

DISSERTATION

submitted to the

Combined Faculty of Mathematics, Engineering and Natural
Sciences of Heidelberg University, Germany

for the degree of

Doctor of Natural Sciences

Put forward by

Matteo Guida

Born in Moncalieri, Italy

Oral examination: 20.10.2025

Low-Energy Electronic Recoils in XENONnT

New Physics Searches,
First Sub-keV Calibration, and Improved Krypton Assay

Referees:

Prof. Dr. Teresa Marrodán Undagoitia

Prof. Dr. Stephanie Hansmann-Menzemer

*To my friends,
the priceless treasure along my way.*

*To my parents and sisters,
with gratitude for your love, roots, and support.*

*To you, Cristina,
my fellow, my soulmate, and my home.*

Abstract

Within the context of the XENONnT experiment, this thesis presents contributions to low-energy electronic recoil (ER) searches and characterizations, advancing the quest for physics beyond the Standard Model (BSM).

Analysis of 1.16 tonne-years of XENONnT data from the first science run achieves a record-low ER background rate of (15.8 ± 1.3) events/(tonne·year·keV) in the 1 to 30 keV range, a 5-fold reduction over XENON1T. Calibration with the 2.82 keV ^{37}Ar K-shell line establishes an accurate detector response in the critical few-keV region. The findings establish world-leading laboratory upper limits on solar axions, bosonic dark matter, and new neutrino physics. Specifically, the limits on new neutrino physics include those on non-standard interactions with vector and scalar mediators, as well as on an enhanced magnetic moment and millicharge. These results effectively rule out a BSM explanation for the previously observed XENON1T low-energy excess.

This work also reports the first calibration and observation of sub-keV ERs in liquid xenon (LXe), detecting the 0.27 keV ^{37}Ar L-shell decay with 11.9σ significance. This is realized by combining a lowered energy threshold, resulting in a signal acceptance increase by a factor of about 66, with robust data-driven accidental coincidence background modeling and an improved suppression strategy. Validated full-chain simulations are used to understand sub-keV detector response and to determine the photon yield. From ^{37}Ar L-shell events, a scintillation photon yield of $\text{PY} = 2.96 \pm 0.08 \text{ (stat.)}^{+0.57}_{-0.33} \text{ (syst.)} \gamma/\text{keV}$ is measured, providing essential data for sub-keV ER emission modeling.

Looking towards the next generation of LXe observatories like XLZD, this thesis includes developments for the automated rare gas mass spectrometer (Auto-RGMS), an instrument combining a cryogenic gas chromatography stage and a mass spectrometer for ultra-sensitive parts-per-quadrillion (ppq) quantification of krypton traces in xenon gas. Controlling and characterizing the beta-decaying ^{85}Kr isotope is critical, as it represents a significant ER background in low-energy analyses. By using the new adsorbent, HayeSep Q, a 12-fold improvement in chromatographic resolution is achieved. This enhancement enables processing of much larger samples and helps achieve the designed few-ppq detection limit needed for XLZD's future high-precision measurements of solar pp neutrinos and the weak mixing angle.

Furthermore, sensitivity projections demonstrate that the newly validated sub-keV ER analysis framework can enhance XLZD's searches for BSM physics, improving sensitivity to neutrino electromagnetic properties by up to 19% and to boosted dark matter by up to a factor of two, thereby maximizing the scientific output of future multi-ten-tonne scale LXe experiments.

Zusammenfassung

Im Kontext des XENONnT-Experiments beschreibt diese Doktorarbeit Beiträge zur Suche nach und Charakterisierung von niederenergetischen elektronischen Rückstößen (ER) und treibt damit die Suche nach Physik jenseits des Standardmodells (BSM) voran.

Die Analyse von 1.16 Tonnenjahren an XENONnT-Daten aus der ersten Datennahmephase erreicht eine rekordtiefe ER-Untergrundrate im Bereich von 1 bis 30 keV von (15.8 ± 1.3) Ereignissen/(Tonne·Jahr·keV), eine 5-fache Reduzierung gegenüber XENON1T. Eine Kalibrierung mit der ^{37}Ar K-Schalen-Linie bei 2.82 keV etabliert eine genaue Detektorantwort im kritischen Energiebereich von wenigen keV. Die Ergebnisse führen zu weltweit führenden laborbasierten Obergrenzen für solare Axionen, bosonische Dunkle Materie und neue Neutrino-physik. Insbesondere umfassen die Grenzen für neue Neutrino-physik solche für nicht-standardmäßige Wechselwirkungen mit Vektor- und Skalarmediatoren sowie für ein erhöhtes magnetisches Moment und Milliladungen. Diese Resultate schließen eine BSM-Erklärung für den zuvor beobachteten niederenergetischen Ereignisüberschuss in XENON1T effektiv aus.

Diese Arbeit präsentiert auch die erste Kalibrierung und Beobachtung von Sub-keV-ERs in LXe, wobei der 0.27 keV ^{37}Ar L-Schalen-Zerfall mit einer Signifikanz von 11.9σ nachgewiesen wird. Dies geschieht durch die Kombination einer abgesenkten Energieschwelle, die zu einer Erhöhung der Signalakzeptanz um einen Faktor von etwa 66 führt, mit einer robusten datengestützten Modellierung des zufälligen Koinzidenzuntergrunds und einer verbesserten Unterdrückungsstrategie. Validierte vollständige Simulationsketten werden verwendet, um die Detektorantwort im Sub-keV-Bereich zu verstehen und die Photonenausbeute zu bestimmen. Aus ^{37}Ar L-Schalen-Ereignissen wird eine Szintillationsphotonenausbeute von $\text{PY} = 2.96 \pm 0.08 \text{ (stat.)}^{+0.57}_{-0.33} \text{ (syst.)} \gamma/\text{keV}$ gemessen, was wesentliche Daten für die Modellierung der Sub-keV-ER-Emission liefert.

Mit Blick auf die nächste Generation von LXe-Observatorien wie XLZD umfasst diese Arbeit Entwicklungen für das automatisierte Edelgas-Massenspektrometer (Auto-RGMS), ein Instrument, das eine kryogene Gaschromatographiestufe und ein Massenspektrometer zur ultrasensitiven Quantifizierung von Kryptonspuren in Xenongas im Parts-per-Quadrillion (ppq)-Bereich kombiniert. Die Kontrolle und Charakterisierung des beta-zerfallenden ^{85}Kr -Isotops ist entscheidend, da es einen signifikanten ER-Untergrund in Niederenergieanalysen darstellt. Durch die Verwendung des neuen Adsorbens HayeSep Q wird eine 12-fache Verbesserung der chromatographischen Auflösung erreicht. Diese Verbesserung ermöglicht die Verarbeitung wesentlich größerer Proben und trägt dazu bei, die angestrebte Nachweisgrenze von wenigen ppq zu realisieren, die für zukünftige hochpräzise Messungen von solaren pp-Neutrinos und des schwachen Mischungswinkels mit XLZD erforderlich ist.

Darüber hinaus zeigen Sensitivitätsprojektionen, dass das neu validierte Sub-keV-ER-Analyse-Framework die Suchmöglichkeiten von XLZD für BSM-Physik erweitern kann, indem es die Sensitivität für elektromagnetische Eigenschaften von Neutrinos um bis zu 19% und geboostete Dunkle Materie um bis zu einem Faktor zwei verbessert und so den wissenschaftlichen Ertrag zukünftiger Multi-Zehn-Tonnen-LXe-Experimente maximiert.

Contents

Abstract

1	The XENON project	1
1.1	Dark matter and physics beyond the Standard Model (BSM)	1
1.2	Xenon as a radiation detection medium for rare-event searches	2
1.3	Signal generation: scintillation and ionization	3
1.4	Dual-phase TPC XENONnT experiment	5
1.5	Main physics goals of XENONnT	7
1.6	XLZD: the next-generation xenon observatory	8
2	Auto-RGMS: automated krypton tracing in xenon at ppq levels	9
2.1	^{85}Kr in liquid xenon experiments	9
2.2	The rare gas mass spectrometer (RGMS)	10
2.3	Auto-RGMS: The evolution of the RGMS detector	11
2.4	Gas-solid chromatography principles	13
2.5	Experimental setup for adsorbents testing	15
2.6	Adsorbents for krypton-xenon separation	16
2.7	Effect of temperature and helium flow on resolution	19
2.8	Memory effect evaluation	20
2.9	Conclusion and discussion	23
3	Data analysis in XENONnT	25
3.1	Data processing pipeline and event building	25
3.2	S1 detection efficiency and peak reconstruction bias	27
3.3	Vertex reconstruction and field distortion corrections	28
3.4	Signal corrections	30
3.5	Selection criteria overview	33
4	Search for BSM physics in keV-scale electronic recoils	35
4.1	First XENONnT science run	35
4.2	^{37}Ar 2.82 keV calibration for energy reconstruction	36
4.3	Energy resolution	40
4.4	Background models	42
4.5	Signal models	45
4.6	Event selection and cut acceptance	47
4.7	Inference and results	48
4.8	Conclusion and discussion	51
5	First sub-keV electronic recoil calibration with ^{37}Ar L-shell	53
5.1	Modeling near the energy threshold	53
5.2	Sub-keV energy smearing	55
5.3	Characterization of the accidental coincidence background	56
5.4	S1/S2 BDT classifiers for accidental coincidence mitigation	60

5.5	Selection criteria and efficiencies for ^{37}Ar L-shell search	62
5.6	Inference and fit results	65
5.7	Scintillation photon yield via simulation	66
5.8	Conclusion and discussion	69
6	Sensitivity to BSM physics in XLZD's low-energy electron recoils	71
6.1	Modeling signals: 2-fold vs. 3-fold coincidences	71
6.2	Background projections for XLZD	73
6.3	Accidental coincidence projections for XLZD	74
6.4	XLZD sensitivity for different accidental coincidence scenarios	76
6.5	Conclusion and discussion	77
Summary and Conclusion		79
Appendix		83
A	Author's publications used in this work	83
B	Acknowledgements	85
C	Bibliography	87

Chapter 1

The XENON project

The XENON project, the focus of this chapter, is the culmination of over half a century of conceptual and technological advancements in particle detection. The potential of liquid xenon (LXe) as a radiation detection medium was initially highlighted by Luis Alvarez in 1968 [1]. A few years later, in 1974, David Nygren proposed for the first time the use of the time projection chamber (TPC) as a central track detector in the PEP4 experiment [2]. In 1977, Carlo Rubbia pioneered the concept of a large-scale noble liquid TPC for probing rare phenomena such as proton decay and solar neutrino scattering [3]. Building on this foundation, the landscape for dark matter (DM) detection was transformed in 2007 when Elena Aprile et al. [4] demonstrated the advantages of dual-phase xenon TPCs for this purpose. Since then, dual-phase xenon TPC technology has firmly established itself as the worldwide leading technology for the direct detection of particle DM for masses above a few GeV [5]. In the last two decades, the XENON DM project has operated four successful stages of detectors based on the LXe dual-phase TPC technology, each with a progressively larger mass, a lower background, and improved sensitivity. All of these experiments (XENON10 [4, 6], XENON100 [7, 8], XENON1T [9, 10], and XENONnT [11, 12]) were built at the Laboratori Nazionali del Gran Sasso in Italy, located beneath 3600 meters of water-equivalent rock to provide shielding from cosmogenic backgrounds. Further information on these detectors and their DM search results can be found in the respective references.

1.1 Dark matter and physics beyond the Standard Model (BSM)

The Standard Model (SM) of particle physics has achieved remarkable success in describing fundamental particles and forces [5]. However, it is considered an incomplete theory, unable to address several questions arising from both experimental observations and theoretical considerations, thereby necessitating physics beyond the Standard Model (BSM). To list a few of these shortcomings, the SM offers no explanation for neutrino masses and oscillations [13], the universe's matter-antimatter asymmetry [14], the hierarchy problem and the Higgs mass fine-tuning [15], the strong CP problem [16], a consistent description of gravity within a quantum framework [17], nor the existence of DM [18] and dark energy [19].

The existence of DM stands out as one of the most significant of these unresolved issues. A vast range of cosmological and astrophysical observations, such as galaxy rotation curves, gravitational lensing, and the cosmic microwave background, indicates that about 26.4% [5] of the universe's mass-energy, according to the Λ CDM cosmology model, consists of an unknown, non-luminous, non-relativistic, and non-baryonic form of matter inferable through its gravitational effects. Despite over four decades of intensive research through direct detection, indirect detection, and collider searches, the precise nature of DM remains an open question. The possible range of DM masses

extends across more than 90 orders of magnitude, from ultra-light scalar particles to macroscopic objects, and it is not yet known whether the hypothetical DM particle(s) interacts through anything other than gravity. Among the plethora of theoretical particle candidates, Weakly Interacting Massive Particles (WIMPs) have been the most extensively studied and represent the primary scientific target for the XENON project. To observe the rare interactions of WIMPs with detector nuclei, direct detection experiments like XENON leverage the effective “DM wind”. This wind is a direct consequence of the Sun’s motion through the stationary galactic DM halo, resulting in an apparent stream of DM particles that seems to originate from the direction of the Cygnus constellation.

While alternatives to the particle paradigm, such as theories that modify gravity, have been proposed to explain some observed phenomena, they generally face challenges in reproducing the full suite of cosmological and astrophysical observations attributed to DM.

A comprehensive treatment of DM, including the topics outlined above, along with a thorough overview of experimental and theoretical approaches and results, is provided in [18] and references therein. Additional BSM scenarios testable by XENONnT, including solar axions, bosonic DM, and new physics in the neutrino sector specifically from solar neutrino interactions, are motivated and detailed in chapter 4.

1.2 Xenon as a radiation detection medium for rare-event searches

The choice of liquid noble gases, particularly LXe and liquid argon (LAr), is common for building DM detectors. These materials excel at both generating scintillation light and releasing ionization electrons as radiation passes through. This feature is fundamental for developing dual-phase TPCs, which are covered later. Since xenon and argon are normally gases at room temperature and pressure, they need to be cooled to cryogenic temperatures to turn into liquids for use in detectors. Table 1.1 summarizes some selected important features of these two noble gases.

Property	Xenon	Argon
Atomic number (Z)	54	18
Atomic mass (A)	131.3	40.0
Boiling point at 1 atm (T_b)	165.0 K	87.3 K
Liquid density at T_b	2.94 g/cm ³	1.40 g/cm ³
Atmospheric amount	0.09 ppm	9340 ppm
Cost	\$\$\$	\$ [*]
Scintillator	✓	✓
W	13.7 eV/quantum	19.5 eV/quantum
Ionizer	✓	✓
Scintillation wavelength	174.8 nm	128.1 nm

Table 1.1: Selected properties of liquid xenon and liquid argon, which are commonly used as target materials for direct dark matter detection. W represents the mean energy required to create either scintillation or ionization quanta. Values are from references [20–23].

^{*}This refers to atmospheric argon, which is not suitable for low-energy rare event searches due to the presence of the radioactive isotope ^{39}Ar . Underground argon, on the other hand, is significantly more expensive [24].

Choosing LXe as a detector material offers several benefits, among which there are:

- **Large cross-section:** Xenon has a high atomic mass, which is important because the expected cross-section for spin-independent (SI) WIMP interactions scales with the atomic mass squared ($\sigma_{\text{SI}} \propto A^2$).
- **Low radioactivity:** With the exception of two exceptionally long-lived isotopes (^{136}Xe [25] and ^{124}Xe [26]), natural xenon is composed entirely of stable isotopes, resulting in very low internal background radiation. In contrast, natural argon contains the radioactive isotope ^{39}Ar ($\sim 1 \text{ Bq/kg}$), which needs to be removed for effective use.
- **Self-shielding:** As shown in table 1.1, LXe has a high density. This allows it to block external background radiation, creating a clean inner volume (fiducial volume) for particle detection.
- **Spin-dependent sensitivity:** About 47.6% of natural xenon isotopes have a non-zero nuclear spin (^{129}Xe with spin 1/2 and ^{131}Xe with spin 3/2). This feature allows for studies of spin-dependent DM interactions, which is valuable for understanding DM's nature.
- **High signal yield:** As detailed in table 1.1, the average energy needed to generate an electron-ion pair or a scintillation photon is lower in xenon. This means that for the same amount of deposited energy, more photons and electrons are produced, leading to a stronger signal.
- **Easier cryogenics:** Compared to LAr, LXe requires less extreme cooling, making its cryogenic system simpler to operate.
- **No wavelength-shifter:** Since xenon emits scintillation light in the vacuum ultraviolet (VUV) range and is transparent to it, photomultiplier tubes (PMTs) can detect it directly. In contrast, in LAr detectors, wavelength shifters are used.

The high cost of xenon is due to its scarcity; it accounts for just 1 part in 10^7 of Earth's atmosphere. It is extracted from air in industrial plants, with prices typically from 1000 to 3000 EUR/kg [27]. Acquiring enough xenon for large-scale detectors faces significant challenges, notably limited global production, much of which is tied up in long-term contracts, and increasing geopolitical tensions impacting supply chains.

1.3 Signal generation: scintillation and ionization

When a particle interacts with LXe, it deposits energy either by scattering off electrons (electron recoils, ER) or by scattering off xenon nuclei (nuclear recoils, NR). In both cases, the total deposited energy is split into two main parts: the inelastic channel, which includes ionization (creating $e^- + \text{Xe}^+$ pairs) and excitation (creating Xe^* excited atoms), and the elastic channel, which involves atomic motion (heat):

$$\frac{dE}{dx} = \left(\frac{dE}{dx} \right)_{\text{inel}} + \left(\frac{dE}{dx} \right)_{\text{el}}. \quad (1.1)$$

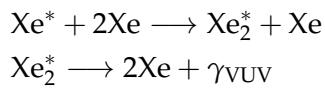
The energy deposited into the inelastic channel, $E_0 L$, can be described by the Platzman equation [28], which relates the number of resulting electron-ion pairs (N_i) and excited atoms (N_{ex}):

$$E_0 L = N_i E_i + N_{ex} E_{ex} + N_i \epsilon, \quad (1.2)$$

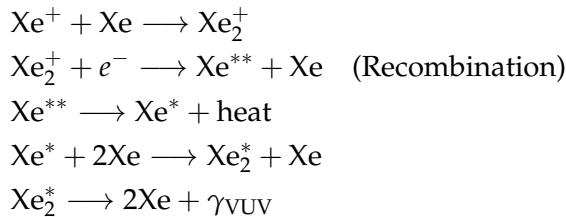
where E_0 is the total deposited energy, L the Lindhard factor, E_i and E_{ex} are the average energies to create an ion pair and an exciton, respectively, and ϵ is the average kinetic energy of sub-excitation electrons, the energy of which goes into heat. For LXe, the value of ϵ is about 5 eV [29]. The Lindhard factor L describes how the total deposited energy E_0 is "quenched", or reduced, depending on the type of particle interaction. For an ER, the electron's mass is very small compared to the target xenon atoms. This means that almost no energy is lost through elastic collisions, so $L \simeq 1$ [21]. However, for NRs, as the recoiling nucleus loses a significant fraction of its energy through elastic collisions, the Lindhard factor is less than 1, typically around 0.15 to 0.2 in the range of 3–100 keV [30].

Scintillation refers to the emission of light by liquids or solids when they luminesce. This process involves atoms emitting photons to achieve a lower energy state. In noble gases like xenon, scintillation photons are produced either through excitation (Xe^*) or ionization (Xe^+). Both mechanisms result in the creation of an excited dimer (Xe_2^*) via collisions. The dimer subsequently returns to a dissociated ground state by emitting VUV photons, γ_{VUV} . The detailed reaction pathways for both channels are shown below:

Excitation channel



Ionization channel



The wavelength of the scintillation light produced in LXe is around 175 nm [23], which is larger than that of LAr (~ 128 nm) and allows for light detection without the use of a wavelength shifter. In pure LXe, scintillation light consists of two decay components, originating from the de-excitation of the Xe_2^* dimer via its spin singlet ($s = 0$) and triplet ($s = 1$) states. These states differ in molecular angular momentum and produce fast (singlet) and slow (triplet) light emission as they relax to the ground state. In principle, the difference in pulse decay shapes can be used to distinguish ERs from NRs. In practice, this is not very effective in xenon because the time separation between the two decay components is small: about 2.2 ns for the singlet and 27 ns for the triplet [31].

The average energy required to generate a single photon, W_{ph} , is given by:

$$W_{\text{ph}} = \frac{E_0}{N_{\text{ex}} + N_i} = \frac{W_i}{1 + N_{\text{ex}}/N_i}, \quad (1.3)$$

where N_{ex} is the number of excitons and N_i is the number of electron-ion pairs. The total number of produced photons is:

$$N_{\text{ph}} = N_{\text{ex}} + rN_i, \quad (1.4)$$

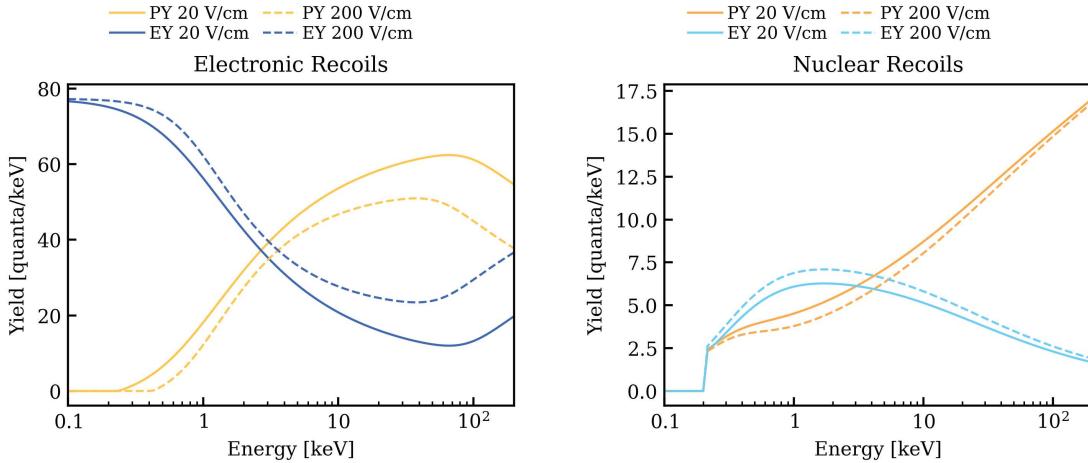


Figure 1.1: Electron (EY) and photon yields (PY) in liquid xenon as a function of the initial energy of the incident particle, for both electronic recoils (left) and nuclear recoils (right), obtained from NEST [32] (nestpy v2.4.0 [33]). The dependence of the yields on the applied electric drift field is illustrated by plotting the curves for 20 V/cm (solid lines) and 200 V/cm (dashed lines).

where r is the recombination fraction. If an external electric field is applied, the number of electrons not undergoing recombination and thus extracted as charge is:

$$N_q = (1 - r)N_i. \quad (1.5)$$

Energy partitioning between excitation and ionization channels varies with the density of electron-ion pairs formed along the particle track, thereby yielding distinct responses for NRs and ERs. Moreover, the recombination fraction r is influenced by both the applied electric field and the ionization density along the track. Figure 1.1 shows the resulting photon and electron yields as a function of recoil energy for both ERs and NRs, demonstrating the clear dependence on the applied electric field. With these relations, the total deposited energy can be written as:

$$E_0 = W_{\text{ph}}(N_q + N_{\text{ph}}). \quad (1.6)$$

The energy needed to generate a single detectable quantum (N_q or N_{ph}) is known as the W -value, where $W_{\text{ph}} = W$. The standard value adopted is $W = (13.7 \pm 0.2)$ eV/quantum [21]. More information is given in section 4.2.

1.4 Dual-phase TPC XENONnT experiment

A dual-phase TPC operates by detecting two distinct signals that arise when a particle interacts within its active liquid volume. The first is a prompt scintillation light signal, known as S1. These are the VUV photons from the decay of diatomic excimers, as described in section 1.3. Emitted isotropically from the interaction point, these S1 photons are collected by two arrays of PMTs located at the top and bottom of the TPC. The TPC's internal surfaces are covered with highly reflective polytetrafluoroethylene (PTFE) panels, which efficiently reflect these VUV photons, enhancing overall light collection. The second signal, S2, originates from the ionization of xenon atoms, which liberates free electrons. A uniform electric drift field guides these electrons upwards

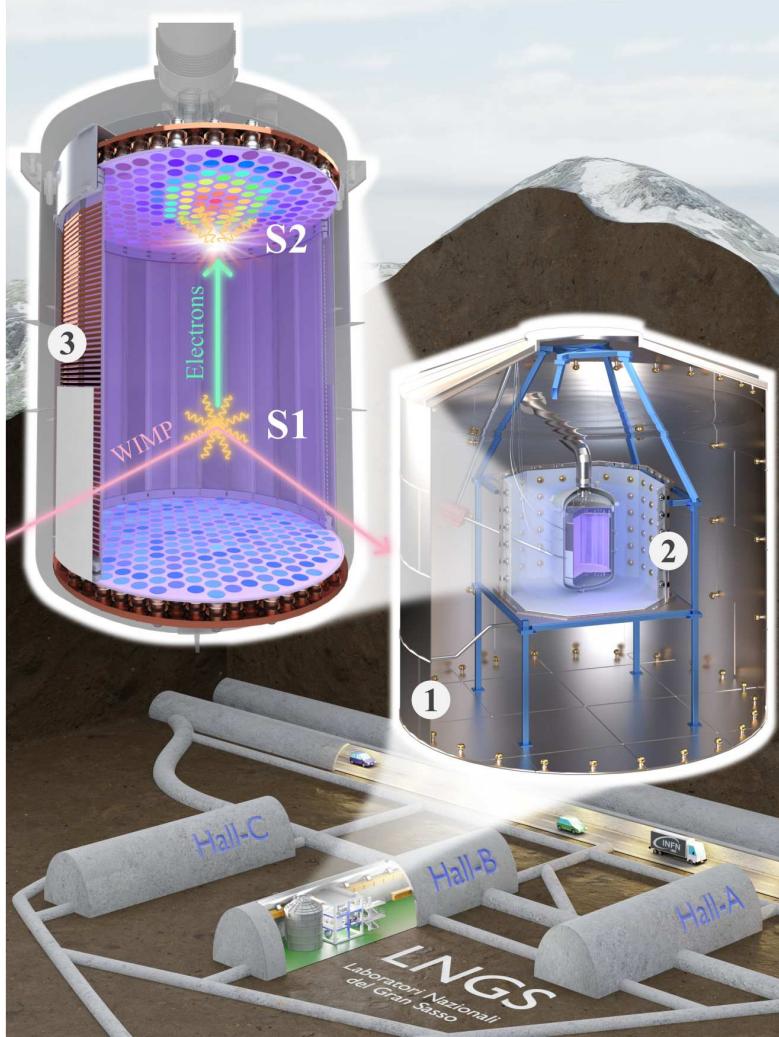


Figure 1.2: A visual representation of the XENONnT experiment located at the Laboratori Nazionali del Gran Sasso (LNGS), Italy, illustrating its deep underground site accessible through vehicle tunnels. The experiment consists of three nested detectors: a water tank (1) serving as a muon veto, enclosing the neutron veto (2), which in turn surrounds the central dual-phase time projection chamber (TPC) (3). The working principle of the dual-phase TPC, which is discussed in more detail in the text, is also illustrated. Figure from [34].

through the liquid, limiting their recombination with ions. Upon reaching the liquid-gas interface, a much stronger electric extraction field pulls the electrons into the gaseous xenon (GXe) phase. Here, they are accelerated, inducing a secondary proportional scintillation signal via electroluminescence, which is primarily detected by the top PMT array. The working principles of the dual-phase TPC technology, specifically as implemented in XENONnT, are outlined in figure 1.2.

The simultaneous readout of the prompt S1 and the delayed S2 signals provides the foundational capabilities of the detector. Three-dimensional position reconstruction is achieved by using the time difference between the S1 and S2 signals to infer the interaction's depth (Z) and the S2 light pattern on the top PMT array to determine its horizontal (X, Y) position. This allows for the definition of an inner fiducial volume, effectively shielding the analysis region from external backgrounds.

Furthermore, the ratio of the two signals ($S2/S1$) serves as a powerful tool for particle discrimination, as it is proportional to the linear energy transfer (see section 1.3). ERs,

which can originate from various background radiation sources like beta and gamma decays, as well as from solar neutrinos or BSM signals (see section 4.5), generally produce a higher S2/S1 ratio. In contrast, NRs, generated by solar neutrinos via coherent elastic neutrino-nucleus scattering (CEvNS), neutron background, and WIMP DM interactions, exhibit a lower S2/S1 ratio, enabling efficient discrimination [35]. Finally, the combination of the S1 and S2 signals allows for a precise reconstruction of the total deposited energy, as detailed in chapter 3. For a comprehensive review of LXe detectors and more about the principles behind dual-phase TPCs, see [1].

The XENONnT experiment is made up of three nested detectors: at its center is a xenon dual-phase TPC encased in a cryostat, which is enclosed by a neutron veto (NV) detector [36]. This assembly sits inside, yet remains optically separate from, a muon veto (MV) detector [37]. Both veto systems are housed together inside a 700-tonne water tank and act as water Cherenkov detectors to mitigate radiogenic neutrons and cosmogenic muon-induced backgrounds. With a diameter of 1.33 m and a drift length of 1.49 m, the TPC’s sensitive volume contains about 5.9 tonnes of LXe. To maintain purity, xenon is continuously purified by a gas and liquid purification system [38], effectively eliminating electronegative impurities detrimental to electron drift. To suppress low-energy ER background, two cryogenic distillation columns are employed: one for online radon mitigation [39], and another for removing krypton and argon [40].

The drift field is maintained between the cathode electrode at the bottom of the active volume and the gate electrode just below the LXe surface; together with the field cage, these electrodes establish and shape the electric field in the drift region [41]. During initial XENONnT operations, a discharge that occurred during cathode ramp-up likely caused a short with the bottom screening mesh, capping the maximum cathode voltage and consequently limiting the maximum achievable drift field to approximately 23 V/cm, below the original design value [42]. Similarly, the extraction field had to be reduced due to intense single-electron emission, which resulted in an extraction efficiency of around 50% [11]. All electrodes are built from sets of parallel wires. To avoid sagging, the gate and anode wires are supported by two and four transverse wires, respectively. Additional parallel-wire screening electrodes shield the PMT arrays, which consist of 494 Hamamatsu 3" PMTs arranged with one array at the top and one at the bottom of the detector [43]. The top array contains 253 PMTs and the bottom array has 241, both arranged in a hexagonal pattern to maximize light collection efficiency. All construction materials were chosen through a radioassay program that focused on low radioactivity to minimize background [44]. Further details of the detector’s components and performance can be found in [11].

1.5 Main physics goals of XENONnT

As a low-threshold, low-background tonne-scale detector, XENONnT is a multi-purpose observatory for rare-event searches [52], with its primary physics goal being the search for WIMP DM. This section briefly reports two major outcomes from XENONnT NR analyses that are not discussed further in this thesis, but have notable scientific relevance. Other important results obtained from low-energy ER data are presented in the following chapters. The first achievement is the anticipated detection of NRs from solar ^8B neutrinos via CEvNS [48]. This analysis provides the first detection of astrophysical neutrinos in a DM detector, the first measurement of CEvNS from an astrophysical neutrino source, and the first measurement of CEvNS with a xenon target. The result rejects the background-only hypothesis with a statistical significance of 2.7σ .

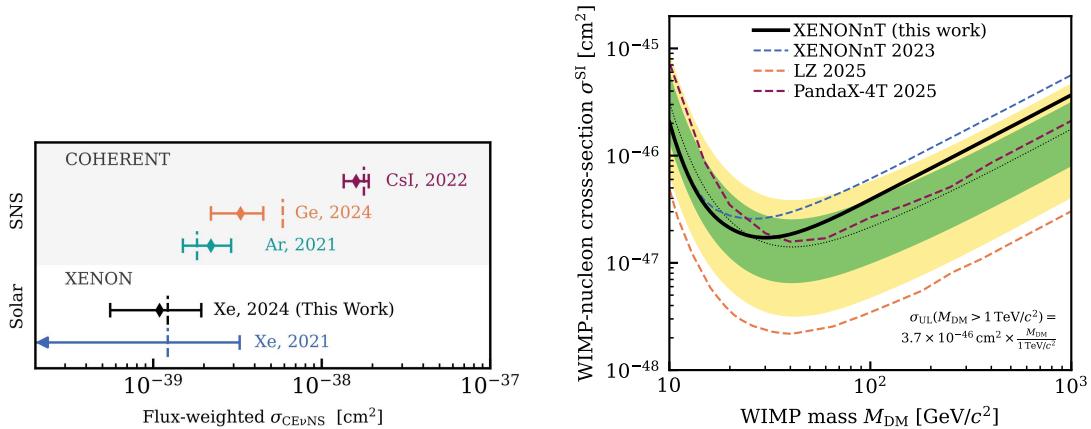


Figure 1.3: Left: Flux-weighted coherent elastic neutrino-nucleus scattering (CE ν NS) cross section measurements ($\sigma_{CE\nu NS}$). The result from the XENONnT analysis from solar 8 B neutrinos is shown in black, alongside the 90% CL upper bound from XENON1T (blue). COHERENT Collaboration results using neutrinos produced at the Spallation Neutron Source (SNS) are shown for CsI (red) [45], Ar (green) [46], and Ge (orange) nuclei [47]. Vertical dashed lines mark Standard Model predictions. Figure from [48]. **Right:** Upper limits (90% C.L.) on the spin-independent WIMP-nucleon cross-section as a function of WIMP mass (black line), with sensitivity bands for 68% (green) and 95% (yellow) of expected upper limits under the background-only hypothesis, plus the median (dotted). Published limits from XENONnT SR0 [49], LZ [50], and PandaX-4T [51] are also plotted for comparison. Figure from [12].

The second result is the search for WIMP DM using the combined dataset from the first and second science runs of XENONnT [49]. No significant excess of events is observed, resulting in an upper limit on the SI WIMP-nucleon scattering cross-section of $1.7 \times 10^{-47} \text{ cm}^2$ for a 30 GeV WIMP mass. A summary of these two findings is shown in figure 1.3.

1.6 XLZD: the next-generation xenon observatory

XLZD is the next generation in LXe dual-phase TPC observatory technology with a nominal active mass of 60 tonnes and dimensions of 2.98 m in diameter and 2.97 m in height [53]. Designed to push the sensitivity frontier for WIMPs across a broad mass spectrum, XLZD aims to probe interaction rates down to the neutrino fog where astrophysical neutrinos become an irreducible background and deliver a 3σ evidence potential for WIMP-nucleon cross sections as low as $3 \times 10^{-49} \text{ cm}^2$ at 40 GeV WIMP mass. The observatory will also offer leading sensitivity to a broad range of alternative DM models and boost sensitivity to all BSM signals described in detail in chapter 4 through the analysis of low-energy ER data. These data will also allow high-statistics measurements of the solar pp neutrino flux and precise determinations of the weak mixing angle at low energies, reaching a momentum transfer region over an order of magnitude smaller than what atomic parity violation experiments have explored, in a region sensitive to BSM physics [54]. Delivering this measurement requires frequent quantification of krypton contamination at ppq levels, enabled by the technology and R&D detailed in Chapter 2. XLZD will also advance leading investigations of neutrinoless double beta decay of ^{136}Xe , with projected sensitivity to half-lives up to 5.7×10^{27} years, achieved through the analysis of high-energy (MeV) electron recoils.

Chapter 2

Auto-RGMS: automated krypton tracing in xenon at ppq levels

2.1 ^{85}Kr in liquid xenon experiments

Krypton contamination, specifically the beta-decaying isotope ^{85}Kr , poses a significant challenge for low-background liquid xenon (LXe) experiments searching for rare processes. One significant source of this background is the presence of trace amounts of natural krypton ($^{\text{nat}}\text{Kr}$) in commercially available gaseous xenon (GXe) distilled from air, along with outgassing from detector materials, such as polytetrafluoroethylene (PTFE) panels that were exposed to air prior to assembly. ^{85}Kr , generated through nuclear fission, enters the atmosphere predominantly as a result of nuclear fuel reprocessing and past nuclear weapons tests. In Europe, reprocessing plant activities can double ^{85}Kr levels compared to the baseline, with the La Hague plant in France being the most significant source. To ensure a representative assessment of the European background level, ^{85}Kr samples are measured at contrasting altitudes and locations, specifically at the high-altitude site of Jungfraujoch (3500 m) and in Freiburg (280 m), and the results are consistent [55]. As a result, Europe currently has an average baseline atmospheric ^{85}Kr activity concentration of $(1.39 \pm 0.08) \text{ Bq/m}^3$, corresponding to a $^{85}\text{Kr}/^{\text{nat}}\text{Kr}$ ratio of $(2.33 \pm 0.14) \times 10^{-11} \text{ mol/mol}$.

The decay of ^{85}Kr , which has a half-life of 10.76 years, primarily occurs through β^- emission (99.56%) to the ground state of ^{85}Rb , with an end-point energy of 687 keV. In 0.44% of disintegrations, the β^- emission is followed by a delayed γ emission from the de-excitation of ^{85}Rb ($J^\pi = 9/2^+$) with a relaxation time of $1.02 \mu\text{s}$ [56]. The uniform distribution of $^{\text{nat}}\text{Kr}$ in xenon prevents spatial discrimination of this background, as both are noble gases that mix homogeneously. Consequently, techniques like shielding or fiducialization of the LXe volume are ineffective, classifying ^{85}Kr as an intrinsic background. A cryogenic distillation column is used in the XENON dark matter program to minimize krypton contamination. This process exploits the higher vapor pressure of krypton compared to xenon at -98°C , allowing a single pass to achieve a reduction factor of $(6.4^{+1.9}_{-1.4}) \times 10^5$ [40]. The system can process around 43 kg of xenon per day and can operate in both offline and online modes. In offline mode, xenon is purified before entering the detector, whereas in online mode, krypton is continuously removed during data taking. The lowest krypton concentration recorded, at the end of XENON1T, was below 26 parts-per-quadrillion (ppq) [40].

A xenon gas sample is extracted from the running experiment approximately every 1-2 months to assess the krypton concentration. The quantification of krypton at the required ppq level in these samples is a significant analytical challenge. While techniques such as atom trap trace analysis (ATTA) can detect individual ^{85}Kr atoms, with demonstrated sensitivities of 10^{-17} mol/mol in atmospheric samples [57], their

application to krypton-in-xenon measurements has so far achieved only a sensitivity of about 10 ppt [58, 59], which is inadequate for the needs of current and future LXe experiments.

The leading approach at the moment for indirect quantification of ^{85}Kr contamination in xenon involves separating krypton and xenon via gas-solid chromatography (GSC), followed by counting krypton ions in a mass spectrometer. The rare gas mass spectrometer (RGMS) [60] operated at MPIK Heidelberg currently has the best sensitivity for these measurements.

2.2 The rare gas mass spectrometer (RGMS)

To quantify the concentration of a trace component at $\mathcal{O}(\text{ppt-ppq})$ levels in a mixture sample of $\mathcal{O}(\text{cm}^3)$, the RGMS system uses a two-stage approach: GSC separates krypton from xenon, and the mass spectrometer measures ^{84}Kr and ^{86}Kr to infer ^{85}Kr . This design is necessary because direct analysis of xenon-rich samples would disrupt the spectrometer's vacuum and damage the detector. To prevent contamination, the system maintains an ultra-high vacuum below 10^{-10} mbar and uses ultra-pure helium as a carrier gas. Before reaching the sample, helium passes through the column labeled T1 (see figure 2.1) to remove any impurities. Helium's chemical inertness ensures it does not interact with the highly porous adsorbent powder filling the columns. The xenon-krypton sample is injected into the system after T1 and is pushed by the helium flow through column T2, where the gases are separated. Krypton and xenon atoms interact with the highly porous adsorbent via weak van der Waals forces, which are electromagnetic interactions between fluctuating charges on two neutral surfaces. As helium pushes the sample through the column, krypton and xenon undergo repeated adsorption and desorption cycles. By precisely controlling the helium flow and raising the temperature of T2 from -196°C to -90°C , the two gases are separated based on their differing adsorption and desorption rates. After around 20 minutes krypton is released from T2, whereas xenon appears after about 100 minutes. Krypton is then directed to column T3, which contains the same adsorbent material as T2 and is maintained at liquid nitrogen (LN_2) temperature. A valve between T2 and T3 is precisely timed to close once all the krypton is captured in T3, preventing xenon, which elutes later, from contaminating the krypton. T3 is then heated to release the krypton gas, which is transferred to the mass spectrometer using cryogenic pumping. This involves cooling a "cold finger" in the mass spectrometer with activated carbon, where the krypton gas condenses and is collected. In the mass spectrometer, krypton is ionized using a hot filament that emits electrons, which collide with krypton atoms creating ions. Ions are accelerated and differentiated by their mass-to-charge ratio in a magnetic field. A secondary electron multiplier (SEM) measures the ion rate at a fixed angle. Isotope selection in the SEM is achieved by varying the magnetic field strength, which usually falls within the 0.3-0.4 T range.

