

Extending the physics reach of the LUX and LZ dark matter experiments by lowering the scintillation threshold

Nellie Marangou

Imperial College London

Department of Physics

A dissertation submitted to Imperial College London

for the degree of Doctor of Philosophy.

June 2020

Abstract

Dark matter experiments searching for weakly interacting massive particles (WIMPs), which are hypothesised to constitute the dark matter content of the Universe, probe other important physical processes producing very low energy (\sim keV) nuclear recoils. These include the scattering of (non-WIMP) light dark matter particles and the coherent nuclear scattering of ^8B and ^3He -proton fusion (hep) solar neutrinos. Noble liquid dual-phase detectors, and particularly the Large Underground Xenon experiment (LUX), have placed stringent upper limits on the spin-independent WIMP-nucleon scattering cross-section at masses above ~ 6 GeV. LUX was located at the Sanford Underground Research Facility (SURF), one mile underground, where it completed two successful runs between 2013 and 2016, utilising 300 kg of liquid xenon (LXe). The site will now be used by the forthcoming LUX-ZEPLIN (LZ) experiment, which is planned to start running in late 2020 and will utilise 10 tonnes of LXe.

In this work we present a novel analysis technique for LXe time projection chambers (TPCs) that allows for a lower threshold by relying on events with a prompt scintillation signal consisting of single detected photons. The energy threshold of the LUX dark matter experiment was primarily determined by the smallest scintillation response detectable, which previously required a 2-fold coincidence signal in its photomultiplier arrays, enforced in data analysis. The technique presented here exploits the double photoelectron emission effect observed in some photomultiplier models at vacuum ultraviolet wavelengths. We demonstrate this analysis using an electron recoil calibration dataset and place new constraints on the spin-independent scattering cross-section of WIMPs down to 2.5 GeV WIMP mass using the 2013 LUX dataset. This new technique is promising to enhance light WIMP and astrophysical neutrino searches in next-generation liquid xenon experiments. To this end, we investigate the possible improvement in sensitivity brought about by applying this technique in LZ.

Declaration

I declare that the material in this thesis is my own work. Any work that is not my own is appropriately stated and referenced. The specific contributions in each chapter are listed below.

Chapters 2 and 3 review the literature on dark matter and direct searches for its interactions in scattering experiments. The latter also describes the LUX detector, and as such, includes work by the LUX collaboration.

Chapter 4 reviews photomultiplier tube technologies from background information and describes a specific calibration of the LUX photosensors. The calibration presented here is my contribution, but data used have been pre-processed by the LUX collaboration.

Chapter 5 and 6 include extending results published by LUX to lower energies, both for a calibration study and a WIMP search. The original results are presented and compared to the new proposed technique, which I was responsible for. The work discussed in these chapters was published in Ref. [1].

The initial sensitivity study in Chapter 7 is my own work, while the result of a complete analysis shown in this chapter was performed by an LZ collaborator.

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

- Nellie Marangou

Acknowledgements

First and foremost, I would like to acknowledge and thank my supervisor Professor Henrique Araújo for guiding an engaging project and being my mentor during the duration of my PhD. Your guidance has been essential, inspiring and motivating and I am grateful for having the chance to work with you.

Thanks to my colleagues at Imperial, I enjoyed working and learning from all of you. I would also like to thank the LUX and LZ collaborations for all the opportunities, interesting discussions, experiences and knowledge they shared with me.

Special thanks to my friends at Imperial, that shared this journey with me in all of its different eras. Through our many coffee breaks, you have been a constant source of motivation, understanding and fun and I cannot imagine this experience without you. I feel fortunate to have met you and learned something valuable from each of you.

This thesis would not have been possible without my parents, Katerina and Akis. I would like to thank them for endorsing and driving my dreams and making their fulfilment possible through their unfailing support and love.

I am grateful to my partner, Filippos, for believing in me, encouraging me and always being there for me.

Contents

List of figures	xiii
List of tables	xvii
1 Introduction	1
2 The search for the missing mass	5
2.1 Evidence for dark matter	5
2.1.1 Galactic scale	5
2.1.2 Cosmological scale	8
2.2 Dark matter phenomenology	10
2.2.1 General characteristics	10
2.2.2 Weakly Interacting Massive Particles	11
2.3 Detection channels	11
2.4 Rare event searches at low energies	13
2.4.1 Light dark matter	13
2.4.2 Coherent elastic neutrino-nucleus scattering of solar, atmospheric and supernova neutrinos	18
2.5 Summary	20
3 Particle detection in LXe-TPCs	21
3.1 WIMP-nucleon scattering	21
3.1.1 Differential event rate	21
3.1.2 Spin-dependent and spin-independent interactions	23
3.2 Coherent elastic neutrino-nucleus scattering	25
3.3 Signal production in liquid xenon	27
3.3.1 Properties of liquid xenon as a detector medium	27
3.3.2 Scintillation and ionisation from particle interactions in liquid xenon	28
3.3.3 Scintillation and ionisation yields in liquid xenon	30
3.3.4 Noble Element Simulation Technique	31
3.4 Dual-phase time projection chambers	33
3.4.1 Operating principles	33
3.4.2 ER-NR discrimination	34

3.5	The LUX experiment	35
3.5.1	Overview of the experimental setup	35
3.5.2	Data acquisition and reduction	41
3.6	The LUX-ZEPLIN experiment	44
3.7	Low energy and low mass sensitivity techniques	45
3.7.1	S2-only analyses	45
3.7.2	Migdal and Nuclear Bremsstrahlung	46
3.7.3	Doping with light elements	48
3.8	Summary	49
4	Understanding the single VUV photon response of the LUX photomultipliers	51
4.1	Photomultiplier tubes	51
4.1.1	Detection principles	51
4.1.2	Detection Efficiency	53
4.1.3	Double photoelectron emission	54
4.1.4	Afterpulsing	55
4.1.5	Dark current	56
4.2	Calibrating the single photoelectron response in LUX	57
4.2.1	Gain calibration using the LED system	57
4.2.2	Single photoelectron distributions using dark counts	58
4.3	Characterising the single photon response in LUX	64
4.3.1	Calibration using S1 light	64
4.4	Summary	72
5	Exploiting the single VUV photon response in LUX tritium data	73
5.1	Reproducing the standard tritium calibration result	74
5.2	Lowering the scintillation threshold for the tritium dataset	78
5.2.1	DPE cut acceptance	79
5.2.2	Fake single-photon rate from over-fluctuating dark counts	82
5.2.3	Single VUV photon events in the tritium calibration dataset	82
5.2.4	Electron recoil scintillation and ionisation yields at low energies	87
5.2.5	Removing the DPE cut	88
5.3	Summary	91
6	Extending light WIMP searches to single scintillation photons in LUX	93
6.1	Reproducing the original Run 3 result	93
6.2	Background model	95
6.2.1	Nuclear and electron recoil backgrounds	96
6.2.2	Accidental coincidence rate	98

6.3	WIMP search	99
6.3.1	Searching for single-photon S1 events	99
6.3.2	Signal model	107
6.3.3	Upper limit on spin independent cross section	110
6.3.4	Effect of model assumptions on upper limit	115
6.3.5	Evaluating the effect of the DPE cut	115
6.4	Summary	117
7	Lowering the scintillation threshold in LZ	119
7.1	Single photon calibration of the LZ TPC PMTs	119
7.2	Estimated sensitivity improvement	123
7.2.1	Detecting CE ν NS of ^8B neutrinos in LZ	124
7.2.2	Considering low mass WIMPs	129
7.3	Result of a full analysis with a lower threshold	131
7.4	Summary	132
8	Conclusions and outlook	133
A	Confidence intervals: Rolke's method	137
	Bibliography	141

List of figures

2.1	Rotation curve of the spiral galaxy NGC 6503.	7
2.2	Rotation curve of the Milky Way.	7
2.3	CMB temperature map.	8
2.4	CMB temperature anisotropy power spectrum.	9
2.5	Dark matter detection channels.	12
2.6	Dark matter direct detection response channels.	13
2.7	Evolution of dark matter abundance.	15
2.8	FIMP and LOSP Feynman diagram vertices.	16
2.9	Annihilation cross-section for thermal and non-thermal DM.	16
2.10	Solar neutrino production chain.	20
3.1	Event rate for different target media.	25
3.2	Neutrino flux and recoil spectra.	26
3.3	Scintillation and ionisation yield for nuclear recoil interactions.	31
3.4	Scintillation and ionisation yield for electron recoil interactions.	32
3.5	Operating principle of dual-phase xenon detectors.	34
3.6	Discrimination parameter for ER and NR interactions in LUX.	36
3.7	Schematic of the LUX detector.	37
3.8	Example LUX event from a ^{83m}Kr interaction.	38
3.9	Doke method of extracting the g_1 and g_2 gain factors.	39
3.10	Schematic diagram of the LUX DD calibration experimental setup.	40

3.11	The LUX Data Processing Framework.	43
3.12	Pulse classifier decision tree.	43
3.13	Schematic of the LZ detector.	45
3.14	Recent S2-only analyses results.	47
3.15	Migdal and nuclear Bremsstrahlung upper limits.	48
4.1	Schematic of a photomultiplier tube.	52
4.2	The Hamamatsu R8778 PMT.	53
4.3	Measurement of the single photon response for an LZ PMT.	55
4.4	Example of single photoelectron pulse and spectrum.	57
4.5	Dark count distributions for the LUX top and bottom PMT arrays.	58
4.6	Fit parameters for each PMT, acquired through fitting of the SPE distributions.	61
4.7	Example distributions and fits of the SPE response.	62
4.8	Poisson mean and contamination ratio of two-phd to two-phe signals.	65
4.9	PMT fit parameters for single photon distributions.	67
4.10	PMT fit parameters for single photon distributions II.	68
4.11	Example distributions and fits of the single photon distribution I.	69
4.12	Example distributions and fits of the single photon distribution II.	70
4.13	Single photoelectron and single photon response for an example PMT.	72
5.1	Standard search analysis cuts illustrated on the tritium calibration dataset.	76
5.2	Radial position correction effect for the tritium 2-fold dataset.	77
5.3	Tritium S1 and S2 spectra with a 2-fold coincidence requirement.	77
5.4	Tritium energy spectrum for a 2-fold coincidence requirement.	78
5.5	A single-photon S1 event from the Dec. 2013 LUX tritium dataset.	79
5.6	DPE signal region for each active channel.	80
5.7	Light collection efficiency study using $^{83\text{m}}\text{Kr}$ data.	81

5.8	Analysis cuts on the tritium calibration dataset for single-photon S1 events.	83
5.9	Visualux display of a contamination event.	84
5.10	Electron lifetime correction map for a set of tritium single scatters.	86
5.11	S2 spectrum for tritium single-photon S1 events including background expectation.	86
5.12	Scatter plot for tritium events selected for this analysis.	87
5.13	The ER ionisation and scintillation yields.	88
5.14	S2 spectrum for tritium single-photon S1 events with NEST v2.0.0 prediction.	89
5.15	S2 spectrum for tritium single-photon S1 events, with no DPE cut.	90
5.16	Validation of DPE cut acceptance.	90
6.1	Comparison of the WIMP search result for the two analysis chains.	94
6.2	Example event in which the S1 pulse is due to a mis-classification.	95
6.3	S2-only spectrum following quality cuts.	98
6.4	Cleanliness cut on single-photon S1 events appearing in the WIMP search.	100
6.5	Position of single-photon S1 events appearing in the WIMP search.	101
6.6	Example pathological events removed using the S2 pulse quality cut.	104
6.7	Example event with mis-classified single detected photon pulse.	104
6.8	Illustrations of final single-photon S1 WIMP search events.	105
6.9	Electron lifetime correction map examples for WIMP search.	106
6.10	Single-photon S1 events in the 2013 LUX WIMP search exposure.	106
6.11	Differential recoil energy spectra for three different WIMP masses.	107
6.12	The NR ionisation and scintillation yields.	108
6.13	WIMP signal acceptance as a function of S2 threshold.	109
6.14	Background expectation and signal acceptance for WIMP mass dependent cuts.	110
6.15	Signal model along with defined region of interest for three WIMP masses.	111
6.16	Overall signal acceptance as function nuclear recoil energy.	111
6.17	Upper limit on the number of counts at each mass.	113

6.18	90% C.L. upper limits obtained using single-photon S1 events.	114
6.19	Effect of model assumptions on 90% C.L. upper limits.	116
6.20	Discovery significance for a 4 GeV WIMP and variable DPE cuts.	117
7.1	Example LED distributions for the LZ PMTs.	121
7.2	Example VUV distributions for the LZ PMTs.	122
7.3	LZ PMT map of quantum efficiency.	123
7.4	DPE fraction measurement for LZ PMTs.	123
7.5	Comparison of DPE fraction measurement to MC truth for LZ PMTs.	124
7.6	Event rate for CE ν NS of ^8B neutrinos and LZ simulated backgrounds.	125
7.7	Comparison of reconstructed to true energy for ^8B neutrinos.	126
7.8	LZ simulated dataset for the full exposure.	127
7.9	Expected ^8B and background counts for a variable DPE cut.	128
7.10	Confidence intervals for the observation on ^8B neutrinos.	128
7.11	Discovery frequency for a variable live-time and DPE cut.	129
7.12	LZ simulated dataset for the full exposure using the DPE technique.	130
7.13	Event rate for a 3 GeV WIMP and signal improvement for variable WIMP masses.	130
7.14	LZ WIMP sensitivity for standard and combined DPE analysis.	132
A.1	Coverage for a 4 GeV WIMP.	138

List of tables

4.1	Average dark count rates for the top and bottom PMT arrays.	63
6.1	Expected background for various ranges in S2 space for the single-photon WIMP-search analysis of the 2013 LUX WIMP search exposure.	99
6.2	Expected and observed background for the single-photon WIMP-search analysis of the 2013 LUX WIMP search exposure.	103
7.1	Signal and background counts for ^8B study.	126

Chapter 1

Introduction

Over the last few decades, the nature of dark matter has become one of the greatest unanswered questions in physics. Even though an overwhelming number of indirect observations indicates that most of the matter density in the Universe is dark—not composed of ordinary particles—, fundamental questions remain unanswered. Can a modified theory of gravity explain dark matter or is there a dark matter particle? If so, do dark matter particles interact with anything other than gravitationally? What is their nature and what was their production mechanism in the early Universe? A well-motivated class of particle that would resolve the dark matter problem in all its dimensions is the Weakly Interacting Massive Particle (WIMP). These neutral particles, with masses in the GeV to TeV scale, would have been produced as a cold (non-relativistic) thermal relic in the early Universe and survive to the present era with the right abundance to explain the astrophysical and cosmological evidence. A new particle with weak interactions is also appealing for particle physicists — and it would provide a mechanism for their direct detection.

A vast array of experimental methods have been developed and used to search for dark matter, with many focusing on the search for WIMPs, in particular their direct detection by elastic scattering in terrestrial detectors. In the absence of evidence for such an electroweak scale particle to date, the spotlight recently turned to light dark matter particles, in the few-GeV range and below. Chapter 2 of this thesis reviews the galactic and cosmological scale evidence for dark matter as well as the expected phenomenology of dark matter particles. The chapter continues with a discussion of light dark matter models such as non-thermal dark matter, asymmetric dark matter and strongly interacting massive particles. Other interesting rare processes that take place at low energies, and are accessible by the same detector technologies as those searching for low mass dark matter, are also presented, with the most notable one being the coherent nuclear elastic scattering of solar ^8B neutrinos.

In Chapter 3 we begin by discussing the direct detection of dark matter, outlining the theoretical framework behind WIMP-nucleon scattering and paying particular attention to scattering with a liquid xenon target. The discussion then turns more specifically to dual-phase

xenon time projection chambers and their operating principles. Such detectors achieve optimal sensitivity in the 20–1000 GeV mass range. As there is an increasing interest in searching also for lighter particles, other methodologies have succeeded in lowering the energy threshold of dual-phase xenon detectors and hence enhanced their sensitivity to models proposing lower mass candidates. Other avenues, such as the use of the Migdal and nuclear Bremsstrahlung effects and doping the liquid xenon target with light species as hydrogen have also been suggested and are reviewed in Chapter 3. This chapter concludes with an overview of the LUX experiment and its instrumentation.

LUX detects interactions through two different signals: the scintillation and the ionisation response channels, which measure prompt scintillation in the liquid phase (S1) and secondary scintillation (electroluminescence) in the gas phase (S2), respectively. During its operation, it was very successful at probing low energy interactions down to O(keV) energies, with the lowest accessible energy determined by the S1 signal threshold. Even though scintillation signals down to 1 photon are detectable, dark noise in the 122 LUX photomultipliers would overwhelm their detection. Thus, only scintillation signals with more than two detected photons have been considered in most LUX analyses (already a very low threshold compared to what has been achieved by similar experiments). The work presented in this thesis focuses on using an artifact of photomultipliers when responding to the VUV xenon scintillation, whereby double photoelectron emission (DPE) can take place with sizeable probability; this effect is exploited to lower the scintillation and hence energy threshold of LUX. Chapter 4 describes my work to calibrate the DPE effect for the LUX PMTs and the study of their dark count rates.

In Chapter 5 I describe my own analysis of LUX data from an injection of tritiated methane to calibrate low-energy electron recoil interactions, and attempt to lower the scintillation threshold of the equivalent published analysis. The chapter involves an investigation of whether any new background processes exist in this regime, and a search for events where the S1 pulse consisted of single-photon detections exhibiting the DPE effect. This analysis serves as a benchmark study for the use of the DPE technique to lower the energy threshold of these detectors also for nuclear recoil interactions (as the S1 pulse characteristics are the same for both types of recoil in this instance). The single-photon S1 analysis is compared against a Monte Carlo simulation of the same calibration and the effect of systematic uncertainties, originating in our incomplete understanding of signal production in liquid xenon at these energies, is discussed.

In Chapter 6 we proceed to apply the proposed technique to search for single-photon S1 events in the LUX WIMP search dataset acquired in 2013, accumulating 95 live-days of search exposure. The background model for this search is discussed, including both standard backgrounds and those originating from PMT dark counts, and the search is used to arrive at a new competitive cross-section limit for light WIMPs. Subsequently, the effects of model

assumptions on the new result are outlined. The chapter concludes with an evaluation of the effect of the DPE cut on the sensitivity. This analysis is my own work, with input from the LUX collaboration.

The successful demonstration of the proposed technique in LUX naturally motivates its application to the next generation experiments soon to come online such as LZ. The projected sensitivity improvement that LZ can benefit from when lowering its scintillation threshold using the DPE effect is presented in Chapter 7. Specifically, this improvement is investigated in the context of the ability of LZ to speed up a first observation of ^8B neutrinos. I performed some of the early sensitivity studies in this context, which have now been extended through detailed modelling by the Imperial team.

Finally, in Chapter 8 the main findings of the work presented in this thesis are summarised.

Chapter 2

The search for the missing mass

Most of the mass in the Universe is not composed of ordinary particles. And even though the the existence of this dark matter (DM) has become apparent through the structure of galaxies and all the way to the evolution of the Universe at the largest scales, the nature of this elusive substance is poorly understood with direct evidence yet to be observed. Theorists have come up with several dark matter candidates that could explain the astrophysical and cosmological observations with special focus on WIMPs.

The missing mass question has intrigued scientists for decades, with dozens of dedicated experiments built for its search. In this chapter we review the evidence for the existence of dark matter and the three experimental strategies used for its detection. The lack of an observation at the electroweak scale to date has led to the construction of other dark matter models, at lower masses, some of which are also accessible to scattering experiments and are presented here. We also discuss the equally elusive coherent nuclear scattering of astrophysical neutrinos that are accessible by the same detector technologies as those searching for WIMPs.

2.1 Evidence for dark matter

2.1.1 Galactic scale

The history of dark matter began with early indications that there is more in the Universe than could be seen by visible light. In the late 19th century scientists hypothesised the existence of objects that might emit little or no light, but could still be inferred by their gravitational pull on bright objects such as nearby stars. In 1932 Oort observed discrepancies in the vertical kinematics of stars, with their observed speed in our galactic neighbourhood being greater than predicted [2]. Even though he hinted at the existence of dark matter, he mostly attributed his observations to common astrophysical explanations such as an abundance of dim stars.

In 1933 a very different type of observation was made by Zwicky who applied the virial theorem to the Coma galaxy cluster. His findings indicated that its luminous mass was not sufficient to support the velocity dispersion of galaxies observed, with a large component of matter missing [3]. The virial theorem states that, for an equilibrium system of a gravitationally bound set of masses, the total kinetic energy is equal to (minus) half the gravitational potential energy. Similar measurements have since been performed to determine the mass of other galaxy clusters with comparable results [4, 5]. Following his observation, Zwicky famously wrote that *“If this should be verified, it would lead to the surprising result that dark matter exists in much greater density than luminous matter”* [3].

On a much smaller scale, strong evidence for dark matter also exists from the rotation curves of spiral galaxies. Experimental observations from Rubin and Ford [6] in the early 1970s first indicated that the majority of the luminous mass in spiral galaxies is contained within the bulge of the galaxy, with a less massive gaseous component spanning larger distances. Under this assumption and a Newtonian dynamics approximation, it is expected that the rotation speed of stars (v) would fall as the distance of the object to the centre of the galaxy (r) increases ($v \propto 1/\sqrt{r}$). Experimentally, though, it is observed that the rotation speed remains constant for much larger radii [6, 7], as seen in Figure 2.1 for an example spiral galaxy and in Figure 2.2 for the Milky Way. The discrepancy between the expected and observed velocity curves can be resolved with the addition of a dominant dark matter component, distributed in a roughly spherical shaped halo extending above and below the plane of the spiral and to much larger galactic radii [8]. The dark matter density in the Milky Way at the position of the Sun, ρ_o , is estimated to be between $(0.235 \pm 0.030) \text{ GeV/cm}^3$ and $(0.389 \pm 0.025) \text{ GeV/cm}^3$, with a reference value of $\rho_o = 0.3 \text{ GeV/cm}^3$ usually adopted.

Very different type of evidence comes from gravitational lensing, the bending of light from distant sources in the Universe as it travels past massive objects in the foreground. The amount of distortion is directly related to the mass of the object and hence can be used to accurately measure the mass of distant galaxies and galaxy clusters. Measurements agree with rotation curve observations and hence disagree with the amount of luminous mass, further supporting the theory of a major dark matter component [12, 13, 14, 15, 16].

The luminous mass of galaxies and galaxy clusters can also be determined using X-ray emission from the hot inter-galactic plasma [17]. The notable Bullet Cluster observation, involving the collision of two galaxy clusters, has further supported the dark matter theory [18]. X-ray imaging was used to locate the baryonic centre-of-mass of the system. Gravitational lensing was used to map the centre of mass of the system, including both luminous and any possible non-luminous components. The two locations were found to be significantly separated. Presumably, the lack of non-gravitational interaction between the dark matter component of the two clusters allowed these to pass through each other faster than the baryonic components, which were slowed down to create the observed spatial segregation. The cold dark matter

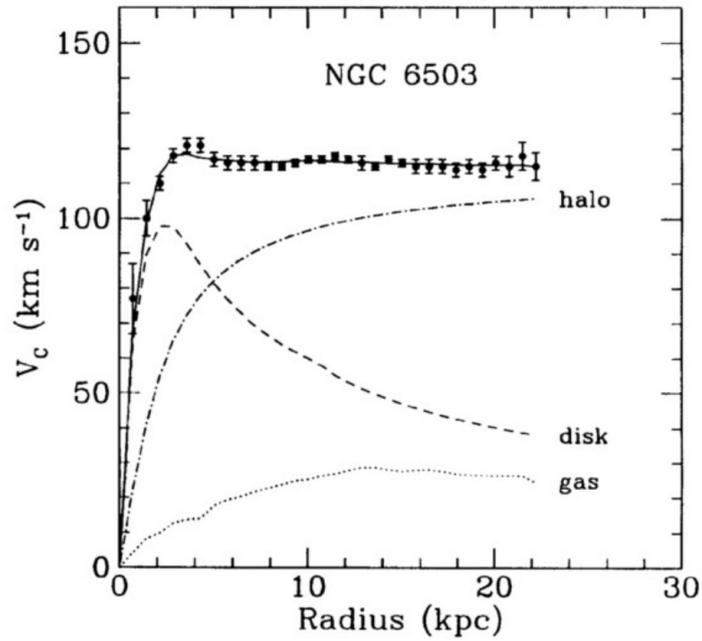


Figure 2.1: Rotation curve of the spiral galaxy NGC 6503. The dotted, dashed and dash-dotted lines indicate the contributions of the gas, disk and dark matter components, respectively [8].

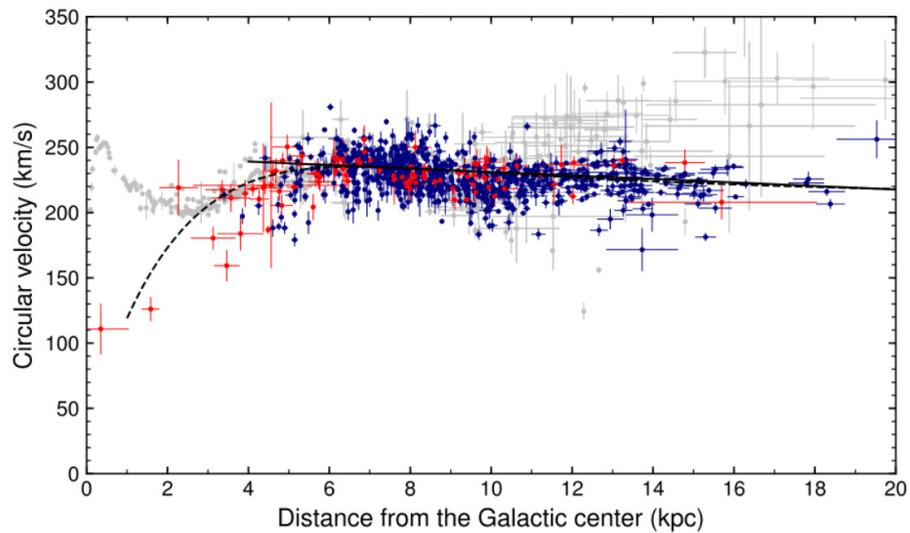


Figure 2.2: Measurements of the rotational velocities in the Milky Way shown for three different datasets in grey [9], red [10] and blue [11]. The Sun is located at ~ 8 kpc.

model is successful in explaining this behaviour while alternative models, such as a modified gravity theory, cannot adequately describe the phenomenon [19].

2.1.2 Cosmological scale

About 370,000 years after the Big Bang, ($T \sim \text{eV}$), electrons and protons combined to form the first hydrogen atoms, after which point the Universe became transparent to photons, which have existed as "relic" radiation ever since [20]. This radiation fills the Universe today, in the form of microwave photons, at a blackbody temperature of 2.7 K and is known as the Cosmic Microwave Background (CMB) radiation. The CMB radiation is mostly isotropic and homogeneous in nature at the 10^{-5} level. It was first observed accidentally [21] and has since been measured to great accuracy (Figure 2.3) [22].

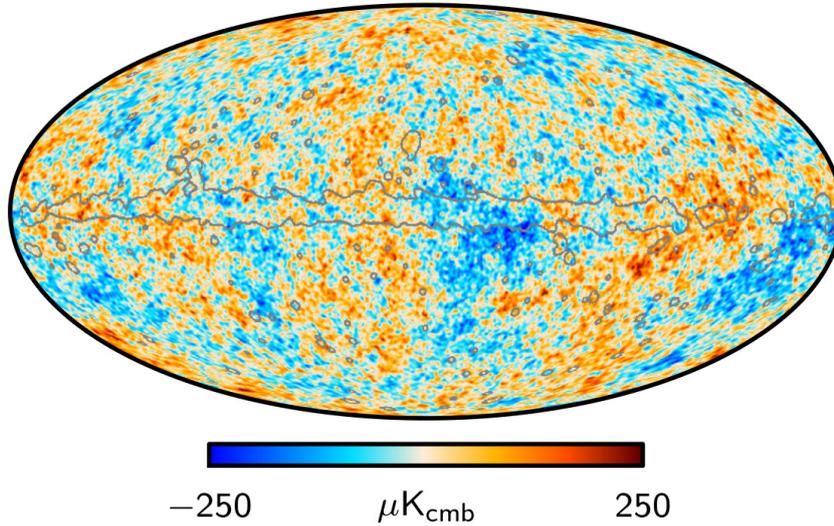


Figure 2.3: CMB temperature map derived from the Planck 2015, 9-year WMAP, and 408 MHz Haslam et al. observations [22].

The small temperature anisotropies in the CMB provide a powerful view into the history of the Universe. In particular, the Wilkinson Microwave Anisotropy Probe (WMAP) and Planck satellites gathered and placed tight constraints on the abundance of baryons and non-baryonic matter in the Universe [22, 23]. This was performed by mapping the size of fluctuations as a function of angular scale, in what is known as the CMB temperature anisotropy power spectrum. The temperature fluctuations are projected in a 2D spherical surface sky, so it is common to expand the temperature field using spherical harmonics. Using this approach, the power spectrum of the Gaussian fluctuations, C_l , can be defined as the variance of the harmonic coefficients (α_{lm}):

$$C_l = \frac{1}{2l+1} \sum_{m=-l}^l \langle |\alpha_{lm}|^2 \rangle \quad , \quad (2.1)$$

where $l = 0, \dots, \infty$ and $-l \leq m \leq l$ and the multipole moment, l , represents a given angular scale in the sky [24].

The oscillating nature of the power spectrum, seen in Figure 2.4, is due to density perturbations which have grown and formed acoustic standing waves before recombination. During this time density fluctuations seeded by inflation attracted matter gravitationally, before being driven to expand from the rising photon-baryon plasma density. This oscillation of gravitational collapse counteracted by thermal pressure (baryon acoustic oscillations) gave rise to the nature of the power spectrum [25]. Initially, only small scales oscillated, with larger and larger scales taking part until recombination occurred. At this time the CMB was frozen with anisotropies corresponding to hotter (low density) and cooler (high density) regions.

In particular, the third peak of the power spectrum contains information on the density of matter and is sensitive to how much of the mass is baryonic. This is due to the fact that even though both standard and dark matter experience the gravitational collapse, dark matter does not interact with light and hence does not experience the light pressure that causes baryon acoustic oscillations. Hence, the shape of the third peak can be used to infer the amount of matter that is non-baryonic in nature. A fixed number of parameters (usually 6 or 7) is used to fit the spectrum with the present mass density of matter (Ω_M) in the Universe extracted. It is essentially impossible to fit the data without the presence of cold dark matter (neutral particles with non-zero mass). The value of the relic density of dark matter is $\Omega_\chi = 0.266$, five times higher than the baryon density, $\Omega_B = 0.049$. The result, indicating that 83% of the matter content of the early Universe is dark, agrees with galaxy and cluster observations [22].

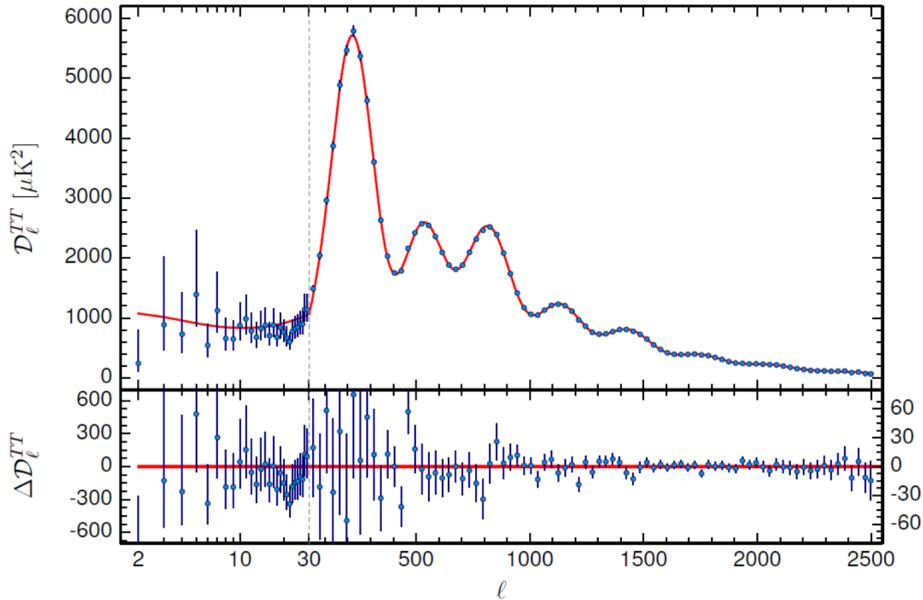


Figure 2.4: The observed CMB temperature anisotropy power spectrum from Planck [22], where ℓ represents the angular scale in the sky (see equation 2.1).

CMB anisotropies eventually resulted in the formation of large scale structure in the Universe as it cooled and expanded further. The precise mechanisms and processes leading

to structure formation from seed inhomogeneities are an active area of research. A common approach involves the use of n -body simulations. Recent studies suggest the existence of a universal dark matter profile, with the same shape for all masses, epochs and input power spectra [26, 27].

The abundance of baryons in the Universe can also be estimated independently from Big Bang Nucleosynthesis (BBN). Long before recombination, BBN occurred, 10 s to 20 minutes after the Big Bang ($T \sim \text{MeV}$), and heavier nuclei such as helium, deuterium and lithium were formed [28, 29]. The measured amount of deuterium suggests a baryon density in agreement with the CMB power spectrum and baryon acoustic oscillation measurements and hence acts as another powerful confirmation for the existence of dark matter [30].

2.2 Dark matter phenomenology

2.2.1 General characteristics

The dark matter observational evidence we discussed, both at cosmological and galactic scales, encodes important information on the nature of dark matter that is essential when constructing viable particle models.

One of the main conclusions on the nature of dark matter is that it should be electrically neutral. If dark matter is coupled to standard model particles this is expected to be via a (sub)weak coupling, either via the weak force or through other portals (e.g. Higgs, dark photon) connecting the visible and dark matter sectors. Dark matter self-interactions are also constrained in strength such that they are compatible with structure formation of galaxy cluster systems. As previously discussed, CMB data confirm that most of dark matter is non-baryonic in nature [24]. Finally, the DM particles should be either stable or extremely long lived. We note that there are both fermionic and bosonic candidates that can be consistent with the astrophysical and cosmological observations.

Apart of the inherent nature of dark matter particles, we also have information on how relativistic they were when they fell out of thermal equilibrium in the early Universe—although there is no particular evidence favouring a specific mass range. Models of hot dark matter, consisting of relativistic particles at the time of decoupling, have been unable to predict the correct structure formation in numerical simulations [27]. An example of a hot dark matter candidate existing in the standard model (SM), the neutrino, was hence ruled out from being the primary dark matter in our Universe. Non-baryonic cold dark matter models assume that DM density perturbations at the time of decoupling provided the early seeds triggering the growth of baryonic matter density perturbations, and thus lead to structure as we know it from large scale structure surveys [27, 31, 32].

These arguments have led to the establishment of a popular paradigm that the dominant fraction of DM is cold, massive, electrically neutral and interacting (sub)weakly with ordinary matter. Such models exist from the $\sim 10^{-10}$ eV scale, with the most popular model in this category being the axion, to WIMPs in the electroweak scale, all the way to primordial black holes in the $\sim 10^{20}$ kg scale.

2.2.2 Weakly Interacting Massive Particles

In thermal dark matter models the particles are thermal relics produced in the early Universe via interactions with the thermal bath. In this paradigm, the final relic density of DM is determined by its abundance at the time of decoupling from the thermal bath. The most widely known such models propose that DM particles reach full equilibrium with the plasma before decoupling, known as thermal freeze-out [33, 34].

In freeze-out models, dark matter particles in the early Universe annihilate producing standard model particles until the time of decoupling, at which point the Universe has expanded and cooled sufficiently for interactions to cease. Thus, the final relic density is determined by the freeze-out abundance and hence the freeze-out temperature. The freeze-out abundance is in this case inversely proportional to the annihilation cross-section: a smaller cross-section means fewer dark matter particles have annihilated into standard model particles before freeze-out occurred, leading to a larger relic density. Given the measured relic density, a limit can be placed on its annihilation cross-section, assuming the thermal freeze-out scenario.

Proposing such a fermionic dark matter particle (the WIMP) with weak scale strength interactions in the ~ 10 GeV – 10 TeV mass range famously produces the correct DM relic abundance observed in today's Universe. The so-called *WIMP miracle* captures the fact that particle physics and cosmological parameters of very different magnitudes conspire to yield a sensible relic abundance in the present era. This prejudice has made WIMPs the most favoured DM candidates for a few decades. It is important to note though that the WIMP miracle assumes a weak interaction strength which is not necessarily required for beyond the standard model theories, such as dark matter particle models. In this thesis we focus on WIMPs and a class of DM models with masses right below the electroweak scale (light dark matter), which we review further in Section 2.4.1.

2.3 Detection channels

Given the vast array of dark matter models and the energy range they span, several methodologies have been developed aiming at dark matter detection, which can be classified in one of three detection channels: indirect detection of dark matter annihilation or decay products, dark

matter production at colliders, and direct detection of DM-nucleus scattering. A schematic of the three detection channels is presented in Figure 2.5.

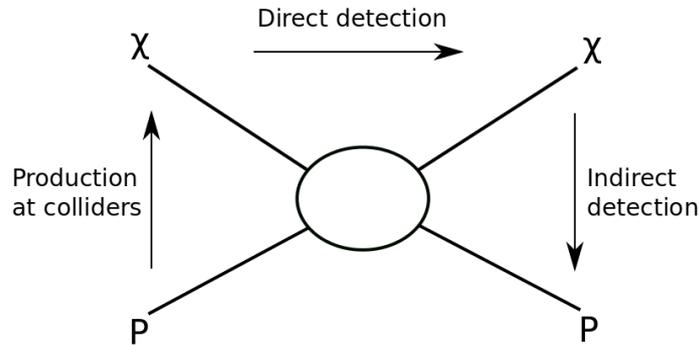


Figure 2.5: Schematic of the three detection channels for dark matter [35].

Particle colliders such as the Large Hadron Collider search for dark matter products following proton-proton collisions ($pp \rightarrow \chi\chi$). These are expected to appear as missing energy, usually in the form of transverse momentum, signalling a dark matter candidate leaving the detector. Missing momentum is also the signature of neutrinos and hence collider searches often focus on events with no leptons [36].

Indirect searches look for annihilation or decay products of dark matter, such as gamma rays ($\chi\chi \rightarrow \gamma\gamma$), electron positron pairs ($\chi\chi \rightarrow e^+e^-$) and neutrino anti-neutrino pairs ($\chi\chi \rightarrow \nu\bar{\nu}$), using either space-based or Earth-based telescopes. These searches often focus in galactic centre searches or dwarf galaxies, where the concentration of dark matter is expected to be higher. Looking for such excesses is challenging and requires a good understanding and modelling of known astrophysical sources of the annihilation products [37].

Direct detection experiments search for the interaction of Milky Way halo dark matter particles with atomic nuclei on Earth [38]. These searches often look for weak-force mediated elastic interactions between dark matter and nuclei, even though alternative interaction models exist. In direct detection detectors the signal from a DM-nucleon interaction can be deposited in three channels: ionisation, scintillation or heat. Different experimental methodologies utilise one or a combination of these response channels, as summarised in Figure 2.6. The biggest challenge for this type of detectors is achieving ultra-low background rates so that detecting a signal of an extraordinarily small rate, such as the one expected from dark matter interactions, becomes possible. Liquid xenon experiments employed deep underground, such as LUX [39], Panda-X-II [40] and Xenon-1T [41], have placed or are expected to place the most stringent constraints on the scattering cross-section of WIMP-nucleon scattering above a few-GeV. A more detailed discussion on noble liquid detectors can be found in Chapter 3.

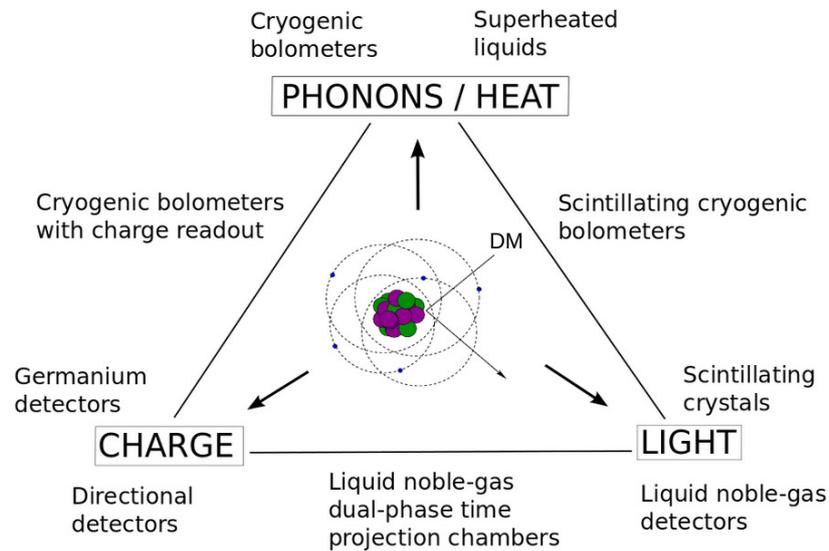


Figure 2.6: Dark matter direct detection experiments depending on the technology in use and the response channels they utilise. Adapted from [35].

2.4 Rare event searches at low energies

2.4.1 Light dark matter

Considering the absence of a dark matter signal at the electroweak scale to date, there has been an increasing interest in DM particles in the few-GeV and below mass range, especially from direct detection experiments. Although some detector technologies are well placed to address this mass range, especially through the use of lighter target elements, liquid xenon time projection chambers can also probe part of this mass region, through the implementation of low energy sensitivity techniques as we discuss further in Section 3.7. As mentioned in Section 2.2, the MeV to GeV mass range is known as the light dark matter (LDM) region [34].

This region is favoured by several models, with the main discriminating feature being their production mechanisms in the early Universe: from thermal freeze-out and freeze-in to completely decoupled dark matter and asymmetric dark matter [33, 42, 43, 44]. We set out to review these models, placing our focus on one of the most crucial aspects in constructing a concrete DM candidate, that is whether it can be produced at the correct abundance in the early Universe.

Freeze-out of thermal light dark matter

As discussed above, freeze-out of thermal dark matter for a particle at the electroweak scale, with weak force interactions, yields an annihilation cross-section that agrees with the current observed value. Masses below a few GeV are forbidden by the Lee-Weinberg limit. Essentially, lighter particles result in a smaller annihilation cross-section and hence a higher relic abundance that would over-close the Universe [45]. This predicament can be avoided if the dark matter particle was a stable scalar, interacting with SM through exchanges of a new neutral gauge boson [46]. In this case, masses above ~ 100 MeV can result in correct cosmological conditions, with scalar dark matter models shown to naturally emerge as a result of gauge-mediated supersymmetry breaking [47, 48]. More complicated cases, in which DM is a combination of two neutral and stable particles, a heavy fermion and a light MeV scale scalar particle, have also been successful in predicting the correct relic abundance [49]. Light scalar dark matter models have attracted significant attention following the realisation that they might explain the 511 keV emission line which has been detected from the bulge of our galaxy, as annihilations of light dark matter particles into e^+e^- pairs [49, 50, 51, 52]. More recently though, it has been shown that this observation could be explained in terms of classical sources of positrons [53].

Freeze-in of thermal dark matter

Concrete thermal dark matter models can also be constructed considering freeze-in of DM. Freeze-in refers to the case in which the rates of matter anti-matter annihilation to form DM are not large enough, so the DM never reaches full equilibrium before decoupling [34]. In this case dark matter consists of Feebly Interacting Massive Particles (FIMPs), which are very weakly coupled to the thermal bath [33, 54, 55]. In the beginning, and at very high temperatures, negligible initial dark matter abundance is assumed while dark matter particles are produced from collisions or decays of SM particles. This production is exponentially suppressed with decreasing temperature. Eventually, decoupling occurs and the final relic abundance is set. In this case, contrary to the freeze-out scenario, a smaller interaction cross-section means a smaller final abundance and hence lower mass particles that would not over-close the Universe are allowed. The evolution of the DM abundance as a function of the temperature of the Universe (represented as $x = m_\chi/T$) is shown in Figure 2.7 for a 100 GeV WIMP, both for the freeze-in and freeze-out cases and for three velocity averaged interaction cross-sections. The observed Planck relic abundance is marked with a horizontal line and can hence be predicted by a velocity averaged interaction cross section, $\langle\sigma v\rangle$, of $10^{-47} \text{ cm}^3\text{s}^{-1}$ for freeze-in and $2.05 \times 10^{-26} \text{ cm}^3\text{s}^{-1}$ for freeze-out, for the 100 GeV WIMP mass assumed.

Several candidates for FIMP dark matter have been proposed, from moduli and modulinos to dirac neutrino masses and other hidden sector models. Most of these models predict that the

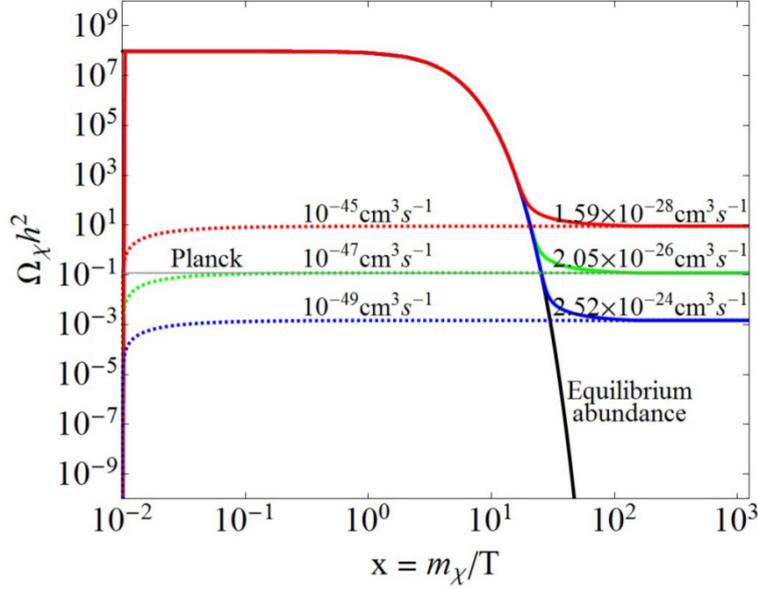


Figure 2.7: Evolution of dark matter abundance as a function of the Universe temperature for three different interaction cross-sections shown in solid lines for the freeze-out scenario and in dashed lines for freeze-in of thermal dark matter [33]. The velocity averaged interaction cross-section for each of the six sub-cases is also presented in the Figure. The final relic density increases for increasing annihilation cross-section in the case of freeze-in while the opposite happens for freeze-out.

FIMP is the direct decay particle of a SM fermion or that a Lightest Ordinary Supersymmetric Particle (LOSP) exists to which FIMPs decay, as shown in Figure 2.8. In the first case the relic abundance is proportional to the dark matter mass, while in the second it is inversely proportional [54].

Non-thermal dark matter

For very small interaction cross-sections the DM particles are produced through the inflaton field, already decoupled from the thermal bath and hence the final abundance is completely determined by the initial one, which in turn depends on the inflaton branching fraction and mass, the DM mass and the reheat temperature [33, 34]. Even though some constraints on the branching fraction exist in order to have a radiation dominated epoch immediately after reheating, followed by matter domination, the constraints on the DM mass are lenient, allowing DM masses from $10^{-4} - 10^{13}$ GeV. A map of the DM annihilation cross-section constraints for thermal and non-thermal DM is shown in Figure 2.9.

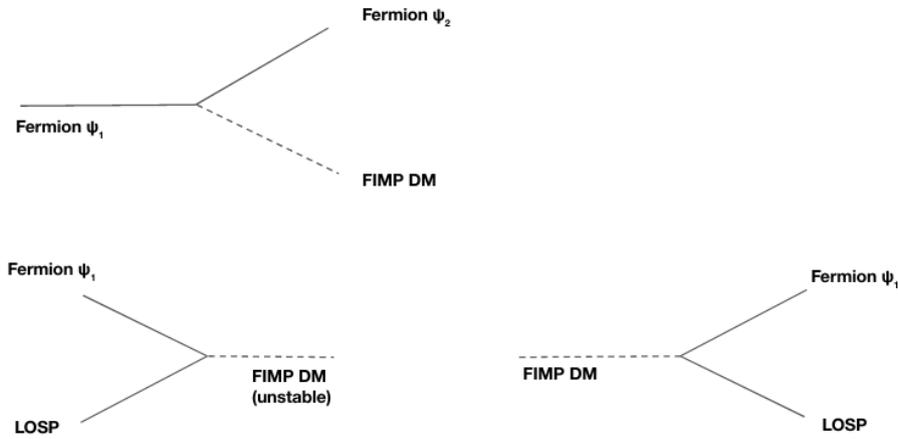


Figure 2.8: FIMP and LOSP Feynman diagram vertices, for the production process of dark matter in the early Universe.

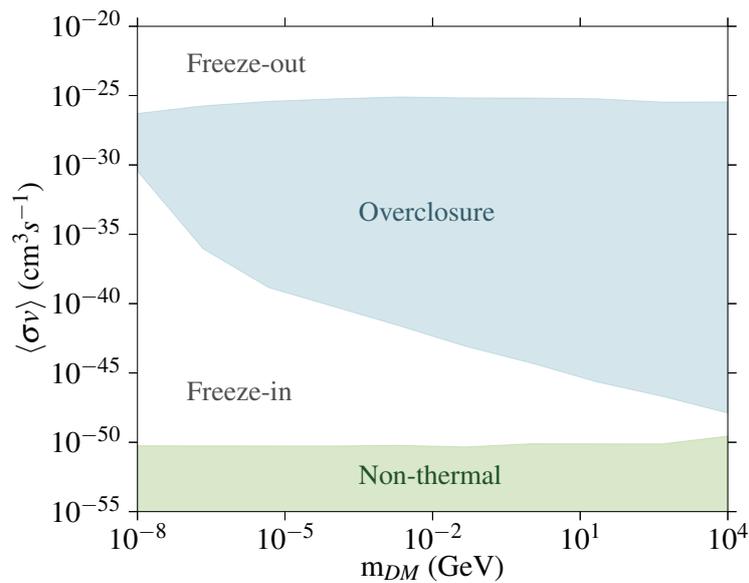


Figure 2.9: Annihilation cross-section as a function of DM mass for both thermal and non-thermal production mechanisms. The blue-shaded region is excluded from relic density constraints, while the white regions represent the thermal cases of freeze-in and freeze-out. The green-shaded region shows the non-thermal dark matter allowed range. The lower bound of the over-closure region and upper bound of the non-thermal region depend on the reheat temperature, inflaton mass and branching fraction and here we show only one representative scenario. Figure adapted from [33].

