dark matter experiment*, [Phys. Rev. Lett.](#) **97** (2006) 081302 [[arXiv:astro-ph/0601552](#)].
- [153] XENON collaboration, *Emission of single and few electrons in xenon1t and limits on light dark matter*, [Phys. Rev. D](#) **106** (2022) 022001.
- [154] N. Ishida et al., *Attenuation length measurements of scintillation light in liquid rare gases and their mixtures using an improved reflection suppresser*, [Nucl. Instrum. Meth. A](#) **384** (1997) 380.
- [155] E. M. Gushchin, A. A. Kruglov and I. M. Obodovskii, *Electron dynamics in condensed argon and xenon*, [Sov. Phys. JETP](#) **55** (1982) 650.
- [156] XENON collaboration, *XENONnT analysis: Signal reconstruction, calibration, and event selection*, [Phys. Rev. D](#) **111** (2025) 062006 [[arXiv:2409.08778](#)].

- [157] E. M. Gushchin, A. A. Kruglov, V. V. Litskevich et al., *Electron emission from condensed noble gases*, Soviet Physics JETP **49** (1979) 856.
- [158] S. Farrell, C. Tunnell, J. R. Angevaare, N. Carrara, G. R. Rischbieter, mszydagis et al., *Nestcollaboration/nestpy: Update to v2.0.2, sync with nest v2.4.0*, Zenodo, Jan., 2024, 10.5281/zenodo.10582363.
- [159] XENON collaboration, *XENON1T Dark Matter Data Analysis: Signal Reconstruction, Calibration and Event Selection*, Phys. Rev. D **100** (2019) 052014 [[arXiv:1906.04717](#)].
- [160] LUX collaboration, *First results from the LUX dark matter experiment at the Sanford Underground Research Facility*, Phys. Rev. Lett. **112** (2014) 091303 [[arXiv:1310.8214](#)].
- [161] PandaX-II collaboration, *Dark Matter Results From 54-Ton-Day Exposure of PandaX-II Experiment*, Phys. Rev. Lett. **119** (2017) 181302 [[arXiv:1708.06917](#)].
- [162] MACRO collaboration, *Measurement of the residual energy of muons in the Gran Sasso underground laboratories*, Astropart. Phys. **19** (2003) 313 [[arXiv:hep-ex/0207043](#)].
- [163] PandaX collaboration, *PandaX-xT— A deep underground multi-ten-tonne liquid xenon observatory*, Sci. China Phys. Mech. Astron. **68** (2025) 221011 [[arXiv:2402.03596](#)].
- [164] XENON collaboration, *The neutron veto of the XENONnT experiment: results with demineralized water*, Eur. Phys. J. C **85** (2025) 695 [[arXiv:2412.05264](#)].
- [165] M.-M. Bé, V. Chisté, C. Dulieu, E. Browne, V. Chechev, N. Kuzmenko et al., *Table of radionuclides (Vol. 1 - A = 1 to 150)*, vol. Vol. 1 - A = 1 to 150 of *Table of radionuclides*. Bureau International des Poids et Mesures, 2004.
- [166] XMASS collaboration, *Distillation of Liquid Xenon to Remove Krypton*, Astropart. Phys. **31** (2009) 290 [[arXiv:0809.4413](#)].
- [167] XENON collaboration, *Design and Performance of the XENON10 Dark Matter Experiment*, Astropart. Phys. **34** (2011) 679 [[arXiv:1001.2834](#)].
- [168] E. Aprile, K. Arisaka, F. Arneodo, A. Askin, L. Baudis, A. Behrens et al., *The xenon100 dark matter experiment*, Astroparticle Physics **35** (2012) 573.
- [169] XENON collaboration, *Lowering the radioactivity of the photomultiplier tubes*