To determine the ^{85}Kr content, the rates of ^{84}Kr and ^{86}Kr are measured due to their high natural abundance (57.0% and 17.3%, respectively [61]). In addition, ^{36}Ar is measured, with its abundance typically about 50 times greater than that of krypton isotopes, consistent with atmospheric concentrations and natural abundances. However, its higher vapor pressure complicates freezing it onto the cold finger without losses, resulting in limited measurement precision. As a result, ^{36}Ar quantification currently serves mainly as a cross-check rather than for precise argon determination. To monitor for possible xenon leaks in the mass spectrometer, $^{130}\text{Xe}^{2+}$ is also measured. If xenon is present, this dication is produced by electrons from the hot filament ionizing xenon atoms. $^{130}\text{Xe}^{2+}$ is

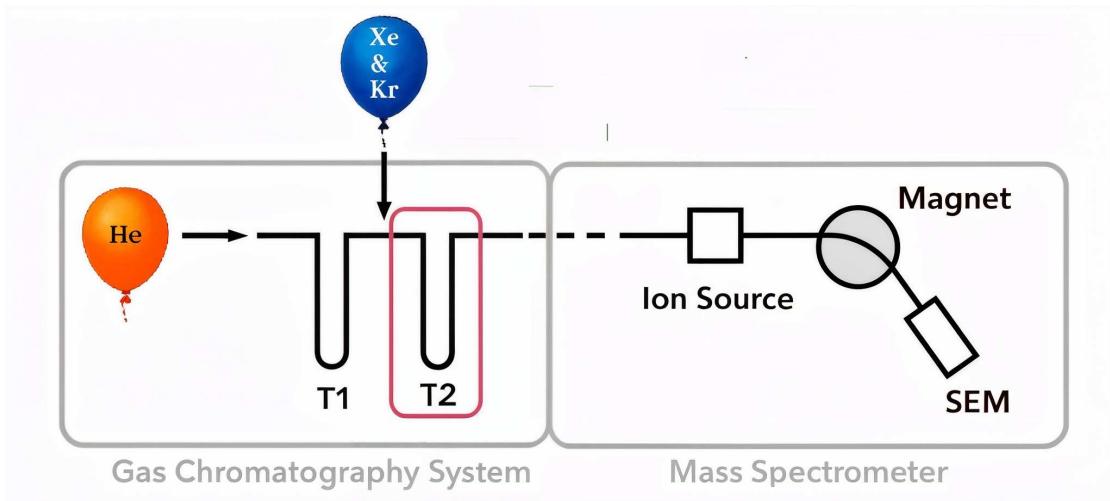


Figure 2.1: Schematic of the rare gas mass spectrometer (RGMS) system, illustrating the two-stage process of gas chromatographic separation followed by isotopic analysis via mass spectrometry. Key subsystems are shown and described in more detail in the main text.

preferred over $^{130}\text{Xe}^+$ because its double charge-to-mass ratio allows measurement at a lower magnetic field, which is easily achievable by the mass spectrometer capabilities.

To assess the background contamination of krypton not originating from the sample, a blank measurement is routinely conducted. This involves following the same measurement procedure as a regular sample measurement, except no sample is injected. The detection limit of RGMS has decreased over the past decade, from approximately 8 ppq under ideal conditions to tens of ppq in recent measurements. The reasons for this decline will be explained in the following section.

2.3 Auto-RGMS: The evolution of the RGMS detector

For future low-background LXe detectors, such as DARWIN/XLZD, a $^{nat}\text{Kr}/\text{Xe}$ concentration of 0.1 ppt is targeted for the dark matter WIMP search [53, 62]. Simultaneously, achieving DARWIN/XLZD's scientific goal of measuring the solar pp neutrino flux with per-mille precision and probing the weak mixing angle at momentum transfers an order of magnitude lower than current measurements [5] requires a $^{nat}\text{Kr}/\text{Xe}$ concentration of a few ppq [63], necessitating equally precise external measurements. The nearly identical electron scattering spectrum of the solar pp neutrinos and the beta-decaying ^{85}Kr , coupled with ^{85}Kr being the second most significant background after ^{214}Pb , means that reducing and precisely quantifying ^{85}Kr is crucial for enhancing the sensitivity. During XENONnT SR0, a $^{nat}\text{Kr}/\text{Xe}$ molar concentration of (60 ± 40) ppq was achieved and measured with the existing RGMS [64]. However, the current RGMS system, originally developed for isotopic studies in geosciences and not optimized for trace impurity control, has several limitations and does not meet the requirements of next-generation LXe detectors.

Firstly, performing a measurement requires numerous manual steps by a single operator, which introduces substantial variability and systematic uncertainties, resulting occasionally in deviations between measurements from the same sample and demanding considerable human resources. Operational errors can also cause downtime and subsequent increases in background levels. Secondly, valve aging has reduced their tightness,

and operational incidents (e.g., burning of adsorbent powders) have left residual carbonized particles in the system, further increasing the background. Minute changes in leak tightness anywhere in the large system also impact the blank background, hindering reproducibility. Thirdly, in recent years, the RGMS has exhibited a delayed response, requiring approximately 15 minutes for the ^{86}Kr count rate to stabilize [65]. Ideally, upon gas inlet, the detected count rate should immediately reflect the actual krypton concentration in the sample. However, the observed “grow-in” phase, where the count rate initially increases before stabilizing, introduces uncertainty in the measurement. This effect could be due to either a delayed response of the instrument to the sample or the release of residual krypton from the instrument’s walls. Until systematic tests clarify the cause and allow for a correction, both higher and lower interpretations of the measured concentration are reported, with an uncertainty range covering both cases [66]. These factors, combined with long connection pipes and dead-ends that complicate the removal of residual krypton and xenon, have degraded the RGMS’s detection limit from its original 8 ppq to tens of ppq. Furthermore, each measurement necessitates a one-day bake-out and helium flushing to remove residual gases from the chromatography adsorbent, limiting throughput.

To address these limitations and meet the demands of next-generation xenon observatories, which require improved krypton detection limits, increased measurement frequency for prompt leak detection, and time-dependent krypton concentration characterization, a new fully automated system, Auto-RGMS, is being developed and commissioned at MPIK. The redesigned chromatography stage of Auto-RGMS brings notable improvements over the current RGMS setup. Its more compact design minimizes internal surface areas, lowering the risk of gas redeposition and residual contamination. Importantly, to greatly improve leak-tightness and maintain ultra-high vacuum conditions, all-metal valves with fully welded bellows and metal gasket seals are employed, utilizing 16 mm ConFlat flanges. To ensure a consistently clean vacuum environment, turbo molecular pumps continuously operate 24/7, and vacuum pipes are heated to approximately 80 °C. This heating minimizes residual gas adsorption onto internal surfaces, stabilizing background levels for reproducible measurements. Additionally, a getter pump is installed to effectively remove chemically active contaminants, further improving measurement purity and accuracy. Multiple bypasses integrated into the system offer considerable flexibility for debugging and operational maintenance. Automation of the Auto-RGMS system is implemented using LabVIEW software. The measurement sequences are fully programmable via XML recipes, allowing standardized procedures and minimizing systematic uncertainties that were previously introduced by manual operation. All operational and measurement data are stored in a PostgreSQL database. An uninterruptible power supply is integrated to ensure continuous system operation, even during power outages. One key improvement is the development of a custom-built, all-metal double-valve with a precisely calibrated internal volume of $(0.09610 \pm 0.00034) \text{ cm}^3$ [67]. The precise knowledge of this valve’s volume is crucial for accurately introducing known amounts of calibration gas, which is essential for calibrating the system response and ensuring consistent measurement accuracy. Therefore, this component’s systematic contribution to the total uncertainty is expected to become subdominant. Furthermore, the new chromatography column features precise temperature control over a wide range from -196 °C to +200 °C with a precision better than 0.5 °C, and the system maintains a highly stable helium carrier gas flow rate, varying by less than 1% over periods exceeding 10 hours [68]. Such precise temperature and flow regulation is essential, as chromatographic separation reproducibility depends on stable separation temperatures and carrier gas flow rates, as detailed in section 2.7.

Although these upgrades enhance many aspects of the system, the chromatography column still lacks the separation power needed to process large xenon samples for low-ppq detection limits in next-generation LXe observatories. To overcome this, a range of alternative adsorbents is being systematically investigated, as detailed in the following sections.

2.4 Gas-solid chromatography principles

A short general introduction to GSC principles relevant to krypton-xenon separation is given in the following. GSC is a technique that separates gas-phase atoms or molecules based on their differing adsorption and desorption kinetics on a solid stationary phase (the adsorbent). In this technique, a gas mixture is transported by an inert carrier gas through a column packed with adsorbent material. Due to varying adsorption affinities, components of the mixture travel at distinct velocities, resulting in separation. This method is effective also for noble gases such as krypton and xenon, which interact with solid surfaces through van der Waals forces, a type of electromagnetic interaction arising from correlated fluctuating charges on two electrically neutral surfaces. In GSC, the stationary phase typically consists of a column packed with a powder of a material exhibiting a high specific surface area, typically measured in square meters per gram (m^2/g). Porous materials are used because their high surface area, illustrated in figure 2.2, offers numerous interaction sites for gas particles and improves separation. The pore structure of an adsorbent is classified based on the inner diameter (d) of the pores, which are categorized into three types: macropores ($d > 50 \text{ nm}$), mesopores ($2 \text{ nm} < d < 50 \text{ nm}$), and micropores ($d < 2 \text{ nm}$). Additionally, the pore volume represents the total volume of all pores, independent of their size or classification, and is typically determined using a well-established method. Unfortunately this single parameter does not provide information about the pore size distribution, which is crucial for the adsorption capacity of the material.

Manufacturers typically provide only the specific surface area and main pore size category, but almost never include the complete pore size distribution, which is crucial for assessing adsorption performances. However, even with detailed pore size data, predicting separation capabilities remains challenging due to the complexity of the process, making experimental evaluation essential.

A critical parameter in optimizing separation is the selectivity (α), also called the separation factor [69]. Selectivity measures the ability of the stationary phase to differentiate between analytes based on their interactions. It is defined as the ratio of retention factors of the two components, in our application krypton and xenon, and is given by [69]:

$$\alpha = \frac{k_{\text{Xe}}}{k_{\text{Kr}}},$$

where k_{Kr} and k_{Xe} are the retention factors of krypton and xenon, respectively. Retention factors are defined as:

$$k_{\text{Kr}} = \frac{t_{\text{R}}^{\text{Kr}} - t_0}{t_0}, \quad k_{\text{Xe}} = \frac{t_{\text{R}}^{\text{Xe}} - t_0}{t_0},$$

where t_{R}^{Kr} and t_{R}^{Xe} are the retention times of krypton and xenon, respectively, and t_0 is the unretained time, i.e., the time required for a non-interacting compound to pass through the column. Notably, all compounds spend the same amount of time in the mobile phase, equivalent to t_0 . The selectivity is always greater than or equal to 1; if it equals 1, the

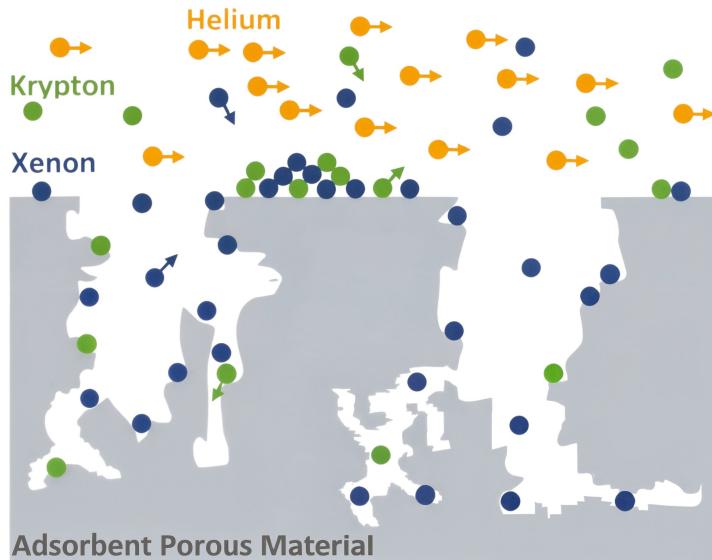


Figure 2.2: Schematic of gas-solid chromatography where krypton and xenon atoms interact with porous adsorbent material during transport by the inert helium carrier gas.

components are not separated. This parameter is highly influenced by temperature and the choice of adsorbent.

While selectivity is important for distinguishing between analytes, the quality of the separation is ultimately quantified by the chromatographic resolution (R). Resolution quantifies the degree of separation between two chromatographic peaks, considering both differences in retention time and the effects of peak broadening and shape. In the general case, including asymmetric peaks, resolution is given by [69]:

$$R = \frac{t_R^{Xe} - t_R^{Kr}}{1.7 (w_{\text{tail}}^{Kr} + w_{\text{front}}^{Xe})},$$

where t_R^{Kr} and t_R^{Xe} denote the retention times of krypton and xenon, respectively, and w_{tail}^{Kr} and w_{front}^{Xe} represent the widths at half-maximum of the tail and front of the corresponding peaks. The factor 1.7 adjusts for the use of half-maximum widths instead of full widths at half maximum. Resolution is not only a function of selectivity but also of the efficiency of the chromatographic system, which can be expressed in terms of the number of theoretical plates (N) and the retention factors. The comprehensive relationship between resolution and these factors is given by [69]:

$$R = \frac{\sqrt{N}}{4} \left(\frac{\alpha - 1}{\alpha} \right) \left(\frac{k_{Xe}}{1 + k_{Xe}} \right),$$

where N is the number of theoretical plates, representing column efficiency, and k_{Xe} is the retention factor of xenon. A large alpha value yet with low resolution frequently indicates that it is feasible to distinguish the two compounds within the system, but the parameters are not optimized. In practice, improving resolution can involve increasing the selectivity by choosing an appropriate adsorbent or adjusting operational parameters like temperature. Additionally, increasing column efficiency (i.e., increasing N) or optimizing the retention factors can further enhance separation.

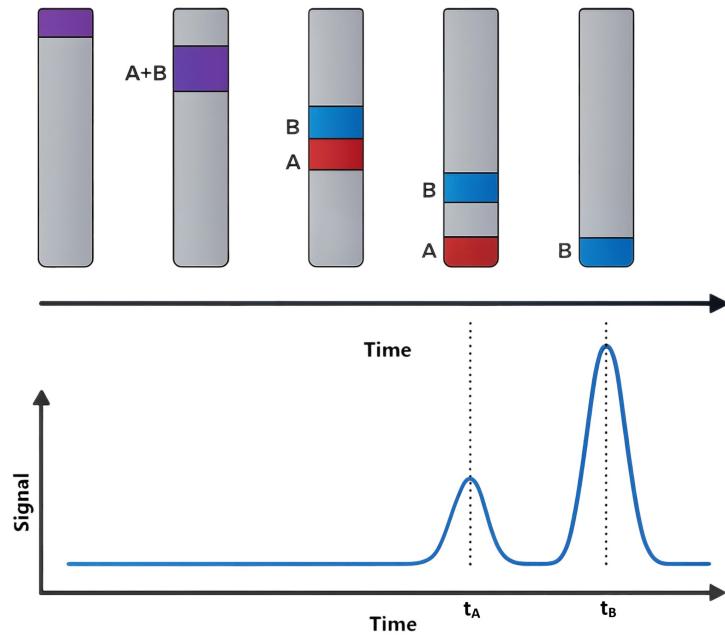


Figure 2.3: Visualization of analyte separation in a chromatography column. As the initial mixture progresses, its components separate into different bands, corresponding to two resolved peaks in the detector with distinct retention times.

The degree of peak asymmetry is characterized by the asymmetry factor (A_s), defined as:

$$A_s = \frac{b}{a},$$

where a represents the front half-width and b the back half-width of the peak at 10% of its maximum height.

2.5 Experimental setup for adsorbents testing

A prototype of an automated GSC system was developed to assess the feasibility and automation principles of the Auto-RGMS project. This demonstrator proved capable of maintaining temperature, pressure, and flow rate stability within one percent for over 10 hours [68], ensuring highly consistent gas separation processes. The demonstrator is employed here to assess new adsorbent materials; its key components are summarized below and shown in figure 2.4.

- Adsorbent column with temperature regulation: The adsorbent column is a U-tube-shaped stainless steel structure containing the test adsorbents. It is enclosed within a copper cylinder composed of two vertical sections, which are clamped together and secured with screws. A key aspect of its design is the precise temperature control, covering a range from -196°C (liquid nitrogen) up to 200°C . This is achieved through an automatically refilling LN_2 dewar, regulated using a magnetic valve (MV1) and monitored by a load cell. Two 350 W heaters provide counter-heating, ensuring that the full temperature range is accessible. To prevent instability caused by gaseous nitrogen bubbles forming at the surface, the copper

cylinder is insulated with a 5 mm-thick PTFE mantle. This thermal insulation enhances the reproducibility of the chromatographic process by preventing direct contact between the nitrogen bath and the copper surface, stabilizing the column temperature. Two PT1000 temperature sensors (PT1–2) continuously monitor the column temperature.

- Sample loading apparatus: Samples are precisely injected using a glass pipette featuring two calibrated volumes ($V_1 = (3.560 \pm 0.005) \text{ cm}^3$, $V_2 = (276.580 \pm 0.007) \text{ cm}^3$), along with three manual valves (V0–V2).
- Supply of carrier gas: High-purity (grade-6) helium from a gas cylinder is delivered at a stable rate, regulated by a mass flow controller (MFC) and tracked with a mass flow meter (MFM).
- Vacuum control system: The downstream vacuum is sustained by a scroll pump. Pressure sensors P1 (covering the entire range) and P2 monitor the pressure upstream, while P3 monitors downstream. Additionally, a control valve (CV1) ensures precise adjustments of flow and pressure at the outlet.
- Gas detection mechanism: Gas detection is performed by a thermal conductivity detector (TCD), consisting of two thermistor-containing tubes designed for incoming and outgoing gases. While the TCD's sensitivity is substantially lower compared to the RGMS mass spectrometer, it offers operational simplicity and sufficient detection capabilities for demonstration purposes.
- Data acquisition: A compact data acquisition (DAQ) system interfaces sensors and actuators, with operations managed by custom software developed in LabVIEW.

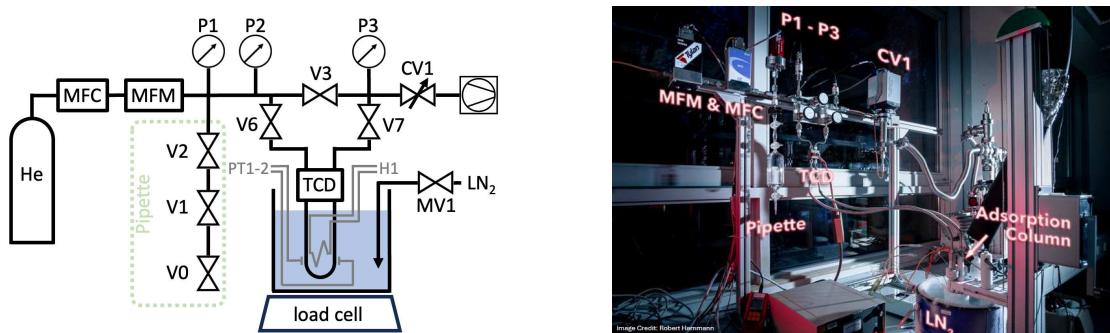


Figure 2.4: Overview of the Auto-RGMS demonstrator used for adsorbent evaluation. The system schematic, with labeled components, is presented on the left, and a photograph of the setup is shown on the right.

2.6 Adsorbents for krypton-xenon separation

In chromatography, extending the separation column's length is a common method to improve chromatographic resolution, as resolution generally scales with the square root of the column length. However, in the Auto-RGMS system, the column length from its RGMS predecessor has been preserved due to strict purity constraints. Achieving high purity after each measurement requires prolonged heating and evacuation procedures to remove residual gases trapped within the adsorbent. If residual gases remain, a "memory

effect" occurs, resulting in elevated background during subsequent measurements. This phenomenon is discussed further in section 2.8. Longer columns worsen the memory effect, as xenon, the dominant component, tends to become trapped in the adsorbent's porous structure, prolonging the cleaning process of baking, pumping, and flushing, and potentially affecting subsequent analyses. Additionally, column length optimization is limited by the availability of suitable hardware variations to test.

In light of these challenges, efforts to enhance chromatographic separation have focused on selecting higher-performing adsorbents and optimizing key operational parameters, such as carrier gas flow and column temperature. Several porous adsorbent candidates for the Auto-RGMS system were evaluated using the Auto-RGMS prototype described in section 2.5, with the goal of surpassing the performance of Chromosorb 102, currently adopted in RGMS. Chromosorb 102, an organic compound characterized by relatively large pore sizes, served as the benchmark for comparing the performance of the other tested materials. Among the organic polymers, two variants from the HayeSep family, specifically HayeSep Q and HayeSep D, were selected due to their frequent recommendation in chromatographic applications as improved replacements for Chromosorb 102, reflecting recent developments in organic porous materials. ShinCarbon-ST was also tested, standing out as a rare but promising carbon molecular sieve suitable for chromatography. Additionally, the inorganic silica MCM-48 was examined for its expected structural stability and favorable pore dimensions. Organic adsorbents tend to degrade more easily and can release contaminants through outgassing. In contrast, MCM-48 exhibits strong resistance to environmental variations, including changes in temperature and humidity, which could allow for higher baking temperatures, potentially improving and accelerating the cleaning phase of the measurement process. A summary of the main properties of the tested materials is provided in table 2.1.

Adsorbent	Surface Area [m ² /g]	Max Temp. [°C]	Avg. Pore Diameter [nm]	Density [g/mL]
Chromosorb 102	350	230 [*]	8.5	0.30
HayeSep D	795	230 [*]	30.0	0.33
HayeSep Q	582	260 [*]	7.5	0.35
ShinCarbon-ST	1500	280	0.5–1.5	0.51
Silica MCM-48	1400–1600	>500	3.0	0.5–1.0

Table 2.1: Properties of the tested adsorbents, listing Chromosorb 102 (Chromatography Research Supplies, Inc.), ShinCarbon-ST (Shinwa Chemical Industries LTD), HayeSep Q (Supelco, catalog 10300-u), HayeSep D (Supelco, catalog 10291), and mesoporous silica MCM-48 (Sigma-Aldrich, catalog 805467-5G). Maximum temperatures noted with ^{*} are based on experimental observations where the materials started showing a yellow tint, rather than manufacturer specifications. See section 2.8 for more details. Data also referenced from [70] and [71].

Each measurement presented in this section followed the same procedure, as outlined below:

1. Column preparation: Load approximately 2 grams of adsorbent material into the stainless steel U-tube. Secure the adsorbent by placing small plugs of glass wool at both ends of the powder bed to prevent material displacement. Heat the column to 150°C and flush with helium (20 cm³ STP/min, P3 = 1200 mbar) for at least 10 minutes.
2. Cool-down: Immerse the column slowly in LN₂ until temperature stabilizes at approximately -196 °C.

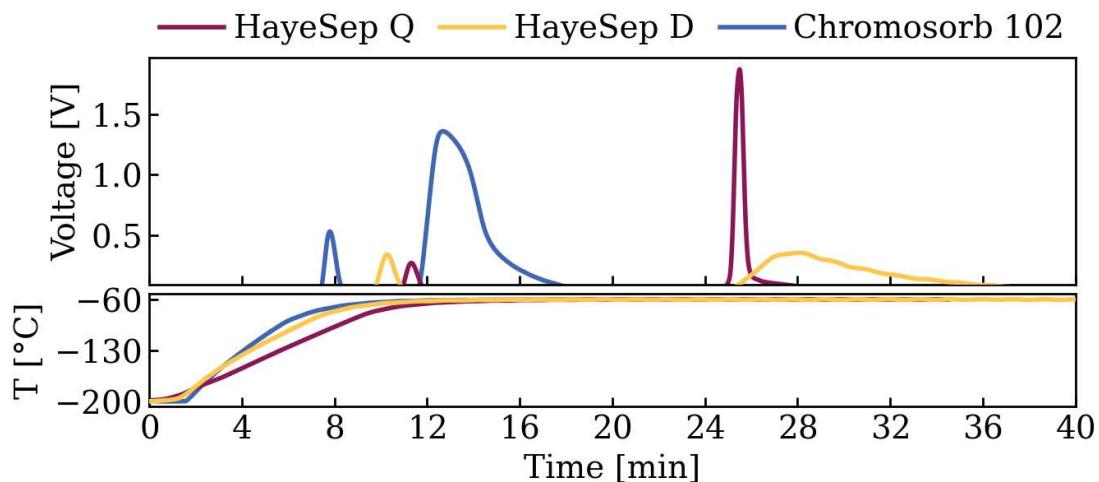


Figure 2.5: Elution profiles from the Auto-RGMS demonstrator at -60°C , comparing the separation of krypton and xenon for HayeSep Q, HayeSep D, and Chromosorb 102 adsorbents.

3. Sample injection: Inject sample gas ($2\text{--}3\text{ cm}^3$ STP) into calibrated volume ($V_{\text{cal}} = (51.3 \pm 0.4)\text{ cm}^3$). Gas samples have a krypton-to-xenon ratio (Kr:Xe = 1:10) to match the detection sensitivity of the TCD.
4. Sample freezing: Freeze the injected sample onto the cold adsorbent column. Allow approximately 10 minutes until the pressure stabilizes.
5. Separation initiation: Establish helium carrier gas flow at 150 sccm with $P_3 = 1200$ mbar, begin column warm-up, and enable TCD signal acquisition.
6. Chromatogram acquisition: Maintain stable temperature using proportional-integral-derivative (PID) control, record chromatogram from TCD output.
7. System reset: After xenon elution, heat the column back to 150°C and flush with helium at 150 sccm for approximately 10 minutes to remove residual gases and prepare the system for the subsequent measurement.

Typical chromatographic profiles obtained using adsorbent materials HayeSep Q, HayeSep D, and Chromosorb 102 are displayed in figure 2.5. The chromatograms distinctly highlight the elution order of krypton, appearing first, followed by xenon. Differences in the shapes and retention positions of these peaks reflect each adsorbent's separation properties. The chromatograms offer key insights into peak symmetry and tailing, both of which significantly affect chromatographic resolution. Notably, xenon exhibits substantial tailing in the case of Chromosorb 102 and HayeSep D. The observed effect is likely due to xenon's distinct adsorption properties on these materials, which provide stronger and more varied binding sites than those available to krypton [72], as well as to column overloading, where xenon exceeds its linear adsorption regime [69]. A stronger xenon peak tailing effect was noted at larger sample sizes, along with a corresponding reduction in retention time [65, 68]. Krypton, however, maintained consistent peak symmetry and stable retention times, in agreement with the expectation that symmetrical peaks are unaffected by varying analyte concentrations [73]. While smaller sample injections could address nonlinearities [69], this strategy does not align with the needs of Auto-RGMS, where larger sample sizes are essential for improved sensitivity.

2.7 Effect of temperature and helium flow on resolution

The resolution measurement results for the tested adsorbents are summarized in table 2.2.

Adsorbent	T_T [°C]	F_T [sccm]	T_S [°C]	F_S [sccm]	t_R^{Kr} [hh:mm:ss]	A_S^{Kr}	t_R^{Xe} [hh:mm:ss]	A_S^{Xe}	R
Chromosorb 102	-60	150	-62(3)	150.09(4)	0:07:48(1)	1.37(3)	0:12:41(5)	4.3(3)	3.1(2)
	-75	150	-76(1)	149.96(5)	0:09:41(1)	1.42(3)	0:19:40(10)	9.0(6)	3.7(3)
	-90	150	-90.0(3)	149.79(8)	0:08:49(2)	1.6(5)	0:43:20(80)	7.7(2)	4(1)
HayeSep D	-60	150	-60(1)	150.53(3)	0:10:17(1)	1.25(3)	0:28:00(10)	3.7(2)	4.5(3)
	-60	100	-60.1(9)	99.50(8)	0:09:21(1)	1.66(4)	0:38:14(8)	3.7(2)	5.7(2)
	-60	50	-60.0(2)	50.35(6)	0:16:06(2)	1.32(6)	1:10:50(30)	N/A	6.6(6)
	-75	150	-75.0(3)	153.71(8)	0:10:41(1)	1.22(3)	1:01:56(40)	6.2(6)	6.0(8)
HayeSep Q	-90	150	-89.9(2)	149.90(9)	0:15:00(3)	1.31(5)	3:22:00(160)	N/A	8(1)
	-60	150	-61(2)	149.92(9)	0:11:19(1)	1.32(3)	0:25:30(4)	1.1(2)	15(2)
	-75	150	-75.3(4)	150.10(3)	0:14:24(1)	1.32(2)	0:45:05.3(4)	0.91(2)	27.7(5)
ShinCarbon-ST	-90	150	-90.0(2)	151.91(5)	0:12:53(2)	1.41(4)	1:43:19.3(3)	1.27(3)	51(1)
	0	150	0(1)	149.74(5)	0:13:17(2)	1.7(2)	0:37:10(20)	12(4)	14(4)
Silica MCM-48	-75	15	-75.00(4)	16(1)	0:15:29(4)	1.5(5)	0:59:00(30)	5.0(3)	3.7(2)

Table 2.2: Comparison of separation metrics for adsorbent candidates tested with the Auto-RGMS demonstrator. The table presents the target temperature (T_T) and helium carrier gas flow rate (F_T), as well as the actual values achieved during operation (T_S and F_S). It further lists the measured retention times for krypton (t_R^{Kr}) and xenon (t_R^{Xe}), the asymmetry factors for krypton (A_S^{Kr}) and xenon (A_S^{Xe}), and the resulting chromatographic resolution (R). Reported uncertainties in operational parameters are the sample standard deviation of values observed between the krypton and xenon peaks.

Figure 2.6 illustrates the temperature-dependent behavior of the krypton/xenon chromatographic resolution for the different adsorbents. In general temperature influences the adsorption strength of krypton and xenon onto the stationary phase, increasing the retention times. Resolution is usually enhanced, except when peak broadening is so significant that it cancels out the improvement in separation. The data reveal that HayeSep Q surpasses other adsorbents not only in its separation performance at set temperatures, but also demonstrates remarkably enhanced performance with decreasing temperature conditions. While the improvement slope of HayeSep Q is very pronounced, indicating a significant enhancement in performance with decreasing temperature, HayeSep D also shows relevant improvements, though to a lesser extent. In contrast, the performance improvement for Chromosorb 102 is very limited, highlighting its lower effectiveness compared to the other adsorbents. This occurs because HayeSep Q remains in the linear regime for our sample size, leading to greater symmetry and sharpness of the xenon peak, which is clearly visible in figure 2.5, while this is not the case for HayeSep D and Chromosorb 102.

While ShinCarbon-ST achieved favorable separation at higher temperatures (see table 2.2), its memory effect, discussed in section 2.8, disqualified it for Auto-RGMS. For this reason, no further measurements at lower temperatures were taken. Tests with Silica MCM-48 demonstrated that the commercially available grain size was inadequate, as it failed to sustain the standard helium flow rate of 150 sccm without causing a significant pressure increase before the column. To achieve normal pressure conditions, the flow rate had to be reduced to 15 sccm (see table 2.2). However, even though this reduction should have improved resolution, as seen for HayeSep D in figure 2.7, the performance was still inferior to the other powders tested.

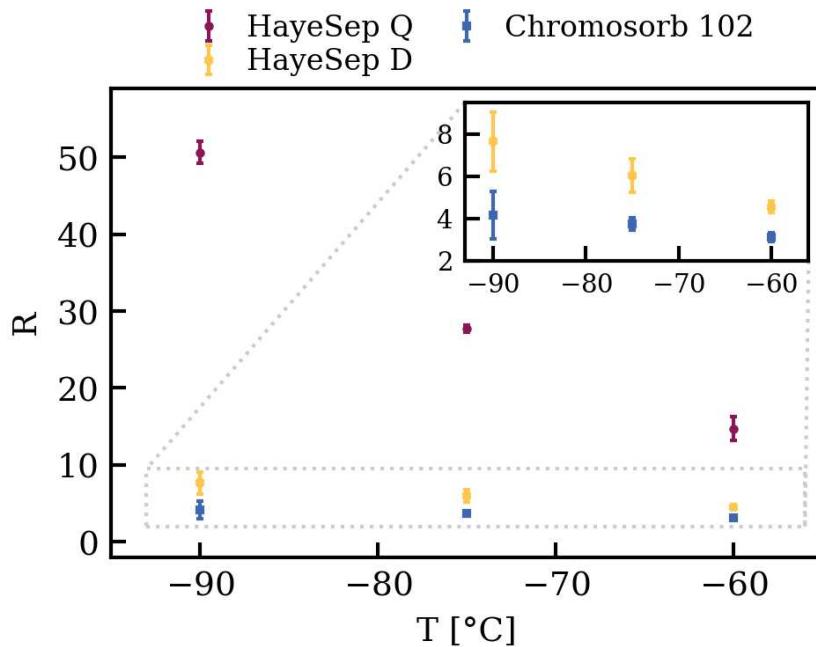


Figure 2.6: The effect of separation temperature on chromatographic resolution (R) across the adsorbents shown in table 2.1, highlighting the superior performance of HayeSep Q. ShinCarbon-ST and Silica MCM-48 are not included since their tests were performed under different experimental conditions. Further information is provided in the main text.

The separation performance is strongly dependent on the helium flow rate. In general, increasing the helium flow rate pushes compounds through the column more quickly, reducing retention times, but often limiting the time available for complete adsorption-desorption equilibrium and consequently lowering resolution, whereas a lower flow rate prolongs retention times and can improve resolution unless it becomes so low that diffusion-driven peak broadening dominates. Figure 2.7 illustrates how the resolution of HayeSep D at $-60\text{ }^{\circ}\text{C}$ decreases in a linear manner as the helium flow increases. In the case under examination, at a flow rate of 150 sccm, the xenon peak appears after approximately 28 minutes. Lowering the flow to 50 sccm shifts the xenon peak to 1 hour and 11 minutes, which remains a reasonable duration for our objectives. Additionally, in this flow range, it was not observed a saturation of resolution improvement with decreasing flow rate due to the peak broadening effect. The helium flow rate in Auto-RGMS will be set according to the final operational and performance constraints, with a preference for the lowest feasible value to achieve superior chromatographic resolution.

2.8 Memory effect evaluation

To ensure accurate measurements using RGMS, a standard procedure involves eliminating any residual gas from the adsorbent material after each analysis. This is achieved through a combination of baking, vacuum pumping, and purging with purified helium. Baking temperatures were adjusted based on the adsorbent type. For instance, organic HayeSep materials are sensitive to heat and exhibit significant outgassing at higher temperatures. Furthermore, excessive heating can lead to a color change from white

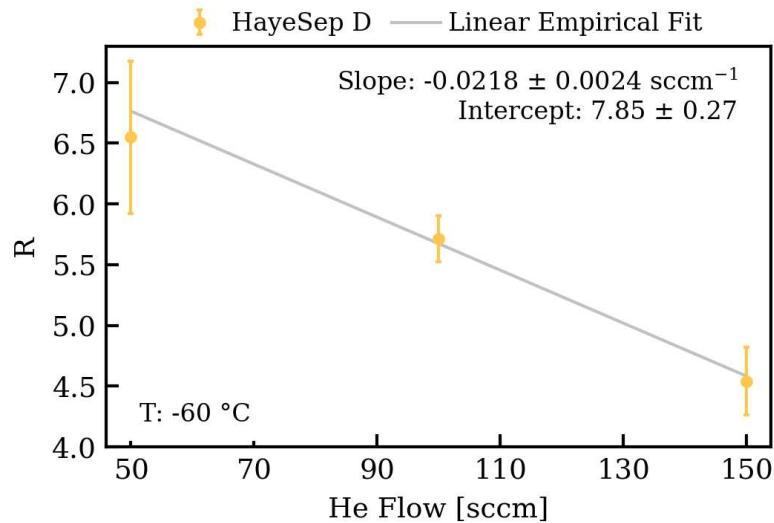


Figure 2.7: Effect of helium carrier gas flow rate on the chromatographic resolution for the HayeSep D adsorbent at -60 °C.

to yellow. Although literature suggests this color shift does not alter chromatographic properties [70], we interpret it as a sign of chemical composition changes, which is undesirable given our system's stringent purity needs. Considering these factors, HayeSep D and HayeSep Q were baked at 110°C and 130°C, respectively. ShinCarbon-ST, as a carbon-based material, was able to withstand baking at 250°C. MCM-48 was not tested for memory effects, as it is not a possible candidate for Auto-RGMS. Direct RGMS measurements were conducted to effectively evaluate the memory effect. This involved connecting the identical adsorbent columns used in separation experiments directly to the RGMS system through a dedicated spare port at the chromatography stage's outlet. Adsorbent cleaning for these tests was limited to baking and vacuum pumping, as helium flushing was not feasible due to the spare port's configuration. Consequently, the results from these measurements should be considered conservative, as implementing helium flushing would drastically further reduce the observed memory effect. The procedure began with a blank measurement following baking and pumping to confirm that the column achieved a sufficiently low background. Subsequently, the column was cooled to LN₂ temperature and exposed to a calibration gas, as described in [74], containing known quantities of krypton, xenon, and argon. After a one-hour adsorption period, the column was sealed and heated for an hour to ensure even distribution of the calibration gas across the adsorbent. To mimic a standard cleaning cycle, the column was then subjected to vacuum pumping and baking for 24 hours. Following this, the column was closed, and any residual calibration gas outgassing was collected for 40 minutes, mirroring the acquisition time for XENON sample measurements. Finally, RGMS analysis was performed on the collected gas. A persistent high signal would suggest a significant memory effect. In such cases, further cycles of baking and pumping are performed, followed by repeated measurements to quantify the memory effect's magnitude. This methodology presented several challenges. Firstly, achieving a leak-tight system and then reducing the background to levels suitable for RGMS analysis after air exposure proved difficult. Secondly, the availability of time slots for memory effect measurements was limited because they had to be scheduled around the regular RGMS operations for the XENONnT experiment, resulting in a limited number of data points. Figure 2.8 displays the blank-corrected results of these memory effect measurements for

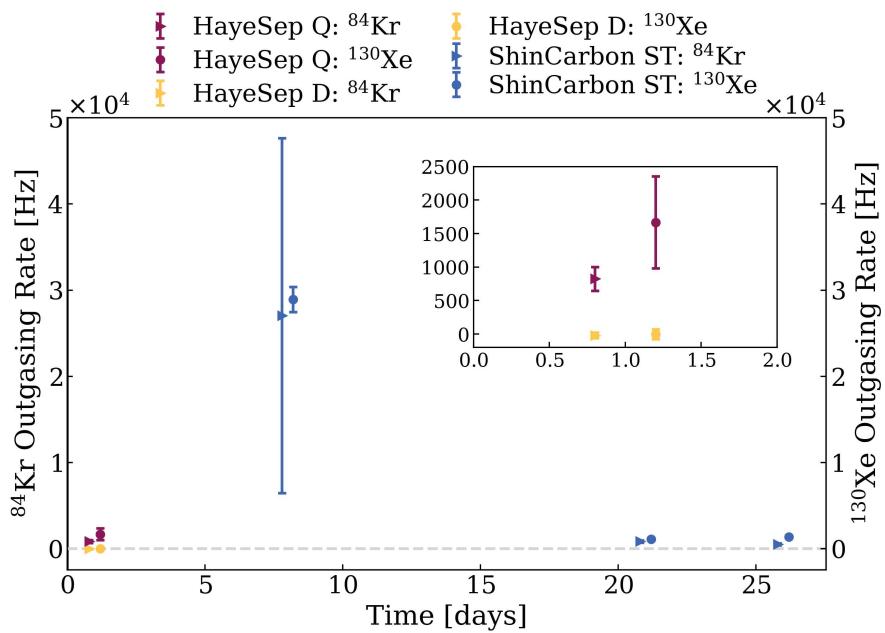


Figure 2.8: Assessment of memory effects for various adsorbents. Residual outgassing rates for krypton (^{84}Kr , left y-axis, triangles) and xenon (^{130}Xe , right y-axis, circles) from different adsorbent materials are shown. The x-axis shows the duration of heat and pumping after calibration gas injection, with intentional offsets applied between ^{84}Kr and ^{130}Xe data points to improve clarity.

HayeSep Q, HayeSep D, and ShinCarbon-ST, plotted against pumping duration. For each data point, error bars are calculated based on conservative estimates of systematic uncertainties in detector response and time-dependent variability associated with each adsorbent. The ShinCarbon-ST measurement exceeded the RGMS detection limit after the first day and was stopped to prevent potential damage to the mass spectrometer detector.

The ^{84}Kr measurement serves as an indicator of typical RGMS performance. Incomplete removal of ^{84}Kr between sample analyses would lead to a significant krypton contamination in subsequent measurements. Due to the high xenon abundance in samples, $^{130}\text{Xe}^{2+}$ was measured in those tests to reduce the signal rate and protect the detector. Furthermore, significant xenon presence in the mass spectrometer can interfere with the krypton signal. This detector-related phenomenon is evident in the large error bar for ShinCarbon-ST in figure 2.8. Significant memory effects, originating from either krypton or xenon retention in an adsorbent, present a major limitation for its suitability in Auto-RGMS. The data show that ShinCarbon-ST, known for its high microporosity, exhibited a persistent memory effect and could not reach the required blank background level even after a month of baking and pumping, making it unsuitable for Auto-RGMS. Conversely, HayeSep D demonstrated a negligible memory effect. While measurements for HayeSep Q were prematurely terminated due to altered RGMS conditions, the data indicated a discernible memory effect, which detracted from its otherwise promising performance. It is important to reiterate that these memory effect tests, conducted without helium flushing, represent a conservative assessment. The Auto-RGMS system, however, is designed with a continuous purified helium supply, ensuring helium flushing during adsorbent column preparation for each measurement. A dedicated test, utilizing a temporary helium supply to flush HayeSep Q for a workday, confirmed that helium flushing can effectively reduce its memory effect to negligible

levels. With the continuous flushing capability of Auto-RGMS, we anticipate achieving blank background levels more rapidly. Despite the sub-optimal conditions of the test, the findings suggest that HayeSep Q is the most suitable adsorbent for the Auto-RGMS application.

2.9 Conclusion and discussion

Next-generation low-background LXe experiments, such as DARWIN/XLZD, require robust, high-frequency krypton monitoring with exceptional low-ppq sensitivity. This is crucial for quantifying the beta-decaying isotope ^{85}Kr , as it represents a major electron recoil background source in low-energy studies. This level of sensitivity is indispensable for achieving scientific objectives like the precise determination of the solar pp neutrino flux and the measurement of the weak mixing angle at unexplored low energies. This chapter presented R&D for the Auto-RGMS system, an automated rare gas mass spectrometer engineered to meet these stringent requirements. Auto-RGMS represents a substantial advancement over its manually operated predecessor RGMS through complete automation of operational procedures and precise control of critical parameters including temperature and carrier gas flow. Combined with optimized surface design, these improvements reduce systematic uncertainties, stabilize backgrounds, and eliminate manual operational requirements. A key limitation of the current system is that the injectable sample size is constrained by the poor chromatographic resolution of the used adsorbent material. A larger sample size provides more krypton in the post-separation sample, allowing the mass spectrometer to detect lower concentrations of krypton, thereby enhancing the overall measurement sensitivity. The systematic evaluation of adsorbent materials identified HayeSep Q as optimal for krypton-xenon separation, delivering up to a 12-fold improvement in chromatographic resolution compared to the previously used Chromosorb 102. This superior resolution directly addresses the primary limitation on injectable sample size. Furthermore, the implementation of a new custom all-metal double-valve system significantly mitigates the systematic uncertainty associated with calibration gas injection. The influence of separation temperature and helium flow on chromatographic resolution was also studied, providing important data for the optimisation of Auto-RGMS operational parameters. While HayeSep Q exhibits memory effects, these are proven to be effectively mitigated through continuous helium flushing integrated into the Auto-RGMS design. The Auto-RGMS system is currently undergoing commissioning at MPIK in Heidelberg, representing a critical step toward deployment in future xenon observatories.