SIMPs and ELDERS

Usually, when discussing a DM candidate created as a thermal relic from freeze-out, we assume a relic abundance set by $2 \rightarrow 2$ annihilations of DM into SM particles. As previously

mentioned, dark matter particles of mass smaller than a few GeV are prohibited under these assumptions, as they result in a relic abundance that would have led to over-closure of the Universe. A way to overcome this problem is to propose that the dark matter relic abundance is set thermally by the freeze-out of $3 \rightarrow 2$ self annihilations. This class of models introduces the idea of Strongly Interacting Massive Particles (SIMPs), in the MeV to GeV mass range, with strong self-interactions and subdominant annihilations into SM particles [56, 57]. SIMPs are in kinetic equilibrium with the thermal bath but exist in a dark sector, that is only weakly coupled to the visible sector.

A special case of SIMPs, known as Elastic Coupling Relic (ELDER) dark matter, combines both the weak scale freeze-out scenario and the SIMP paradigm [34, 58, 59]. In this case, the dark matter particles are thermalised with standard model particles initially. Because of the lower mass of DM in this model, the relic abundance immediately after the ordinary freeze-out phase is too high. A self heating period follows during which the SIMPs, which have now decoupled from SM, interact with-themselves in a $3 \rightarrow 2$ process, restoring the correct abundance. Finally, as the temperature further decreases they become completely decoupled and are too dilute to self-interact.

Asymmetric dark matter

The WIMP miracle has been one of the biggest motivators for thermal dark matter. It does nevertheless suffer from an important drawback, that is often overlooked: it ignores the fact that the relic densities of dark matter and baryonic matter today are very similar ($\Omega_{DM} \approx 5\Omega_B$). The baryonic matter in our Universe is known to have been generated by an asymmetry between baryons and anti-baryons and not by a thermal freeze-out. There is no reason why a similar — or even the same — production mechanism did not apply for DM production. On the contrary, the fact that a baryon - anti-baryon asymmetry in the order of one-in-a-billion created a relic abundance of the same order of magnitude as the dark matter abundance motives a common origin and cannot easily be overlooked as a coincidence. Assuming that an asymmetry was the mechanism behind dark matter production gives rise to the idea of asymmetric dark matter (ADM) [34, 43, 44].

Most ADM models exist in a hidden dark sector. A special example is mirror dark matter, that proposes a dark sector isomorphic to the visible sector [60, 61]. The two sectors interact gravitationally through photon mirror-photon kinetic mixing and Higgs mirror-Higgs mixing. The photon kinetic mixing interaction can be observed in direct detection experiments [62] while the Higgs mixing is accessible by collider searches.

2.4.2 Coherent elastic neutrino-nucleus scattering of solar, atmospheric and supernova neutrinos

There also exist other types of rare processes creating low energy interactions detectable by the same detector technologies as those for low mass WIMPs. Coherent elastic neutrino-nucleus scattering (CE ν NS) is the process through which a neutrino interacts coherently with a nucleus through a neutral current interaction and the nucleus is not excited to a higher energy state [63]. Since CE ν NS is a neutral current interaction, all neutrino flavours are expected to participate with the same cross-section. Additionally, this process is concretely predicted by the standard model. Any deviation to the cross-section could be a signature of beyond the standard model physics. Specifically, the cross-section can be used as a complementary approach in constraining the weak mixing angle, as this would change the weak nuclear charge, on which the magnitude of the recoil spectrum depends (see Section 3.2). Moreover, CE ν NS measurements provide a probe into non-standard interactions of neutrinos with quarks that would result in anomalous interactions rates. An example of this would be the vector couplings for ν_e with up and down quarks, that are currently not highly constrained. Constraints on the neutrino magnetic moment can also be obtained using CE ν NS since for a magnetic moment a term is added to the cross-section, producing a measurable effect on the recoil spectrum at low energies [64].

Desirable neutrino sources would benefit from a high flux and well understood spectra, as well as being multi-flavor, to allow testing for new physics. Artificial candidates include reactors, as they benefit from enormous fluxes of $\bar{\nu}_e$. This range up to energies of only ~ 7 MeV and coupled with the significant backgrounds make an experimental observation a challenging process [65, 66]. A more interesting source for CE ν NS detection is a stopped-pion source. In this type of source, a dense target is irradiated with a proton beam, to cause the resulting pions to decay at rest producing ν_μ , ν_e and $\bar{\nu}_\mu$ with well understood spectra and in an energy range (0–50 MeV) that would cause detectable coherent interactions. After several decades of experimental effort by several teams, the COHERENT collaboration has finally observed CE ν NS using neutrinos produced at the Spallation Neutron Source (SNS) at ORNL [67].

Apart from Earth-based artificial sources, there exist many other natural contexts in which CE ν NS is expected to occur, including solar, atmospheric, and supernova neutrino detection and searches for sterile neutrinos. A nearby supernova star would create an intense burst of neutrinos, accessible by dark matter direct detection experiments, with the magnitude of CE ν NS encoding interesting information on the core explosion mechanism, and especially the role of neutrinos in the process. An experiment such as LZ could register up to 100 CE ν NS events from a supernova 10 kpc away [68]. The signals registered in dark matter detectors searching for WIMPs, can be used to constrain the explosion energy and reconstruct the supernova neutrino light curve, with sensitivities reaching competitive levels compared to

dedicated neutrino telescopes, while at the same time providing complementary information inaccessible by these technologies [69]. Even though all types of neutrinos emerge from the supernova core, matter enhanced flavor transformation and effects such as neutrino self-interaction could mix the spectra. A combination of charged and neutral current measurements can be used to disentangle the flavour transformation effects from the original neutrino signal [70, 71, 72].

The insensitivity of $CE\nu NS$ to active flavour transformation is also ideal for studying sterile neutrino oscillations using a disappearance experiment [73, 74]. More specifically, sterile neutrinos do not interact with a nucleus through the regular weak interaction and hence show up as a decrease in the $CE\nu NS$ rate. If such a decrease is observed in this neutral current process, it can be inferred that this is not due to flavor mixing, to which these interactions are insensitive. Hence, $CE\nu NS$ can become a probe to parameter-space regions of current interest for sterile neutrinos.

In addition, the most prominent neutrino backgrounds for WIMP searches in near-future experiments are expected to come from $CE\nu NS$ of solar neutrinos [75]. The solar neutrino production chain is shown in Figure 2.10. In particular, direct detection experiments will be sensitive to two sources of coherent neutrino-nuclear scattering of solar neutrinos, namely ${}^8\text{B}$ and *hep* neutrinos. Specifically, the neutrinos from the ${}^8\text{B}$ decay in the ppIII branch have just about sufficient energy to produce events detectable in near-future direct detection experiments while testing specific WIMP mass hypotheses. This will be an irreducible background that needs to be subtracted from the spectra of future experiments. For example, the ${}^8\text{B}$ solar neutrinos induce a differential spectrum equivalent to a WIMP of ~ 5.5 GeV with a spin-independent cross-section of $\sim 10^{-45}$ cm² [75]. The measurement of ${}^8\text{B}$ neutrinos through $CE\nu NS$ is sensitive to all neutrino flavours forming an interesting result on its own, apart from a background to these experiments. A clear first detection of this flux would also demonstrate the sensitivity to supernova neutrinos, which produce a similar spectrum (but lasting only for a fraction of a second).

Atmospheric neutrinos result from muon and pion decays in the atmosphere. Even though $CE\nu NS$ of atmospheric neutrinos is expected to produce O(1) event in near-future direct detection experiments, its recoil spectrum is indistinguishable from that of a typical WIMP and will hence need to be studied extensively for next generation experiments [77]. The atmospheric neutrino background induces a differential spectrum that is similar to a ~ 100 GeV WIMP with a spin-independent cross-section of $\sim 10^{-48}$ cm².

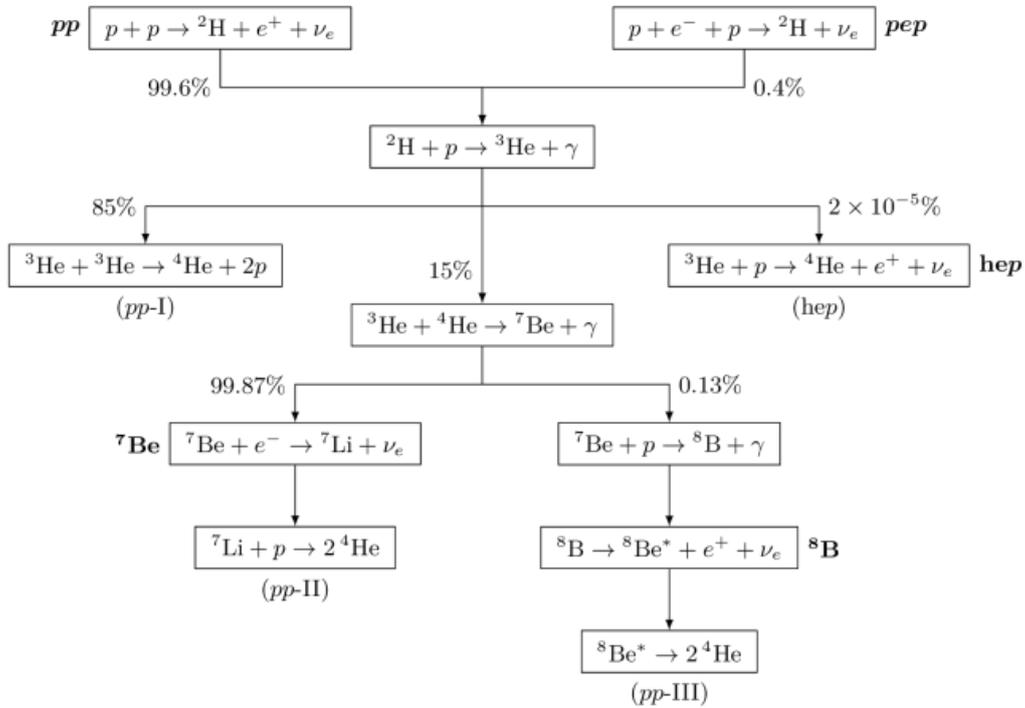


Figure 2.10: Solar neutrino production chain [76].

2.5 Summary

In this chapter we summarised the main cosmological and galactic evidence for dark matter and reviewed the insight they give us into the nature of this elusive particle. We also presented the detection channels probing different regions of the dark matter phase space, including direct detection methods, relevant to the work in this thesis. Finally, we motivated the existence of several dark matter models in the few-GeV and sub-GeV mass scale as well as other physical processes of interest involving the scattering of neutrinos, both producing very soft nuclear recoil spectra which are central to this thesis.

Chapter 3

Particle detection in LXe-TPCs

We have previously motivated the existence of particle dark matter and introduced some of the leading candidate models, and went on to review experimental searches dedicated to observing and understanding dark matter interactions. In this chapter, we focus on the direct (scattering) detection of WIMPs using dual-phase xenon detectors.

We begin by introducing the theoretical framework behind WIMP-nucleon scattering, before focusing our discussion on nuclear and electron recoil interactions in LXe. The physics of particle interactions in this medium leads to a discussion of how signal generation is accomplished in dual-phase (liquid/gas) time projection chambers and their sensitivity for WIMP searches. An exploration of the low-energy sensitivity of direct detection experiments is also included in this chapter. Finally, we introduce the main components of the LUX and LZ experiments.

3.1 WIMP-nucleon scattering

3.1.1 Differential event rate

A plethora of experiments aim at the direct detection of WIMPs by searching for signals from dark matter particles scattering off Earth-based detectors. A canonical model in which WIMPs in our galaxy are gathered in a spherical isothermal halo (standard halo model — SHM), much larger than the visible baryonic disc, is usually adopted, even though more complicated models exist that provide better agreement with galactic n-body simulations. In particular, recent data for the Milky Way suggest that a substantial strongly radially anisotropic DM population exists in combination with the spherical isotropic halo, resulting in an overall speed distribution somewhat discrepant with the SHM [78, 79].

For the purposes of this thesis we adopt the SHM approximation, in which the dark matter in our galaxy is arranged in an isothermal, isotropic halo with a dark matter velocity

distribution ($f(\mathbf{v}, \mathbf{v}_E)$) following a Maxwell-Boltzmann distribution truncated at the galactic escape velocity ($v_{esc} = 544$ km/s [80]), above which WIMPs are no longer gravitationally bound to the Milky Way. In the detector (Earth) rest frame, the dark matter velocity distribution depends on the velocity of the Earth relative to the galactic rest frame (\mathbf{v}_E), which is given by the sum of three components: the circular velocity with which the solar system moves through the halo ($v_o = 230$ km/s), the peculiar velocity of the Sun, and the Earth's orbital velocity around the Sun. The average value of the latter component is $v_{orb} = 30$ km/s with the projection along the direction of motion varying (sinusoidally) throughout the year. The SHM approximation we adopt results in only a small bias on our final results. The differential event rate, in the detector reference frame, is then given by:

$$\frac{dR}{dE_R} = \frac{\rho_o \sigma_A}{2m_\chi \mu_A^2} F^2(E_R) \int_{v_{min}}^{v_{esc} + v_E} \frac{f(\mathbf{v}, \mathbf{v}_E)}{|\mathbf{v}|} d^3v \quad , \quad (3.1)$$

where σ_A is the elastic WIMP-nucleus scattering cross-section, m_χ is the WIMP mass, μ_A is the WIMP-nucleus reduced mass, and $F^2(E_R)$ is the nuclear form factor [81, 82, 83]. The lower limit of integration corresponds to the lowest speed that can cause a recoil of energy E_R :

$$v_{min} = \sqrt{\frac{m_A E_R}{2\mu_A^2}} \quad . \quad (3.2)$$

The recoil energy is given by [84]:

$$E_R = \frac{1}{2} m_\chi v^2 \frac{4m_\chi m_A}{(m_\chi + m_A)^2} \frac{1 - \cos \Theta}{2} \quad , \quad (3.3)$$

where Θ is the scattering angle in the centre-of-mass frame, related to the lab scattering angle, θ , through:

$$\cos \Theta = 1 - 2 \cos^2 \theta \quad . \quad (3.4)$$

Evidently, the maximum energy transfer occurs for backscattering ($\Theta = \pi$) and when the target and dark matter mass are similar.

The differential spectrum is a decreasing exponential with decay constant determined by the WIMP mass, and it can be integrated to predict the number of events to be observed in an experiment with exposure $e = \text{live-time} \times \text{target mass}$, an energy-dependent efficiency (ϵ), and an energy region of interest between E_{thr} and E_{max} :

$$N = \int_{E_{thr}}^{E_{max}} e \epsilon \frac{dR}{dE_R} dE_R \quad . \quad (3.5)$$

The differential rate in equation 3.1, and hence the direct detection of WIMPs, depends not only on the properties of the dark matter halo, but also on the properties of the dark matter particles themselves, as well as the nuclear physics of the target medium. In particular, the momentum transfer (q) dependence of the WIMP-nucleus interaction rates is encoded within the nuclear form factor, $F^2(q)$, which represents a correction which accounts for the finite size of the nucleus. Usually, the nucleus is approximated as a solid sphere of radius r_n with a smooth nucleon density transitioning to zero over a nuclear skin thickness s , with a particularly popular expression due to Helm [85]:

$$F_{Helm}^2(q) = \left(\frac{3j_1(qr_n)}{qr_n} \right)^2 e^{(-qs)^2} \quad , \quad (3.6)$$

where j_1 is a spherical Bessel function. The form factor is normalised to unity at zero momentum transfer, $F^2(0) = 1$.

It is also interesting to consider how the lowest WIMP mass that is observable by a direct detection experiment depends on the target nucleus and energy threshold of the detector. For low WIMP mass we may approximate $\mu_A \sim m_\chi$, and hence the minimum velocity that can cause a recoil at threshold becomes:

$$v_{min} = \sqrt{\frac{m_A E_{thr}}{2m_\chi^2}} \quad . \quad (3.7)$$

This means that we are only probing the upper tail of the WIMP velocity distribution function for a given DM mass. The elastic WIMP-nucleus scattering cross-section (σ_A) is inversely proportional to the velocity distribution integral, which diverges when v_{min} approaches $v_{esc} + v_E$ (equation 3.1). Thus, the combination of an experimental energy threshold with the fact that DM particles have a maximum velocity in the halo imposes a limit to the lowest WIMP mass that experiments are sensitive to:

$$m_\chi > \sqrt{\frac{m_A E_{thr}}{2(v_{esc} + v_E)^2}} \quad . \quad (3.8)$$

For example, for an experiment utilising xenon ($m_A = 122.2$ GeV) and with a "hard" energy threshold at O(1 keV), the lowest detectable WIMP mass is ~ 3.3 GeV.

3.1.2 Spin-dependent and spin-independent interactions

The WIMP-nucleus differential scattering rate in equation 3.1 can also be expressed as:

$$\frac{dR}{dE_R} = \frac{\rho_o}{m_\chi m_A} \int_{v_{min}}^{v_{esc} + v_E} v f(\mathbf{v}, \mathbf{v}_E) \frac{d\sigma_A}{dE_R} d^3v \quad . \quad (3.9)$$

The WIMP-nucleus differential cross-section, $d\sigma_A/dE_R$ is typically divided into two components:

$$\frac{d\sigma_A}{dE_R} = \frac{m_A}{2v^2\mu_A^2} (\sigma_A^{SI} F_{SI}^2 + \sigma_A^{SD} F_{SD}^2) \quad , \quad (3.10)$$

where σ_A^{SD} and σ_A^{SI} represent the zero momentum transfer terms for spin-dependent and spin-independent interactions. Writing the differential WIMP-nucleus cross-section in terms of the WIMP-nucleon cross-section is useful in comparing different target materials or constraints from collider results. The spin-independent term describes scalar ($s \sim \bar{\chi}\chi\bar{q}q$) and vector ($v \sim \bar{\chi}\gamma^\mu\chi\bar{q}\gamma^\mu q$) coupling interactions, while the spin-dependent term represents the axial-vector ($v \sim \bar{\chi}\gamma^\mu\gamma_5\chi\bar{q}\gamma^\mu\gamma_5 q$) component [86].

At low momentum transfer, all partial waves from the scattering off individual nucleons participate in the interaction, while for higher recoil energies the WIMP does not scatter off the whole nucleus coherently. The coherence effect is described through the energy-dependent form factors F_{SI} and F_{SD} [87], with the Helm factor that was previously defined often used for the SI case. The nuclear form factor incorporates the relationship between the momentum of the incoming particle or the momentum transfer through the mediator and the size of the target nucleus. When the incoming particle's de Broglie wavelength ($\lambda = h/p$) is large compared to the size of the nucleus, the particle scatters coherently with the entirety of the nucleus, leading to an enhancement in the interaction rate [87]. For higher particle energies, when the wavelength becomes comparable to the size of the nucleus, coherence is lost [84].

The differing coherence enhancement in the scattering rate between spin-dependent and spin-independent interactions is encoded within the WIMP-nucleus cross-section at zero momentum transfer (also known as the "standard" cross-section, σ_A). The spin-independent WIMP-nucleon cross-section relates to the WIMP-nucleus cross-section through [87]:

$$\sigma_A^{SI} = \sigma_n \frac{\mu_A^2}{\mu_n^2} A^2 \quad , \quad (3.11)$$

where μ_A is the reduced mass of the WIMP-nucleus system, μ_n is the WIMP-nucleon reduced mass and σ_n the WIMP-nucleon cross-section. The A^2 enhancement of the rate means that experiments utilising heavier targets are expected to observe substantially higher event rates. At the same time, the recoil energies produced for heavier elements are necessarily lower, posing challenges for light WIMP sensitivity. The SI event rate for three different target media (Xe, Ge and Ar) at a cross-section of $1.0 \times 10^{-45} \text{ cm}^2$ is shown in Figure 3.1 for 6 and 60 GeV WIMPs, respectively. The effect of the nuclear form factor is shown in the same figure — and it is negligible for the WIMP masses of relevance in this thesis.

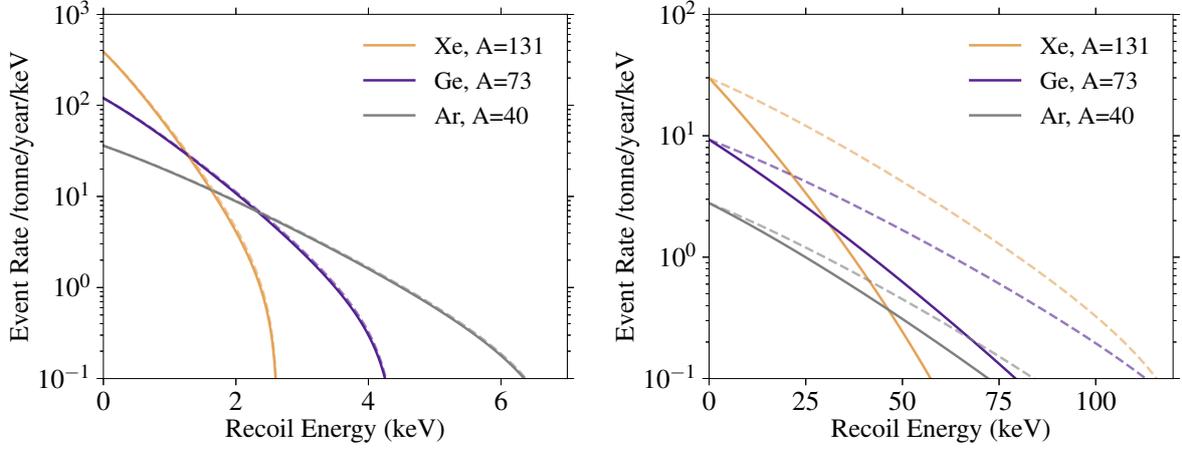


Figure 3.1: Differential recoil energy spectra for three different targets (Xe, Ge and Ar) and for (*left*) a 6 GeV and (*right*) a 60 GeV WIMP at a cross-section per nucleon of $1.0 \times 10^{-45} \text{ cm}^2$. The dashed lines present the event rates for a unity form factor.

For spin-dependent interactions the cross-section is given by [86]:

$$\sigma_A^{SD} = 8G_F^2 \mu_A^2 \frac{J+1}{J} [\alpha_p \langle S_p \rangle + \alpha_n \langle S_n \rangle]^2, \quad (3.12)$$

where α_p and α_n represent the spin-dependent couplings to individual protons and neutrons, J is the total nuclear spin, and $\langle S_p \rangle$ and $\langle S_n \rangle$ represent the expectation values for proton and neutron spins in the nucleus, respectively [88], and G_F is the Fermi constant. In simplified terms, the spin-dependent scattering happens with essentially a single (unpaired) nucleon, while the spin contribution from the proton or neutron pairs cancels out [84, 89]. In reality, more sophisticated nuclear shell model calculations have highlighted that contributions from 2-body currents become significant. In xenon — which has odd-neutron isotopes but even proton number — this confers sensitivity also to spin dependent WIMP-proton interactions.

In this section we introduced the two main types of interaction between WIMPs and nuclei: SI and SD. These are in fact special cases of the possible elastic DM-nucleus response, with a complete set of response functions described by non-relativistic effective field theory (EFT) — through the use of a full set of more fundamental EFT operators that mediate the DM-ordinary matter interaction [90]. In this thesis we pursue only the calculation of the SI interaction.

3.2 Coherent elastic neutrino-nucleus scattering

As discussed in Section 2.4.2, CE ν NS of solar, atmospheric and supernova neutrinos is another rare process creating low energy nuclear recoils, one which may can act as a background to

dark matter searches — but is also of physical interest itself. This neutrino interaction has now been detected for the first time with beam neutrinos in CsI [91] and, very recently, in LAr [67] by the COHERENT collaboration. The coherent neutrino-nucleus scattering cross-section per unit recoil energy E_R is given by:

$$\frac{d\sigma(E_\nu, E_R)}{dE_R} = \frac{G_F^2}{4\pi} Q_W^2 m_A \left(1 - \frac{m_A E_R}{2E_\nu^2}\right) F^2(E_R) \quad , \quad (3.13)$$

where $Q_W = N - (1 - 4 \sin^2 \theta_W)Z$ is the weak charge, θ_W is the weak mixing angle and E_ν is the neutrino energy [85]. The differential event rate is then given by:

$$\frac{dR}{dE_R} = \int_{E_{min}}^{\infty} \frac{dN}{dE_\nu} \frac{d\sigma(E_\nu, E_R)}{dE_R} dE_\nu \quad , \quad (3.14)$$

where $\frac{dN}{dE_\nu}$ is the differential neutrino flux and $E_{min} = \sqrt{m_A E_R / 2}$ is the minimum neutrino energy that can produce a nuclear recoil at energy E_R [75]. The neutrino spectra, and corresponding nuclear recoil spectra, for the prominent neutrino background processes for a direct detection experiment with a xenon target are shown in Figure 3.2.

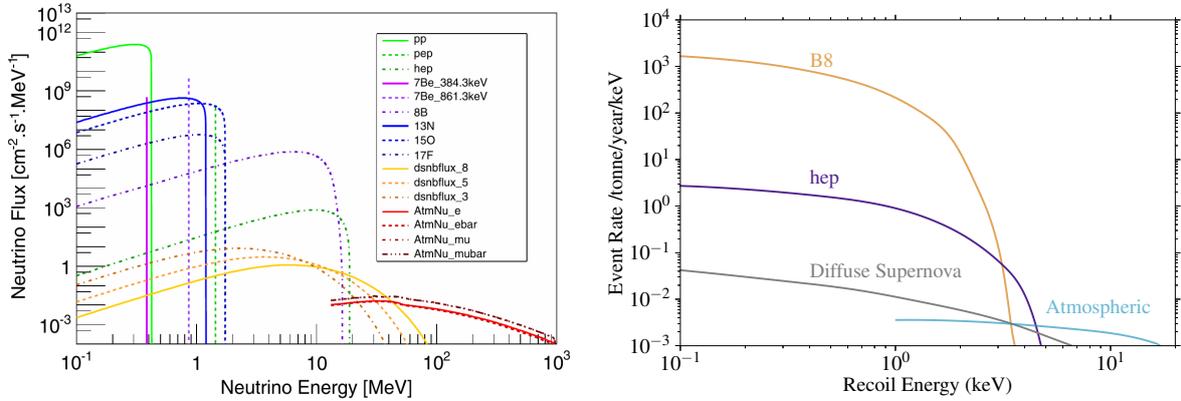


Figure 3.2: Differential spectra for the (*left*) neutrino flux [75] and (*right*) the event rate for coherent neutrino-nucleus elastic scattering, for the most prominent neutrino fluxes for future direct detection experiments with a xenon target. These include ${}^8\text{B}$ and *hep* solar neutrinos, diffuse supernova neutrinos and atmospheric neutrinos that have been discussed in Section 2.4.2.

Similarly to WIMP scattering, $\text{CE}\nu\text{NS}$ can occur when the momentum transfer to the nucleus is smaller than the inverse nuclear radius; in xenon, coherent scattering leads to nuclear recoils in the keV range for neutrinos with energies $\lesssim 50$ MeV. The cross-section remains highly enhanced by an $\sim N^2$ factor in this region (since $Q_W \sim N$), while for recoil energies above 10 keV the loss of coherence reduces the neutrino-nucleus cross-section very dramatically. At the same time, the recoil energy has to be high enough to fall above the threshold of the detector, creating a range of energies in which $\text{CE}\nu\text{NS}$ is an observable process by current detection technologies. $\text{CE}\nu\text{NS}$ signals for ${}^8\text{B}$ and *hep* neutrinos will be accessible

to the new generation of experiments soon to come online (such as LZ) due to this major enhancement effect.

3.3 Signal production in liquid xenon

3.3.1 Properties of liquid xenon as a detector medium

The high atomic mass of Xe is of benefit to the signal rate for spin independent dark matter models, as discussed previously. Natural Xe contains many isotopes at non-negligible abundance (^{124}Xe , ^{126}Xe , ^{128}Xe , ^{129}Xe , ^{130}Xe , ^{131}Xe , ^{132}Xe , ^{134}Xe , ^{136}Xe); the two odd-neutron isotopes, ^{129}Xe and ^{131}Xe , which together constitute 47.5% of natural Xe, make it a leading medium for SD WIMP-neutron scattering, and WIMP-proton interactions are enhanced by the 2-body currents mentioned previously to make this a competitive channel too [92, 93, 94].

Some of the other Xe isotopes allow detectors conventionally built for WIMP direct detection to probe competitively into several rare event processes of interest. An example would be the search for neutrinoless and two-neutrino double beta decay of ^{134}Xe and ^{136}Xe . The two-neutrino process has already been observed in ^{136}Xe [95, 96], while there are excellent prospects for the search for neutrinoless double beta decay using LXe in experiments such as nEXO [97] and, in fact, in dual-phase systems such as LZ [98]. This is a physical process of great interest, with many experiments dedicated to its search [99, 100, 101]; if the decay is followed by the emission of zero neutrinos, then neutrinos are Majorana particles and lepton number conservation is violated [102]. This is also true for double electron capture in ^{124}Xe , the two-neutrino mode of which has recently been detected by the Xenon collaboration [103].

Liquid xenon has long been used as a radiation detection medium in particle physics, astrophysics and medical imaging experiments. There are several characteristics of LXe that have made it such an interesting target for so many applications, and specifically for rare event searches — despite the low temperatures required, as xenon is only liquid over a narrow temperature range around around -100°C .

The sensitivity of rare event searches is highly dependent on the experimental achievement of very low background rates. LXe has the highest stopping power (dE/dx) among the liquid rare gases, thus blocking the most penetrating radiation effectively within a few cm (neutrons, MeV gamma-rays). This is a result of its high atomic number ($Z = 54$) and high density ($\sim 3 \text{ g/cm}^3$) [104, 105]. At the same time there is no intrinsic contamination from long-lived radioisotopes.¹

¹An exception is the two-neutrino double-beta decay of ^{136}Xe with $T_{1/2} = 2.165 \times 10^{21}$ years and $Q_{\beta\beta} = 2.458 \text{ keV}$, but this does not pose a problem for dark matter searches with natural abundance xenon.

3.3.2 Scintillation and ionisation from particle interactions in liquid xenon

Signal production in LXe begins with a collision between an incoming particle and the xenon nucleus (nuclear recoil, NR) or an atomic electron (electron recoil, ER). The scattering results in an energy deposition (E_{dep}) translated into excitation, ionisation and heat signals [106], with the energy partition given by [107]:

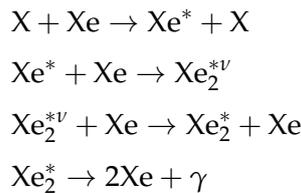
$$E_{dep} = N_i W_i + N_{ex} W_{ex} + W_{heat} \quad , \quad (3.15)$$

where N_i and W_i are the mean number of ionised atoms created per deposition and the mean energy required to produce an electron-ion pair (the so-called "W-value" for ionisation), and N_{ex} represents the mean number of excited atoms at an average energy expended of W_{ex} . We begin by describing the interactions of electron recoils. In this case, omitting the heat deposition term (W_{heat}), which is negligible for ER interactions, and denoting the ratio of excitons to ions as $\alpha = N_{ex} / N_i$, yields [106, 107]:

$$E_{dep} = N_i (\alpha W_{ex} + W_i) \quad . \quad (3.16)$$

For NR interactions, the particle-nucleus collision causes a cascade of secondary recoils with neighbouring xenon atoms. For this type of interaction, energy is deposited into the heat channel but LXe detectors of interest to this thesis are unable to detect it. Hence, the two observable signal channels are due to scintillation (due to the exciton creation) and ionisation (due to the ion creation).

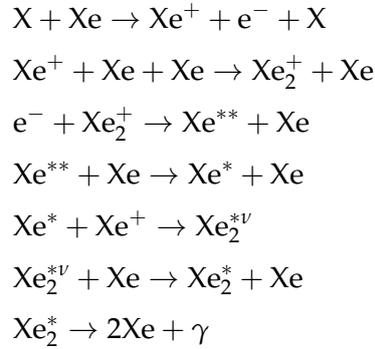
We begin by considering the scintillation signal production mechanism in LXe. The scattering of an electron or a nucleus (denoted by X) off a xenon atom causes a recoil, which in turn may result in the formation of a large number of excited atoms (Xe^*) as the recoil loses energy along its track. Each of these excited atoms collides with neutral atoms to form an excited diatomic molecule (Xe_2^{*v}) [106]. The excited dimer relaxes through additional collisions, before decaying to a dissociative ground state, emitting a 175 nm scintillation photon (γ) [108]. This process can be summarised in the following steps [107]:



The superscripts * and v indicate electronically and vibrationally excited states, respectively. The vacuum ultraviolet (VUV) scintillation is a result of the decay of the Xe_2 molecule, which has singlet and triplet state configurations. These generate fast and slow decay components of

the scintillation light with decay times of 4.3 ns and 22 ns, respectively [107]. It should be noted that the ground state of the molecular dimer is only weakly bound and for much larger atomic distances than the excited dimer (hence termed "dissociative"), which prevents the 175 nm scintillation photons from being reabsorbed. This is an essential property of a good scintillator.

Recoiling nuclei or electrons also produce ions (Xe^+) along their track, releasing free electrons that can carry the ionisation signal. However, additional scintillation can also be produced through recombination of ionisation electrons, which consequently leads to Xe_2^* production and decay. These are the relevant processes following ionisation events [106]:



The ** superscript here denotes an atom in a highly-excited electronic state. Even though the de-excitation of Xe_2^* is the same as in the first process leading to scintillation, the recombination process introduces a significant time delay (~ 45 ns).

Under an applied electric field, ionisation electrons that are not captured by positive Xe ions will escape the interaction as free charge carriers. A small recombination fraction (r) results in a large number of electrons escaping the interaction and hence a higher ionisation signal accompanied by a lower scintillation signal. This is the reason behind the anticorrelation observed between the scintillation and ionisation signals in LXe-TPCs. The number of photons (n_γ) and of electrons (n_e) escaping the interaction site can hence be expressed as [106]:

$$n_\gamma = N_{ex} + rN_i \quad (3.17)$$

$$n_e = (1 - r)N_i \quad (3.18)$$

The recombination fraction is smaller for low recoil energies and reduces with the application of an electric field — and it depends also on the type of recoil. When a field is applied, ionisation electrons are removed from the recoil track and hence their recombination probability is reduced [109]. In any case, the total number of quanta produced is equal to the total number of ions and excitons generated originally, for all values of r [106]:

$$N_i + N_{ex} = n_\gamma + n_e \quad (3.19)$$

This can then be used to calculate the recoil energy (E) of an ER interaction using:

$$E = W(n_\gamma + n_e) \quad , \quad (3.20)$$

where $W = (13.7 \pm 0.2)$ eV/quantum is the average energy needed to produce a quantum of scintillation or ionisation [110]. Defining a single W -value for production of either photons or electrons was found to be a good approximation to the use of distinct values (equations 3.15 and 3.16) from keV to MeV energies, as discussed later in this thesis.

For nuclear recoils, the fact that a considerable fraction of the deposited energy is transferred into atomic motion and eventually heat needs to be considered. The fraction of the energy transferred to the electronic systems of the atoms (excitation and ionisation) as opposed to atomic motion, is known as the Lindhard factor (\mathcal{L}) [111]. The recoil energy in this case becomes:

$$E = \frac{W}{\mathcal{L}}(n_\gamma + n_e) \quad . \quad (3.21)$$

The Lindhard factor is energy dependent and has been calculated to be between 0.15 and 0.3 for nuclear recoil energies between 1 and 100 keV [112]. Because of this difference between ER and NR scattering, reconstructing the energy of an interaction requires knowledge of the type of recoil.

It is important to point out at this stage that the track topology of electron and nuclear recoil interactions in LXe is different. In NR interactions a lot of the energy may be transferred in each collision leading to a small cascade with the distinction between the projectile and the target atom quickly lost. On the other hand, ER tracks are longer and well defined. Additionally, the energy partition between ions and excitons differs between NR and ER interactions with ~ 20 times fewer excitons than ions produced by ER interactions ($\alpha_{ER} = 0.06$) and a comparable number of excitons and ions seen in NRs ($\alpha_{NR} \sim 1$) [110]. The importance of these differences between NR and ER interactions for dark matter searches is discussed in the next section.

3.3.3 Scintillation and ionisation yields in liquid xenon

Having discussed the fundamental processes behind signal production in LXe, two quantities become of interest. The ionisation or charge yield (Q_y) is defined as the number of free electrons produced per keV of recoil energy, while the scintillation or light yield (L_y) is defined as the number of photons emitted in scintillation per keV:

$$n_e = Q_y E \quad , \quad (3.22)$$

$$n_\gamma = L_y E \quad . \quad (3.23)$$

Measurements of the light and charge yields at different energies and for a range of electric fields are shown in Figure 3.3 for NR and Figure 3.4 for ER interactions. The mean light and charge yields are anticorrelated since for each ion produced in an interaction, the electron may either escape and contribute to the ionisation response, or recombine and contribute to the scintillation response. For ER interactions, the ionisation channel is nearly an order of a magnitude more sensitive than the scintillation channel at the lowest of energies (~ 1 keV) while the two responses are comparable for NR interactions. The field variation is seen to have little effect on the NR yields and is much more significant for ER interactions, where the charge yield increases with an increasing drift field, while the light yield decreases since more electrons escape recombination. This dependence is stronger for higher energies and becomes negligible for energies below ~ 6 keV.

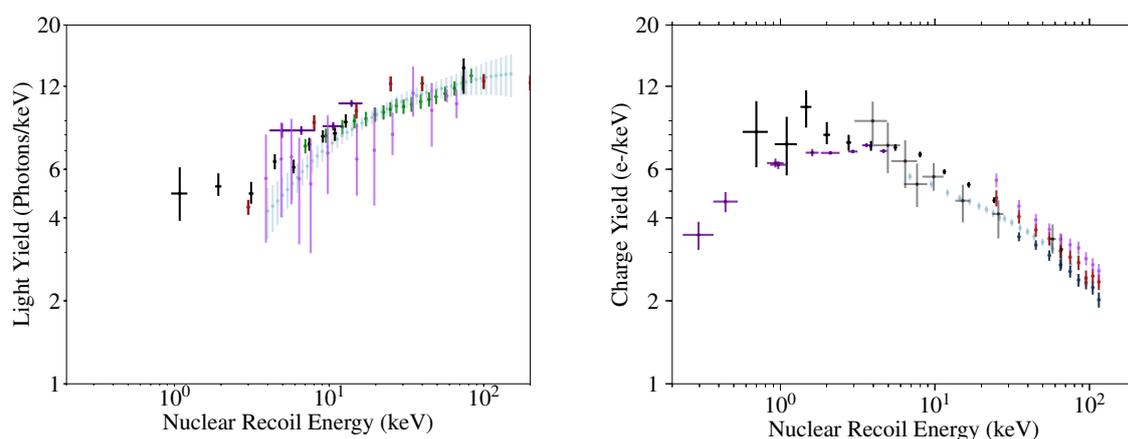


Figure 3.3: *Left:* Selection of representative LXe NR scintillation yield measurements for energies between 1.1 keV and 200 keV and electric fields of 60 V/cm in green (Dahl 2011 [110]), 180 V/cm in black (LUX 2016 [113]), 530 V/cm in red (Xenon100 2013 [114]), 1000 V/cm in violet (Manzur 2010 [115]) and 3400 V/cm in lightblue (ZEPLIN-III 2011 [116]). *Right:* Selection of representative LXe NR ionisation yield measurements for energies between 0.3 keV and 115 keV and electric fields of 60 V/cm in light blue (Dahl 2011 [110]), 100 V/cm in dark blue (Case 2006 [117]), 180 V/cm in black (LUX 2016 [113]), 220 V/cm in purple (LLNL 2019 [118]), 270 V/cm in red (Columbia 2006 [117]), 1000 V/cm in grey (Manzur 2010 [115]) and 2000 V/cm in lilac (Columbia 2006 [117]).

3.3.4 Noble Element Simulation Technique

The Noble Element Simulation Technique (NEST) is a simulation package to model the scintillation, ionisation and electroluminescence processes in noble elements [123]. It has been specifically tailored to dual-phase xenon time projection chambers and has been found to produce accurate results for both ER and NR interactions over a wide range of energies. It has since been extended to liquid argon detectors.

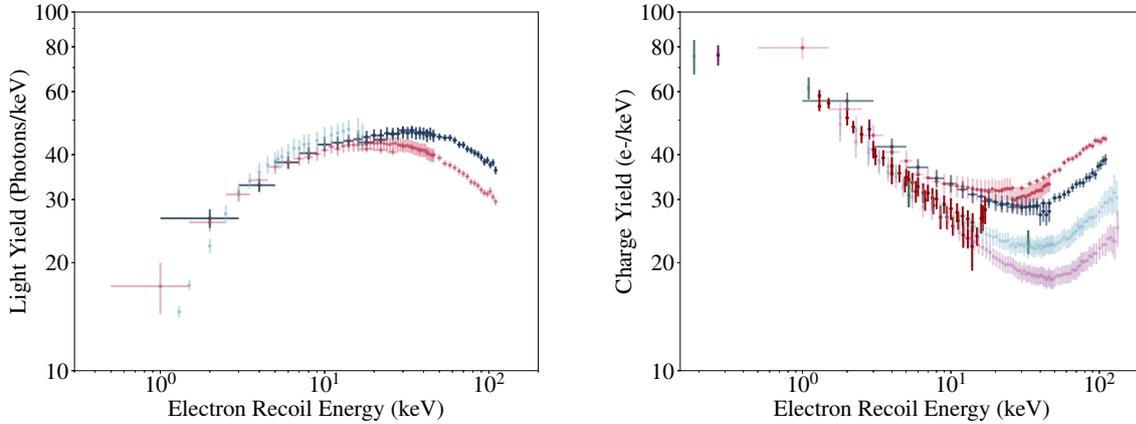


Figure 3.4: *Left:* Selection of representative LXe ER scintillation yield measurements for energies between 1 keV and 115 keV and electric fields of 180 V/cm in light blue (LUX 2016 [119]), 190 and 480 V/cm in black and red, respectively (PIXeY 2017 [120]). *Right:* Selection of representative LXe ER ionisation yield measurements for energies between 0.186 eV and 115 keV and electric fields of 98 V/cm in lilac (LUX 2019 [121]), 180 V/cm in magenta (LUX 2016 [119]), green (LUX [122]) and light blue (LUX 2019 [121]), 190 and 480 V/cm in black and red, respectively (PIXeY 2017 [120]).

In essence, in NEST the equations presented in Section 3.3 are parametrized, with values taken from global fits to available data. The model uses the fits to predict the number of electrons and photons produced in an interaction. Material parameters such as the density and temperature of the medium are taken into account. The electric field applied is also incorporated in the model and affects the recombination fraction and hence the ion-to-exciton ratio for a given interaction. Moreover, the package gives the ability to model custom detectors including efficiencies and operational parameters. For NR interactions NEST models the number of quanta produced during an interaction using a power-law model, guided by the Lindhard theory. The model is allowed to depart from the Lindhard approximation at large recoil energies (allowing for bi-excitonic and Penning quenching).

Data describing the response of noble liquids to particle interactions published in 2019 have been used to motivate the latest NEST release (v2.0). This version includes two sub-versions: NEST v2.0.0 and v2.0.1 [124]. The ER models are identical in the two versions, while for NRs a key new dataset [118] has been included in v2.0.1 that affects the ionisation model for nuclear recoils at low energy (critical for this work); these new data reveal a pronounced roll-off at low energies (which is not unexpected), as shown in Figure 3.3. More detail on the low energy behaviour of the NR models is presented in Section 6.3.2 (and for the ER model in Section 5.2.4).

3.4 Dual-phase time projection chambers

3.4.1 Operating principles

Dual-phase xenon detectors have the ability to record two different optical response signals produced by interactions in LXe. The S1 signal arising from the detection of the prompt scintillation response includes the combined photon output from both the exciton and the recombination luminescence described previously. A delayed ionisation response is produced by the electrons released during the interaction, which are drifted upwards to a gas phase with the use of electric fields, where they produce a secondary scintillation signal, S2, via the process of electroluminescence. In most experiments of this type, including LUX and LZ, the two signals are detected by two arrays of photomultiplier tubes (PMTs), placed at the top and bottom of the detector — this is illustrated in Figure 3.5.

Two main regions of electric field are typically needed in these detectors. A drift field applied to the bulk of the liquid drifts the electrons up to the liquid surface; a stronger field region near the surface helps emit the electrons to the gas phase. In LXe a field of a few kV/cm is required to overcome the electron affinity of the liquid and hence achieve a significant emission probability into the gas (extraction efficiency of 0.85 at 5 kV/cm in the liquid [125]). Due to the dielectric constant of the liquid phase, the field in the gas is nearly twice as strong (~ 10 kV/cm) and this helps generate a large electroluminescence yield. The electroluminescence photon yield is in fact directly proportional to the gas-phase electric field.

The emission of electroluminescence in the gas phase proceeds via the formation of Xe_2^* excited dimers. The accelerating electrons collide inelastically with xenon atoms causing excitations, which leads to Xe_2^* dimers which then decay emitting VUV photons, with spectra similar to the primary scintillation. There is though a small difference in the mean wavelength and spectral width, due the slightly lower exciton energy levels in the liquid compared to the gas [107, 126]. We will ignore these differences here.

The detected signals can then be defined as:

$$S1 = g_1 n_\gamma \tag{3.24}$$

$$S2 = g_2 n_e \tag{3.25}$$

The quantities g_1 and g_2 are the detection efficiencies for S1 and S2, or more specifically, the number of photons detected (phd) per scintillation photon or ionisation electron, respectively [107].

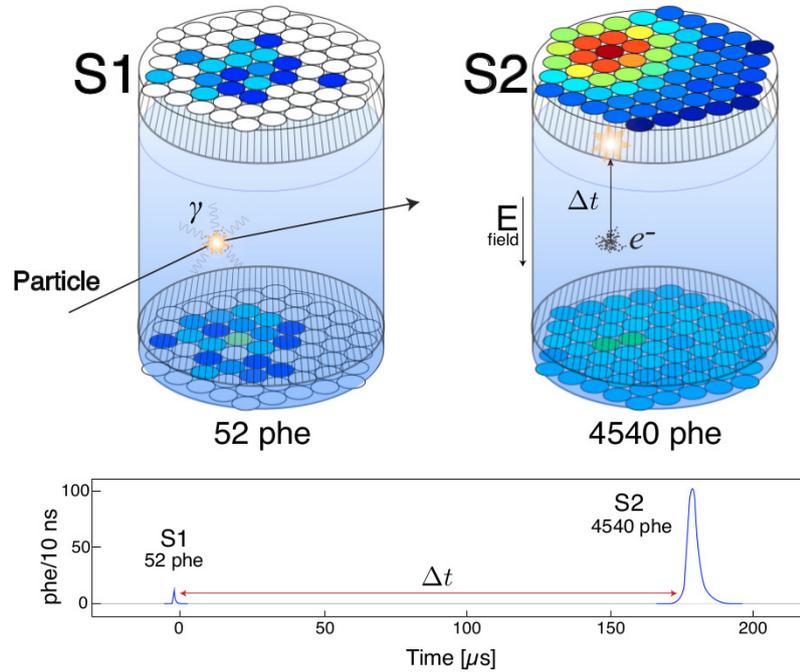


Figure 3.5: Representation of the operating principle of LXe-TPCs such as LUX along with a schematic of an example event [77].

Dual-phase TPCs utilise the time separation between S1 and S2 and the (x,y) resolution, which comes from analysing the signal pattern in the top PMT array, for a full position reconstruction of the event and hence the exclusion of interactions taking place in the outer layers of the detector becomes possible, thus providing self shielding from radioactive detector backgrounds [68, 107]. In principle, integrating the pulse area of the S1 and S2 signals and dividing by the average response produced by one photoelectron gives the number of scintillation and electroluminescence photons that were detected.

3.4.2 ER-NR discrimination

Nuclear and electron recoils are governed by different microscopic mechanisms for signal production, related to the different nature of the respective tracks, and hence their signals in dual-phase xenon detectors exhibit different characteristics. The discrimination between these two type of interactions is of particular interest as in many dark matter models the hypothesised dark matter particle (such as a WIMP) causes NR interactions, while the prominent backgrounds generate ER interactions.

Because of the increased number of electron-ion pairs observed in electron recoils and taking into account the fact that recombination is similar between ER and NRs at very low energies, a higher number of electrons is produced and hence detected in ERs. The ratio of

the delayed to prompt signals, $S2/S1$, is hence used as a discrimination parameter for the two interactions, with higher ratios seen for electron recoils — see Figure 3.6. To quantify the discrimination ability of a particular detector, the ER leakage past the median of the NR population is used, averaged over some region of interest in energy or in $S1$: LUX achieved a value of 99.8% for the Run 3 (2013) re-analysis [127], while LZ has defined a minimum requirement at 99.5% [98] for this parameter. The discrimination values quoted are dependent on the recoil energy since the ion-to-excitation ratio is dependent on the energy; in addition, it depends also on electric field, which influences the recombination dynamics. A comprehensive study of discrimination in liquid xenon has just been published by the LUX collaboration [128].

In addition to the $S2/S1$ ratio, an additional mechanism is in principle available for ER-NR discrimination. Even though prompt scintillation occurs through the de-excitation of Xe_2^* , the ratio of triplet to singlet state excitations is dissimilar between the two types of interaction. Thus, the scintillation decay time will contain varying fractions of slow ($\tau \sim 22$ ns) and fast (~ 4.3 ns) components. At the same time, the recombination time, which applies mostly to higher energy ER, is even slower (~ 45 ns). The difference between the two average time constants in LXe for ER (~ 32 ns) and NR interactions (~ 18 ns) is rather small, making the effect on $S1$ shape difficult to observe but nevertheless possible [129, 130, 131]. Detectors utilising liquid argon, where the decay times between the triplet (7 ns) and single (1700 ns) states differ significantly, have more effectively made use of pulse shape discrimination to differentiate electron from nuclear recoils [132, 133]. However, LAr analyses require large $S1$ signals for differences in the time constant to be well measured with thresholds usually placed at 50–100 photoelectrons. This is significantly higher than the 2- or 3-photoelectron threshold used in LXe-TPCs.