- for the XENON1T dark matter experiment, *Eur. Phys. J. C* **75** (2015) 546 [[arXiv:1503.07698](#)].
- [170] LVD collaboration, *Muon ‘Depth intensity’ relation measured by LVD underground experiment and cosmic ray muon spectrum at sea level*, *Phys. Rev. D* **58** (1998) 092005 [[arXiv:hep-ex/9806001](#)].
- [171] XENON collaboration, *The XENON1T Dark Matter Experiment*, *Eur. Phys. J. C* **77** (2017) 881 [[arXiv:1708.07051](#)].
- [172] XENONnT collaboration, *Design and performance of the field cage for the XENONnT experiment*, *Eur. Phys. J. C* **84** (2024) 138 [[arXiv:2309.11996](#)].
- [173] L. Althüser, S. Lindemann, M. Murra, M. Schumann, C. Wittweg and C. Weinheimer, *VUV Transmission of PTFE for Xenon-based Particle Detectors*, *JINST* **15** (2020) P12021 [[arXiv:2006.05827](#)].
- [174] V. C. Antochi et al., *Improved quality tests of R11410-21 photomultiplier tubes for the XENONnT experiment*, *JINST* **16** (2021) P08033 [[arXiv:2104.15051](#)].
- [175] E. Aprile, K. Abe, F. Agostini, S. Ahmed Maouloud, M. Alfonsi, L. Althueser et al., *Material radiopurity control in the xenonnt experiment*, *The European Physical Journal C* **82** (2022) .
- [176] Iwatani Industrial Gas Co., *Model pc150*, <https://www.iwatani.co.jp/eng/>.
- [177] SAES, *Ps4-mt series rare gas purifier, model ps4-mt50-r*, http://www.saespuregas.com/Library/specifications-brochures/s110-233_a_521.pdf.
- [178] E. Brown et al., *Magnetically-coupled piston pump for high-purity gas applications*, *Eur. Phys. J. C* **78** (2018) 604 [[arXiv:1803.08498](#)].
- [179] Barber Nichols, *Pump model bnep-32c-000*, <https://barber-nichols.com>.
- [180] BASF, *Cu-0226s (cuo on al₂o₃)*, <https://catalysts.basf.com/products/cu-0226-s-1>.
- [181] XENON collaboration, *Radon Removal in XENONnT down to the Solar Neutrino Level*, *Phys. Rev. X* **15** (2025) 031079 [[arXiv:2502.04209](#)].
- [182] International Atomic Energy Agency, *Database of Prompt Gamma Rays from Slow Neutron Capture for Elemental Analysis*. IAEA, Vienna, 2007.
- [183] Y. Koshio, M. Nakahata, H. Sekiya and M. R. Vagins, *Upgrade of Super-Kamiokande with Gadolinium*, [arXiv:2511.03921](#).

- [184] Hamamatsu Photonics K.K., *Photomultiplier Model R5912-100-10 WA-D30 SEL Assy*. Hamamatsu Photonics.
- [185] 3M, *3MTM Specular Film DF2000MA Technical Data Sheet*. 3M, 2025.
- [186] Hamamatsu Photonics K.K., *Photomultiplier Model R5912ASSY*. Hamamatsu Photonics.
- [187] C. Geis, C. Grignon, U. Oberlack, D. Ramírez García and Q. Weitzel, *Optical response of highly reflective film used in the water Cherenkov muon veto of the XENON1T dark matter experiment*, [JINST **12** \(2017\) P06017](#) [[arXiv:1706.03687](#)].
- [188] XENON collaboration, *Low-Energy Nuclear Recoil Calibration of XENONnT with a ⁸⁸YBe Photoneutron Source*, [arXiv:2412.10451](#).
- [189] A. Manalaysay et al., *Spatially uniform calibration of a liquid xenon detector at low energies using 83m-Kr*, [Rev. Sci. Instrum. **81** \(2010\) 073303](#) [[arXiv:0908.0616](#)].
- [190] F. Jörg, S. Li, J. Schreiner, H. Simgen and R. F. Lang, *Characterization of a ²²⁰Rn source for low-energy electronic recoil calibration of the XENONnT detector*, [JINST **18** \(2023\) P11009](#) [[arXiv:2306.05673](#)].
- [191] R. F. Lang, A. Brown, E. Brown, M. Cervantes, S. Macmullin, D. Masson et al., *A ²²⁰Rn source for the calibration of low-background experiments*, [JINST **11** \(2016\) P04004](#) [[arXiv:1602.01138](#)].
- [192] XENON collaboration, *Low-energy calibration of XENON1T with an internal ³⁷Ar source*, [Eur. Phys. J. C **83** \(2023\) 542](#) [[arXiv:2211.14191](#)].
- [193] XENON collaboration, *Light Dark Matter Search with 7.8 Tonne-Year of Ionization-Only Data in XENONnT*, [arXiv:2601.11296](#).
- [194] XENON collaboration, *Neutrinoless double beta decay search in XENON1T and XENONnT*, [PoS ICRC2021 \(2021\) 540](#).
- [195] L. Baudis, A. Ferella, A. Kish, A. Manalaysay, T. Marrodan Undagoitia and M. Schumann, *Neutrino physics with multi-ton scale liquid xenon detectors*, [JCAP **01** \(2014\) 044](#) [[arXiv:1309.7024](#)].
- [196] P. Antonioli et al., *SNEWS: The Supernova Early Warning System*, [New J. Phys. **6** \(2004\) 114](#) [[arXiv:astro-ph/0406214](#)].