Chapter 3

Data analysis in XENONnT

3.1 Data processing pipeline and event building

Reconstructing a fully characterized (S1, S2) event in the XENONnT experiment from raw photomultiplier tube (PMT) signals involves a multi-stage process. The XENONnT data acquisition system (DAQ) runs without a global hardware trigger and instead uses independent, per-PMT self-triggers. Any PMT pulse that crosses its channel-specific threshold is digitized and stored as a waveform, which dynamically extends for the duration that the pulse remains above threshold and includes both pre- and post-threshold samples. For over 98% of TPC channels, the threshold is set at 2 mV, ensuring an average single-photoelectron (PE) acceptance above 90% [75]. The lowest level of stored data is referred to as **raw-records**.

Offline software reconstruction begins by retrieving the raw-records from all channels. The initial step, hit finding, selects a time window of 30 ns before and 200 ns after the interval during which the signal remains above the threshold. Next, **peaklet** building clusters these per-PMT hits. Neighboring hits from any PMT are grouped into clusters if their temporal separation is 700 ns or less. Isolated hits that do not meet this criterion are classified as **lone hits**, primarily attributed to PMT dark counts or after-pulses. Peaklets initially formed may contain overlapping components. A sub-clustering step, known as peaklet splitting, is performed using a natural break algorithm. This involves analyzing a summed waveform of the peaklet's PMT hits and splitting it at local minima if the drop in pulse amplitude exceeds predefined thresholds, producing refined peaklets that better represent distinct physical signals. Further details can be found in [34, 75]. After peaklet formation, signals are classified as S1 or S2 based on their distinct waveform characteristics. The classification leverages several shape and spatial distribution features of peaklets along the data-driven classification boundaries shown in figure 3.1.

- **Rise time:** The time interval between 10% and 50% of the integrated area in the waveform. S1 signals exhibit faster rise times due to prompt light emission, while S2 signals have slower rise times reflecting the electron drift and extraction process.
- **Area fraction top (AFT):** The ratio of signal detected by the top PMT array to the total signal area. S1 signals yield low AFT values as scintillation light is primarily detected by the bottom array due to total internal reflection of photons at the liquid-gas interface. S2 signals produce higher AFT values as proportional scintillation occurs in the gas phase directly beneath the top PMT array.
- **Tight coincidence (TC):** The number of distinct PMTs registering hits within a ± 50 ns window centered on the peaklet's maximum amplitude. This criterion substantially reduces the rate of false S1 identifications arising from random coincidences of PMT dark counts, which might otherwise then be paired with

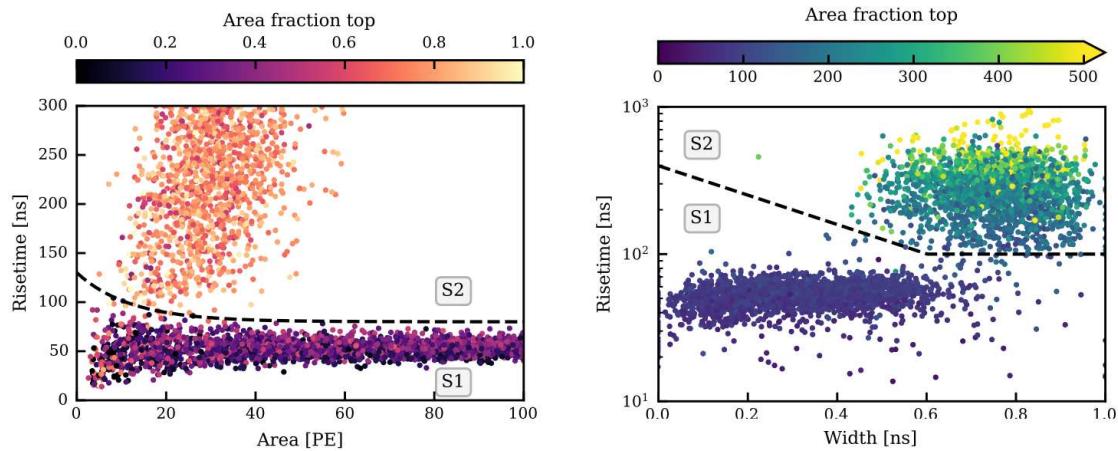


Figure 3.1: Classification boundaries for S1 and S2 signals in the first science run. Left panel: Rise time (time between 10% and 50% area quantiles) versus peak area, with color indicating area fraction top (fraction of signal detected by top photomultiplier tube array). Dashed lines represent the classification boundaries for S1 and S2 signals. Data shown include S1 signals from ^{220}Rn calibration and single-electron S2 signals. Right panel: Same dataset plotted as AFT versus rise time, with color showing the width (central 50% of integrated area). Figure published in [34].

isolated S2 signals. S1 classification requires $\text{TC} \geq 2$, but standard analyses only consider S1 signals with $\text{TC} \geq 3$, a threshold that sets the S1 detection efficiency and energy threshold at low energies (see section 3.2).

- **Area:** the total gain-corrected charge integrated across all PMT hits within the peaklet's duration.
- **Width:** The time interval containing the central 50% of the integrated signal area.
- **Number of contributing PMTs:** Peaklets with fewer than 5 contributing PMTs that fail to meet S1 classification criteria are categorized as unidentified. This helps in identifying potential dark count pile-up.

Following peaklet classification, S2 peaklets require further processing known as S2 peaklet merging. Electron drift causes longitudinal diffusion that temporally extends S2 signals. This phenomenon can cause signals from just a few electrons, separated by microseconds, to be incorrectly divided during initial signal processing. To accurately capture the original electrons from the same electron cloud, neighboring S2 peaklets are evaluated for potential aggregation. Merging occurs when the time gap meets specific criteria, which is dynamically adjusted based on the integrated area of the resulting S2 signal, see [34, 76] for additional details. The procedure continues iteratively until either the total duration of the newly formed S2 would exceed a defined maximum of 50 μs , or no other proximal S2 peaklets are identified for combination. The time limit mitigates the inclusion of unrelated signals from undesired processes like delayed electron extraction, where electrons temporarily trapped at the liquid-gas interface are released with significant time delay, or xenon impurity-induced photoionization, which arises from photoemission due to xenon impurities. After the S2 peaklet merging process is complete, the resulting S1 and S2 signals are collectively termed **peaks**.

The next critical step is *event building*, where individual S1 and S2 peaks are temporally associated to form **events**, representing the hypothesis of a single particle interaction within the TPC. Events are constructed around an S2 peak, which is selected if its integrated area exceeds 100 PE and fewer than 8 neighboring peaks with area greater

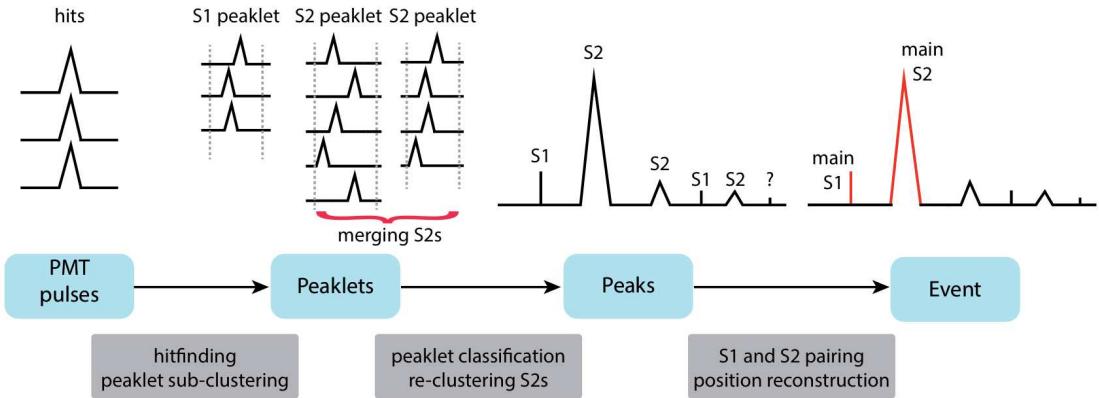


Figure 3.2: Schematic of the straxen reconstruction process. Starting from photomultiplier tube pulses, intermediate data objects (blue) are processed (gray) to produce (S1, S2) events. Unclassified peaklets are marked with "?". Figure published in [34].

than 50 PE are present within a ± 10 ms window. Since the S1 signal precedes the S2 due to electron drift time within the TPC, a broad temporal window of 2.45 ms before and 0.25 ms after the S2 trigger is utilized for event building, with the latter buffer ensuring the full S2 pulse and capturing potential secondary signals, such as a second S2 from multiple scatters, without truncation. Within each defined event window, the main S1 is identified as the largest S1 peak. In instances where event windows overlap, they are subsequently merged to ensure a unique temporal segmentation of the data. After event building, the second-largest S1 and S2 peaks are identified as alternative signals, with the alternative S2 also required to follow the main S1 temporally. This identification of alternative signals is particularly valuable for analyzing multi-scatter events, where a particle interacts at multiple sites within the detector volume, producing multiple S1-S2 pairs.

3.2 S1 detection efficiency and peak reconstruction bias

A key criterion for accepting an S1 signal in the analysis is the tight coincidence (TC) requirement: a minimum number of PMTs (typically $TC \geq 3$ in standard analyses) must detect photons within a short coincidence window of $\Delta t = 100$ ns. This requirement effectively suppresses backgrounds from random coincidences of PMT dark counts, which are particularly relevant for isolated S1 signals at low energies, where signals generate few detected photons. The S1 reconstruction efficiency for genuine S1 signals, primarily governed by the TC requirement, is determined through two cross-validating methods: full-chain simulation and a data-driven technique. The simulation approach utilizes a detailed model of S1 pulse generation, including scintillation physics, light propagation, and PMT response. The S1 reconstruction efficiency represents the probability that a simulated S1 signal, conditional on generating a particular number of detected photons in PMTs, will be properly reconstructed by the analysis pipeline described in section 3.1. Figure 3.3 shows the result for the two methods and their agreement.

The reconstruction of signal peak areas is subject to inaccuracies, quantified by the relative difference between the reconstructed area and the true area from simulations; these systematic deviations are referred to as peak reconstruction bias. They are observed for both S1 and S2 signals and stem from several factors. Underestimation, particularly for smaller signals, can result from insufficient single PE amplification and signal

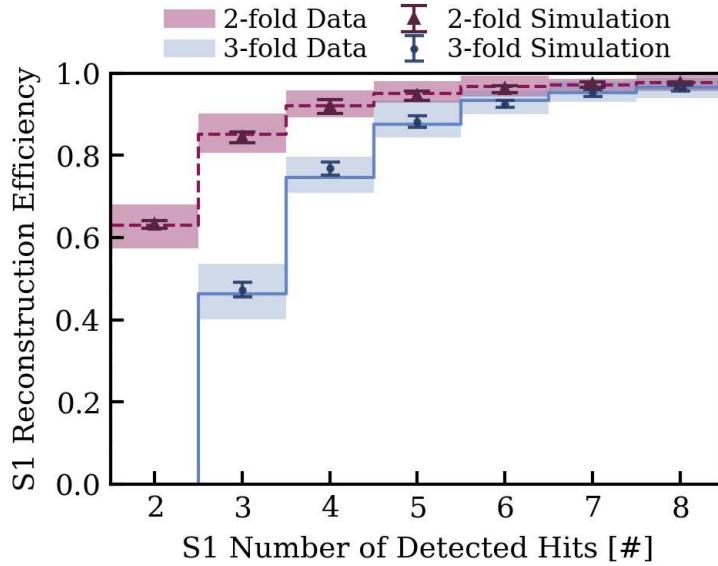


Figure 3.3: S1 reconstruction efficiency is shown as a function of the number of detected S1 signal hits, for selection criteria requiring a coincidence of at least 2 (2-fold) or 3 (3-fold) photomultiplier tubes within ± 50 ns of the S1 peak. Estimations from data-driven analysis and full-chain simulation are shown, demonstrating their agreement. The shaded bands and error bars represent 68% confidence intervals.

thresholds (DAQ and hit-finding), leading to a negative shift. Conversely, phenomena like PMT afterpulsing and photoionization, when merged into the main signal by the reconstruction, inflate the reconstructed area, causing a positive shift. Electronic noise introduces random area distortions, and the reconstruction algorithm may incorrectly define peak boundaries. The systematic bias, characterized by the median of the error distribution, was evaluated using simulations of S1 and S2 signals uniformly distributed in the detector. These simulations reveal distinct bias trends and associated uncertainties. For S1 signals, the bias is negative, around -2%, at low energies below 50 PE, where a wide 1σ uncertainty of $\sim 6\%$ reflects the impact of noise and limited statistics. It transitions to a positive value at higher energies above 10^3 PE due to afterpulse inclusion, with the uncertainty band narrowing to $\sim 1\%$. S2 signals exhibit analogous behavior: a bias of approximately -1% with a $\sim 2\%$ uncertainty band is observed for signals below 10^3 PE, while it becomes typically +0.5% or less at high energies above 10^4 PE with the uncertainty band decreasing below $\sim 1\%$. A more detailed description and investigation of these effects is provided in [77].

3.3 Vertex reconstruction and field distortion corrections

Accurate 3D position reconstruction of interaction vertices within the TPC is crucial for signal corrections (see section 3.4) and background reduction and modeling (see section 3.5 and 4.4), as it permits fiducialization (selecting a low-background inner volume) and allows position-dependent corrections to S1 and S2 signals. The vertical position (Z) is determined from the drift time (t_{drift}) between the S1 and S2 signals. This time difference is proportional to the depth of the interaction and is converted to a physical distance using the measured in-situ electron drift velocity of (0.675 ± 0.006)

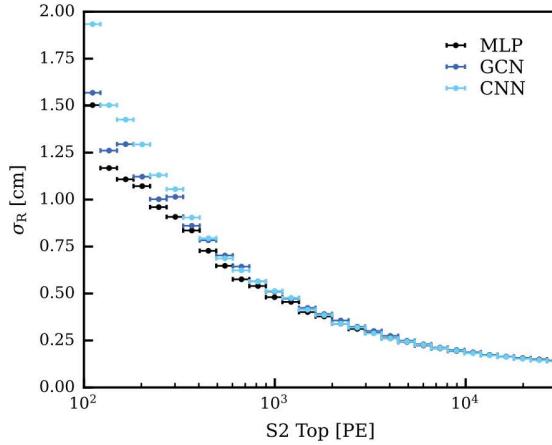


Figure 3.4: Simulation-based evaluation of radial position resolution for three machine learning vertex reconstruction algorithms: multilayer perceptron (MLP), convolutional neural network (CNN), and graph constrained network (GCN), within the inner detector volume (radius < 60 cm). Figure published in [34]

mm/ μ s [34] at a drift field of 23 V/cm. The reference depth ($Z = 0$ cm) is defined at the bottom of the gate electrode, with the cathode located at $Z = -148.65$ cm.

The horizontal coordinates (X, Y) are reconstructed from the S2 light pattern detected by the top PMT array. XENONnT employs three machine learning algorithms for this reconstruction: a multilayer perceptron (MLP), a convolutional neural network (CNN), and a graph constrained network (GCN). These models are trained on full-chain simulation events incorporating realistic detector conditions, including PMT gain variations and non-operational channels. The position resolution primarily depends on $S2_{top}$, the total photoelectrons detected by the top PMT array. Figure 3.4 shows the simulation-based estimate of the radial resolution (σ_R) as a function of $S2_{top}$ for these algorithms. For low-energy events (~ 100 PE), the resolution reaches 1.5 cm for both GCN and MLP, while CNN achieves 1.9 cm. The resolution improves significantly for larger signals ($\geq 10^4$ PE), with all algorithms achieving sub-0.25 cm precision. Near the edge of the active volume ($R \sim 65$ cm), the position resolution degrades to approximately 1.5 \times worse for small S2 signals and 2 \times worse for large signals compared to central events. This degradation is caused by S2 light reflections on the PTFE walls, which alter the hit pattern distribution and thus the reconstruction accuracy.

Due to field non-uniformities from the relatively low drift field, field shaping ring discreteness, and potential charge accumulation on PTFE walls, the electric field can become distorted, leading to biased position reconstruction. To correct for these effects, a three-dimensional field distortion correction (FDC) map is applied:

$$R_{corr} = R_{obs} + c_{FDC}(R_{obs}, t_{drift}, \phi_{obs}), \quad (3.1)$$

where R_{obs} is the observed radius, t_{drift} is the drift time, and ϕ_{obs} is the azimuthal angle. Unless otherwise specified, all coordinate references (e.g., R) in the following denote the corrected values.

^{83m}Kr is periodically introduced into the TPC via the gas xenon recirculation system, typically every two weeks, to provide accurate calibration. This isotope decays through a two-step cascade, emitting conversion electrons of 32.1 keV and 9.4 keV, thereby producing electronic recoils. The initial decay has a half-life of 1.83 hours, which is long enough to ensure the atoms distribute uniformly throughout the volume before

decaying. The second decay follows rapidly, with a half-life of just 157 ns [78]. As a result, the S2 signals are almost always merged, having a typical time separation of $\mathcal{O}(100)$ ns and widths of $\mathcal{O}(1)$ μ s, and the corresponding S1 signals are also typically observed as a single combined pulse. Unless otherwise specified, in this work the ^{83m}Kr calibration denotes the 41.5 keV events in which both S1 and S2 signals are merged into a single event and are not resolved by the splitting algorithm, as this is the most abundant and robust population [79]. These well-defined, monoenergetic events are essential for monitoring detector stability and characterizing the detector response across space and time. Figure 3.5 illustrates the spatial distributions of ^{83m}Kr events before and after applying this correction, demonstrating improved uniformity across the detector volume. After position reconstruction, signal corrections are applied to account for position-dependent detector effects.

3.4 Signal corrections

The S1 and S2 signals require correction due to spatial and temporal variations introduced by detector effects such as electric field non-uniformity, light absorption within the liquid xenon (LXe), and changes in xenon purity over time. Without proper correction, an interaction with a fixed energy could produce S1 and S2 signals of different sizes depending on its location and the time of occurrence within the detector. This would degrade the energy resolution, ultimately reducing the ability to distinguish rare signals from background.

S1 corrections

The photon count measured at a specific position in the detector, described by cylindrical coordinates (R, φ, Z) , is influenced by multiple factors:

- **Photon yield (PY):** The number of photons generated per unit of deposited energy E_0 under a given electric drift field $E_{\text{drift}}(R, \varphi, Z)$.
- **Light collection efficiency (ϵ_L):** The fraction of emitted photons that reach a PMT photocathode, with values dependent on event location.
- **Quantum efficiency ($\epsilon_{QE} \approx 0.34$):** The probability that an incident photon produces a PE at the photocathode.
- **PMT collection efficiency (ϵ_{CE}):** The fraction of PEs that are successfully collected and recorded by the PMT and the DAQ system.

Combined, these efficiencies determine the **light yield (LY)**, defined as the mean number of detected S1 PE per unit of deposited energy E_0 in the medium.

$$\text{LY}(R, \varphi, Z, E_0) = \frac{\langle \text{S1} \rangle(R, \varphi, Z, E_0)}{E_0} \quad (3.2)$$

$$= \text{PY}(E_0, E_{\text{drift}}(R, \varphi, Z)) \cdot \epsilon_L(R, \varphi, Z) \cdot \epsilon_{QE} \cdot \epsilon_{CE}, \quad (3.3)$$

where $\langle \text{S1} \rangle$ denotes the mean S1 signal amplitude in PEs at energy E_0 , corrected for PMT gain and averaged over all selected events at position (R, φ, Z) . To account for the spatial variations, a three-dimensional correction map c_1 is constructed using the 41.5 keV events from ^{83m}Kr calibration data [80]. This correction map incorporates both light collection and drift field dependencies:

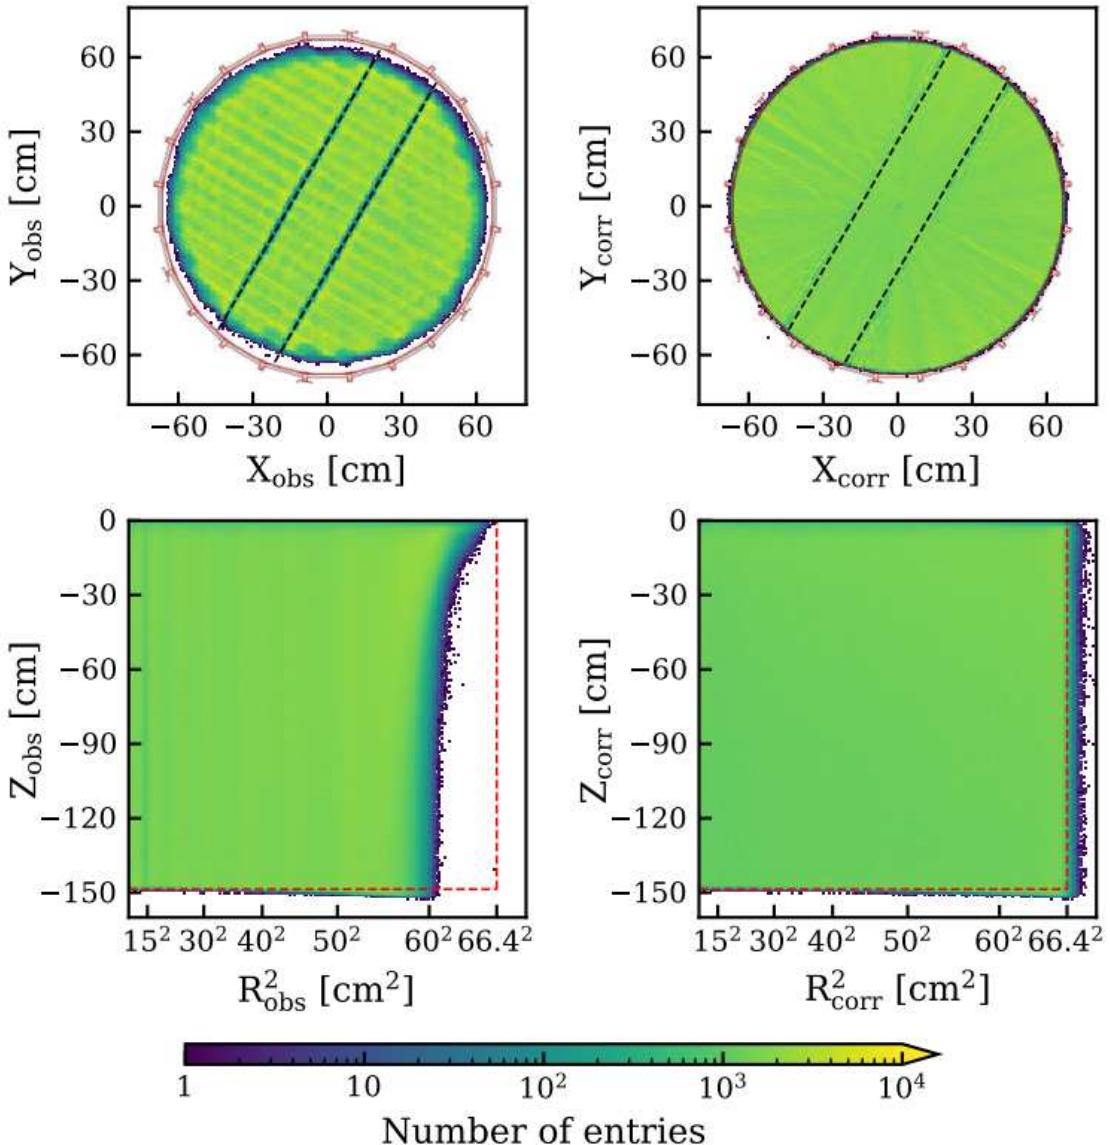


Figure 3.5: Comparison of pre- and post-correction (left/right columns) spatial distributions of ^{83m}Kr calibration events reconstructed using the multilayer perceptron (MLP) algorithm in the first science run: (top) (X, Y) view showing PTFE structures (red) and gate wires (black dashed), and (bottom) (R^2, Z) view with active volume boundaries (cathode and PTFE walls in red dashed). Figure published in [34].

$$c_1(R, \varphi, Z) = \frac{\langle \text{LY}(E_0) \rangle}{\text{LY}(R, \varphi, Z, E_0)}, \quad (3.4)$$

where the angle brackets denote spatial averaging. The corrected S1 signal (cS1) is then calculated as:

$$\text{cS1} = \text{S1}(R, \varphi, Z) \cdot c_1(R, \varphi, Z). \quad (3.5)$$

Thanks to the presented correction, the cS1 spatial non-uniformities are significantly reduced, achieving a maximum deviation within $\pm 1.5\%$ across the fiducial volume, as demonstrated in the voxel analysis with ^{37}Ar K-shell calibration at 2.82 keV (see section 4.2).

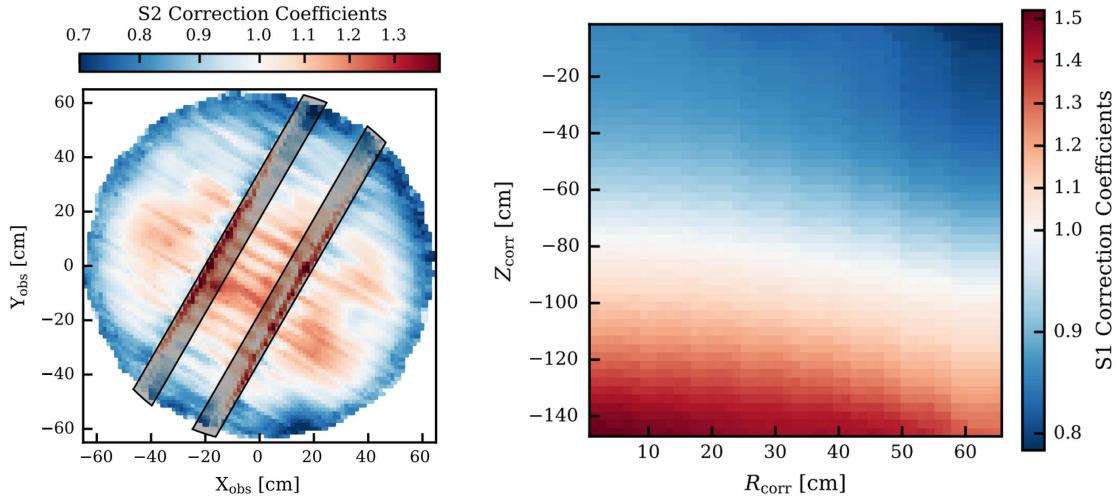


Figure 3.6: (Left) S2 correction map from $^{83\text{m}}\text{Kr}$ data, applied to signals detected by the top photomultiplier tube array in XENONnT. The white region marks the normalization point (correction factor = 1), with a linear color scale. The near-wire regions are marked, with the perpendicular gate and anode grid wires (installed to mitigate sagging) showing distinct features. Along these wires and in the center, enhanced single-electron gain and extraction efficiency yield correction factors ≥ 1 . A similar map is used for the bottom array. (Right) S1 light yield map from the same calibration, used for the first science run corrections. The color scale and normalization follow the same convention.

S2 corrections

To account for systematic variations in the S2 detector response, three primary corrections must be applied to the S2 signals:

- **Electron attenuation during drift:** Under the external electric field, drifting electrons in the xenon volume may be captured by electronegative impurities. This loss follows an exponential dependence on drift time t_{drift} , characterized by the so-called electron lifetime τ_{EL} .
- **Position-dependent effects:** Non-uniformities in the detector, including PMT response variations, electric field distortions, and geometric asymmetries introduce spatial biases in S2 amplitudes. These are mitigated using a calibration map $c_2(X, Y)$, derived from mono-energetic $^{83\text{m}}\text{Kr}$ decays, which normalizes the observed S2 signals across the active region.
- **Temporal variations in gain and extraction:** Due to electrode voltage fluctuations caused by mechanical sagging, the detector experiences time-dependent variations in the single-electron gain (SEG) and the electron extraction efficiency (EE) at the liquid-gas interface. These dynamic changes are corrected using empirically derived time-scaling factors, $\text{RSEG}(t)$ and $\text{REE}(t)$, which represent the relative deviation of SEG and EE from their nominal values at time t .

The fully corrected S2 (cS2) is then computed as:

$$\text{cS2}(X, Y, t_{\text{drift}}, t) = \frac{\text{S2}(X, Y, t_{\text{drift}}) \cdot \exp(t_{\text{drift}}/\tau_{\text{EL}})}{c_2(X, Y) \cdot \text{RSEG}(t) \cdot \text{REE}(t)}, \quad (3.6)$$

where the spatial calibration c_2 is applied in the observed (X, Y) coordinate system, as S2 non-uniformities arise primarily from processes in the gas phase. Given the negligible

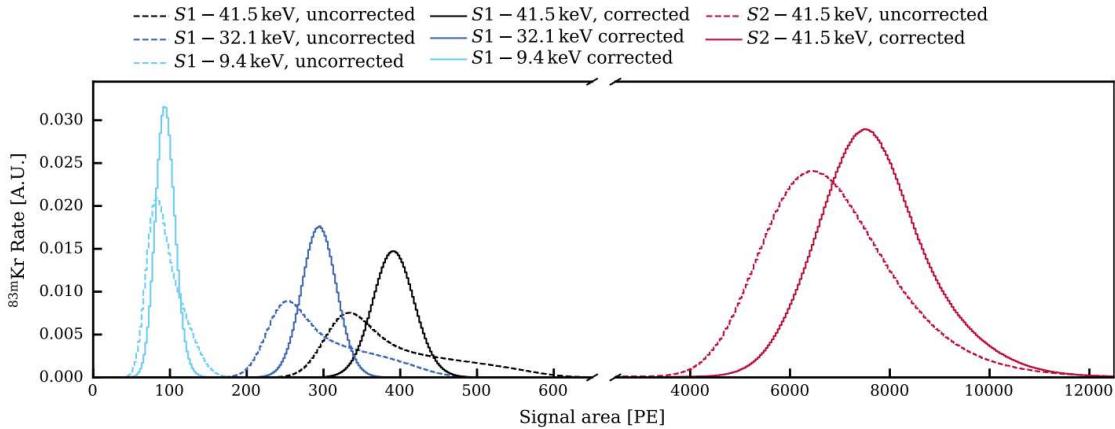


Figure 3.7: Comparison of the ^{83m}Kr S1 and S2 signals before and after correction, showing that Gaussianity is restored in the latter case. Figure published in [34].

influence of spatial dependencies in SEG and EE, they are omitted for simplicity (see [34] for details). Also in the case of cS2, the spatial uniformity achieved after applying the correction is evaluated using voxel analysis with ^{37}Ar K-shell calibration at 2.82 keV (see section 4.2), showing deviations within -2% to $+3\%$ across the fiducial volume. The largest deviations are observed in a fraction of voxels near the wire regions, where the modeling is more complex and challenging; nevertheless, this level of uniformity robustly fulfills the required accuracy for the analyses.

The effectiveness of the corrections in improving signal quality and achieving Gaussian-like distributions is illustrated in figure 3.7 for both S1 and S2 signals from ^{83m}Kr calibrations.

3.5 Selection criteria overview

In XENONnT, event selection criteria, commonly also referred to as cuts in particle physics, are analysis-specific rather than universal. Each physics search applies a distinct set of cuts, optimized for its target signal and specific background conditions. Nonetheless, most low-energy analyses, particularly the WIMP search via nuclear recoils and the electronic recoil analysis, share a common set of selection criteria, which are detailed below. These criteria are designed to isolate single-scatter events within the region of interest and to reject unphysical or misreconstructed events. They can be grouped into four main categories:

- **Time-based cuts:** These cuts reject events occurring near muon or neutron veto triggers, as well as those partially acquired due to DAQ busy states. They also exclude data taken at the edges of run boundaries or improperly triggered peaks. These cuts reduce the effective live time.
- **S2 quality:** To select single-scatter events, an S2 single-scatter cut is applied. This cut rejects events where a significant alternative S2 signal is present relative to the main S2, indicating a likely secondary interaction within the detector. Furthermore, information derived from the spatial distribution of light detected by the top PMT array for the main S2, known as the S2 pattern, is utilized. This cut employs a chi-squared goodness-of-fit test to compare the observed light distribution with expected patterns derived from calibration data and simulations. This helps in rejecting multi-scatter events where the S2 signals might merge or be distorted.

Finally, events exhibiting anomalous S2 AFT are rejected as they are likely caused by interactions occurring in the gas phase above the anode electrode rather than in the LXe target.

- **S1 quality:** For an S1 signal to be accepted, at least three distinct PMTs must detect photons within a short coincidence window (± 50 ns), effectively discriminating against backgrounds from random PMT noise. Events are rejected if the primary S2 signal can be plausibly matched to any S1 signal preceding it within the same event, other than the main S1. An alternative S1-S2 pairing is considered plausible if it satisfies certain event quality criteria, such as a consistent S2 width based on the drift time, an S1 signal not dominated by a single PMT, and an S1 AFT compatible with a physical interaction. The S1 hit pattern, reflecting the distribution of light among PMTs, is compared to the expected pattern for a single scatter event derived from optical simulations. A likelihood ratio is used for this comparison to identify and reject anomalous cases. To suppress background from PMT afterpulses, a selection criterion is applied based on the temporal shape of the S1 signal. This criterion identifies deviations in the pulse profile, including total width and peak quantiles, that are inconsistent with physical scintillation interactions in LXe. Additional discrimination against PMT-related artifacts is achieved by analyzing the distribution of light among the PMTs contributing to the S1 signal. Events in which a single PMT contributes an anomalously depth-dependent large fraction of the total S1 signal are rejected.
- **S1-S2 matching and event quality:** Events exhibiting an S2 signal width that deviates significantly from the expected width based on the electron diffusion law during the drift time are removed. This selection helps to filter out accidental pairings of S1 and S2 signals that are not physically related and can also identify potential multi-scatter events where the electron clouds might not have diffused as expected for a single interaction. A selection is applied based on the AFT of the S1 signal. Events with unusually high or low AFT values are rejected. This helps in discriminating against poorly formed S1 signals or potential background events. The overall quality of an event is evaluated by analyzing its ambient environment within the defined event time window. Events occurring in periods or regions with a high rate of small isolated S1 and S2 peaks, also called lone hits, can lead to distorted reconstruction results and are therefore rejected. To mitigate backgrounds from accidental coincidences and delayed emission processes, event selection criteria are applied based on the temporal and spatial proximity of the candidate event to prior high-energy depositions (of at least 10^3 PE). This selection includes considering the time elapsed since a prior large signal and the spatial distance from previous interaction sites. Events occurring too soon after or too close to these previous signals are rejected to suppress backgrounds related to delayed electrons or photo-ionization. Further details on these selection criteria are provided in section 5.3.

Chapter 4

Search for BSM physics in keV-scale electronic recoils

4.1 First XENONnT science run

Following an initial commissioning phase, the first XENONnT science run, named SR0, occurred from July 6 to November 10, 2021. In this period, a raw livetime of 97.1 days of science data was accumulated. The operational timeline, including calibration periods, is illustrated in figure 4.1.

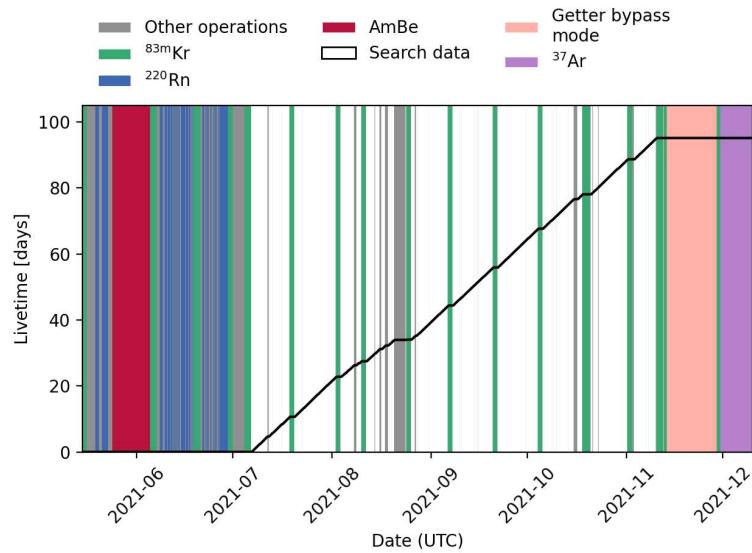


Figure 4.1: Deadtime-corrected livetime accumulation (black line) for the beyond Standard Model search in XENONnT’s first science run (SR0). Colored segments indicate calibration periods and operational downtimes.

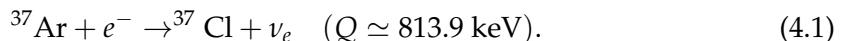
Through the meticulous radioassay campaign and the innovative cryogenic radon distillation column, significant radon mitigation was achieved, reducing the critical low-energy background from ^{214}Pb . The ^{222}Rn concentration was lowered from $11.1 \mu\text{Bq}/\text{kg}$ in the previous low-energy electronic recoils (low-ER) analysis in XENON1T [81] to $1.7 \mu\text{Bq}/\text{kg}$ in XENONnT SR0. A short circuit between the cathode and bottom screen, occurring in November 2020, necessitated operating SR0 with non-optimal electric fields: the average drift field was about 23 V/cm (compared to a design value of 200 V/cm) and the electron extraction field near the liquid-gas interface ranged from around 2.9 kV/cm to 3.7 kV/cm , depending on the radial position due to electrode sagging. However, continuous xenon purification achieved an excellent electron lifetime exceeding 10 ms ,

enabling effective charge drift even at this lower field (refer to section 3.4). Crucially, key performance metrics for the low-ER analysis, such as position and energy resolution, as well as the energy threshold, remained comparable to XENON1T, allowing XENONnT to acquire science data and exploit its unprecedentedly low background level. The reduced fields pose greater challenges for the nuclear recoil analysis, as they impact the discrimination power and significantly increase the rate of accidental coincidences (AC) [12, 49]. Regular detector calibration was performed via biweekly injections of ^{83m}Kr (refer to chapter 3), with calibration periods indicated by colored segments in figure 4.1. For ER calibration, events from ^{212}Pb (a progeny of ^{220}Rn), as well as ^{83m}Kr and ^{37}Ar , are utilized. Nuclear recoil calibrations and neutron-veto tagging efficiency are determined using the $^{241}\text{AmBe}$ source.

During SR0, the light yield (LY) and charge yield (CY) demonstrated remarkable stability, with fluctuations from the mean remaining below 1% and 1.9%, respectively. Further details, all reported numerical values, and in-depth explanations can be found in [34, 64] as well as in the references cited therein.

4.2 ^{37}Ar 2.82 keV calibration for energy reconstruction

Precise calibration of the detector's response to low-ERs across the entire Time Projection Chamber (TPC) volume is essential for obtaining accurate results, particularly in the few-keV energy region. This range is crucial because it is where XENON1T observed an excess [81] (see figure 4.7), and it's also where XENONnT has world-leading sensitivity to various beyond the Standard Model (BSM) signals. Due to the self-shielding of liquid xenon (LXe), external radioactive sources cannot effectively calibrate the inner detector regions. As a result, it is necessary to use internal sources, evenly distributed within the LXe, to achieve high-statistics calibration. XENONnT utilizes ^{83m}Kr and ^{220}Rn as its primary internal ER calibration sources. However, each has limitations for precisely mapping the detector response at the lowest energies. The ^{83m}Kr source provides clean monoenergetic lines (9.4 keV and 32.1 keV), but the short half-life of its intermediate state (154.4 ns) leads to signal pile-up. As a result, the combined 41.5 keV peak, where both transitions are typically merged, is the one used for calibration (see section 3.3), but it lies well above the few-keV region of most interest. The ^{220}Rn calibration provides a continuous β -spectrum from ^{212}Pb , but the lack of distinct low-energy peaks and lower statistics make it difficult to achieve high-precision, spatially resolved calibration in the few-keV range. To address this gap, the radioactive isotope ^{37}Ar is employed. First employed as an internal calibration in XENON1T [82], ^{37}Ar is produced by thermal neutron irradiation of ^{36}Ar -enriched gas. Through the purification system, a suitable dose is injected into the LXe TPC. ^{37}Ar decays via electron capture with a half-life of 35.0 days [83]:



This decay involves electron capture from atomic shells, leading to electron rearrangement and subsequent emission of Auger-Meitner electrons and X-rays. The energy deposits correspond to the shell binding energies, with characteristic energies and branching ratios listed in table 4.1.

The ^{37}Ar calibration is carried out at the end of a science run. This timing is set by the need for about 60 days of cryogenic distillation to remove the ^{37}Ar and restore background conditions suitable for science data taking. An order of magnitude reduction in ^{37}Ar rate occurs approximately every 8 days during this process.

Shell	Energy [keV]	Probability
K-shell	2.822	90.21%
L-shell	0.270	8.72%
M-shell	0.018	1.06%

Table 4.1: ^{37}Ar decay characteristics: shell-specific energies and branching ratios. Data from [84].