3.5 The LUX experiment

3.5.1 Overview of the experimental setup

The LUX experiment completed two WIMP-search runs between 2013 and 2016, utilising 300 kg of LXe in a dual-phase xenon detector [127, 134]. The LUX TPC was contained within a low intrinsic radioactivity titanium cryostat. The cryostat itself was placed within a water tank shielding the xenon detector volume from external radioactivity. The water tank was located 4,850 ft underground in the Davis Campus of the Sanford Underground Research Facility (SURF), in Lead, South Dakota, USA [135]. The $\sim 270,000$ litres of ultra-pure water surrounding the cryostat along with its placement deep underground reduced cosmogenic and laboratory backgrounds to negligible levels [136]. Key to the unprecedented level of background achieved by LUX was the careful selection of materials and components for very

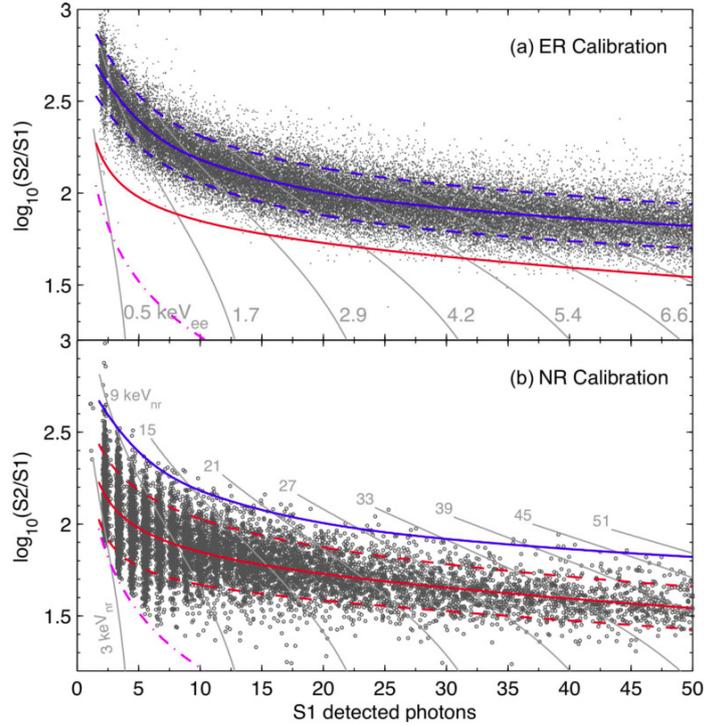


Figure 3.6: Discrimination parameter $\log(S2/S1)$ as a function of S1 signal obtained with LUX; the upper panel shows the response to a tritiated methane ER calibration [119] and the lower panel shows the response to NR calibration using DD neutrons [113]. Solid lines show band means, while dashed lines show the 1.28σ contours (blue for ER and red for NR). The dashed magenta line presents the S2 threshold applied in the analysis.

low intrinsic radioactivity and the ^{85}Kr chromatography removal campaigns [137, 138]. Further detail on LUX backgrounds and their relevance to this thesis is discussed in Section 6.2.

The LXe target was surrounded by 12 high-reflectance panels made from polytetrafluoroethylene (PTFE), and viewed by two arrays each with 61 PMTs (2-inch Hamamatsu R8778). The active region of the TPC contained 250 kg of active mass, within which interactions produced S1 and S2 signals. This region was defined vertically between a gate grid located 0.5 cm below the liquid surface and a cathode grid at the bottom, 48 cm below the gate grid. The diameter of the active region was 47 cm, enclosed by the PTFE panels [39]. The top PMT array was located in the vapour phase, while the bottom array was immersed in the liquid below the cathode grid, leading to an enhanced S1 collection efficiency observed at the bottom array, due to the similar refractive index between the PMT quartz windows and LXe. A schematic of the LUX detector is shown in Figure 3.7.

In the first underground run (relevant to this thesis), ionisation electrons produced in the LXe target were drifted toward the gas phase using an electric field created between the cathode grid at -10.0 kV and the gate grid at -1.5 kV. There were 47 copper field-shaping rings connected by a resistive ladder between the gate and cathode grids, placed outside of

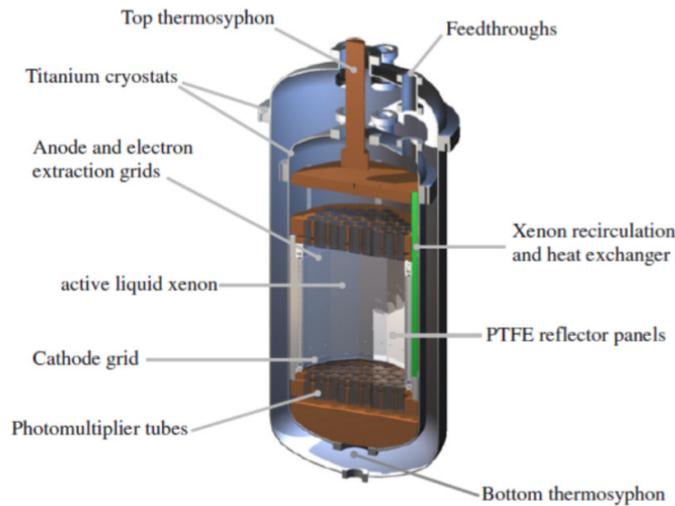


Figure 3.7: Schematic of the LUX detector, illustrating the main components of the TPC, cryostat and thermosyphons [77].

the PTFE panels, with the average drift field being 180 V/cm. For this field and dimensions, electrons starting out from near the cathode grid exhibited a maximum drift time of 322 μ s. A further potential difference was applied between the gate and anode grid (0.5 cm above the LXe surface) creating a high enough field to extract and accelerate electrons in the gas phase, giving rise to the secondary scintillation signal. The electron extraction field was 2.84 ± 0.16 kV/cm in the liquid and 5.53 ± 0.3 kV/cm in the gas. A single electron emitted from the liquid produced a total of ~ 25 detected photons. Two additional grids were used, known as the top and bottom shield grids. The top shield lied 4 cm above the anode and 2 cm below the top PMT array while the bottom shield was 2 cm above the bottom PMT array and 2 cm below the cathode. These grids were held at low potentials compared to the cathode and anode, counteracting their high electric fields and protecting the PMTs [39].

During operation, the xenon was continuously vapourised and circulated at a rate of 229 kg/day in order to remove impurities emanating from detector surfaces. This was achieved utilising a heated Zr getter [39]. Electronegative impurities quickly kill off the free electron lifetime, and hence the S2 signal, as drifting electrons can attach to species such as O₂, N₂ and others on their way to the liquid surface. The getter material reacts with these electronegative impurities (oxidizing) and subsequently heats to remove the reacted layers. Getters can only be used in the gas phase so Xenon is pumped out, evaporated, gettered and then condensed and re-introduced. This process, along with the dual-phase nature of LUX required excellent control of the thermal parameters of the detector which was achieved through dedicated nitrogen thermosyphons and heat exchangers [139, 140], with the detector temperature held a small fraction of one degree near 175 K at an absolute pressure of 1.8 bar.

A good understanding of the performance of the LUX detector, and especially the low energy response, was accomplished through an extensive set of calibrations [135]. These included dedicated calibrations of the all-important single photon response of all PMTs using an LED system, of the response of the detector to ER interactions through external sources as well as dispersible radioactive sources, and to NR interactions using external neutron sources.

LUX has made extensive use of ^{83m}Kr for calibration purposes, with this radioisotope used regularly (approximately weekly) during its operation [135, 141]. ^{83}Rb was used as the source material, as it decays to ^{83m}Kr with a half life of 86.2 days. The produced ^{83m}Kr diffuses from the generator material into the detector volume (being a noble element, it is not removed by the getter), where it decays through two transitions to the ^{83}Kr ground state with a half-life of 1.83 h releasing a total energy of 41.5 keV in two conversion electrons (with the intermediate state having $T_{1/2} = 154$ ns. This calibration was used for a vast array of studies. The electrons created in the interaction in LXe deposit their energy within $O(10 \mu\text{m})$ of the decay site, making it a good source for calibrating (x,y) geometric non-uniformities (used to correct the S1 and S2 signals) and for understanding the electric field in the target [141, 142]. It was also used to calibrate the response of LUX to single ionisation electrons, to measure the electron lifetime, to understand the light yield stability, to measure detector levelling, etc. An example LUX event from a ^{83m}Kr interaction is shown in Figure 3.8. LUX also deployed another novel ER calibration to map the response to lower ER interactions using tritiated methane – this is described in Chapter 5.

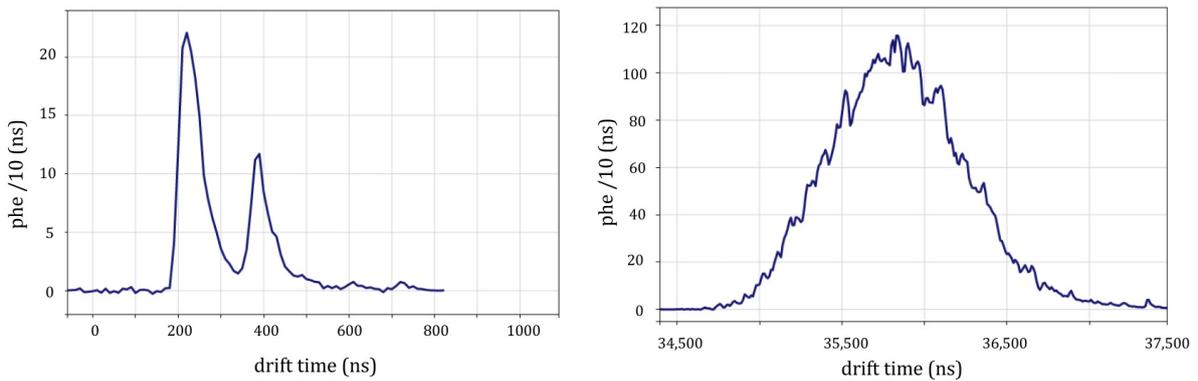


Figure 3.8: The (left) S1 and (right) S2 pulses for an example LUX event from a ^{83m}Kr interaction with 41.5 keV combined energy. Two distinct S1 pulses can be seen in this event, produced by the conversion electron cascade — these are merged together into a single S1 pulse in a significant fraction of events.

The gain factors g_1 and g_2 were extracted using the Doke method [122, 135], which utilises a set of ER line sources in the energy range of 5.3 to 662 keV (^{83m}Kr , ^{127}Xe , ^{129m}Xe , ^{83m}Xe , ^{214}Bi , ^{137}Cs , ^{131m}Xe). The sources have known energies and different mean S1 and S2 signals between them, so the combined energy model (equation 3.21) can be fitted for the gain factors to be determined, as shown in Figure 3.9.

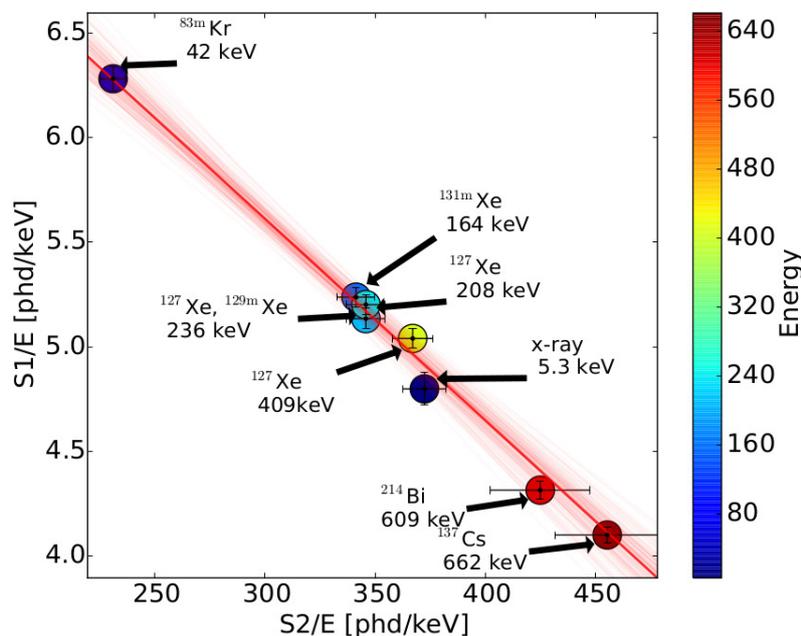


Figure 3.9: The Doke method of extracting the g_1 and g_2 gain factors, using a set of ER line sources [135].

Finally, neutrons from a deuterium-deuterium (DD) generator [113] and from an Am-Be source were used to study the response of the detector, and of its medium in general, to nuclear recoils. The DD calibration was particularly important. Mono-energetic 2.45 MeV neutrons were collimated through an air-filled conduit spanning the distance from the water tank wall to the LUX detector cryostat. A key population of neutrons entering the active liquid xenon volume would scatter twice and then leave the target medium, as shown in Figure 3.10. Using the full 3D position of both scatters allowed a reconstruction of the recoil energy and hence the DD calibration offered a unique ability of calibrating the ionisation yield as a function of energy in LXe (with LUX achieving a low energy calibration down to 0.7 keV). A smaller population of single scatter events was then used to calibrate the light yield, with the S2 charge used as a measure of energy. The DD calibration was also used to calibrate the location of the nuclear recoil band in LXe (as shown in Figure 3.6).

LUX was decommissioned in 2016, making way for the LUX-ZEPLIN (LZ) dark matter detector, following 3.5-years of operation. During this time, two successful science runs were completed, WS2013 (Run 3) and WS2014-2016 (Run-4), placing world-leading exclusion limits on WIMP interactions. The initial dark matter search for WS2013 consisted of 85.3 live-days using 118 kg fiducial mass [136]. An improved re-analysis of the same dataset increased the exposure to 95 live-days of search data and, mostly through improved calibrations, lowered the energy threshold, improved event reconstruction and the background model, achieving a more stringent and more robust SI result, especially at low WIMP masses [127]. Several other

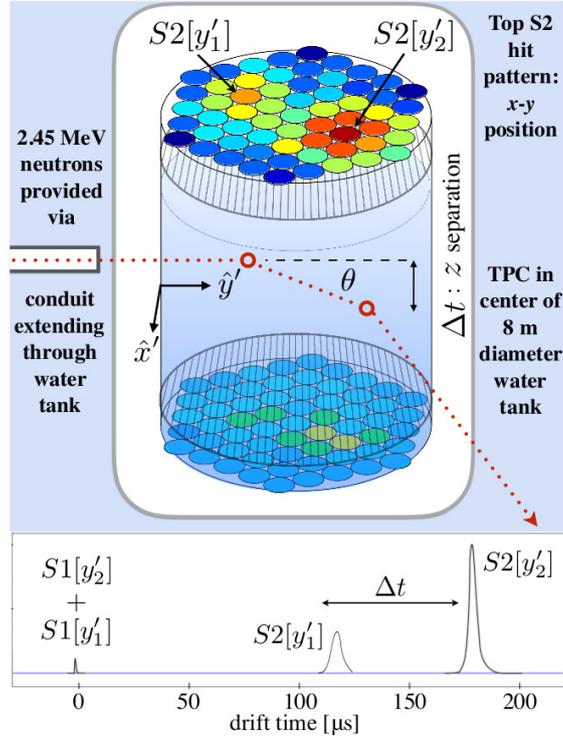


Figure 3.10: Schematic diagram of the LUX DD calibration experimental setup [113].

physics results were also obtained using the WS2013 dataset for SD scattering of WIMPs [143], searches for Axion-Like Particles (ALPs) [144] and sub-GeV dark matter limits using the Migdal effect [145].

Following the WS2013 data taking, the detector grids were "conditioned" to try to improve the high-voltage performance of the instrument; this was done by maintaining high fields between the electrodes just past the onset of sustained micro-discharge through the gas (after draining the liquid). The conditioning was performed in an attempt to increase both the drift and the electroluminescence fields, and hence improve the electron extraction efficiency from the liquid to the gas. The extraction efficiency was successfully increased from $49 \pm 3\%$ in WS2013 to $73 \pm 4\%$ in WS2014-2016 [134]. However, the intense VUV flux produced during this operation led to major charging of the PTFE panels (it is important to note that conditioning of the LZ grids will not take place and as such PTFE charging issues will be avoided). Unfortunately, PTFE is a phenomenal insulator and the charge could not be neutralised, which led to a major distortion of the electric field in WS2014-2016. Contrary to WS2013 operation, the electric field lines now varied in azimuth, depth and evolved in time making (x,y) position reconstruction very challenging. Eventually, time varying field maps were created using a combination of calibrations with dispersed sources and electrostatic modelling [141], for WS2014-2016 but the analysis suffered significant difficulties, including the need for a more conservative fiducialisation, varying light and charge yields and the need to partition the

exposure (live-time \times target mass) in 16 approximately uniform partitions (four for drift time and four for time period) for the final statistical analysis. This analysis technique was used to place limits on SI [134] and SD [94] WIMP scattering. In this thesis we only consider data obtained using WS2013 and present detector operating conditions and parameters for this time. The field variations in WS2014-2016 would increase the complexity of our analysis, the purpose of which was to benchmark the new techniques, and would not have led to major improvements in sensitivity.

3.5.2 Data acquisition and reduction

The raw data in LUX are essentially the PMT waveforms digitised at 100 MS/s. The output from each PMT was amplified and shaped before being digitised into 10 ns samples [146]. The digitisation mode used is known as pulse-only digitisation (POD), which records only a limited span of the waveform around a triggered signal. The acceptance for this threshold is discussed in detail in Section 4.2.1. The POD trigger operated on PMT-channel pairs (the PMTs were instrumented in pairs to reduce financial cost), hence when one channel of a pair went above threshold the data of the partner channel was also recorded. The DAQ was operated in multi-event mode recording all digitised pulses from the PMTs into ".dat" files, and later sorting in events (see next paragraph). These files were stored underground and transferred to the LUX Primary Mirror at Brown University.

The LUX trigger chain selected events in the DAQ data stream for writing to file [147]. The 122 PMTs were summed into 16 trigger channels with no adjacent PMTs in the same sum. The trigger system was made from two internal filters, the S1 and S2 filters, which acted as integrators and had a response proportional to the area of the input pulse optimised separately to each pulse shape. Events were then chosen based on their pulse shape characteristics. In WS2013 the trigger system was operated in "S2 search" mode [135], in which only the output of the S2 filter was used in the trigger decision process. During this time a low filter threshold of 8 phe was chosen. The trigger condition selected events for which at least two trigger channels had an S2 filter output greater than the low filter threshold within a 2 μ s window. Once the trigger signal was sent to the DAQ, a hold-off time of 1 ms was established, to ensure that subsequent pulses in the same event did not re-trigger the DAQ. The S2 trigger efficiency for S2 pulses was found to be 98% for $S2 > 90$ phe, which corresponds to a mean of ~ 4 ionisation electrons emitted into the gas [148].

Offline software, called the event builder, used the ".dat" files (output of the DAQ) together with the trigger information to find the PMT signals that fell within a specified window around the trigger pulse (S2 pulse in WS2013). These final triggered LUX events were stored as ".evt" files and sorted by event, then by channel and then by pulse.

The LUX data processing framework (DPF) extracts relevant features from raw digitised PMT data and returns a set of reduced quantities (RQs), operating on each event in the ".evt" files. To do that it employs several algorithm modules. A schematic of the data acquisition and processing framework is shown in Figure 3.11. A set of (~ 400) RQ quantities are produced at this stage, constituting the set of processed data (".rq.root" files). The DPF modules are divided in 8 sections:

- *Initialising* the RQ files.
- *Calibrating* the pulses from mV to detected photons per (10 ns) sample based on the PMT gains (see Section 4.2.1).
- *A pulse finding* algorithm, which uses a fixed-width (400 samples) box filter. The algorithm returns a start time and end time for each pulse and saves the 10 largest pulses per event, in decreasing area, with the classification prioritising pulses appearing prior (in time) to the largest pulse in area.
- *Pulse feature reduction* where the waveform data are reduced to RQs, such as pulse width, area, rise time, area fractions, etc.
- *The pulse classification* algorithm classifies pulses into S1, S2, single detected photon (SDP), single electron (SE) or other, using pulse RQs calculated in the previous step, related to the pulse shape and properties. The logic followed by the pulse classifier is shown in Figure 3.12. Initially, pulses are divided into S1-like and S2-like groups. This selection is based on width information (S1s are much narrower than S2s — see Figure 3.8) and area fractions (to utilise the shape of S1-like signals, with fast rise compared to S2s) and the PMT hit-distributions (most light seen at bottom array for S1s and the opposite for S2s). Subsequently, S1-like pulses are divided into S1s and SDP pulses. A valid S1 must include detection of one or more photoelectrons in at least two non-partner PMTs. Pulses with single photon hits are classified as SDPs. Similarly, S2-like pulses with more than 55 phd are classified as S2s and with less than 55 phd as single electrons. Finally, two modules are responsible for the position reconstruction and application of corrections. These corrections are later applied to the pulse areas and ensure a uniform response across the detector.
- *Event classification* during which events are classified as single scatters, multiple scatters, S2-only or other.
- *Position reconstruction* of the event. The top PMT array is used to reconstruct the (x,y) location of an interaction, using the S2 photon hit pattern, through the Mercury algorithm [149]. This utilises a set of empirical functions, known as light response functions (LRFs), determined by an iterative procedure involving all PMT responses. This achieves a resolution of 1 cm at the S2 threshold, improving rapidly for larger signals.

- The *detector response corrections* module calculates pulse area corrections given the (x,y,z) position of the event and creates a set of correction RQs that can later be used to produce geometrically unbiased results.

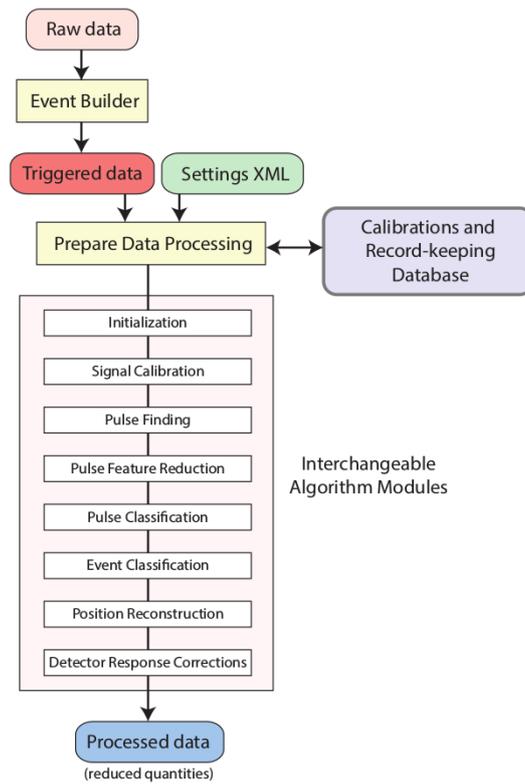


Figure 3.11: Schematic view of the LUX Data Processing Framework [135].

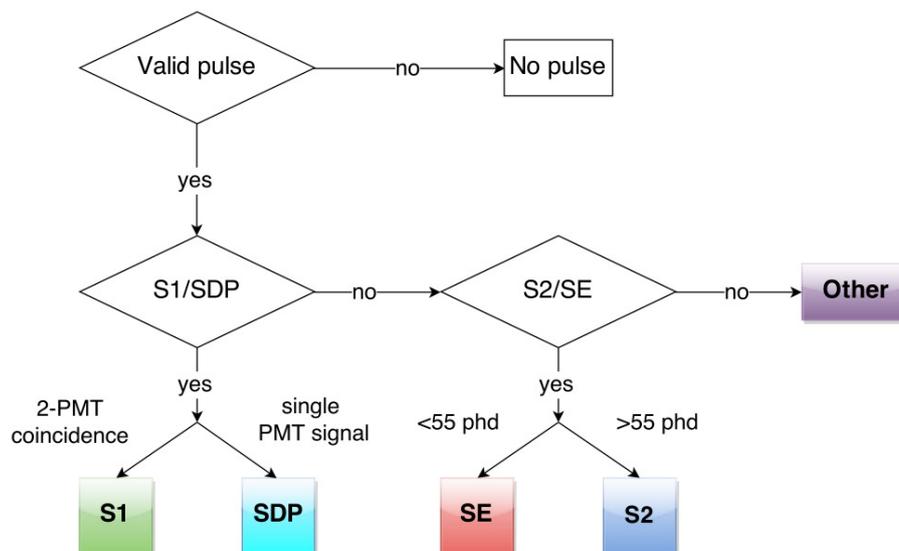


Figure 3.12: Logic followed by the pulse classifier algorithm [135].

As noted above, the data processing framework results in a large set of RQs. Most analyses only require a small subset of these RQs, in a more processed format. The output of the DPF includes PMT-by-PMT area and time information for each pulse, RQs holding different corrections and several detailed pulse quality quantities. Some analyses work on summed (over all channels) and corrected pulse quantities and use a selection of single scatter events. This last step in data reduction is performed by the *event filtering and summed quantity* algorithm, known as the "filter code". At this stage, events of interest are selected and the total corrected S1 and S2 areas, drift time and z-corrected (x,y) information is produced along with a set of other needed RQs (at pulse rather than PMT level), as specified by the user. Events of interest for most analyses are single scatters and known as "golden" events:

- one valid S2 follows a valid S1 (with additional S2s occurring before the valid S1 being allowed),
- only one valid S1 signal before the S2 signal (but additional S1s can follow the S2 signal),
- the area of the S2 signal must be larger than the area of the S1.

The analysis described in this thesis requires reprocessing of the full science dataset and calibrations because it requires PMT level information and a lowered threshold on the S1 signal. The alternative processing chain is discussed and validated in Section 6.1.

3.6 The LUX-ZEPLIN experiment

At the heart of the LZ experiment is a traditional dual-phase xenon detector, like LUX. LZ benefits from a size increase compared to LUX (7,000 kg of active mass), along with an increase of PMT number from 122 to 493 (3-inch Hamamatsu R11410-20). The TPC is located inside the inner cryostat vessel of ~ 1.6 m diameter and 2.6 m height. Thin PTFE panels are attached to the inner cryostat vessel wall and bottom dome to enhance light collection efficiency [64]. The LXe between the inner cryostat wall and the outer TPC wall is known as Xe "skin" layer and will act as a gamma-ray veto detector. The Xe skin is viewed by 131 side and dome PMTs of 1-inch (93 PMTs) and 2-inch (38 PMTs) diameter. The outer cryostat vessel is surrounded by a nearly hermetic liquid organic scintillator outer detector, acting as a neutron veto system. The 17.3 tonnes of scintillator is encapsulated in clear acrylic vessels viewed by an additional 120 PMTs installed ~ 1 m away in the water space. The integrated veto system is capable of rejecting several backgrounds such as gamma-rays and neutrons generated within the detector (cryostat materials, PMTs, etc.) [98]. LZ is located within the same water tank that hosted LUX at SURF. A schematic of the LZ detector is shown in Figure 3.13. LZ is soon to start final commissioning and is due to start taking data in late 2020.

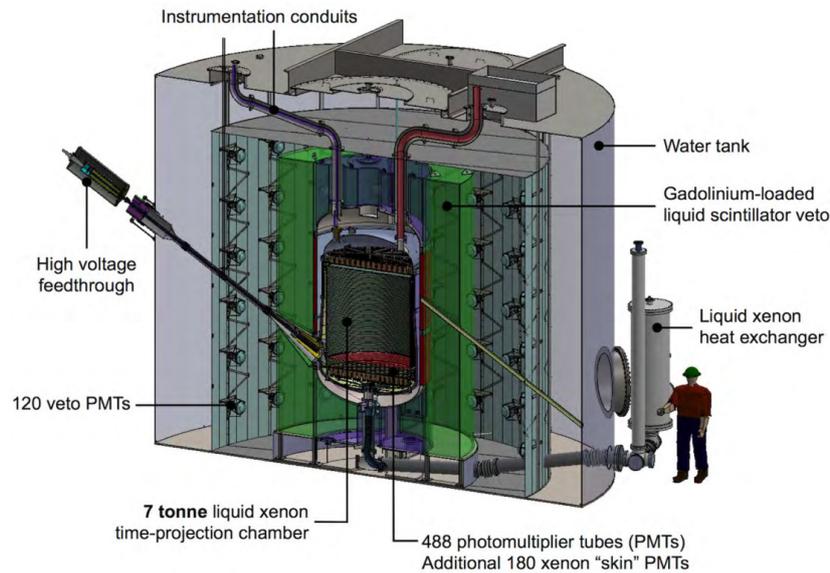


Figure 3.13: Schematic of the LZ detector illustrating the LXe-TPC contained within the titanium cryostat, located inside the water tank at SURF [64].

3.7 Low energy and low mass sensitivity techniques

The presence of the S1 and S2 signals is crucial in standard WIMP search analyses. Naturally, there exists a threshold energy below which any signal is beyond the reach of a given experiment. We conclude this chapter with a survey of techniques that can be applied to lower these thresholds, before describing the new technique developed in this thesis in Chapter 4.

For standard S1+S2 analyses, the smallest detectable energy is determined by the S1 signal, since the ionisation channel is approximately an order of a magnitude more sensitive than the scintillation channel. Even though S1 signals down to a single detected photon are observable by LXe-TPCs, PMT dark noise motivates a coincidence requirement involving multiple channels. In smaller detectors, such as LUX, only S1 signals with 2 or more photons detected in different PMTs and within 100 ns are selected (2-fold coincidence). It is essential in removing a source of background events in which PMT dark counts coincide in time and occur a short time before an S2 pulse, faking a valid signal. The corresponding requirement for LZ is 3 phd, due to with the increased number (and size) of PMTs there is now a substantial probability that two dark counts are recorded in coincidence (see Chapter 7).

3.7.1 S2-only analyses

The most common technique to search for lower energy deposits than possible with standard S1+S2 analyses is known as an "S2-only" analysis, where events containing an S2 pulse but

no discernible S1 pulse are used, exploiting the lower threshold of the ionisation channel. This type of analysis provides a better sensitivity for WIMP masses below ~ 10 GeV but, nevertheless, suffers from significant limitations. The first arises from the fact that ER-NR discrimination through the use of the S2/S1 ratio as the discriminating factor is not possible. At the same time, drift time information is not available and hence the position of the event cannot be properly reconstructed; this prevents identification of signals arising in the TPC grids, which are easily removed in standard analyses. Thus, experiments which rely significantly on self-shielding of external backgrounds, such as LXe-TPCs, will need to contend with much higher background levels.

Even though LXe-TPCs cannot achieve ER-NR discrimination in S2-only analyses, they can still place competitive limits for WIMP-nucleon scattering. Xenon10 and Xenon100 were the first experiments to successfully publish such analyses, producing competitive results [150, 151]. A recent analysis published by Xenon1T reached an energy threshold of 0.4 keV, compared to the 3.5 keV threshold achieved in standard analyses. The analysis employed S2-width cuts (the width of the S2 pulse changes as electrons diffuse before reaching the surface) to remove events originating on the cathode wires or decays on the electrodes at the top [152]. Even though the method was viable, it is not as effective as using both S1 and S2 signals. Recent limits placed by Xenon1T and DarkSide-50 are shown in Figure 3.14. The S2 analysis threshold, which in this case determined the energy threshold, was placed at 12 produced electrons for Xenon1T and 4-7 extracted electrons for the DarkSide-50 analysis. A LUX publication is in preparation on a similar analysis, with a threshold of ~ 5 extracted electrons. In both the Xenon1T and DarkSide-50 analyses, the background model fails below their respective threshold with a large increase observed in the event rate. Sources of the single and few-electron increased background are thought to be due to electrons trapped in the liquid and later released (incomplete extraction), building up to several electron signals and due to spurious emission from the grids [153, 154].

Experiments utilising liquid argon, such as DarkSide-50, use pulse shape discrimination (PSD) on the S1 signal to identify ER and NR interactions and hence also lose the ability to discriminate ER from NR interactions in S2-only analyses. DarkSide-50 published an analysis where the energy threshold was lowered from 13 to 0.6 keV using ionisation signals [155], but it is noted that the NR ionisation yield in LAr has not been measured over most of the energy range of interest for such an analysis (the lowest energy measured is ~ 7 keV [156]).

3.7.2 Migdal and Nuclear Bremsstrahlung

Noble liquid TPCs are most successfully used to search for WIMPs above ~ 5 GeV through NR interactions (see Section 3.1.1 for the mass and energy threshold dependence). Recently, though, it was proposed that direct detection experiments can utilise alternative channels to achieve

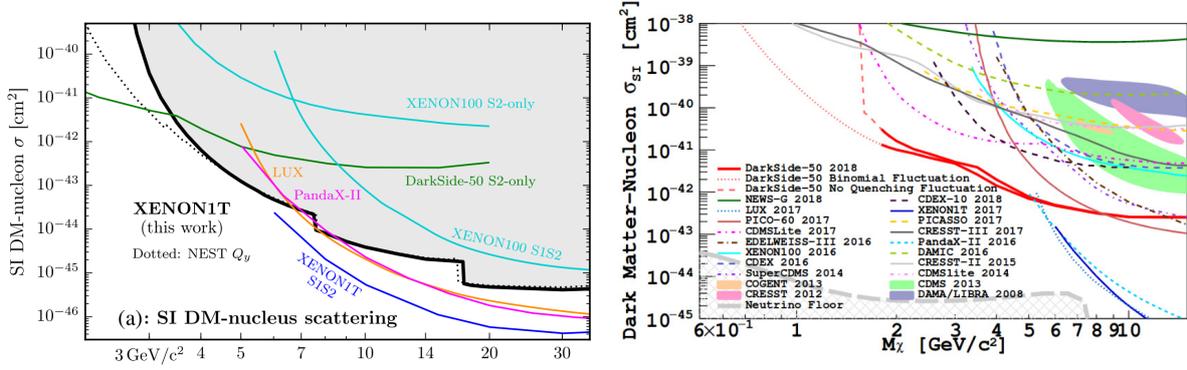


Figure 3.14: Recent 90% confidence level upper limits on SI WIMP-nucleon scattering, using an S2-only analysis, for (left) Xenon1T [152] and (right) DarkSide-50 [155].

sensitivity to lighter dark matter particles scattering off atomic nuclei. These mechanisms are known as the nuclear Bremsstrahlung [157] and Migdal effects [158]. Essentially, both techniques attempt to infer low energy nuclear recoils (below the standard NR energy threshold of the detector) through an accompanying ER signal that is created (albeit rarely) during the scattering event.

During DM-nuclear scattering the atomic nucleus recoil is assumed to be followed by an instantaneous movement of the electron cloud surrounding the nucleus. In reality, there is a time delay in this process, during which the atom remains polarised. During this time, the atom may emit a low energy electron through the so-called Migdal effect [158]. Nuclear Bremsstrahlung is thought to occur when the recoiling atom emits a low-energy photon due to its instantaneous polarisation. In both cases the ER signal due to the emitted electrons or photons co-occurs with an NR signal but, significantly, while the NR may have such low energy as to be undetectable, the ER component may fall above the ER threshold and produce a measurable signal.

Unfortunately, both processes have very small probabilities, but they confer sensitivity in a region where few other technologies are competitive. Upper limits obtained from a recent analysis performed by the LUX experiment for the Migdal and nuclear Bremsstrahlung channels are shown in Figure 3.15 [145]. Competitive sensitivity is achieved even though the underlying effects have small probabilities, because the inner fiducial volume of LUX (and other LXe-TPCs) benefits from low ER backgrounds. Next generation experiments are expected to be even more sensitive due to their increased target mass. Even though this method is powerful in extending the sensitivity of direct detection experiments to sub-GeV masses, the Migdal and nuclear Bremsstrahlung processes have yet to be directly verified experimentally.

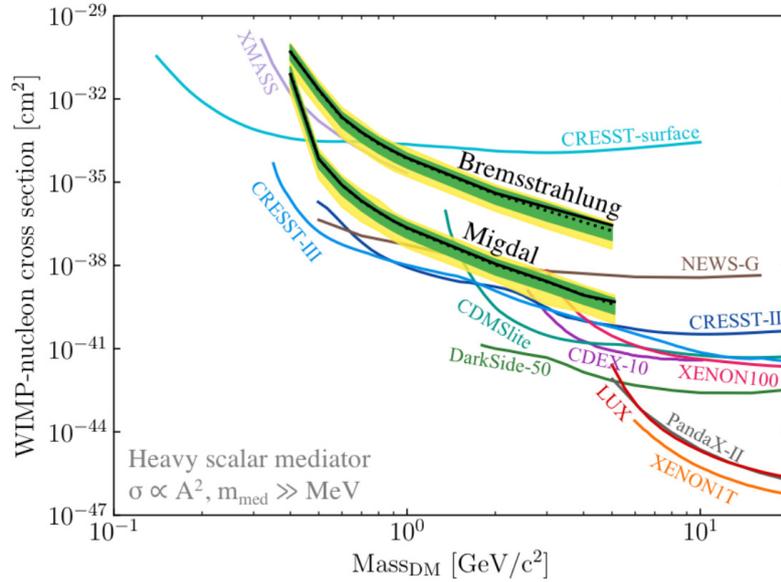


Figure 3.15: Upper limits on the SI DM-nucleon cross-section at 90% C.L. as calculated using the nuclear Bremsstrahlung and Migdal effect signal models assuming a heavy scalar mediator [145] for the LUX experiment, shown in black. Standard results from other experiments are also presented.

3.7.3 Doping with light elements

Another method proposed to enable searches for very light dark matter involves adding hydrogen or helium to the xenon, with hydrogen being the preferred proposal as helium can diffuse through the PMT glass and cause significant afterpulsing. The light species will produce much larger recoil energies than the xenon and hence lead to enhanced sensitivity.

There exist practical challenges to the amount of hydrogen that can be loaded into xenon or argon, with estimates limiting the amount to $O(1\%)$ mole fraction. Even this small percentage of light mass target could bring about significant sensitivity, in a region of parameter space that is relatively unexplored. This is especially true for SD WIMP-proton interactions, where hydrogen is projected to lead to world leading sensitivity in LZ. Doping with deuterium is also possible, which would bring SD WIMP-neutron sensitivity. Additionally, the background level at the core of LXe experiments is very low, much lower than can be achieved in a detector with a light element target (especially since those are typically very small). Additional studies are ongoing to determine how the liquid xenon properties are affected by such doping. In specific, any possible wavelength change on the S1 and S2 photons in doped LXe needs to be understood, especially as the PMT detection efficiency is dependent on it. At the same time the light and charge yields, as well as the Lindhard factor in the doped liquid, would need to be determined not only so that the energy of interactions can be accurately reconstructed but

also to establish whether any changes result in signals falling below threshold or in the S2/S1 discriminating factor losing its power [159].

3.8 Summary

In this chapter we discussed the mechanisms behind signal production in LXe, and specifically signal production and detection in dual-phase xenon detectors. Such experiments have the ability to detect several rare processes of interest, including low mass WIMPs and astrophysical neutrino fluxes. We also described briefly the LUX and LZ experiments, focusing on aspects which are relevant for this work. Additionally, we discussed the low energy threshold of LXe-TPCs and techniques that are or could be employed to extend the low energy sensitivity. This is of particular interest in this thesis, as we present an alternative way to extend the low-energy performance of such detectors.

Chapter 4

Understanding the single VUV photon response of the LUX photomultipliers

LUX has been very successful at probing low energy nuclear and electron recoil interactions, achieving a threshold recoil energy of 1–2 keV for ERs and below 5 keV for NRs. Although both the S1 and the S2 channels are sensitive to very low recoil energies, the ionisation response is more sensitive than the scintillation channel due to the higher ionisation yield at low energies and the higher efficiency and gain of the S2 channel. Thus, the S1 2-fold coincidence requirement determines, to a large degree, the energy threshold of the experiment. This requirement was applied to eliminate a major source of background events, whereby time-coincident PMT dark counts (DC) appearing up to a few hundred microseconds before an S2 pulse can fake a valid event.

The work presented in this thesis is focused on utilising the DPE effect observed in photomultiplier tubes at VUV wavelengths to extend the standard analysis to a single scintillation photon in LUX — and, more widely, to lower the energy threshold and extend the sensitivity of direct detection experiments using LXe targets. In order to proceed with such an analysis the single photon response of the LUX PMTs was carefully calibrated, with the individual PMT DPE fractions and dark count rates quantified. This work is described in this chapter.

4.1 Photomultiplier tubes

4.1.1 Detection principles

Photomultiplier tubes are the most practical and sensitive detectors available for a large number of applications. These vacuum tubes detect photons in the ultraviolet, visible and near-infrared regions of the electromagnetic spectrum. A photomultiplier comprises a semi-transparent photocathode, in which a primary photoelectron (phe) is emitted following the absorption of an incident photon. An electric field directs the photoelectron to an electrode, or dynode, which

amplifies the signal, producing a number of secondary electrons. These secondary electrons are in turn accelerated through a chain of dynodes building up the avalanche multiplication of electrons. The number of electrons increases until they reach the anode, numbering typically over a million or more per primary photoelectron, at which point the signal current is read out. These components are sealed inside an evacuated envelope [160]. A schematic of the internal components of a photomultiplier tube is shown in Figure 4.1.

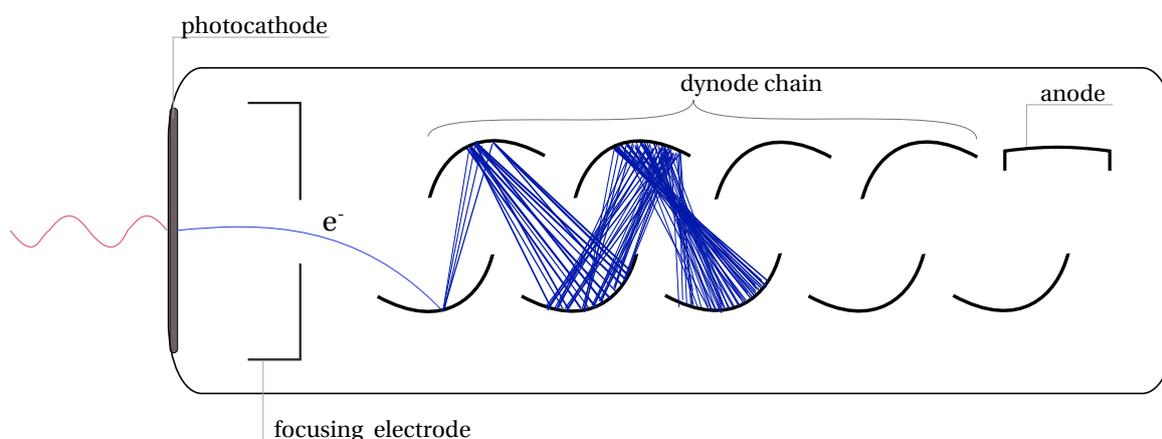


Figure 4.1: Schematic of a photomultiplier tube.

The LUX photomultipliers (Hamamatsu R8778 – 56 mm PMT diameter) feature a VUV-transparent synthetic quartz window (SiO_2) and a bialkali photocathode (Rb-Cs-Sb), a material with a typical work function of ~ 1 eV. An underlying Al strip pattern is also used, to decrease the average resistivity of the photocathode layer upon cooling, thus improving low temperature operation at 170 K. The photocathode diameter is 45 mm. The PMTs were made with low radioactivity materials through careful component selection and have been measured to have 9.8 ± 0.7 mBq/PMT and 2.3 ± 0.5 mBq/PMT for ^{238}U and ^{232}Th , respectively [161]. The PMTs were negatively biased using voltage divider bases made from low background components.

The gain of a PMT is defined as the total electron multiplication factor at the output for a single photoelectron emitted by the photocathode, and depends both on the dynode voltage distribution and the electrostatic design. The dynode chain involves 12 multiplication stages and was operated at the typical bias voltage of 1200 V in LUX, making the total charge detected at the anode 4×10^6 electrons per detected photon. The statistical nature of the multiplication process introduces gain fluctuations and, for the R8778 PMT model at the LUX operating conditions, the single photoelectron response has been generally found to be well described by a Gaussian distribution with a typical resolution of $\sim 40\%$. The R8778 PMT entrance window with its Al strip pattern is shown in Figure 4.2.

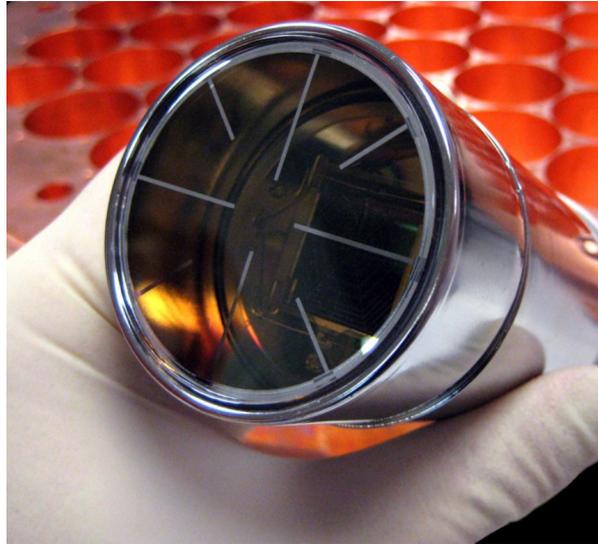


Figure 4.2: The Hamamatsu R8778 PMT used in LUX, showing the PMT window with the Al pattern [162].

4.1.2 Detection Efficiency

Even though PMTs are particularly sensitive instruments, a photon hitting the photocathode does not always result in successful detection. The detection efficiency (DE) of a photomultiplier tube is defined as the product of the quantum efficiency (QE) and collection efficiency (CE) on the first dynode. Quantum efficiency is the probability that an incoming photon will be converted into a photoelectron emitted into the vacuum once it hits the PMT photocathode. This probability is defined as:

$$QE = \frac{\text{Number of photoelectrons emitted}}{\text{Number of incident photons}} . \quad (4.1)$$

As expected through this definition, QE depends on the energy of the incident photons and hence their wavelength. The process of successful photoemission involves three steps: (1) absorption of a photon resulting in the transfer of energy from photon to electron, (2) diffusion of the electron toward the material-vacuum interface, and (3) escape of the electron into the vacuum. The first process is linked to the photocathode material, while for the second process both the quality of the sensitive layer and the temperature have a significant effect, as drifting electrons scatter off phonons and lattice imperfections. For the final process, the remaining energy of the electron has to be sufficient to overcome the potential barrier, which again depends on the material.

For instance, metals exhibit low QE ($\sim 0.1\%$) as a significant amount of energy is lost in all three stages. In the photoemission stage a large proportion of light is reflected, while

in the electron drift stage, photoelectrons lose energy rapidly in collisions with the large number of free electrons. Finally, the work function of most metals is higher than that of semiconductors, further reducing the chance of successful electron emission. Conversely, semiconductors absorb a much higher fraction of the incident light and energy loss by electron-electron scattering is low. The mean QE for the LUX PMTs at the mean scintillation wavelength of LXe (175 nm [163]) is $\sim 33\%$, with a standard deviation of 2% for the LUX sample [135].

As mentioned, an electrostatic field arrangement exists within the PMT which is responsible for directing the ejected photoelectrons towards the first dynode and then along the dynode chain to built up the multiplication process. The field in the input optics section at the front of the PMT is designed to increase the collection efficiency (CE) of the first dynode. This parameter largely depends on the electrostatic design and hence the PMT size and geometry. For the LUX PMTs a typical CE value is 0.9 [135].

4.1.3 Double photoelectron emission

Usually, when an incident photon is successfully detected by a PMT, a single phe is produced. If that were always the case, the quantum efficiency given in Equation 4.1 would be equal to the probability of a photoelectron being generated by a photon, $\eta(\lambda)$. Recently, it was found that at VUV wavelengths DPE — whereby two "photoelectrons" are detected in response to a single incident photon — can take place with sizeable probability [164]. The quantum efficiency is then:

$$QE(\lambda) = [1 + F_{DPE}(\lambda)]\eta(\lambda) , \quad (4.2)$$

where F_{DPE} denotes the DPE fraction, defined as the number of DPE events divided by the sum of single and double photoelectron emission events.

Studies have shown that for visible light wavelengths only single photoelectron (SPE) signals are observed. For shorter wavelength illumination, the impact ionisation probability, whereby the initial photoelectron releases a second electron within the photocathode, rises. The QE wavelength dependence is due to both the impact ionisation probability and the ability of the emitted electron(s) to reach the vacuum. More specifically, the QE initially slowly decreases as photon wavelengths decrease into the UV region. In this energy range two electrons are likely to be created but it is unlikely both will have enough energy to escape to the vacuum. For even shorter wavelengths QE increases again, with DPE emission taking place. Naturally, if the secondary electron is produced near the surface of the material, the probability of escape of the two electrons is greater. As with standard SPE emission, a lower temperature results in lower energy loss while the electrons drift through the material and hence the DPE fraction is observed to increase for decreasing temperatures. Naturally, DPE

emission interacts with the first dynode collection efficiency (CE), although this is a complex dependence: the emission point on the photocathode is the same for both electrons, but their loss through elastic scattering off the first dynode is uncorrelated.

The ~ 1 eV bandgap of the best photocathode materials means that for VUV photons (~ 7 eV) the initial photoelectron has sufficient energy to produce one or even more electron-hole pairs through impact ionisation. Then the electrons have to drift to the surface and overcome the potential barrier to the vacuum for DPE response to take place. For the detection of liquid xenon scintillation in cryogenic conditions the DPE fraction is $F_{DPE} \approx 17\%$ for the LUX PMTs [164], and as much as $\approx 23\%$ for the Hamamatsu R11410 model [165] used in forthcoming experiments such as LZ. In general, this probability varies significantly with PMT model and, in particular, with photocathode technology. Our group's measurement of this VUV effect for an LZ PMT is shown in Figure 4.3.

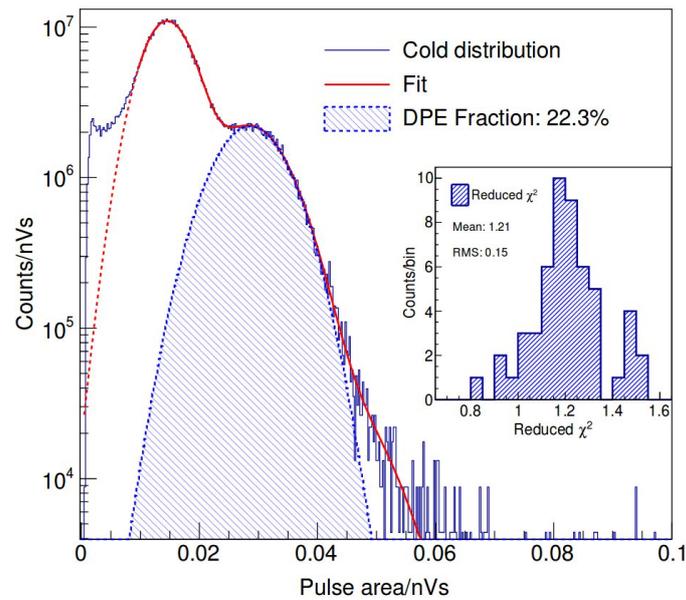


Figure 4.3: Measurement of the single photon response for an R11410-22 PMT used in LZ [165]. A hint of triple-photoelectron emission can be seen at large pulse areas. The inset shows the distribution of reduced χ^2 for the respective fit for the sample of 35 PMTs measured in that study. The measurement shown was performed at -97.4°C .