- [197] P. Barrow et al., *Qualification Tests of the R11410-21 Photomultiplier Tubes for the XENON1T Detector*, [JINST **12** \(2017\) P01024](#) [[arXiv:1609.01654](#)].
- [198] L. Baudis, A. Behrens, A. Ferella, A. Kish, T. Marrodan Undagoitia, D. Mayani et al., *Performance of the Hamamatsu R11410 Photomultiplier Tube in cryogenic Xenon Environments*, [JINST **8** \(2013\) P04026](#) [[arXiv:1303.0226](#)].
- [199] C. H. Faham, V. M. Gehman, A. Currie, A. Dobi, P. Sorensen and R. J. Gaitskell, *Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes*, [JINST **10** \(2015\) P09010](#) [[arXiv:1506.08748](#)].
- [200] Hamamatsu Photonics, *High energy photomultiplier tubes*, https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/High_energy_PMT_TPMZ0003E.pdf, 2023.
- [201] R. Saldanha, L. Grandi, Y. Guardincerri and T. Wester, *Model Independent Approach to the Single Photoelectron Calibration of Photomultiplier Tubes*, [Nucl. Instrum. Meth. A **863** \(2017\) 35](#) [[arXiv:1602.03150](#)].
- [202] A. Bismark, *Tests of the Fundamentals of Quantum Mechanics with Low-Background Experiments and Detector Technologies for Liquid Xenon Time Projection Chambers*, Ph.D. thesis, University of Zurich (2024).
- [203] Y. Mosbacher, L. Sanchez, S. Shi, D. Xu, A. Ravindran, G. Volta et al., *XENONnT/xedocs: V0.2.36*, Zenodo, Jan., 2025, 10.5281/zenodo.14737183.
- [204] K. Chodorow et al., *MongoDB: The Definitive Guide*. O'Reilly Media, Inc., 1st ed., Aug., 2010.
- [205] J. Aalbers, *AxFoundation/strax: v0.11.1*, June, 2020, 10.5281/zenodo.3894675.
- [206] J. R. e. a. Angevaare, *XENONnT/straxen: V2.0.4*, Zenodo, 10.5281/zenodo.7541292.
- [207] T. Matsubara et al., *Evaluation of 400 low background 10-in. photo-multiplier tubes for the Double Chooz experiment*, [Nucl. Instrum. Meth. A **661** \(2012\) 16](#) [[arXiv:1104.0786](#)].
- [208] XENON collaboration, *XENONnT WIMP search: Signal and background modeling and statistical inference*, [Phys. Rev. D **111** \(2025\) 103040](#) [[arXiv:2406.13638](#)].