Energy scale determination

The observable energy E_0 for ER events is reconstructed from the anticorrelated scintillation (cS1) and ionization (cS2) signals. For ER interactions, the Lindhard factor L is taken to be 1.0 [21]. This assumption does not imply the absence of heat loss, but rather a near-constant energy loss as a function of the observable energy E_0 . This energy loss is eventually absorbed into the work function $W = (13.7 \pm 0.2)$ eV/quantum [21], which is defined as the average energy required to produce a photon or electron in LXe. Therefore, the energy reconstruction for ER interactions follows from equation 1.6:

$$E_0 = (N_\gamma + N_e) \cdot W = \left(\frac{\text{cS1}}{g_1} + \frac{\text{cS2}}{g_2} \right) \cdot W, \quad (4.2)$$

where cS1 and cS2 are the signals observed for a monoenergetic interaction with observable energy E_0 , and g_1 and g_2 are the detector-specific photon and electron gains (in PE/quantum). It is noteworthy that other experiments have reported different values for the work function W . Notably, the EXO-200 experiment measured $W = (11.5 \pm 0.5)$ eV/quantum [85], and the Xurich II dual-phase TPC reported $W = 11.5^{+0.2}_{-0.3}$ eV/quantum [86]. While adopting a different W slightly alters the inferred light and charge yields, the energy reconstruction remains unchanged. Assuming W is constant with energy, any variation in W is absorbed into the fitted detector gains g_1 and g_2 , leaving the reconstructed energy unaffected.

Equation 4.2 shows that determining the energy requires knowledge of the gains g_1 and g_2 . These are obtained by calibrating the detector response using monoenergetic processes of known energy. Due to light-charge anticorrelation, monoenergetic events form ellipses in the cS1-cS2 space. Their mean signals (μ_{S1}, μ_{S2}) are extracted by modeling the ellipses as rotated 2D Gaussians:

$$A(cS1, cS2) = A_0 \cdot \exp \left(-\frac{f_1(cS1, cS2)^2}{2a^2} - \frac{f_2(cS1, cS2)^2}{2b^2} \right), \quad (4.3)$$

where A_0 is the peak amplitude, a and b are related to the standard deviations along the rotated ellipse axes, and the rotation is defined by:

$$f_1(cS1, cS2) = ((cS1 - \mu_{S1}) \cdot \cos \theta - (cS2 - \mu_{S2}) \cdot \sin \theta), \quad (4.4)$$

$$f_2(cS1, cS2) = ((cS1 - \mu_{S1}) \cdot \sin \theta + (cS2 - \mu_{S2}) \cdot \cos \theta), \quad (4.5)$$

where θ is the rotation angle of the ellipse. Additional background terms may be included if required by the specific dataset. The few-keV range is the primary focus of low-ER analyses due to the XENON1T excess [81] and its highest sensitivity to BSM physics. Thus, accurate energy calibration in this region using ^{37}Ar K-shell calibration data is essential. Being close to the energy threshold, determined by the requirement that at least three distinct PMTs detect photons within a ± 50 ns coincidence window (see section 3.2), it is necessary to account for the S1 detection efficiency in the fit. The fit is performed in S1 and cS2 space, as the S1 detection efficiency is most accurately

accounted for by considering the raw S1 signal rather than the corrected cS1 signal; this is particularly relevant for fitting high-statistics datasets such as the ^{37}Ar K-shell calibration. Multiple experiments have shown that the S1 signal in the few-keV range is better described by a skew-Gaussian distribution [87–89]. Based on the aspects described, a dedicated analysis is performed for the ^{37}Ar K-shell calibration data, dividing the detector volume into bins defined in cylindrical coordinates. Specifically, the detector is segmented into 12 bins along Z and 6 radial bins, with the number of azimuthal bins increasing with radius to ensure uniform XY cross-sectional areas. This results in a total of 1008 voxels, chosen to optimize spatial resolution and maintain adequate event statistics for robust fitting in each voxel. The correction $c_1(R, \varphi, Z)$, as defined in equation 3.4, was manually applied to the S1 mean obtained from the fit at the barycenter of each voxel. A binned Poisson log-likelihood fit is performed in each voxel, factorizing the 2D distribution of events in S1-cS2 space as:

$$f_{2\text{D}}(S1, cS2) = P_1(S1) \cdot P_2(cS2|S1) = \epsilon(S1) \cdot f_{\text{skew}}(S1; a, w, e) \cdot f_{\text{gauss}}(cS2|S1; \mu, \sigma, \beta), \quad (4.6)$$

where the S1 probability density $P_1(S1)$ combines two components: a skewed Gaussian distribution $f_{\text{skew}}(S1; a, w, e)$ with skewness a , width w , and location e ; and the data-driven estimation of the S1 detection efficiency $\epsilon(S1)$. To account for the light-charge anti-correlation, the conditional cS2 probability density $P_2(cS2|S1)$ is modeled as a Gaussian distribution whose parameters depend on S1. Specifically, the mean $\mu_{S2|S1}$ and standard deviation $\sigma_{S2|S1}$ vary linearly with the deviation of S1 from the skew-Gaussian mean μ_{S1}^{skew} .

$$\mu_{S2|S1} = \mu \cdot (1 - \beta \cdot (S1 - \mu_{S1}^{\text{skew}})), \quad (4.7)$$

$$\sigma_{S2|S1} = \sigma \cdot |1 - \beta \cdot (S1 - \mu_{S1}^{\text{skew}})|, \quad (4.8)$$

where μ and σ represent the baseline cS2 parameters, while β controls the strength of the anti-correlation. The expected number of events, M_{ij} , in each bin (i, j) is given by:

$$M_{ij}(\vec{\theta}) = N \cdot f_{2\text{D}}(S1_i, cS2_j; a, w, e, \mu, \sigma, \beta) \cdot \Delta S1 \cdot \Delta cS2, \quad (4.9)$$

where $f_{2\text{D}}$ is the fit function from equation 4.6 evaluated at the bin center $(S1_i, cS2_j)$, $\Delta S1$ and $\Delta cS2$ are the bin widths, and N is the overall normalization parameter. Using this model, the optimal parameter set for $\vec{\theta} = \{a, w, e, \mu, \sigma, \beta, N\}$ is determined by minimizing the negative Poisson log-likelihood:

$$-\ln \mathcal{L}(\vec{\theta}) = \sum_{i,j} \left[M_{ij}(\vec{\theta}) - N_{ij} \ln(M_{ij}(\vec{\theta})) \right] + \text{constant}, \quad (4.10)$$

where N_{ij} represents the observed counts in each bin (i, j) .

The left panel of figure 4.2 shows the two-dimensional histogram of S1 and cS2 for the ^{37}Ar K-shell calibration data in a single voxel with the best fit. Despite localized variations observed in the voxel analysis, where the maximum observed deviation of the mean cS1 across all voxels was $\pm 1.5\%$, and for the mean cS2 was between -2% and +3%, these fluctuations effectively average out over the detector's larger scales, resulting in a robust overall energy scale with deviations contained within $\pm 1\%$ across the fiducial volume (FV) at the 2.82 keV of ^{37}Ar K-shell. Averaging over the results of all voxels

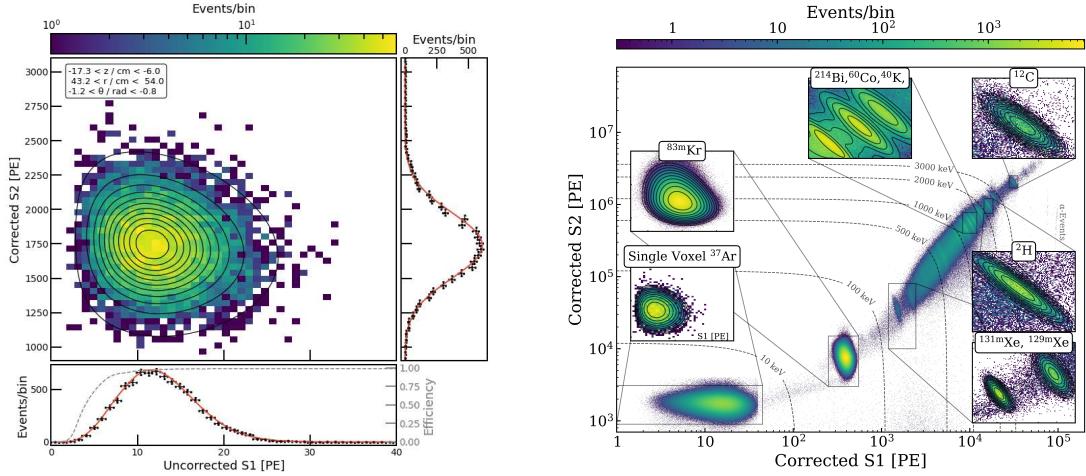


Figure 4.2: Left: Two-dimensional histogram of uncorrected S1 and corrected S2 for the ^{37}Ar K-shell (2.8 keV) calibration data in a single voxel with the best fit. Projection onto the corrected S2 (right) axis shows the best-fit model (red line), while the uncorrected S1 projection (bottom) displays both its best-fit model (red line) and the S1 detection efficiency (dashed line). **Right:** Compilation of various XENONnT SR0 calibration datasets alongside background data in the cS1-cS2 space. The insets emphasize the monoenergetic peaks and their corresponding fits used for energy calibration.

provides the following values for the light and charge yields:

$$\text{LY}_{^{37}\text{Ar}} = (5.3325 \pm 0.0014) \text{ PE/keV},$$

$$\text{CY}_{^{37}\text{Ar}} = (620.38 \pm 0.28) \text{ PE/keV}.$$

The other monoenergetic lines used in the g_1 - g_2 calibration are ^{83m}Kr (41.5 keV), ^{131m}Xe (163.9 keV) and ^{129m}Xe (236.2 keV). The ^{83m}Kr calibration, introduced in section 3.3, exhibits a tail toward larger cS2 values, most likely due to imperfect correction for field inhomogeneities [34]. Consequently, its cS2 distribution is modeled with an empirical skew-Gaussian, while a standard Gaussian is used for cS1. This artifact is also observed in other monoenergetic lines but has a negligible impact on their LY/CY measurements. ^{131m}Xe and ^{129m}Xe are isomeric states produced by neutron activation of ^{129}Xe and ^{131}Xe in nuclear recoil calibration with $^{241}\text{AmBe}$ neutron source, having half-lives of 11.86 and 8.88 days, respectively. They are uniformly distributed in the TPC. This dataset is modeled as a sum of two independent 2D elliptical Gaussian peaks for the xenon gamma lines plus an additive empirical 2D Gaussian band for the background from Compton scatter and β -decay events.

Several high-energy gamma lines (>1000 keV) were also analyzed for completeness, despite falling outside the energy region of interest (ROI). These gamma sources are predominantly found at the edges of the TPC volume, areas typically mostly excluded from the FV due to increased surface and material backgrounds, charge loss effects, and field inhomogeneities. The analyzed lines include ^{60}Co (1173.2 keV and 1332.5 keV), ^{40}K (1460.8 keV), ^2H (2224.5 keV), and ^{12}C (4438.9 keV). The gamma sources derive from either radiogenic decays in detector materials (^{60}Co , ^{40}K) or neutron interactions during calibration runs (^2H , ^{12}C), providing valuable supplementary validation of the energy scale. The right panel of figure 4.2 shows the two-dimensional histogram of cS1 and cS2 for the previously described monoenergetic lines combining all the datasets together with the best fits.

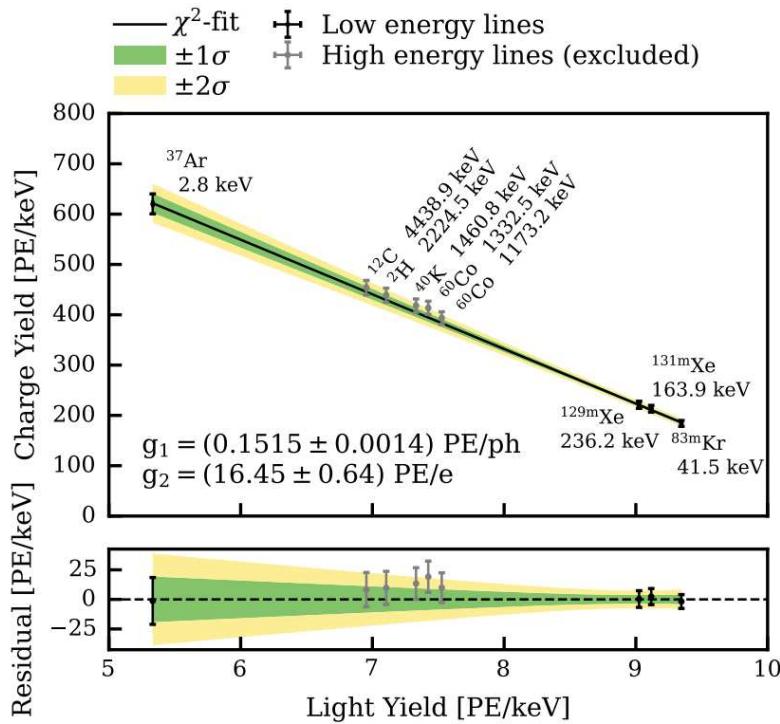


Figure 4.3: Linear fit of selected calibration data in the space of anti-correlated light and charge yield for the determination of g_1 and g_2 , commonly referred to as the Doke plot. The gray high-energy lines are omitted to avoid an energy-dependent bias and are only shown for reference. Although excluded from the fit, these high-energy data points are still compatible with the fitted model within their respective uncertainties. Figure published in [34].

Before determining the gain parameters g_1 and g_2 , the measured LY and CY are adjusted to account for the energy-dependent peak reconstruction bias described in section 3.2. This correction is applied only to the LY/CY measurements, not to the cS1/cS2 values, leading to a residual bias in the energy scale. To mitigate this effect at low energies, the reconstruction bias correction is normalized such that the ^{37}Ar line exhibits zero bias. The remaining energy bias is quantified and later included in the analysis.

A systematic uncertainty of 3.2% is assigned to the CY, based on the observed discrepancies between the low-energy and high-energy calibration lines. The resulting gain parameters are determined as:

$$\begin{aligned} g_1 &= (0.151 \pm 0.001) \text{ PE/ph}, \\ g_2 &= (16.5 \pm 0.6) \text{ PE/e}^-. \end{aligned} \quad (4.11)$$

4.3 Energy resolution

Following the energy reconstruction, the detector's energy resolution, or energy smearing, must be characterized. Energy smearing describes how the detector response transforms monoenergetic energy depositions into finite-width peaks and modifies the shape of continuous spectra. Accurate characterization is paramount as it is required to build realistic signal and background models for the analysis from the literature-based deposited energy spectra.

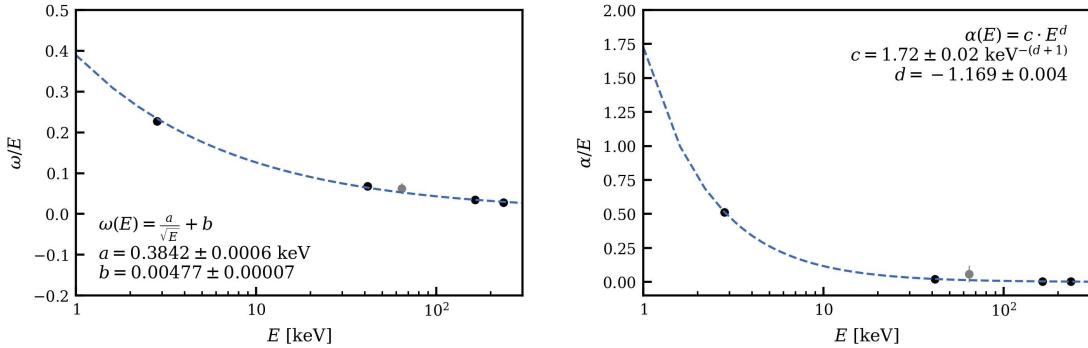


Figure 4.4: Left: Width parameter $w(E)$ of the skewed Gaussian as a function of observable energy E in XENONnT low-energy calibration data. **Right:** Skewness parameter $\alpha(E)$ versus observable energy E for the same dataset. The gray point (excluded from fits) corresponds to the KK-capture peak from ^{124}Xe two-neutrino double electron capture decay, shown as a cross-check with science data.

In the simplest model, the distribution of reconstructed energies (E_R) for a monoenergetic source is approximated by a Gaussian function, reflecting the statistical fluctuations inherent in signal generation and the detector response, which is the result of many independent processes, thereby justifying the application of the central limit theorem [90]. The energy resolution is empirically determined by calibrating the detector with monoenergetic sources, using the same datasets as for the g_1 - g_2 calibration (section 4.2). Applying the determined g_1 and g_2 values, E_R is calculated for each event (equation 4.2). The distribution of E_R values is then histogrammed and fitted with a probability distribution (e.g., Gaussian or skewed Gaussian), yielding a measure of the energy resolution, $\sigma(E_0)$, at the source's observable energy E_0 .

A widely used empirical model, motivated by the underlying physics of signal generation, describes the relative energy resolution as:

$$\frac{\sigma(E_0)}{E_0} = \frac{a}{\sqrt{E_0}} + b \quad \Leftrightarrow \quad \sigma(E_0) = a\sqrt{E_0} + bE_0, \quad (4.12)$$

where the $a/\sqrt{E_0}$ term primarily reflects statistical fluctuations in the number of quanta produced, $N_Q \propto E_0$, which follows Poisson statistics ($\sigma_{N_Q} \propto \sqrt{N_Q}$), while the constant b term captures energy-independent contributions, such as systematic effects from spatial non-uniformities or position reconstruction uncertainties. An additional term, c/E_0 , representing the contribution from electronic noise, is often considered in detector physics; however, for low-ER events in XENONnT analyses, including this term yielded non-physical negative values for c when unconstrained, and showed no significant fit improvement when constrained to be positive, leading to its exclusion for simplicity. While physically motivated, this is an empirical model; the parameters a and b are determined by fitting and are not fundamental predictions, meaning the model's extrapolation accuracy outside the calibration range is not guaranteed with good accuracy. This limitation was, for example, observed in XENON1T analyses comparing fits including or excluding the low-energy ^{37}Ar 2.82 keV point [91].

To obtain accurate analyses at low energies, a skew-Gaussian function better describes the energy distribution, as implemented and motivated for the ^{37}Ar analysis in section 4.2. In this more refined approach, the energy dependence of both the width

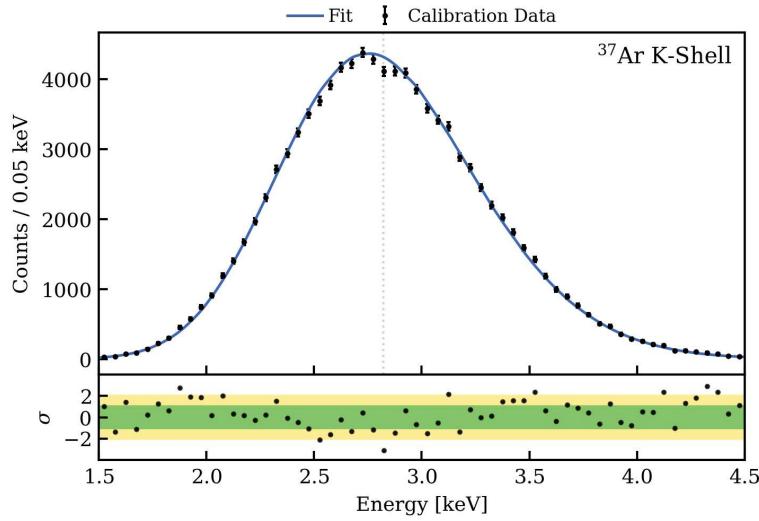


Figure 4.5: Fit with the SR0 low-energy electronic recoil model to the ^{37}Ar K-shell calibration data. The expected mean energy for the K-shell transition is indicated by the dotted vertical line at 2.82 keV [92].

parameter w (analogous to σ) and the skewness parameter α is modeled empirically by:

$$\frac{w(E)}{E} = \frac{a}{\sqrt{E}} + b, \quad \frac{\alpha(E)}{E} = c \cdot E^d. \quad (4.13)$$

Figure 4.4 shows the energy dependence of the relative width and skewness, with the parameters a, b, c, d obtained by fitting the empirical functions to the calibration data. The final result of the calibration is an energy-dependent skewed Gaussian probability density function, $f_{\text{skew}}(E_R | \xi = \xi(E_0), w = w(E_0), \alpha = \alpha(E_0))$, which describes the probability of reconstructing energy E_R given an observable energy E_0 . The location $\xi(E_0)$ is calculated such that the mean of the skewed Gaussian distribution equals E_0 . This function is then used as a smearing kernel in the analyses to convolve theoretical observable energy spectra, $R_0(E_0)$, into the reconstructed energy spectra, $R_R(E_R)$, that can be compared with the measured data:

$$R_R(E_R) = \int R_0(E_0) f_{\text{skew}}(E_R | \xi(E_0), w(E_0), \alpha(E_0)) dE_0. \quad (4.14)$$

To validate the model, an unbinned log-likelihood fit is performed on the high-statistics ^{37}Ar K-shell data, resulting in a good agreement as shown in figure 4.5. The statistical framework is identical to that of the final analysis, but the likelihood includes only the monoenergetic ^{37}Ar component. In this fit, the peak energy and resolution are fixed, and efficiency is modeled via a nuisance parameter. An equivalent fit is performed using the continuous spectrum from ^{212}Pb beta decays in ^{220}Rn calibration data, as shown in [64].

4.4 Background models

The XENONnT low-ER analysis benefits significantly from a comprehensive background model, without which sensitivity to the rare processes targeted by the physics program

could not be established. This framework provides the understanding necessary for determining discovery potential and ultimately claiming a discovery.

Covering all known and anticipated ER sources within the 1–140 keV reconstructed energy range and defined FV, the background model consists of nine components, which are detailed in the following paragraphs (more information in [64]).

^{214}Pb : The beta decay of ^{214}Pb (1.02 MeV Q-value), a progeny of ^{222}Rn continuously emanating into the xenon target from detector materials, is the dominant background source at the lowest energies. Its spectral shape is derived from theoretical calculations considering it as a forbidden transition [93]. The rate is constrained by measured alpha activities of ^{218}Po and ^{214}Po , resulting in ^{214}Pb activity bounds of $(0.777 \pm 0.006_{\text{stat}} \pm 0.032_{\text{sys}})$ to $(1.691 \pm 0.006_{\text{stat}} \pm 0.072_{\text{sys}})$ $\mu\text{Bq}/\text{kg}$. In the background model, a nuisance parameter is uniformly constrained within this range.

^{85}Kr : The beta decay of ^{85}Kr (687 keV Q-value, 10.8 years half-life), is an intrinsic background source that persists as residual trace amounts despite cryogenic distillation [40]. Quantifying this at parts-per-quadrillion (ppq) levels is challenging, a difficulty addressed by sophisticated techniques (detailed in chapter 2). The ^{85}Kr rate used in the background model is Gaussian constrained based on measured ^{85}Kr concentrations (60 ± 40) ppq from SR0 gas samples, assuming a standard European atmospheric ^{85}Kr abundance [55].

Materials: Gamma rays from radioactive decays (^{40}K , ^{60}Co , U/Th) in detector materials induce ER events via Compton scattering in the LXe. Unlike external X-rays, with $\mathcal{O}(10 \mu\text{m})$ penetration depth, these gammas reach the inner volume. Depending on the analysis needs, this background can be suppressed by reducing the FV. Geant4 simulations [94], informed by material radioassays [44], predict a flat ER background spectrum below 140 keV in the FV. The estimated rate is (2.1 ± 0.4) events $/(\text{keV} \cdot \text{tonne} \cdot \text{year})$, with uncertainty from simulation statistics and radioassay measurements. In the background model, a nuisance parameter for this rate is constrained by a Gaussian.

^{136}Xe two-neutrino double beta decay ($2\nu\beta\beta$ decay): The natural isotope ^{136}Xe undergoes $2\nu\beta\beta$ decay (2.17×10^{21} years half-life [25]), yielding a continuous ER spectrum up to the 2458 keV Q-value. The decay rate, a dominant background above ~ 30 keV, is constrained by the half-life and the in-situ measurement of the ^{136}Xe abundance $(8.97 \pm 0.16)\%$ [95]. Theoretical uncertainty on the spectral shape exists due to nuclear matrix elements, leading to differing predictions from the two models SSD [96] and HSD [97]. The analysis baseline is the average spectrum, with shape uncertainty incorporated by allowing variations based on model differences.

Solar neutrinos: Solar neutrinos originating from the Sun constitute an irreducible background in LXe-based detectors. They can undergo elastic scattering off electrons (neutrino-electron scattering), a process with a well-characterized cross-section in the Standard Model (SM) [98]. Within the ROI for the low-ER analysis, the solar neutrino flux is dominated by pp neutrinos, followed by the ^7Be component. Interactions involve both charged-current (CC) and neutral-current (NC) processes for electron neutrinos (ν_e), while only NC processes occur for muon (ν_μ) and tau (ν_τ) neutrinos. The cross-section is convolved with the solar neutrino flux, derived from standard solar model calculations incorporating the large mixing angle Mikheyev-Smirnov-Wolfenstein solution for

oscillations [99, 100]. To model the expected event spectrum correctly, it's important to include the effects of electron binding energies. The stepping approximation handles this by defining $Z_{\text{eff}}^{\text{Xe}}(E_r)$, which represents how many xenon electrons can be ionized for a given recoil energy E_r . This is calculated as:

$$Z_{\text{eff}}^{\text{Xe}}(E_r) = \sum_{i=1}^{54} \Theta(E_r - B_i), \quad (4.15)$$

where the sum extends over the 54 electrons of the xenon atom, B_i represents the binding energy of the i -th electron (taken from [101]), and Θ is the Heaviside step function. The relativistic random-phase approximation (RRPA), a more sophisticated model, is not considered here because, at the time of writing, it lacks predictions above 30 keV [102]. Based on the Borexino measurement [103], the solar neutrino rate is constrained by a Gaussian with 10% uncertainty. The $\sim 7\%$ peak-to-peak annual modulation from Earth's orbit is ignored as negligible in the context of this analysis.

^{124}Xe two-neutrino double electron capture ($2\nu\text{ECEC}$): The natural isotope ^{124}Xe contributes via $2\nu\text{ECEC}$. In this process, two orbital electrons are captured and two neutrinos emitted, releasing atomic binding energy as X-rays and Auger electrons. The theoretical spectrum [26], computed considering captures from various atomic shells, exhibits distinct peaks (KK, KL, LL), with the most prominent KK capture peak at around 64.3 keV. The rate of this background is left unconstrained in the fit, treated as a free parameter, allowing precise determination of the $2\nu\text{ECEC}$ half-life from the data.

Accidental coincidences (AC): Occasionally, an unrelated S1 signal can be erroneously paired with an unrelated S2 signal. These fake events are termed ACs. This background source is described in detail in section 5.3. The rate and spectrum of AC events are estimated using a data-driven method: isolated S1 and S2 signals recorded during normal data taking are randomly paired to create synthetic events, which are then processed through the standard analysis chain. This procedure yields a predicted AC rate of (0.61 ± 0.03) events/(tonne · year) in the ROI. This prediction is used with a Gaussian constraint in the background model.

^{133}Xe : ^{133}Xe is a radioactive isotope (5.3 days half-life) produced via neutron activation of stable ^{132}Xe , from the $^{241}\text{AmBe}$ calibration source used for nuclear recoil calibrations. ^{133}Xe decays via beta emission and a prompt 81 keV gamma ray. These rapid, localized emissions are detected as unresolved ER events with combined energy, producing a continuous spectrum extending up to about 427 keV. Due to the lack of direct measurements of the trace ^{133}Xe concentration at the beginning of SR0, its rate is left unconstrained (free parameter) in the background model fit.

^{83m}Kr : As detailed in section 3.3, ^{83m}Kr is a residual background from calibration runs that persists despite continuous purification. Its decay results in monoenergetic ER events at 41.5 keV. Lacking ancillary measurements, its rate is an unconstrained nuisance parameter in the fit, affecting sensitivity exclusively in that narrow energy region.

Tritium: The β decay of tritium (18.6 keV Q-value, 12.3 years half-life [104]) was considered a potential explanation for the XENON1T excess [81]. Consequently, XENONnT implemented stringent measures against HT/HTO contamination [11]. To verify the

impact of these procedures, a “tritium-enhanced” dataset was collected after SR0 by bypassing the gaseous xenon getter, an arrangement designed to amplify any tritium presence by a factor of 10–100 times. The analysis of this data showed no tritium-like signal, confirming contamination reduction effectiveness [64]. Thus, tritium was excluded from the background model.

In summary, the nine components constituting the background model B_0 for the XENONnT SR0 low-ER analysis are listed in table 4.2, along with their prior constraints. Their spectral shapes and rate constraints (or lack thereof) are incorporated into the global likelihood function used for statistical inference, allowing for robust estimation of signal strengths or limits while accounting for background fluctuations and uncertainties.

4.5 Signal models

A brief overview of the BSM signal models considered in the XENONnT SR0 low-ER analysis is provided in the following. More details are given for the models that were integrated into the analysis pipeline by the author. Detailed descriptions for the others, including solar axions and bosonic dark matter, along with their implementation, can be found in [64, 81] and the references therein. The following candidates are investigated in this search.

Solar axions: The strong CP problem in quantum chromodynamics motivated the introduction of axions [105]; for a detailed review of axion models, see [106]. If they exist, these hypothetical particles can be produced in the Sun through various mechanisms, including axio-recombination, Compton, and Bremsstrahlung processes (ABC). They can be detected in XENONnT via the axioelectric effect, which is analogous to the photoelectric effect but involves axions instead of photons.

Solar neutrinos with an enhanced magnetic moment: The SM assumes neutrinos are massless, meaning they should not have a magnetic dipole moment. However, observing neutrino oscillations indicates that neutrinos do have mass [5]. This finding shows that the SM must be extended. Neutrino mass implies a magnetic moment of $\mu_\nu \sim \mathcal{O}(10^{-20}) \mu_B$ [107], where μ_B is the Bohr magneton. Both theoretical studies and experimental searches have considered larger magnitudes for μ_B [108–110]. Such a property increases the probability of solar neutrinos scattering off electrons via an electromagnetic interaction, leading to a significant enhancement of the event rate, particularly at low recoil energies. The differential cross section for this process is given by:

$$\left(\frac{d\sigma_\nu}{dE_r} \right)_{\text{NMM}} = Z_{\text{eff}}^{\text{Xe}}(E_r) \frac{\pi \alpha_{\text{em}}^2 \mu_\nu^2}{m_e^2} \left(\frac{1}{E_r} - \frac{1}{E_\nu} \right), \quad (4.16)$$

where μ_ν is the neutrino magnetic moment, α_{em} is the fine-structure constant, m_e is the electron mass, E_r is the electron recoil energy, and E_ν is the incident neutrino energy. The term $Z_{\text{eff}}^{\text{Xe}}(E_r)$, which accounts for atomic binding effects, is defined in section 4.4. This cross-section’s $1/E_r$ dependence leads to a rising spectrum at low energies.

Neutrino millicharge: While within the SM, known elementary particles have charges that are integer multiples of $e/3$, there is no fundamental principle strictly enforcing charge quantization for new particles, nor one requiring massive neutrinos to be electrically neutral [111]. This theoretical flexibility permits models where neutrinos possess

small electric charges [112, 113], together with searches for millicharged particles by dedicated experiments [114, 115]. Neutrinos with even a small electric charge would electromagnetically couple to electrons, generating a recoil spectrum different from that predicted by the SM. The cross section is modeled using the equivalent photon approximation [116], which relies on the measured photoelectric cross section for xenon, $\sigma_\gamma(E_r)$. This properly accounts for electron binding effects in xenon atoms, and it only applies when the interaction is photon-mediated. The photoelectric cross section values are obtained by interpolating the experimental data points from [117], which include data below 1 keV. The differential cross section reads:

$$\left(\frac{d\sigma_\nu}{dE_r}\right)_{\text{NMC}} = \frac{2\alpha_{\text{em}}}{\pi} \frac{\sigma_\gamma(E_r)}{E_r} \log\left(\frac{E_\nu}{m_\nu}\right) q_\nu^2, \quad (4.17)$$

where m_ν is the neutrino mass, q_ν is the neutrino's electric charge. The approximate E_r^{-2} scaling behavior of the cross section, arising from both the explicit E_r^{-1} term and the energy dependence of $\sigma_\gamma(E_r)$, results in a substantial increase in the event rate at low recoil energies.

Neutrino non-standard interactions: Neutrino non-standard interactions (NSIs) between solar neutrinos and electrons, mediated by new, light (sub-MeV) bosons are considered in the following. NSIs are a well-motivated phenomenological approach for exploring new physics in the neutrino sector [13, 118, 119]. Light mediator scenarios typically result in a differential cross section with E_r^{-2} dependence, allowing new physics to be significant at low recoil energies in direct detection low-threshold (keV or less) experiments, yet remain consistent with constraints from higher energy neutrino measurements. Unlike heavy-mediator scenarios, which can be described by effective field theory, the low-mass regime requires explicit inclusion of the mediator propagator. This analysis considers vector/axial-vector, and scalar mediator cases; pseudoscalar and tensor interactions are not included for brevity, but would yield similar results and conclusions. Following standard practice for direct detection experiments, higher-order kinematic terms for the differential cross sections are ignored as their contributions are negligible [98, 120].

- **Vector and axial-vector mediators:** An NSI between neutrinos and electrons can be mediated by a new vector (V) or axial-vector (A); these cases are treated together here since the differential cross section differs only in the specific couplings and mediator masses. These are the most commonly studied types of NSIs and arise in models that extend the SM gauge group with an additional U(1) symmetry [121, 122]. Such extensions introduce a new, massive gauge boson (often denoted Z' or dark photons) that mediates the new interaction. The differential cross section includes a term purely from the new interaction and an interference term with the SM neutral current amplitude [120]:

$$\frac{d\sigma_{\nu,e}}{dE_r} \Big|_{V/A} = \frac{m_e g_{\nu V/A}^2 g_{eV/A}^2}{2\pi(2E_r m_e + m_{V/A}^2)^2} + \frac{\sqrt{2} G_F m_e g_{\nu V/A} g_{eV/A}}{\pi(2E_r m_e + m_{V/A}^2)}, \quad (4.18)$$

where G_F is the Fermi constant, $g_{\nu V/A}$ and $g_{eV/A}$ are the couplings of the new mediator to neutrinos and electrons, respectively, $m_{V/A}$ is its mass, and $g_{V/A}$ are the SM neutrino-flavour dependent electron vector and axial-vector couplings.

- **Scalar mediators:** Scalar-mediated NSI, which may emerge from expanding the scalar part of the Standard Model (like the Higgs sector), has gained growing

interest over the past few years [122, 123]. If the interaction is mediated by a scalar boson (S), it doesn't interfere with the Standard Model vector current. The corresponding differential cross section is given by [120]:

$$\frac{d\sigma_{\nu,e}}{dE_r} \Big|_S = \frac{g_{\nu S}^2 g_{eS}^2 E_r m_e^2}{4\pi E_\nu^2 (2E_r m_e + m_S^2)^2}, \quad (4.19)$$

where $g_{\nu S}$ and g_{eS} are the couplings of the new scalar mediator to neutrinos and electrons, respectively, and m_S is its mass.

Bosonic dark matter: axion-like particles and dark photons: Light bosons such as axion-like particles (ALPs) and dark photons (DPs) arise naturally in various extensions of the SM, like string theory or models with extra U(1) gauge symmetries, and are considered viable candidates for cosmological dark matter [124]. These particles could be absorbed by xenon atoms, producing a monoenergetic ER signal corresponding to their rest mass.

4.6 Event selection and cut acceptance

The event selection criteria described previously in section 3.5 form the basis for low-energy physics analyses. This section details the additional requirements and specific tuning applied for this particular search, which focuses on the reconstructed energy range from 1 to 140 keV. Specifically, the following additional criteria are applied.

- **ER band:** Events are selected within the -2σ to $+4\sigma$ range of the ER band in cS1-cS2, based on fit to ^{220}Rn and ^{37}Ar calibration [125]. This cut removes backgrounds like surface events, gas-phase events and ACs. The -2σ lower bound was chosen to maintain the blinding for the WIMP search. This selection is illustrated in figure 4.7.
- **S2 energy threshold:** A threshold of $S2 > 500$ PE is applied. This is higher than the default straxen S2 reconstruction threshold of 100 PE and serves primarily to reduce the AC background prevalent at low S2 values, without affecting the signal efficiency from ERs above 1 keV.
- **Fiducial volume (FV):** The analysis is performed within a (4.37 ± 0.14) tonne liquid xenon FV, shaped as a polygon to optimize the signal-to-background ratio and minimize instrumental effects. The FV is defined by a maximum Z of -6 cm (excluding gas interactions) and a maximum radius of 63 cm (removing surface events). Within these bounds, further regions are excluded: areas with $>10\%$ discrepancy between the electric field simulation and data, and the charge-insensitive bottom edge caused by field non-uniformities [126]. The simulation-based estimation of the charge-insensitive bottom edge exclusion size dominates the fiducial mass uncertainty.
- **Refined S2 width:** A selection based on the S2 signal width relative to expected electron diffusion is crucial for rejecting ACs, multi-scatter, and unphysical events as detailed in section 3.5. For the low-ER analysis, specific boundaries are defined based on event position relative to perpendicular wires: boundaries derived from simulations and ^{220}Rn data are used for events far from wires, while observed ^{220}Rn data provides boundaries near the wires to account for anomalous diffusion.

- **Additional criteria:** A naive Bayes classifier (NBC) is developed [127] to quantitatively measure how well the observed S1 and S2 waveforms match standard templates derived from a combination of simulation and calibration data. Events with low NBC scores for either S1 or S2, indicating deviations from expected shapes, potentially due to noise, pile-up, or non-physical origins, are rejected. Moreover, to guard against events where reconstruction might be unreliable, the positions determined by the three independent machine learning algorithms (section 3.3) are required to be consistent within a defined tolerance based on the 99% quantile of differences observed in calibration data.

Event selection efficiency

The overall efficiency of event selection, also known as cut acceptance, is the fraction of true signal events expected to pass all criteria. Determining this acceptance and its energy dependence is crucial to be able to define the expected distributions of signal and background events in the likelihood analysis. Total cut acceptance is the product of individual efficiencies, assuming independence; correlated cuts are evaluated jointly as groups. The acceptance and its energy dependence are determined using ^{220}Rn calibration data, which provides coverage across the entire ROI. High-statistics ^{37}Ar K-shell data serves as a validation point in the low-energy region. Individual cut acceptances are estimated using the N-1 method, which involves evaluating the acceptance of the Nth selection criterion after applying the other N-1 criteria. The uncertainties are determined using the Clopper-Pearson method [128].

The cumulative acceptance is approximately 80% at low energies, showing little variation. A peak increase of $\sim 2\%$ occurs at 10 keV, due to the nuclear recoil blinding cut. Above 10 keV, acceptance shows a modest gradual decline, reaching around 78% by 140 keV. Accurately determining acceptance below about 5 keV using the N-1 method on ^{220}Rn data proved challenging, particularly for the S2 width cut. This is because the N-1 method assumes a pure signal dataset, which fails at low energy where many AC events are rejected only by the S2 width cut. This leads to underestimated efficiencies and inconsistencies, as evidenced by the $\sim 7\%$ discrepancy between the ^{37}Ar and ^{220}Rn acceptance estimates at low energies. Consequently, a conservative $\pm 4\%$ uncertainty is assigned to the low-energy acceptance, which gradually narrows as energy increases up to 10 keV.

4.7 Inference and results

The statistical inference is performed using an unbinned profile likelihood method in reconstructed energy space. The likelihood function incorporates both potential signal (μ_s) and background (μ_b) components:

$$\mathcal{L}(\mu_s, \mu_b, \boldsymbol{\theta}) = \text{Poiss}(N_{\text{obs}} | \mu_{\text{tot}}) \prod_{i=1}^{N_{\text{obs}}} \left[\frac{\mu_s}{\mu_{\text{tot}}} f_s(E_i, \boldsymbol{\theta}) + \sum_{j=1}^9 \frac{\mu_{b_j}}{\mu_{\text{tot}}} f_{b_j}(E_i, \boldsymbol{\theta}) \right] \cdot \prod_{m=1}^6 C_{\mu_{b_m}}(\mu_{b_m}) \cdot \prod_{n=1}^2 C_{\theta_n}(\theta_n), \quad (4.20)$$

where $\mu_{\text{tot}} = \mu_s + \sum_j \mu_{b_j}$ is the total number of events. The index i runs over the $N_{\text{obs}} = 3658$ observed events in the ROI, E_i is the reconstructed energy of the i -th event, and j indexes the nine background components detailed in section 4.4. The

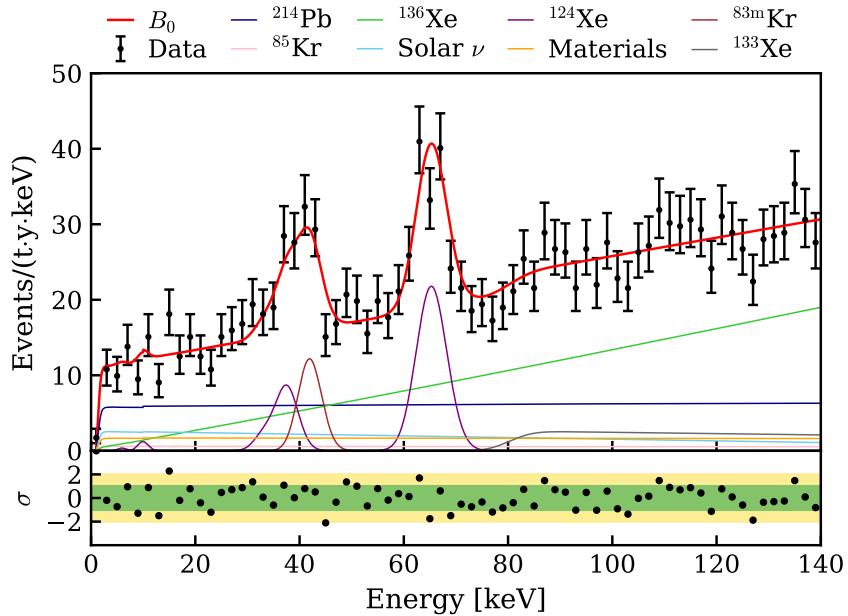


Figure 4.6: Energy spectrum of low-electronic recoil data from XENONnT SR0 after selection, compared to the best-fit background model B_0 (red curve). Individual background components are shown, except for the subdominant accidental coincidence background. The discontinuity at 10 keV arises because the WIMP search region was excluded (blinded) from this analysis. Figure published in [64].

functions f_s and f_{b_j} are the probability density functions for signal and background, respectively, which depend on energy and on nuisance parameters θ . These parameters account for shape uncertainties, including the detector efficiency at low energy and the ^{136}Xe spectrum uncertainty. The C_{μ_m} and C_{θ_n} terms incorporate constraints on background rates and shape parameters from auxiliary measurements or theoretical priors (summarized in table 4.2). Hypothesis testing and upper limit setting are based on the profile log-likelihood ratio test statistic:

$$q(\mu_s) = -2 \ln \frac{\mathcal{L}(\mu_s, \hat{\mu}_b(\mu_s), \hat{\theta}(\mu_s))}{\mathcal{L}(\hat{\mu}_s, \hat{\mu}_b, \hat{\theta})}, \quad (4.21)$$

where single-hat parameters represent the overall best-fit values, and double-hat parameters are those that maximize the likelihood for a fixed μ_s .