4.1.4 Afterpulsing

Ideally, the PMT vacuum would be held at $10^{-7} - 10^{-6}$ Torr. However, outgassing of the PMT materials or vacuum micro-leaks often compromise the vacuum of photomultipliers. Electrons accelerating within the PMT occasionally collide with atoms or molecules in this residual gas generating positive ions. Due to the applied voltage, these positive ions accelerate towards the front of the tube and eventually hit the photocathode, producing additional electron signals.

The ion-induced signals appear several tens of nanoseconds and up to a few microseconds after the main pulse and are known as afterpulses, and this ion-feedback effect is the main cause of afterpulsing in PMTs. The time delay (τ) between the original pulse and an afterpulse is the time taken for the ion to travel between its creation site and the photocathode, and is characteristic of the ion species and more specifically its mass-to-charge ratio. The afterpulsing time delay for each species has been measured to be constant as collisions usually take place in front of the first dynode, where the electron energy is higher [166]. The main afterpulsing populations observed in measurements using LUX PMTs include H^+ ($\tau \sim 0.3 \mu\text{s}$), He^+ ($\tau \sim 0.5 \mu\text{s}$), N^+O^+ ($\tau \sim 0.8 \mu\text{s}$), Ar^+ ($\tau \sim 0.6 \mu\text{s}$), Xe^+ ($\tau \sim 2.8 \mu\text{s}$) [162].

4.1.5 Dark current

The response of a photomultiplier tube in the absence of light is known as its dark current. Dark counts are produced through a number of different processes. Thermionic emission is the process through which single electrons are spontaneously emitted from the photocathode (due to its low band gap and work function) and amplified by the gain of the multiplier section. This rate (r_{DC}) scales with temperature (T) and photocathode surface area (A_c) according to Richardson's law:

$$r_{DC} \propto A_c T^2 \exp\left(-\frac{1}{T}\right). \quad (4.3)$$

At the LUX operating conditions, thermionic emission is the main contributor to the PMT dark rate, seen at a pulse size of 1 photoelectron (along with a different effect which will be discussed later in this chapter).

Dark counts can also be the result of intrinsic radiation from PMT component materials (such as ^{40}K , ^{238}U , ^{232}Th , ^{60}Co); radioactivity interacting in the PMT window causes Cherenkov radiation and other forms of luminescence. The small PMT radioactivity of the R8778, combined with the small probability of occurrence of this process, makes intrinsic radiation one of the smaller contributors to the dark count rate, especially at small pulse sizes. For very high operating voltages (higher than typical voltages used in LUX) regenerative effects such as dynode glow and glass charging effects become the dominant contributor to the dark rate.

In most LUX analyses a pulse must create signals in at least two PMTs within a certain time window (100 ns) to be classified as a valid scintillation pulse. The majority of dark counts are hence dismissed due to this condition and do not pose an issue for standard analyses. Our work focuses entirely on single photon signals and hence dark counts are of particular importance in our study. We will return to this topic in the next section.

4.2 Calibrating the single photoelectron response in LUX

4.2.1 Gain calibration using the LED system

In LUX PMT signals were constantly recorded by the DAQ system. Nominally, an SPE generated a pulse with 4 mV amplitude for the average PMT gain of 4×10^6 . Gain calibration refers to the process by which the mean of the single photon response is determined. Once this is known, the PMT can be used for photon counting. The process involves dividing the area of the raw signal (in pVs) by the calibrated mean area of the response for a single photon, yielding the number of detected photons in the signal. For this calibration, a source of single photons is needed. In LUX the in situ PMT gain calibrations were performed with 12 blue (430 nm) LEDs, placed inside the detector. The LEDs had PTFE diffusers to ensure a more uniform illumination of the PMT arrays. These were pulsed at a rate of 1 kHz and a pulse width of 100 ns. The 430 nm light pulses used do not produce double photoelectron emission, causing a bias that was corrected in data analysis and is further discussed at a later stage.

A POD threshold was enforced to ensure that noise originating from the data acquisition electronics chain was not mis-recorded as an SPE. A typical SPE pulse is illustrated in Figure 4.4 along with the 1.5 mV POD threshold and the height of a 5σ noise fluctuation. The SPE height distribution (using the LED system) for PMT 31 is shown in Figure 4.4 along with the SPE acceptance for this channel. Only two PMTs had an SPE digitisation acceptance below 0.90, with the mode of the SPE acceptance distribution falling above 0.95.

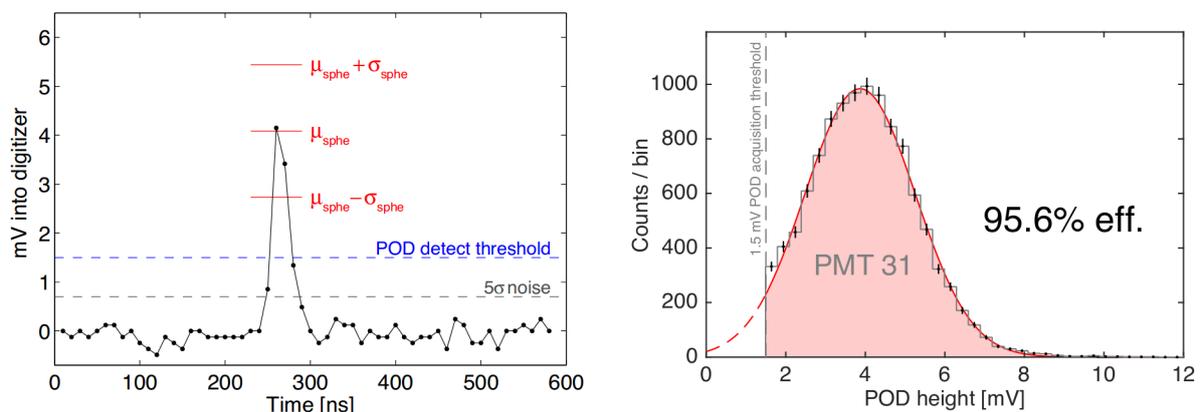


Figure 4.4: *Left:* Typical SPE pulse along with the POD threshold and noise expectation [135]. *Right:* Single photoelectron pulse height spectrum for an example PMT, where the 1.5 mV POD threshold is enforced and indicated by the vertical dashed line. The SPE acceptance, based on the Gaussian fit shown here, was estimated to be 95.6% for this example PMT [135].

4.2.2 Single photoelectron distributions using dark counts

In this work we used data from a tritium calibration (Dec. 2013 injection) with a low event rate (peaking at ~ 0.1 cts/s), allowing both the dark count population and the VUV response to be studied in the same dataset. Tritium data is ideal for this since it produces very clean, low-energy single scatter events, with the vast majority of waveforms containing a single S1 pulse and a single S2 pulse. In this section we focus on the dark counts, which constitute a population of waveforms consisting mostly of SPE pulses and hence can be used to calibrate the SPE responses on a PMT-by-PMT basis. This calibration, along with the dark count rate measurements, are essential for our analysis as dark counts constitute the main background in our study.

Pulses classified as SPE-like were selected when found either preceding an S1 or identified between the S1 and the S2 pulses and hence consisting mostly of PMT dark counts. For pulses following the S1 pulse an afterpulse timing cut was also included. A distribution is shown in Figure 4.5, where all channels in the top and bottom PMT arrays have been summed for the purpose of display.

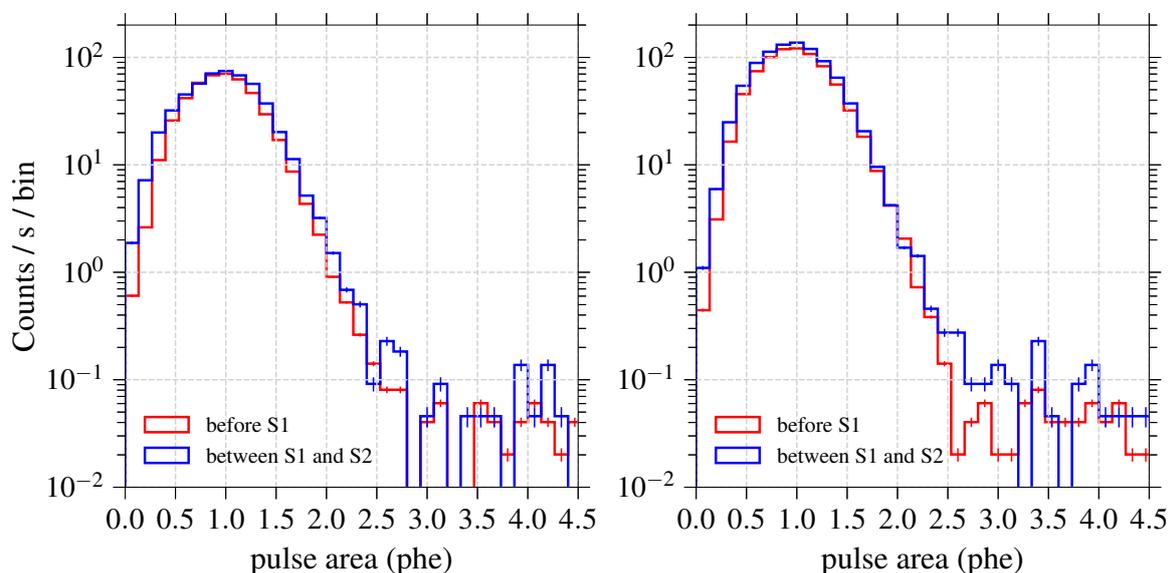


Figure 4.5: Pulse area distributions (in phe) of SPE-like pulses summed across the (*left*) top and (*right*) bottom arrays. These were identified in the quiet waveforms preceding tritium S1 pulses (in red), or between the tritium S1 and S2 pulses (in blue). These distributions do not include contributions from 10 channels that were removed for this analysis.

The distributions for individual channels were fitted with a Gaussian model using the likelihood method to obtain the mean and width of the SPE response (μ_{DC}, σ_{DC}), as well as the DC rate for each PMT. The fit was restricted to start at 0.3 phe to allow for the POD digitisation threshold (and clearly the Gaussian model will no longer be appropriate for very

small pulse areas). Ten PMT channels with outlier fit parameters, goodness-of-fit estimators (χ^2), or abnormally high dark count rates — often found to be detecting spurious light either due to PMT malfunction or from occasional electrical discharging leaking into the TPC — were excluded from this analysis for similar reasons. There were also three PMTs that were previously "switched off" in the standard analysis. The remaining 109 channels will thereafter be referred to as "good PMTs". The mean and resolution of the SPE distribution along with the goodness of fit ($\frac{\chi^2}{\text{dof}}$) for the good PMTs are shown in Figure 4.6. Example distributions along with the Gaussian fits are shown in Figure 4.7.

It is worth noting that four PMTs with $\chi^2/\text{dof} > 3$ are included in the set of good PMTs. The Gaussian model does not fit the peak of the distribution well for these particular channels (not plotted in Figure 4.7); however, the overall fit was deemed satisfactory and these PMTs did not exhibit spurious light emission, thus they were used for the analysis. PMT 9 also observed an outlier dark count rate compared to the mean DC rate of the top array, but given that the distribution was still indicative of SPE response and this rate was not higher than that accepted for the bottom PMTs it was included in the analysis.

The average dark count rates for the "good" PMTs measured using the tritium calibration are summarised in Table 4.1. We repeated the dark count rate measurement using the WS2013 search dataset and the results are also presented in Table 4.1. Several interesting observations result from these data, which are also apparent in Figure 4.5. The DC rate recorded at the bottom array is significantly higher than that at the top array; and both are much lower than the ~ 40 cts/s/PMT observed in low temperature in tests on the surface. These effects are also observed in other LXe-based experiments [167]. The latter effect is presumably due to the low radiation environment in LUX (e.g. reducing Cherenkov emission in PMT windows caused by Compton electrons and other charged particles related to ambient radioactivity). The difference between the two arrays is more interesting; since the bottom array has a higher photon collection efficiency for interactions in the liquid, we believe that the rate measured there consists partly (or even mostly) of detected light, and is not attributed solely to thermionic emission from the photocathodes. This light did not undergo double photoelectron emission and hence was not of VUV wavelength. A likely explanation is that these photons are due to PTFE fluorescence at optical wavelengths, which is induced when VUV photons are absorbed in the PTFE (see [168, 169]). Other manifestations of this effect have previously been observed in LUX. This ability to distinguish between VUV and visible photon fluxes is a useful application of the DPE effect.

The above observations allow us to decompose the observed dark rates into a thermionic, R_{th} , and a visible light component, R_{vis} :

$$\text{Total top rate} = N^{\text{T}} R^{\text{T}} = N^{\text{T}} R_{\text{th}} + R_{\text{vis}}^{\text{T}} = N^{\text{T}} R_{\text{th}} + \phi g_1^{\text{T}} , \quad (4.4)$$

$$\text{Total bottom rate} = N^{\text{B}} R^{\text{B}} = N^{\text{B}} R_{\text{th}} + R_{\text{vis}}^{\text{B}} = N^{\text{B}} R_{\text{th}} + \phi g_1^{\text{B}} , \quad (4.5)$$

where N represents the number of PMTs, ϕ represents the flux of visible light component contributing to the dark rate, which we assume constant inside the (high-reflectance) detector, and superscripts "T" and "B" refer to the top and bottom arrays, respectively. The top array and bottom array specific g_1 components represent the average light collection efficiency of each array. Using these equations and our knowledge of the different coupling to light between the top and bottom arrays we can estimate the rate of thermionic emission. Specifically, the ratio of the rate detected at the top to that detected at the bottom array is given by:

$$\frac{R_{\text{vis}}^{\text{T}}}{R_{\text{vis}}^{\text{B}}} = \frac{g_1^{\text{T}}}{g_1^{\text{B}}} = \frac{S1^{\text{T}}}{S1^{\text{B}}} \approx 0.26 . \quad (4.6)$$

The ratio of S1 light observed at the top to the light observed at the bottom ($\frac{S1^{\text{T}}}{S1^{\text{B}}}$) was estimated using $^{83\text{m}}\text{Kr}$ calibration data. Using this information, and the rates given in Table 4.1, the thermionic emission component was estimated to be $R_{\text{th}} \approx 6$ cts/s. This is an important observation, as we can deduce that, for the bottom array, detected light and not thermionic emission is the main source of the "dark" count rate. The photocathode diameter for the LUX PMTs is 45 mm. Given the linear dependence of photocathode area to the thermionic rate, we can express this dark rate as $R_{\text{th}} \approx 0.48$ /s/cm². Xenon1T has reported a measured dark count rate of ~ 12 and ~ 24 cts/s for their top and bottom arrays of 3-inch R11410 PMTs, respectively [167]. Assuming the same ratio of light seen at the top over bottom array as in LUX (also a PTFE-lined TPC), and taking into consideration the photocathode diameter of the R11410 model (64 mm), we estimate a thermionic rate of $R_{\text{th}} \approx 0.28$ /s/cm². This is a reasonable agreement considering the different photocathode technologies employed by the two models. More generally, the very low dark count rate per unit sensitive area in these PMTs is key to achieving low scintillation threshold in LXe detectors.

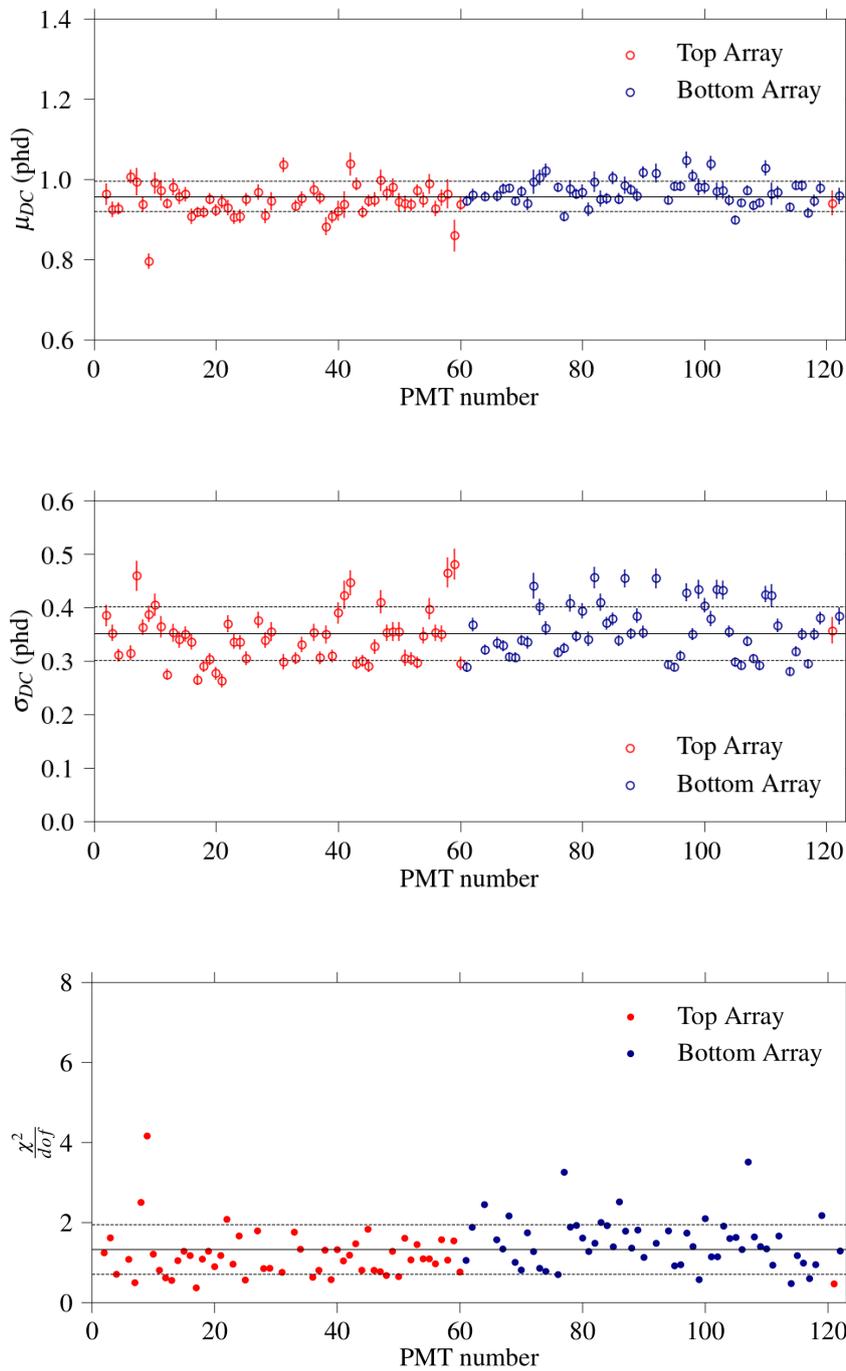


Figure 4.6: Scatter plots illustrating the mean and resolution of the SPE pulse area distribution as well as the reduced χ^2 for each PMT for distributions of dark count pulses. Ten channels with outlier fit parameters are excluded from the analysis and not shown here. The horizontal lines indicate the mean and standard deviation of each presented quantity for the remaining 109 channels.

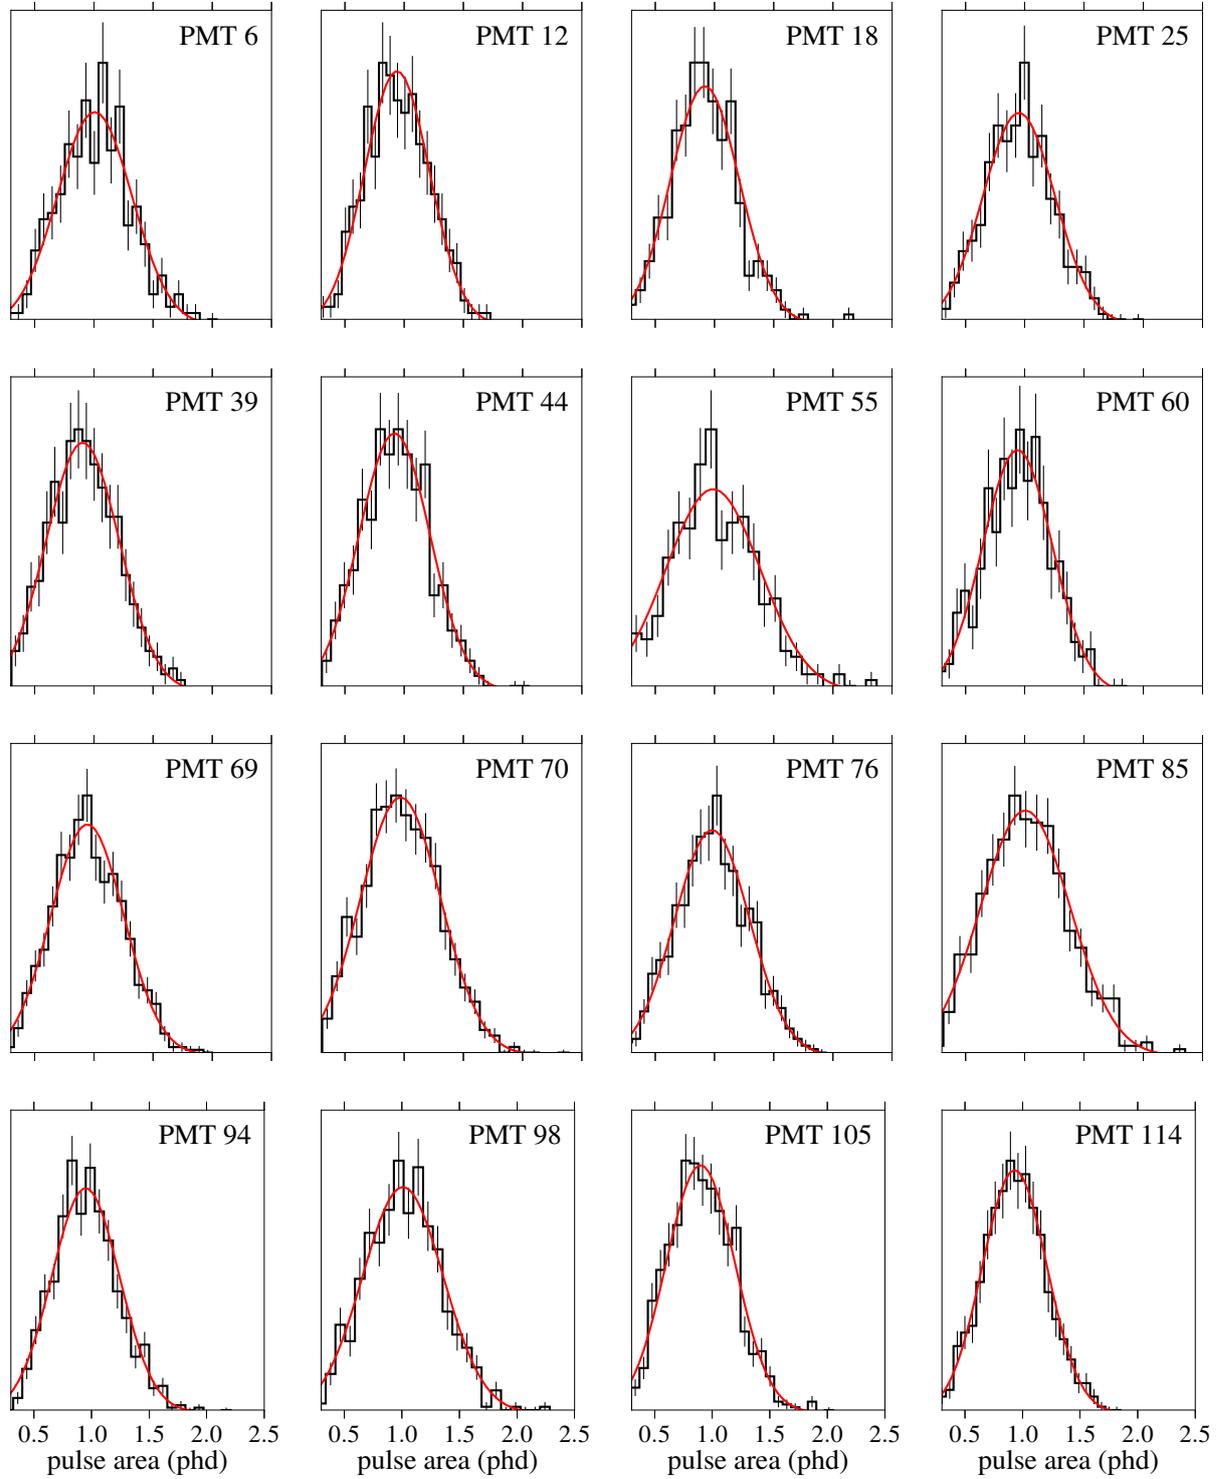


Figure 4.7: Pulse area distributions and fits for the single photoelectron response obtained using dark count populations for a representative set of PMTs.

Table 4.1: Average dark count rates for the PMTs used in this analysis as measured both via the tritium calibration and the 2013 WIMP search dataset. These rates were measured using populations of single scatter events, searching before the S1 pulse (<S1 window) and between the S1 and S2 pulses (S1–S2).

Window	Array	Tritium (cts/s)	WS data (cts/s)
<S1	Top	8.2 ± 0.4	8.0 ± 0.7
	Bottom	14.7 ± 0.5	14.0 ± 0.9
S1–S2	Top	9.4 ± 0.4	9.6 ± 3.1
	Bottom	16.8 ± 0.6	17.2 ± 1.7

Finally, there is a small difference between the rates measured in the two time windows (before S1 and between S1 and S2 pulses), both in measurements using the tritium dataset and during the WIMP search, which is unlikely to be statistical in nature. This appears at a small time window following the S1 and is more pronounced after larger S1 pulses, indicating it is likely due to S1-induced PTFE fluorescence. In the remainder we continue to refer to these SPE signals as "dark counts", but bearing in mind that most are actually photon-induced. There are currently ongoing measurements on the LUX (and other) photomultiplier tubes, in which the photocathode is blocked, to identify the true dark count rate (without a PTFE fluorescence component) and hence establish the robustness of the determinations presented here.

Observing Figure 4.5, we conclude that there is reasonable agreement between the distributions before and after the S1, with no significant excess observed near two photoelectrons. To be more exact, even though the distribution of SPE-like pulses between S1 and S2 is overall broader than the distribution for pulses recorded before the S1 pulse, the difference does not agree in terms of shape with a 2-phe peak but instead with a higher σ of the distribution. Even if we attribute any difference around the 2-phe area to DPE emission fully it would still only indicate an additional DPE component contributing less than *sim*1%. This indicates that no significant VUV photon sources exist outside of the main xenon luminescence mechanisms associated with S1 and S2 photon production, and hence no major backgrounds are expected that could make our analysis non-viable — this is one of the main results from our study. Nonetheless, the SPE distributions in Figure 4.5 do include a modest tail to large pulse areas, which we attribute to the way in which the LUX pulse classifier identifies SPE-type pulses, which is a loose requirement: those must be S1-like in shape and recorded in a single PMT within a 100 ns window, regardless of size [135]. They would allow, for example, light pulses emitted within the PMTs themselves [170], and Cherenkov signals in the PMT windows — in underground experiments the latter may be due to Compton electrons generated by background gamma-rays [171].

4.3 Characterising the single photon response in LUX

We have established that, through the use of LED photons or dark counts, we can effectively calibrate the single *photoelectron* response of each photomultiplier tube. Neither of these processes includes the effect of double photoelectric emission and are hence not representative of the PMT response to xenon scintillation *photons*. Once the existence of double photoelectron emission was recognised, it was realised that some early LUX analyses were biased by not accounting fully for this effect, with the number of photons in calibrated signals overestimated. The LUX pulse-parametrization algorithms were corrected to account for this effect, with the size of S1 and S2 pulses expressed in 'photons detected' (phd) instead of the traditional phe unit obtained by pulse charge integration. For larger pulses this involved an average correction factor per PMT, while for smaller (S1) pulses an algorithm for "spike" counting was introduced, where each spike was an excursion of the waveform above some (low) threshold. These corrections brought about significant improvements in both the linearity and the energy resolution of the detector [119, 127].

A full characterisation of the single photon response, including both SPE and DPE components, has not previously been pursued for the LUX PMTs, as obtaining sources of single photons with adequate statistics was difficult and not necessarily needed for most analyses. Our study is based purely on single photon signals and hence a complete understanding of the single photon response, including the DPE response was essential.

4.3.1 Calibration using S1 light

The full single VUV photon response was measured using scintillation light from tritium interactions. Tritium β^- decays are an excellent calibration source for the single photon signals, as they are an internal source that becomes uniformly distributed throughout the detector and provide high statistics in all areas of the TPC. Hence, it is generally unlikely for individual PMTs to detect several photons from one tritium decay, especially when restricting the S1 size to obtain only the smallest of interactions and the ER location to be far away from the PMT in question.

However, to correctly determine the single photon response, we must make sure our acquired distributions are overwhelmingly composed of single photon signals. Poisson statistics were used to understand and evaluate the contribution of two-photon hits to our single photon populations. The Poisson mean μ for the signal generation in this process was inferred using the fraction of events that do not have a signal, $P(0; \mu)$ – i.e. by counting the number of

interactions for which no photons are detected in a particular PMT, as follows:

$$\mu = -\ln(P(0; \mu)) = -\ln\left(\frac{\# \text{ triggers with no signal}}{\# \text{ triggers}}\right). \quad (4.7)$$

In our study two sub-volumes were defined, within the fiducial volume of the TPC, with the lower sub-volume used for the top array calibration and the upper sub-volume for the bottom array calibration. The PMTs were divided into six different groups (three for each array), according to their location in the array. Different S1 size cuts were applied for each of these groups. The area cut on the S1 had to be different for the top and bottom PMT arrays since the light collection efficiency is different. Additionally, because of their specific position some peripheral PMTs are less well coupled optically to the centre of the detector and hence required different cuts in order to detect enough light for the distributions to be fitted. The study also included two different S1 cuts, for each of the six groups, to assess the systematic effects due to the amount of light seen by the arrays. We will refer to the two cuts as the aggressive and relaxed S1 selection cuts. Overall, the selection criteria were optimised such that the Poisson mean expectation per channel was both small and constant ($\mu \lesssim 0.03$ phd) and that the ratio of two-photon hits to one-photon hits producing 2 phe signals was lower than $\sim 5\%$. These conditions allowed us to recover populations of single photons for each PMT. The Poisson mean expectation and the contamination ratio are given in Figure 4.8.

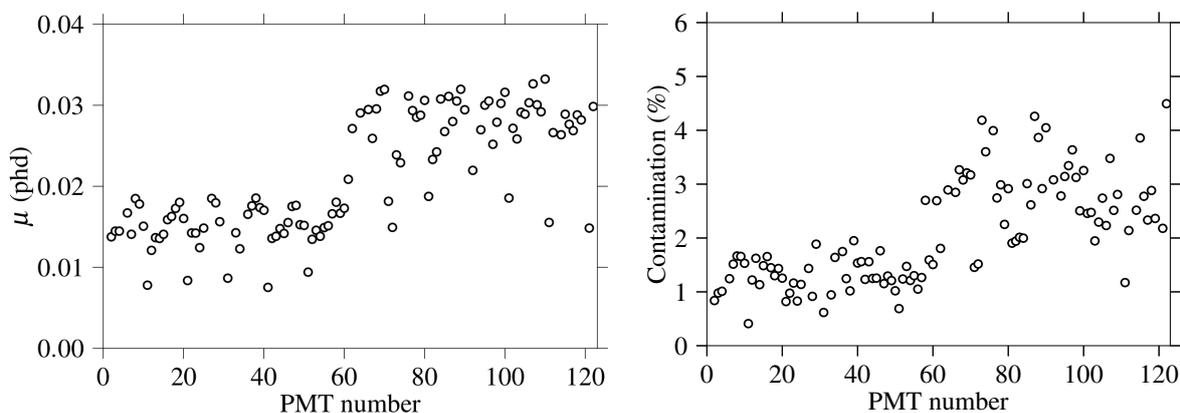


Figure 4.8: The (*left*) Poisson mean and (*right*) contamination ratio of two-photon signals to 2 phe signals for each good PMT.

The full single photon response for each PMT was subsequently fitted using a double-Gaussian model with PMT calibration parameters that were previously fitted using the dark count distributions used as initial estimators. The fit parameters in this case consisted of the SPE mean (μ_1), the SPE resolution (σ_1), the DPE mean and resolution (μ_2 and σ_2) and the DPE fraction (F). The DPE fraction was corrected using the expected number of two-photon hits. The fit parameters, the corrected DPE fraction and the goodness of fit estimator are shown in

Figures 4.9 and 4.10. We present the parameters recovered using the relaxed S1 selection cuts, as they benefit from greater statistics. The DPE fraction presented uses both measurements (the aggressive and relaxed S1 cut approaches), corrected from the expected contamination, with uncertainties including both statistical and systematic components (estimated using the method discussed below). A range of example fits with the high and low S1 allowances are illustrated in Figures 4.11 and 4.12, respectively.

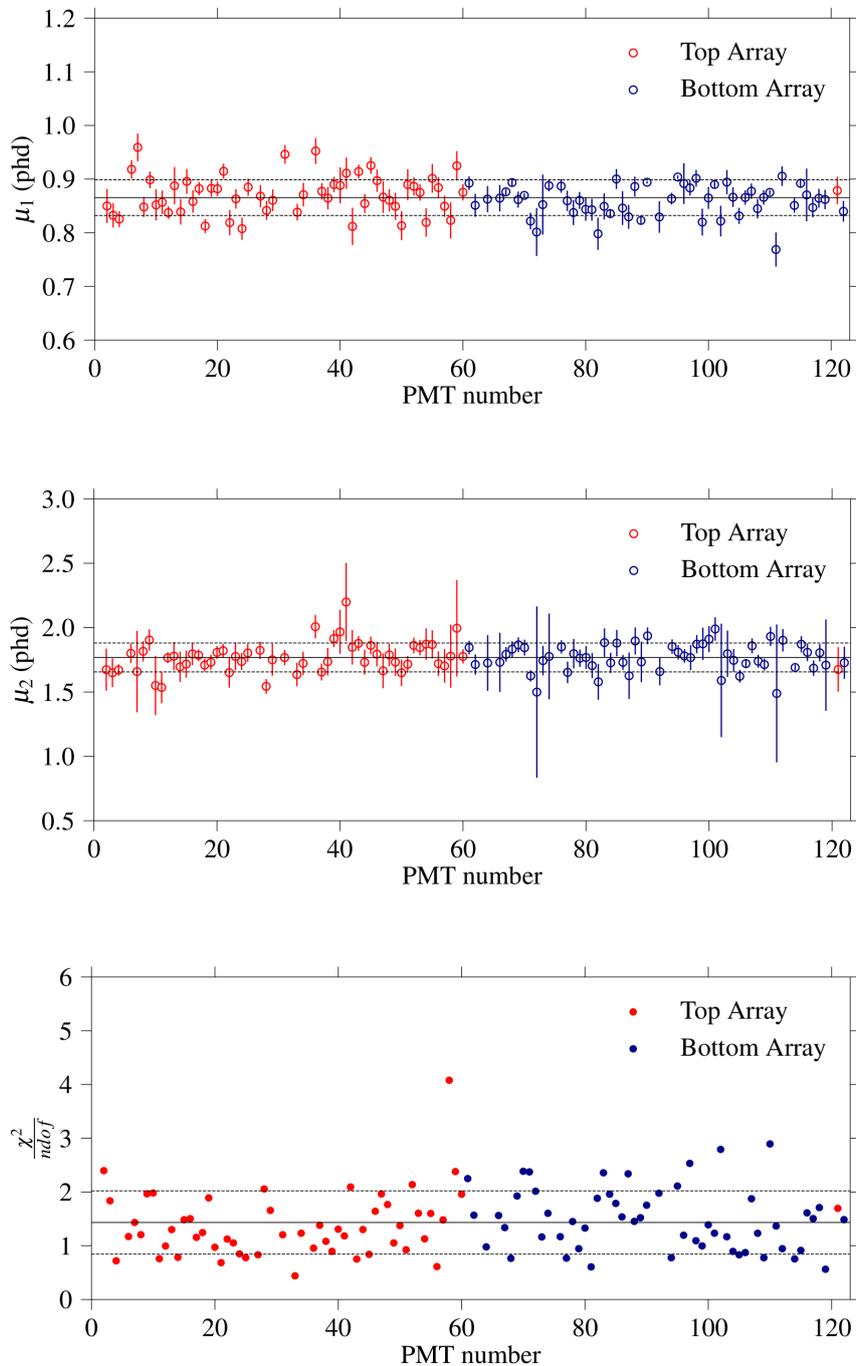


Figure 4.9: Scatter plots illustrating the mean of the SPE and DPE pulse area distributions and the goodness of fit for each good PMT, for distributions of single photons obtained using tritium interactions (including both an SPE and DPE response).

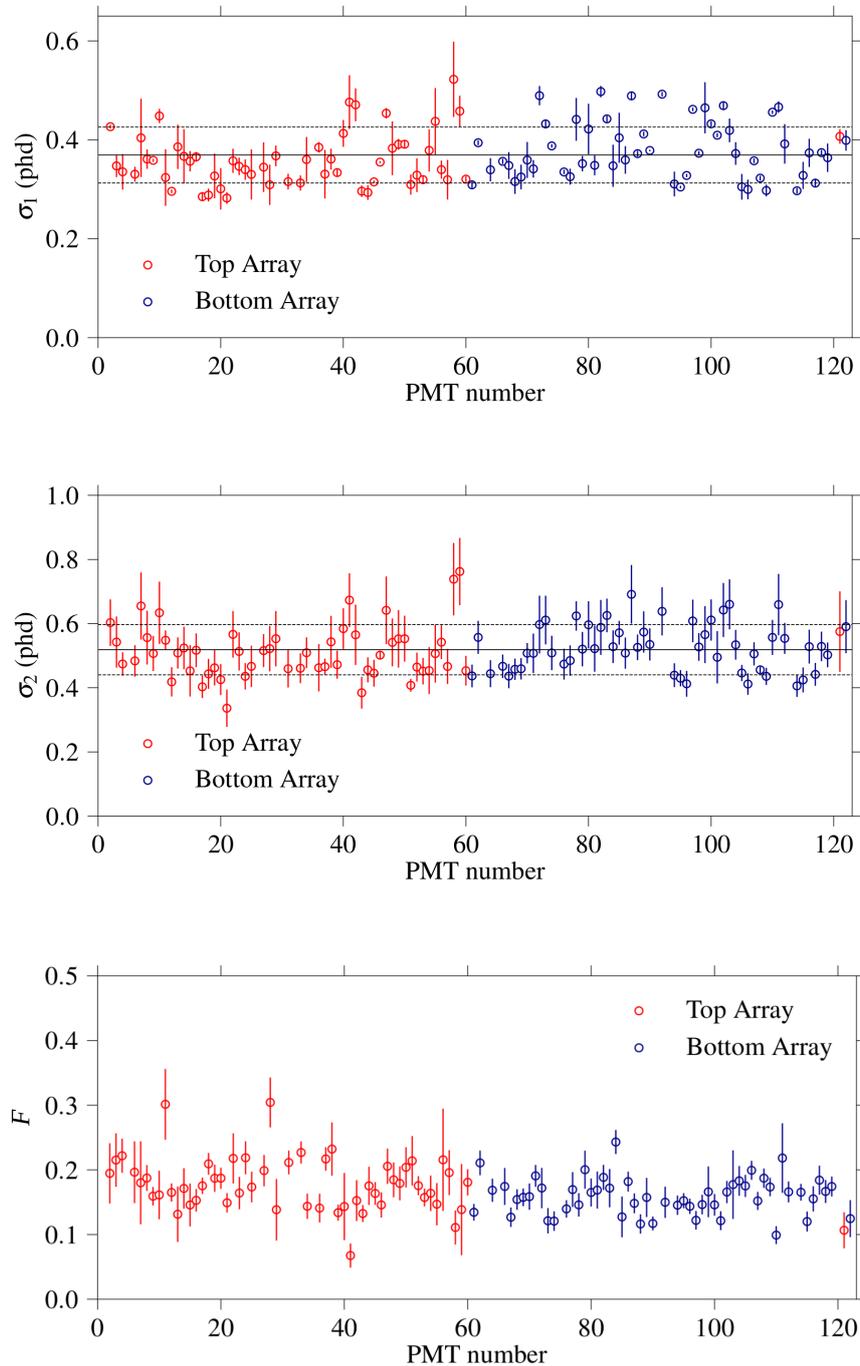


Figure 4.10: Scatter plots illustrating the resolution of the SPE and DPE pulse area distributions and the DPE fraction for each good PMT, for distributions of single photons obtained using tritium interactions (including both an SPE and DPE response).

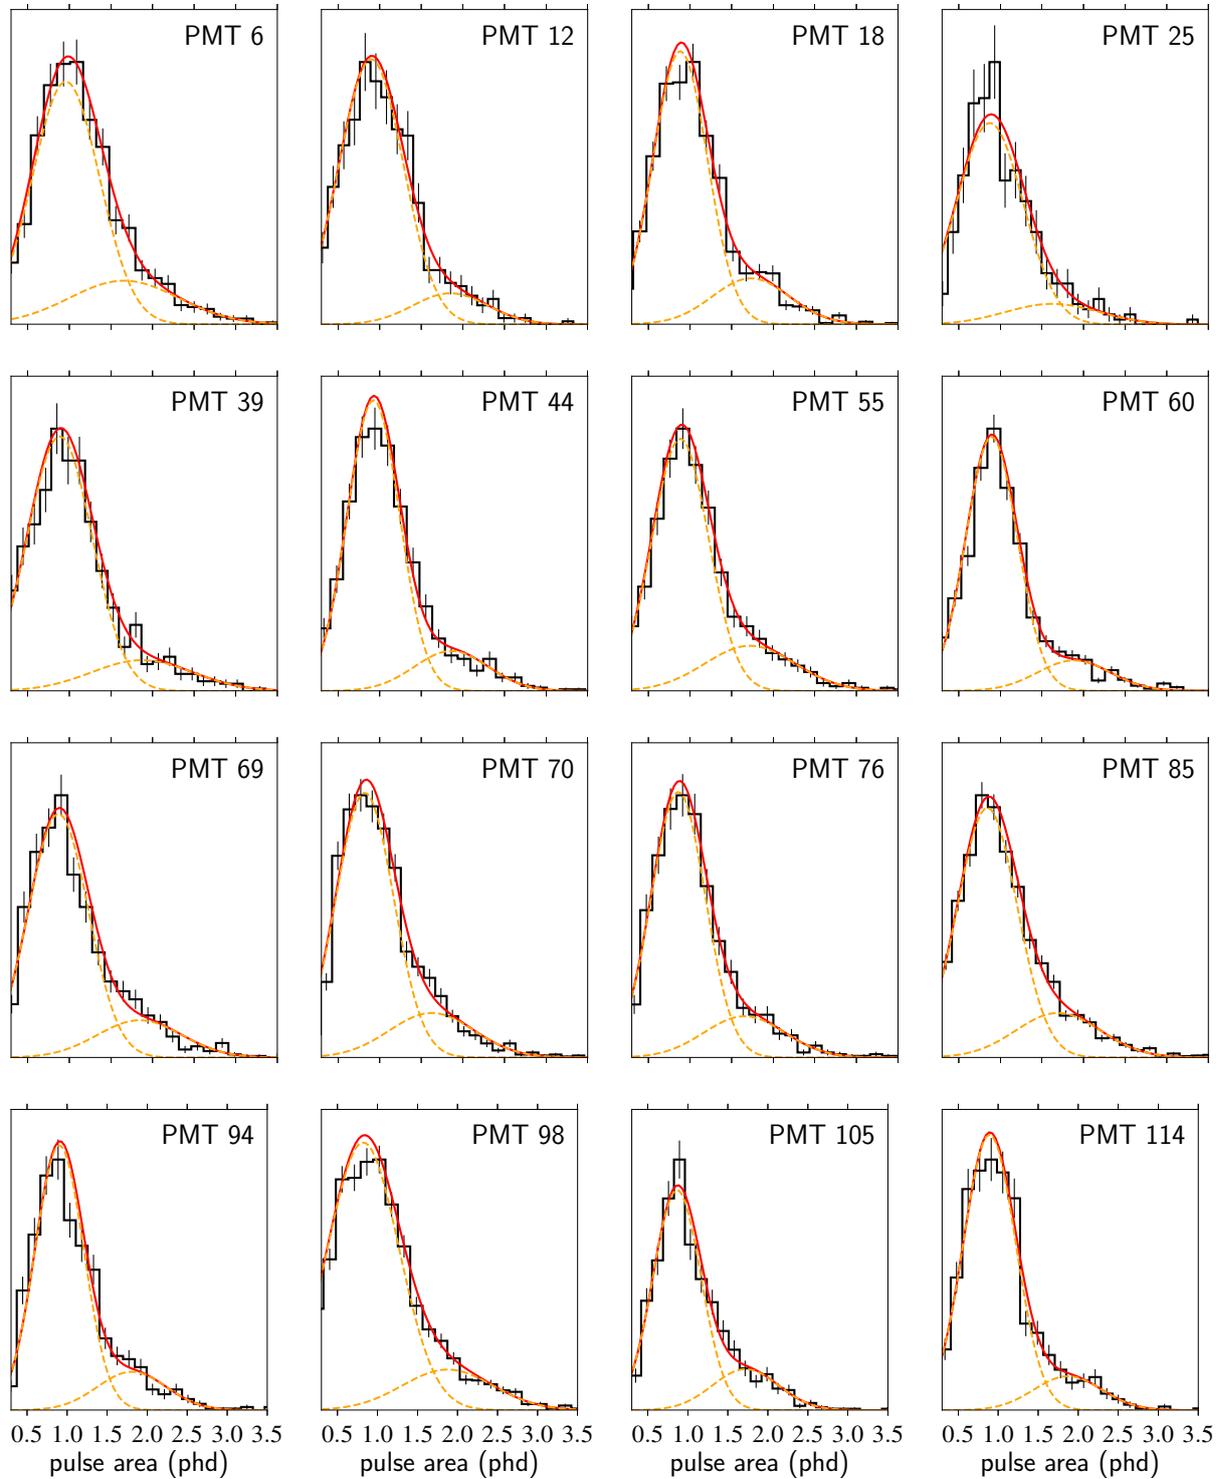


Figure 4.11: Pulse area distributions for the single VUV photon response obtained using S1 pulses in the tritium calibration dataset. Distributions illustrated here were obtained using the higher S1 allowance, as discussed in text. The DPE fraction for all PMTs is given in Figure 4.10.

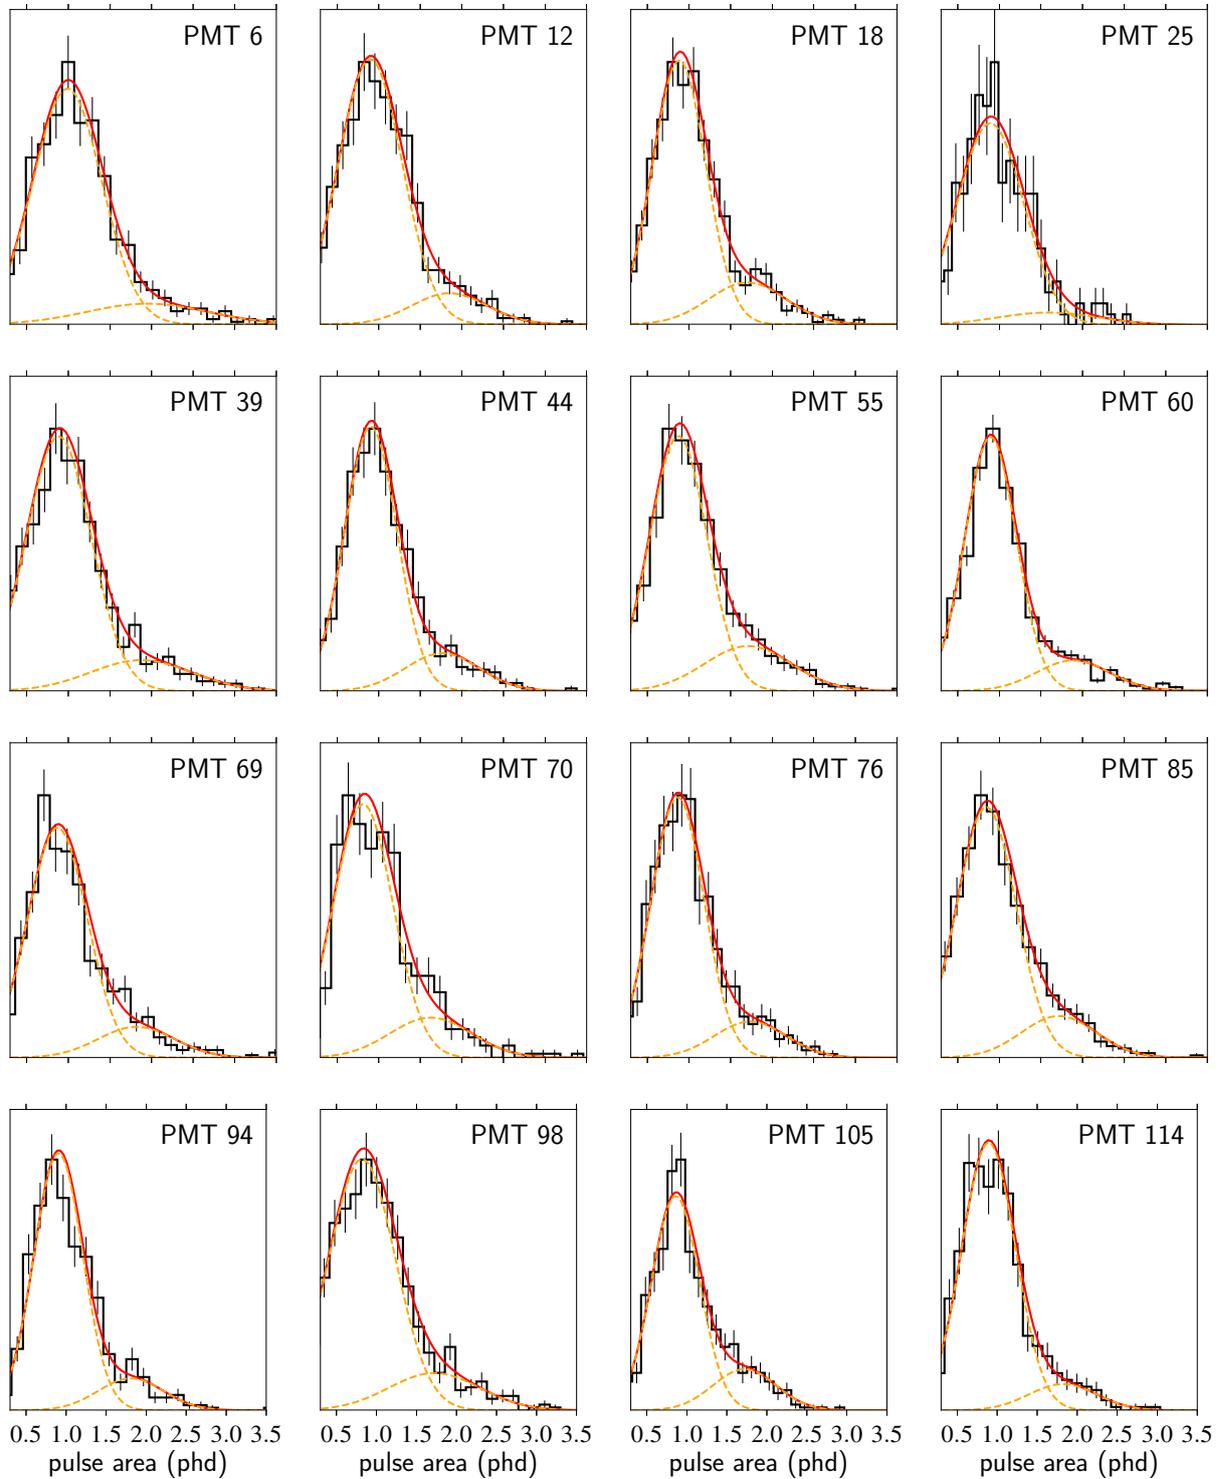


Figure 4.12: Pulse area distributions for the single VUV photon response obtained using S1 pulses in the tritium calibration dataset. Distributions illustrated here were obtained using the lower S1 allowance, as discussed in text. The DPE fraction for all PMTs is given in Figure 4.10.