- [209] XENON collaboration, *Excess electronic recoil events in XENON1T*, [Phys. Rev. D **102** \(2020\) 072004](#) [[arXiv:2006.09721](#)].
- [210] XENON collaboration, *First Dark Matter Search with Nuclear Recoils from the XENONnT Experiment*, [Phys. Rev. Lett. **131** \(2023\) 041003](#) [[arXiv:2303.14729](#)].
- [211] F. Toschi, *Design of the Field Cage and Charge Response of the XENONnT Dark Matter Experiment*, Ph.D. thesis, Albert-Ludwigs-Universität Freiburg (2022).
- [212] F. Jörg, D. Cichon, G. Eurin, L. Höttsch, T. Undagoitia Marrodán and N. Rupp, *Characterization of alpha and beta interactions in liquid xenon*, [Eur. Phys. J. C **82** \(2022\) 361](#) [[arXiv:2109.13735](#)].
- [213] LUX collaboration, *Investigation of background electron emission in the LUX detector*, [Phys. Rev. D **102** \(2020\) 092004](#) [[arXiv:2004.07791](#)].
- [214] A. Kopec, A. L. Baxter, M. Clark, R. F. Lang, S. Li, J. Qin et al., *Correlated single- and few-electron backgrounds milliseconds after interactions in dual-phase liquid xenon time projection chambers*, [JINST **16** \(2021\) P07014](#) [[arXiv:2103.05077](#)].
- [215] XENON collaboration, *Projected WIMP sensitivity of the XENONnT dark matter experiment*, [JCAP **11** \(2020\) 031](#) [[arXiv:2007.08796](#)].
- [216] D. Ramírez García, *Simulating the XENONnT dark matter experiment: backgrounds and WIMP sensitivity*, Ph.D. thesis, Albert-Ludwigs-Universität Freiburg (2022). 10.6094/UNIFR/228338.
- [217] SNO collaboration, *Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory*, [Phys. Rev. C **88** \(2013\) 025501](#) [[arXiv:1109.0763](#)].
- [218] PandaX collaboration, *Search for Solar B8 Neutrinos in the PandaX-4T Experiment Using Neutrino-Nucleus Coherent Scattering*, [Phys. Rev. Lett. **130** \(2023\) 021802](#) [[arXiv:2207.04883](#)].
- [219] XENON collaboration, *Search for Coherent Elastic Scattering of Solar 8B Neutrinos in the XENON1T Dark Matter Experiment*, [Phys. Rev. Lett. **126** \(2021\) 091301](#) [[arXiv:2012.02846](#)].
- [220] XENON collaboration, *Search for New Physics in Electronic Recoil Data from XENONnT*, [Phys. Rev. Lett. **129** \(2022\) 161805](#) [[arXiv:2207.11330](#)].

- [221] J. N. Bahcall and C. Pena-Garay, *Solar models and solar neutrino oscillations*, [New J. Phys.](#) **6** (2004) 63 [[arXiv:hep-ph/0404061](#)].
- [222] W. C. Haxton, R. G. Hamish Robertson and A. M. Serenelli, *Solar Neutrinos: Status and Prospects*, [Ann. Rev. Astron. Astrophys.](#) **51** (2013) 21 [[arXiv:1208.5723](#)].
- [223] J.-W. Chen, H.-C. Chi, C. P. Liu and C.-P. Wu, *Low-energy electronic recoil in xenon detectors by solar neutrinos*, [Phys. Lett. B](#) **774** (2017) 656 [[arXiv:1610.04177](#)].
- [224] PandaX collaboration, *Searching for Two-Neutrino and Neutrinoless Double Beta Decay of $Xe134$ with the PandaX-4T Experiment*, [Phys. Rev. Lett.](#) **132** (2024) 152502 [[arXiv:2312.15632](#)].
- [225] PandaX collaboration, *Searching for neutrinoless double-beta decay of ^{136}Xe with PandaX-4T*, [Sci. Bull.](#) **70** (2025) 1779 [[arXiv:2412.13979](#)].
- [226] LZ collaboration, *Projected sensitivity of the LUX-ZEPLIN experiment to the $0\nu\beta\beta$ decay of ^{136}Xe* , [Phys. Rev. C](#) **102** (2020) 014602 [[arXiv:1912.04248](#)].
- [227] LZ collaboration, *Projected sensitivity of the LUX-ZEPLIN experiment to the two-neutrino and neutrinoless double β decays of ^{134}Xe* , [Phys. Rev. C](#) **104** (2021) 065501 [[arXiv:2104.13374](#)].
- [228] XENON collaboration, *Spectral Measurement of the ^{214}Bi beta-decay to the ^{214}Po Ground State with XENONnT*, [arXiv:2510.04846](#).
- [229] XENON1T collaboration, *Conceptual design and simulation of a water Cherenkov muon veto for the XENON1T experiment*, [JINST](#) **9** (2014) P11006 [[arXiv:1406.2374](#)].
- [230] OSG, *Open science data federation*, OSG, 2015, 10.21231/0KVZ-VE57.
- [231] R. Pordes, D. Petravick, B. Kramer, D. Olson, M. Livny, A. Roy et al., *The open science grid*, in *J. Phys. Conf. Ser.*, vol. 78 of 78, p. 012057, 2007, [DOI](#).
- [232] Geant4 collaboration, *Geant4: A Simulation toolkit*, [Nucl. Instrum. Meth. A](#) **506** (2003) 250.
- [233] Geant4 collaboration, *Recent developments in Geant4*, [Nucl. Instrum. Meth. A](#) **835** (2016) 186.
- [234] P. Sorensen, *Anisotropic diffusion of electrons in liquid xenon with application to*