Following the definition of the analysis procedure, data above 20 keV were initially unblinded and showed good agreement with the background model (p -value ~ 0.2). Subsequently, the full ER dataset was unblinded, while the WIMP search region (NR band) below 10 keV was left blinded (see figure 4.7). The background-only ($\mu_s = 0$) hypothesis shows excellent agreement with observed data, indicating no statistically significant excess. The fitted spectrum is presented in figure 4.6, with best-fit results for each background component detailed in table 4.2. Given that XENONnT does not observe the excess previously seen in XENON1T [81] (see figure 4.7), the various BSM scenarios are excluded; for instance, the monoenergetic peak hypothesis at 2.3 keV is rejected at $\sim 4\sigma$ significance. This leaves trace tritium as the most plausible explanation for the XENON1T excess.

Correcting for selection efficiency, the average ER background rate in the energy range (1, 30) keV is found to be (15.8 ± 1.3) events/(tonne·year·keV). This marks a

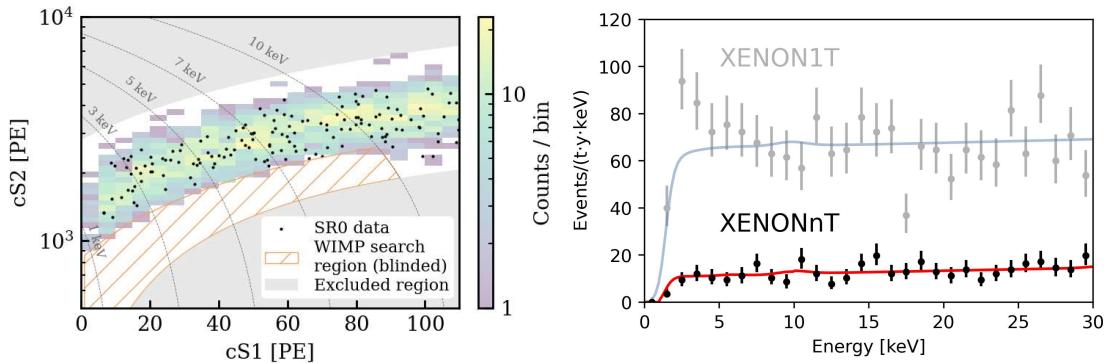


Figure 4.7: Left: Electronic recoil (ER) events (black dots) from the XENONnT first science run dataset superimposed on the 2D ^{220}Rn calibration histogram. The data are displayed in the plane of corrected scintillation (cS1) and ionization (cS2) signals. The hatched orange area, denoting the blinded weakly interacting massive particle (WIMP) search region, and the gray shaded area are both excluded from the analysis. **Right:** Comparison of the measured ER background spectra for XENONnT (this work) and XENON1T [81]. A direct comparison demonstrates the approximate five-fold background reduction relative to XENON1T, due to improved radon mitigation. The XENON1T data reveal the previously observed excess at around 2.3 keV, which is not present in the XENONnT data. Solid lines represent the best-fit background models for each experiment.

Component	Constraint [Events]	Best Fit [Events]
$^{136}\text{Xe } 2\nu\beta\beta$	1560 ± 60 (Gaussian)	1550 ± 50
^{214}Pb	(570, 1200) (Uniform)	960 ± 120
Solar neutrino (pp)	300 ± 30 (Gaussian)	300 ± 30
Materials	270 ± 50 (Gaussian)	270 ± 50
$^{124}\text{Xe } 2\nu\text{ECEC}$	Unconstrained	250 ± 30
^{133}Xe	Unconstrained	150 ± 60
^{85}Kr	90 ± 60 (Gaussian)	90 ± 60
^{83m}Kr	Unconstrained	80 ± 16
AC	0.70 ± 0.04 (Gaussian)	0.71 ± 0.03
Total observed events	3658	

Table 4.2: Summary of best-fit event counts for background components in the XENONnT SR0 low-energy analysis. Constraints on background rates are applied via the C_{μ_m} terms in equation 4.20.

factor of about 5 reduction compared to XENON1T [81] and is the lowest ER background achieved in a dark matter direct detection experiment to date. The fit provides also the valuable measurement of the $^{124}\text{Xe } 2\nu\text{ECEC}$ half-life: $T_{1/2}^{2\nu\text{ECEC}} = (1.18 \pm 0.13_{\text{stat}} \pm 0.14_{\text{sys}}) \times 10^{22}$ y, where the main sources of systematic uncertainty are the selection efficiency and the theoretical modeling of the capture fraction [26]. This finding is consistent with the previous XENON1T measurement [26]. Given the consistency between the observed data and the background-only scenario, 90% C.L. upper limits are set on the parameters of the models introduced in section 4.5. The results for solar axions, neutrinos with an enhanced magnetic moment ($\mu_\nu < 6.4 \times 10^{-12} \mu_B$), ALPs, and DPs are published in [64] and are displayed in figure 4.8. The results for neutrino millicharge ($q_\nu < 1.5 \times 10^{-13} e_0$) and neutrino NSIs with scalar and vector mediators are obtained by the author and are shown in figure 4.9. These figures demonstrate that these limits represent world-leading laboratory constraints across multiple regions of

the parameter spaces of the different signal models.

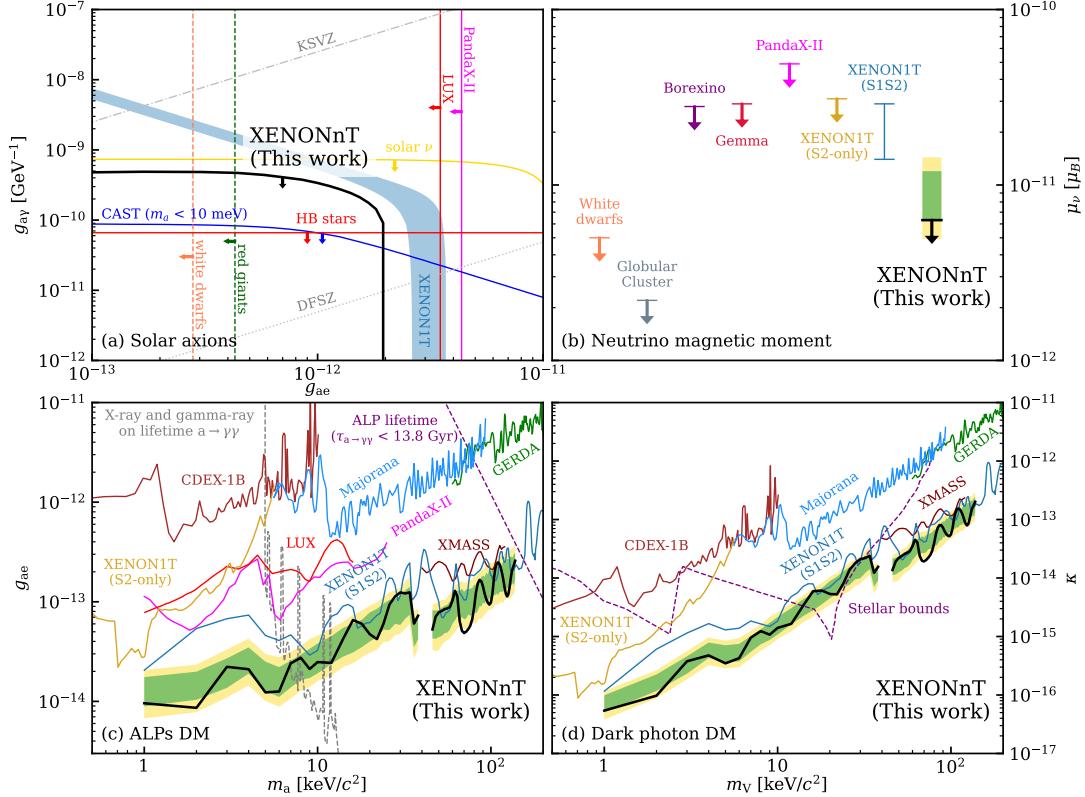


Figure 4.8: 90% C.L. upper limits derived from the XENONnT SR0 low-electronic recoil data for various beyond the Standard Model physics models. (a) Constraints on the axion-electron coupling g_{ae} (via ABC solar axions) and axion-photon coupling $g_{a\gamma}$ (via Primakoff solar axions). (b) Limit on the effective neutrino magnetic moment μ_ν . (c) Limit on the axioelectric coupling g_{ae} for axion-like particle dark matter. (d) Limit on the kinetic mixing parameter κ for dark photon dark matter. Expected sensitivity bands (1 σ green, 2 σ yellow) are shown. The region 39–44 keV/c 2 is excluded due to residual ^{83m}Kr background. For details of limits of other experiments and astrophysical bounds, see [64] the references therein, where the figure was published.

4.8 Conclusion and discussion

The analysis of the first XENONnT science run, covering a 1.16 tonne-years exposure, reveals no statistically significant excess beyond the established background model in the low-energy ER spectrum. A record-low ER background rate of (15.8 ± 1.3) events/(tonne-year·keV) in the 1–30 keV range has been achieved, representing an about 5-fold reduction compared to XENON1T. ^{37}Ar calibration with the analysis of the 2.82 keV K-shell line provided a precise and robust modelling of the low-energy response including energy scale and resolution determination. The spatial partitioning (voxelization) approach for the ^{37}Ar K-shell line analysis allowed for an accurate accounting and validation of the S1 detection efficiency. Leveraging these precise characterizations, the findings of the analysis rule out a BSM interpretation for the XENON1T low-energy excess [81] (excluding a 2.3 keV monoenergetic peak at around 4 σ), reinforcing the hypothesis of a tritium origin. The limits derived for solar axions, neutrinos with an enhanced magnetic moment, neutrino millicharge, neutrino NSIs (vector and scalar mediators), ALPs, and DPs represent the most competitive laboratory constraints

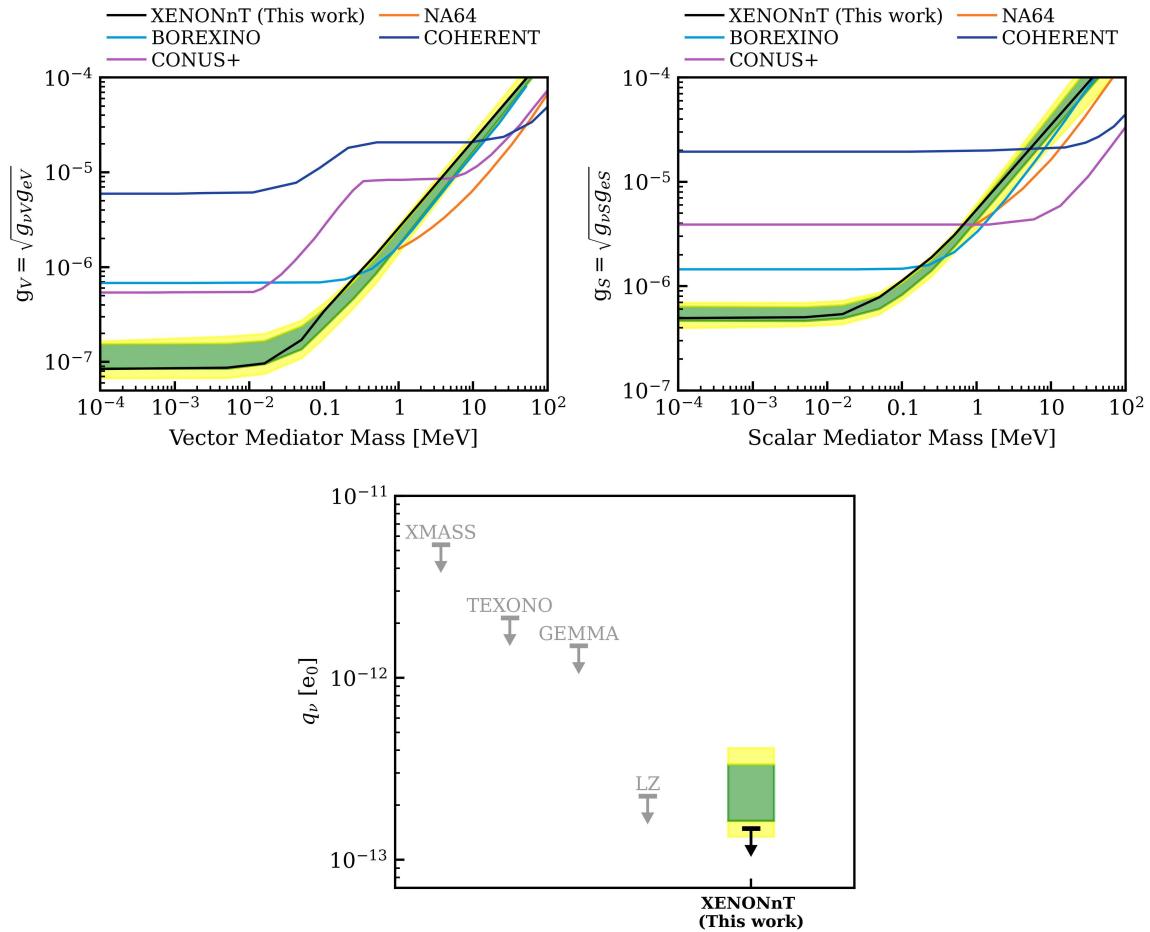


Figure 4.9: 90% C.L. upper limits from XENONnT data. For all three panels, the sensitivity band is indicated by the region containing 68% (green) and 95% (yellow) of expected upper limits under the background-only hypothesis. **Top left and right:** Limits on neutrino interactions with a light vector mediator (left) and a light scalar mediator (right). For both top panels, selected results from other experiments are also shown, including Borexino [129], COHERENT [130], CONUS+ [131], and NA64 [132]. **Bottom:** Limits on neutrino effective millicharge. Selected results from other experiments are also shown, including GEMMA [133], LZ [134], TEXONO [135], and XMASS [136].

currently available across most of the probed parameter spaces. The continued operation of XENONnT post-SR0, benefiting from further ^{222}Rn background reduction via a combined gaseous and liquid xenon flow radon distillation system [137], is expected to soon enable the detection of the solar pp neutrino component and further improve sensitivity in searches for BSM physics.

Chapter 5

First sub-keV electronic recoil calibration with ^{37}Ar L-shell

The search for beyond the Standard Model (BSM) physics in low-energy electronic recoil (ER) data, as described in chapter 4, is limited by an analysis threshold of about 1 keV, where the signal efficiency is approximately 10%. Lowering this threshold could improve sensitivity to different signals, particularly those with energy distributions that rise towards lower energies. To the best of the author’s knowledge, no sub-keV ER processes had previously been observed in liquid xenon (LXe) detectors using both ionization and scintillation signals. This chapter introduces a novel analysis framework that successfully measured the monoenergetic 0.270 keV L-shell electron capture decay using an internal ^{37}Ar calibration source. Prior to this work, the photon yield (PY) for sub-keV ERs in LXe was unknown, and it was uncertain whether any events from ^{37}Ar L-shell decays could be detected.

5.1 Modeling near the energy threshold

The characterization of sub-keV ER signals, including the 0.270 keV depositions from ^{37}Ar L-shell decays, poses significant challenges. Neither a first-principles model nor direct measurements exist for the photon emission at these energies, leaving the PY unknown and the model not well-established. This means that it is not known whether and how many photons are produced by an interaction of a given energy. Moreover, the intrinsically low photon counts expected per event significantly amplify the impact of stochastic Poissonian fluctuations. Consequently, while the Noble Element Simulation Technique (NEST) provides a state-of-the-art semi-empirical model [32, 138], its predictive power in this uncalibrated sub-keV domain is severely limited.

To systematically investigate the potential range of PYs and their impact on the detector response, a dedicated full-chain simulation strategy was employed. The XENONnT experiment’s Monte Carlo (MC) simulation framework [139] was utilized. This framework integrates outputs from generic packages like `nestpy` [33] with XENONnT-specific detector modeling, extending up to photomultiplier tube (PMT) and data acquisition (DAQ) responses (further details on the waveform simulator can be found in [34, 140]). The methodological approach involved systematically varying the quanta production within this simulation pipeline. Specifically, for 0.270 keV energy depositions, the entire distribution of NEST-generated photons was shifted by applying a constant offset to its median. This procedure preserved the inherent shape of the photon distribution while allowing for a controlled variation for each event. Following these modifications, the full-chain MC simulation was executed.

The simulation results, shown in figure 5.1, illustrate the direct consequence of varying the PY on the number of S1 photons produced and detected from a 0.270 keV

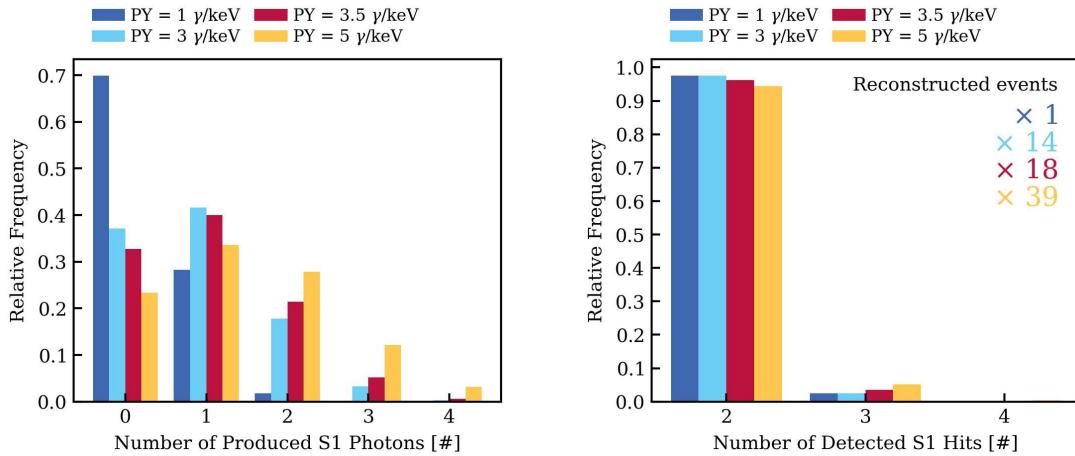


Figure 5.1: Left: Relative frequency distribution of the number of produced S1 photons from 0.270 keV electronic recoils, simulated for various photon yields (PY). **Right:** The corresponding relative frequency of detected S1 hits, after full-chain simulation and application of a 2-fold S1 coincidence selection. This demonstrates that while higher PY values increase the overall number of reconstructed events (indicated by the “x” scaling factors relative to 1 γ/keV), the range of detected S1 n -hits remains fundamentally unchanged.

energy deposition. As PY increases from 1 to 5 γ/keV , the mean of the produced S1 photon distribution shifts to higher values. However, even at 5 γ/keV , the number of produced photons predominantly remains very low (0-4 photons). Consequently, as can be seen in the right panel of figure 5.1, the number of events with {2, 3, 4} reconstructed photons shows a drastic increase when varying the PY (1-5) γ/keV , but the range of S1 n -hits remains the same despite the increased PY. This allows for analyses like this one, seeking a peak in reconstructed energy space, to model the signal by centering it on the average energy from MC simulations. This approach is valid because the MC average reconstructed energy is stable across the wide PY range (1-5) γ/keV (see figure 5.2).

Ultimately, the 0.270 keV energy deposition from ^{37}Ar L-shell decays is reconstructed in S1-S2 exclusively when there is an upward fluctuation in S1 production, coupled with favorable photon propagation. A distribution centered on the literature value for this decay (0.270 keV) is obtained only if the entire distribution of the energy depositions is represented in the reconstructed sample. This produces a bias in the reconstructed energy, known as Eddington bias [141, 142]. At 0.270 keV, this bias manifests as a reconstructed energy that is approximately 1.9 times the deposited energy, as shown in the right panel of figure 5.2.

The presented simulation studies clearly indicate that the standard requirement of at least three (3-fold) distinct PMTs registering photons within a ± 50 ns window, introduced in section 3.2, is too restrictive for the low photon counts from ^{37}Ar L-shell decays, effectively discarding all signal (see figure 5.3). To overcome this, the requirement is lowered to at least two (2-fold) distinct PMTs. This change is key to achieving a detectable signal rate, increasing the number of expected events by a factor of ~ 66 . The immediate consequence is an increase in the accidental coincidence (AC) background, which requires the development of dedicated mitigation techniques and criteria, detailed later in this chapter.

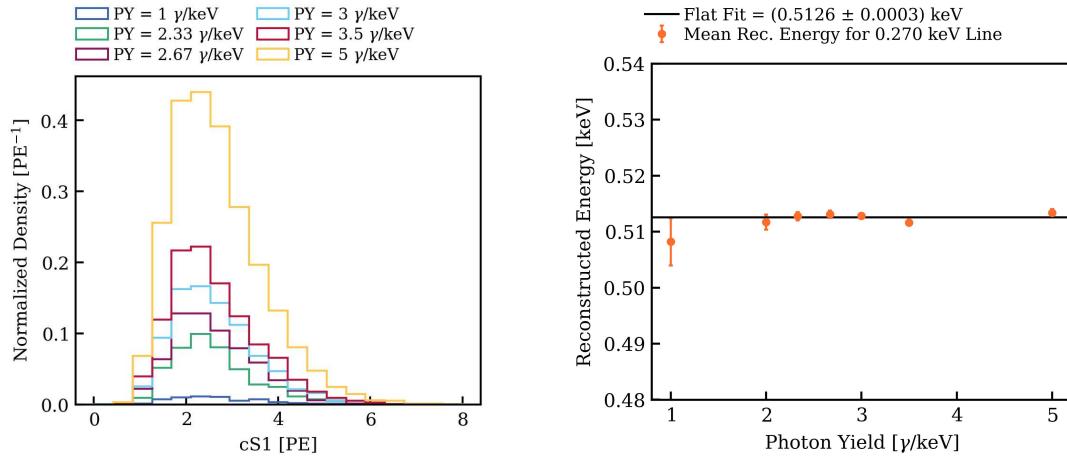


Figure 5.2: Left: Corrected S1 area (cS1) histograms for different photon yields (PY), demonstrating a consistent distribution shape that scales in normalization due to the increased number of reconstructed events at higher PY. **Right:** The mean reconstructed energy as a function of the input PY. The big difference between deposited energy and reconstructed energy for ^{37}Ar L-shell interactions is explained in detail in the text. The data points exhibit minimal variation, indicating that the average reconstructed energy is largely independent of the photon yield within the simulated range.

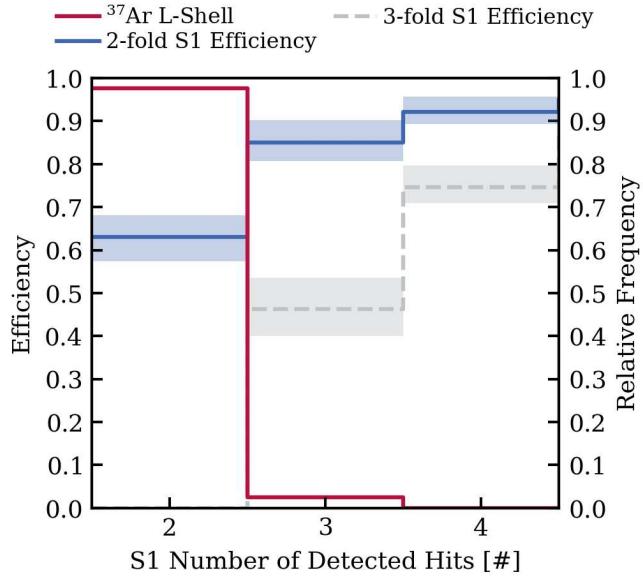


Figure 5.3: Impact of S1 tight coincidence selection on ^{37}Ar L-shell signal acceptance. The bar plot shows the simulated distribution of detected S1 photons from 0.270 keV depositions with a nominal photon yield of 3 γ/keV . The remarkable improvement is evident when comparing the 2-fold criterion, requiring at least 2 PMTs to register hits within ± 50 ns of the S1 peak, against the 3-fold criterion, which requires at least 3 PMTs. While the 2-fold selection keeps a substantial fraction of events, the 3-fold selection leads to basically no signal acceptance.

5.2 Sub-keV energy smearing

It is a known fact that energy resolution and PY are degenerate [143]. If more photons are produced for an interaction of a given energy, then the energy resolution will improve. However, as shown in section 5.1, variations in PY between (1-5) γ/keV affect the

event counts within the S1 n-hit range $\{2, 3, 4\}$, but do not alter the extent of the S1 n-hit distribution. Consequently, the energy resolution-PY dependence does not pose a difficulty for the signal modeling presented in the following.

As detailed in section 4.3, the reconstructed energy distribution for a monoenergetic interaction is described by a skewed Gaussian function. Although the parametric model itself is applicable in the sub-keV range, its parameters, fitted empirically from higher energies and lacking first-principles meaning, are not directly transferable (see figure 5.4). This was also observed in XENON1T with and without the inclusion of the ^{37}Ar K-shell point [91]. Given the lack of sub-keV monoenergetic calibration sources, the followed approach is to generate simulated datasets for monochromatic peaks with full-chain simulation. The parametric model is then fitted to these synthetic datasets. After verifying good agreement between the model and simulation, the width and skewness parameters of the skewed Gaussian are extracted as functions of energy.

In the plot in figure 5.4, the width of the skewed Gaussian can be observed versus the observable energy. The full-chain simulator correctly models the smearing effect, as validated by the comparison between parameters derived from ^{37}Ar K-shell (2.8 keV) and ^{83m}Kr (41.5 keV) calibration data. In fact, these parameters show agreement between data and simulations within 2%. It is observed that around 1.5 keV, the data points from the simulation deviate from the model presented in section 4.3. Since the full-chain simulation models both the microphysics and the detailed detector response, the behavior of these (brown) simulation points is deemed trustworthy. For this reason, the model for energies below 1.5 keV was established using a linear fit, representing the lowest-order fit that ensures good agreement with the data. The model incorporates an uncertainty band that comprises both statistical and systematic uncertainties. The systematic uncertainty includes the uncertainty on the charge yield in the sub-keV range, quantified by varying the charge yield by $\pm 1\sigma$ of literature values in simulation.

5.3 Characterization of the accidental coincidence background

AC background results when the event building algorithm, described in section 3.1, mistakenly pairs unrelated S1 and S2 signals within a time window T_{\max} , set by the detector's maximum drift time of about 2.2 ms. Accurate characterization is crucial, as the resulting AC events can mimic the signatures of genuine low-energy physical interactions. The expected AC rate between t_0 and t_1 is given by:

$$R_{\text{AC}} = \int_{t_0}^{t_1} R_{\text{isoS1}}(t) \times R_{\text{isoS2}}(t) \times T_{\max} dt, \quad (5.1)$$

where $R_{\text{isoS1}}(t)$ and $R_{\text{isoS2}}(t)$ are the rates of isolated S1 and S2 events, respectively. An “isolated” peak is a signal recorded by the detector without a physically associated partner from the same particle interaction. These signals appear constantly in the data stream and have origins in detector physics and instrumental effects. Despite the identification and estimation of several mechanisms producing isolated S1 and S2 signals, the existence of uncharacterized sources and the inability to isolate all the individual contributions from the total rates preclude a direct validation of these component estimates. Consequently, the LXe dual-phase time projection chamber (TPC) community currently lacks a fully validated, first-principles model for the AC background.

Isolated S1 peaks: These peaks typically do not originate from single-particle interactions. Instead, they are most often artifacts caused by the accidental temporal overlap

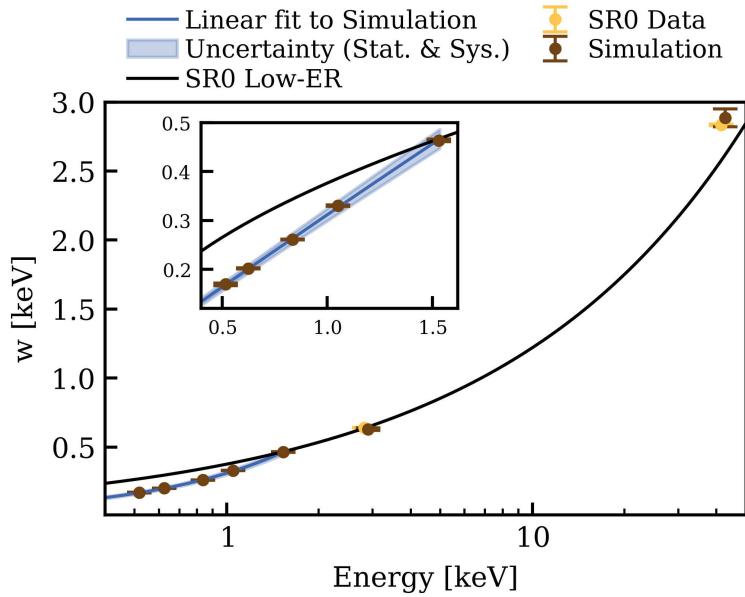


Figure 5.4: Energy resolution model for low-energy electronic recoils. The width parameter, w , of the skewed Gaussian energy resolution function is shown versus energy. The black curve represents the empirical model from section 4.3, which deviates from full-chain simulations (blue line) below ~ 1.5 keV. The SR0 ^{37}Ar K-shell and ^{83m}Kr calibration data points (yellow points) are in good agreement with the simulation results. Simulation uncertainties include systematic effects from charge yield variation in the sub-keV region, quantified by varying the charge yield by $\pm 1\sigma$ of literature values in simulation. The inset zooms in on the sub-keV region, where a linear fit to the simulation points is used to describe the width parameter below 1.5 keV. More details can be found in the text.

of lone hits, with the exception of signals from charge-insensitive regions. Lone hits refer to single-photoelectron pulses in the individual PMTs that are not sufficiently close in time to other hits to be clustered by the reconstruction algorithm, which uses a 700 ns window. Two primary sources contribute to lone hits. The first is considered a stable component, with a rate estimated to be around 8.5 kHz, thought to be caused by PMT dark counts, which are thermally induced emissions of photoelectrons that occur without an incident photon hitting the photocathode. In SR0, this rate originates from the 248 active top-array PMTs, with an average dark rate estimated at 12 Hz each, and the 229 active bottom-array PMTs, with an average dark rate estimated at 24 Hz each [9, 64]. The second source is correlated in time with high-energy (HE) interactions in the detector and is believed to result from PTFE fluorescence, which occurs due to VUV photon absorption [144, 145]. Isolated S1-like peaks can also arise from misclassified single-electron S2s or from events occurring in charge-insensitive regions. These regions include an approximately 200 kg LXe volume below the cathode and an about 112 kg volume near the detector edges, where local field configurations prevent electron extraction, precluding S2 signal generation [11]. Because of the relatively high rate of lone hits, the probability of two or more randomly coinciding within the S1 coincidence window of 100 ns is non-negligible, generating spurious S1 signals. In science data, and before any mitigation strategies are applied, the rate of these accidental 2-fold isolated S1s is ~ 15 Hz.

In the context of the ^{37}Ar calibration data, which is the main subject of this chapter, the most significant source of isolated S1s arises from K-shell interactions occurring in the

charge-insensitive volumes below the cathode and near the detector edges. With a total rate of such interactions throughout the active volume of about 19 Hz, this component in the charge-insensitive regions contributes ~ 1 Hz to the isolated S1 rate.

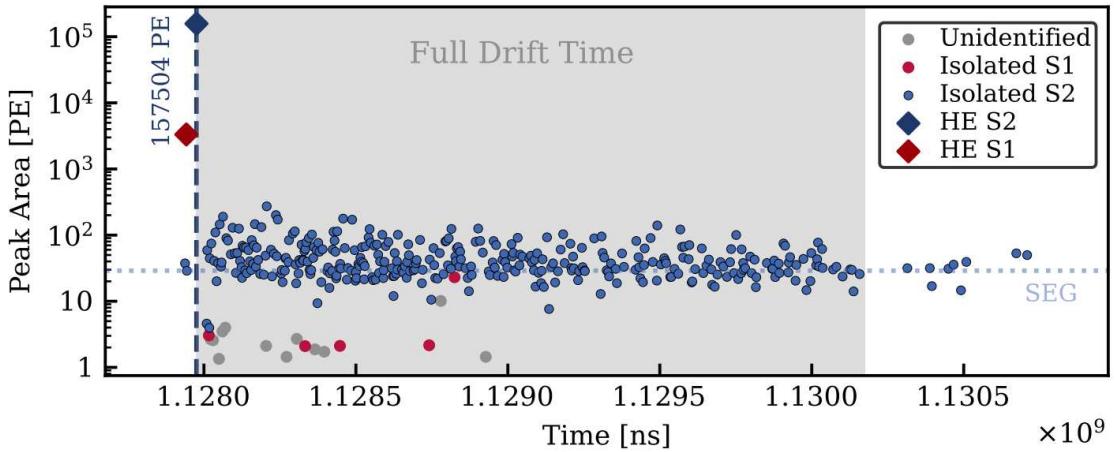


Figure 5.5: Example of single hits and small signals following a randomly selected high-energy interaction during a science run. A clear temporal relationship is observed, with a noticeable increase in the rate of isolated peaks after the high-energy interaction, followed by a significant drop in the rate of isolated S2 signals after one full drift time. Many of the isolated S2 peaks cluster near the single-electron gain (SEG) level. Refer to the text and in particular to section 3.1 for additional details on the classification between S1, S2, and unidentified peaks. Figure taken from [146].

Isolated S2 peaks: These signals are typically genuine secondary scintillations from small numbers of electrons. A prominent population of isolated S2s occurs immediately following HE interactions (>0.1 MeV), which are mainly caused by gamma rays from material radioactivity and produce S2 areas exceeding 10^5 PE. This component is characterized by three key features: its rate diminishes rapidly after one drift time (as can be seen in figure 5.5); its signals are spatially correlated with the primary interaction; and its rate scales positively with drift time and negatively with LXe purity [147]. The hypothesized origin is the temporary trapping and subsequent release of electrons by impurities within the LXe volume. A second, spatially uncorrelated component is also observed. These events occur far (>20 cm) from the primary S2 and show only a weak dependence on drift time and purity. Their origin, possibly related to longer-lived trapping states, is not yet fully understood [147]. A third component involves delayed electron extraction that persists for over 100 ms after a HE event. These signals are also spatially correlated with the initial interaction. Their physical origin is under investigation, with one possible mechanism being electrons trapped just below the liquid-gas interface [148, 149]. Moreover, intense, localized single-electron emission near anode transverse wires can create "hot spots" or "warm spots." In SR0, these effects were mitigated by temporarily reducing the anode voltage, at the cost of 11% live time [11]. The affected data-taking periods were omitted from the analysis. The total rate of all isolated S2s is on the order of 0.1–0.2 Hz in science data, making them a significant background despite being less frequent than isolated S1s.

For the ^{37}Ar calibration data specifically, the dominant source of isolated S2s is signals from L-shell interactions where the corresponding S1 scintillation light was either not produced or not detected (as explained in section 5.1 and shown in figure 5.1). These events, which dominate the isolated S2 population for this dataset, have a rate of ~ 1.7 Hz.

Suppression of correlated isolated signals

A few selection criteria are applied to suppress isolated peaks correlated with preceding HE events or with high local detector activity.

Time shadow: The temporal correlation between isolated signals and preceding HE interactions is evaluated using a "time shadow" metric. For any isolated peak, all HE events within a 1 s window are considered, and the one with the maximum ratio of $S2_{\text{pre}}/\Delta t_{\text{pre}}$ is identified as the dominant source of potential contamination. Here, $S2_{\text{pre}}$ is the S2 area of the HE event and Δt_{pre} is the time difference. Based on this quantity, thresholds of $<10.1 \text{ PE}/\mu\text{s}$ and $<38.2 \text{ PE}/\mu\text{s}$ are applied to two-hit and three-hit isolated S1 signals, respectively. This cut maintains signal efficiencies of 87% for two-hit and 96% for three-hit signals while reducing the isolated S1 rate by over 80% and 50%, respectively [48].

Position shadow: To reduce isolated S2s caused by delayed electron extraction, which retain spatial correlation with the parent HE event, a "position shadow" cut is applied. A Half-Cauchy distribution, $f(\Delta X, \Delta Y)$, empirically captures the spatial separation between an isolated S2 and a preceding large S2, with higher f values indicating a tighter spatial connection to the previous HE interaction. A two-dimensional selection is then defined using both the time shadow metric and the spatial distance f , rejecting events that exhibit strong correlations in both time and position. This selection is tuned to keep about 97.5% of signal events while removing roughly 53% of the correlated isolated S2 background [34].

Ambience cut: This selection criterion targets events that occur in a "busy" detector environment, regardless of any single preceding HE event. It defines a spatio-temporal region around an event of interest and rejects it if the local activity is too high. Specifically, it vetoes events if a 2 ms window before the S1 contains more than 100 lone hits, 3 unidentified peaks, 1 S1, or 1 S2 (area $< 60 \text{ PE}$). It also imposes a stricter spatial veto, rejecting events if any small S2 is found within a 6.7 cm radius. The signal acceptance is approximately 99% [34].

Model construction and validation

As a first-principles simulation of the AC background is not currently available, a data-driven approach developed by the XENON Collaboration using the axidence software package [150] is employed. This method constructs the model by directly sampling real isolated S1 and S2 peaks from the specific data category of interest (for example, using ^{37}Ar calibration data for ^{37}Ar analysis) and combining them to generate a synthetic AC dataset. This approach ensures that the instrument-specific characteristics of the isolated peaks are accurately replicated in the model for each data category. The method is inherited from XENON1T [151], with two improvements to create a more realistic model and mitigate potential biases. Firstly, to prevent bias from mismatched detector activity, isolated S1 and S2 peaks are paired based on similar time shadow values ($S2_{\text{pre}}/\Delta t_{\text{pre}}$), which serve as a proxy for local detector conditions. This ensures that simulated AC events reflect the ambient environment present during real coincidences. Secondly, the rates of isolated S1 and S2 peaks vary by $\sim 15\%$ and $\sim 20\%$, respectively, over timescales of a few days due to changing detector conditions. To capture this time dependence, the data are divided into shorter, sequential periods (typically a few hours each), and

artificial AC events are generated separately for each period, accurately reflecting the evolving background rates throughout the campaign. Once the isolated S1 and S2 peaks are paired, the resulting synthetic events are processed using the standard pipeline described in chapter 3, providing events that have the same analysis dimensions as events from genuine physical interactions.

The AC model's accuracy and systematic uncertainty are validated against AC-dominated control samples [48]. The most important test for this analysis uses ^{37}Ar calibration data, specifically events created by accidental pairings of S1s from K-shell decays in charge-insensitive regions and S2s from L-shell decays without associated S1s. The dataset, comprising approximately 500 events, shows strong agreement with the AC model's predictions, enabling a 5% constraint on the AC rate's systematic uncertainty.

5.4 S1/S2 BDT classifiers for accidental coincidence mitigation

While the analysis in chapter 4 relied primarily on a 3-fold S1 coincidence requirement to limit the AC background, lowering the threshold to 2-fold necessitates new and additional mitigation strategies. Two boosted decision tree (BDT) machine learning classifiers, based on the LightGBM framework [152], are employed for this purpose. To determine the optimal operating point for each classifier, the cut on the BDT score is chosen by maximizing the median discovery significance, Z . This significance is calculated using the Poissonian counting experiment approximation as outlined in [153]:

$$Z = \sqrt{2 \cdot (s + b) \cdot \ln \left(1 + \frac{s}{b} \right) - 2 \cdot s}$$

where s represents the number of signal events and b is the number of background events.

S2 BDT

The S2 BDT is developed to distinguish genuine S1-S2 pairs from isolated S1s and S2s based on electron drift and diffusion physics in LXe. For a genuine S2 signal, the electrons generated at an interaction site drift upwards due to the presence of the drift field, and their spatial distribution broadens to a Gaussian distribution along Z . This spatial smearing translates to a temporal width upon electron extraction from the liquid surface. The standard deviation of this temporal width, σ_{diffu} , exhibits a dependence on the electron drift time, t_{drift} , following the relation $\propto \sqrt{t_{\text{drift}}}$ [154]. For an isolated S2 forming an AC event, its reconstructed Z position (determined by the time difference to an unrelated paired S1) is an artifact and bears no physical relation to its intrinsic S2 width. The S2 BDT is trained using full-chain simulations of few-keV energy interactions as signal and a synthetic, data-driven AC background sample generated according to the procedure in section 5.3. It utilizes the following S2 signal characteristics as input features:

- **S2 rise time:** Reflects the initial spread and arrival characteristics of electrons at the gas phase. It is the time between the 10% and the 50% area quantiles of the summed waveform.
- **S2 50% width:** The time interval covering the central 50% of the S2 signal's area.
- **S2 90% width:** The time interval covering the central 90% of the S2 signal's area.