After correcting for contamination from two-photon (2 phd) events, the mean DPE fraction was found to be $0.169 \pm 0.015_{sys} \pm 0.005_{stat}$, consistent with measurements from other LUX analyses. The statistical error presented here was propagated from the error on the fit value of the DPE fraction for each PMT (which are uncorrelated), while the systematic error was calculated by studying how the DPE fraction varied with S1 cut definition. This was proven useful as it allowed us to establish that the DPE fraction was not significantly affected by the amount of light observed by the PMTs, evident through the acceptable level of systematic error obtained for the accuracy of the presented analysis. Nevertheless, there was a small difference in the DPE fraction observed using the two measurements, indicating a small pileup component and hence bias in the measurement. In future analyses, this can be improved by implementing longer or more concentrated tritium calibration runs so that enough statistics can be obtained with even more stringent S1 cuts. Additionally, other calibration sources with smaller S1 signals could be considered (as smaller S1s would lead to a smaller photon pileup probability on individual PMTs).

Overall, the DPE fraction was found to vary significantly between different PMTs, confirming the need for the independent channel by channel calibration, with the sample standard deviation found to be 0.04.

An example of the single photoelectron and single photon responses (including both SPE and DPE components) for a typical PMT, fitted with Gaussian and double-Gaussian models, respectively, is shown in Figure 4.13. We note that both the mean (μ_{DC}, μ_1) and the resolution (σ_{DC}, σ_1) of the two SPE responses are somewhat different, and this was observed for other channels. This effect is primarily attributed to the fact that the response measured with photons includes direct photon hits to the first dynode, biasing the mean of the distribution to lower values [164]. Additionally, in LUX a maximum of 10 pulses per event were parameterised in the standard data reduction, and small-area DCs may be lost; this effect was found to account for an increase on the mean of the dark count SPE distribution by $\sim 2.5\%$. In any case, this inefficiency for tagging SPE-like pulses has little effect on our DPE analysis ($< 1\%$), as for the low-energy events of interest the 10 pulse limit is rarely reached.

The DPE fraction mean for the LUX PMTs might be small, but for DPE emission with large pulse areas there is little contamination from dark counts. The optimal region for which the contamination is small and efficiency for DPE events large must hence be defined. One such region is indicated in Figure 4.13 and this optimisation is discussed in the next chapter.

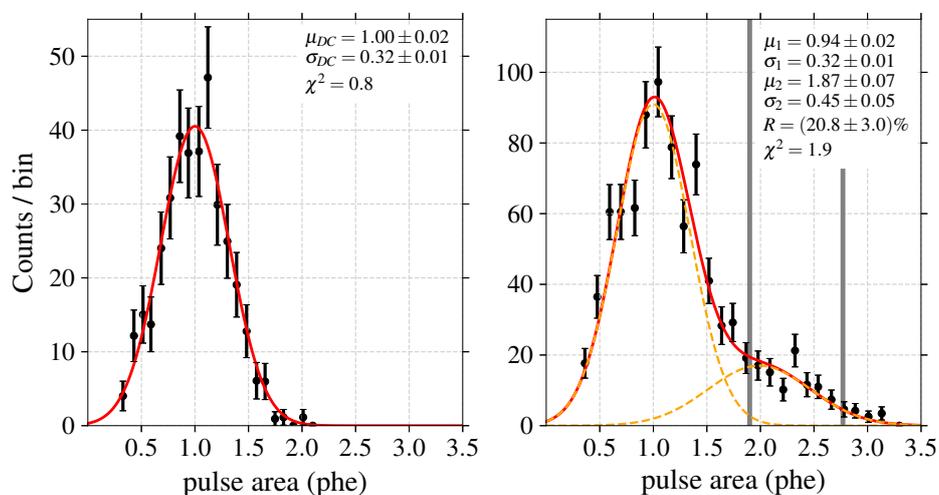


Figure 4.13: Pulse area distributions (in phe) for the (*left*) single photoelectron response and the (*right*) single VUV photon response of an example PMT, along with the fit parameters. These include the mean and standard deviation obtained from the dark count population (μ_{DC}, σ_{DC}), and the mean and standard deviation of both the SPE (μ_1, σ_1) and DPE responses (μ_2, σ_2) obtained from the single scintillation photon response. The DPE probability for this PMT is $F = (20.8 \pm 3.0)\%$.

4.4 Summary

In this chapter, we discussed the calibration of the single photoelectron and single VUV photon response for the LUX PMTs. We began by using dark counts to characterise the SPE response and dark count rates for each PMT. Important observations include the different dark count rates between the top and bottom arrays, likely due to the contribution of PTFE fluorescence to the dark count populations. A good understanding of the dark count responses is essential for our understanding of the background rate from accidental coincidences that contaminates our single-photon S1 searches. We observed no significant sources of VUV photons outside the main LXe mechanisms, validating our belief that lowering the scintillation threshold of LUX was viable. This is discussed in detail in Chapters 5 and 6. Additionally, the full VUV response and DPE fraction for each LUX PMT were calibrated using small S1 signals. These are used in Chapters 5 and 6 to estimate our signal acceptance. Both the calibration of the full single photon response and our search for possible backgrounds in that area could be improved in accuracy through the use of larger and more varied calibration datasets, a process we suggest for future experiments. In detail, a longer or more concentrated tritium dataset or an alternative source with smaller S1 signals could allow us to calibrate the single photon response with a smaller pileup probability, yielding more precise results.

Chapter 5

Exploiting the single VUV photon response in LUX tritium data

The ER response of LUX has been calibrated using both internal and external sources. One of the primary ER calibration sources deployed in LUX was tritiated methane (CH_3T), providing a high statistics calibration dataset of uniformly distributed events. Tritium is a β^- emitter with a Q-value of 18.6 keV and half life of 12.3 years. The energy range of this beta spectrum is not too dissimilar to that for WIMP searches, and beta decays in this range have similar properties to Compton scatters, which constitute the main background in LUX – making this source ideal for studying the ER response and ER-NR discrimination of the detector. After injection through the gas system, the tritiated methane was removed by the Zr getter. A main concern when using this source for the first time was the diffusion of tritium activity into detector parts in a way that could re-contaminate the LXe during data taking. To reduce the risk involved, initial studies involved pure methane to verify that this species could be removed effectively by the getter; this was monitored by a high-sensitivity gas sampling system. This was followed by the tritium calibration campaign which was performed following the first run of data taking, with a small dataset (~ 7000 events) initially obtained for prudence, before injecting a larger activity (10 Bq) of tritium in December 2013 ($\sim 170,000$ fiducial events).

In this chapter we use the December 2013 ER calibration dataset to extend the original analysis described in Ref. [119] to lower energy signals. We have previously described the calibrations of single and double photoelectron responses of individual PMTs and their dark count rates and with the methods used no significant new backgrounds were identified that could jeopardise a rare event search. We now set out to search for single-photon S1 events in the tritium calibration dataset in an attempt to optimise and validate the proposed single photon technique. This is part of the work published in Ref. [1].

5.1 Reproducing the standard tritium calibration result

In Section 3.5.2 we summarised the standard data processing procedures followed in LUX. As discussed, the last step of the data reduction procedures involved event filtering and calculation of final summed pulse quantities, corrected for geometric or other effects, through the use of the "filter code". The S1 2-fold coincidence requirement, the single scatter selection, the pulse area corrections and a cut on high activity periods during operation were enforced at this stage. Following a pass through the filter code, the data is reduced from a large set of RQs, at PMT level, to selection of summed over all channels quantities for a subset of events (for most analyses these are single scatters).

Most LUX studies used the resulting filtered datasets in their analyses. For our single photon work a less filtered dataset is needed, with the 2-fold coincidence requirement not enforced and individual PMT information present. We hence developed an alternative processing chain. To validate our chain and the appropriate application of corrections, we begun our work by reproducing the standard tritium calibration analysis. We must note at this point that there was an implementation of small improvements on the correction algorithms following the tritium calibration publication and the original code was not fully preserved, which is why we expected minor differences between our alternate and standard chain results.

The new (alternative) processing chain started with the standard pre-selection of single scatter events applied in most LUX analyses and discussed in Section 3.5.2. The efficiency of this event selection was found to be 98.8% [135]. Following this selection, the 2-PMT coincidence requirement was imposed on the S1 signals, selecting events in which two non-adjacent (i.e. non-partner) PMTs exceeded a 0.3 phd area threshold, within a 100 ns coincidence window. Subsequently, adapted versions of the algorithms discussed in Section 3.5.2 for S1 and S2 summed-pulse parameter estimation were implemented.

The filtered dataset was then analysed, with standard search analysis cuts applied. These cuts were originally optimised for the tritium calibration first presented in Ref. [135]. Firstly, a dataset "quiet time" cut was applied. The quiet time cut was applied such that short periods of livetime were excluded from the data if an abnormal pulse detection rate was observed. This was often the case following large energy depositions, where high "activity" was observed in the detector [153]. Subsequently, a fiducial volume selection was performed, with the radial cut placed at 18 cm, as in the standard tritium analysis and original WS2013 analysis. Optimisation of this cut was driven by the leakage of background interactions occurring on the detector walls. The height of the fiducial volume was defined to extend from 38 to 305 μ s in drift time to reduce backgrounds from the PMT arrays and electrodes.

Finally, two sets of pulse quality cuts were applied. One of the main targets of these cuts was background events due to electron emission. This process can involve either one or few-electron

signals and has been extensively studied, both in LUX and other xenon TPCs, with several viable explanations for the background populations observed proposed [153, 172]. The main categories involve photoionization electrons released by S1 or S2 photons absorbed by impurity atoms in the liquid, clustered electron emission that consists of multiple electron pulses in short periods of time following a small percentage (1%) of genuine particle interactions (thought to originate from unextracted electrons trapped below the liquid surface that can emerge into the gas together under certain conditions), delayed emission of individual electrons (trapped by impurities or at the liquid surface) and spurious emission from the grids. We further discuss electron emission backgrounds and their relevance to this thesis in Chapter 6.

For low energy analyses, these signals are particularly dangerous as a few-electron pulses can be incorrectly classified as an S2 and accidentally coincide with either S1-only events or misclassified S1 pulses to produce fake single scatter events. To eliminate such pathologies, an S2 pulse area threshold cut was applied to the uncorrected raw pulse areas and was optimised to exclude events during periods in which a high rate of single electrons was recorded, causing misidentified S2 pulses. This was placed at 165 phd, which is equivalent to ~ 6 electrons in the gas phase.

The second pulse quality cut, known as the waveform "cleanliness" cut, was applied on the "good" and "bad" area quantities. Good area is defined as the summed S1 and S2 pulse area in an event and bad area as the summed area of any other pulses present in the event window. During particularly high electron emission periods, several small pulses exist within events and hence the bad area increases. There is an expectation that the frequency of spurious signals increases with S1 and S2 pulse areas, and this cut reflects this behaviour — as shown in Figure 5.1. The full set of cuts including the quiet time, fiducial volume, cleanliness and pulse area cuts are shown in Figure 5.1.

Following the application of the quality cuts, the S1 and S2 corrected pulse areas were calculated. This involved correcting the S1 and S2 signals for the main sources of non-uniformity in the detector response. The S1 raw signal was corrected for the (x,y) dependent probability for signal photons to be detected by the PMTs, such that for each S1 pulse the corrected area matched the S1 response in the centre of the active region. Conversely, the S2 signal was corrected for both the spatial variation in (x,y) but also for the finite electron lifetime, associated with electronegative impurities in the xenon. The corrected S2 signal matched the S2 response in the centre of the TPC radially (x,y) and just below the liquid surface in depth (z) .

Maps for these corrections have been created previously, using ^{83m}Kr calibration data, and these maps are used to calculate a "flat-fielding" correction for each event, given its position and area. These corrections are calculated in the data processing framework and are available as arrays of individual corrections for each event. A comparison of the uncorrected (x,y) positions to the positions following correction for our dataset is shown in Figure 5.2. Here,

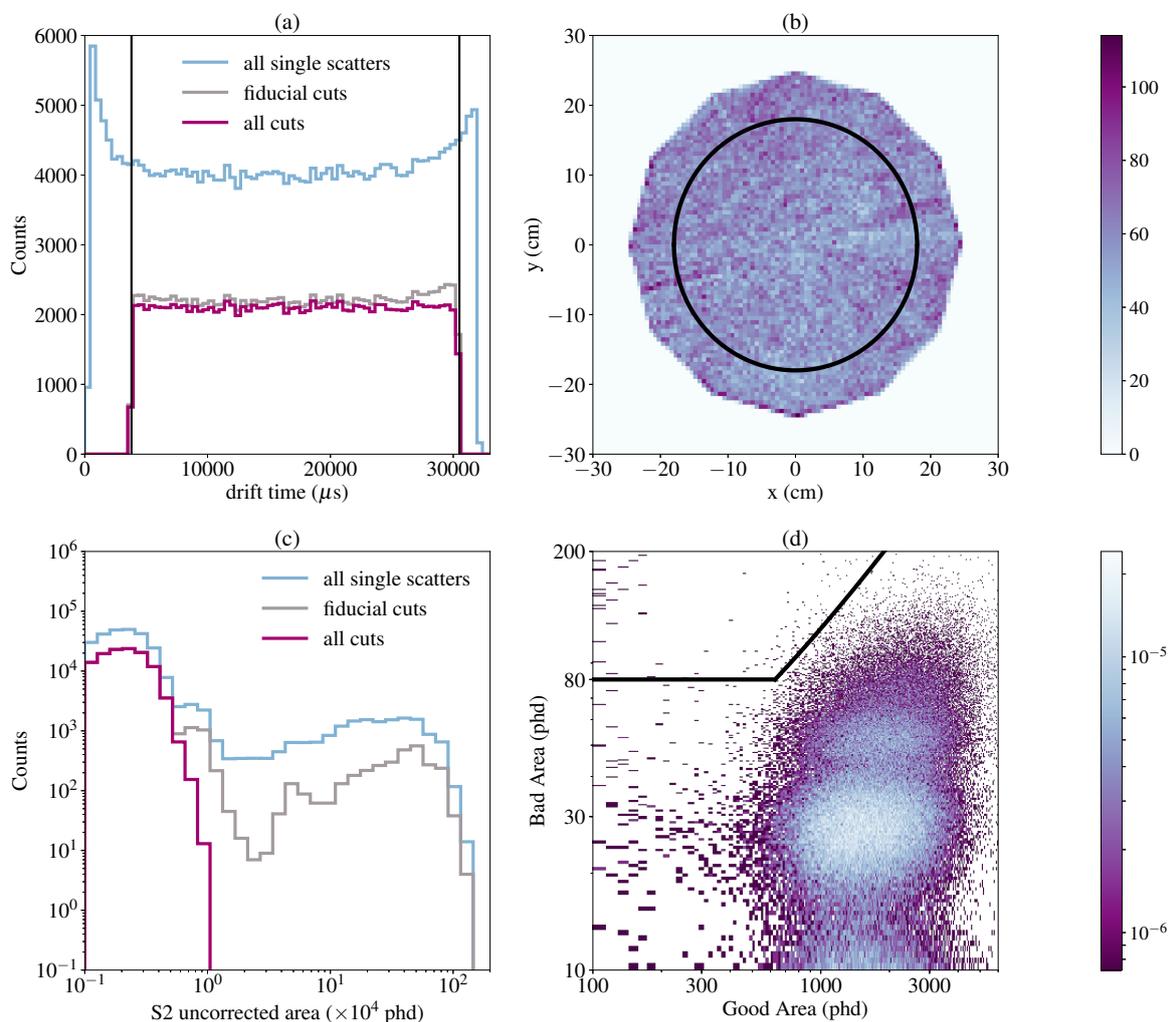


Figure 5.1: Standard search analysis cuts illustrated on the tritium calibration dataset. (a) Drift time of single scatter interactions for all events (blue), events passing fiducial cuts (grey) and events recovered after all cuts are applied (magenta). The drift time cut region is indicated by the two black lines. (b) Radial position of single scatter events with the fiducial region shown in black. (c) Uncorrected area of the S2 signal for single scatter interactions for all events (blue), events passing fiducial cuts (grey) and events recovered after all cuts are applied (magenta). (d) Good and bad area quantities for single scatter interactions with the cut region shown in black. Events plotted here represent the full single scatter event dataset for the tritium calibration, before the application of the quality cuts discussed in the text.

we use these corrections to calibrate the S1 and S2 signal areas so that they equal the pulse areas of the signals if they had occurred at the detector centre and no signal loss had occurred. For the S2 signal, the summed PMT pulse area response is used, while for the S1 signal an algorithm taking advantage of the spike counting information and correcting for photon pile-up (overlap) at low S1 values (below 80 phd) is used. After application of the quality cuts previously defined, the resulting S1 and S2 calibrated signals are shown in Figure 5.3,

along with the standard chain result. The result presented here is shown to be in reasonable agreement with the standard filter code result, with differences assumed to be due the lack of analysis preservation as previously mentioned.

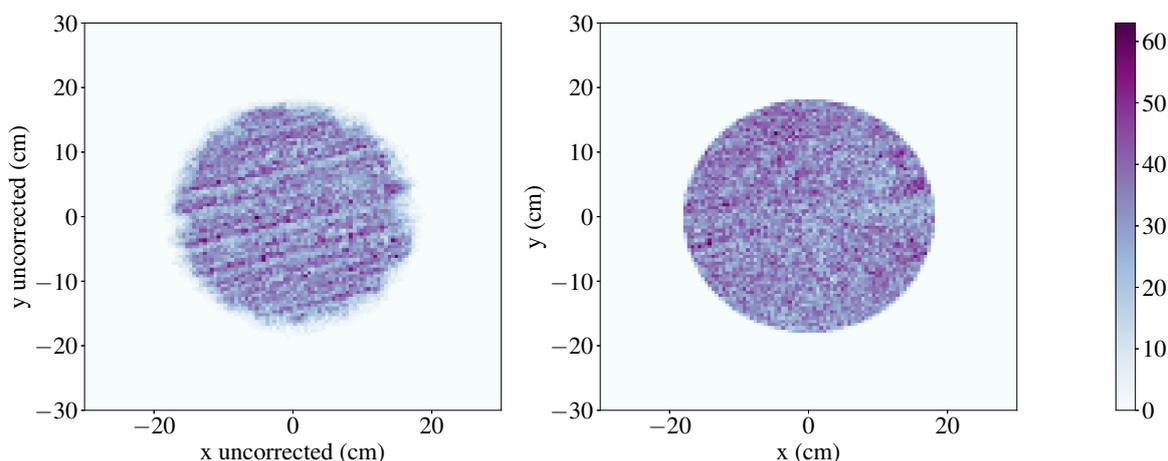


Figure 5.2: *Left:* Uncorrected (x,y) positions for CH_3T data with drift times between 38 to 305 μs and corrected radius smaller than 18 cm. All other cuts discussed in text were also applied to the data presented here. The striped pattern is observed due to electron focusing near the gate wires, due to the field difference in the drift and electron extraction regions. *Right:* The same data, after the correction to (x,y) positions has been applied.

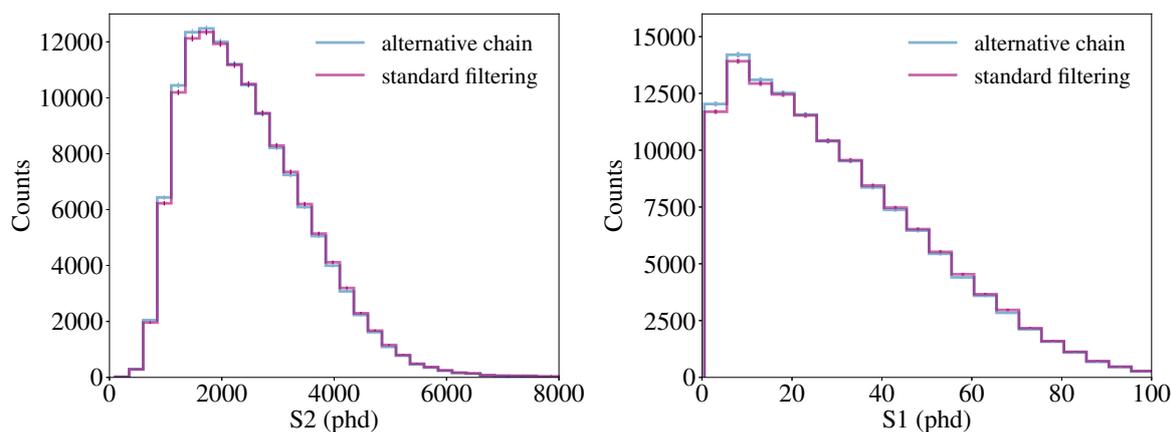


Figure 5.3: Tritium (*left*) S1 and (*right*) S2 spectra for a 2-fold coincidence requirement for the December 2013 data. The pulse area distributions are presented using the standard filtering and reconstruction chain (purple) and our alternative chain (blue) along with the count ratios for comparison.

Naturally, we also present the data in terms of the combined energy of interactions and compare with the well known tritium beta spectrum convolved with detector resolution (see Figure 5.4). The reconstructed energy for each event is directly proportional to the number of quanta produced and calculated using equation 3.21. For the dataset presented here, g_1 ,

g_2 , and the extraction efficiency used for energy reconstruction were 0.115 phd/photon, 12.1 phd/electron (liquid), and 50.9%, as reported for the original analysis of this dataset [119] and in agreement with other LUX analyses [135].

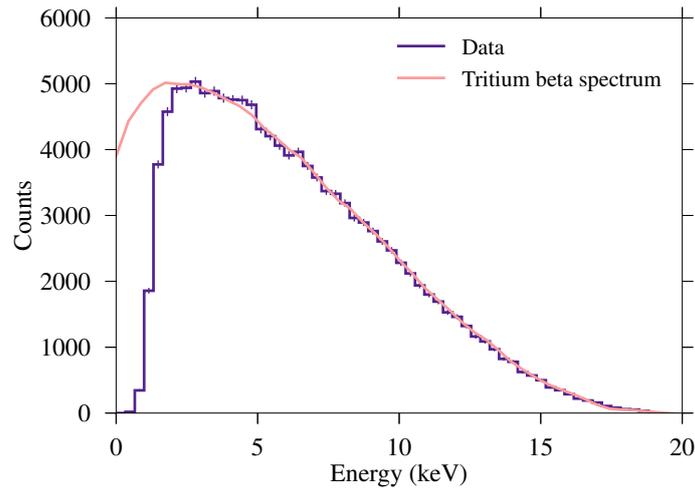


Figure 5.4: Tritium energy spectrum for a 2-fold coincidence requirement for the December 2013 data in purple with events presented here processed through the standard processing framework, while event filtering and reconstruction has taken place in our alternative chain. The combined energy model discussed in text was used for energy reconstruction. The tritium beta spectrum convolved with detector resolution (taken from [119]) is shown in pink.

5.2 Lowering the scintillation threshold for the tritium dataset

Following the successful reproduction of the tritium calibration result using the new processing chain, we set out to lower the scintillation threshold for this dataset. Specifically, we extend the standard "2-fold" analysis by including single photon pulses (1 phd) in those instances where one VUV photon had produced two photoelectrons (2 phe) through DPE. This allowed lower energies to be detected by recovering a number of events that were previously discarded by the 2-fold requirement. PMT dark counts consist almost entirely of 1 phe signals and hence do not pose a significant background. An example event that fell below the 2-fold requirement but in which the S2 was preceded by a single detected photon of large pulse area (indicative of DPE emission) is shown in Figure 5.5. Note the significant S2 pulse size for this small (~ 1 keV) ER signal; although NR signals generate proportionally smaller S2s, they still represent several emitted electrons that are easily detectable by electroluminescence.

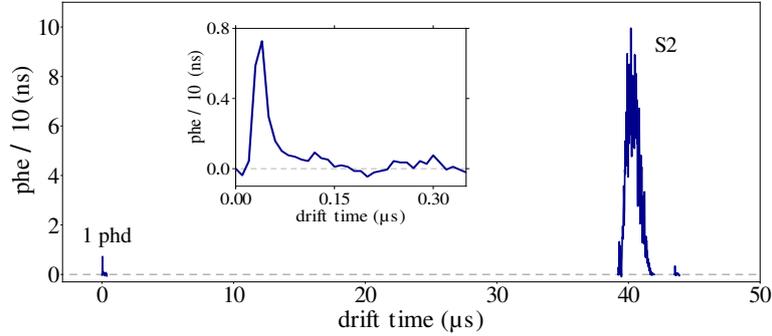


Figure 5.5: A single-photon S1 event from the Dec. 2013 LUX tritium dataset. The S2 pulse (~ 35 extracted electrons) was preceded by a pulse integrating to 2.4 phe — likely due to DPE emission in response to a real S1 detected photon (the relevant PMT waveform is shown in inset). In the standard S1+S2 analysis this would have been classified as an "S2-only" event, as the S1 candidate has failed the 2-fold requirement.

5.2.1 DPE cut acceptance

Given that the LUX PMTs did not fully resolve 1 phe from 2 phe signals, a cut must be applied on the candidate single-photon S1s that accepts DPE signals while rejecting dark count signals (we favoured the transparency of this simple cut-and-count method over more sophisticated procedures in this analysis). Figure 4.13 in the previous chapter suggests that the optimal cut would be located near 2.0 phe. A signal-to-background ratio optimisation designed to limit the total number of background events to $\mathcal{O}(1)$ (given the DC and S2-only rates in the full 2013 WIMP search dataset) yielded the following acceptance region: $[\mu_1 + 3.0\sigma_1, \mu_2 + 2.0\sigma_2]$, where μ_i and σ_i are the means and widths of the two Gaussians measured from the fits to the VUV dataset. Figure 4.13 (right) indicates this region for the PMT in question. The signal regions for all channels are shown in Figure 5.6.

The acceptance of the DPE cut for each PMT was calculated by integrating the modelled single photon response (as presented in Section 4.3.1) within its signal region. Then the relative contribution of each PMT was used to estimate the mean acceptance, using the individual channel quantum efficiency and the different collection efficiencies of the top and bottom arrays for scintillation light. We measured the top-bottom photon detection efficiency ratio using $^{83\text{m}}\text{Kr}$ calibration data to be 0.26. In particular, we divided the means of the S1 uncorrected spectra as seen by the top and bottom arrays (shown in Figure 5.7(a)). The mean QE for the good PMTs is $33 \pm 2\%$, with the error here representing the standard deviation of all channels on the sample. The mean weighted acceptance for the single detected photon population recovered after the DPE cut for all good PMTs was calculated for the two S1 cut allowances we defined in the previous chapter, with the variation between the two values presented as a systematic error.

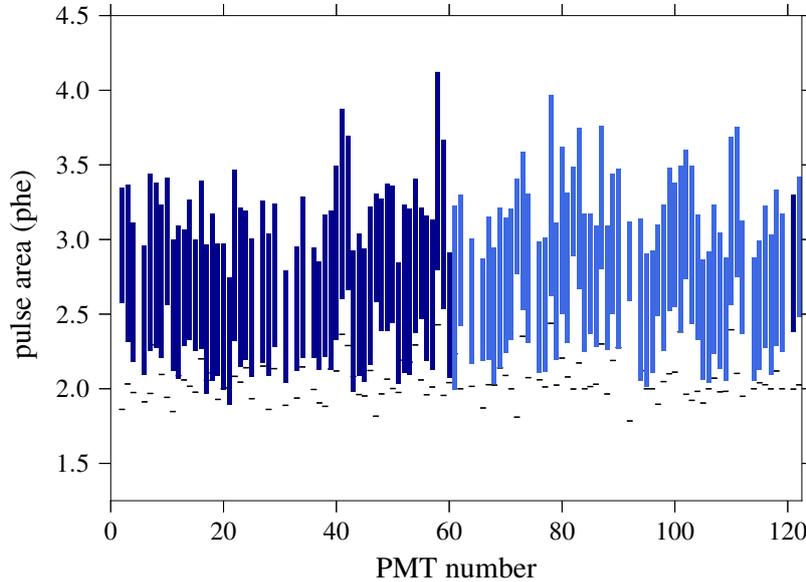


Figure 5.6: DPE signal region for each active channel. Blue bars illustrate the $[\mu_1 + 3.0\sigma_1, \mu_2 + 2.0\sigma_2]$ region for each PMT, while missing bars indicate where channels were not used for this analysis. The dark blue bars indicate channels located in the top PMT array, while the lighter blue bars show bottom array channels. Black lines show the mean of the DPE response (μ_2) for each PMT.

At this point it is essential to further discuss the simulation framework used to support this study. As introduced in Section 3.3.4, NEST is used to model the number of quanta produced during an interaction in LXe for ER and NR interactions. This simulation package also allows to set detector parameters such as the electric field, light collection and extraction efficiencies but is not a full simulation with photon propagation, and detector geometric effects are not modelled. Instead, we manipulate the experimental data to reach corrected quantities which we can then compare to the simulated quantities.

Our simulated events will hence be on corrected quantities for all PMTs. So far, we have presented the mean acceptance for the good PMTs. Hence, we need to estimate the effect between the corrected S1 for all PMTs (simulated) and an uncorrected S1 for the good PMTs (what data will contain). More specifically, the removal of PMTs from our analysis results in a reduction in g_1 and a commensurate reduction in single photon event acceptance between good and all PMTs. To estimate the effect of this reduction we used the ^{83m}Kr calibration dataset and investigated the effect of "turning PMTs off" on the overall S1 light collection efficiency. We performed this study by comparing the mean S1 corrected size for all PMTs and the mean uncorrected S1 size for the good PMTs, as shown in Figure 5.7. Note that NEST does allow for total reduction in the number of PMTs but does not allow for a choice on PMT location and as such studying the effect of "turning PMTs off" is more accurate using data, where the location of PMTs is accounted for. Using the comparison outlined, we estimate a

6.0% reduction in g_1 relative to the value indicated previously. Hence, with this applied, the final DPE cut acceptance was estimated to be 0.052 ± 0.005 . This acceptance translates directly to the analysis efficiency of the DPE cut, and hence represents the fraction of events with an S1 consisting of a single detected VUV photon that are recovered after its application.

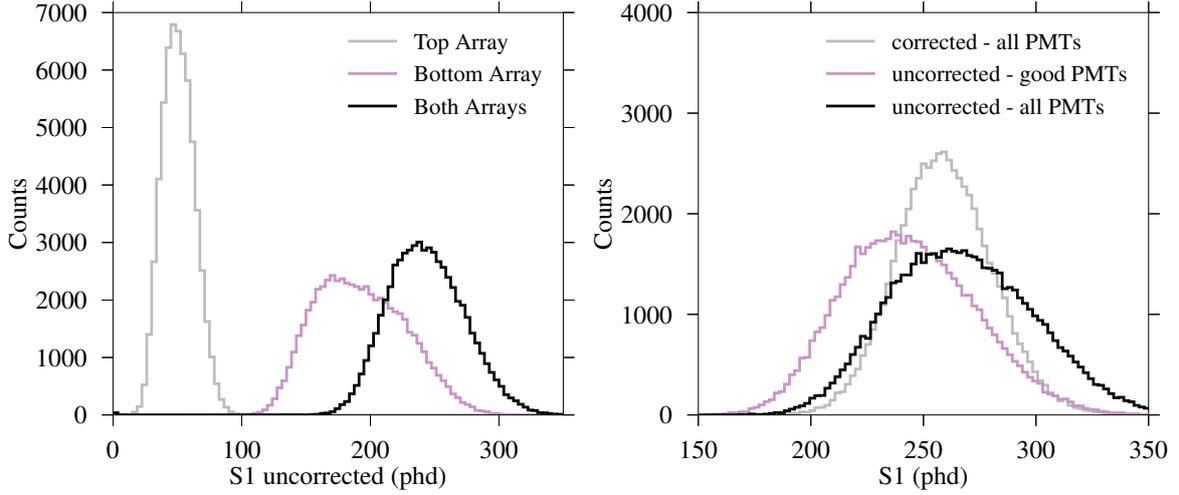


Figure 5.7: Light collection efficiency study using $^{83\text{m}}\text{Kr}$ data. *Left:* S1 uncorrected spectra for $^{83\text{m}}\text{Kr}$ data for the top, bottom and both arrays illustrating the top-bottom asymmetry ratio. *Right:* S1 spectra for corrected and uncorrected quantities for all and only good PMTs, illustrating the light loss due to turning PMTs off.

We now discuss the efficiency of identifying single photons, that is, how often does the pulse classifier miss single photon pulses. By hand-scanning ~ 200 single photon pulses in the tritium S2-only dataset — a dataset of events that have one S2 and no S1 pulses — we found that 86% of the single photon pulses occurring before an S2 were successfully found. About 5-6% of the 14% inefficiency is due to the POD threshold while the remainder was almost always seen in events with more than 10 pulses (reaching the LUX pulse limit for pulse finding). In such events, the pulses that remain unidentified are mostly those with smaller area. To quantify the efficiency of the pulse classifier for larger single photon pulses, we looked at the percentage of such pulses with area above 1.0 phd that were unidentified, which was found to be $(1.5 \pm 1.5)\%$. The efficiency is seen to increase with increasing pulse size and, in the DPE signal region that interests us (>1.5 phd), the efficiency loss due to the pulse finder should be negligible. Additionally, events in which single photon pulses of large size remained untagged were events with many more than 10 pulses that are possibly pile-up events following some very large energy deposition (which can exhibit significant electron emission for several ms) and hence it is unlikely that the very small rate of pulses lost are signal single photon pulses. Hence, we estimate the efficiency of the pulse finder in identifying single photons, that appear before the S2 in an event, with pulse areas that fall within our DPE signal region to be ~ 1 .

5.2.2 Fake single-photon rate from over-fluctuating dark counts

In standard analyses, the 2-fold coincidence requirement eliminated a major source of background whereby PMT dark counts recorded up to a few hundred microseconds before an isolated S2 pulse faked a valid single scatter event possessing one S1 and one S2. The dark count response consists of mostly 1 phe signals as discussed in Chapter 4. Nevertheless, a small rate of the area distribution of 1 phe DC pulses leaks into our defined DPE signal region and can coincide with a valid drift time window preceding S2-only events, producing an accidental coincidence background.

The rate of pulses classified as single photons that fall within the defined DPE signal region, over the entire analysed array, was found to be (2.0 ± 0.2) cts/s, with the error representing the statistical uncertainty on the counting measurement. This observation, obtained using the data distributions shown in Figure 4.5, is in agreement with the expected rate due to dark count leakage into the signal region, that was calculated to be (2.2 ± 0.2) cts/s using the model fits (Figure 4.7), with the error propagated from the fitted distributions. Using the counting observation, which could include a small contribution from random VUV single photon sources, ensures that we consider all observed leakage rate in background expectation calculations.

5.2.3 Single VUV photon events in the tritium calibration dataset

Following the definition of the single photon signal region and the fake single-photon rate, we proceeded to perform a single-photon S1 search using the 2013 CH₃T dataset. Significantly, the S1 event selection at 1 phd is identical for ER and NR datasets, and so it is directly applicable to the WIMP search analysis presented in the next chapter.

To perform this search, events with a single valid S2 pulse and no identified S1 pulse were initially selected and the standard S2 pulse quality and quiet time cuts were applied (defined in Section 5.1). Selected events were searched for a candidate single photon pulse preceding the S2 and falling within the DPE acceptance region of the firing PMT — such as the event already shown in Figure 5.4. For those fulfilling these requirements the single photon pulse was taken as the S1 signal, and all additional standard event selection cuts (fiducial, cleanliness and pulse area) were applied. To benefit from the extension to lower energies, the (uncorrected) S2 threshold was lowered from 165 phd to 100 phd (~ 4 emitted electrons). The effect of the standard tritium analysis cuts on our single photon search dataset is shown in Figure 5.8. At this very low S2 pulse threshold the (x,y) position resolution remains small ($\sigma \simeq 0.8$ cm [149]), and the trigger efficiency is already $\approx 100\%$ [148].

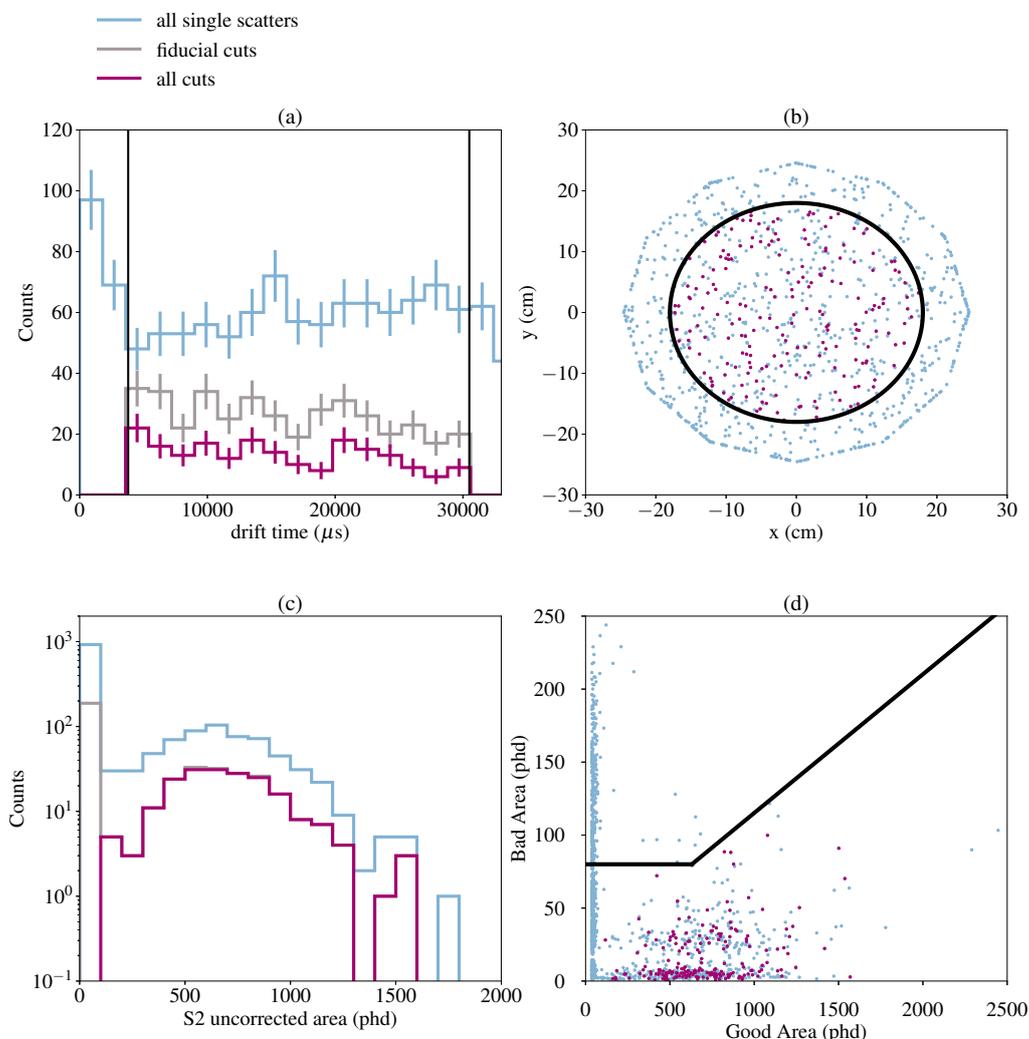


Figure 5.8: Selection cuts for the tritium single photon dataset. (a) Drift time of single-photon S1 events for all events (blue), events passing fiducial cuts (grey) and events recovered after all cuts are applied (magenta). The drift time cut region is indicated by the two black lines. (b) Radial position of single-photon S1 events for all events (blue), events passing all cuts (magenta) and the radial cut region shown in black. (c) Uncorrected area of the S2 signal of single-photon S1 events for all events (blue), events passing fiducial cuts (grey) and events recovered after all cuts are applied (magenta). (d) Good and bad area quantities of single-photon S1 events for all events (blue), events passing all cuts (magenta) and the cut region shown in black.

Two additional cuts were required to remove background sources that did not appear in the standard 2-fold analysis. These cuts aimed at reducing the background rate by tackling known pathologies that result in measurable rates of accidental coincidences between S2 pulses and

over-fluctuating dark count pulses. The first cut removed interactions occurring very near the liquid surface, such that the S1 was essentially merged with the S2 pulse (such events do not appear in the standard analysis because a valid S1 preceding the S2 is required in that case). To successfully remove such events, we enforced a two-dimensional cut in the prompt fraction (fraction of pulse area recorded at the beginning of pulse) and top-bottom array asymmetry ratio. The second cut was enforced following a hand-scanning of events, which identified a new background source. By visually inspecting events a contamination source was observed in which the identified single-photon S1 appeared at fixed time window following either a SE or following a pulse that was classified as "other-unidentified". The latter category of pulses contained two large area pulses detected by the same PMT, such as that depicted in Figure 5.9. The fact that in each case the single photon candidate pulse was within a few μs from the unidentified or SE pulse and appeared in the same PMTs suggests afterpulsing or other PMT artefact, and hence a cut at $5 \mu\text{s}$ was used to eliminate these contamination events. We also evaluated the effect of removing these events from the accidental coincidence rate presented in Section 5.2.2 and found that the rate was identical when the class of events disuccsed here was excluded, validating that we did not bias our analysis by removing events from our final dataset that were somehow included in our background estimate. Nevertheless, we recognise that removing events using hand-scanning was only acceptable in this analysis due to its benchmarking nature, which allowed the identification of new PMT related background sources. Now that these have been properly identified, future analyses could impose cuts to eliminate them in a more rigorous way, without the need to review individual events.

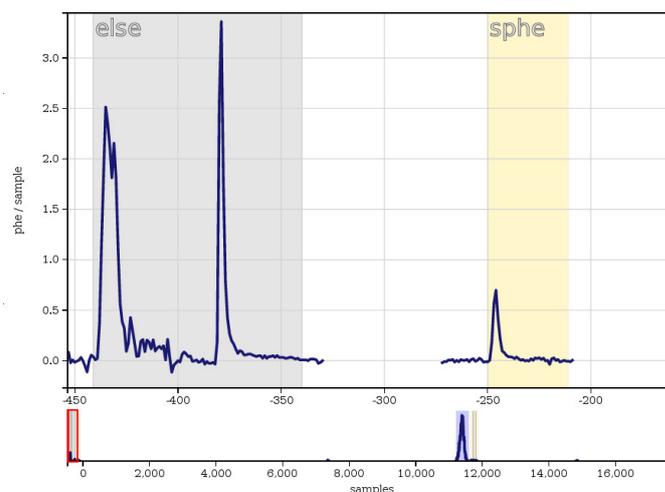


Figure 5.9: Visualux display of a contamination event in which a large-area single photon (here labelled "sphe") follows an odd pulse topology ("else") detected in the same PMT, and likely due to some form of afterpulsing. The time axis in this figure is presented in samples, with 1 sample equivalent to 10 ns.

We also investigated if in our final set of events there was indication of a PMT high-rate period, and in particular whether there were additional single photon like pulses from the same PMT in an event, possibly associated with electrical micro-discharging. We enforced a cut eliminating events in which the number of single photon pulses detected by the same PMT, appearing before the S2, was greater than two. Given the measured dark count rates, the number of events in which more than two dark counts from the same PMT appear in the same event before the S2 should be negligible. In our single-photon S1 dataset no such events were observed.

Following the application of analysis cuts, we proceeded to flat-field the S2 response across the detector. In particular, we correct for (x,y) variations and z (electron lifetime) effects. Unlike the 2-fold analysis, correction arrays for the single photon events recovered did not exist. This is because our set of events had no identified S1 above the 2-fold requirement and hence, when they originally passed through the data processing framework, they were tagged as S2-only events. The geometric corrections (x,y,z) are a function of drift time, and since the events are classified as S2-only, with no drift time information present, corrections were not calculated. For the (x,y) corrections, we adapted the algorithm used within the data processing framework (note that the algorithm uses the newly recovered drift time to calculate the corrected radial positions) so that it could be directly used for our set of events. We also benchmarked this correction application by studying its effect on single scatter events with known corrections.

For the electron lifetime correction we measured the electron lifetime using single-scatter events from the standard analysis. To achieve this, we created a database of the corrections for standard tritium events to create a fit on the correction for each drift time, as shown in Figure 5.10. This was then directly applied to our single-photon S1 events, given their drift time.

In total $\sim 15,000$ S2-only events passed the S2 quality and cleanliness cuts alone. Given the simulated tritium spectrum we estimate that these are largely ($\approx 75\%$) due to genuine tritium decays with an S1 failing the 2-fold coincidence requirement (either 0- or 1-fold S1), followed by background sources of S2-only events. Of these 15,000 events, 247 were found with an S1 consistent with DPE emission on 1 phd within the appropriate drift time window.

As mentioned previously, S2-only events may coincide with large-area DC pulses to produce a viable background topology. In this case, where the biggest proportion of the S2-only distribution originates from genuine interactions with truly no S1 detectable signals, the biggest proportion of the random coincidence rate involves genuine tritium S2 pulses, coinciding with large-area dark counts. Using the rate of (2.0 ± 0.2) cts/s for the latter and the fact that each of the $\sim 15,000$ events that pass the quality cuts may coincide with these dark counts, an expectation of 10.8 ± 1.1 random coincidences was estimated for this tritium dataset.

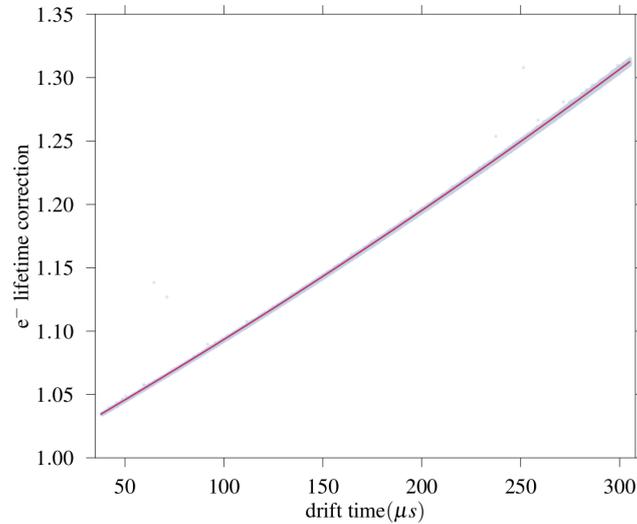


Figure 5.10: Electron lifetime correction at each drift time from a set of single scatter tritium events fitted with a third degree polynomial function.

The S2 pulse area spectrum for the single scintillation photon events passing all cuts is shown in Figure 5.11. We also present the expected distribution for the accidental background population. For this distribution, the S2-only spectrum was scaled to reflect the expected coincidence rate and S2 pulse areas were corrected by assigning a random (x,y) and z position to events. The predicted spectrum for pulse overlap events in which S2-only events present as both a background and signal count is subtracted from the background expectation.

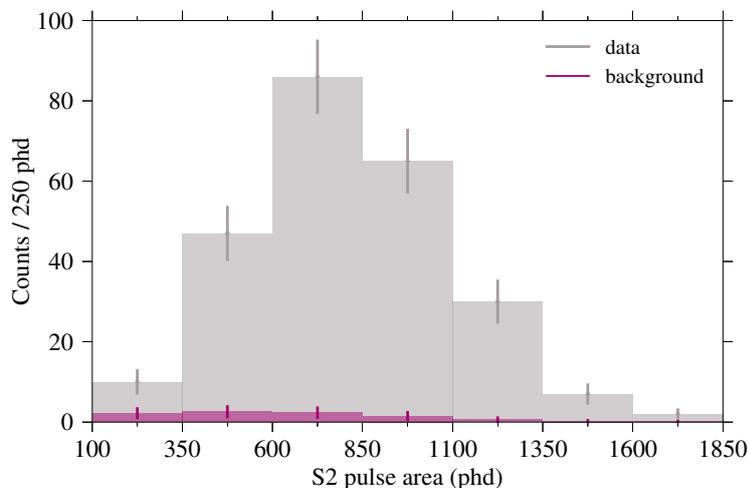


Figure 5.11: S2 pulse area spectrum for tritium dataset events in which the S1 consisted of a single detected VUV photon with pulse area within the DPE signal region, along with the background expectation from DC+S2-only coincidence events.

Finally, the scatter plot for 2-fold events recovered using the standard analysis along with the single-photon S1 detected events recovered here is shown in Figure 5.12. The 50th (solid), 10th and 90th (dashed) percentiles of S2 at given S1 are shown using black lines.