- improving the sensitivity of direct dark matter searches*, *Nucl. Instrum. Meth. A* **635** (2011) 41 [[arXiv:1102.2865](#)].
- [235] COMSOL AB, *COMSOL Multiphysics® v.5.4 User's Guide*. COMSOL AB, Stockholm, Sweden.
- [236] F. Toschi, *PyCOMes - Package to load and handle electric field data coming from electrostatic simulations in COMSOL*, <https://github.com/ftoschi/PyCOMes>.
- [237] J. B. A. et al. (EXO-200 Collaboration), *Measurement of the drift velocity and transverse diffusion of electrons in liquid xenon with the exo-200 detector*, *Phys. Rev. C* **95** (2017) 025502.
- [238] Garfield++ Collaboration, *Garfield++*, <https://garfieldpp.web.cern.ch/garfieldpp/>.
- [239] M. Ester, H.-P. Kriegel, J. Sander and X. Xu, *A density-based algorithm for discovering clusters in large spatial databases with noise*, in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD'96)*, pp. 226–231, AAAI Press, 1996.
- [240] M. Szydagis et al., *A Review of Basic Energy Reconstruction Techniques in Liquid Xenon and Argon Detectors for Dark Matter and Neutrino Physics Using NEST*, *Instruments* **5** (2021) 13 [[arXiv:2102.10209](#)].
- [241] G. Knoll, *Radiation Detection and Measurement (4th ed.)*. John Wiley, Hoboken, NJ, 2010.
- [242] E. Aprile, J. Aalbers, F. Agostini et al., *Energy resolution and linearity of xenon1t in the mev energy range*, *The European Physical Journal C* **80** (2020) 785.
- [243] N. Barros, J. Thurn and K. Zuber, *Double beta decay searches of ^{134}Xe , ^{126}Xe and ^{124}Xe with large scale Xe detectors*, *J. Phys. G* **41** (2014) 115105 [[arXiv:1409.8308](#)].
- [244] R. D. Deslattes, E. G. Kessler, P. Indelicato, L. de Billy, E. Lindroth and J. Anton, *X-ray transition energies: new approach to a comprehensive evaluation*, *Rev. Mod. Phys.* **75** (2003) 35.
- [245] O. A. Ponkratenko, V. I. Tretyak and Y. G. Zdesenko, *The Event generator DECAY4 for simulation of double beta processes and decay of radioactive nuclei*, *Phys. Atom. Nucl.* **63** (2000) 1282 [[arXiv:nucl-ex/0104018](#)].

- [246] M. Be, V. Chiste, C. Dulieu, E. Browne, V. Chechev, N. Kuzmenko et al., *Table of radionuclides (Vol. 4 - A = 133 to 252)*. Editions BIPM, 2008.
- [247] XENON collaboration, *Material radioassay and selection for the XENON1T dark matter experiment*, [Eur. Phys. J. C **77** \(2017\) 890](#) [[arXiv:1705.01828](#)].
- [248] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, [Eur. Phys. J. C **71** \(2011\) 1554](#) [[arXiv:1007.1727](#)].
- [249] C. Bishop, *Pattern Recognition and Machine Learning*, vol. 16, pp. 140–155. 01, 2006. 10.1117/1.2819119.
- [250] L. Breiman, J. Friedman, R. A. Olshen and C. J. Stone, *Classification and Regression Trees*. Chapman and Hall/CRC, 10, 2017, [10.1201/9781315139470](#).
- [251] Y. Coadou, *Boosted Decision Trees and Applications*, [EPJ Web Conf. **55** \(2013\) 02004](#).
- [252] J. H. Friedman, *Greedy function approximation: A gradient boosting machine.*, [Annals Statist. **29** \(2001\) 1189](#).
- [253] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel et al., *Scikit-learn: Machine learning in python*, *Journal of Machine Learning Research* **12** (2011) 2825.
- [254] A. S. Cornell, W. Doorsamy, B. Fuks, G. Harmsen and L. Mason, *Boosted decision trees in the era of new physics: a smuon analysis case study*, [JHEP **04** \(2022\) 015](#) [[arXiv:2109.11815](#)].
- [255] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel et al., *Machine learning at the energy and intensity frontiers of particle physics*, [Nature **560** \(2018\) 41](#).
- [256] J. Aalbers, F. Agostini, M. Alfonsi, F. Amaro, C. AMSler, E. Aprile et al., *Darwin: Towards the ultimate dark matter detector*, [Journal of Cosmology and Astroparticle Physics **2016** \(2016\)](#) .
- [257] PandaX-4T collaboration, *Dark Matter Search Results from the PandaX-4T Commissioning Run*, [Phys. Rev. Lett. **127** \(2021\) 261802](#) [[arXiv:2107.13438](#)].
- [258] L. Baudis and et al., *Design and construction of Xenoscope — a full-scale vertical demonstrator for the DARWIN observatory*, [JINST **16** \(2021\) P08052](#) [[arXiv:2105.13829](#)].