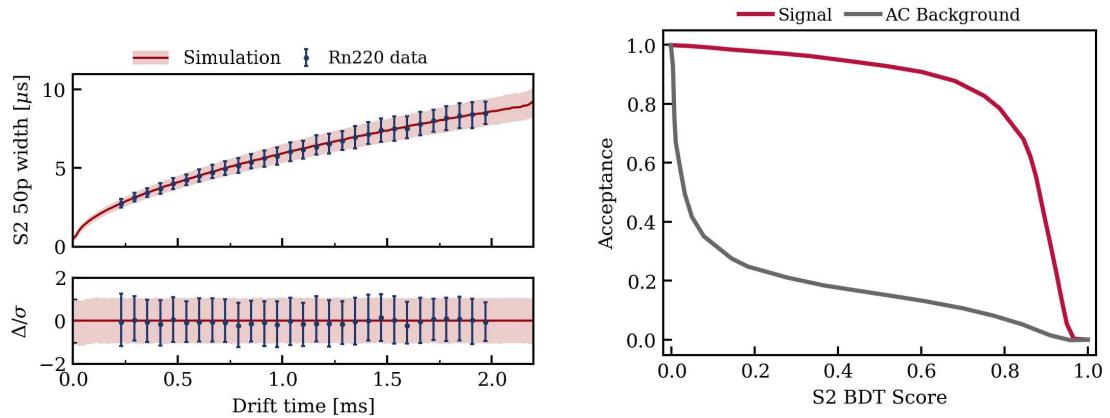


Figure 5.6: **Left:** Validation of the S2 signal width model. The median 50% width of the S2 pulse is plotted against electron drift time, showing excellent agreement between ^{220}Rn calibration data and full-chain simulations, used to train the S2 boosted decision tree (BDT). The error bars and band correspond to 1σ of the distributions. The bottom panel displays the normalized residuals. **Right:** Performance of the S2 BDT. Fraction of signal and accidental coincidence background events passing a given boosted decision tree score cut.

- **Z position:** Derived from the S1-S2 time difference.

The S2 BDT provides a score from 0 to 1 for each input event, where higher scores indicate a higher likelihood of the S1-S2 pair being from a physical interaction. Since the signal dataset is based on full-chain simulations, it is essential that the input features used by the S2 BDT, particularly the S2 temporal features and their correlation with Z , accurately reproduce those observed in real data. For the AC dataset, this condition is automatically met since it is generated directly from data. Figure 5.6 shows a validation of the simulation’s fidelity using ^{220}Rn calibration data. Similar validations were also performed using ^{37}Ar L-shell, exploiting the uniformity of the distribution of events.

The performance of the S2 BDT can be quantified by the area under the receiver operating characteristic (ROC) curve (AUC), a standard metric for evaluating binary classifiers, which is 0.95 for the S2 BDT. Applying the significance maximization procedure described above, a selection cut requiring an S2 BDT score > 0.80 is chosen. This cut effectively removes about 95% of AC events, while retaining approximately 80% of the physical interactions. The permutation scores [155, 156] indicate the relative importance of the features for the classification in the trained model. A higher number means the variable is more important. The scores for the S2 BDT are: 50% width (0.20), rise time (0.15), 90% width (0.14), Z (0.09). Since the S2 pulse shape parameters already encode some information about Z , its importance for the model is relatively low.

S1 BDT

To improve the discrimination between genuine physical S1 signals and temporal pile-up of S1 lone hits, the S1 BDT is used. It exploits distinct features between the two types of signals:

- **Maximum hit area:** Photons from physical S1 interactions can induce double photoelectron emission (DPE) at the PMT photocathode with a probability around 0.22, where a single VUV photon results in two photoelectrons being emitted [157]. This leads to a larger integrated charge area for that specific photon detection in

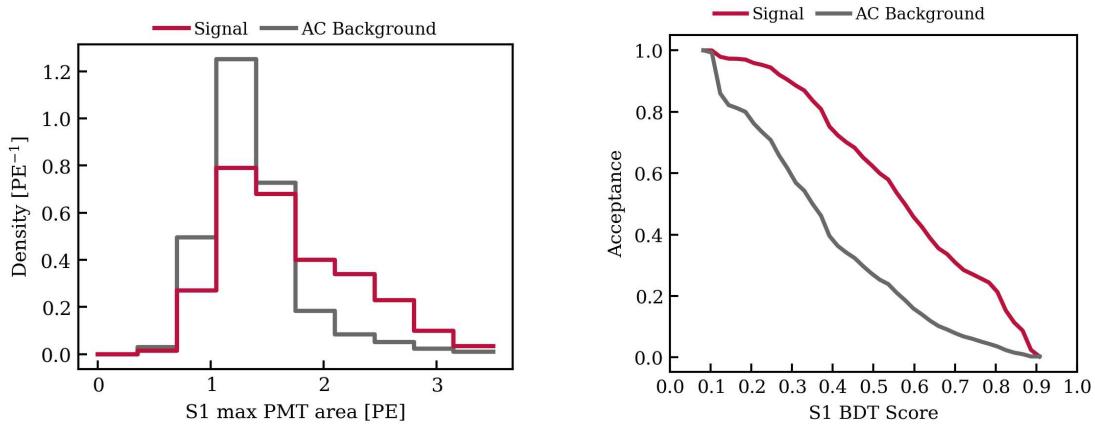


Figure 5.7: **Left:** Distribution of the S1 area in the largest-contributing photomultiplier tube channel, the most important discriminating variable for the S1 BDT. Genuine signals show a tail to larger areas due to double photoelectron emission, which is absent in accidental coincidence background. **Right:** Performance of the S1 BDT classifier. Signal and background acceptance as a function of the cut applied on the S1 BDT score.

the PMT compared to a standard single-photoelectron. Lone hits, which form AC S1s, do not exhibit the DPE phenomenon.

- **Minimum hit time difference:** The temporal profile of the hits that constitute a physical S1 signal is determined by the scintillation physics in LXe and optical propagation within the detector. In contrast, AC S1s are formed by the random coincidence of uncorrelated lone hits, resulting in a different inter-hit time difference distribution.
- **Number of contributing channels in top PMT array:** The fraction of S1 light detected by the top PMT array versus the bottom PMT array is a function of the interaction's position due to optical effects. AC S1s do not have this spatially-dependent hit distribution.
- **Hit number:** The number of hits in an S1 from a physical interaction depends on the energy deposition.

The S1 BDT demonstrates moderate discrimination power (AUC-ROC = 0.75); the feature importance ranking is as follows: maximum PMT area (0.10), minimum hit time difference (0.05), number of contributing channels in the top array (0.04), and total number of contributing channels (0.02). Following the optimization procedure outlined above, the optimal S1 BDT score threshold is found to be 0.34. Requiring events to have a score greater than this value preserves approximately 87% of the signals while rejecting around 46% of the AC background.

5.5 Selection criteria and efficiencies for ^{37}Ar L-shell search

This analysis implements variations to the traditional event selection criteria to address its unique challenges.

S2 selection: S2 signals are required to fall within the range of 120 to 500 PE, which corresponds to approximately 4 to 16 extracted electrons. This threshold is reduced

compared to the 200 PE lower bound used in the standard WIMP search [49] and the 500 PE lower bound in the low-ER analysis presented in chapter 4. Lowering the S2 threshold is needed to enable the detection of low-energy events, such as those from ^{37}Ar L-shell decays. A conservative 120 PE lower bound was chosen due to observed mismatches below this threshold in AC background model validation using calibration and sideband science datasets [48]. The effect of the S2 area range on signal acceptance was evaluated and accounted for. The other S2 selection criteria outlined in section 3.5 are also applied, as their effectiveness and validity persist. In addition, the S2 BDT classifier, detailed in section 5.4, is employed to further suppress the AC background.

S1 selection: S1 signals are required to have a total of either 2 or 3 PMT hits. This criterion guarantees full signal acceptance for a broad range of possible, and currently unknown, PY at 0.270 keV, as shown in figure 5.1. Additionally, S1 signals must exceed 1 PE in total area, which suppresses events arising from baseline fluctuations and preserves low-energy signals from interactions occurring in LXe. Events in which a single PMT contributes more than 4 PE to the S1 are excluded, as these are either artifacts or due to PMT-intrinsic radioactivity. As previously discussed, the S1 tight coincidence requirement is relaxed to at least two distinct PMTs registering photons within a ± 50 ns window, increasing by a factor of ~ 66 the expected signal rate compared to the 3-fold requirement. Given the low photon statistics (<4 detected photons) from ^{37}Ar L-shell decays, most of the standard S1 quality criteria described in section 3.5, such as those based on S1 hit pattern or the fraction of light detected by the top PMT array, are not applied in this analysis as they become ineffective. The S1 temporal width criterion is also removed, as its discriminating power is incorporated into the more effective S1 BDT classifier, which is introduced for improved AC background rejection at low energies (see section 5.4).

Fiducial volume (FV): A significant rate of single- and few-electron S2 signals is observed near the perpendicular support wires for the gate and anode electrodes. This results in higher isolated S2 rates and a corresponding rise in AC events. The S2 pulse shape is also modified in the vicinity of these wires [34]. These effects complicate the AC background modeling, which is the only background component in this search and is increased by a factor of roughly 6 near the perpendicular wires [49]. Moreover, applying the S2 BDT selection to events close to these wires would lead to systematic errors in signal acceptance, due to inaccurate modeling of the S2 pulse shape. To mitigate these issues, regions within 15 cm of each perpendicular wire are excluded from the FV. Additionally, events with Z coordinates outside the range -142 cm to -13 cm, as well as those with a radius larger than 60.15 cm, are removed from the analysis due to insufficient understanding of the detector's response and background rates in these regions. After applying these criteria and considering field distortion, the FV mass is (3.97 ± 0.20) t.

Accurate modeling of the expected signal as a function of reconstructed energy, where the ultimate statistical inference is performed, necessitates characterizing the cut acceptance: the fraction of true signal events expected to pass all selections. The acceptance evaluation uses the N-1 method as described in section 4.6, with explicit checks for cut correlations; correlated cuts are combined to avoid overestimating acceptance. The strategy for determining acceptance depends on the specific cut. Calibration data are optimal for this purpose, as they inherently capture all detector effects and the response to authentic, signal-like physical interactions. However, the lack of sub-keV

S1-S2 calibration data, also faced in the smearing model, limits this approach. For cuts defined only on S2 signal features, acceptance can be evaluated using a dataset in which the S2 signals are produced by the calibration source, regardless of whether the paired S1 is the true related partner. High-statistics sub-keV sources, such as ^{37}Ar L-shell, which this analysis primarily aims to detect via S1-S2 signals, most of the time generate S2 signals where the corresponding S1 photons are either not produced or not successfully reconstructed. This leads to instances of an S2 paired with an unrelated S1 or an S2 appearing as an S2-only signal. These events, occurring at approximately 1 Hz, provide a clean sample for calibration with negligible S2 background. This dataset is used to assess three cuts from section 3.5: S2 pattern consistency, S2 single-scatter identification, and anomalous S2 top-array signal fraction rejection, with respective acceptances of about 97%, 99%, and $>99\%$. Acceptance for AC background mitigation cuts based on event proximity (see section 3.5) and the event building efficiency are determined using the salting method implemented in `axidence` [150], yielding $\sim 82\%$ and $\sim 77\%$ acceptance, respectively. This technique inserts simulated signal event waveforms into segments of real data, specifically the summed PMT waveforms from the ^{37}Ar calibration of this analysis, thereby capturing the complex, not fully understood and modeled, ambient noise and isolated hit conditions. Event building efficiency losses are often driven by ambience interference, where unrelated background peaks in real data (like lone hits or single electrons) compromise the reconstruction of signal events; this is a much more pronounced challenge in low-energy analyses targeting signals with two detected photons and fewer than ten extracted electrons. For example, efficiency can be reduced if an event is skipped during building because too many competing peaks surround it (see section 3.1), or if a genuine S1/S2 signal is overshadowed by a larger, uncorrelated peak. Processing these salted datasets through the standard reconstruction chain and comparing results to an ambience-free simulation reference allows for a data-driven quantification of ambience-induced effects, such as event building failures. For the remaining cuts, a full-chain MC simulation of ^{37}Ar L-shell events, assuming a nominal PY of $3\ \gamma/\text{keV}$, is employed. Specifically, this dataset is used to determine the acceptance for the S1 and S2 BDT cuts (described in section 5.4). Among all cuts evaluated, the S1 BDT cut stands out as the only one exhibiting significant non-uniform acceptance across the 0-1 keV energy range; its acceptance averages $\sim 86\%$, varying between $\sim 79\%$ and $\sim 93\%$. This is primarily due to its most important discriminating feature, the maximum area observed by a single PMT (see section 5.4); larger values, suggesting DPE, enhance the BDT's confidence in identifying signals, leading to a higher BDT score and thus greater acceptance for the cut. Events with a single PMT area between 2 and 4 PE tend to reconstruct towards the higher end of the energy range, determining the energy dependence of the S1 BDT acceptance. Due to its 0.23 correlation with the S2 BDT, the S2 width-Z cut was evaluated together with the S2 BDT, resulting in an acceptance of about 88%. The overall cut acceptance for the ^{37}Ar L-shell analysis is then obtained by multiplying the individual efficiencies and it is shown in figure 5.8.

Systematic uncertainties associated with the cut acceptances are evaluated and propagated, to be included in the likelihood analysis. The primary sources are the S2 and S1 BDT cuts, each with an assigned $\pm 5\%$ uncertainty. These values are based on observed differences in cut acceptance between simulation and two independent validation datasets. Cathode-vicinity events, sourced from beta decays of ^{220}Rn daughters near the cathode during calibrations, form the first validation set. A fraction of the ionization electrons produced in each decay can be pulled back toward the cathode surface by the strong local electric field, reducing S2 signal magnitudes into the S2 ^{37}Ar L-shell region of interest (ROI). As these events are distributed across the entire horizontal

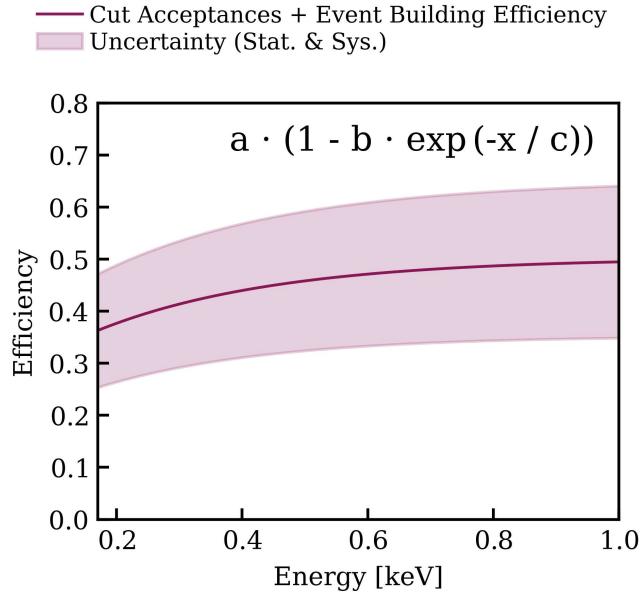


Figure 5.8: The plot shows the total selection efficiency for the ^{37}Ar L-shell search, including event building losses, as a function of energy. The uncertainty band reflects both statistical and systematic uncertainties, as detailed in the text.

plane, they are also used to check the X-Y spatial dependence of the efficiencies. The second validation set consists of surface events, mainly beta decays occurring near the PTFE wall panels of the TPC recorded in standard science runs. During the drift of these events, ionization electrons undergo transverse diffusion and some of them are trapped by neighboring PTFE walls, causing a reduction in S2 signal magnitudes to within the S2 ^{37}Ar L-shell ROI. Given the vertical distribution of these events, they are used to check the Z spatial dependence of cut efficiencies.

5.6 Inference and fit results

The search for ^{37}Ar L-shell decays was conducted using 4.09 days of data from the ^{37}Ar calibration period collected in December 2021. After applying the selection criteria detailed in section 5.5, a total of 1084 events were identified within the ROI. The reconstructed cS1-cS2 energy distribution and the vertical position (Z) distribution of the selected events are shown in figure 5.9. The statistical inference is performed using an unbinned profile likelihood method, which follows the same principles as the framework detailed in section 4.7. The likelihood function is constructed from the individual energies, E_i , of the $N_{\text{obs}} = 1084$ selected events as:

$$\mathcal{L}(\mu_s, \mu_{\text{AC}}, \theta_e) = \frac{(\mu_s + \mu_{\text{AC}})^{N_{\text{obs}}} e^{-(\mu_s + \mu_{\text{AC}})}}{N_{\text{obs}}!} \prod_{i=1}^{N_{\text{obs}}} \left[\frac{\mu_s f_s(E_i | \theta_e) + \mu_{\text{AC}} f_{\text{AC}}(E_i)}{\mu_s + \mu_{\text{AC}}} \right] \quad (5.2)$$

$$\cdot C_{\mu_{\text{AC}}}(\mu_{\text{AC}}) \cdot C_{\theta_e}(\theta_e). \quad (5.3)$$

Here, μ_s and μ_{AC} represent the signal and AC background events. The signal probability density function $f_s(E_i | \theta_e)$ is modeled as a skewed Gaussian with fixed shape parameters based on the pre-determined energy smearing model as described in section 5.2, but

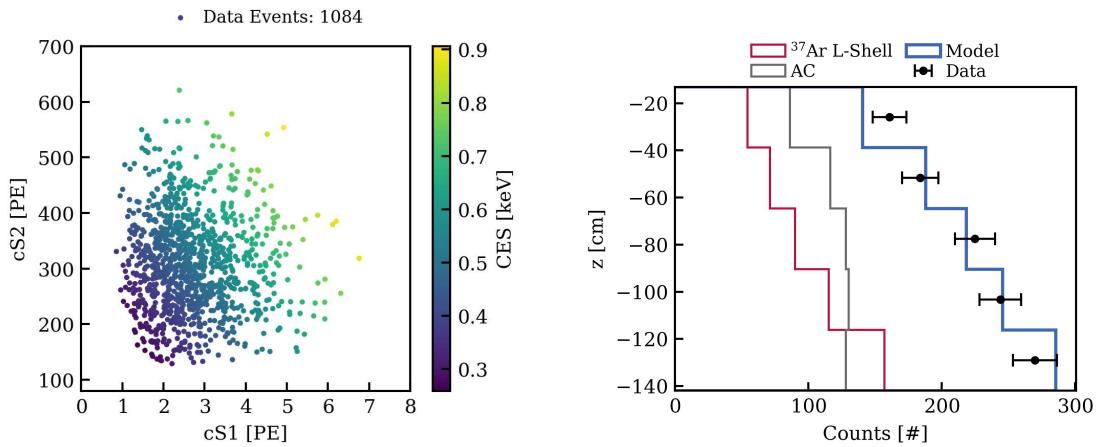


Figure 5.9: **Left:** Distribution of the selected events in the cS1-cS2 plane. The color scale indicates the reconstructed energy (CES). **Right:** Selected events as a function of the vertical position Z . The observed data are compared to the best-fit model components: ^{37}Ar L-shell signal and accidental coincidence background.

adjusted during the fit by the nuisance parameter θ_e , which accounts for signal efficiency uncertainty. The AC background PDF $f_{\text{AC}}(E_i)$ is non-parametric, derived from data-driven simulations that already incorporate efficiency effects. The term $C_{\mu_{\text{AC}}}(\mu_{\text{AC}})$ is a Gaussian constraint on the AC background rate, with a 5% uncertainty as detailed in section 5.3, while $C_{\theta_e}(\theta_e)$ is a Gaussian constraint on the total signal efficiency nuisance parameter θ_e .

Hypothesis testing and the determination of confidence intervals are based on the profile log-likelihood ratio (LLR) test statistic, as defined in equation 4.21. The statistical significance is derived using asymptotic approximations as $Z = \sqrt{q(\mu_s = 0)}$, following [153]. A 5σ significance level, combined with the absence of competing explanations, is the conventional criterion for declaring an observation [90].

The result of this fit is presented in the left panel of figure 5.10. The fit yields a best-fit value of 589 ± 29 AC background events and 490^{+220}_{-120} (stat. $^{+40}_{-40}$) ^{37}Ar L-shell signal events. The total uncertainty, representing the 68.3% confidence interval, is obtained from the LLR by profiling nuisance parameters, while the statistical component is determined with these parameters held fixed. The right panel shows the LLR curve, from which the $Z = 11.9\sigma$ significance is derived.

As an additional consistency check, the spatial distribution of the selected events is analyzed. No suspicious clustering is observed in the X-Y plane. The Z-distribution, presented in the right panel of figure 5.9, is of particular interest due to the anticipated Z-dependence of the signal, a consequence of the variations in light collection efficiency [42]. The observed Z-distribution shows good consistency with the model shape, normalized for the best-fit rates (p -value = 0.255). This agreement provides additional strong confidence in the obtained results.

5.7 Scintillation photon yield via simulation

This section details the procedure to translate the result of the fit, the number of observed signal events, into the physical quantity of interest: the scintillation PY at 0.270 keV. The conversion is nontrivial due to the Eddington bias and because a substantial fraction of ^{37}Ar L-shell energy depositions produce either no photons or only one, as discussed

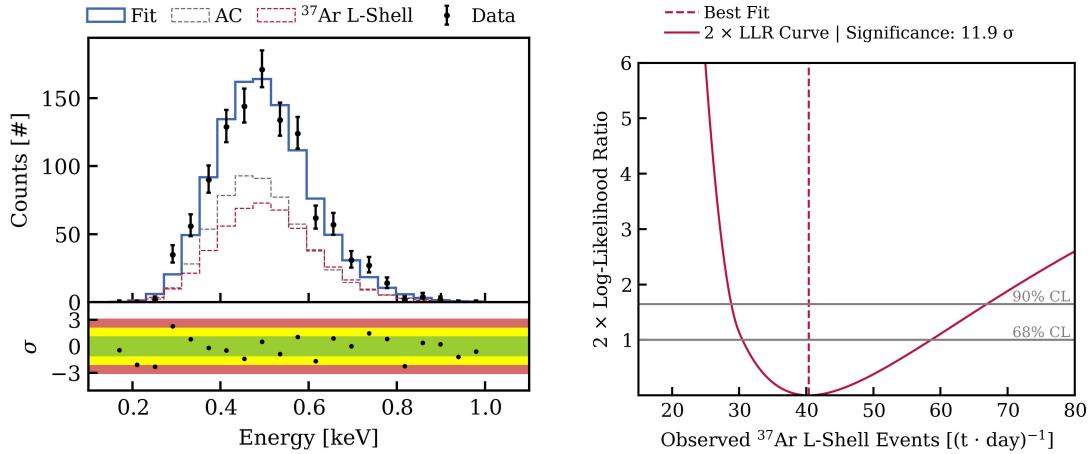


Figure 5.10: Left: Best-fit result to the reconstructed energy spectrum of the selected events. The data are shown as black points with statistical uncertainties. The total fit model is the solid blue line, composed of the ^{37}Ar L-shell signal (dashed red) and the accidental coincidence background (dashed gray). The lower panel shows the normalized residuals. **Right:** The profile log-likelihood ratio as a function of the number of observed ^{37}Ar L-shell events. The minimum of the curve corresponds to the best-fit value. The intersection with the horizontal lines indicates the boundaries of the 68% and 90% confidence intervals. The detection significance is 11.9σ .

in section 5.1. Therefore, the usual straightforward approach involving the gain factors (g_1, g_2) cannot be applied. The adopted methodology relies on a dedicated full-chain simulation [139] to establish a direct mapping between the number of observed signal events and the PY value used in the microphysics part of the simulation. First, the total number of ^{37}Ar L-shell energy depositions, $N_{L,\text{dep}}$, that occurred within the FV during the 4.09-day calibration period is estimated. This is achieved by measuring the number of energy depositions from the well-understood ^{37}Ar K-shell (2.82 keV) decays, described in detail in section 4.2, in the same 4.09-day dataset, and applying the known measured L/K branching ratio of $(10.11 \pm 0.44)\%$ [82]. After accounting for selection efficiencies for the K-shell events, the total number of L-shell depositions is determined to be $N_{L,\text{dep}} = (6.1 \pm 0.4) \times 10^5$. This absolute number of depositions serves as input for full-chain MC simulations. In each simulation, $N_{L,\text{dep}}$ events of 0.270 keV are generated, with the PY varied across the $(1\text{-}5\text{ }\gamma/\text{keV})$ range by applying the same median shifting approach to the NEST-generated photon distribution that was described in section 5.1. Multiple consistent measurements of the electron yield from different LXe dual-phase TPCs have been performed with ^{37}Ar L-shell decays; see [82] and references therein. For this analysis, small variations in the charge yield are not critical, as the S2 signals from ^{37}Ar L-shell decays are well above the detection threshold. The effect of increasing the electron yield from 67 to 73 e^-/keV was tested and found to lead to a variation that is not statistically significant on the reconstructed number of events, with a relative change well below 1%. Despite this negligible impact, the expected charge yield variation with different electric fields was considered. For the full-chain simulations, the value measured in XENONnT was used, which was obtained from an S2-only analysis and the appletree software [158], yielding $70.0 \pm 2.7 \text{ e}^-/\text{keV}$ [159]. Each simulated dataset is then processed through the exact same analysis chain used for the data. Where selection criteria are not fully validated in simulations, a data-derived acceptance curve is applied. This generates a mapping function, $\text{PY} = f(N_{\text{obs}})$, as shown in figure 5.11. The uncertainty on $N_{L,\text{dep}}$ is accounted for by generating

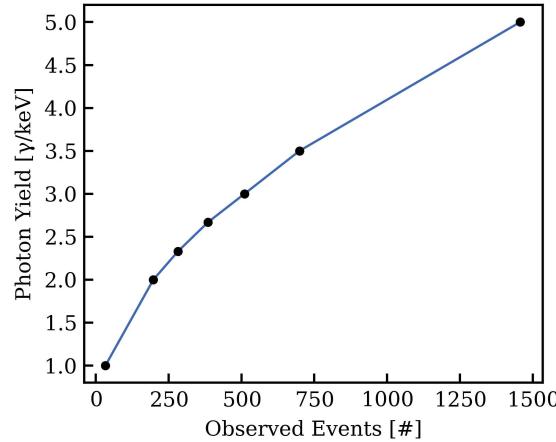


Figure 5.11: Mapping of observed signal events to the scintillation photon yield (PY) for 0.270 keV electronic recoils. This curve is generated using full-chain Monte Carlo simulations, varying the input PY and recording the resulting number of events that pass the entire analysis chain. Additional similar curves are also generated to propagate the systematic uncertainty on the initial estimated number of L-shell energy depositions ($N_{L,\text{dep}}$), refer to the text for further details.

equivalent mapping curves for its $\pm 1\sigma$ variations. The best-fit number of signal events together with the combined uncertainty, 490^{+220}_{-120} , is then translated into a PY using these mapping functions. Combining all the uncertainties, the scintillation PY for 0.270 keV ERs is measured to be:

$$\text{PY} = 2.96 \pm 0.08 \text{ (stat.)} {}^{+0.57}_{-0.33} \text{ (syst.) } \gamma/\text{keV} = 2.96^{+0.58}_{-0.34} \gamma/\text{keV} \quad (68\% \text{ C.L.})$$

A breakdown of the relative contributions from each source of uncertainty is provided in section 5.7.

Source of Uncertainty	Lower Unc.	Upper Unc.	Combined Unc.
Selection criteria	81.4%	90.1%	87.9%
S1 reconstruction	6.0%	6.6%	6.5%
K-derived L-shell energy deposits	7.1%	1.5%	3.0%
Statistical	5.1%	1.7%	2.6%
Energy smearing	0.4%	<0.1%	0.1%

Table 5.1: Percentage contributions to the total uncertainty on the ^{37}Ar L-shell photon yield measurement.

This result represents the first direct calibration of scintillation PY in LXe in the sub-keV energy regime. The energy deposition in ^{37}Ar L-shell is almost entirely ($>99.8\%$ [84]) due to the Auger cascade following an inner-shell vacancy, so X-ray contamination is negligible. This is significant, since it is well known that X-rays and electron interactions produce different PYs given the same experimental conditions, necessitating separate modeling in NEST [32]. Figure 5.12 places this measurement in the context of other experimental data and compares it with the NEST v2 [33] model prediction for ERs from electron scattering. Because NEST’s parameters are based on experimental data, and our measurement explores energies ten times lower than the closest calibration point, comparisons with the model should be made with caution. While the measurement sits

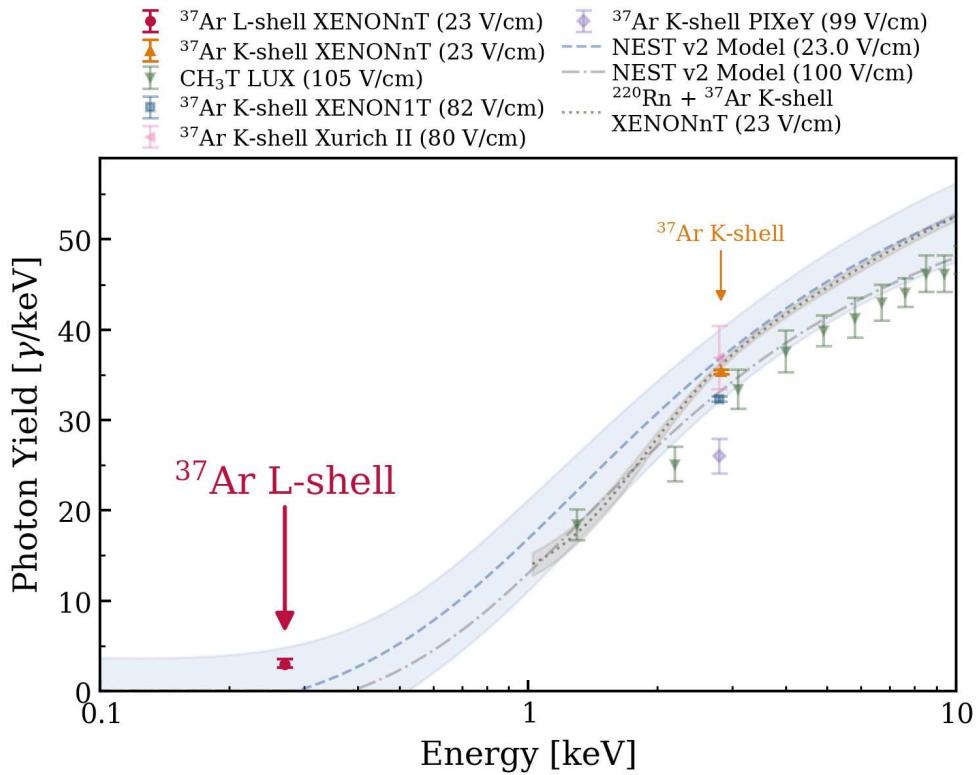


Figure 5.12: First measurement of the scintillation photon yield for sub-keV electronic recoils in liquid xenon. The result of this work for ^{37}Ar L-shell decays at 0.270 keV is shown in red with its statistical and systematic uncertainties. The result at 2.82 keV from ^{37}Ar K-shell decays derived in section 4.2 is also presented. This work’s results are compared to the NEST v2 model prediction for two different electric fields [33], along with other experimental measurements. These include data from LUX [160], from Dahl’s work [21], XENON1T [82], Xurich II [161], PIXeY [162], and the fitted response from XENONnT calibration data [125].

above the NEST central value, it remains compatible within the model’s wide uncertainty band at this low energy. By detecting the first sub-keV scintillation signal in LXe, we demonstrate that traditional S1+S2 signal reconstruction remains viable for sub-keV ER analyses in dual-phase LXe TPCs. This experimental data provides essential input for refining NEST simulation parameters and predictive accuracy at these low energies.

5.8 Conclusion and discussion

This chapter reports the first direct calibration of sub-keV ERs in LXe, successfully detecting ^{37}Ar L-shell decays at 0.270 keV with an 11.9σ significance using a 4.09-day dataset and the full S1+S2 event reconstruction. This analysis determined a scintillation photon yield of $\text{PY} = 2.96^{+0.58}_{-0.34} \gamma/\text{keV}$. The result was permitted by lowering the S1 coincidence criterion to a 2-fold requirement, allowing the detection of two-photon signals. The AC background increased as a result of this relaxation and was precisely modeled through a data-driven approach, then suppressed with an improved selection strategy including S2 and S1 BDTs and additional criteria. Detector response in the sub-keV ER regime, signal generation, and sub-keV PY from reconstructed events were all modeled using validated full-chain simulations taking into account detector and selection effects.

This work delivers a measurement that advances the understanding of the sub-keV ER emission model and serves as a key input for the NEST semi-empirical framework. An accurately constrained emission model is indispensable for robustly characterizing ER events, whether they are treated as a background component in dark matter nuclear recoil searches or when BSM signals are being investigated within them. Recent XENONnT analyses highlight this need: limited understanding of the sub-keV ER emission model forced the adoption of a 100% systematic uncertainty in ER background modeling for both light dark matter searches and solar ^8B neutrino CEvNS measurements [48, 163].

Chapter 6

Sensitivity to BSM physics in XLZD's low-energy electron recoils

Next-generation liquid xenon (LXe) observatories, such as the proposed XLZD experiment introduced in section 1.6, are designed to explore the landscape of rare event searches with unprecedented sensitivity. By scaling up the target mass and further reducing backgrounds, these experiments aim to probe lower cross sections for Weakly Interacting Massive Particle (WIMP) dark matter (DM) reaching the limit set by the neutrino fog, while also investigating other beyond the Standard Model (BSM) physics channels [53]. An important area in this exploration is the low-energy electronic recoils (ER), particularly the few-keV to sub-keV regime, where the expected event rates of several BSM signals are predicted to peak or increase. Previous chapters established the experimental foundation for accessing this regime. Chapter 4 analyzed XENONnT ER data using a standard S1 tight coincidence requirement of three or more (3-fold) photomultiplier tubes (PMTs) within ± 50 ns around the maximal amplitude of the S1 waveform, see section 3.2. This allows for an energy threshold of 1 keV, with a signal efficiency of approximately 10%. Chapter 5 then presented the first-ever calibration of sub-keV monoenergetic ER signals in liquid xenon using ^{37}Ar 0.270 keV L-shell decays. This was made possible by reducing the S1 coincidence requirement to 2-fold, developing machine learning algorithms to address the increased accidental coincidence (AC) background, and modeling the signal with validated full-chain simulations. Building on this work, the current chapter presents a quantitative evaluation of the BSM reach in low-energy ERs for a future XLZD-scale observatory, emphasizing the role of the newly validated analysis framework with a sub-keV energy threshold.

6.1 Modeling signals: 2-fold vs. 3-fold coincidences

As introduced in section 3.2, reducing the S1 tight coincidence requirement from 3-fold to 2-fold improves the energy threshold accessible to dual-phase LXe time projection chambers (TPCs), potentially enhancing sensitivity to different BSM physics channels. As shown in the left panel of figure 6.1, the S1 efficiency decreases rapidly at low energies due to the reduced probability of producing sufficient scintillation photons to trigger the coincidence requirement. The shift between the 2-fold and 3-fold curves is about 0.3 keV, providing access to a previously inaccessible energy regime where many BSM signals with steep low-energy spectra are expected to exhibit increased or peaked event rates.

The signal models considered in this analysis are identical to those detailed in section 4.5, including neutrino millicharge (NMC), neutrino magnetic moment (NMM), neutrino non-standard interactions (NSI) with both vector and scalar mediators, solar axions, and bosonic DM candidates. Each model's theoretical spectrum is convolved with the appropriate detector response functions for both 2-fold and 3-fold coincidence

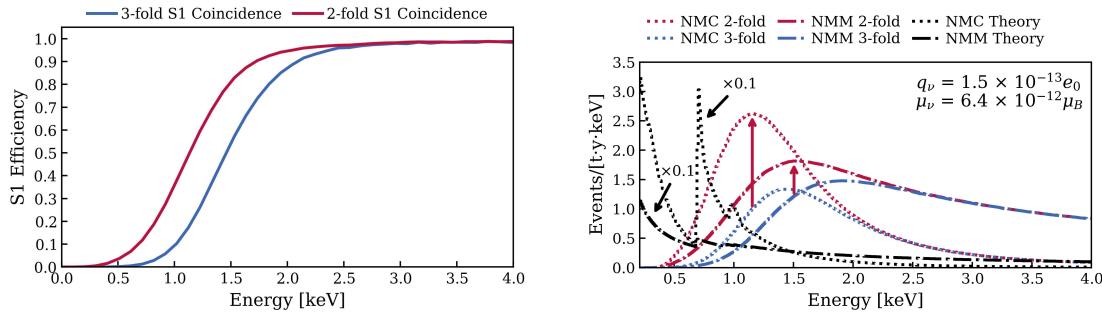


Figure 6.1: **Left:** S1 efficiency, including both photon production and detection, for criteria requiring a coincidence of at least 2 (2-fold) or 3 (3-fold) PMTs within ± 50 ns around the S1 peak. This illustrates the lower energy threshold achievable with the 2-fold selection. **Right:** Event rate spectra for neutrino millicharge (NMC, $q_\nu = 1.5 \times 10^{-13} e_0$) and magnetic moment (NMM, $\mu_\nu = 6.4 \times 10^{-12} \mu_B$) signals. The red and blue curves display the spectra after applying detector resolution and efficiency for 2-fold and 3-fold analyses, respectively. Vertical red arrows highlight the increased signal rates at low energies achieved with the 2-fold selection.

requirements, accounting for S1 efficiency, detector resolution, and selection acceptance. A basic quantitative comparison is presented in table 6.1, which shows the rate enhancements for the 2-fold analysis across all considered models within the (0–3) keV energy range. Due to the approximate E_r^{-2} dependence of its differential cross section, the NMC

Signal model	Rate gain (2-fold vs. 3-fold)
Neutrino millicharge (NMC)	+80%
Neutrino non-standard interaction (NSI) vector	+45%
NSI scalar	+34%
Neutrino magnetic moment (NMM)	+33%
Solar axion	+28%
Bosonic dark matter	(1–0.5) keV

Table 6.1: Relative increase in signal rate for 2-fold vs. 3-fold S1 coincidence selection in the (0–3) keV energy range for different new physics models detailed in section 4.5.

signal shows the largest gain, with the detectable signal rate up by about 80%.

Boosted dark matter

In recent times, interest in boosted DM to probe $\mathcal{O}(\text{keV} - \text{MeV})$ DM candidates in terrestrial experiments has grown significantly. The maximum energy that DM particles from the standard DM halo can deposit in a detector is capped by their velocity, which is limited by the Milky Way's escape velocity of around 544 km/s [164]. For a DM-electron interaction, this corresponds to a maximum energy deposition of about 0.5 keV in LXe. At this energy, the S1 efficiency is so low that the proposed 2-fold analysis cannot effectively probe these signals, as shown in figure 6.2. However, through known boosting mechanisms, such as interactions with high-energy particles like cosmic rays or upscattering by the Sun, DM can acquire substantial velocity boosts, allowing it to deposit more energy in a detector. Considering the very steep rise of event rates at low energies, a 2-fold analysis of boosted DM is particularly appealing. A notable example, which will be considered in the following, is the solar reflected DM (SRDM) model, where halo DM particles are upscattered by the Sun [165, 166], though other boosting

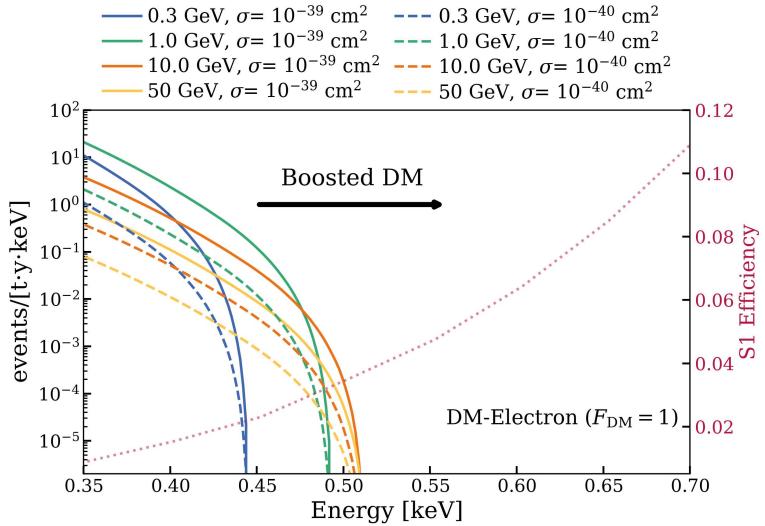


Figure 6.2: Theoretical energy spectra for standard (non-boosted) halo dark matter (DM) interacting with electrons in liquid xenon detectors. The plot shows the differential event rates for various DM masses and interaction cross-sections, assuming a standard halo model and a form factor $F_{DM}(q) = 1$. All spectra exhibit a sharp cutoff around 0.5 keV, which, in this scenario, corresponds to the maximum recoil energy from a DM particle with a velocity limited by the galactic escape velocity of around 544 km/s. The arrow illustrates that boosting mechanisms shift the signal spectrum to higher, more detectable energies. The S1 detection efficiency is also shown (red dotted line, right axis), emphasizing the difficulty of observing events in this low-energy regime.

mechanisms would work along similar lines. The SRDM scenario considered here is based on heavy mediators with a DM form factor $F_{DM}(q) = 1$, which has recently been explored by the XENONnT and PandaX-4T collaborations [167, 168] and other direct detection experiments [169, 170]. Compared to the other signal models listed in table 6.1, the advantages of a 2-fold analysis in this context are much more significant, because of the faster increase in event rate at lower energies. In fact, considering an SRDM particle with DM-electron cross section of $1 \times 10^{-38} \text{ cm}^2$, the gains in the rate between 2-fold and 3-fold analysis in the (0-3) keV range (same as table 6.1) are:

- $\sim 122\%$ for SRDM with DM particle mass of 2 MeV,
- $\sim 514\%$ for SRDM with DM particle mass of 30 keV.

Although boosting the signal rate can improve sensitivity, this only works if the background stays low and well understood; this is discussed in the next section.

6.2 Background projections for XLZD

The background model for low-energy ER searches in the next-generation XLZD observatory is based on the well-established framework used by XENONnT. This section focuses on background components expected to show a change in their rate per unit mass compared to XENONnT. Details on other components can be found in section 4.4. The projected values and assumptions are based on planned technological improvements and the scaling properties of the detector described in [53]. The AC model is a unique challenge for this prediction because there is no first-principles model, as explained in section 5.3. This will be discussed in more detail in section 6.3.