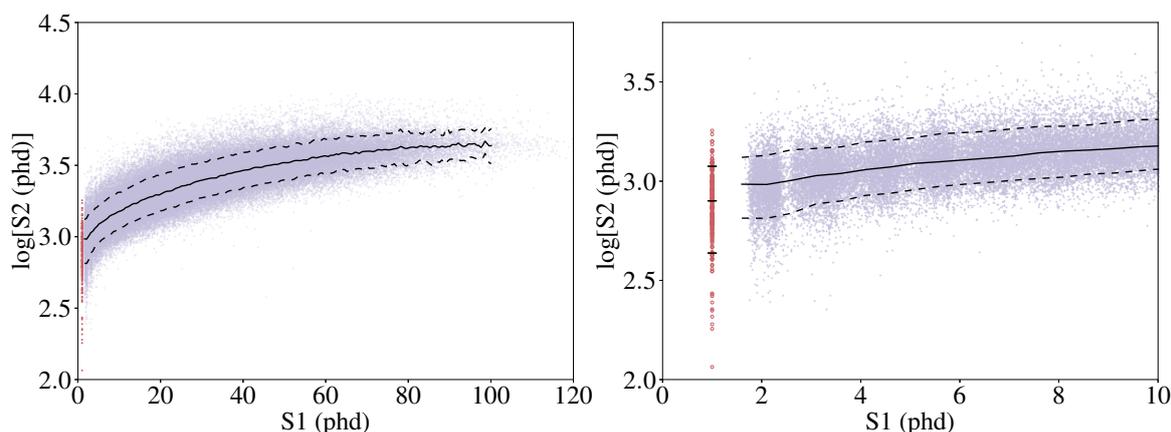


Figure 5.12: *Left:* Scatter plot for tritium selected for this analysis. The single-photon S1 events recovered in this analysis are plotted in red at exactly 1 phd, while for events with higher S1 (lilac), the S1 variable is not an integer because of the application of corrections. The 50th (solid), 10th, and 90th (dashed) percentiles of S2 at given S1 are shown using black lines. *Right:* Same scatter plot presented only for S1 signals smaller than 10 phd.

5.2.4 Electron recoil scintillation and ionisation yields at low energies

In Section 3.3 we discussed signal production in liquid xenon and in particular the light and charge yields at different energies. Here, we focus on measurements and models for low-energy electron recoil interactions. Key published measurements [113, 119, 122, 173] of these yields below 5 keV are summarised in Figure 5.13. We also present the NEST v2.0.0 yield models, obtained using the datasets shown in the figure. The ER model threshold was placed at 0.186 keV, corresponding to the lowest energy ionisation yield measurement [122], and extends below the lowest energy at which the scintillation yield has been measured to date (1.3 keV [119]). When extrapolating the model to this energy threshold, the scintillation signal was assumed to be fully anti-correlated with the charge signal, as was observed at higher energies. The uncertainties presented as bands around the ER yields in Figure 5.13 indicate the variation observed when the free parameters of the model were allowed to vary by 1σ .

The NEST v2.0.0 simulation package incorporating both these yields and the LUX detector geometry and efficiencies was used to simulate the S2 spectra for the tritium interactions under study. The S2 pulse area spectrum for the single scintillation photon events along with the S2 response predicted by NEST v2.0.0 are shown in Figure 5.14. The modelled response includes the (small) background distribution from upward-fluctuating DCs associating with S2-only events. Systematic and yield uncertainties are also shown. The former include errors on g_1 and

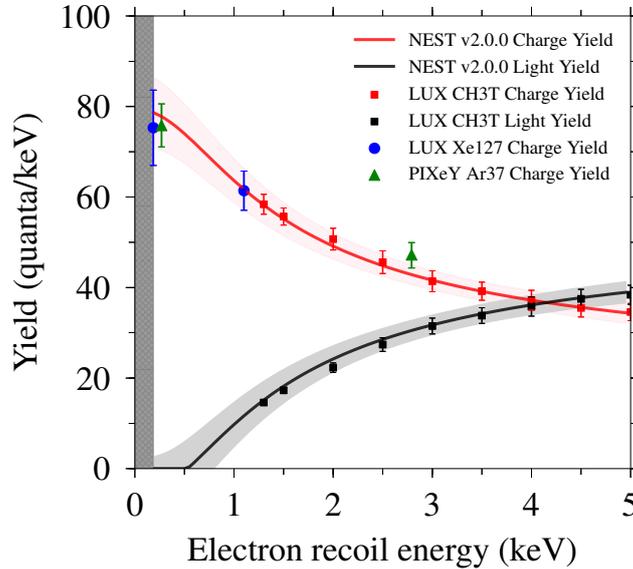


Figure 5.13: The ER ionisation (red) and scintillation (grey) yields as a function of recoil energy given. Solid lines present the yield models given by NEST v2.0.0 with the bands around these indicating model uncertainties corresponding to those presented in Figure 5.14. Red and black points indicate LUX measurements from the tritium calibration [119]. Additional measurements are shown in blue and green for the ER ionisation yield [122, 173].

the DPE cut acceptance added in quadrature, while the yield uncertainty includes ionisation and scintillation yield variations within measurement errors at low energies — see Figure 5.13. The data are in good agreement with the NEST v2.0.0 ER model, and the total number of events predicted (inc. background) is $208 \pm 21_{\text{sys}} \pm 11_{\text{yield}}^{25}$ (cf. 247 observed). The ER model predicts the energy of the single photon tritium events to lie between 0.3 and 2.6 keV, with a mean of 1.1 keV.

The good agreement between data and model for the single photon ER calibration shown in Figure 5.14 suggests that the model extrapolations assumed in NEST, below the lowest light yield measurement, are reasonable for the energy range presented here. This suggests that these yields are well modelled by NEST down to sub-keV energies. In addition, we have shown that a population of single photon events undergoing double-photoelectron emission can be reliably selected to lower the detection threshold for ER interactions. The modest backgrounds mean that this technique can be used in searches for leptophilic DM and other rare interactions producing electronic recoils near threshold in liquid xenon detectors.

5.2.5 Removing the DPE cut

It is worth considering whether a DPE cut is required in the first place, i.e., would a full analysis including all pulses compatible with 1 phd, including dark counts, despite the much higher

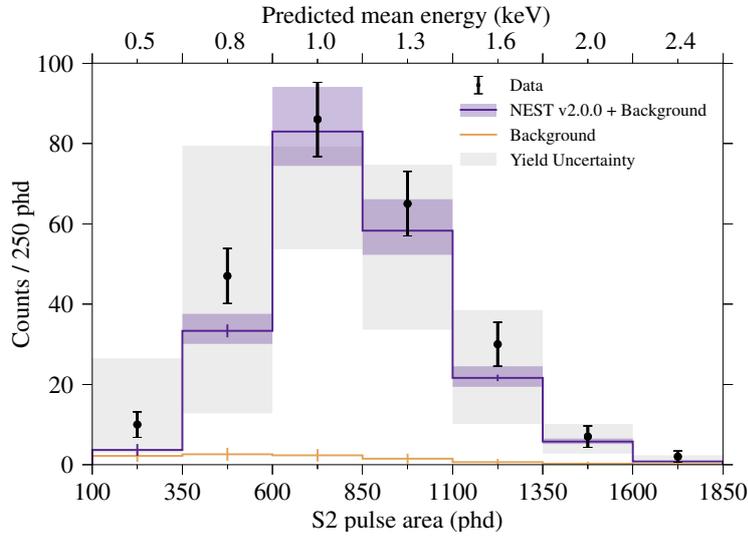


Figure 5.14: S2 pulse area spectrum for tritium dataset events in which the S1 consisted of a single detected VUV photon with pulse area within the DPE signal region, along with NEST v2.0.0 prediction added to the background expectation from DC+S2-only coincidence events. The error bars on the NEST v2.0.0+Background line represent the uncertainty on the background expectation. The shaded purple regions represent the systematic uncertainty due to the g_1 and DPE cut acceptance error, while the shaded grey regions represent the yield uncertainties (produced by incorporating appropriate yield variations in the NEST model, shown in Figure 5.13). The mean energy as predicted by NEST v2.0.0 for the range of S2s in each bin is indicated on the upper x-axis.

background? For the tritium dataset most such events are indeed tritium interactions with sub-threshold S1 signals, and we have confirmed that the background-subtracted S2 spectrum is still in good agreement with the NEST prediction.

To investigate further our understanding of the background calculation, we performed same analysis but removing the DPE cut and allowing all single photon pulses preceding S2s within an appropriate drift time window to be considered as valid. The S2 area spectrum for these events is shown in Figure 5.15.

To compare our results with and without the DPE cut, and hence validate our understanding of the DPE cut acceptance, we compare the data-to-model ratio of counts for both cases. These ratios are presented in Figure 5.16. The data points are used for both ratios, as presented in Figures 5.14 and 5.15, while for the model expectation the mean model response is used (shown as purple lines in the relevant figures). We propagate both the statistical data uncertainty and the model (g_1 and DPE cut) uncertainties to the ratio calculations. The yield uncertainty is not included or presented in Figure 5.16. We observe that the ratios agree for the two studies, within errors, supporting our understanding of the DPE cut acceptance and background rate calculations, which are the only parameters changing between the two studies. In this case, subtracting the NEST-predicted signal spectrum from the S2-only population used

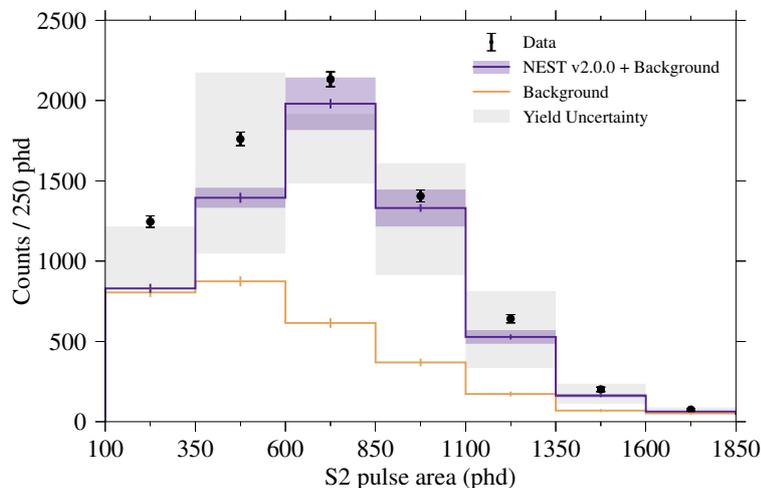


Figure 5.15: S2 pulse area spectrum for tritium dataset events in which the S1 consisted of a single detected VUV photon, along with NEST v2.0.0 prediction added to the background expectation from the DC+S2-only coincidence events. The error bars on the NEST v2.0.0+Background line represent the uncertainty on the background expectation. The shaded purple regions represent the systematic uncertainty due to the g_1 and DPE cut acceptance error, while the shaded grey regions represent the yield uncertainties (produced by incorporating appropriate yield variations in the NEST model, shown in Figure 5.13).

to predict the background spectrum is vital to ensure that events are not double-counted in the NEST+background distribution.

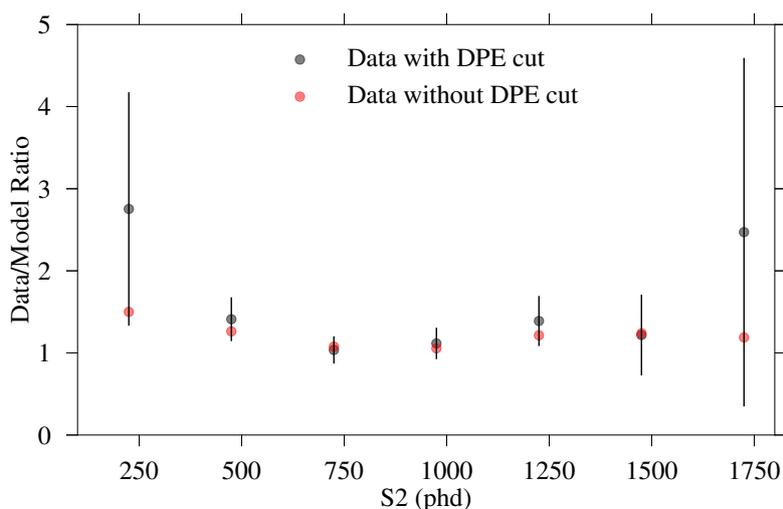


Figure 5.16: Data-to-model counts ratio for dataset including a DPE cut (as presented in Figure 5.14) and without a DPE cut (as presented in Figure 5.15). For the ratio presented here the mean model response is used. Error bars include the statistical uncertainty on the data and the systematic uncertainty due to the g_1 and DPE cut acceptance error in the NEST model.

5.3 Summary

The analysis of the tritium data confirmed that single scintillation photon events can be robustly selected to recover a population of very low energy interactions which had been previously discarded, and it suggested that the dominant random coincidence backgrounds can be calculated accurately. The analysis extended successfully the standard CH₃T calibration threshold and no significant setbacks were discovered. We demonstrated the accurate reconstruction of a population of events where the S1 pulse consisted of a single detected photon recorded in the tritium calibration of the LUX experiment.

This analysis procedure could be used to lower the energy threshold of ER calibration studies for LUX and other LXe detectors. Additionally, several interesting dark matter candidates present themselves as electron recoil signals as discussed in Section 3.7.2. Particularly, recent studies have suggested that DM-nucleus scattering can be accompanied by signal that results in ERs at higher energies than the partner NR signal, extending the reach of LXe detectors to sub-GeV masses. The Migdal and nuclear Bremsstrahlung channels and their sensitivity reach can be extended by lowering the scintillation yield using procedures such as the one described in this chapter.

Chapter 6

Extending light WIMP searches to single scintillation photons in LUX

The first LUX WIMP search dataset was acquired during Run 3 in the period March–October 2013 (WS2013), accumulating 95 live days of search exposure. The results were first published in 2014 [136] with the data being consistent with the background-only hypothesis. Further calibrations and analysis improvements allowed an increase in sensitivity and a re-analysis of the same dataset was published in 2016 [127], producing even more stringent upper limits on the cross-section of spin-independent WIMP-nucleon scattering. Here, given the success of applying the single photon analysis to the LUX (ER) calibration, we applied the same technique to search for light WIMP interactions (NR) in the same dataset. This work was published as Ref. [1]. We note that the second science run of LUX in 2014–16 (WS2014-2016) produced a significantly larger exposure, but this was affected by the electric field non-uniformity issues highlighted in Section 3.5.1; it required a major analysis effort to correct those data. For this reason, we opted to use WS2013 for a clearer demonstration of this analysis technique.

In this chapter we begin by considering the LUX backgrounds in the low-energy regime, and their extension to our single-photon S1 dataset. We then present results for the WIMP search, initially reproducing the original WS2013 published result and then searching for single-photon S1 events, which would not have been considered in the original analysis. Subsequently, we arrive at a new competitive upper limit for light WIMPs, using a simple cut-and-count approach and NEST v2.0.0 to inform our signal estimation. The model assumptions on our result were then studied, with alternative models used to evaluate the systematic errors arising from the scintillation and ionisation yields we have adopted.

6.1 Reproducing the original Run 3 result

Data acquired over 132 calendar days in 2013 were analysed to produce the LUX WS2013 search result [127]. Approximately 22% of this period was used for ER and NR calibrations,

purification after tritiated methane injections, or affected by xenon circulation outages. A further 2.5%, 1.8% and 1.6% were excluded due to detector instability, trigger hold-off and other minor contributions, respectively. Hence, the final live-time for the actual dark matter search corresponded to 95 days (an effective duty cycle of 72%).

Here we set out to reproduce the result published in Ref. [127] using both the standard filtering and analysis chain and our alternative chain. Event filtering (recovering single scatters) and fiducial cuts were applied as discussed in Section 5.1, with the radial cut placed at 18 cm and the drift time cut extending from 38 to 305 μs . The S2 lower threshold was applied to the uncorrected raw pulse areas and placed at 165 phd, to mimic the published result. Later on, when analysing single-photon S1 interactions, this cut was relaxed. Finally, S1 pulses smaller than 50 phd and above the 2-fold coincidence requirement were selected, with the S2 pulses restricted to be lower than 10,000 phd. The cleanliness cut was also enforced, in the same way as that illustrated in Figure 5.1.

A scatter plot of the events remaining following analysis cuts, both recovered using our alternative analysis procedure and using the standard filter code, is shown in Figure 6.1. We present results for S1 signals below 20 phd as this reflects lower energy interactions of interest to our subsequent single-photon S1 search. We observed that five events appear in the alternative chain that are not present in the standard chain. A visual inspection of those showed that the S1 pulse had been in all cases misclassified; an example is shown in Figure 6.2. An additional cut applied in the standard filter chain for the WIMP search (not applied in the tritium analysis) was added in the alternative chain for the WIMP search and hence these events were removed from our alternative chain result too.

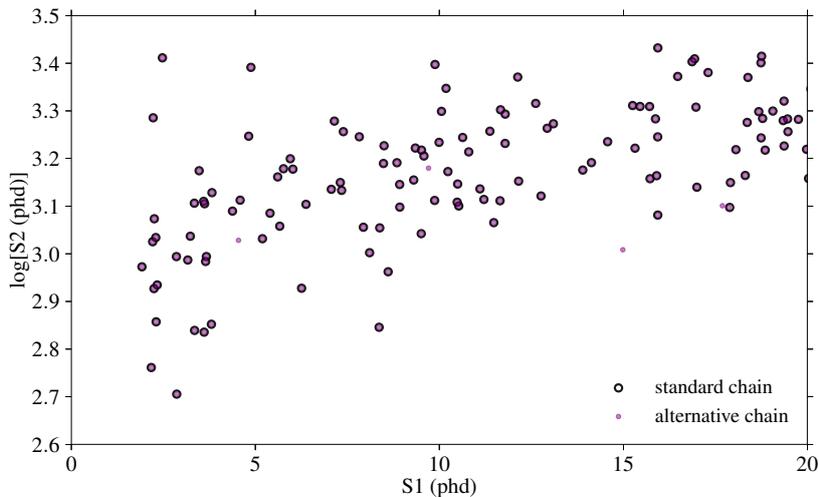


Figure 6.1: Events observed in the 2013 WIMP search exposure of 95 live days and 118 kg fiducial mass, following analysis cuts and pulse area corrections. Events recovered using the standard and alternative analysis chains are shown in black and purple, respectively.

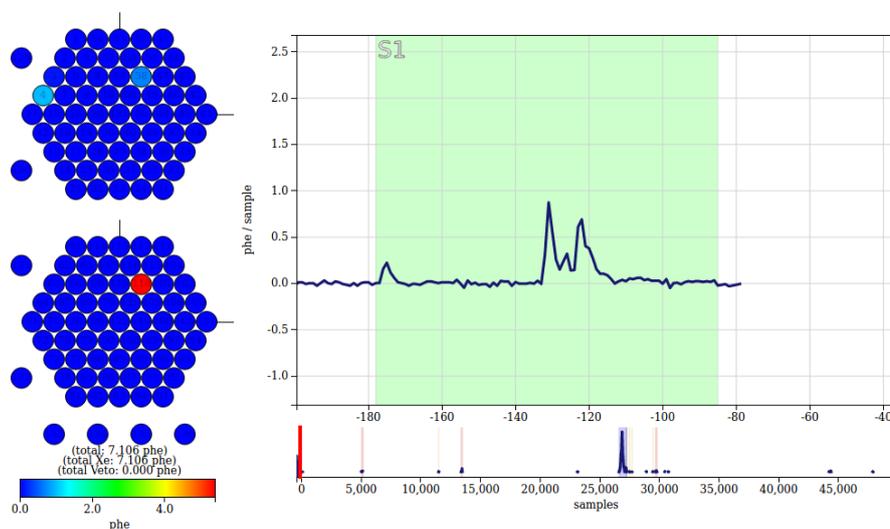


Figure 6.2: Example event in which the S1 pulse is due to mis-classification. This is possibly due to a PMT artifact as the distance between the two signals is much larger than the Xe scintillation time constants.

Given that we have reached agreement between the standard and alternative chain results, we proceeded to compare against the published result. Similarly to the tritium calibration result, the published analysis algorithms were not fully preserved and hence we expected minor differences, which we set out to study and understand. In the published result, we observe five events that do not appear in the standard chain result — these include three events which are identical to the ones also observed and removed from the alternative chain because of mis-identified S1s. Presumably, there were changes and improvements implemented in the standard chain filtering procedure between the time of publication of the original result and the current version in use. Note that this is not a difference in the data, as we select the same files as the ones used in the original analysis. The other two events do not appear either in the standard or the alternative chain and hence cannot be inspected, as we were unable to identify their specifying time stamp. We believe this is possibly due to some fine tuning of cuts that we are unable to retrieve, as with the application of identical cuts between standard and alternative chains these events are not recovered. Altogether, we have found no significant differences between the published, standard analysis and alternative analysis results. Any observed differences involved mostly improvements such as removing events that do not correspond to genuine interactions.

6.2 Background model

A low background rate and good understanding of its different contributions and how they translate into detected signals in the detector is imperative for any rare event search. Back-

grounds are expected to produce electron recoil, nuclear recoil or accidental coincidence types of events. At low energies, the background rate achieved is very low [135, 138]. In this section, we will discuss the dominant backgrounds for our analysis in turn before summarising them in Table 6.1.

6.2.1 Nuclear and electron recoil backgrounds

As discussed in Section 3.4.2, LUX achieved a good discrimination ability between electron and nuclear recoil signals with a $\sim 0.2\%$ leakage of ERs past the NR band mean for an S1 range 2–100 phd [135]. Nevertheless, this small leakage fraction accounted for the most prominent background in standard LUX analyses. Electron recoils within the LUX fiducial volume are caused by beta minus and electron capture decays within the Xe target and by gamma-ray emission from decays in the surrounding materials. The main component of ER backgrounds was due to Compton scattering of gamma-rays generated from the decay of radioisotope impurities (uranium-chain, thorium-chain, ^{40}K and ^{60}Co) in detector materials surrounding the LXe. Due to the density of the target and the energy of these gamma-rays (up to several MeV) their mean free path is of the order of only a few cm and hence they are suppressed within the fiducial volume. Gamma-rays generated in the cavern rock are attenuated to negligible contributions due to the 300-tonne water shield and 20-tonne steel shield existing underneath the water tank.

A spatially uniform population of β^- -decays makes up the second most important contribution to the ER background. These arise from two isotopes present in the LXe target: anthropogenic ^{85}Kr mixed with the Kr contamination, and the radon daughter ^{214}Pb . Anthropogenic ^{85}Kr is a result of nuclear fuel reprocessing and is a significant contamination source for LUX since Kr and Xe are removed together from the atmosphere by fractional distillation of liquid air, and its boiling point is not sufficiently distinct from that of xenon. ^{85}Kr also has a long half-life (10.8 years) and a maximum decay energy of 687 keV. Despite the reduction of natural Kr in LUX through a dedicated effort (to 3 ppt [137]), trace amounts of ^{85}Kr remain in the xenon. ^{214}Pb is the result of ^{222}Rn decay emanating out of materials and into the active detector volume. These can pose a problem either presenting as a "weak" naked beta-decay from ^{214}Pb in the TPC bulk or when accompanying gamma-rays ($Q_{max} = 352$ keV) escape the detector for peripheral interactions.

Finally, electron-capture decays by cosmogenic isotopes dispersed in the LXe contribute to the ER background population. Even though cosmogenic production of such isotopes is negligible in underground laboratories, Xe can become contaminated in the production and storage phase which takes place at the surface where the cosmic-ray flux is significantly higher. In LUX xenon was stored 200 m above sea level at least for 8 months, then shipped to SURF in two batches. Half was stored at the 1.6 km altitude of the SURF surface laboratory for 1 month

and the other half for 7 days. Most of the isotopes predicted to contribute to cosmogenic activation (^{127}Xe , ^{129m}Xe , ^{131m}Xe , ^{133}Xe) have short half-lives and hence result in negligible background contributions in the WS2013 science run, with the exception of ^{127}Xe . This isotope has a half-life of 36.4 days and decays via electron capture, in which its nucleus absorbs one of the atomic electrons. Its daughter nucleus, ^{127}I , further decays via gamma-ray emission (>90% of the time) from its initial state, which is either at 203 keV or 375 keV, to the ground state. The vacancy created during the electron capture is then filled via a cascade of X-rays with total energies of 32.2 keV (33%), 5.2 keV (12%), and 1.1 keV (3%). Even though the gamma-ray tags most events as a background, for more peripheral interactions it may escape from the detector leaving only the low-energy X-ray interactions. Another such cosmogenic source, ^{37}Ar , with a significantly lower rate, has been thought to also contribute a very small number of events to the ER background in LUX.

Most of the ER backgrounds listed, apart from ^{127}Xe , produce spectra which are flat in S1. ^{127}Xe produces energy deposits at three discrete energies, with the lowest one being of interest for this analysis. In any case, ^{127}Xe accounts for a small contribution. Hence, assuming an overall ER background spectrum flat in S1 for small energy signals (approximately consistent with what is observed), we expect a rate of 0.6×10^{-3} counts/phd/kg/day. This translates to 7 ER counts at each S1 photon for the WS2013 dataset, consistent with counting in the Figure 6.1. Given our DPE cut acceptance, this translates to an expectation of $\sim 0.4 \pm 0.2$ counts for the single photon analysis, and hence 0.02 ± 0.01 leakage into the NR signal region (defined here as signal falling below the 10th percentile of an ER band calculated for a flat ER spectrum as discussed below).

A subdominant background is expected from NR signals from neutron scatters. Neutrons are generated in detector materials through (α ,n) reactions or from spontaneous fission of ^{238}U . Due to the 10–20 cm mean scattering length of MeV neutrons in LXe, most neutrons entering the detector will scatter multiple times (producing multiple S2 pulses) and so they can be effectively rejected. Muon-induced neutrons are typically produced far from the detector but can penetrate through the shielding; these constitute a negligible background at the 4850-ft depth of the Davis lab. In total, the NR background expectation is 0.08 events for the WS2013 search, in the S1 range of 1–50 phd [135, 136, 138]. Hence, for our single photon search and taking into account the DPE cut acceptance and the lower S2 threshold, NR interactions are expected to contribute <0.01 events in the NR signal region.

Interactions near the PTFE walls of the TPC field cage, mostly due to radon-daughter plateout, can appear with suppressed S1 and/or S2 signals — these peripheral interactions are known in LUX as "wall events". Due to the suppressed S2 signals, imperfect position reconstruction can bias their location into the TPC and an aggressive radial cut was required in LUX to deal with this significant population (this was especially problematic at the lowest energies). These events are suppressed for a fiducial cut that eliminates interactions outside

a radius of 18 cm. For the lower S2 threshold of 100 phd we adopt in our analysis, ~ 4 such background events are expected between 2 and 25 phd. For single-photon S1s we expect ~ 0.05 counts with the additional efficiencies applied.

6.2.2 Accidental coincidence rate

The largest background for the single photon search is due to accidental coincidences between upward-fluctuating dark counts and S2-only events. There are several detector pathologies that can generate S2-only events. Firstly, S2-only events are often due to a random association of individual electrons leaving the LXe surface following large energy depositions but in a long enough timescale such that the event is not included in the high activity exclusion periods. These electrons can pile up, leaving the LXe within a few μs of each other, and hence producing S2-identified pulses above the S2 threshold. Additionally, grid emission can produce multiple-electron signals, again faking a valid S2 pulse. Furthermore, interactions can occur in the gas phase above the anode, with electrons drifting down and producing S2 light close to the anode wires. Lastly, shallow events in which the S1 and S2 are close together (X-rays interact so close to the surface that their S1 and S2 pulses are effectively merged) and hence not resolved by the pulse finder contribute to the S2-only rate. The S2-only spectrum following all analysis cuts (discussed in detail in Section 6.3.1) is shown in Figure 6.3. The rate within the region presented in this figure is 0.001 c/s .

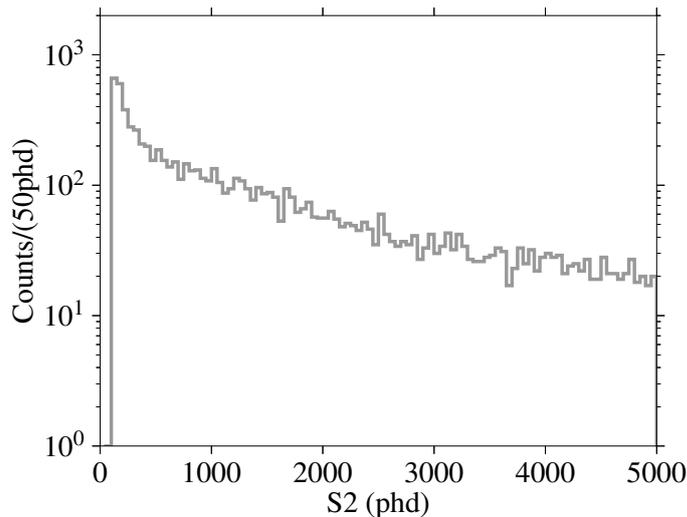


Figure 6.3: S2-only spectrum following quality cuts. The S2 area was corrected for (x,y) geometric variations and with a mean electron lifetime correction applied.

In Section 5.2.2 we calculated the leakage rate of dark counts into the DPE signal region to be $(2.2 \pm 0.2) \text{ cts/s}$ for all channels in the analysis. Given the maximum drift time cut we

Table 6.1: Expected background for various ranges in S2 space for the single-photon WIMP-search analysis of the 2013 LUX WIMP search exposure (95 live days, 118 kg fiducial mass). Entries are related to an ER band calculated for a flat spectrum and defined at the 10th and 90th percentiles. S2 pulse sizes are given for spatially-corrected variables; the 100 phd uncorrected S2 threshold corresponds to $\simeq 120$ phd.

S2 region	S2 size (phd)	Coincidences	ER	NR	Wall events	Total
> 90% ER	1,110–5,000	2.0 ± 0.2	0.02 ± 0.01	< 0.01	< 0.01	2.0 ± 0.2
10% – 90% ER	515–1,110	0.9 ± 0.1	0.4 ± 0.2	< 0.01	< 0.01	1.3 ± 0.2
< 10% ER	$\simeq 120$ –515	1.7 ± 0.3	0.02 ± 0.01	0.01 ± 0.01	0.05 ± 0.02	1.8 ± 0.3
Total	$\simeq 120$ –5,000	4.6 ± 0.4	0.4 ± 0.2	0.01 ± 0.01	0.05 ± 0.02	5.1 ± 0.4

apply, a time window of $267 \mu\text{s}$ exists before each S2 in which an "oversized" dark count pulse can be selected. Hence the probability that one such pulse occurs before an S2 is $P(1) = e^{-\lambda}\lambda = 5 \times 10^{-4}$. To predict the number of coincidences in a specific S2 range, we multiply this probability (which is the probability of one event being a coincidence) with the number of S2-only events that fall within the specified range. The S2-only rate within a region, and hence the number of coincidences, depends on the electron lifetime correction application. In the absence of depth information, the S2-only population is corrected for electron lifetime using a mean correction, by assigning a random drift time to each event. The standard deviation of the mean correction is also used to estimate an error on the number of S2-only events falling inside a specified S2 range, which will later serve as our signal region for different WIMP masses. Finally, this error is combined with the error in DC leakage rate and the two are treated as uncorrelated uncertainties. Further detail on the electron lifetime correction can be found in Section 6.3.1. The number of such dark count coincidences with S2-only events in the dataset is 4.6 ± 0.4 .

Table 6.1 summarises the background expectations calculated for the leading sources along with the random coincidence background for the single photon 2013 WIMP search. The total prediction for single photon events was 5.1 ± 0.4 counts before ER-NR discrimination, with a fraction of these being compatible with NR signals (those near the bottom of the ER band).

6.3 WIMP search

6.3.1 Searching for single-photon S1 events

Having confirmed that our new event selection is robust and consistent with that of the original analysis, we proceeded to perform a search for single-photon S1 events in the WIMP search dataset. Here high activity exclusion periods were removed from the original dataset

and events with S2s preceded by single-photon pulses consistent with DPE emission were selected. In this section we describe each cut applied to the obtained events in turn and how the composition of the dataset changes with its application. Note that even though we did not pursue a blind analysis, the main quality cuts described here have been kept unaltered from the original analysis of the WS2013 dataset, and where variations have been applied these are discussed.

The first cut was the lower and upper S2 thresholds, as applied in the tritium calibration, placed at 100 phd and 5000 phd. Note that the threshold here is lower than in the standard LUX 2-fold analyses [127, 136], as this analysis will necessarily lower the energy threshold. The number of single scatters with a single-photon S1 signal and S2s within the defined region was found at this stage to be 1113, many of which are due to accidental coincidences before the application of quality cuts.

Following the definition of the S2 region of interest, the cleanliness cut, previously defined in the tritium and original WS2013 analysis, was applied as shown in Figure 6.4, removing 534 events. The cleanliness cut has a severe effect on the dataset as it removes a large number of events with misidentified S2 pulses following large energy depositions, events for which the area fraction contained in the S1 and S2 pulses is very small.

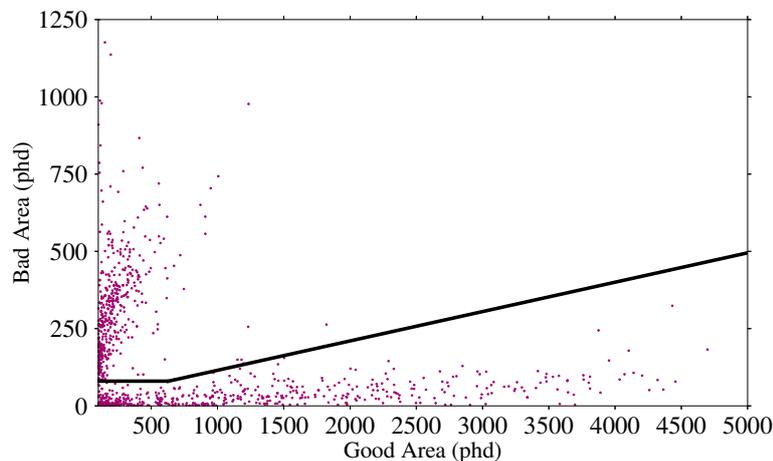


Figure 6.4: Cleanliness cut applied on the good and bad area quantities for events with a single-photon S1 in the WIMP search dataset. We present events following event selection, exclusion of high activity periods and application of the S2 lower and upper thresholds, placed at 100 and 500 phd, respectively.

The (x,y) positions of the resulting events were corrected using the algorithm discussed in Section 5.2.3. The fiducial volume was defined as described in both Chapter 5 and in the original WS2013 analysis. The radial cut was placed at 18 cm (same as the original WS2013 analysis [136] and smaller than the 20 cm radial cut used in the re-analysis [127]). We chose

this smaller fiducial radius as we have lowered the S2 threshold and hence we desire to limit wall events (which appear at small S2s) from our background rate. A fraction of the 534 events ($\sim 50\%$) remaining following the cleanliness cut was seen at drift times of $2\text{--}3\ \mu\text{s}$, consistent with interactions originating near the gate grid, which is located $0.5\ \text{cm}$ below the LXe surface. Of the 534 events, only 16 fall within the $118\ \text{kg}$ fiducial volume, as illustrated in Figure 6.5.

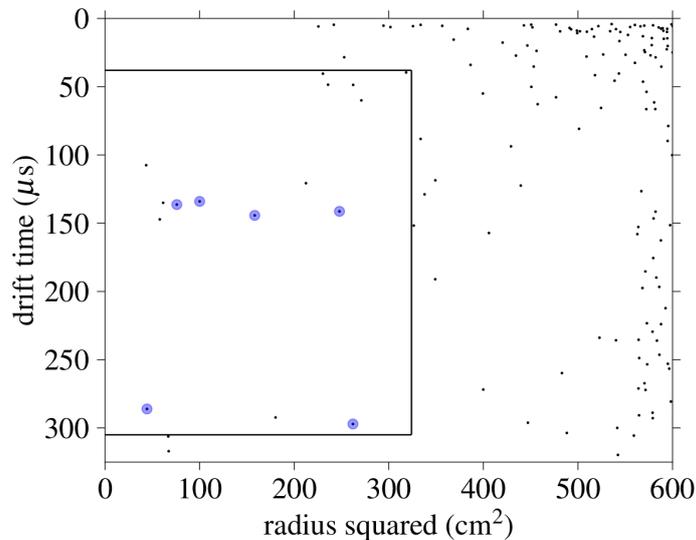


Figure 6.5: Corrected position of single-photon S1 events appearing in the WIMP search following event selection, S2 threshold enforcement and application of the cleanliness cut. The fiducial region used for the WIMP search analysis is also shown. Events illustrated in blue show the final set of events passing all cuts, as defined in text.

The two additional cuts discussed in Section 5.2.3 were also enforced: these tackle interactions in the gas phase and shallow events, as well as those in which the recovered single-photon S1 was an afterpulse of "other" identified pulses. The latter removed zero events while the former removed 4 events, all of which were visually inspected and found to consist of pathological or misidentified S2 pulses, such as those shown in Figure 6.6.

Lastly, a cut was applied removing events in which the S2 was preceded by more than two pulses identified as single photons. This cut removed 2 additional events, in which the S2 was preceded by 3 single-photon-like pulses seen in the same PMT — but separated in time by more than $1\ \mu\text{s}$. The expected number of such accidental coincidences for the full WIMP search dataset is expected to be negligible. The Poisson probability for a given PMT (with an average count rate taken from Table 4.1) to fire three times due to a dark count within the drift time window before an S2 is 5.5×10^{-9} . Hence, the probability that any of the 110 active PMTs gives such a signal in an event, multiplied by the total number of S2-only events passing analysis cuts, gives the expectation of this class of accidental coincidence events to be ~ 0.005 counts for

the full dataset. It is very unlikely that these are generated by real particle interactions in the fiducial volume.

The analysed dataset at this point, following all cuts, consisted of 10 events. These were visually inspected using Visualux. In four of these events the S1 pulse was found to consist of double pulses within a 300–400 ns window, both detected in the same PMT, instead of a time-coincident DPE response. One such example is given in Figure 6.7. Thus, these events are not viable signal events for our single-photon S1 search and were removed from the dataset (in any case this large time separation is not consistent with LXe scintillation). It is important to note that the number of such background events observed does not necessarily reflect a high rate of this process in the WIMP search dataset, but the fact that any such events are likely to show up in our final dataset as they integrate to two-phe and hence appear within the DPE signal region. It is also important to note that this topology was also not included in our background estimate and as such subtracting them does not affect the alignment with our background model. This is different to the accidental background of S2-only pulses coinciding with over-fluctuating dark counts and hence does not bias our accidental coincidence calculation in that sense either. Similarly to our tritium calibration analysis hand-scanning of events was chosen because we were looking to identify any potential new background sources, as the benchmarking and novel nature of this analysis meant that it was impossible to anticipate all possible pathologies prior to the search. The small number of such events also meant that using a minor percentage of our dataset to look for new background sources prior to employing a full analysis could lead to unidentified pathologies and as such hand-scanning was preferred. This procedure could prove useful for future experiments as we have now established a more robust and complete background model for such analyses.

The S2 pulse area of the remaining six events entering our final analysis is illustrated in Figure 6.8. The number of these events is consistent with the expectation presented in Table 6.2, mainly due to accidental coincidences. The presumed origin may explain their non-uniform pulse shape, as these are likely due to gate grid radioactivity. Such particle emission happens in a region of high field that seems to distort the pulse shape. Currently, a LUX study is using Machine Learning to perform pulse shape discrimination on such signals in an attempt to lower the S2-only rate for a full S2-only analysis. In the analysis presented here this background requires coincidence of an S2-only event with a large area DC, significantly reducing — but not eliminating — this rate.

The S2 pulse area for the remaining 6 events was corrected for geometric effects. For the electron lifetime correction, we use maps of the electron lifetime correction at each drift time, as in the tritium calibration analysis. In this case, we are investigating a longer dataset during which the electron lifetime is not constant. As such, we used the single scatter dataset presented in the previous section to create a database of corrections according to the drift time of a particular event and the time during the WIMP search at which it was recorded. Figure 6.9

Table 6.2: Total expected background and observed counts for various ranges in S_2 space for the single-photon WIMP-search analysis of the 2013 LUX WIMP search exposure (95 live days, 118 kg fiducial mass). Entries are related to an ER band calculated for a flat spectrum and defined at the 10th and 90th percentiles. S_2 pulse sizes are given for spatially-corrected variables; the 100 phd uncorrected S_2 threshold corresponds to $\simeq 120$ phd.

S_2 region	S_2 size (phd)	Total expected	Observed
> 90% ER	1,110–5,000	2.0 ± 0.2	3
10% – 90% ER	515–1,110	1.3 ± 0.2	2
< 10% ER	$\simeq 120$ –515	1.8 ± 0.3	1
Total	$\simeq 120$ –5,000	5.1 ± 0.4	6

presents events from three of the time-subsets along with the corresponding fits that were used to correct the newly recovered single-photon S_1 events.

Figure 6.10 shows the observation in (S_1, S_2) space both for interactions at ≥ 2 -fold and for the six single-photon signals. The numbers falling into each S_2 region are consistent with the background expectations as summarised in Table 6.2. Two events were observed below the ER band (expectation value was 1.8 ± 0.3), i.e. below an S_2 pulse area which contains the full NR acceptance for WIMP signals.

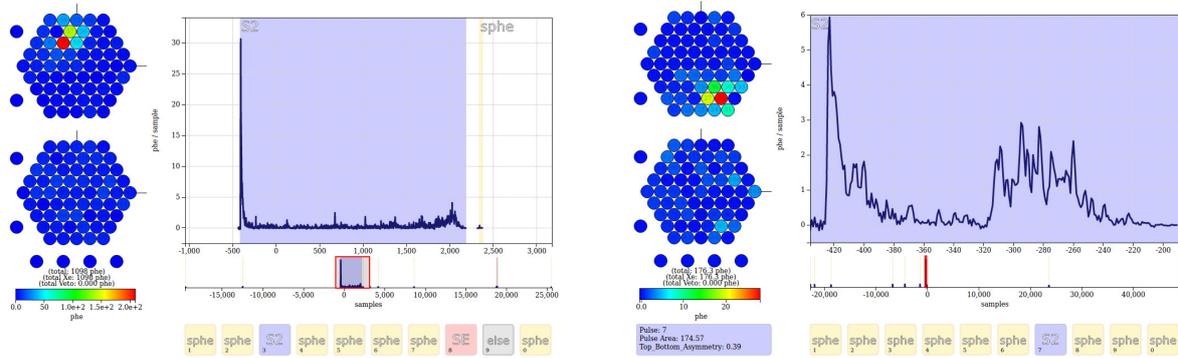


Figure 6.6: Illustrations of example pathological events removed using the S2 pulse quality cut. These are likely gas events due to the initial ionisation above the anode followed by electron drift to the anode.

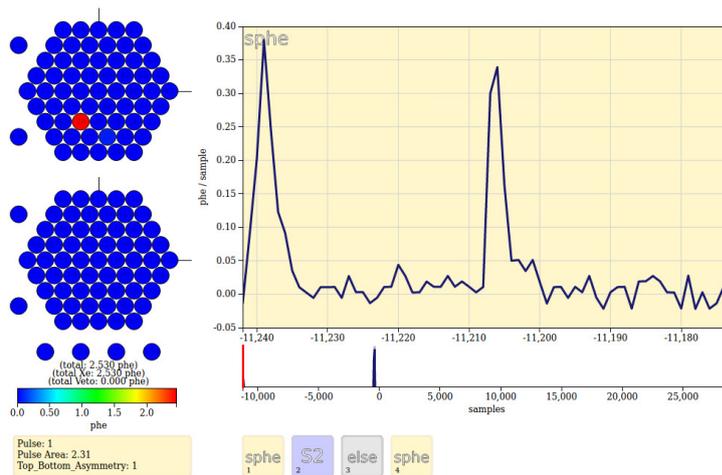


Figure 6.7: Example event in which the identified single detected photon pulse is a mis-classification of a DPE signal seen in one PMT (in which case the detection would be time-coincident), probably due to a PMT artifact. Such events were removed from our dataset.

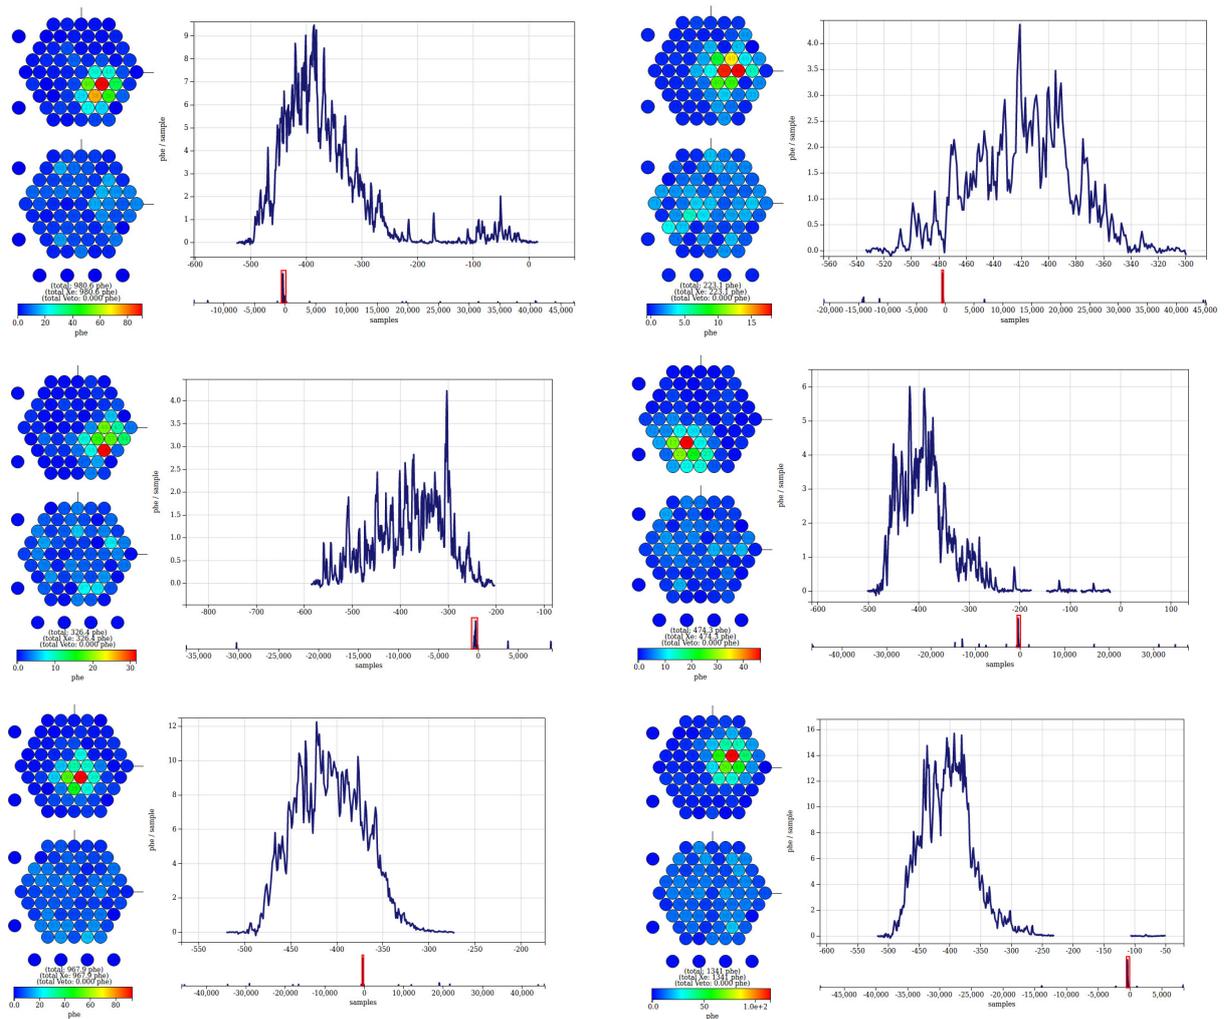


Figure 6.8: Visualux illustrations of the S2 pulse of the final single-photon S1 WIMP search events. Some of the different pulse shapes are suggestive of low-energy particle radioactivity emitted by the electrode grids; such events are accounted for by the coincidence component of the background model.

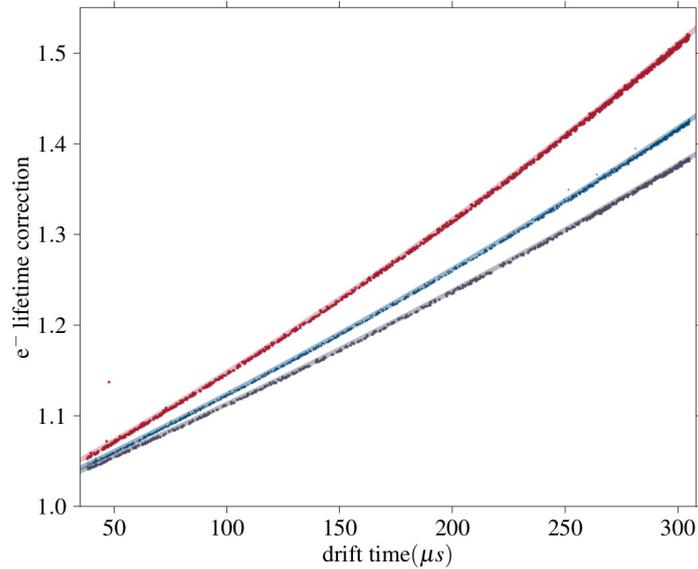


Figure 6.9: Electron lifetime correction at each drift time for single scatter events found in the WIMP search dataset, before application of pulse area cuts, fitted with a third degree polynomial function. The three colors represent three different subsets of the WIMP search dataset, corresponding to different times during the search.

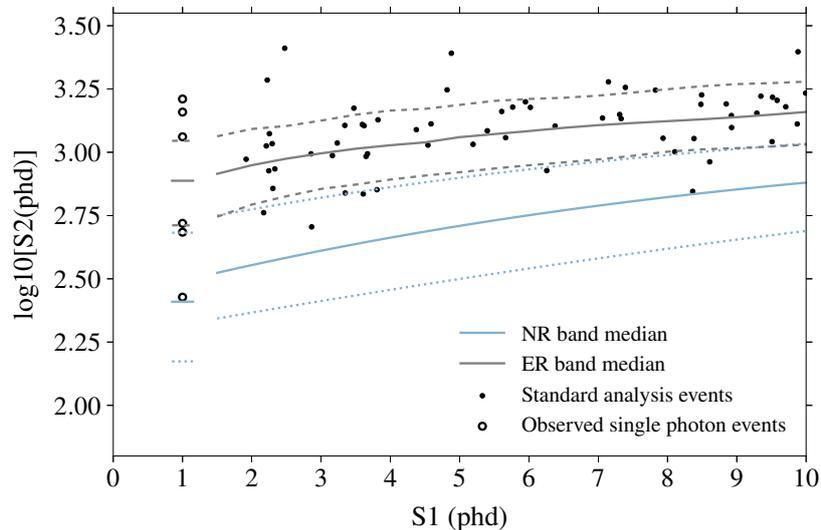


Figure 6.10: Events observed in the 2013 LUX WIMP search exposure of 95 live days and 118 kg fiducial mass. Solid black markers represent events meeting the S1 2-fold coincidence requirement, while the 6 hollow markers indicate those with an S1 of a single photon which are the focus of this analysis. Distribution contours for an ER beta spectrum (grey) and an example 50 GeV WIMP signal (blue) are indicated at the 50th (solid), 10th and 90th (dashed) percentiles of S2 at given S1. These percentiles are shown separately at 1 phd, with the S2 threshold lowered to the 100 phd value (uncorrected) adopted in this analysis.