- [259] SNEWS collaboration, *SNEWS 2.0: a next-generation supernova early warning system for multi-messenger astronomy*, *New J. Phys.* **23** (2021) 031201 [[arXiv:2011.00035](#)].
- [260] L. Baudis, *DARWIN/XLZD: A future xenon observatory for dark matter and other rare interactions*, *Nucl. Phys. B* **1003** (2024) 116473 [[arXiv:2404.19524](#)].
- [261] F. Girard, *Design, Construction, and Operation of Xenoscope and Photosensor Characterisation for the DARWIN Observatory*, Ph.D. thesis, University of Zurich (2023).
- [262] A. Brown et al., *PANCAKE: a large-diameter cryogenic test platform with a flat floor for next generation multi-tonne liquid xenon detectors*, *JINST* **19** (2024) P05018 [[arXiv:2312.14785](#)].
- [263] L. Baudis et al., *Electron transport measurements in liquid xenon with Xenoscope, a large-scale DARWIN demonstrator*, *Eur. Phys. J. C* **83** (2023) 717 [[arXiv:2303.13963](#)].
- [264] K. Thieme, *The Low-Energy and Large-Scale Frontier of Dual-Phase Xenon Time Projection Chambers for Dark Matter Search*, Ph.D. thesis, Universität Zürich (2022).
- [265] Entegris, Inc., *Gas purifiers — entegris*, <https://www.entegris.com/shop/en/USD/Products/Gas-Filtration-and-Purification/Gas-Purifiers/c/gaspurifiers>, 2025.
- [266] Kelvion GmbH, *Kelvion — heat exchanger solutions*, <https://www.kelvion.com>.
- [267] Cryo-Con, *Model 32 cryogenic temperature controller*, <https://www.cryocon.com>.
- [268] Pfeiffer Vacuum Technology AG, *Pfeiffer vacuum — switzerland / pfeiffer-vacuum schweiz*, <https://www.pfeiffer-vacuum.com/ch/de/>.
- [269] KNF Group, *Knf switzerland — pumps & solutions*, <https://knf.com/en/ch>.
- [270] Teledyne Hastings Instruments, *Teledyne hi — vacuum & flow instruments*, <https://www.teledyne-hi.com/>.
- [271] P. Jamshidi, C. Pahl, N. C. Mendonça, J. Lewis and S. Tilkov, *Microservices: The journey so far and challenges ahead*, *IEEE Software* **35** (2018) 24.

- [272] E. M. Gushchin, A. A. Kruglov and I. M. Obodovskii, *Electron dynamics in condensed argon and xenon*, Sov. Phys. JETP **55** (1982) 650.
- [273] L. Baudis, Y. Biondi, C. Capelli, M. Galloway, S. Kazama, A. Kish et al., *A Dual-phase Xenon TPC for Scintillation and Ionisation Yield Measurements in Liquid Xenon*, Eur. Phys. J. C **78** (2018) 351 [[arXiv:1712.08607](#)].
- [274] D. Baur et al., *The XeBRA platform for liquid xenon time projection chamber development*, JINST **18** (2023) T02004 [[arXiv:2208.14815](#)].
- [275] O. Njoya et al., *Measurements of electron transport in liquid and gas Xenon using a laser-driven photocathode*, Nucl. Instrum. Meth. A **972** (2020) 163965 [[arXiv:1911.11580](#)].
- [276] E. Hogenbirk, M. Decowski, K. McEwan and A. Colijn, *Field dependence of electronic recoil signals in a dual-phase liquid xenon time projection chamber*, JINST **13** (2018) P10031.
- [277] B. Rebel et al., *High Voltage in Noble Liquids for High Energy Physics*, JINST **9** (2014) T08004 [[arXiv:1403.3613](#)].
- [278] L. Tvrznikova, E. Bernard, S. Kravitz, K. O’Sullivan, G. Richardson, Q. Riffard et al., *Direct comparison of high voltage breakdown measurements in liquid argon and liquid xenon*, Journal of Instrumentation **14** (2019) P12018.
- [279] J. Watson, I. Olcina, J. Soria, D. N. McKinsey, S. Kravitz, E. E. Deck et al., *Study of dielectric breakdown in liquid xenon with xebra: The xenon breakdown apparatus*, Review of Scientific Instruments **94** (2023) .
- [280] G. F. Steib and E. Moll, *High-voltage conditioning at large gaps in industrial vacuum*, Journal of Physics D: Applied Physics **6** (1973) 243.
- [281] D. C. Faircloth, *Technological Aspects: High Voltage*, in CAS - CERN Accelerator School: Ion Sources, pp. 381–419, 2013, DOI [[arXiv:1404.0952](#)].
- [282] Heinzinger, *Datasheet: Shielded hv cables in various voltage classes up to 200 kv*, https://www.heinzinger.com/assets/uploads/downloads/Datenblatt_HVC_ENG_2024.pdf.
- [283] D. AG, *Conduit 18 mm acier galvanisé métallique “flexa”, produit no. 11011111018*, <https://www.distrelec.ch/fr/conduit-18mm-acier-galvanise-metallique-flexa-11011111018/p/15514228>.