^{214}Pb : The beta decay of ^{214}Pb , a daughter of ^{222}Rn , is the main ER background at the lowest energies in XENONnT. For XLZD, reducing this background is a major objective. The collaboration aims to lower the ^{222}Rn concentration to a target of $0.1 \mu\text{Bq}/\text{kg}$. Achieving this reduction is the focus of an intensive and diverse R&D program, which is exploring several innovations to mitigate ^{222}Rn . A more powerful online radon removal system, capable of purifying xenon at around 10 tonnes per day, is in development [171]. Thin, clean, and tight coatings are being investigated to block recoil- and diffusion-driven ^{222}Rn emanation, especially from the cryostat [172]. A new gas-tight TPC design is being explored to separate the active target from the more radon-contaminated outer volume [173]. Moreover, by optimizing the detector design to limit the convection velocity field, further reduction of radon-induced background can be achieved using software tagging algorithms [50, 174]. Further details can be found in [53].

Materials: Gamma rays from radioactivity in detector components create background events through Compton scattering. With a larger detector and more material surface area, the overall number of events increases, but the rate per unit mass in the fiducial volume decreases. This is due to the reduced surface-to-volume ratio, which approximately follows a mass $^{-1/3}$ dependence. For XLZD, the projected spectrum is obtained by applying a scaling factor ~ 0.39 to the XENONnT spectrum, based on the nominal relevant design parameters. This is a conservative estimate and does not account for possible improvements in material radiopurity compared to XENONnT.

6.3 Accidental coincidence projections for XLZD

Given that the AC background, as explained in section 5.3, lacks a first-principles model, its projection poses a distinct challenge. To address this, its rate, R_{AC} , is projected by scaling the rates of its individual constituent components, building upon the validated XENONnT model. The fundamental relationship is:

$$R_{\text{AC}} \propto R_{\text{isoS1}} \times R_{\text{isoS2}} \times T_{\text{max}}, \quad (6.1)$$

where R_{isoS1} and R_{isoS2} correspond to the rates of isolated S1 and S2 signals, respectively, and T_{max} is the maximum electron drift time. Next, we analyze each component’s scaling.

Maximum drift time (T_{max}): The drift time in XLZD is determined by its significantly longer drift distance and the projected higher electron drift velocity. XLZD is designed with a drift length of 2.97 m, approximately double that of XENONnT. However, it is projected to operate at a higher electric field, with a baseline goal of 100 V/cm, which is notably higher than the approximately 23 V/cm achieved in the XENONnT SR0 analysis, but it has been reached and surpassed by the LZ experiment [50]. This higher field results in an increased electron drift velocity, and the net effect is a maximum drift time of about 2.2 ms, which is slightly shorter than in XENONnT. The scaling factor applied to this component is therefore ~ 0.98 .

Isolated S1 rate (R_{isoS1}): The rate of isolated S1 signals relevant for low-energy searches is dominated by the random coincidence of PMT dark counts, as discussed in section 5.3. Herein, for a baseline comparison between XENONnT and XLZD, the only assumed difference is the number of PMTs: 477 active in XENONnT’s SR0 analysis versus 1182

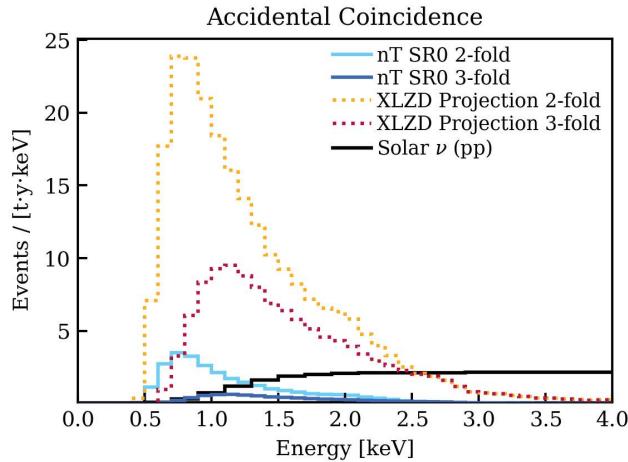


Figure 6.3: Energy distributions of accidental coincidence background. The figure shows the distributions for XENONnT SR0 with 2-fold and 3-fold S1 tight coincidence requirements (solid lines), along with their projections for XLZD (dotted lines). The spectrum from solar neutrinos (pp) is also displayed for comparison.

PMTs planned for XLZD. By treating the total dark count rate from all PMTs as a single Poisson process, the occurrence of these random and independent dark pulses within a defined time window can be modeled using Poisson statistics. Let N_{PMT} be the total number of PMTs, $r_{\text{pmt-dark}}$ be the average dark count rate per PMT, and Δt be the duration of the coincidence window. The rate of such m -fold coincidences reads:

$$R_m = \frac{(N_{\text{PMT}} \cdot r_{\text{pmt-dark}} \cdot \Delta t)^m e^{-N_{\text{PMT}} \cdot r_{\text{pmt-dark}} \cdot \Delta t}}{m! \cdot \Delta t}. \quad (6.2)$$

Based on equation 6.2, this leads to an increase from XENONnT to XLZD in the 2-fold isolated S1 rate by a factor of ~ 6.1 and by a factor of ~ 15.2 in the 3-fold isolated S1 rate. Given that for typical values, the exponential term $e^{-N_{\text{PMT}} \cdot r_{\text{pmt-dark}} \cdot \Delta t} \simeq 1$, the scaling factor for an m -fold coincidence can be well-approximated by $(N_{\text{PMT}}^{\text{XLZD}} / N_{\text{PMT}}^{\text{XENONnT}})^m$. It is important to note that for a given detector with a fixed number of PMTs, the m -fold coincidence rate approximately follows an $(r_{\text{pmt-dark}})^m$ scaling. Therefore, a fractional reduction of a in $r_{\text{pmt-dark}}$ will lead to an approximate fractional reduction of $1 - (1 - a)^m$ in the m -fold rate R_m . For example, in a 2-fold case ($m = 2$), a 10% reduction in $r_{\text{pmt-dark}}$ translates into a $\sim 19\%$ reduction in the isolated S1 rate. This demonstrates that even minor improvements in the dark count rate per PMT in XLZD compared to XENONnT can yield a sizeable impact on the overall AC background rate.

Isolated S2 rate (R_{isoS2}): The origins of isolated S2 signals, believed to be related to surface phenomena and electron trapping, are discussed in section 5.3. A commonly accepted assumption is that the isolated S2 rate scales with the detector’s internal surface area [175], leading to a projected increase by a factor of ~ 5 for XLZD compared to XENONnT. However, another factor influencing the isolated S2 rate is the electron extraction efficiency from the liquid to the gas phase. Given that a primary source of isolated S2s is attributed to delayed extraction of trapped electrons, increasing this efficiency from $\sim 50\%$ in XENONnT to a target of $>90\%$ in XLZD would suppress this background component by a factor of ~ 5 . As a result, the overall isolated S2 rate in XLZD is expected to remain comparable to that in XENONnT SR0.

Selection

The XLZD AC background projection employs a refined selection, derived from the XENONnT criteria with specific modifications adopted from the low-energy ER 3-fold coincidence SR0 search detailed in chapter 4. In that context, an important requirement was to maintain blindness for the subsequent WIMP analysis, so the region below $ER - 2\sigma$ in cS1-cS2 was not considered. While this WIMP-blinding requirement effectively reduces the AC background, it also causes a $\sim 2\%$ loss in signal efficiency below 10 keV [64]. Removing this specific constraint, while keeping the S2 threshold of 500 photoelectrons (PE), results in a 3-fold AC rate of 2.12 events/(tonne · year), an increase by a factor of ~ 3.5 . For the purposes of this chapter's analysis, the XLZD projections assume a standard scenario without the WIMP-blinding cut, implementing an alternative selection strategy for AC reduction. This includes a $\pm 3\sigma$ cut around the ER band in the cS1-cS2 space, calibrated using ^{220}Rn data taken in SR0 with a 2-fold selection. This selection, effective for both 2-fold and 3-fold coincidences, exploits the differing distributions of ER and AC events [125]. Additionally, the S2 boosted decision tree (BDT) described in section 5.4 is tuned to a working point of 0.627 to ensure a signal acceptance of 99.5%, making this choice independent of the signal model used in the analysis. This requirement provides an additional AC background reduction of 48% (2-fold) and 42% (3-fold). The S1 BDT described in section 5.4 is not considered, as its current performance does not significantly reduce the AC background without causing considerable, signal-dependent acceptance loss, which is beyond the scope of this work. With this new selection, the baseline XENONnT AC rates are (0.79 ± 0.04) events/(tonne·year) for the 3-fold coincidence case and (2.74 ± 0.14) events/(tonne·year) for the 2-fold one in XENONnT SR0. Applying the scaling factors described in this section to these rates, the projected AC backgrounds for XLZD are (11.8 ± 0.6) events/(tonne·year) and (23.8 ± 1.2) events/(tonne·year) for the 3-fold and 2-fold analyses, respectively. The obtained distributions in the energy space for the different cases are shown in figure 6.3.

6.4 XLZD sensitivity for different accidental coincidence scenarios

The sensitivity of XLZD is projected using the analysis framework detailed in section 4.7. An exposure of 200 tonne-years is assumed, which is the nominal case for XLZD [53]. The background model is constructed from the components detailed in section 6.2, section 6.3, and section 4.4. The AC background projection, as discussed earlier, is subject to significant uncertainty, including potential, yet unrealized, improvements arising from ongoing and planned intensive R&D in this area. To study its impact on the sensitivities, an "AC background factor" is introduced, which scales the projected AC rate, where a factor of 1 corresponds to the energy spectra shown in figure 6.3. The smearing models applied to all other components are those described in section 4.3 for the 3-fold analysis and section 5.2 for the 2-fold analysis. The selection efficiency is based on the results presented in section 4.6, with the additional criteria described in section 6.3. Upper limits are set at a 90% confidence level (C.L.) for selected signals described in section 4.5 and 6.1.

Figure 6.5 illustrates XLZD's projected 90% C.L. sensitivity to the NMC (q_ν) and NMM (μ_ν), plotted against the AC factor for both the standard 3-fold analysis and the low-energy threshold 2-fold analysis. The 2-fold analysis consistently demonstrates superior sensitivity to both models across the entire range of AC background assumptions. At the baseline AC projection (factor = 1), the NMC median limit is improved

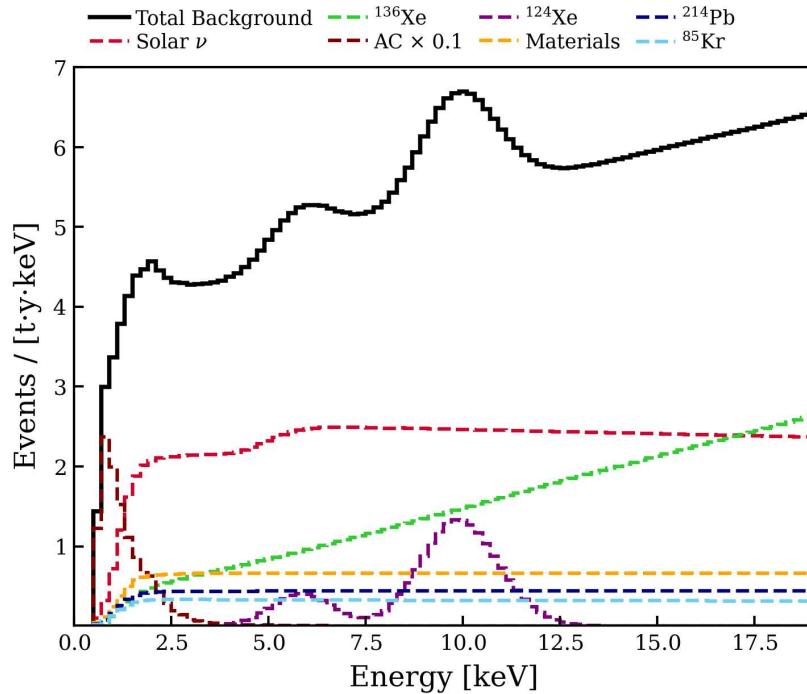


Figure 6.4: Projected differential event rates for the background components in the XLZD low-energy electronic recoil search using a 2-fold S1 coincidence threshold. The total background is shown as a solid black line. The accidental coincidence background component baseline projection has been scaled down by a factor of 10. Further details are provided in the text.

by $\sim 19\%$ with the 2-fold analysis relative to the 3-fold case, while the NMM median limit is improved by $\sim 6\%$. As anticipated, the sensitivity in both instances is visibly dependent on the AC background level. In the NMC case, the sensitivity increases by roughly a factor of 2.5 when the AC background factor is lowered from 2 to 0.1, while for NMM, the improvement is about 1.3 times under the same reduction. A comprehensive sensitivity analysis for the SRDM model introduced in section 6.1 is not included in this work. However, given the substantial rate gains observed and the signal model shape, the 2-fold analysis is expected to enhance sensitivity to this model by up to a factor of 2. This is a conservative estimate based on the other results obtained in this section.

In conclusion, the 2-fold coincidence analysis consistently demonstrates enhanced sensitivity across all considered BSM signal models, with improvements directly correlated to the steepness of their low-energy spectra. Given the minimal additional implementation overhead, this approach constitutes a worthwhile optimization for next-generation LXe observatories like XLZD.

6.5 Conclusion and discussion

XLZD will further improve the world-leading sensitivity of LXe observatories to BSM physics in low-energy ER searches. This chapter quantitatively evaluated the impact of reducing the S1 tight coincidence requirement from 3-fold to 2-fold, allowing access to a sub-keV energy threshold and increasing signal rates for several BSM models. Specifically, the 2-fold analysis yields signal rate increases ranging from 33% to 80% for neutrino electromagnetic properties and up to 514% for boosted DM scenarios. Projections of the AC background rates are derived based on the best available knowledge. The increase

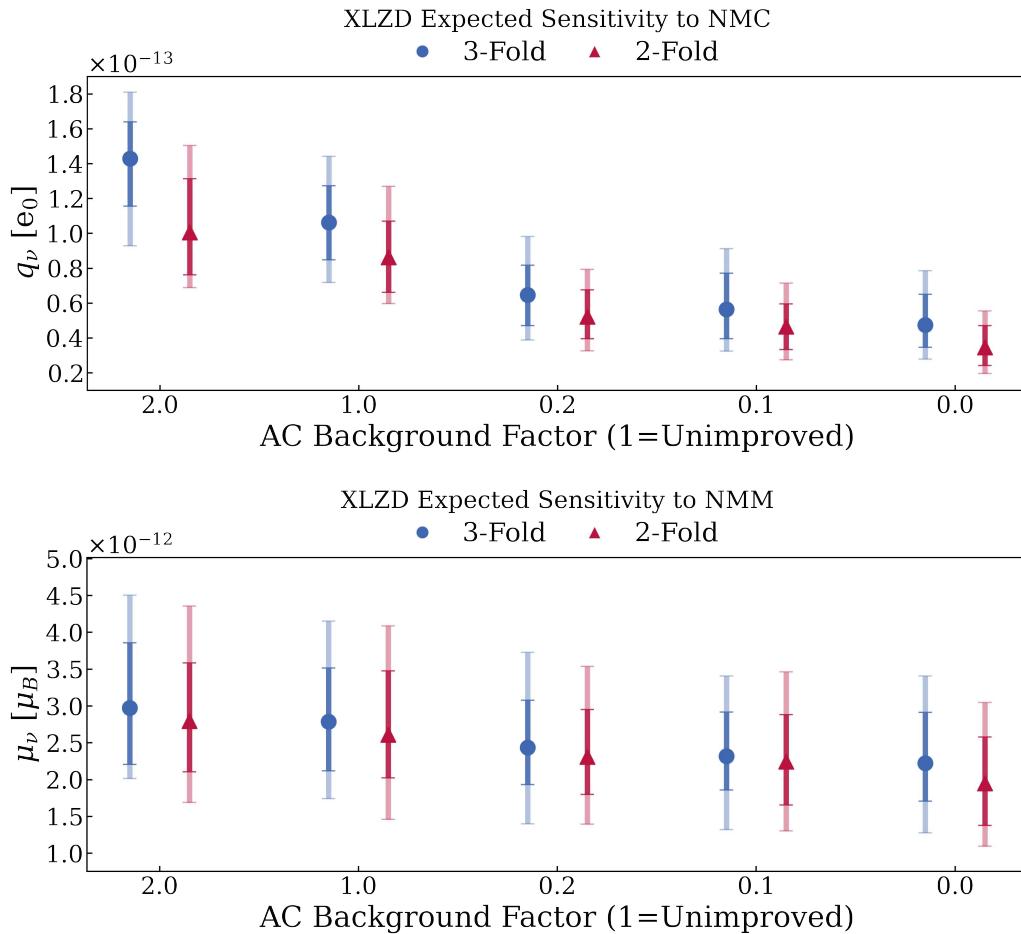


Figure 6.5: Projected 90% confidence level upper limits under the background-only hypothesis as a function of the accidental coincidence background scaling factor for an exposure of 200 tonne-years with XLZD. The red triangles show the median expected upper limits for 2-fold S1 coincidence analysis, and the blue circles are for the 3-fold one. The darker (fainter) error bars represent the 68% (95%) confidence intervals.

Top: Upper limits distributions for the neutrino millicharge (q_ν). **Bottom:** Upper limits distributions for the neutrino magnetic moment (μ_ν).

by a factor of ~ 14.9 for the 3-fold analysis and ~ 8.7 for the 2-fold analysis compared to the XENONnT SR0 cases is mainly due to detector scaling, specifically the larger number of PMTs and the associated rise in total dark count rate. The effects of the updated and improved selection are also taken into account. Sensitivity projections are presented for different AC scenarios, showing for the nominal case an improvement of 19% for the NMC model and 6% for the NMM model with the 2-fold analysis compared to the 3-fold one. Considering its low-energy rate enhancement, the 2-fold analysis is also expected to improve sensitivity to boosted SRDM scenarios by up to a factor of 2. These results demonstrate that the lower energy threshold 2-fold framework provides a practical and effective optimization for XLZD's low-energy ER BSM searches.

Summary and Conclusion

The search for physics beyond the Standard Model (BSM), including understanding the nature of dark matter (DM), investigating neutrino properties, and addressing the strong CP problem, is at the core of current particle physics research. Multi-tonne scale dual-phase liquid xenon time projection chambers, while maintaining their core mission of searching for weakly interacting massive particles (WIMPs), have emerged as premier multi-purpose observatories for rare event searches. This thesis focuses on electronic recoil (ER) events, which are generated when incoming particles interact with atomic electrons in the detector.

Chapter 2 described R&D advancements for the automated rare gas mass spectrometer (Auto-RGMS), a new automated system for weekly measurements of krypton contamination in xenon gas at a few parts-per-quadrillion (ppq). This is essential for controlling and characterizing the beta-decaying isotope ^{85}Kr , which constitutes a critical ER background source in low-energy analyses. Through a systematic evaluation of various gas chromatography adsorbents, HayeSep Q was identified as the best choice, providing a 12-fold improvement in the chromatographic resolution over the material used in the previous system. Automating the measurement workflow and refining calibration with a custom-built valve significantly reduce systematic uncertainties. The superior chromatographic separation achieved with the new adsorbent ensures that significantly larger xenon samples can be processed, and increasing the batch size directly enhances the detectable krypton amount, thereby improving the detection limit and enabling sensitivity at the few ppq level. These improvements are important for the next-generation XLZD science program [53], especially for measuring the solar pp neutrino flux with per-mille precision [63] and probing the weak mixing angle at a momentum transfer an order of magnitude lower than current state-of-the-art measurements [5].

Chapter 4 presented a comprehensive search for BSM physics in low-energy ER data using 1.16 tonne-years of data from XENONnT's first science run. The analysis benefited from an unprecedentedly low background rate of (15.8 ± 1.3) events/(tonne·year·keV) in the 1 to 30 keV range, a 5-fold reduction compared to its predecessor, XENON1T. Central to the analysis's success was the construction of a robust detector response model, precisely calibrated at low energies by the 2.82 keV K-shell line from an internal ^{37}Ar source. This calibration was key for accurately determining the energy scale, monitoring the spatial uniformity of the charge and light yield, and for determining and validating both the selection efficiency and the S1 detection efficiency. Following the unblinding of the data in the final analysis, an excellent agreement with the background model was observed, showing no evidence of new physics. This result was used to set world-leading laboratory constraints on a broad range of BSM models, including solar axions, bosonic DM, and new neutrino physics such as enhanced magnetic moment, neutrino millicharge, and non-standard neutrino interactions (vector and scalar mediators). The findings reject a BSM origin for the XENON1T low-energy excess [81], strengthening the tritium contamination hypothesis.

To push the sensitivity to even lower energies, Chapter 5 detailed a novel analysis framework that successfully achieved the first calibration of sub-keV ERs in liquid xenon. The detection of the 0.27 keV L-shell decay of ^{37}Ar with a statistical significance of 11.9σ

is reported. By relaxing the S1 signal coincidence requirement from three (3-fold) to two (2-fold) photomultiplier tubes detecting photons within a 100 ns window, the signal acceptance for L-shell decays was increased by a factor of about 66. Accurate characterization of the increased accidental coincidence background was obtained with a precise data-driven model, while its rate was effectively reduced with a dedicated selection strategy, including two boosted decision tree classifiers. The analysis yielded a measurement of the scintillation photon yield of $\text{PY} = 2.96 \pm 0.08 \text{ (stat.)} {}^{+0.57}_{-0.33} \text{ (syst.) } \gamma/\text{keV}$ (68% C.L.) at 0.27 keV. Beyond validating the analysis framework, this measurement provides important input for the NEST semi-empirical model, enabling a more accurate ER emission model in the uncalibrated sub-keV energy range. This improved understanding is essential not only when searching for BSM physics within sub-keV ER data, but also for accurately modeling ER events when they constitute a background component in low-energy nuclear recoil analyses. For instance, in the recent XENONnT search for light WIMPs [163] and in the measurement of solar ${}^8\text{B}$ neutrinos via coherent elastic neutrino-nucleus scattering (CE ν NS) [48], a conservative 100% systematic uncertainty on the ER background rate was required due to this lack of knowledge.

Finally, Chapter 6 projected the impact of the newly developed sub-keV ER analysis framework on the BSM physics sensitivity of the future XLZD observatory. By modeling signal and background expectations for both 2-fold and 3-fold S1 coincidence analyses, which incorporate an updated event selection strategy to optimize background rejection, it was demonstrated that the lower energy threshold enhances sensitivity for models with rising low-energy spectra. For a nominal 200 tonne-year exposure, the 2-fold analysis, compared to the standard 3-fold case, improves projected sensitivity to neutrino electromagnetic properties by up to 19% and suggests an improvement of up to a factor of two for boosted DM. Consequently, leveraging this refined analysis strategy is beneficial for maximizing the scientific potential and discovery reach of future-generation ER data.

Appendix

Appendix A

Author's publications used in this work

Many of the results reported in this thesis are based on XENON Collaboration data and would not have been possible without the dedicated work and input of numerous collaborators, for which I am truly thankful. The following list outlines my specific contributions for each chapter and the associated publications.

- **Chapter 2:** This chapter is based on: Guida, M., Lin, YT. & Simgen, H. *Improved and automated krypton assay for low-background xenon detectors with Auto-RGMS*. *Eur. Phys. J. C* **85**, 576 (2025). I carried out all the work, except for developing and measuring the custom valve and analyzing the adsorbents' memory effect.
- **Chapter 4:** I performed the analysis of the 2.82 keV ^{37}Ar K-shell calibration, including voxel analysis for charge and light yield determination, spatial stability of yields, validation of the S1 detection efficiency and cut acceptance. I contributed to the determination of the g1-g2 energy scale, energy resolution, and solar neutrino background modeling. I implemented neutrino non-standard interactions (vector and scalar) and neutrino millicharge signal models, and carried out inference and derived upper limits for these signals. Parts of these results contributed to and appear in the following publications: Aprile, E. *et al.* (XENON Collaboration). *Search for New Physics in Electronic Recoil Data from XENONnT*. *Phys. Rev. Lett.* **129**, 161805 (2022) and Aprile, E. *et al.* (XENON Collaboration). *XENONnT analysis: Signal reconstruction, calibration, and event selection*. *Phys. Rev. D* **111**, 062006 (2025).
- **Chapter 5:** I conceived the idea and developed all aspects of the analysis related to the 0.270 keV ^{37}Ar L-shell line, with the exception of the accidental coincidence background modeling. A XENON Collaboration paper on these results is currently in preparation, and I am part of the writing team.
- **Chapter 6:** All results presented in this chapter were obtained by me and have not been published elsewhere.

Appendix B

Acknowledgements

I'm really glad to have the chance to write a few lines through which I can express my gratitude. I'm truly happy for this PhD journey. It's been an adventure and such a great experience. If I had the chance to start over, I'd make the same choice without a second thought.

First, a huge thanks to Teresa for giving me the space to explore my ideas and learn from my mistakes without too much pressure. That freedom has been one of the best parts of my PhD, and I really appreciate how you were always there with support and good advice. A big thank you to Hardy: it's been a pleasure working with you. I really appreciate your honesty, straightforwardness, and the way you always had something helpful to add whenever I reached out. I'd like to thank Manfred for the opportunity to join the MPIK XENON group, for creating and sustaining a stimulating workplace culture, and for having supported my initiatives. Many thanks to Stephie for your support throughout my master's thesis, during my transition to the PhD program, and for kindly agreeing to be the second referee for this work. I am grateful to Britta and Anja for their consistent kindness and support toward everyone in the workplace and for playing a central role in creating many wonderful events. I would also like to thank Steffen and Jonas for always being approachable, patient, and ready to help. My thanks also go to the dedicated scientists of the XENON collaboration for allowing me to experience what great science and research look like. I am especially grateful for the way you shared your insights and time whenever I reached out, which was fundamental for my PhD.

I honestly can't imagine what my PhD would have been like without all the amazing friends, PhD students and postdocs, who've been part of our MPIK division during these years: you've truly been the heart of my PhD experience. A heartfelt thank you to (names are in no particular order) Robi, Lui, Andrii, Dominick, Gio, Flo, Quim, Kai, Sophie, Tim, Ting, Mona, Veronica, and everyone else in the division I've met, especially Andreas, Sudip, Ting, Horay, and Andreas. Being able to discuss ideas with you, exchange thoughts, and count on your availability was essential at every stage of this work. I'm even more grateful for all the moments we shared outside of work, which for me meant a lot and were more than "social life/activities". Thanks to those among you who opened up parts of their lives to me, for the many evenings we spent in Heidelberg (often at Schwarzer Peter), for the days in Neckarwiese, for days at the lake, for the trips we organized near Gran Sasso, for visiting my family and Moneglia, for the joy of singing a serenade to Cristina with other Italian friends the night before our wedding, and for sharing my wedding day, and so many other occasions. All of this meant a lot to me.

Broadening the circle beyond the division, there are many more people to thank,

so many friends from Germany, Italy, and scattered around the world to whom I am deeply grateful. Rather than listing all of you individually, I'll allow myself to group these thanks as follows: you all know that anyone who was invited to our wedding is someone I hold dear, and who, to varying degrees, has contributed to shaping who I am and making my life beautiful. I'd also like to include those who know that our friendship has grown and deepened over the past year, even if they weren't part of the original wedding crew.

A huge thank you to my parents, Silvia and Robi, and my sisters, La and Mich. Your love, your personalities, and your humanity have shaped so deeply who I am today, and to you I owe everything. Thank you for continuing to support and care for me throughout our long-distance relationship over the last four and a half years. Thank you for never making me feel guilty for wanting to explore the world, even when it meant being far away. Thank you also for giving me the quiet certainty that life is an extraordinary and meaningful adventure.

The best arrives at the end: you, Cris. In our relationship, I keep (re)discovering myself. You are the window that doesn't just show the horizon, but invites me to walk toward it barefoot and unafraid. Your alliance, your company, your understanding, your generosity, your support and above all, you yourself, are the most beautiful and precious gift I have ever received.

Truly yours,
Matte

Appendix C

Bibliography

- [1] E. Aprile and T. Doke. "Liquid Xenon Detectors for Particle Physics and Astrophysics". In: *Rev. Mod. Phys.* 82 (2010), pp. 2053–2097. DOI: [10.1103/RevModPhys.82.2053](https://doi.org/10.1103/RevModPhys.82.2053). arXiv: [0910.4956 \[physics.ins-det\]](https://arxiv.org/abs/0910.4956).
- [2] David R. Nygren. "Proposal to investigate the feasibility of a novel concept in particle detection". In: (1974).
- [3] C. Rubbia. "The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors". In: (May 1977).
- [4] J. Angle et al. "First Results from the XENON10 Dark Matter Experiment at the Gran Sasso National Laboratory". In: *Phys. Rev. Lett.* 100 (2008), p. 021303. DOI: [10.1103/PhysRevLett.100.021303](https://doi.org/10.1103/PhysRevLett.100.021303). arXiv: [0706.0039 \[astro-ph\]](https://arxiv.org/abs/0706.0039).
- [5] S. Navas et al. "Review of Particle Physics". In: *Phys. Rev. D* 110 (3 Aug. 2024), p. 030001. DOI: [10.1103/PhysRevD.110.030001](https://doi.org/10.1103/PhysRevD.110.030001). URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.030001>.
- [6] E. Aprile et al. "Design and Performance of the XENON10 Dark Matter Experiment". In: *Astropart. Phys.* 34 (2011), pp. 679–698. DOI: [10.1016/j.astropartphys.2011.01.006](https://doi.org/10.1016/j.astropartphys.2011.01.006). arXiv: [1001.2834 \[astro-ph.IM\]](https://arxiv.org/abs/1001.2834).
- [7] E. Aprile et al. "The XENON100 Dark Matter Experiment". In: *Astropart. Phys.* 35 (2012), pp. 573–590. DOI: [10.1016/j.astropartphys.2012.01.003](https://doi.org/10.1016/j.astropartphys.2012.01.003). arXiv: [1107.2155 \[astro-ph.IM\]](https://arxiv.org/abs/1107.2155).
- [8] E. Aprile et al. "Dark Matter Results from 225 Live Days of XENON100 Data". In: *Phys. Rev. Lett.* 109 (2012), p. 181301. DOI: [10.1103/PhysRevLett.109.181301](https://doi.org/10.1103/PhysRevLett.109.181301). arXiv: [1207.5988 \[astro-ph.CO\]](https://arxiv.org/abs/1207.5988).
- [9] E. Aprile et al. "The XENON1T Dark Matter Experiment". In: *Eur. Phys. J. C* 77.12 (2017), p. 881. DOI: [10.1140/epjc/s10052-017-5326-3](https://doi.org/10.1140/epjc/s10052-017-5326-3). arXiv: [1708.07051 \[astro-ph.IM\]](https://arxiv.org/abs/1708.07051).
- [10] E. Aprile et al. "Dark Matter Search Results from a One Ton-Year Exposure of XENON1T". In: *Phys. Rev. Lett.* 121.11 (2018), p. 111302. DOI: [10.1103/PhysRevLett.121.111302](https://doi.org/10.1103/PhysRevLett.121.111302). arXiv: [1805.12562 \[astro-ph.CO\]](https://arxiv.org/abs/1805.12562).
- [11] E. Aprile et al. "The XENONnT dark matter experiment". In: *Eur. Phys. J. C* 84.8 (2024), p. 784. DOI: [10.1140/epjc/s10052-024-12982-5](https://doi.org/10.1140/epjc/s10052-024-12982-5). arXiv: [2402.10446 \[physics.ins-det\]](https://arxiv.org/abs/2402.10446).
- [12] E. Aprile et al. "WIMP Dark Matter Search using a 3.1 tonne × year Exposure of the XENONnT Experiment". In: (Feb. 2025). arXiv: [2502.18005 \[hep-ex\]](https://arxiv.org/abs/2502.18005).
- [13] Mohammad Sajjad Athar et al. "Status and perspectives of neutrino physics". In: *Prog. Part. Nucl. Phys.* 124 (2022), p. 103947. DOI: [10.1016/j.ppnp.2022.103947](https://doi.org/10.1016/j.ppnp.2022.103947). arXiv: [2111.07586 \[hep-ph\]](https://arxiv.org/abs/2111.07586).

- [14] Laurent Canetti, Marco Drewes, and Mikhail Shaposhnikov. "Matter and Anti-matter in the Universe". In: *New J. Phys.* 14 (2012), p. 095012. DOI: [10.1088/1367-2630/14/9/095012](https://doi.org/10.1088/1367-2630/14/9/095012). arXiv: [1204.4186 \[hep-ph\]](https://arxiv.org/abs/1204.4186).
- [15] Stephen P. Martin. "A Supersymmetry primer". In: *Adv. Ser. Direct. High Energy Phys.* 18 (1998). Ed. by Gordon L. Kane, pp. 1–98. DOI: [10.1142/9789812839657_0001](https://doi.org/10.1142/9789812839657_0001). arXiv: [hep-ph/9709356](https://arxiv.org/abs/hep-ph/9709356).
- [16] Jihn E. Kim and Gianpaolo Carosi. "Axions and the Strong CP Problem". In: *Rev. Mod. Phys.* 82 (2010). [Erratum: *Rev.Mod.Phys.* 91, 049902 (2019)], pp. 557–602. DOI: [10.1103/RevModPhys.82.557](https://doi.org/10.1103/RevModPhys.82.557). arXiv: [0807.3125 \[hep-ph\]](https://arxiv.org/abs/0807.3125).
- [17] Salvatore Capozziello and Mariafelicia De Laurentis. "Extended Theories of Gravity". In: *Phys. Rept.* 509 (2011), pp. 167–321. DOI: [10.1016/j.physrep.2011.09.003](https://doi.org/10.1016/j.physrep.2011.09.003). arXiv: [1108.6266 \[gr-qc\]](https://arxiv.org/abs/1108.6266).
- [18] Marco Cirelli, Alessandro Strumia, and Jure Zupan. "Dark Matter". In: (June 2024). arXiv: [2406.01705 \[hep-ph\]](https://arxiv.org/abs/2406.01705).
- [19] Joshua Frieman, Michael Turner, and Dragan Huterer. "Dark Energy and the Accelerating Universe". In: *Ann. Rev. Astron. Astrophys.* 46 (2008), pp. 385–432. DOI: [10.1146/annurev.astro.46.060407.145243](https://doi.org/10.1146/annurev.astro.46.060407.145243). arXiv: [0803.0982 \[astro-ph\]](https://arxiv.org/abs/0803.0982).
- [20] Marc Schumann. "Dark Matter Search with liquid Noble Gases". In: (June 2012). arXiv: [1206.2169 \[astro-ph.IM\]](https://arxiv.org/abs/1206.2169).
- [21] Carl Eric Dahl. "The physics of background discrimination in liquid xenon, and first results from Xenon10 in the hunt for WIMP dark matter". PhD thesis. Princeton U., 2009.
- [22] Tadayoshi Doke et al. "Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles". In: *Jap. J. Appl. Phys.* 41 (2002), pp. 1538–1545. DOI: [10.1143/JJAP.41.1538](https://doi.org/10.1143/JJAP.41.1538).
- [23] Keiko Fujii et al. "High-accuracy measurement of the emission spectrum of liquid xenon in the vacuum ultraviolet region". In: *Nucl. Instrum. Meth. A* 795 (2015), pp. 293–297. DOI: [10.1016/j.nima.2015.05.065](https://doi.org/10.1016/j.nima.2015.05.065).
- [24] H. O. Back et al. "First Large Scale Production of Low Radioactivity Argon From Underground Sources". In: (Apr. 2012). arXiv: [1204.6024 \[astro-ph.IM\]](https://arxiv.org/abs/1204.6024).
- [25] J. B. Albert et al. "Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector". In: *Phys. Rev. C* 89.1 (2014), p. 015502. DOI: [10.1103/PhysRevC.89.015502](https://doi.org/10.1103/PhysRevC.89.015502). arXiv: [1306.6106 \[nucl-ex\]](https://arxiv.org/abs/1306.6106).
- [26] E. Aprile et al. "Double-Weak Decays of ^{124}Xe and ^{136}Xe in the XENON1T and XENONnT Experiments". In: *Phys. Rev. C* 106.2 (2022), p. 024328. DOI: [10.1103/PhysRevC.106.024328](https://doi.org/10.1103/PhysRevC.106.024328). arXiv: [2205.04158 \[hep-ex\]](https://arxiv.org/abs/2205.04158).
- [27] Henrique Araujo and Laura Baudis. *Xenon gas for XLZD*. Slide presentation at the XLZD Collaboration Meeting. Restricted access material, available only to XLZD Collaboration members. LNGS, Italy, June 2025. URL: https://internal-collaboration-site.example.com/slides/xenon_gas_for_xlzd_20250630.pdf (visited on 03/08/2024).
- [28] R.L. Platzman. "Total ionization in gases by high-energy particles: An appraisal of our understanding". In: *The International Journal of Applied Radiation and Isotopes* 10.2 (1961), pp. 116–127. ISSN: 0020-708X. DOI: [https://doi.org/10.1016/0020-708X\(61\)90108-9](https://doi.org/10.1016/0020-708X(61)90108-9). URL: <https://www.sciencedirect.com/science/article/pii/0020708X61901089>.

[29] Ulrich Sowada, John M. Warman, and Matthijs P. de Haas. “Hot-electron thermalization in solid and liquid argon, krypton, and xenon”. In: *Phys. Rev. B* 25 (5 Mar. 1982), pp. 3434–3437. DOI: [10.1103/PhysRevB.25.3434](https://doi.org/10.1103/PhysRevB.25.3434). URL: <https://link.aps.org/doi/10.1103/PhysRevB.25.3434>.

[30] Peter Sorensen and Carl Eric Dahl. “Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter”. In: *Phys. Rev. D* 83 (6 Mar. 2011), p. 063501. DOI: [10.1103/PhysRevD.83.063501](https://doi.org/10.1103/PhysRevD.83.063501). URL: <https://link.aps.org/doi/10.1103/PhysRevD.83.063501>.

[31] D. S. Akerib et al. “Liquid xenon scintillation measurements and pulse shape discrimination in the LUX dark matter detector”. In: *Phys. Rev. D* 97.11 (2018). DOI: [10.1103/PhysRevD.97.112002](https://doi.org/10.1103/PhysRevD.97.112002). arXiv: [1802.06162 \[physics.ins-det\]](https://arxiv.org/abs/1802.06162).

[32] M. Szydagis et al. “A Review of NEST Models for Liquid Xenon and Exhaustive Comparison to Other Approaches”. In: (Nov. 2022). DOI: [10.3389/fdest.2024.1480975](https://doi.org/10.3389/fdest.2024.1480975). arXiv: [2211.10726 \[hep-ex\]](https://arxiv.org/abs/2211.10726).

[33] Sophia Farrell et al. *NESTCollaboration/nestpy: Update to v2.0.2, sync with NEST v2.4.0 (v2.0.2)*. 2024. DOI: [10.5281/zenodo.10582363](https://doi.org/10.5281/zenodo.10582363). URL: <https://doi.org/10.5281/zenodo.10582363>.

[34] E. Aprile et al. “XENONnT analysis: Signal reconstruction, calibration, and event selection”. In: *Phys. Rev. D* 111.6 (2025), p. 062006. DOI: [10.1103/PhysRevD.111.062006](https://doi.org/10.1103/PhysRevD.111.062006). arXiv: [2409.08778 \[hep-ex\]](https://arxiv.org/abs/2409.08778).

[35] E. Aprile et al. “Simultaneous measurement of ionization and scintillation from nuclear recoils in liquid xenon as target for a dark matter experiment”. In: *Phys. Rev. Lett.* 97 (2006), p. 081302. DOI: [10.1103/PhysRevLett.97.081302](https://doi.org/10.1103/PhysRevLett.97.081302). arXiv: [astro-ph/0601552](https://arxiv.org/abs/astro-ph/0601552).

[36] E. Aprile et al. “The neutron veto of the XENONnT experiment: results with demineralized water”. In: *Eur. Phys. J. C* 85.6 (2025), p. 695. DOI: [10.1140/epjc/s10052-025-14105-0](https://doi.org/10.1140/epjc/s10052-025-14105-0). arXiv: [2412.05264 \[physics.ins-det\]](https://arxiv.org/abs/2412.05264).

[37] E. Aprile et al. “Conceptual design and simulation of a water Cherenkov muon veto for the XENON1T experiment”. In: *JINST* 9 (2014), P11006. DOI: [10.1088/1748-0221/9/11/P11006](https://doi.org/10.1088/1748-0221/9/11/P11006). arXiv: [1406.2374 \[astro-ph.IM\]](https://arxiv.org/abs/1406.2374).

[38] G. Plante et al. “Liquid-phase purification for multi-tonne xenon detectors”. In: *Eur. Phys. J. C* 82.10 (2022), p. 860. DOI: [10.1140/epjc/s10052-022-10832-w](https://doi.org/10.1140/epjc/s10052-022-10832-w). arXiv: [2205.07336 \[physics.ins-det\]](https://arxiv.org/abs/2205.07336).

[39] M. Murra et al. “Design, construction and commissioning of a high-flow radon removal system for XENONnT”. In: *Eur. Phys. J. C* 82.12 (2022), p. 1104. DOI: [10.1140/epjc/s10052-022-11001-9](https://doi.org/10.1140/epjc/s10052-022-11001-9). arXiv: [2205.11492 \[physics.ins-det\]](https://arxiv.org/abs/2205.11492).

[40] E. Aprile et al. “Removing krypton from xenon by cryogenic distillation to the ppq level”. In: *Eur. Phys. J. C* 77.5 (2017), p. 275. DOI: [10.1140/epjc/s10052-017-4757-1](https://doi.org/10.1140/epjc/s10052-017-4757-1). arXiv: [1612.04284 \[physics.ins-det\]](https://arxiv.org/abs/1612.04284).

[41] E. Aprile et al. “Design and performance of the field cage for the XENONnT experiment”. In: *Eur. Phys. J. C* 84.2 (2024), p. 138. DOI: [10.1140/epjc/s10052-023-12296-y](https://doi.org/10.1140/epjc/s10052-023-12296-y). arXiv: [2309.11996 \[hep-ex\]](https://arxiv.org/abs/2309.11996).

[42] E. Aprile et al. “Projected WIMP sensitivity of the XENONnT dark matter experiment”. In: *JCAP* 11 (2020), p. 031. DOI: [10.1088/1475-7516/2020/11/031](https://doi.org/10.1088/1475-7516/2020/11/031). arXiv: [2007.08796 \[physics.ins-det\]](https://arxiv.org/abs/2007.08796).