6.3.2 Signal model

In this work we consider spin-independent (SI) WIMP-nucleon scattering. The velocity profile used follows a standard Maxwellian velocity distribution with $v_o = 220$ km/s, $v_E = 245$ km/s (taken for the period of the year when the relative velocity is higher) and $v_{esc} = 544$ km/s. The WIMP mass density was taken to be $\rho_o = 0.3$ GeV/cm³. During computation of the event rate for SI scattering, a coherent interaction between the DM particle and the xenon nucleus is assumed. This leads to an enhanced differential rate in recoil energy, compared to spin-dependent rates. The differential energy spectra — some of which are shown in Figure 6.11 — were then processed through NEST with LUX detector parameters, for the WS2013 period of operation, to generate S1 and S2 observables.

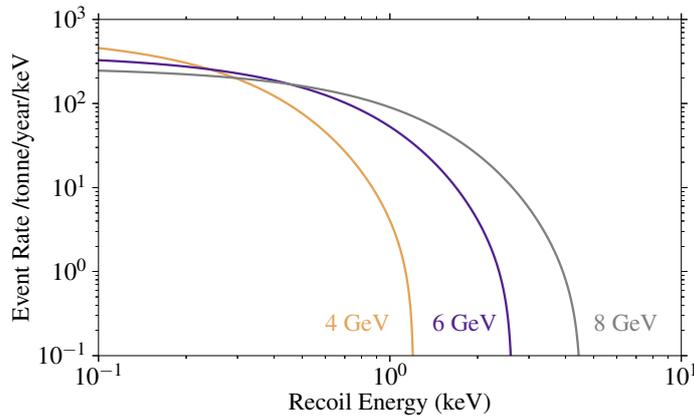


Figure 6.11: Differential recoil energy spectra for three different WIMP masses: 4, 6 and 8 GeV at a cross-section of 1.0×10^{-45} cm².

The NR scintillation and ionisation yields as a function of recoil energy as modelled by NEST v2.0.0 are presented in Figure 6.12. In our analysis, the NR model cut-off was placed at 0.3 keV, which corresponds to the lowest-energy ionisation yield data point in a newly published measurement [174]. Enforcement of the model cut-off means that below this energy a zero light and charge yield are assumed. Even though the newly published NR measurements motivated in part this lower cut-off, they were not included in the NEST v2.0.0 model development. A newer NEST model (NEST v2.0.1) that became available after this analysis had been completed is also shown in Figure 6.12. This was used to assess the effect of yield variations. A more detailed discussion of the NR NEST models and their effect on our result is presented in Section 6.3.4.

NEST v2.0.0 was used to simulate WIMP signals between 2.5 and 15 GeV and the efficiency of those cuts as a function of recoil energy or WIMP mass was investigated. These included efficiency loss due to the selection of events with single-photon S1 pulses, the (lower) S2 analysis threshold and the application of the energy threshold to the simulated signal model

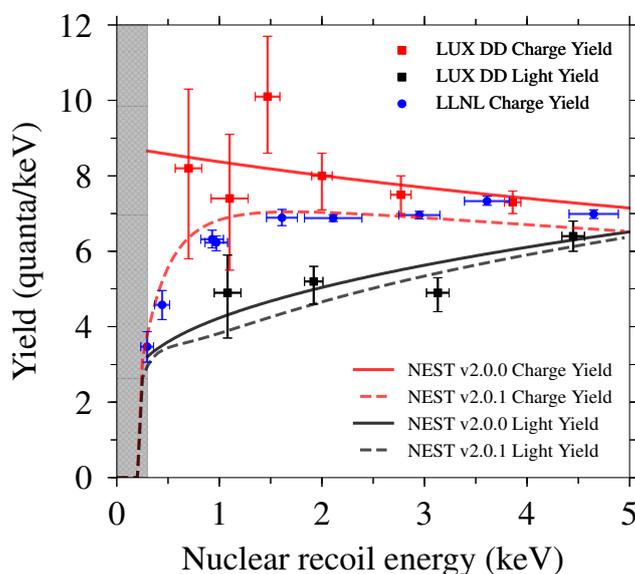


Figure 6.12: The NR ionisation (red) and scintillation (grey) yields as a function of recoil energy given by NEST v2.0.0. Red and black points indicate LUX measurements from D-D neutron calibration [113] and blue points illustrate recent measurements, performed at the Lawrence Livermore National Laboratory (LLNL), were published too late to be incorporated in the NEST v2.0.0 models. The NEST v2.0.1 NR model yields incorporating the new data are shown using dotted lines. The newer model became available after the analysis was complete and hence was not used for the baseline simulation but instead was incorporated to understand the systematic effect of yield variations.

(as discussed in the previous paragraph). To study the effect of the S2 threshold, events in the WIMP simulations in which the S1 consisted of a single photon were selected and the S2 threshold cut was varied between 85 phd (~ 3 emitted electrons) and 175 phd (~ 6 emitted electrons, corresponding to 8 electrons drifting in the liquid). The S2 threshold in our analysis was placed at 100 phd, which resulted in an efficiency improvement of approximately one order of magnitude for a 4 GeV WIMP compared with the cut applied in the original analysis, which was placed at 165 phd. For lower mass WIMPs the sensitivity improvement for a lower S2 threshold is even more pronounced as the S2 distribution shifts to lower values. The results of this study are illustrated in Figure 6.13, where the efficiency of different S2 thresholds, following single-photon S1 event selection, are shown for three example WIMP masses.

The successive effect of selecting single photon events, applying the 100 phd S2 threshold and enforcing a hard signal cut-off applied at 0.3 keV in simulations is shown in Figure 6.14. The 0.3 keV cut-off removes 23% and 0.001% of simulated events at 3 and 6 GeV, respectively. The single photon event selection might bring about increased sensitivity for low mass WIMPs compared to a standard 2-fold analysis, but the lightest WIMPs will produce genuine S2-only events at the lowest masses. Adding to this energy-dependent efficiency, the DPE signal acceptance calculated in Section 5.2.1, and the cleanliness cut efficiency (98.4%) are applied.

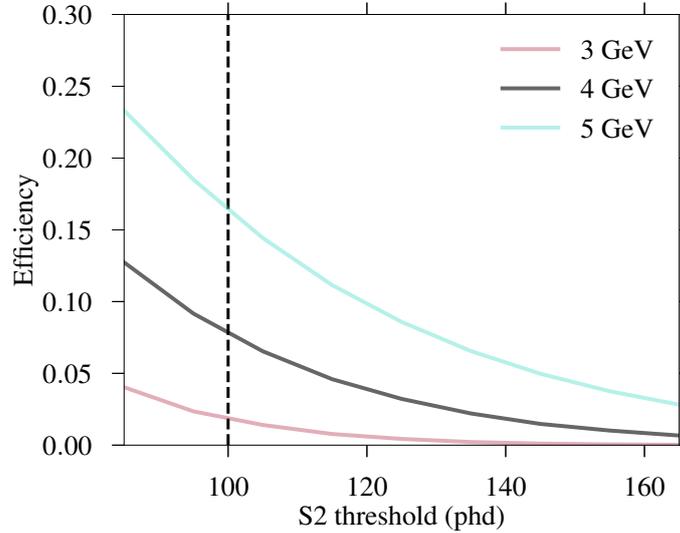


Figure 6.13: WIMP signal acceptance as a function of S2 threshold for single photon events simulated with NEST v2.0.0 before application of cuts and for 3 different WIMP masses.

Finally, for each WIMP signal model we define an S2 region of interest. In general, the NR signal region of interest for light WIMPs lies below the NR band represented in Figure 6.10, which was derived for a 50 GeV WIMP spectrum (which coincides approximately with the NR spectrum obtained with a D-D neutron generator). This is due to the fact that, at a particular S2 size, only events with an over-fluctuating S1 pulse fall above threshold. For the very low masses considered here this effect is even more extreme, with the predicted S2 range falling well below the NR band as defined for higher energy recoils.

Thus, to determine the precise NR acceptance region for a single-photon S1, the S2 acceptance regions were defined between an S2 pulse area (uncorrected) of 100 phd and a maximum value that retains 95% acceptance. This avoids most of the ER background while retaining high signal efficiency — exploiting the ER-NR discrimination that remains effective even for these very small signals. A histogram of the S2 uncorrected pulse area, as simulated for a 4, 6 and 8 GeV WIMPs is shown in Figure 6.15, along with the S2 upper bound of the defined region of interest. The histograms present events following all event selection, analysis cuts and efficiency reduction for an exposure of 10^6 kg·days and cross-section of 10^{-40} cm²; the integrated number of events for the 4, 6 and 8 GeV models is $\sim 8,000$, 80,000 and 200,000, respectively.

Background expectations for each WIMP mass, within this region of interest, were estimated and the final acceptance and background expectation are also presented in Figure 6.14. The uncertainty on the background expectation is determined by the position corrections on the S2 spectrum, which affect the number of events falling within the S2 acceptance region. The overall final acceptance includes all analysis and threshold cuts as well as acceptance in the

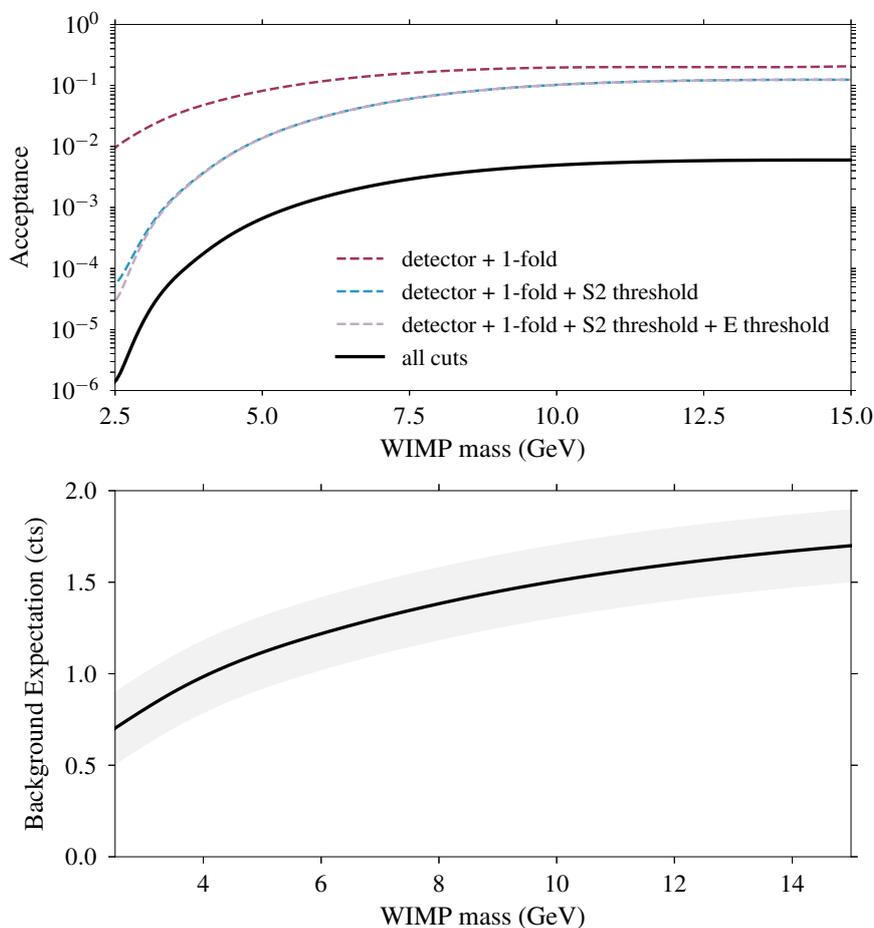


Figure 6.14: *Top:* Efficiency of analysis cuts for a range of WIMP masses for events simulated with NEST v2.0.0. *Bottom:* Background expectation for a range of WIMP masses for the 2013 WIMP search exposure calculated for the appropriate acceptance region for each mass. Shaded regions represent corresponding statistical and systematic uncertainties combined.

region of interest. The upper boundary of this region moves upward with increasing particle mass, and captures the first observed event for a 5.3 GeV WIMP model at $\log_{10}(S2) = 2.4$.

Figure 6.16 shows the overall acceptance as a function of nuclear recoil energy and in comparison with the published acceptance for the 2-fold WIMP search analysis of the 2013 dataset. The single photon analysis presented here becomes more sensitive than the published result below ~ 1.2 keV.

6.3.3 Upper limit on spin independent cross section

The calculation of upper limits on a parameter or theory is one of the most important problems in experimental physics. As previously mentioned, in the analysis presented here we chose a cut-and-count analysis, which was preferred to a full profile likelihood ratio analysis for

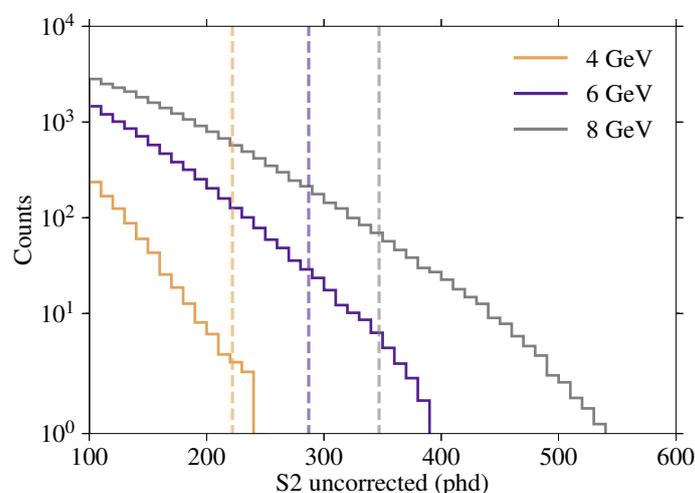


Figure 6.15: Expected WIMP signal for a mass of 4, 6 and 8 GeV for single-photon S1 pulses only, following all analysis cuts, for an exposure of 10^6 kg·days and cross-section of 10^{-40} cm². The lower bound of the defined signal region is placed at an uncorrected S2 of 100 phd while the lower bound, shown here in the dashed lines was chosen so that it contains 95% of the signal.

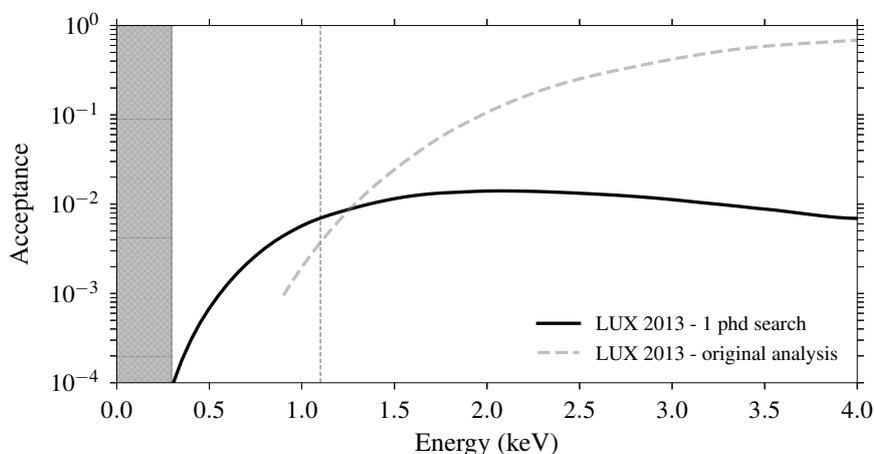


Figure 6.16: Efficiencies for NR event detection as presented in the original analysis of the LUX 2013 data [127] (dashed gray) and as estimated using NEST v2.0.0 for the single-photon search presented here (black). These efficiencies include detection of single scatter events passing the S1 and S2 thresholds. In the original analysis the S1 threshold required at least two PMTs detecting photons and a minimum uncorrected S2 of 165 phd. The single photon analysis includes events with a single-photon S1 producing double-photoelectron emission and an uncorrected S2 threshold of 100 phd.

its transparency. In any case the two are expected to give comparable results in the case of background free experiments, such as ours over most of our defined signal region. A unified cut-and-count approach proposed by Feldman and Cousins [175] is often used, but suffers from a few practical drawbacks such as its inability to deal with background rate uncertainty. An

alternative method, first published by Rolke [176], based on the likelihood ratio test statistic, has been shown to yield correct coverage rates with the ability to include both efficiency and background uncertainties. A description of Rolke's method and confirmation of its statistical coverage in our particular case are presented in Appendix A.

As hinted through Table 6.2 our observation is consistent with the background only hypothesis and patently there is no clear signal (lower limit is consistent with zero at all masses). The ROOT implementation of the Rolke calculation was used to set 90% C.L. upper limits on the number of signal counts at each mass. We model the background and efficiencies as Gaussian and the signal as Poisson.

The uncertainty on the signal acceptance included both the uncertainty on g_1 and the DPE cut acceptance error and combined to 11%. Yield uncertainties were not included but the effect of yield variation is presented through the use of different yield models and is discussed in Section 6.3.4. The expected background for an example 4 GeV WIMP is 1.0 ± 0.1 as seen in Figure 6.14. A small amount of undercoverage is observed for small signals, as shown in Appendix A, but this is deemed acceptable (e.g., a similar effect is reported in the original publication by Rolke [176]).

The upper limit on the number of counts at each mass is shown in Figure 6.17, along with the median (background only) result and the 1σ and 2σ bands of background-only trials shown in green and yellow. The observed upper limit is seen to deviate from the median expectation below 5.3 GeV. This is due to the fact that no observed events fall within the S2 region of interest for WIMP masses below 5.3 GeV, while the background expectation is still non-zero. From these results we calculated the corresponding upper limits on the spin-independent elastic scattering WIMP-nucleon cross-section, assuming the same astrophysical parameters considered in [127] and summarised earlier in this chapter. These limits are plotted in Figure 6.17. The step seen at 5.3 GeV marks the WIMP mass beyond which the signal region includes the first observed event.

These limits were capped at 2.3 counts below 5.3 GeV to prevent our result for surpassing that of a background-free experiment; hence, this produces a conservative result. The final cross section result is shown in Figure 6.18 along with previous competitive measurements. This analysis utilised exclusively single-photon events and not those at 2-fold and above, producing competitive limits below 5.5 GeV. There are, however, differences between this and the original analysis that contribute to this improvement.

Overall, at 4 GeV we observe an improvement between the single photon analysis and the previously published result (LUX 2014) of approximately two orders of magnitude. It is important to note that this improvement is due to several effects. The lowering of the S1 threshold using the DPE effect allowed probing lower energy recoils, and hence it improved

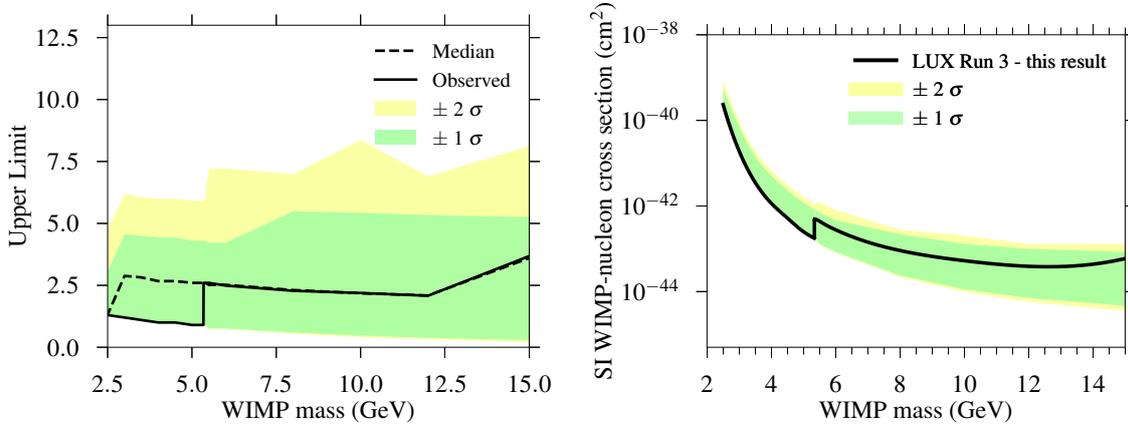


Figure 6.17: *Left:* Upper limit on the number of counts at each mass along with background only median, 1σ and 2σ bands of background-only trials shown in green and yellow. *Right:* 90% C.L. upper limits on the spin-independent WIMP-nucleon cross-section obtained using the single-photon population producing double-photoelectron emission in the LUX 2013 WIMP search. The observed limit with a 0.3 keV NR energy cut-off is shown in solid black, with 1σ and 2σ bands of background-only trials shown in green and yellow.

the sensitivity at lower masses. Clearly, the background suppression due to the low acceptance of single photon events enabled a decrease in S2 threshold.

Additionally, the energy cut-off applied in simulation was lowered from 1.1 keV to 0.3 keV and hence also accounts for part of the sensitivity improvement. The result of applying this higher energy cut-off to the single photon analysis is further discussed in Section 6.3.4. As observed in Figure 6.19, at 4 GeV, applying the higher (1.1 keV) cut-off to our analysis, results in an upper limit approximately equal to the previously published 2-fold analysis. This does not mean that the improvement at 4 keV is alone due to the relaxation of the cut-off but that any improvement between the 2-fold and single photon analyses lies below 1.1 keV. In other words, if the 2-fold limit were to be relaxed to include events below 1.1 keV it would not reach the single photon result. This is also supported and represented more clearly by the efficiencies shown in Figure 6.16, where we can see that any efficiency improvement happens below ~ 1.2 keV.

It is impractical to assess the effect of lowering the energy cut-off in the standard analysis instead, as this employed a different NEST model and a more powerful statistical analysis, with selection cuts optimised for that analysis. For this reason it is also difficult to assess the effect of lowering S2 threshold in the 2-fold analysis and hence deduce how much of the improvement arises from the use of single photons alone. For an approximate assessment of these effects, we produced an upper limit with a 2-fold requirement, using the same model as the one used for the single photon analysis. For this approximate study we assumed a background free experiment for the 2-fold analysis. It was found that the improvement in

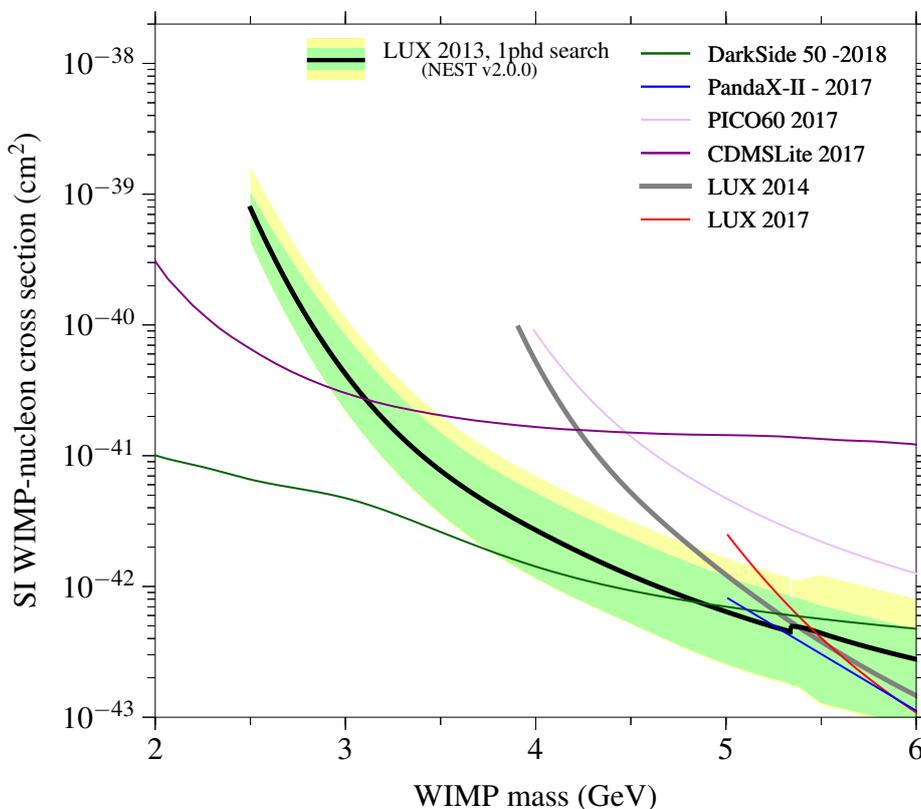


Figure 6.18: 90% C.L. upper limits on the spin-independent WIMP-nucleon cross-section obtained using the single-photon population producing double-photoelectron emission in the LUX 2013 WIMP search. The observed limit with a 0.3 keV NR energy cut-off is shown in solid black, with 1σ and 2σ bands of background-only trials shown in green and yellow. Also shown are the previous results from the LUX 2013 search [127] (gray), the LUX complete exposure result [134] (red), and from DarkSide-50 [155] (green), PandaX-II [177] (blue), PICO60 [178] (lilac) and CDMSLite [179] (purple).

sensitivity is approximately equally split in four consecutive improvements, due the use of a different yield model and statistical analysis, the reduction of the energy cut-off, the lowering of the S2-threshold, and the use of single photons (with DPE cut). In reality, reducing the S2 threshold in the standard search results in 6 events observed at low S2 values, and hence our assumption of a background-free experiment yields a better upper limit for the 2-fold case than the one we would have following a full re-analysis of the published 2-fold result. Hence, the estimate of an improvement equally split in four consecutive improvements is conservative and in reality reducing the S2 threshold in the standard result would yield a smaller improvement than that of our conservative estimate, leading to an even stronger improvement due to the use of single photon events. Note also that the comparisons here are for the upper limit capped at 2.3 counts.

6.3.4 Effect of model assumptions on upper limit

Inevitably, the new analysis demands a good understanding of scintillation and ionisation yields at very low energies, and in particular the light yields for both ER and NR rely on models reaching below the lowest measurements at present. The LUX collaboration and others have been working for over a decade to extend these measurements and improve their systematic uncertainty, and major progress has been made for liquid xenon — and this is set to continue.

The lower (0.3 keV) model cutoff for NR interactions was adopted here as this is approximately the energy at which the first ionisation electron is expected from extrapolating to very low energies the power-law behaviour that is assumed in the Lindhard model [111, 111, 112]. Recent measurements [174] confirm electron release at 0.3 keV in liquid xenon but are consistent with a lower ionisation yield than the one assumed in the NEST v2.0.0 model. To understand how such yield variations affect the WIMP sensitivity, we also present an upper limit result obtained using the newer NEST v2.0.1 NR model, shown using dashed lines in Figure 6.12. The NEST v2.0.1 NR model incorporates data from the recent publication and is also consistent at the 1σ level with previous measurements conducted with LUX (as shown in Figure 6.19). While the NEST v2.0.0 model assumes total scintillation and ionisation yield anti-correlation below the smallest light yield measurement, the new NEST v2.0.1 model allows the total number of quanta to decrease at lower energies. Use of the NEST v2.0.1 model resulted in a slightly higher observed limit (shown as a dotted black line in Figure 6.18) but within the 2σ band of the NEST v2.0.0 upper limit.

6.3.5 Evaluating the effect of the DPE cut

It is worth considering whether a DPE cut is required in the first place, i.e., would a full analysis including all "single photon" pulses (including dark counts) be equally sensitive, despite the much higher background? Even though both analyses may be useful, the signal-to-background ratio is markedly better with the DPE cut. For the masses explored in Figure 6.18, accepting all single photon pulses would increase the efficiency by ~ 20 -fold, while the background expectation would increase by a factor of 300–500 over the numbers in Figure 6.16. To study the significance of both searches (with and without use of a DPE cut) we investigated the discovery significance (Z) for each DPE cut value and a range of signal counts. The result of this study is shown in Figure 6.20 for a 4 GeV WIMP. The discovery significance shown represents the median discovery significance (Asimov formula) with uncertain background, using a profile likelihood ratio test, and is calculated using:

$$Z = \left[2 \left((s+b) \ln \left[\frac{(s+b)(b+\sigma^2)}{b^2 + (s+b)\sigma^2} \right] - \frac{b^2}{\sigma^2} \ln \left[1 + \frac{\sigma^2 s}{b(b+\sigma^2)} \right] \right) \right]^{1/2}, \quad (6.1)$$

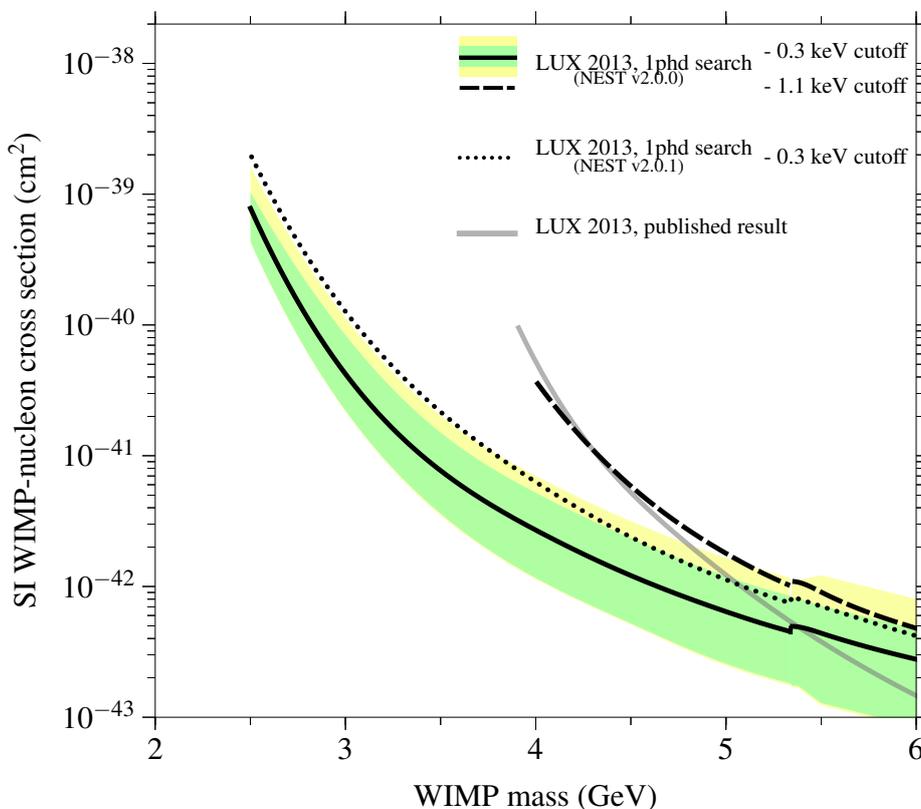


Figure 6.19: 90% C.L. upper limits on the spin-independent WIMP-nucleon cross-section obtained using the single-photon population producing double-photoelectron emission in the LUX 2013 WIMP search. We show the observed limit with a 0.3 keV NR energy cut-off in solid black, with 1σ and 2σ bands of background-only trials shown in green and yellow. The dashed black line is derived from the same analysis but with a model cut-off at 1.1 keV. Both of these results correspond to the NEST v.2.0.0 model shown in Figure 5.13. The upper limit using a 0.3 keV NR energy cut-off with the newer NEST v.2.0.1 model is shown using a dotted black line.

where s represents signal counts, b is background expectation and σ is the uncertainty on the background.

According to Figure 6.20 a full analysis of the LUX dataset, including all single-photon S1 events would yield a 90% C.L. upper limit of the same magnitude to that including the DPE cut (with the analysis without a DPE cut found to be a factor of ~ 3 less sensitive at 4 GeV). In this case, the sensitivity is now dominated by the systematic uncertainty on the background expectation. This is largely determined by the population of S2-only events that are predicted to fall within the S2 signal region for each mass: in the absence of meaningful depth (z) information, correcting the S2 pulse (e.g. for finite electron lifetime) cannot be done accurately event-by-event. In addition, for such an analysis the possibility of events presenting both signal and background (one or more dark counts coinciding with a real single-photon S1)

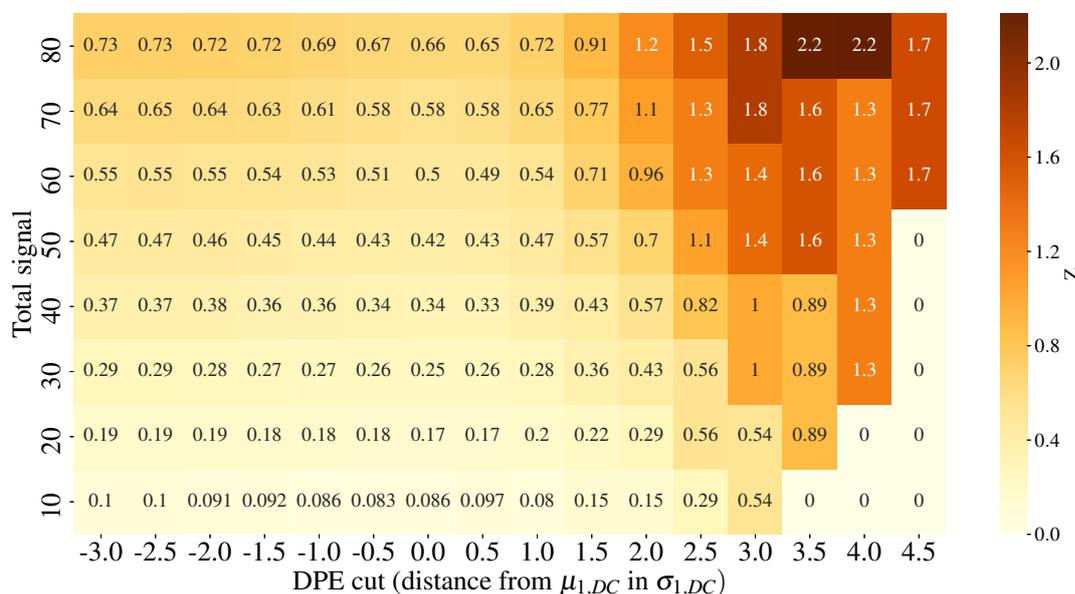


Figure 6.20: Discovery significance for a 4 GeV WIMP and variable DPE cuts and signal expectations. The x-axis represents the placement of the DPE cut compared to the mean and standard deviation of the DC distributions presented in Chapter 4.

becomes non-negligible and needs to be addressed. Moreover, systematic uncertainties play a bigger role in a discovery situation, and the new analysis is better able to control these.

6.4 Summary

In summary, we applied the single photon methodology to a search for low-energy NRs in the LUX WS2013 search dataset. In this process we used hand-scanning to identify potential new background sources. Even though this is not ideal for a robust and unbiased analysis, it was necessary given the benchmarking nature of our analysis and lead to a more complete background model that future experiments looking to employ a similar search could benefit from. Following our search we performed a statistical analysis improving the spin-independent scattering cross-section limits significantly between 2.5 GeV and 5 GeV WIMP mass compared with the previous analysis — where a standard 2-fold threshold had been applied to the S1 signal. We also studied the effect of model assumptions on the cross-section limits and specifically the effect of varying the light and charge yield models in the energy range between 0.3 and 1.1 keV, where limited measurements of the yields exist.

Chapter 7

Lowering the scintillation threshold in LZ

In this chapter we begin by demonstrating our ability to calibrate the single photon response of LZ PMTs using simulations of early datasets which is an essential step for a successful DPE analysis. To do this we make use of the simulation framework and, in particular, a mock dataset of background interactions expected to be recorded in the first few days of operation. Subsequently, we investigate the application of the DPE technique to LZ and estimate the potential improvement it can bring to its sensitivity to the rare nuclear recoil interaction caused by ^8B neutrinos.

7.1 Single photon calibration of the LZ TPC PMTs

The LZ simulation framework incorporates a variety of software packages. It begins with BACCARAT [180] built on the GEANT4 toolkit [181] and is specifically tuned for noble liquid detectors. At this stage energy deposits are recorded on a geometry-component basis, and NEST is used to convert them to scintillation and electroluminescence photons. Optical processes in GEANT4 are then used to propagate the primary photons to the PMT faces. The Detector Electronics Response (DER) software is then used to produce full digitised waveforms. The resulting simulations are saved in the DAQ event format and later analysed using the data processing framework, LZap. Further detail on the simulation framework can be found in Refs. [180, 182].

The simulation chain and other analysis tools and techniques were evaluated and validated with a series of three mock data challenges (MDC). These challenges have been designed to ensure that LZ is physics-ready as soon as data taking begins. In this study we will use simulated data from the final MDC to calibrate the single photon response of the LZ PMTs using realistic datasets.

In LZ PMTs are expected to operate at a gain of 3.5×10^6 and single photon pulses will generate waveform amplitudes of a few mVs. Similarly to the calibration procedure presented

in Chapter 4 for LUX, LZ PMTs require single photon calibration such that the digitised pulse voltage is converted into physical units of phd. Our calibration is again defined such that the mean area of a pulse due to a single photon at the wavelength of Xe scintillation light is 1 phd. In reality, these *in situ* calibrations can be compared with the measurements performed on a sample of 35 LZ PMTs in Ref. [165].

A characterisation of the full single VUV photon response in LZ requires obtaining sources of single photons with adequate statistics. Traditionally, LED signals were used for the SPE calibration, but the blue LED photons do not create DPE response. Hence, we calibrate using both an LED calibration, to understand the SPE response which is similar to that produced by dark counts, and a source of VUV single photons, to characterise the full response including DPE.

The TPC top and bottom arrays contain 24 LEDs each, that will be used to perform periodic single photoelectron and afterpulsing calibrations. The LEDs incorporated in one array shine on the opposite array. An LED source is incorporated in BACCARAT, allowing the production of simulations using selectable LEDs and photon numbers. Here we use a dataset produced using 48 LEDs at 18.75 Hz each and 10 photons per pulse. This results in PMT distributions mostly composed of single photoelectron signals — with contributions from multiple photon hits accounting for $\ll 1\%$ —, as shown in Figure 7.1 for a set of example PMTs.

To obtain sources of single VUV photons we use a simulated dataset of the background only expectation for the first 8 days of data taking in LZ and select low energy events (< 10 keV), with small S1 signals ($\sim 5,000$ events following selection cuts). A selection of cuts is applied, optimised to balance having enough statistics to characterise the response of each PMT and to minimise the multiple overlap possibility. The resulting distributions per PMT are fitted to obtain the SPE and DPE mean, resolution, and DPE fraction, with a component for double-photon hits also taken into account in the fitting process. It is important to note though that using the Poisson counting method in this case to account for such pile-up pulses suffers from significant limitations as the broad cut enforced on the S1 area and the large size of the TPC volume in which interactions can take place often mean that the Poisson mean is not in fact constant across the selected region. Thus, it is important to make sure we enforce strict enough cuts such that the Poisson mean expectation per channel is very small (< 0.05 phd) and hence contamination from two-photon signals negligible. Example fits for the VUV response using S1 light from the selected low-energy background interactions are shown in Figure 7.2.

As expected the S1 light calibration includes smaller number of events compared to the LED calibration. To study the effect of possible systematic effects we compare the mean and sigma of the SPE distribution obtained from the fits to S1 light to the LED parameters. The mean single photoelectron area over all channels is higher when using S1 light than the LED calculated mean by only $\sim 1\%$. The resolution of the single photoelectron distribution as

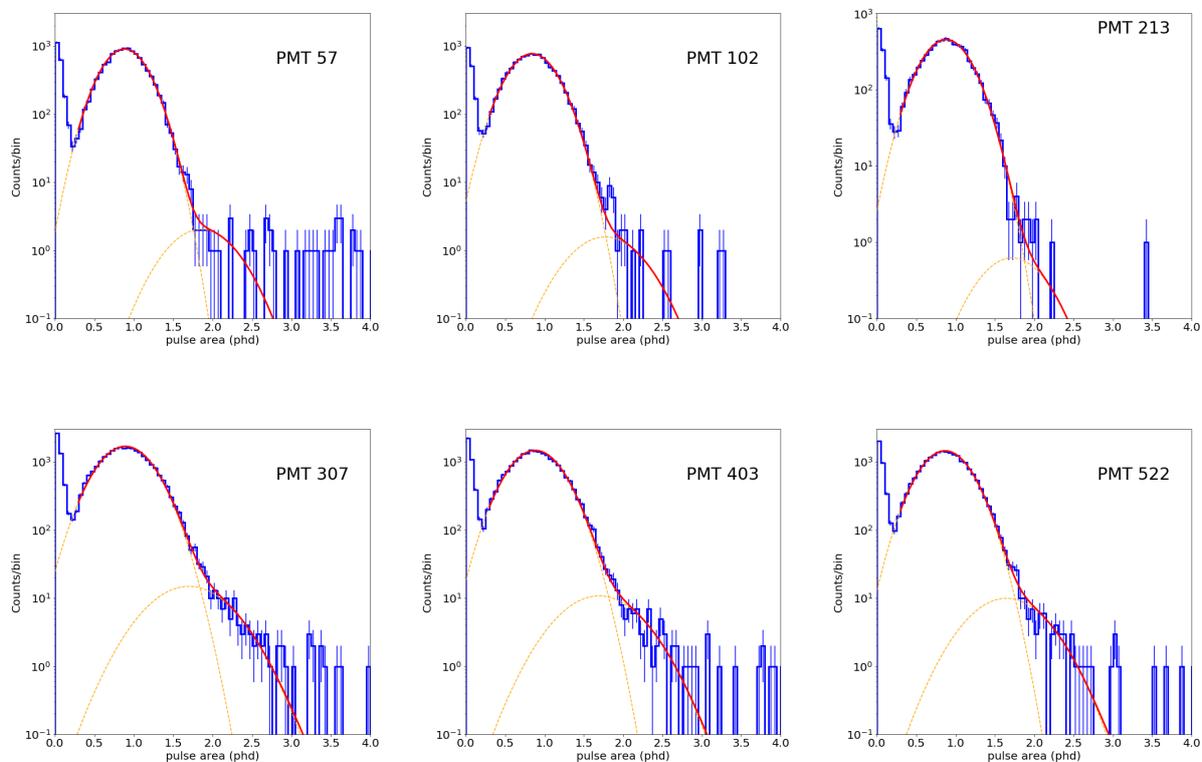


Figure 7.1: Distributions and fits of the single photoelectron response for a set of example PMTs, obtained using the LED MDC3 dataset. The distribution includes a two-photon (and multiple) photon hit component guided by the Poisson counting method discussed in the text.

calculated using S1 light is on average $\sim 2\%$ smaller than the LED calculated. Note that in reality these will differ by a bigger percentage due to the effects discussed in Chapter 4, such as direct dynode hits.

The DPE fraction for each PMT obtained using the VUV distributions is shown in Figure 7.4. The average DPE fraction was found to be 0.26. The geometric structure observed is a result of the array placement of PMTs according to their QE, as it is found to be correlated to the DPE fraction (the correlation found in Ref. [165] was implemented in the simulations). A map of the placement of PMTs along with their QE (from manufacturer data for the phototubes actually installed in the detector) is shown in Figure 7.3.

Finally, we compared our DPE fraction measurement to the MC truth, as presented in Figure 7.5. The mean ratio of our measurement to MC truth was only 1.004 overall but there is some bias in opposite directions for both arrays. This is due to the distinct cuts applied for the two arrays — for the top array less strict cuts were applied, in order to obtain enough statistics, due to its the lower S1 collection efficiency. Because of this, evaluating the efficiency

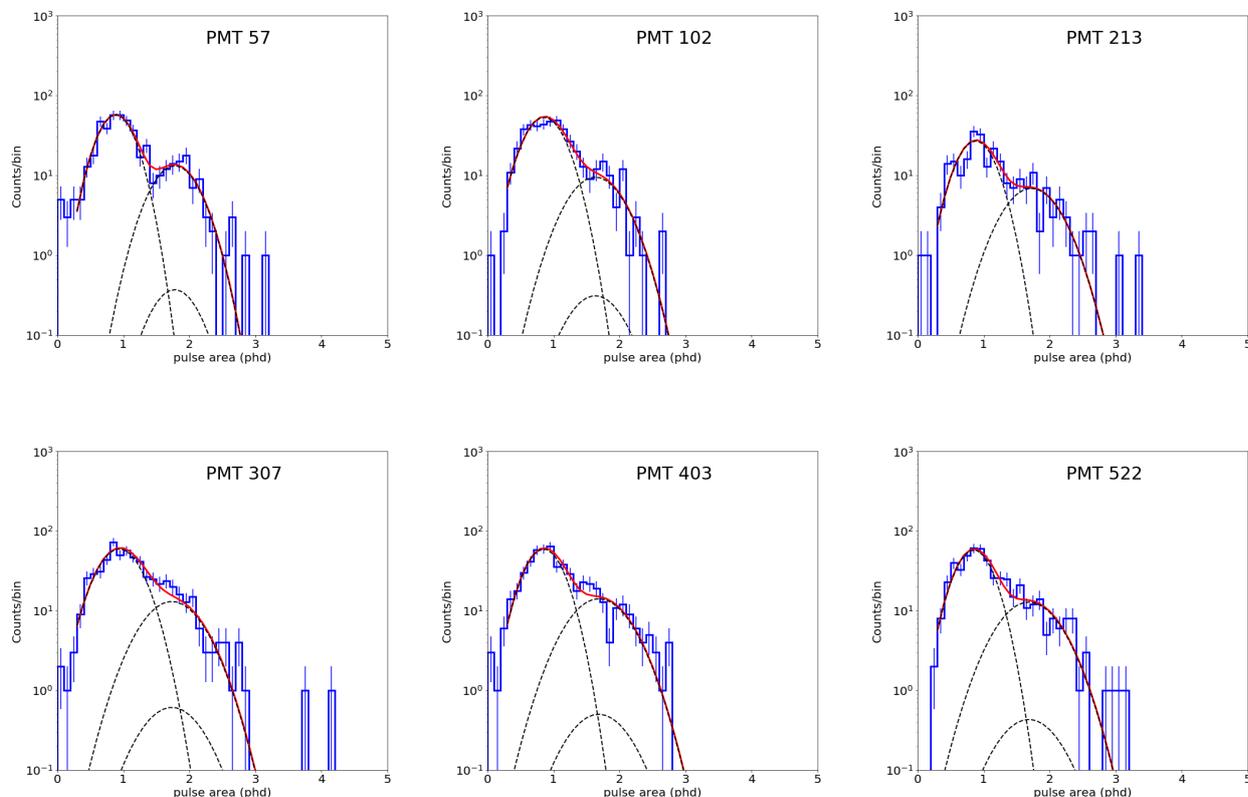


Figure 7.2: Example distributions and fits of the full VUV single photon response for a set of example PMTs, obtained using the S1 light from the commissioning datasets. The fit includes an SPE component and a DPE component from single photons and a two-photon hit (overlap) component normalised to the rate obtained in Figure 7.1.

of a DPE cut on a PMT-by-PMT basis — important for a DPE analysis — would involve large uncertainties. For the average array estimate to be sufficiently accurate to use for an early analysis, such as the one connected to a ^8B observation, the systematic array bias would need to be addressed, for which further commissioning data would be required. Dedicated calibrations such as that with tritiated methane employed in LUX would significantly improve precision due to their relatively high activity and uniform distribution. In LUX five days of tritiated methane calibration with 200 Bq total activity resulted in 170,000 events (most of which would have the appropriate S1 and S2 size for the single photon calibration in LZ). Conversely, the background datasets used here are not uniform so the number of events used for the calibration is only a few percent of the numbers used in LUX. Tritiated methane may not be among the first calibration sources to be deployed until the risk of contamination is perfectly understood in the new detector, and $^{83\text{m}}\text{Kr}$ events are too energetic for this purpose.

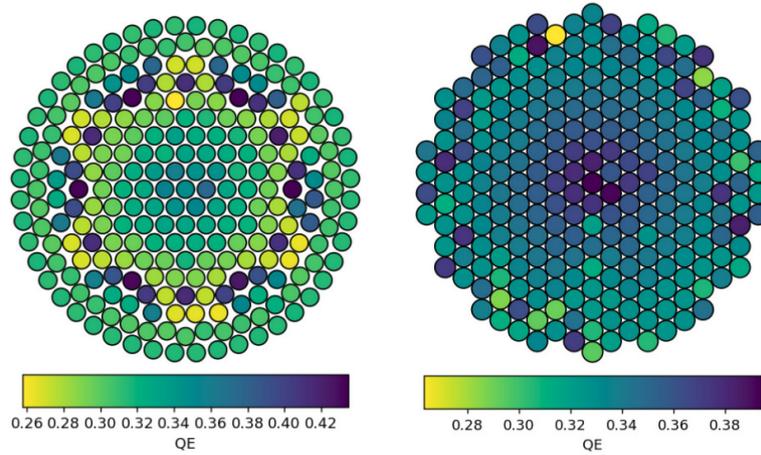


Figure 7.3: PMT map for (*left*) top and (*right*) bottom PMT arrays with the colormap representing the PMT QE.

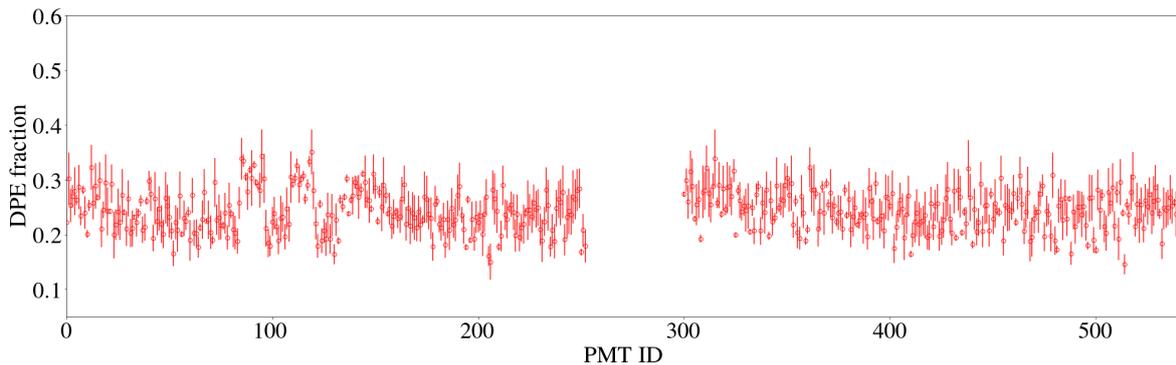


Figure 7.4: Scatter plot of DPE fraction as given through the fit of distributions obtained using S1 light for each PMT.