- [284] K. J. L. Company, *Torr seal low vapor pressure epoxy* [product page], <https://www.lesker.com/newweb/fluids/sealants-epoxy/torrseal/>, 2025.
- [285] CeramTec, *Datasheet: Advanced Ceramic-To-Metal and Glass-Ceramic Sealing Technology*, https://www.ceramtec-group.com/fileadmin/user_upload/Corporate/11_Downloads/05_Ceramaseal/CeramTec_us_ceramaseal_feedthroughs.pdf.
- [286] H. P. K.K., “60 w xenon flash lamp series l7684, l7685 / l6604, l6605 – built-in reflective mirror” [datasheet], https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/L7684_L6604_TLSX1029E.pdf, Feb., 2006.
- [287] L. C. Ltd, *Fibre optic cable assembly components & accessories*, <https://www.lewvac.co.uk/product/fibre-optic-cable-assembly-components-accessories/>, 2025.
- [288] I. Thorlabs, *0.22 na solarization-resistant step index multimode optical fibers (na 0.22)* [product web page], https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=6840, 2013.
- [289] Quorum Technologies Ltd., *Q150t plus — turbomolecular pumped coater (turbomolecular pumped coater: Sputter and evaporation)* [product brochure], https://www.emme3-srl.it/wp-content/uploads/2024/04/Q150T-Plus_Quorum.pdf, 2018.
- [290] Y. Biondi, *Sensitivity of DARWIN to Rare Events and the Purity Monitor for Xenoscope*, Ph.D. thesis, University of Zurich (2022).
- [291] O. Njoya, T. Tsang, M. Tarka, W. Fairbank, K. Kumar, T. Rao et al., *Measurements of electron transport in liquid and gas xenon using a laser-driven photocathode*, [Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment](#) **972** (2020) 163965.
- [292] C. S.p.A., *V1724 – 8 channel 14-bit 100 ms/s digitizer*, <https://www.caen.it/products/v1724/>, 2025.
- [293] R. Peres, *Pylars v0.3.1*, 2024, 10.5281/zenodo.13756802.

Acknowledgments

Emprender un doctorado se parece, en muchos sentidos, a escalar una montaña. Es un trayecto largo e incierto, lleno de esfuerzo silencioso y de momentos en los que uno se pregunta por qué empezó. La cima, cuando llega, es breve, pero lo que realmente transforma es el camino. Para mí, este proceso estuvo además marcado por la experiencia de estar lejos, muy lejos, de casa. Migrar implica un desprendimiento profundo, físico y también emocional. Es, en cierta forma, como arrancar un árbol de raíz para plantarlo en otra tierra. Ese desarraigo duele y desorienta, pero también obliga a crecer de nuevas formas, a mirar lo que antes era propio desde la distancia y a reconstruirse con mayor conciencia. Y esa posibilidad, aunque exigente, es profundamente enriquecedora. Este proceso no habría sido posible sin las personas que, de una u otra forma, acompañaron ese trasplante y ayudaron a que ese árbol volviera a echar raíces y crecer.