- [43] V. C. Antochi et al. "Improved quality tests of R11410-21 photomultiplier tubes for the XENONnT experiment". In: *JINST* 16.08 (2021), P08033. DOI: [10.1088/1748-0221/16/08/P08033](https://doi.org/10.1088/1748-0221/16/08/P08033). arXiv: [2104.15051](https://arxiv.org/abs/2104.15051) [physics.ins-det].
- [44] E. Aprile et al. "Material radiopurity control in the XENONnT experiment". In: *Eur. Phys. J. C* 82.7 (2022), p. 599. DOI: [10.1140/epjc/s10052-022-10345-6](https://doi.org/10.1140/epjc/s10052-022-10345-6). arXiv: [2112.05629](https://arxiv.org/abs/2112.05629) [physics.ins-det].
- [45] D. Akimov et al. "Measurement of the Coherent Elastic Neutrino-Nucleus Scattering Cross Section on CsI by COHERENT". In: *Phys. Rev. Lett.* 129.8 (2022), p. 081801. DOI: [10.1103/PhysRevLett.129.081801](https://doi.org/10.1103/PhysRevLett.129.081801). arXiv: [2110.07730](https://arxiv.org/abs/2110.07730) [hep-ex].
- [46] D. Akimov et al. "First Measurement of Coherent Elastic Neutrino-Nucleus Scattering on Argon". In: *Phys. Rev. Lett.* 126.1 (2021), p. 012002. DOI: [10.1103/PhysRevLett.126.012002](https://doi.org/10.1103/PhysRevLett.126.012002). arXiv: [2003.10630](https://arxiv.org/abs/2003.10630) [nucl-ex].
- [47] S. Adamski et al. "Evidence of Coherent Elastic Neutrino-Nucleus Scattering with COHERENT's Germanium Array". In: *Phys. Rev. Lett.* 134.23 (2025), p. 231801. DOI: [10.1103/PhysRevLett.134.231801](https://doi.org/10.1103/PhysRevLett.134.231801).
- [48] Elena Aprile et al. "First Indication of Solar B8 Neutrinos via Coherent Elastic Neutrino-Nucleus Scattering with XENONnT". In: *Phys. Rev. Lett.* 133.19 (2024), p. 191002. DOI: [10.1103/PhysRevLett.133.191002](https://doi.org/10.1103/PhysRevLett.133.191002). arXiv: [2408.02877](https://arxiv.org/abs/2408.02877) [nucl-ex].
- [49] E. Aprile et al. "First Dark Matter Search with Nuclear Recoils from the XENONnT Experiment". In: *Phys. Rev. Lett.* 131.4 (2023), p. 041003. DOI: [10.1103/PhysRevLett.131.041003](https://doi.org/10.1103/PhysRevLett.131.041003). arXiv: [2303.14729](https://arxiv.org/abs/2303.14729) [hep-ex].
- [50] J. Aalbers et al. "Dark Matter Search Results from 4.2 Tonne-Years of Exposure of the LUX-ZEPLIN (LZ) Experiment". In: *Phys. Rev. Lett.* 135.1 (2025), p. 011802. DOI: [10.1103/4dyc-z8zf](https://doi.org/10.1103/4dyc-z8zf). arXiv: [2410.17036](https://arxiv.org/abs/2410.17036) [hep-ex].
- [51] Zihao Bo et al. "Dark Matter Search Results from 1.54 Tonne · Year Exposure of PandaX-4T". In: *Phys. Rev. Lett.* 134 (1 Jan. 2025), p. 011805. DOI: [10.1103/PhysRevLett.134.011805](https://doi.org/10.1103/PhysRevLett.134.011805). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.134.011805>.
- [52] J. Aalbers et al. "A next-generation liquid xenon observatory for dark matter and neutrino physics". In: *J. Phys. G* 50.1 (2023), p. 013001. DOI: [10.1088/1361-6471/ac841a](https://doi.org/10.1088/1361-6471/ac841a). arXiv: [2203.02309](https://arxiv.org/abs/2203.02309) [physics.ins-det].
- [53] J. Aalbers et al. "The XLZD Design Book: Towards the Next-Generation Liquid Xenon Observatory for Dark Matter and Neutrino Physics". In: (Oct. 2024). arXiv: [2410.17137](https://arxiv.org/abs/2410.17137) [hep-ex].
- [54] M. S. Safronova et al. "Search for new physics with atoms and molecules". In: *Rev. Mod. Phys.* 90 (2 June 2018), p. 025008. DOI: [10.1103/RevModPhys.90.025008](https://doi.org/10.1103/RevModPhys.90.025008).
- [55] Andreas Bollhöfer et al. "85Kr Activity Determination in Tropospheric Air". In: *International Foundation HFSJG Activity Report* (2020) 55–56.
- [56] H. Sievers. "Nuclear data sheets update for A = 85". In: *Nuclear Data Sheets* 62 (1991), pp. 271–325. ISSN: 0090-3752. DOI: [https://doi.org/10.1016/0090-3752\(91\)80016-Y](https://doi.org/10.1016/0090-3752(91)80016-Y). URL: <https://www.sciencedirect.com/science/article/pii/009037529180016Y>.
- [57] Chao Gao et al. "Monitoring atmospheric 85Kr by atom counting". In: *Journal of Environmental Radioactivity* 233 (July 2021), p. 106604. DOI: [10.1016/j.jenvrad.2021.106604](https://doi.org/10.1016/j.jenvrad.2021.106604).

[58] E. Aprile et al. "An atom trap trace analysis system for measuring krypton contamination in xenon dark matter detectors". In: *Review of Scientific Instruments* 84.9 (Sept. 2013), p. 093105. ISSN: 0034-6748. DOI: [10.1063/1.4821879](https://doi.org/10.1063/1.4821879). eprint: https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/1.4821879/16100011/093105_1_online.pdf. URL: <https://doi.org/10.1063/1.4821879>.

[59] E. Aprile et al. "Atom trap trace analysis of krypton in xenon for the XENON Dark Matter Project". In: *JINST* 13.10 (2018), P10018. DOI: [10.1088/1748-0221/13/10/P10018](https://doi.org/10.1088/1748-0221/13/10/P10018).

[60] Sebastian Lindemann and Hardy Simgen. "Krypton assay in xenon at the ppq level using a gas chromatographic system and mass spectrometer". In: *Eur. Phys. J. C* 74 (2014), p. 2746. DOI: [10.1140/epjc/s10052-014-2746-1](https://doi.org/10.1140/epjc/s10052-014-2746-1). arXiv: [1308.4806 \[physics.ins-det\]](https://arxiv.org/abs/1308.4806).

[61] D. R. Lide. *CRC Handbook of Chemistry and Physics*. 87th. CRC Press, 2006.

[62] J. Aalbers et al. "DARWIN: towards the ultimate dark matter detector". In: *JCAP* 11 (2016), p. 017. DOI: [10.1088/1475-7516/2016/11/017](https://doi.org/10.1088/1475-7516/2016/11/017). arXiv: [1606.07001 \[astro-ph.IM\]](https://arxiv.org/abs/1606.07001).

[63] J. Aalbers et al. "Solar neutrino detection sensitivity in DARWIN via electron scattering". In: *Eur. Phys. J. C* 80.12 (2020), p. 1133. DOI: [10.1140/epjc/s10052-020-08602-7](https://doi.org/10.1140/epjc/s10052-020-08602-7). arXiv: [2006.03114 \[physics.ins-det\]](https://arxiv.org/abs/2006.03114).

[64] E. Aprile et al. "Search for New Physics in Electronic Recoil Data from XENONnT". In: *Phys. Rev. Lett.* 129.16 (2022), p. 161805. DOI: [10.1103/PhysRevLett.129.161805](https://doi.org/10.1103/PhysRevLett.129.161805). arXiv: [2207.11330 \[hep-ex\]](https://arxiv.org/abs/2207.11330).

[65] D. Winkler. "Determination of the Krypton-Impurity Evolution in the XENON1T Detector and Optimization of Separation Columns for Krypton in Xenon Assays". MA thesis. Heidelberg University, 2020. DOI: <https://hdl.handle.net/21.11116/0000-0007-CFA8-3>.

[66] Hardy Simgen. *SR0 Kr-in-Xe measurements with RGMS*. XENONnT internal note. Access restricted. Accessed 02.08.2025. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenonnt:dsg:clean:analytics:rgms:rgms_sr0.

[67] Matteo Guida, Ying-Ting Lin, and Hardy Simgen. "Improved and automated krypton assay for low-background xenon detectors with Auto-RGMS". In: *Eur. Phys. J. C* 85.5 (2025), p. 576. DOI: [10.1140/epjc/s10052-025-14262-2](https://doi.org/10.1140/epjc/s10052-025-14262-2). arXiv: [2501.10993 \[physics.ins-det\]](https://arxiv.org/abs/2501.10993).

[68] Robert Hammann. "Investigation of a Charge Insensitive Volume in XENONnT, Analysis of Goodness-of-Fit Techniques, and Feasibility Studies for an Automated Krypton Assay System". MA thesis. Heidelberg University, 2022. DOI: [10.5281/zenodo.13946149](https://zenodo.13946149).

[69] K. Dettmer-Wilde and W. Engewald. "Practical Gas Chromatography: A Comprehensive Reference". In: Springer Berlin Heidelberg, 2014. DOI: [10.1007/978-3-642-54640-2](https://doi.org/10.1007/978-3-642-54640-2).

[70] T.J. Bruno and P.D.N. Svoronos. *CRC Handbook of Basic Tables for Chemical Analysis: Data-Driven Methods and Interpretation*. CRC Press, 2020. ISBN: 9781351613156. URL: <https://books.google.de/books?id=LGPxDwAAQBAJ>.

[71] Analytics Shop. *Packed GC Columns*. <https://www.analytics-shop.com/gb/chromatography/gc-columns/packed-gc-columns>. (Visited on 09/25/2024).

[72] Torgny Fornstedt, Guoming Zhong, and Georges Guiochon. "Peak tailing and slow mass transfer kinetics in nonlinear chromatography". In: *Journal of Chromatography A* 742.1 (1996), pp. 55–68. ISSN: 0021-9673. DOI: [https://doi.org/10.1016/0021-9673\(96\)00323-8](https://doi.org/10.1016/0021-9673(96)00323-8).

[73] R.L. Grob and E.F. Barry. "Modern Practice of Gas Chromatography". In: Wiley, 2004. DOI: <10.1002/0471651141>.

[74] R. Weiser. "Aufbau einer Ultrahochvakuum-Probenaufschlußapparatur für die Massenspektrometrische Analyse von Edelgasen". MA thesis. Heidelberg University, 1993.

[75] E. Aprile et al. "The triggerless data acquisition system of the XENONnT experiment". In: *JINST* 18.07 (2023), P07054. DOI: <10.1088/1748-0221/18/07/P07054>. arXiv: [2212.11032 \[physics.ins-det\]](2212.11032).

[76] Daniel Wenz. "Commissioning of the world's first water Cherenkov neutron veto and first WIMP dark matter search results of the XENONnT experiment". PhD thesis. Mainz U., Mainz U., 2023. DOI: <10.25358/openscience-9654>.

[77] Joran Remco Angevaare. "First WIMP results of XENONnT and its signal reconstruction". PhD thesis. Amsterdam U., 2023.

[78] Brookhaven National Laboratory, National Nuclear Data Center. *National Nuclear Data Center*. Accessed 31.07.2025. URL: <http://www.nndc.bnl.gov>.

[79] Maxime Pierre. "Neutrinoless double beta decay search with XENONnT". PhD thesis. Ecole nationale supérieure Mines-Télécom Atlantique, Dec. 2022. URL: <https://theses.hal.science/tel-04216774>.

[80] Florian Jörg et al. "Characterization of alpha and beta interactions in liquid xenon". In: *Eur. Phys. J. C* 82.4 (2022), p. 361. DOI: <10.1140/epjc/s10052-022-10259-3>. arXiv: [2109.13735 \[physics.ins-det\]](2109.13735).

[81] E. Aprile et al. "Excess electronic recoil events in XENON1T". In: *Phys. Rev. D* 102.7 (2020), p. 072004. DOI: <10.1103/PhysRevD.102.072004>. arXiv: [2006.09721 \[hep-ex\]](2006.09721).

[82] E. Aprile et al. "Low-energy calibration of XENON1T with an internal ^{37}Ar source". In: *Eur. Phys. J. C* 83.6 (2023), p. 542. DOI: <10.1140/epjc/s10052-023-11512-z>. arXiv: [2211.14191 \[physics.ins-det\]](2211.14191).

[83] M.-M. et al. Be. *Table of Radionuclides*. Vol. 5, volume 7. Monographie BIPM. Pavillon de Breteuil, Sèvres, France: Bureau International des Poids et Mesures, 2013.

[84] V. P. Chechev. "The evaluation of half-lives and other decay data used in nuclear astrophysics and cosmochronology". In: *Phys. Atom. Nuclei* 74.12 (Dec. 2011), pp. 1713–1717. DOI: <10.1134/S106377881111007X>.

[85] G. Anton et al. "Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment". In: *Phys. Rev. C* 101.6 (2020), p. 065501. DOI: <10.1103/PhysRevC.101.065501>. arXiv: [1908.04128 \[physics.ins-det\]](1908.04128).

[86] Laura Baudis, Patricia Sanchez-Lucas, and Kevin Thieme. "A measurement of the mean electronic excitation energy of liquid xenon". In: *Eur. Phys. J. C* 81.12 (2021), p. 1060. DOI: <10.1140/epjc/s10052-021-09834-x>. arXiv: [2109.07151 \[physics.ins-det\]](2109.07151).

[87] V. N. Lebedenko et al. “Results from the first science run of the ZEPLIN-III dark matter search experiment”. In: *Phys. Rev. D* 80 (5 Sept. 2009), p. 052010. DOI: [10.1103/PhysRevD.80.052010](https://doi.org/10.1103/PhysRevD.80.052010). URL: <https://link.aps.org/doi/10.1103/PhysRevD.80.052010>.

[88] D. S. Akerib et al. “Discrimination of electronic recoils from nuclear recoils in two-phase xenon time projection chambers”. In: *Phys. Rev. D* 102 (11 Sept. 2020), p. 112002. DOI: [10.1103/PhysRevD.102.112002](https://doi.org/10.1103/PhysRevD.102.112002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.112002>.

[89] 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”. In: *Instruments* 5.1 (2021), p. 13. DOI: [10.3390/instruments5010013](https://doi.org/10.3390/instruments5010013). arXiv: [2102.10209 \[hep-ex\]](https://arxiv.org/abs/2102.10209).

[90] L. Lista. *Statistical Methods for Data Analysis: With Applications in Particle Physics*. Lecture Notes in Physics. Springer, 2023. ISBN: 9783031199349. URL: <https://books.google.de/books?id=PA28EAAAQBAJ>.

[91] Evan Shockley. “Study of Excess Electronic Recoil Events in XENON1T”. PhD thesis. Chicago U., 2021. DOI: [10.6082/uchicago.2743](https://doi.org/10.6082/uchicago.2743).

[92] V. I. Barsanov et al. “Artificial neutrino source based on the ^{37}Ar isotope”. In: *Physics of Atomic Nuclei* 70.2 (Feb. 2007), pp. 300–310. ISSN: 1562-692X. DOI: [10.1134/S1063778807020111](https://doi.org/10.1134/S1063778807020111).

[93] S. J. Haselschwardt et al. “Improved calculations of beta decay backgrounds to new physics in liquid xenon detectors”. In: *Phys. Rev. C* 102 (2020), p. 065501. DOI: [10.1103/PhysRevC.102.065501](https://doi.org/10.1103/PhysRevC.102.065501). arXiv: [2007.13686 \[hep-ex\]](https://arxiv.org/abs/2007.13686).

[94] J. Allison et al. “Recent developments in Geant4”. In: *Nucl. Instrum. Meth. A* 835 (2016), pp. 186–225. DOI: [10.1016/j.nima.2016.06.125](https://doi.org/10.1016/j.nima.2016.06.125).

[95] Joaquim Palacio. *Xe124-Xe136 ratio final result for SR0 XENONnT*. XENONnT internal note. Access restricted. Accessed 31.07.2025. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenonnt:jpalacio:rga_xenonnt_xe124_xe136_result.

[96] J. Kotila and F. Iachello. “Phase space factors for double- β decay”. In: *Phys. Rev. C* 85 (2012), p. 034316. DOI: [10.1103/PhysRevC.85.034316](https://doi.org/10.1103/PhysRevC.85.034316). arXiv: [1209.5722 \[nucl-th\]](https://arxiv.org/abs/1209.5722).

[97] J. Kotila. (private communication to XENONnT Collaboration).

[98] David G. Cerdeño et al. “Physics from solar neutrinos in dark matter direct detection experiments”. In: *JHEP* 05 (2016). [Erratum: *JHEP* 09, 048 (2016)], p. 118. DOI: [10.1007/JHEP09\(2016\)048](https://doi.org/10.1007/JHEP09(2016)048). arXiv: [1604.01025 \[hep-ph\]](https://arxiv.org/abs/1604.01025).

[99] L. Wolfenstein. “Neutrino Oscillations in Matter”. In: *Phys. Rev. D* 17 (1978), pp. 2369–2374. DOI: [10.1103/PhysRevD.17.2369](https://doi.org/10.1103/PhysRevD.17.2369).

[100] S. P. Mikheyev and A. Yu. Smirnov. “Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos”. In: *Sov. J. Nucl. Phys.* 42 (1985), pp. 913–917.

[101] Albert C. Thompson et al. *X-Ray Data Booklet*. October 2009. Center for X-Ray Optics, Lawrence Berkeley National Laboratory. Berkeley, CA, USA, 2009. URL: <https://xdb.lbl.gov>.

[102] Jiunn-Wei Chen et al. "Low-energy electronic recoil in xenon detectors by solar neutrinos". In: *Physics Letters B* 774 (2017), pp. 656–661. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2017.10.029>.

[103] M. Agostini et al. "Comprehensive measurement of pp -chain solar neutrinos". In: *Nature* 562.7728 (2018), pp. 505–510. DOI: <10.1038/s41586-018-0624-y>.

[104] L.L. Lucas and M.P. Unterweger. "Comprehensive Review and Critical Evaluation of the Half-Life of Tritium". In: *J. Res. Natl. Inst. Stand. Technol.* 105.4 (2000), pp. 541–549. DOI: <10.6028/jres.105.043>.

[105] R. D. Peccei and Helen R. Quinn. "CP Conservation in the Presence of Pseudoparticles". In: *Phys. Rev. Lett.* 38 (25 June 1977), pp. 1440–1443. DOI: <10.1103/PhysRevLett.38.1440>. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.38.1440>.

[106] Luca Di Luzio et al. "The landscape of QCD axion models". In: *Phys. Rept.* 870 (2020), pp. 1–117. DOI: <10.1016/j.physrep.2020.06.002>. arXiv: [2003.01100 \[hep-ph\]](2003.01100[hep-ph]).

[107] Kazuo Fujikawa and Robert E. Shrock. "Magnetic Moment of a Massive Neutrino and Neutrino-Spin Rotation". In: *Phys. Rev. Lett.* 45 (12 Sept. 1980), pp. 963–966. DOI: <10.1103/PhysRevLett.45.963>. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.45.963>.

[108] Nicole F. Bell et al. "Model independent bounds on magnetic moments of Majorana neutrinos". In: *Phys. Lett. B* 642 (2006), pp. 377–383. DOI: <10.1016/j.physletb.2006.09.055>. arXiv: <hep-ph/0606248>.

[109] Nicole F. Bell et al. "How magnetic is the Dirac neutrino?" In: *Phys. Rev. Lett.* 95 (2005), p. 151802. DOI: <10.1103/PhysRevLett.95.151802>. arXiv: <hep-ph/0504134>.

[110] M. Agostini et al. "Limiting neutrino magnetic moments with Borexino Phase-II solar neutrino data". In: *Phys. Rev. D* 96.9 (2017), p. 091103. DOI: <10.1103/PhysRevD.96.091103>. arXiv: [1707.09355 \[hep-ex\]](1707.09355[hep-ex]).

[111] M. I. Dobroliubov and A. Yu. Ignatiev. "Millicharged particles". In: *Phys. Rev. Lett.* 65 (6 Aug. 1990), pp. 679–682. DOI: <10.1103/PhysRevLett.65.679>. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.65.679>.

[112] K. S. Babu and R. R. Volkas. "Bounds on minicharged neutrinos in the minimal Standard Model". In: *Phys. Rev. D* 46 (1992), R2764–R2768. DOI: <10.1103/PhysRevD.46.R2764>. arXiv: <hep-ph/9208260>.

[113] Sudip Jana, Michael Klasen, and Vishnu P. K. "How Charged Can Neutrinos Be?" In: (Apr. 2025). arXiv: [2504.20044 \[hep-ph\]](2504.20044[hep-ph]).

[114] A. A. Prinz et al. "Search for Millicharged Particles at SLAC". In: *Phys. Rev. Lett.* 81 (6 Aug. 1998), pp. 1175–1178. DOI: <10.1103/PhysRevLett.81.1175>. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.1175>.

[115] A. Ball et al. "Search for millicharged particles in proton-proton collisions at $\sqrt{s} = 13$ TeV". In: *Phys. Rev. D* 102 (3 Aug. 2020), p. 032002. DOI: <10.1103/PhysRevD.102.032002>. URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.032002>.

[116] Chung-Chun Hsieh et al. "Discovery potential of multiton xenon detectors in neutrino electromagnetic properties". In: *Phys. Rev. D* 100.7 (2019), p. 073001. DOI: <10.1103/PhysRevD.100.073001>. arXiv: [1903.06085 \[hep-ph\]](1903.06085[hep-ph]).

[117] Wm. J. Veigle. "Photon cross sections from 0.1 keV to 1 MeV for elements Z = 1 to Z = 94". In: *Atom. Data Nucl. Data Tabl.* 5 (1973), pp. 51–111. DOI: [10.1016/S0092-640X\(73\)80015-4](https://doi.org/10.1016/S0092-640X(73)80015-4).

[118] O. G. Miranda and H. Nunokawa. "Non standard neutrino interactions: current status and future prospects". In: *New J. Phys.* 17.9 (2015), p. 095002. DOI: [10.1088/1367-2630/17/9/095002](https://doi.org/10.1088/1367-2630/17/9/095002). arXiv: [1505.06254 \[hep-ph\]](https://arxiv.org/abs/1505.06254).

[119] Y. Farzan and M. Tortola. "Neutrino oscillations and Non-Standard Interactions". In: *Front. in Phys.* 6 (2018), p. 10. DOI: [10.3389/fphy.2018.00010](https://doi.org/10.3389/fphy.2018.00010). arXiv: [1710.09360 \[hep-ph\]](https://arxiv.org/abs/1710.09360).

[120] Thomas Schwemberger and Tien-Tien Yu. "Detecting beyond the standard model interactions of solar neutrinos in low-threshold dark matter detectors". In: *Phys. Rev. D* 106 (1 July 2022), p. 015002. DOI: [10.1103/PhysRevD.106.015002](https://doi.org/10.1103/PhysRevD.106.015002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.015002>.

[121] Roni Harnik, Joachim Kopp, and Pedro A. N. Machado. "Exploring nu Signals in Dark Matter Detectors". In: *JCAP* 07 (2012), p. 026. DOI: [10.1088/1475-7516/2012/07/026](https://doi.org/10.1088/1475-7516/2012/07/026). arXiv: [1202.6073 \[hep-ph\]](https://arxiv.org/abs/1202.6073).

[122] Bhaskar Dutta et al. "Non-standard neutrino interactions in light mediator models at reactor experiments". In: *JHEP* 03 (2023), p. 163. DOI: [10.1007/JHEP03\(2023\)163](https://doi.org/10.1007/JHEP03(2023)163). arXiv: [2209.13566 \[hep-ph\]](https://arxiv.org/abs/2209.13566).

[123] Bhaskar Dutta et al. "Non-standard neutrino interactions mediated by a light scalar at DUNE". In: *JHEP* 07 (2024), p. 213. DOI: [10.1007/JHEP07\(2024\)213](https://doi.org/10.1007/JHEP07(2024)213). arXiv: [2401.02107 \[hep-ph\]](https://arxiv.org/abs/2401.02107).

[124] Paola Arias et al. "WISPy Cold Dark Matter". In: *JCAP* 06 (2012), p. 013. DOI: [10.1088/1475-7516/2012/06/013](https://doi.org/10.1088/1475-7516/2012/06/013). arXiv: [1201.5902 \[hep-ph\]](https://arxiv.org/abs/1201.5902).

[125] E. Aprile et al. "XENONnT WIMP search: Signal and background modeling and statistical inference". In: *Phys. Rev. D* 111.10 (2025), p. 103040. DOI: [10.1103/PhysRevD.111.103040](https://doi.org/10.1103/PhysRevD.111.103040). arXiv: [2406.13638 \[physics.data-an\]](https://arxiv.org/abs/2406.13638).

[126] Francesco Toschi. "Design of the field cage and charge response of the XENONnT dark matter experiment". PhD thesis. Freiburg U., 2023. DOI: [10.6094/UNIFR/234323](https://doi.org/10.6094/UNIFR/234323).

[127] E. Aprile et al. "Detector signal characterization with a Bayesian network in XENONnT". In: *Phys. Rev. D* 108.1 (2023), p. 012016. DOI: [10.1103/PhysRevD.108.012016](https://doi.org/10.1103/PhysRevD.108.012016). arXiv: [2304.05428 \[hep-ex\]](https://arxiv.org/abs/2304.05428).

[128] C. J. Clopper and E. S. Pearson. "The use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial". In: *Biometrika* 26.4 (1934). DOI: [10.1093/biomet/26.4.404](https://doi.org/10.1093/biomet/26.4.404).

[129] Pilar Coloma et al. "Constraining new physics with Borexino Phase-II spectral data". In: *JHEP* 07 (2022). [Erratum: *JHEP* 11, 138 (2022)], p. 138. DOI: [10.1007/JHEP07\(2022\)138](https://doi.org/10.1007/JHEP07(2022)138). arXiv: [2204.03011 \[hep-ph\]](https://arxiv.org/abs/2204.03011).

[130] V. De Romeri et al. "Physics implications of a combined analysis of COHERENT CsI and LAr data". In: *JHEP* 04 (2023), p. 035. DOI: [10.1007/JHEP04\(2023\)035](https://doi.org/10.1007/JHEP04(2023)035). arXiv: [2211.11905 \[hep-ph\]](https://arxiv.org/abs/2211.11905).

[131] Valentina De Romeri, Dimitrios K. Papoulias, and Gonzalo Sanchez Garcia. "Implications of the first CONUS+ measurement of coherent elastic neutrino-nucleus scattering". In: *Phys. Rev. D* 111.7 (2025), p. 075025. DOI: [10.1103/PhysRevD.111.075025](https://doi.org/10.1103/PhysRevD.111.075025). arXiv: [2501.17843 \[hep-ph\]](https://arxiv.org/abs/2501.17843).

[132] Yu. M. Andreev et al. “Search for Light Dark Matter with NA64 at CERN”. In: *Phys. Rev. Lett.* 131.16 (2023), p. 161801. DOI: [10.1103/PhysRevLett.131.161801](https://doi.org/10.1103/PhysRevLett.131.161801). arXiv: [2307.02404 \[hep-ex\]](https://arxiv.org/abs/2307.02404).

[133] A. G. Beda et al. “GEMMA experiment: Three years of the search for the neutrino magnetic moment”. In: *Phys. Part. Nucl. Lett.* 7 (2010), pp. 406–409. DOI: [10.1134/S1547477110060063](https://doi.org/10.1134/S1547477110060063). arXiv: [0906.1926 \[hep-ex\]](https://arxiv.org/abs/0906.1926).

[134] J. Aalbers et al. “Search for new physics in low-energy electron recoils from the first LZ exposure”. In: *Phys. Rev. D* 108.7 (2023), p. 072006. DOI: [10.1103/PhysRevD.108.072006](https://doi.org/10.1103/PhysRevD.108.072006). arXiv: [2307.15753 \[hep-ex\]](https://arxiv.org/abs/2307.15753).

[135] M. Deniz et al. “Measurement of $\bar{\nu}_e$ -electron scattering cross section with a CsI(Tl) scintillating crystal array at the Kuo-Sheng nuclear power reactor”. In: *Phys. Rev. D* 81 (7 Apr. 2010), p. 072001. DOI: [10.1103/PhysRevD.81.072001](https://doi.org/10.1103/PhysRevD.81.072001). URL: <https://link.aps.org/doi/10.1103/PhysRevD.81.072001>.

[136] K. Abe et al. “Search for exotic neutrino-electron interactions using solar neutrinos in XMASS-I”. In: *Phys. Lett. B* 809 (2020), p. 135741. DOI: [10.1016/j.physletb.2020.135741](https://doi.org/10.1016/j.physletb.2020.135741). arXiv: [2005.11891 \[hep-ex\]](https://arxiv.org/abs/2005.11891).

[137] E. Aprile et al. “Radon removal in XENONnT down to the solar neutrino level”. In: *Phys. Rev. X* (Aug. 2025). DOI: [10.1103/zc1w-88p6](https://doi.org/10.1103/zc1w-88p6). URL: <https://link.aps.org/doi/10.1103/zc1w-88p6>.

[138] M. Szydagis et al. *Noble Element Simulation Technique*. Version v2.4.0beta. Aug. 2023. DOI: [10.5281/zenodo.8215927](https://doi.org/10.5281/zenodo.8215927). URL: <https://doi.org/10.5281/zenodo.8215927>.

[139] Peter Gaemers et al. *XENONnT/WFSim: v1.0.2*. Version v1.0.2. Oct. 2022. DOI: [10.5281/zenodo.7216324](https://doi.org/10.5281/zenodo.7216324). URL: <https://doi.org/10.5281/zenodo.7216324>.

[140] Diego Ramirez Garcia. “Simulating the XENONnT dark matter experiment: backgrounds and WIMP sensitivity”. PhD thesis. Freiburg U., 2022. DOI: [10.6094/UNIFR/228338](https://doi.org/10.6094/UNIFR/228338).

[141] M. Szydagis et al. “Investigating the XENON1T low-energy electronic recoil excess using NEST”. In: *Phys. Rev. D* 103.1 (2021), p. 012002. DOI: [10.1103/PhysRevD.103.012002](https://doi.org/10.1103/PhysRevD.103.012002). arXiv: [2007.00528 \[hep-ex\]](https://arxiv.org/abs/2007.00528).

[142] E. Aprile et al. “Light Dark Matter Search with Ionization Signals in XENON1T”. In: *Phys. Rev. Lett.* 123.25 (2019). Supplemental Material, p. 251801. DOI: [10.1103/PhysRevLett.123.251801](https://doi.org/10.1103/PhysRevLett.123.251801). arXiv: [1907.11485 \[hep-ex\]](https://arxiv.org/abs/1907.11485).

[143] Marc Schumann et al. “Dark matter sensitivity of multi-ton liquid xenon detectors”. In: *JCAP* 10 (2015), p. 016. DOI: [10.1088/1475-7516/2015/10/016](https://doi.org/10.1088/1475-7516/2015/10/016). arXiv: [1506.08309 \[physics.ins-det\]](https://arxiv.org/abs/1506.08309).

[144] Ryan Gibbons and Peter Sorensen. *On track for discovery of sub-GeV dark matter with liquid xenon TPCs*. Talk presented at the TeV Particle Astrophysics Conference (TeVPA 2024). Presentation slides. Talk date: August 26, 2024. Chicago, IL, USA, Aug. 2024. URL: <https://indico.uchicago.edu/event/427/contributions/1344/attachments/371/566/tevpa%202024.pdf>.

[145] D. S. Akerib et al. “Extending light WIMP searches to single scintillation photons in LUX”. In: *Phys. Rev. D* 101.4 (2020), p. 042001. DOI: [10.1103/PhysRevD.101.042001](https://doi.org/10.1103/PhysRevD.101.042001). arXiv: [1907.06272 \[astro-ph.CO\]](https://arxiv.org/abs/1907.06272).

[146] Lanqing Yuan. “First Search for Light Dark Matter in the Neutrino Fog with XENONnT”. PhD thesis. Chicago U., 2025. DOI: [10.6082/uchicago.14815](https://doi.org/10.6082/uchicago.14815).

[147] E. Aprile et al. “Emission of single and few electrons in XENON1T and limits on light dark matter”. In: *Phys. Rev. D* 106.2 (2022). [Erratum: *Phys. Rev. D* 110, 109903 (2024)], p. 022001. DOI: [10.1103/PhysRevD.106.022001](https://doi.org/10.1103/PhysRevD.106.022001). arXiv: [2112.12116](https://arxiv.org/abs/2112.12116) [hep-ex].

[148] Peter Sorensen. “Electron train backgrounds in liquid xenon dark matter search detectors are indeed due to thermalization and trapping”. In: (Feb. 2017). arXiv: [1702.04805](https://arxiv.org/abs/1702.04805) [physics.ins-det].

[149] D. S. Akerib et al. “Investigation of background electron emission in the LUX detector”. In: *Phys. Rev. D* 102.9 (2020), p. 092004. DOI: [10.1103/PhysRevD.102.092004](https://doi.org/10.1103/PhysRevD.102.092004). arXiv: [2004.07791](https://arxiv.org/abs/2004.07791) [physics.ins-det].

[150] XENONnT Collaboration and Xu, Dacheng. *axidence*. Version 0.3.3. Code repository. May 16, 2025. URL: <https://github.com/XENONnT/axidence>.

[151] E. Aprile et al. “Search for Coherent Elastic Scattering of Solar ^8B Neutrinos in the XENON1T Dark Matter Experiment”. In: *Phys. Rev. Lett.* 126 (2021), p. 091301. DOI: [10.1103/PhysRevLett.126.091301](https://doi.org/10.1103/PhysRevLett.126.091301). arXiv: [2012.02846](https://arxiv.org/abs/2012.02846) [hep-ex].

[152] Guolin Ke et al. “LightGBM: A Highly Efficient Gradient Boosting Decision Tree”. In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon et al. Vol. 30. Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/6449f44a102fde848669bdd9eb6b76fa-Paper.pdf.

[153] Glen Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *Eur. Phys. J. C* 71 (2011). [Erratum: *Eur. Phys. J. C* 73, 2501 (2013)]. DOI: [10.1140/epjc/s10052-011-1554-0](https://doi.org/10.1140/epjc/s10052-011-1554-0). arXiv: [1007.1727](https://arxiv.org/abs/1007.1727) [physics.data-an].

[154] Peter Sorensen. “Anisotropic diffusion of electrons in liquid xenon with application to improving the sensitivity of direct dark matter searches”. In: *Nucl. Instrum. Meth. A* 635 (2011), pp. 41–43. DOI: [10.1016/j.nima.2011.01.089](https://doi.org/10.1016/j.nima.2011.01.089). arXiv: [1102.2865](https://arxiv.org/abs/1102.2865) [astro-ph.IM].

[155] Scikit-learn developers. *Permutation importance*. Web page. Accessed on 08.05.2025. URL: https://scikit-learn.org/stable/modules/permuation_importance.html.

[156] Leo Breiman. “Random Forests”. In: *Machine Learning* 45.1 (Oct. 2001), pp. 5–32. DOI: [10.1023/A:1010933404324](https://doi.org/10.1023/A:1010933404324). URL: <https://doi.org/10.1023/A:1010933404324>.

[157] C. H. Faham et al. “Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes”. In: *JINST* 10.09 (2015), P09010. DOI: [10.1088/1748-0221/10/09/P09010](https://doi.org/10.1088/1748-0221/10/09/P09010). arXiv: [1506.08748](https://arxiv.org/abs/1506.08748) [physics.ins-det].

[158] XENON Collaboration. *XENONnT/appletree: v0.3.2*. Zenodo. 2024. DOI: [10.5281/zenodo.10784631](https://doi.org/10.5281/zenodo.10784631). URL: <https://doi.org/10.5281/zenodo.10784631>.

[159] Dacheng Xu and Zihao Xu. *Charge yield of ^{37}Ar L-shell with S2-Only events*. XENONnT internal note. Access restricted. Accessed 12.06.2025. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenonnt:ar37_l_shell_sensi:charge_yield.

[160] D. S. Akerib et al. “Tritium calibration of the LUX dark matter experiment”. In: *Phys. Rev. D* 93.7 (2016), p. 072009. DOI: [10.1103/PhysRevD.93.072009](https://doi.org/10.1103/PhysRevD.93.072009). arXiv: [1512.03133](https://arxiv.org/abs/1512.03133) [physics.ins-det].

[161] L. Baudis et al. "The first dual-phase xenon TPC equipped with silicon photo-multipliers and characterisation with ^{37}Ar ". In: *Eur. Phys. J. C* 80.5 (2020), p. 477. DOI: [10.1140/epjc/s10052-020-8031-6](https://doi.org/10.1140/epjc/s10052-020-8031-6). arXiv: [2003.01731](https://arxiv.org/abs/2003.01731) [astro-ph.IM].

[162] E. M. Boulton et al. "Calibration of a two-phase xenon time projection chamber with a ^{37}Ar source". In: *JINST* 12.08 (2017), P08004. DOI: [10.1088/1748-0221/12/08/P08004](https://doi.org/10.1088/1748-0221/12/08/P08004). arXiv: [1705.08958](https://arxiv.org/abs/1705.08958) [physics.ins-det].

[163] E. Aprile et al. "First Search for Light Dark Matter in the Neutrino Fog with XENONnT". In: *Phys. Rev. Lett.* 134.11 (2025), p. 111802. DOI: [10.1103/PhysRevLett.134.111802](https://doi.org/10.1103/PhysRevLett.134.111802). arXiv: [2409.17868](https://arxiv.org/abs/2409.17868) [hep-ex].

[164] D. Baxter et al. "Recommended conventions for reporting results from direct dark matter searches". In: *Eur. Phys. J. C* 81.10 (2021), p. 907. DOI: [10.1140/epjc/s10052-021-09655-y](https://doi.org/10.1140/epjc/s10052-021-09655-y). arXiv: [2105.00599](https://arxiv.org/abs/2105.00599) [hep-ex].

[165] Timon Emken. "Solar reflection of light dark matter with heavy mediators". In: *Phys. Rev. D* 105.6 (2022), p. 063020. DOI: [10.1103/PhysRevD.105.063020](https://doi.org/10.1103/PhysRevD.105.063020). arXiv: [2102.12483](https://arxiv.org/abs/2102.12483) [hep-ph].

[166] Timon Emken, Rouven Essig, and Hailin Xu. "Solar reflection of dark matter with dark-photon mediators". In: *JCAP* 07 (2024), p. 023. DOI: [10.1088/1475-7516/2024/07/023](https://doi.org/10.1088/1475-7516/2024/07/023). arXiv: [2404.10066](https://arxiv.org/abs/2404.10066) [hep-ph].

[167] E. Aprile et al. "Search for Light Dark Matter in Low-Energy Ionization Signals from XENONnT". In: *Phys. Rev. Lett.* 134.16 (2025), p. 161004. DOI: [10.1103/PhysRevLett.134.161004](https://doi.org/10.1103/PhysRevLett.134.161004). arXiv: [2411.15289](https://arxiv.org/abs/2411.15289) [hep-ex].

[168] Guofang Shen et al. "Search for Solar Boosted Dark Matter Particles at the PandaX-4T Experiment". In: *Phys. Rev. Lett.* 134.16 (2025), p. 161003. DOI: [10.1103/PhysRevLett.134.161003](https://doi.org/10.1103/PhysRevLett.134.161003). arXiv: [2412.19970](https://arxiv.org/abs/2412.19970) [hep-ex].

[169] Z. Y. Zhang et al. "Experimental Limits on Solar Reflected Dark Matter with a New Approach on Accelerated-Dark-Matter–Electron Analysis in Semiconductors". In: *Phys. Rev. Lett.* 132.17 (2024), p. 171001. DOI: [10.1103/PhysRevLett.132.171001](https://doi.org/10.1103/PhysRevLett.132.171001). arXiv: [2309.14982](https://arxiv.org/abs/2309.14982) [hep-ex].

[170] Prakruth Adari et al. "First Direct-Detection Results on Sub-GeV Dark Matter Using the SENSEI Detector at SNOLAB". In: *Phys. Rev. Lett.* 134.1 (2025), p. 011804. DOI: [10.1103/PhysRevLett.134.011804](https://doi.org/10.1103/PhysRevLett.134.011804). arXiv: [2312.13342](https://arxiv.org/abs/2312.13342) [astro-ph.CO].

[171] Christian Weinheimer. *Low radon and low internal radioactivity for dark matter and rare event xenon detectors*. ERC Advanced Grant project 101055063; Start date: 1 November 2022; End date: 31 October 2027. DOI: [10.3030/101055063](https://doi.org/10.3030/101055063). (Visited on 07/03/2025).

[172] Florian Jörg. "From 222Rn measurements in XENONnT and HeXe to radon mitigation in future liquid xenon experiments". PhD thesis. Heidelberg U., 2022. DOI: [10.11588/heidok.00031915](https://doi.org/10.11588/heidok.00031915).

[173] Julia Dierle et al. "Reduction of ^{222}Rn -induced backgrounds in a hermetic dual-phase xenon time projection chamber". In: *Eur. Phys. J. C* 83.1 (2023), p. 9. DOI: [10.1140/epjc/s10052-022-11151-w](https://doi.org/10.1140/epjc/s10052-022-11151-w). arXiv: [2209.00362](https://arxiv.org/abs/2209.00362) [physics.ins-det].

[174] E. Aprile et al. "Offline tagging of radon-induced backgrounds in XENON1T and applicability to other liquid xenon time projection chambers". In: *Phys. Rev. D* 110.1 (2024), p. 012011. DOI: [10.1103/PhysRevD.110.012011](https://doi.org/10.1103/PhysRevD.110.012011). arXiv: [2403.14878](https://arxiv.org/abs/2403.14878) [hep-ex].

[175] Tina Pollmann. *Accidental coincidence background modelling*. Access restricted. Accessed 24.06.2025. DARWIN collaboration. May 2023. URL: https://indico.physik.uni-heidelberg.de/event/695/contributions/2057/attachments/943/1339/ColMeeting2023_WG1_Pollmann.pdf.