I would like to thank Prof. Dr. Laura Baudis for inviting me to join her group, for her trust and flexibility throughout these years, and for bringing together people from different backgrounds and cultures, creating a stimulating and enriching place to work. Thank you also for the time you dedicated to reading and commenting on this thesis, and for the discussions and encouragement along the way. Thank you to Prof. Dr. Patrick Decowski and Prof. Dr. Teresa Marrodán Undagoitia for carefully reading this thesis and for their insightful comments which helped improve its quality. A big thanks to Dr. Luisa Höttsch, Dr. Chiara Capelli, Dr. Christian Wittweg, Dr. Jose Cuenca, and Dr. Diego Ramírez for kindly taking the time to read parts of this thesis. In particular, a special thanks to Dr. Diego Ramírez García, for the generous guidance, the constant support and the companionship throughout most of the projects presented in this thesis.

Thank you to all the members of the Astroparticle Physics group. Being part of such a diverse group, with people from many different backgrounds and perspectives, has certainly been a valuable experience from which I have learned a great deal. Thank you to Alex, Ricardo, and Jose for making the Y36 K 38 a genuinely *cool office*.

Thanks to the XENON Collaboration for providing an exciting scientific environment to work in, as well as for the many enjoyable moments during in-person meetings. Thanks to the high-energy group, specially to Dr. Maxime Pierre for the fruitful discussions, and to Carlo, Saad and Henning for the valuable feedback over the last year.

Gracias a Diego, esta vez por todo lo que no cabe arriba. No sólo acompañaste el trasplante del arbolito que llegó a Zúrich hace casi cinco años: te tomaste el tiempo de averiguar cuánta luz y agua necesitaba y estuviste atento cuando algo no iba bien. Gracias por tu dedicación y atención incluso después de haber dejado el grupo, y por estar durante los meses más difíciles de escritura. Mucho de lo que aprendí estos años lleva la huella de nuestras conversaciones y del trabajo compartido. Gracias siempre.

Gracias a todos los *Mexicanos en Irchel*, por devolverle la vida y el color a los almuerzos en la Mensa y por recordarme que aunque estamos a miles de kilómetros de México, uno puede sentirse como en casa con la gente adecuada. Un agradecimiento muy especial va para César, David y Fátima, quienes aparecieron en mi vida en un momento muy oscuro, y su compañía en el último año me salvo de no perderme en el intento de escribir una tesis para la que no tenía motivación. Gracias por los breaks de media tarde, por las cenas, las salidas, el mezcal y la cumbia.

Gracias a mis amigos en México, Nany, Sam, Berni y Sofí, que ocupan un lugar especial en mi corazón, y que me reciben en sus vidas como si los años no hubieran pasado y como si la distancia no existiera. Gracias también a mis amigos físicos, *Los Pichones*: Fabs, Felipe, Edgar y Chino, que sin saberlo, alegraron muchos de mis días en el exilio con su particular mezcla de lógica y humor masculino. Gracias por construir y sostener un espacio seguro para ser nosotros mismos, a pesar de los años. Gracias a Fabs: nos precede un año muy duro, lleno de pérdidas y redefiniciones. No ha sido fácil y la distancia lo ha complicado todo, pero seguimos aprendiendo y me alegra estar y que estés. Gracias por ayudarme a ajustar superelipsoides cuando estoy al borde de un ataque de nervios y por hacerme reír cuando me hace falta. Gracias por tanto y perdón por tan poco.

A mis hermanos, Jackie y Alejandro, y a mis padres, Adriana y Alejandro. Familia, la gratitud que siento hacia ustedes no cabe en estas líneas. Estoy convencida de que las circunstancias que ahora me rodean obedecen al esfuerzo de ustedes y al apoyo que inagotablemente me han dado. Gracias por todas las idas y venidas a través del Atlántico, por el cuidado y la ternura que le profesan a Ramiro, por la paciencia conmigo y por confiar en que algo bueno podía salir de todo esto, incluso cuando el panorama era sombrío. Gracias por todo el afecto y por los tantísimos episodios felices; si tengo que elegir uno, gracias por el levantón de febrero 2026: por llevarme al mar y por darme la fuerza necesaria para ese último estirón que hizo posible terminar esta tesis. Gracias totales.

A mi Rami... por dónde empezar? no tengo palabras para expresar mi admiración a tu paciencia y madurez por mi ausencia parcial en casa a lo largo de estos años. Confío

en poder sembrar en ti pasión y constancia, para que un día sientas el entusiasmo que me impulsó a trazar este camino. Gracias por darme el estímulo y tu sensible cariño siempre, todo es pensando en ti.