7.2 Estimated sensitivity improvement

Because of their increased cathode diameter, the LZ PMTs (3-inch Hamamatsu R11410-22) have a proportionally higher cold dark count rate (~ 18 c/s [167]) compared to the LUX PMTs. Moreover, LZ utilises 494 PMTs compared to 120 tubes used in LUX. The total DC rate is hence expected to produce a significant number of fake S1 signals in which two DCs coincide within the specified S1 window. These pathological "S1" signals can in turn become randomly paired with S2-only events, giving rise to a background which we will refer to as the "DCDC accidental rate" in the remainder of this chapter. Due to this pathology, enforcing a 3-fold PMT coincidence requirement for a valid S1 signal is considered essential for LZ analyses, while a 2-fold coincidence level was sufficient in LUX. Considering the S2-only rate to be of order 1 mHz in LZ, we expect 86 000, 380 and 0.17 such accidental coincidence events for a 1-fold,

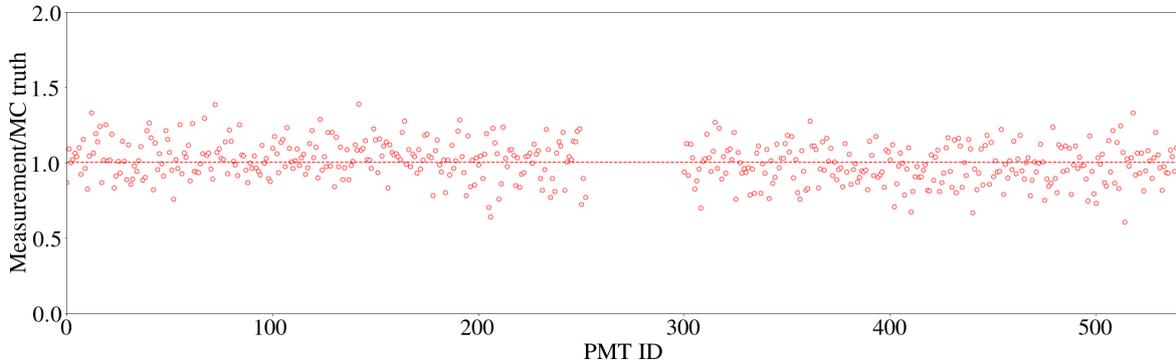


Figure 7.5: Scatter plot of the ratio of our measurement of the DPE fraction (using the simulation data described in text) with the MC truth DPE value for each PMT.

2-fold and 3-fold S1 requirement for the complete LZ exposure (1,000 live-days, 5.6 tonnes fiducial mass, 740 μ sd rift time window).

In this section we investigate the possibility of lowering the S1 threshold by exploiting the DPE response in LZ, assessing the impact on the observation of ${}^8\text{B}+{}_{\text{hep}}$ neutrinos and low mass WIMPs. This analysis would benefit from the increased DPE fraction of the LZ PMTs, with a mean of 22.6% [165] (compared with 16.9% for the LUX PMTs). At the same time, the 2-fold requirement limits the DCDC accidental rate while benefiting from an increased signal acceptance compared to the 1-fold LUX analysis, since when two photons are detected in LZ each of these has a probability of producing a DPE response. In LUX the ratio of accepted events to 1-fold S1, following the DPE cut, was 5.5%. Here the ratio of 2-fold DPE to 2-fold is $\sim 10\%$ – which comes from placing the DPE cut at 2.0 phe and adopting a standard coincidence window of 150 ns. We present a more careful optimisation of this cut for the ${}^8\text{B}$ case below. In this study we do not attempt to lower the coincidence requirement to 1-fold as we anticipate that such an analysis will have overwhelming backgrounds.

7.2.1 Detecting $\text{CE}\nu\text{NS}$ of ${}^8\text{B}$ neutrinos in LZ

In Sections 2.4.2 and 3.2 we discussed the process of $\text{CE}\nu\text{NS}$ of solar, atmospheric and diffuse supernova neutrinos, and specifically the LZ sensitivity to ${}^8\text{B}$ neutrinos from the sun. As seen in Figure 3.2, this is a very low energy NR interaction and hence extending the energy threshold of the detector, with techniques such as lowering the scintillation threshold, can result in an earlier or more robust observation of the process.

To study this possibility we use NEST v2.0.0 to convert the ${}^8\text{B}$ energy spectrum to S1 and S2 observables and investigate the potential improvement from lowering the S1 threshold from the standard 3-fold threshold to a 2-fold requirement involving DPE. The event rate of the ${}^8\text{B}$ process is shown in Figure 7.6. This is presented as a function of true recoil energy for a

0-fold, 2-fold and 3-fold S1 requirement (where the requirements are inclusive, i.e. the 3-fold requirement includes 2-fold observations), with the number of events expected to be observed by LZ predicted to be 9 246, 128 and 28, respectively. It is important to note that most of the signal detectable by LZ is due to the fluctuating nature of the S1, with only over-fluctuations falling above threshold (it should be noted that the NR detection efficiency reaches 50% at 5-6 keV in LZ). This leads to a considerable bias if energy reconstruction is to be attempted using the method presented in equation 3.21, with alternative methods such as using the S2 signal alone yielding more accurate results (see Figure 7.7). In any case, a bias is present using both methods and an improved energy reconstruction, perhaps using an equation calibrated for small S1 and S2 pulses, has to be carefully considered in the case of a discovery analysis.

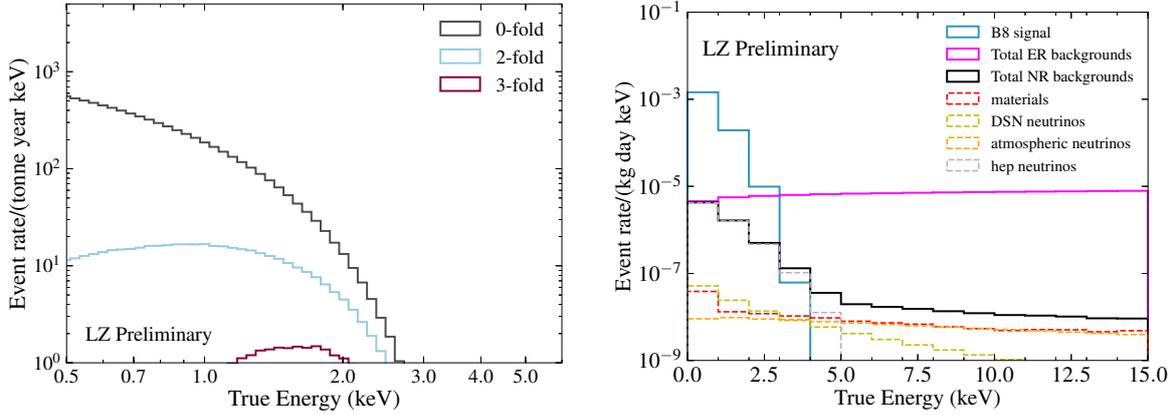


Figure 7.6: *Left:* Event rate for CEνNS of ^8B neutrinos for true recoil energy in three S1 coincidence requirement scenarios: 0-fold (grey), 2-fold (blue) and 3-fold (magenta). *Right:* LZ simulated ER and NR backgrounds for recoil energies below 13 keV along with the expected signal from ^8B neutrinos.

For our study we also need to consider both the standard ER and NR backgrounds (Figure 7.6) and the DCDC accidental rate. It is important to note that *hep* neutrinos are not distinguishable from ^8B neutrinos for such small numbers, and hence these are treated as the same signal in this study. More detail on these backgrounds can be found in Ref. [182]. The mean observations expected by LZ for a background-only experiment and the full exposure are shown in Figure 7.8, both for the 3-fold and 2-fold cases. The 2-fold scatter plot also includes the DCDC accidental rate, which is the leading source of background in this case. However, the random association background includes all two-photon S1 pulses without considering a DPE cut on their area. Before discussing and attempting to optimise the DPE cut, we define an S2 and energy region of interest, that optimises the signal to background ratio for ^8B neutrinos. The number of signal and background counts within this ROI are summarised in Table 7.1.

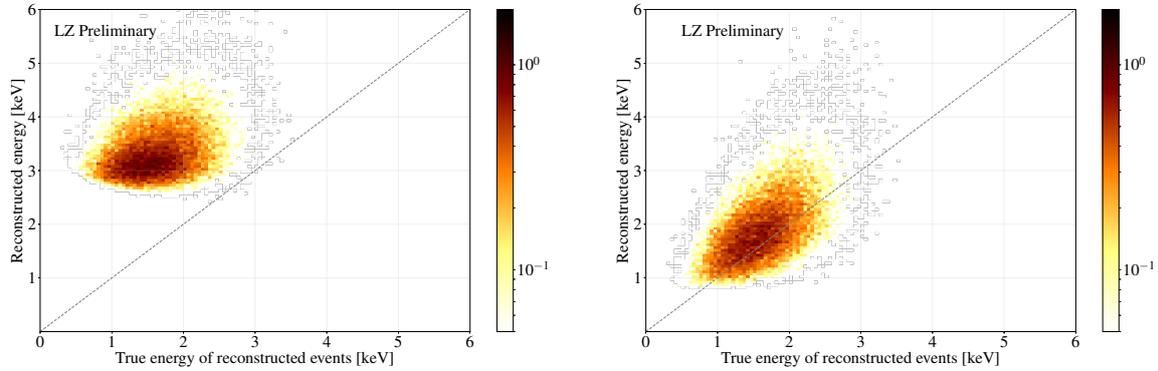


Figure 7.7: Comparison of reconstructed to true energy for CE ν NS of ^8B neutrinos using the (*left*) standard reconstruction method using both the S1 and S2 signals (equation 3.21), and (*right*) only the S2 signal and the yield model at each energy. Both methods suffer from a bias and an improved reconstruction method will be required for the actual analysis.

Table 7.1: Signal ($^8\text{B}+hep$) and background counts (ER, NR and DCDC accidentals) for a 3-fold and 2-fold S1 coincidence requirement, in the ROI, after applying S2 and energy region criteria. The 2-fold expectations, which include 3-fold counts, are presented both before and after the application of a DPE cut, placed at 1.8 phd (discussion on the optimisation of this cut follows in text).

	^8B signal	ER backgrounds	NR backgrounds	DCDC accidentals
3-fold	28	0.9	1.0	0.2
2-fold	128	1.7	2.4	191
2-fold DPE	39	1.2	1.4	3.4

Naturally, at this stage we want to establish whether the application of a DPE cut on two-photon S1 signals can result in an increased sensitivity compared to a 3-fold analysis. For the DPE cut study we enforce a cut such that when exactly two photons are detected the area of at least one of the two must be bigger than the DPE cut. The number of signal and DCDC accidental background events we expect for a range of DPE cuts are presented in Figure 7.9. As expected the number of backgrounds reduces rapidly past the mean of the DC distribution. Beyond ~ 2.3 phd, both the background and signal expectations plateau to the 3-fold expectations, as all two-photon S1 events have been removed by the DPE cut.

To understand the sensitivity effect of applying the DPE technique, the discovery frequency is calculated, defined as the number of times a 5σ or 3σ discovery of ^8B neutrinos is achieved, for variable DPE cuts and live-times. For this calculation expected signal and background event numbers at each DPE cut value (Figure 7.9) are used as the Poisson mean to draw random event numbers over 1,000 trials. At each trial we query whether the relevant confidence interval departs from zero to classify "successful" discovery. The confidence intervals for the full 2- and 3-fold populations are shown in Figure 7.10.

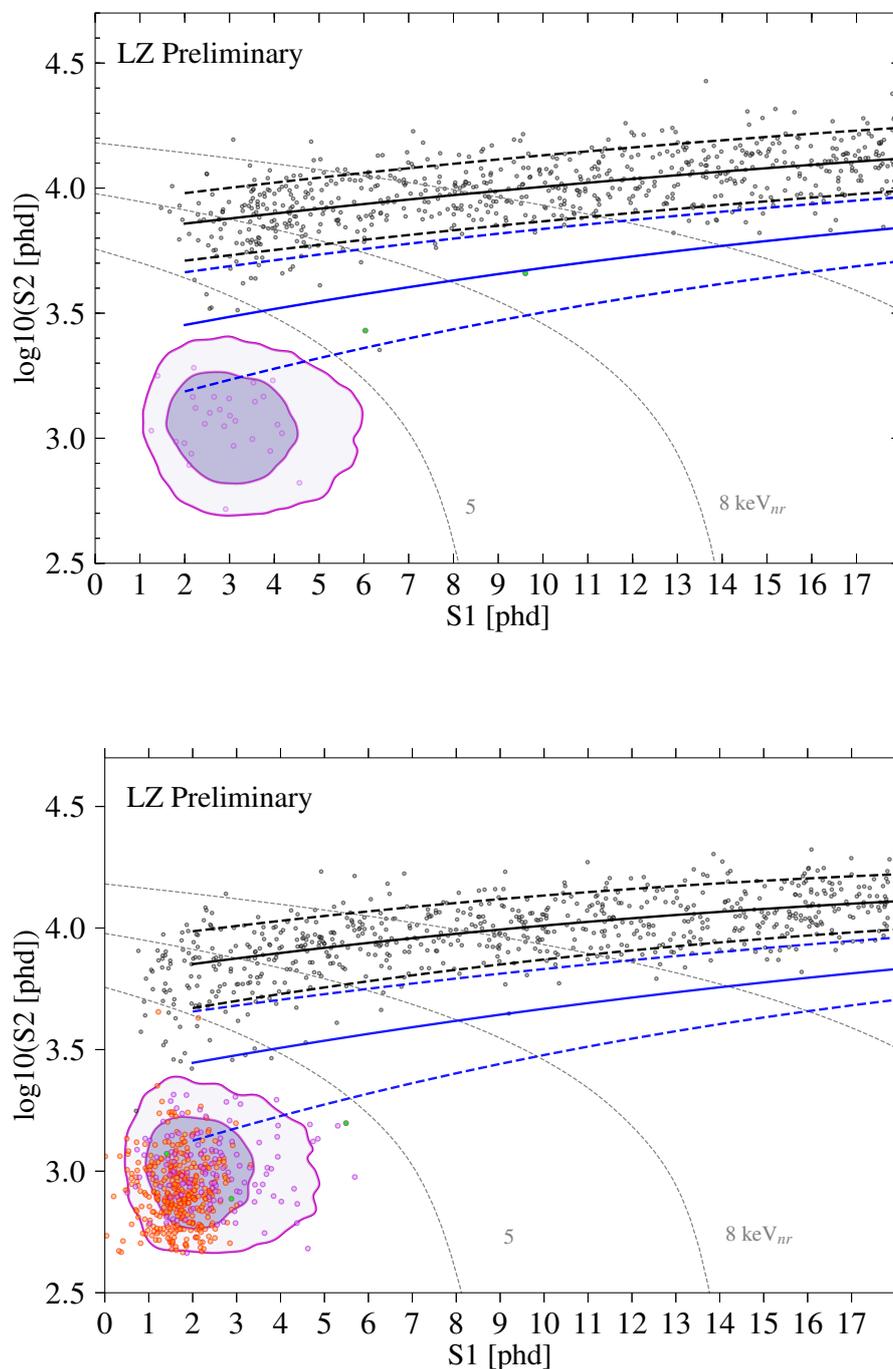


Figure 7.8: Simulated dataset for a background-only 1,000 live-day run and a 5.6 tonne fiducial mass. ER and NR bands are indicated in black and blue, respectively (solid: mean, dashed: 10% and 90%). Energy isocontour lines are shown in dashed grey. The $^8\text{B}+\text{hep}$ observation for the mean experiment is shown in purple with the 1σ and 2σ contours shown as shaded regions. Background counts due to ER and NR interactions are shown in grey and green respectively. The *upper* plot shows the dataset with a 3-fold S1 coincidence requirement applied while the *lower* plot has a 2-fold requirement applied (includes all S1 pulses above 2-fold without a DPE cut). In this case, the background expectation from DCDC accidentals is shown in orange.

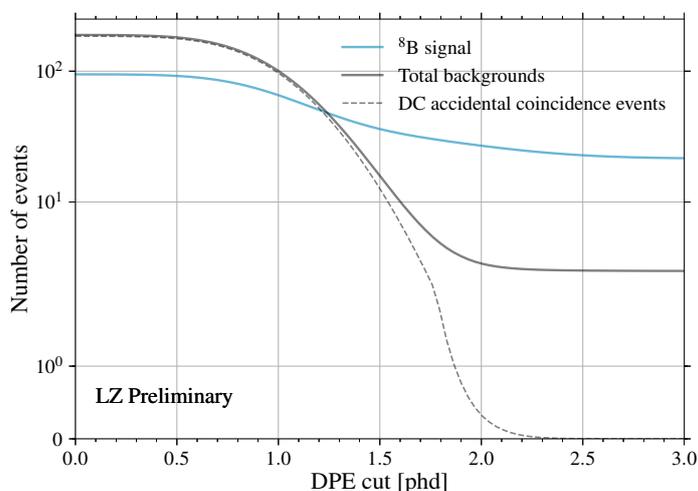


Figure 7.9: Expected ^8B signal and background counts for a variable DPE cut. The number of events presented here includes both 2- and 3-fold S1 coincidences.

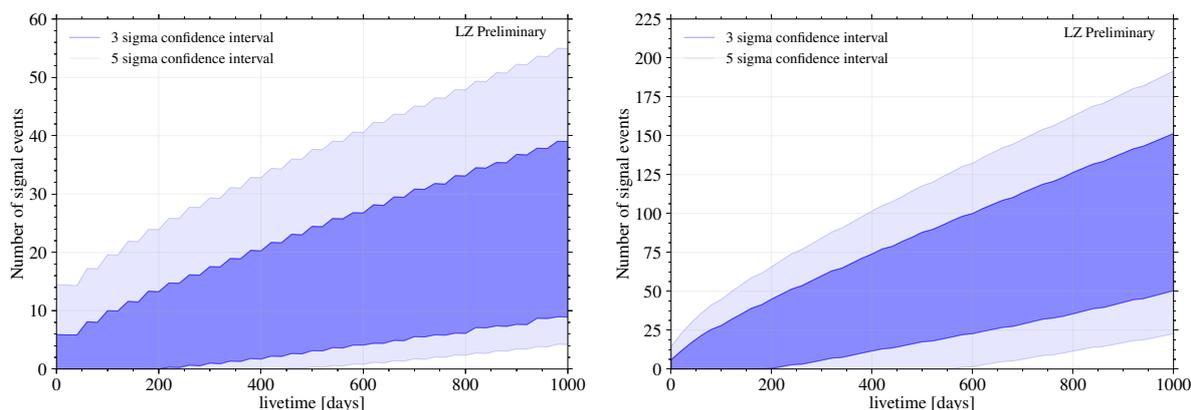


Figure 7.10: Feldman Cousins C.L. intervals for the number of signal events at varying live-times. The "measured" signal at each live-time assumes a realisation of the mean LZ experiment for a fiducial mass of 5.6 tonnes for a (left) 3-fold and (right) 2-fold S1 coincidence requirement. The 2-fold expectation includes 3-fold coincidences.

The 5σ discovery frequency for a range of DPE cuts and three example live-times is shown in Figure 7.11(left). The 5σ discovery frequency is maximised for a DPE cut of ~ 1.8 phd and is improved both compared to the full 2-fold and 3-fold cases. Figure 7.11 (right) presents the 5σ and 3σ discovery frequency as a function of live-time for a full 3-fold analysis, a full 2-fold analysis, and for a 2-fold analysis with the DPE cut that maximises the discovery frequency. The mean observation expected by LZ for this optimised DPE cut is shown in Figure 7.12. The number of events that appear in this plot are also summarised in Table 7.1. The number of DCDC accidental coincidences is reduced from 191 to 3.6 following the DPE cut and the 2-fold DPE case shows a signal improvement factor of 1.4 compared to the 3-fold analysis. For shorter runs (< 400 days), lowering the S1 threshold to 2-fold and using a DPE cut results

in an improved discovery frequency. Eventually, a discovery is expected for any of the three scenarios as enough ^8B events are observed, which is why the three cases are seen to plateau for the 3σ case. Hence, using a lower threshold with a DPE cut is of particular interest if an early observation is desired.

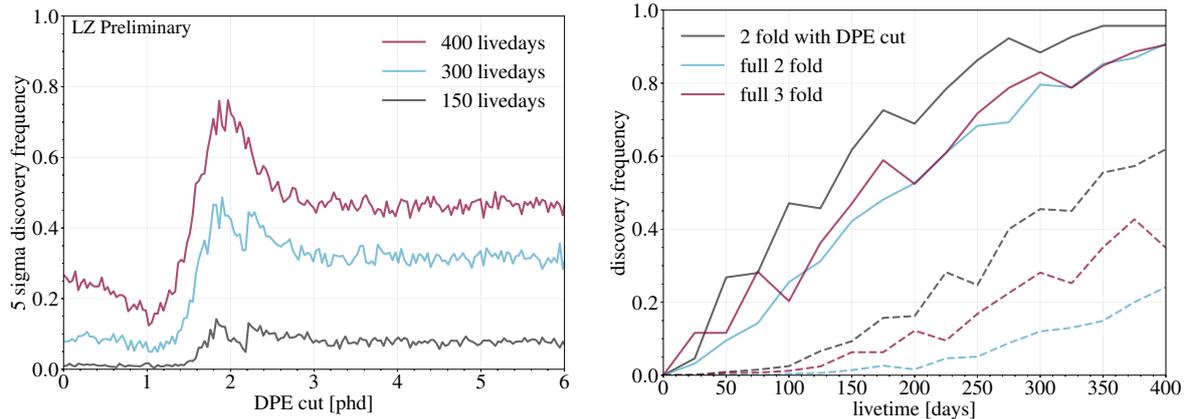


Figure 7.11: *Left:* 5σ discovery frequency for a variable DPE cut and three example live-times of 150, 300, 400 days. *Right:* 3σ (solid) and 5σ (dashed) discovery frequency for a variable live-time and three coincidence requirement scenarios: 3-fold analysis, full 2-fold analysis without a DPE cut and a 2-fold analysis for which the DPE cut is chosen such that it maximises the discovery frequency shown in the *left* panel.

7.2.2 Considering low mass WIMPs

The ^8B study shares various similarities to the search for lower mass WIMPs even though a number of important differences exist. For WIMP masses lower than the ^8B equivalent (which is similar to a 5.5 GeV WIMP) and hence lower recoil energies, a larger proportion of the signal falls below the 3-fold requirement. This becomes obvious in Figure 7.13 (*right*) which illustrates the expected increase in signal counts when lowering the coincidence requirement from 3-fold to 2-fold for a range of low WIMP masses (this is before any application of a DPE cut). For ^8B the increase in signal counts from 3-fold to 2-fold is ~ 4.5 while for a 3 GeV WIMP, the improvement is a factor ~ 18 . The event rate for a 3 GeV WIMP is presented in Figure 7.13 (*left*).

Even though the signal rate improvement increases for smaller WIMP masses, so does the background rate. The DCDC accidental rate itself does not increase for lower mass WIMPs but the lower threshold comes with an increase in the number of $^8\text{B}+\text{hep}$ neutrinos, which pose the real background to this measurement. Despite the real background to this measurement, the signal to background ratio increases when lowering the scintillation threshold, even before applying a DPE cut. To accurately estimate the sensitivity improvement a full PLR analysis

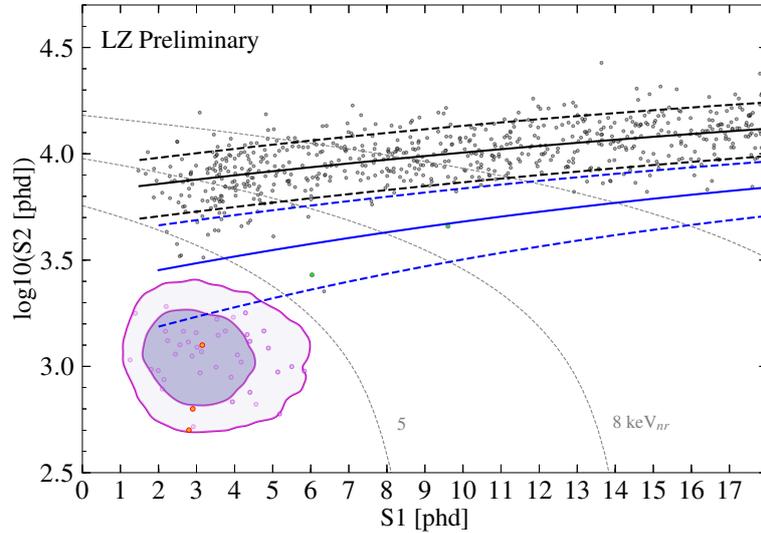


Figure 7.12: Simulated dataset for a background-only 1,000 live-day run and a 5.6 tonne fiducial mass. ER and NR bands are indicated in black and blue, respectively (solid: mean, dashed: 10% and 90%). Energy isocontour lines are shown in dashed grey. The ${}^8\text{B}+\text{hep}$ observation for the mean experiment is shown in purple with the 1σ and 2σ contours shown as shaded regions. Background counts due to ER and NR interactions are shown in grey and green respectively. The plot shows the dataset with the coincidence requirement lowered to a 2-fold for S1 pulses indicative of DPE emission. The background expectation from DCDC accidentals is shown in orange, following the DPE cut.

that considers the effect of scintillation and ionisation yield uncertainties and the uncertainty introduced by the assumed S2-only spectrum (which was scaled from LUX) is needed.

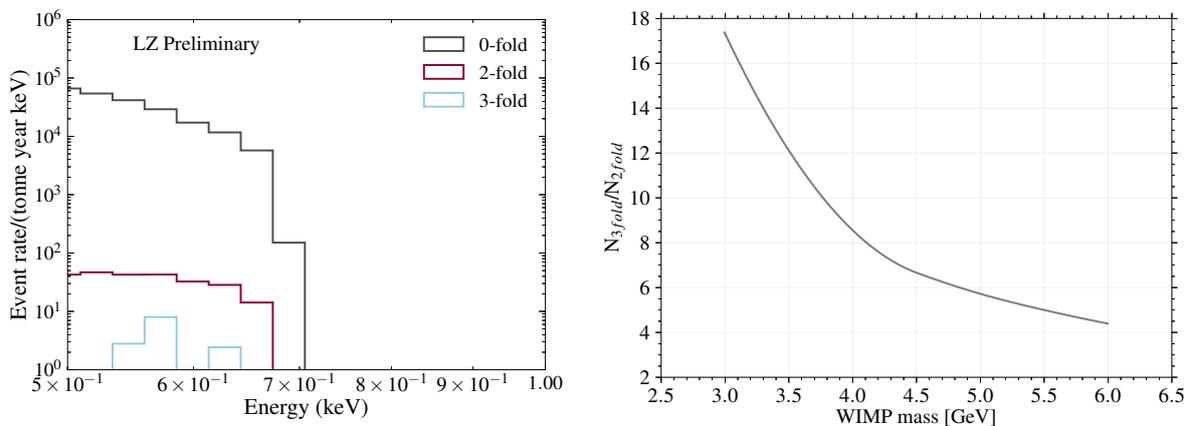


Figure 7.13: *Left:* Event rate for a 3 GeV WIMP at an exposure of $2.3 \times 10^{-42} \text{ cm}^2$ illustrated in true recoil energy following a 0-fold (grey), 2-fold (magenta) and 3-fold (blue) S1 coincidence requirement. *Right:* Increase in signal counts for a range of WIMP masses when lowering the S1 coincidence requirement from 3-fold to 2-fold, prior to the application of a DPE cut.

7.3 Result of a full analysis with a lower threshold

The work presented in the previous sections served as an initial motivation for a complete LZ study on this topic performed by the Imperial team, which we summarise here for completeness. The projected spin-independent sensitivity for a standard 3-fold analysis and an analysis using a 2-fold threshold including DPE events is presented in Figure 7.14. The LZ DPE analysis represents a factor 3 improvement at 2.5 GeV compared to the standard LZ sensitivity. As shown in the lower panel of Figure 7.14, the signal improvement benefits from a factor ~ 5 increase, but the final sensitivity increase is smaller due to the increasing number of ^8B neutrinos as previously discussed. The ratio of signal efficiency between the DPE and standard LZ analysis is specifically highlighted (in the lower panel of the figure) for ^8B neutrinos to be 1.7. If we assume a DPE cut at ~ 1.8 phd in our preliminary study in Section 7.2.1 we find a similar (but slightly smaller) improvement factor relative to the complete analysis result. The complete LZ analysis includes more sophisticated methods of optimising the DPE cut, both in area and photon detection timing, and makes use of a full PLR method to reach the final sensitivity. The complete analysis also assesses systematic effects due to yield and S2-only rate uncertainties, and the effect of the electron lifetime.

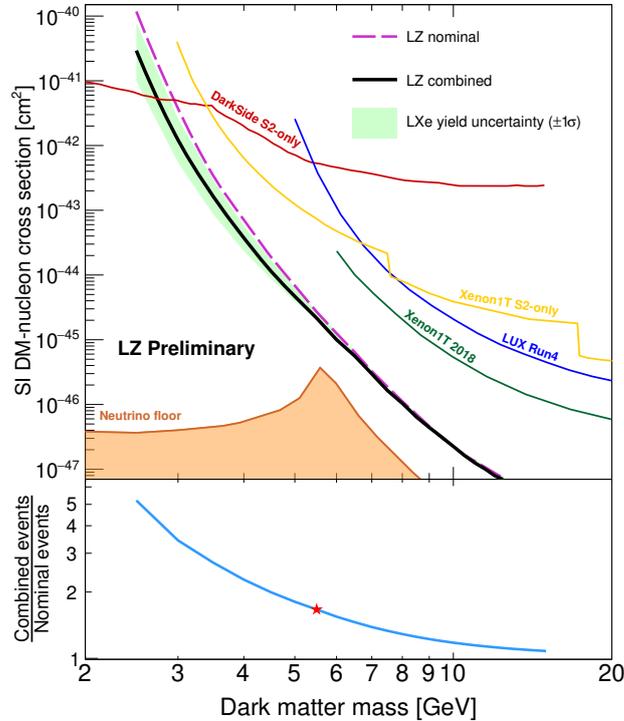


Figure 7.14: Improvement in standard spin-independent LZ WIMP sensitivity with the use of a combined DPE and standard analysis for a lower 2-fold S1 coincidence threshold. The lower panel presents the improvement in signal rate with the ^8B case marked with a star.

7.4 Summary

In this chapter we presented a study on the potential sensitivity improvement that LZ can benefit from when lowering the scintillation threshold from a 3-fold to a 2-fold level with the use of a DPE cut to remove DCDC accidental events. Specifically, we focused on the CE ν NS of $^8\text{B}+hep$ neutrinos and studied the discovery frequency for a range of live-times and DPE cut placements. Lowering the S1 threshold was found to improve the discovery frequency when a DPE cut was used, especially for short live-times and therefore it is clearly worth performing this analysis on the early datasets. The application of this technique to WIMP signals at low masses was also discussed and a full LZ projected sensitivity study, motivated by this work, was briefly presented. This type of analysis depends strongly on our ability to accurately calibrate the single VUV photon response of the LZ PMTs. We demonstrated also that a preliminary calibration of the DPE effect can be conducted on a limited background dataset, but ideally a tritium calibration would be available early on for a more rigorous determination of the PMT response functions.

Chapter 8

Conclusions and outlook

In this thesis we presented a new data analysis technique to search for rare electron and nuclear recoil interactions at sub-keV energies in LXe-TPCs, which we demonstrated using LUX data. Our analysis was based on the efficient detection of single VUV photons that occasionally generate two photoelectrons in some photomultiplier tube models.

In Chapter 4 we focused on understanding the single photon response of the LUX photomultiplier tubes, using data from a CH₃T injection. This involved calibrating the VUV response and studying the dark count response of the LUX PMTs, which determines the main background topology in the analysis presented in this thesis. Although the DPE response was found to be a modest fraction of the total response to single VUV photons, there is essentially a very small dark count rate competing with such signals and low backgrounds can therefore be achieved — this is a key conclusion of this study.

In Chapter 5 we proceeded by benchmarking the proposed technique using the CH₃T ER calibration dataset. We successfully lowered the threshold by accurately recovering events where the S1 pulse consisted of a single detected photon that would have previously been discarded. No significant drawbacks or unexpected backgrounds were discovered and hence we propose that this technique could be used to lower the energy threshold of electron recoil analyses, both for calibration purposes and for dark matter candidates that appear as ER signals, such as those due to Migdal and nuclear Bremsstrahlung effects.

We then applied a similar methodology to a search for low-energy nuclear recoils in the LUX 2013 WIMP search dataset, as presented in Chapter 6. The single-photon S1 analysis led to an improved spin-independent scattering cross-section limit, especially between a 2.5 GeV and 5 GeV WIMP mass compared to the standard 2-fold result. This technique can be used to extend the reach of other rare NR event searches using LXe-TPCs, such as those searching for asymmetric dark matter scattering or improving the sensitivity to neutrino processes that are expected to produce measurable rates in future experiments.

Various groups around the world are pursuing the measurement of scintillation and ionisation yields for ER and NR interactions and to establish the energy required to release the first quantum of ionisation and of scintillation in liquid xenon. The technique proposed in this thesis depends on these yield measurements and the effect of the model assumptions we adopt were discussed both in Chapters 5 and 6.

Finally, there are good prospects for applying this analysis to larger experiments such as LZ, where the improvement can be more significant owing to several factors. In Chapter 7 we find that this could be an important enhancement in the search for very light WIMP interactions and the coherent nuclear scattering of solar ^8B neutrinos. This is more pronounced for early searches, where using the DPE effect to lower the S1 threshold can result in a more rapid ^8B discovery.

Appendix A

Confidence intervals: Rolke's method

The method of maximum likelihood is widely used in the particle physics community [183]. Assuming a probability model with parameters $\theta = (\theta_1, \dots, \theta_l)$, a probability density function $f(x|\theta)$, and independent observations $\mathbf{X} = (X_1, \dots, X_n)$, the likelihood function is given by:

$$\mathcal{L}(\theta|\mathbf{X}) = \prod_{i=1}^n f(X_i|\theta). \quad (\text{A.1})$$

The maximum likelihood estimators for the parameters are then those that maximise the likelihood function:

$$\frac{\partial \mathcal{L}}{\partial \theta_i} = 0, \quad i = 1, \dots, l. \quad (\text{A.2})$$

In this analysis we set confidence intervals for the signal rate in the presence of an uncertain background and efficiency, and we adopted the method by Rolke [176] for its simplicity of implementation. If we denote μ as the signal count, ϵ as the efficiency with uncertainty given by σ_ϵ , and b as the background rate with uncertainty σ_b and assume that both the background and efficiency can be modelled as Gaussians (N), our probability model becomes:

$$X \sim P(\epsilon\mu + b), \quad Y \sim N(b, \sigma_b), \quad Z \sim N(\epsilon, \sigma_\epsilon), \quad (\text{A.3})$$

where P denotes the Poisson distribution. The likelihood function is then given by:

$$\mathcal{L}(\mu, b, \epsilon|x, y, z) = f(x, y, z|\mu, b, \epsilon) = P(\epsilon\mu + b) \times N(b, \sigma_b) \times N(\epsilon, \sigma_\epsilon) \quad . \quad (\text{A.4})$$

In this case, the derivatives of the likelihood function (and hence also its logarithm), assuming a known signal rate, μ_0 , provide the maximum likelihood estimators, \hat{b} and $\hat{\epsilon}$:

$$\frac{\partial}{\partial b} \log \mathcal{L}(\mu, b, \epsilon|x, y, z) = \frac{x}{\hat{\epsilon}\mu + \hat{b}} - 1 + \frac{(y - \hat{b})}{\sigma_b} = 0, \quad (\text{A.5})$$

$$(\text{A.6})$$

$$\frac{\partial}{\partial \epsilon} \log \mathcal{L}(\mu, b, \epsilon | x, y, z) = \frac{x}{\hat{\epsilon}\mu + \hat{b}} - \mu + \frac{(z - \hat{\epsilon})}{\sigma_\epsilon} = 0.$$

Hypothesis testing can then be used to construct confidence intervals using the likelihood ratio test statistic, which represents the ratio between the best explanation for the data if the hypothesis is true and if no assumption is made. This ratio is called the profile likelihood and in the context of the assumptions made here is given by:

$$\lambda(\mu | x, y, z) = \frac{\mathcal{L}(\mu, \hat{b}, \hat{\epsilon} | x, y, z)}{\mathcal{L}(\hat{b}, \hat{\epsilon} | x, y, z)}. \quad (\text{A.7})$$

According to Wilks' theorem, $-2 \log \lambda$ follows a χ^2 distribution in the asymptotic limit of large number of counts, which can be used to extract limits. To find a $(1 - \alpha)100\%$ confidence interval the profile likelihood function is scanned, starting at its minimum (maximum likelihood estimators) and moving to the right and left until the function rises to the α percentile.

Special cases to this formalism exist and are of particular interest here. In the case of a higher background expectation compared to the observed rate, an acceptance region is created based on the null hypothesis and the values of μ for which the profile likelihood enters and leaves the origin. In this case the log profile likelihood curve is no longer parabolic [176].

A coverage study for an example 4 GeV WIMP at 90% confidence level is shown in Figure A.1. Coverage was calculated for a given set of true physics parameters (b, μ) and repeated trials with measured count varied each time. The frequency for which the interval covers the signal level was then checked (it should fail on the upper side 10% of the time). The method discussed is generally found to have good coverage with the small undercoverage at small signals seen as acceptable (and discussed in detail in Ref. [176]).

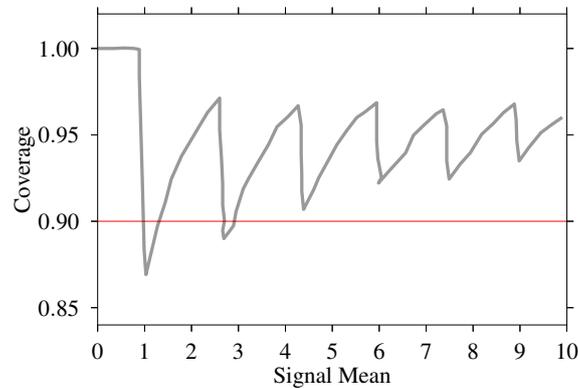


Figure A.1: Coverage for a 90% C.L. upper limit using the Rolke method for a 4 GeV WIMP.

Bibliography

- [1] D. Akerib et al., “Extending light WIMP searches to single scintillation photons in LUX”, *Physical Review D* **101** (2020), no. 4, 042001.
- [2] J. H. Oort et al., “The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems”, *Bulletin of the Astronomical Institutes of the Netherlands* **6** (1932) 249.
- [3] F. Zwicky, “Die rotverschiebung von extragalaktischen nebeln”, *Helvetica Physica Acta* **6** (1933) 110–127.
- [4] A. B. Newman et al., “The density profiles of massive, relaxed galaxy clusters. I. The total density over three decades in radius”, *The Astrophysical Journal* **765** (2013), no. 1, 24.
- [5] N. A. Bahcall et al., “A lightweight universe?”, *Proceedings of the National Academy of Sciences* **95** (1998), no. 11, 5956–5959.
- [6] V. C. Rubin et al., “Rotation of the Andromeda nebula from a spectroscopic survey of emission regions”, *The Astrophysical Journal* **159** (1970) 379.
- [7] J. H. Oort et al., “Some Problems Concerning the Structure and Dynamics of the Galactic System and the Elliptical Nebulae NGC 3115 and 4494.”, *The Astrophysical Journal* **91** (1940) 273–306.
- [8] K. Begeman et al., “Extended rotation curves of spiral galaxies: Dark haloes and modified dynamics”, *Monthly Notices of the Royal Astronomical Society* **249** (1991), no. 3, 523–537.
- [9] Y. Sofue et al., “Unified Rotation Curve of the Galaxy—Decomposition into de Vaucouleurs Bulge, Disk, Dark Halo, and the 9-kpc Rotation Dip—”, *Publications of the Astronomical Society of Japan* **61** (2009), no. 2, 227–236.
- [10] M. Reid et al., “Trigonometric parallaxes of high mass star forming regions: the structure and kinematics of the Milky Way”, *The Astrophysical Journal* **783** (2014), no. 2, 130.

- [11] P. Mróz et al., “Rotation curve of the Milky Way from classical Cepheids”, *The Astrophysical Journal* **870** (2019), no. 1, L10.
- [12] R. Massey et al., “The dark matter of gravitational lensing”, *Reports on Progress in Physics* **73** (2010), no. 8, 086901.
- [13] P. Schneider, “Detection of (dark) matter concentrations via weak gravitational lensing”, *Monthly Notices of the Royal Astronomical Society* **283** (1996), no. 3, 837–853.
- [14] D. Wittman et al., “Detection of weak gravitational lensing distortions of distant galaxies by cosmic dark matter at large scales”, *Nature* **405** (2000), no. 6783, 143–148.
- [15] H. Shan et al., “Offset between dark matter and ordinary matter: evidence from a sample of 38 lensing clusters of galaxies”, *Monthly Notices of the Royal Astronomical Society* **406** (2010), no. 2, 1134–1139.
- [16] R. B. Metcalf et al., “Compound gravitational lensing as a probe of dark matter substructure within galaxy halos”, *The Astrophysical Journal* **563** (2001), no. 1, 9.
- [17] N. A. Bahcall et al., “Where is the dark matter?”, *The Astrophysical Journal* **447** (1995), no. 2, L81.
- [18] D. Clowe et al., “A direct empirical proof of the existence of dark matter”, *The Astrophysical Journal* **648** (2006), no. 2, L109.
- [19] J. Brownstein et al., “The Bullet Cluster 1E0657-558 evidence shows modified gravity in the absence of dark matter”, *Monthly Notices of the Royal Astronomical Society* **382** (2007), no. 1, 29–47.
- [20] R. Durrer, “The cosmic microwave background”, volume 401. Cambridge University Press Cambridge, 2008.
- [21] A. A. Penzias et al., “A measurement of excess antenna temperature at 4080 Mc/s.”, *The Astrophysical Journal* **142** (1965) 419–421.
- [22] N. Aghanim et al., “Planck 2015 results-XI. CMB power spectra, likelihoods, and robustness of parameters”, *Astronomy & Astrophysics* **594** (2016) A11.
- [23] C. Bennett, “First year WMAP observations”, in *Symposium-International Astronomical Union*, volume 216, pp. 18–27, Cambridge University Press. 2005.
- [24] G. Bertone et al., “Particle dark matter: Evidence, candidates and constraints”, *Physics reports* **405** (2005), no. 5-6, 279–390.
- [25] B. Bassett and R. Hlozek, “Baryon acoustic oscillations”, *Dark energy: observational and theoretical approaches* (2010) 246.

- [26] S. Dodelson et al., “Can cosmic structure form without dark matter?”, *Physical Review Letters* **97** (2006), no. 23, 231301.
- [27] C. S. Frenk et al., “Dark matter and cosmic structure”, *Annalen der Physik* **524** (2012), no. 9-10, 507–534.
- [28] K. Lind et al., “The lithium isotopic ratio in very metal-poor stars”, *Astronomy & Astrophysics* **554** (2013) A96.
- [29] J. Bernstein et al., “Cosmological helium production simplified”, *Reviews of Modern Physics* **61** (1989), no. 1, 25.
- [30] K. Jedamzik et al., “Big Bang nucleosynthesis and particle dark matter”, *arXiv:0906.2087* (2009).
- [31] X. Xue, H. Rix, G. Zhao et al., “The Milky Way’s Circular Velocity Curve to 60 kpc and an Estimate of the Dark Matter Halo Mass from the Kinematics of ~ 2400 SDSS Blue Horizontal-Branch Stars”, *The Astrophysical Journal* **684** (2008), no. 2, 1143.
- [32] V. N. Lukash et al., “Formation of the large-scale structure of the Universe”, *Physics-Uspokhi* **54** (2011), no. 10, 983.
- [33] P. Dev et al., “Constraining non-thermal and thermal properties of dark matter”, *Frontiers in Physics* **2** (2014) 26.
- [34] J. Alexander et al., “Dark sectors 2016 workshop: community report”, *arXiv:1608.08632* (2016).
- [35] T. M. Undagoitia et al., “Dark matter direct-detection experiments”, *Journal of Physics G: Nuclear and Particle Physics* **43** (2015), no. 1, 013001.
- [36] A. Boveia et al., “Dark matter searches at colliders”, *Annual Review of Nuclear and Particle Science* **68** (2018) 429–459.
- [37] J. M. Gaskins, “A review of indirect searches for particle dark matter”, *Contemporary Physics* **57** (2016), no. 4, 496–525.
- [38] M. W. Goodman and E. Witten, “Detectability of certain dark-matter candidates”, *Physical Review D* **31** (1985), no. 12, 3059.
- [39] D. Akerib et al., “The Large Underground Xenon (LUX) experiment”, *Nuclear Instruments and Methods in Physics Research Section A* **704** (2013), no. Supplement C, 111 – 126, doi:10.1016/j.nima.2012.11.135.
- [40] X. Cao et al., “PandaX: a liquid xenon dark matter experiment at CJPL”, *Science China Physics, Mechanics & Astronomy* **57** (2014), no. 8, 1476–1494.

- [41] E. Aprile, J. Aalbers, F. Agostini et al., “The XENON1T dark matter experiment”, *The European Physical Journal C* **77** (2017), no. 12, 881.
- [42] L. Roszkowski et al., “WIMP dark matter candidates and searches—current status and future prospects”, *Reports on Progress in Physics* **81** (2018), no. 6, 066201.
- [43] K. Petraki et al., “Review of asymmetric dark matter”, *International Journal of Modern Physics A* **28** (2013), no. 19, 1330028.
- [44] D. E. Kaplan et al., “Asymmetric dark matter”, *Phys. Rev. D* **79** (2009), no. 11, 115016.
- [45] B. W. Lee et al., “Cosmological lower bound on heavy-neutrino masses”, *Physical Review Letters* **39** (1977), no. 4, 165.
- [46] C. Boehm et al., “Scalar dark matter candidates”, *Nuclear Physics B* **683** (2004), no. 1-2, 219–263.
- [47] D. Hooper et al., “Natural supersymmetric model with MeV dark matter”, *Physical Review D* **77** (2008), no. 8, 087302.
- [48] J. L. Feng et al., “Dark-matter particles without weak-scale masses or weak interactions”, *Physical Review Letters* **101** (2008), no. 23, 231301.
- [49] C. Boehm et al., “Light and heavy dark matter particles”, *Physical Review D* **69** (2004), no. 10, 101302.
- [50] J. Silk et al., “Cosmic-ray antiprotons as a probe of a photino-dominated universe”, *Physical Review Letters* **53** (1984), no. 6, 624.
- [51] J. Knödseder et al., “Early SPI/INTEGRAL constraints on the morphology of the 511 keV line emission in the 4th galactic quadrant”, *Astronomy & Astrophysics* **411** (2003), no. 1, L457–L460.
- [52] C. Boehm et al., “Can Annihilating dark matter be lighter than a few GeVs?”, *Journal of Physics G: Nuclear and Particle Physics* **30** (2004), no. 3, 279.
- [53] R. Lingenfelter et al., “Is there a dark matter signal in the galactic positron annihilation radiation?”, *Physical Review Letters* **103** (2009), no. 3, 031301.
- [54] L. J. Hall et al., “Freeze-in production of FIMP dark matter”, *Journal of High Energy Physics* **2010** (2010), no. 3, 80.
- [55] N. Bernal et al., “The dawn of FIMP dark matter: a review of models and constraints”, *International Journal of Modern Physics A* **32** (2017), no. 27, 1730023.

- [56] Y. Hochberg et al., “Mechanism for thermal relic dark matter of strongly interacting massive particles”, *Physical Review Letters* **113** (2014), no. 17, 171301.
- [57] Y. Hochberg et al., “Model for thermal relic dark matter of strongly interacting massive particles”, *Physical Review Letters* **115** (2015), no. 2, 021301.
- [58] E. Kuflik et al., “Elastically decoupling dark matter”, *Physical Review Letters* **116** (2016), no. 22, 221302.
- [59] E. Kuflik et al., “Phenomenology of ELDER dark matter”, *Journal of High Energy Physics* **2017** (2017), no. 8, 78.
- [60] R. Foot, “Mirror matter-type dark matter”, *International Journal of Modern Physics D* **13** (2004), no. 10, 2161–2192.
- [61] R. Foot, “Mirror dark matter: Cosmology, galaxy structure and direct detection”, *International Journal of Modern Physics A* **29** (2014), no. 11n12, 1430013.
- [62] D. Akerib et al., “First direct detection constraint on mirror dark matter kinetic mixing using LUX 2013 data”, *Physical Review D* **101** (2020), no. 1, 012003.
- [63] G. McLaughlin, “Theory and phenomenology of coherent neutrino-nucleus scattering”, in *AIP Conference Proceedings*, volume 1666, p. 160001, AIP Publishing LLC. 2015.
- [64] LZ Collaboration, “LUX-ZEPLIN (LZ) Technical Design Report”, *arXiv:1703.09144* (2017).
- [65] P. Barbeau et al., “Toward coherent neutrino detection using low-background micropattern gas detectors”, *IEEE Transactions on Nuclear Science* **50** (2003), no. 5, 1285–1289.
- [66] C. Hagmann et al., “Two-phase emission detector for measuring coherent neutrino-nucleus scattering”, *IEEE transactions on Nuclear Science* **51** (2004), no. 5, 2151–2155.
- [67] D. Akimov, J. Albert, P. An et al., “Observation of coherent elastic neutrino-nucleus scattering”, *Science* **357** (2017), no. 6356, 1123–1126.
- [68] LZ Collaboration, “LUX-ZEPLIN (LZ) Conceptual Design Report”, *arXiv:1509.02910* (2015).
- [69] R. F. Lang, C. McCabe, S. Reichard et al., “Supernova neutrino physics with xenon dark matter detectors: A timely perspective”, *Physical Review D* **94** (2016), no. 10, 103009.
- [70] C. Lunardini et al., “Probing the neutrino mass hierarchy and the 13-mixing with supernovae”, *Journal of Cosmology and Astroparticle Physics* **2003** (2003), no. 06, 009.

- [71] H. Duan et al., "Collective neutrino flavor transformation in supernovae", *Physical Review D* **74** (2006), no. 12, 123004.
- [72] Y. Pehlivan et al., "Invariants of collective neutrino oscillations", *Physical Review D* **84** (2011), no. 6, 065008.
- [73] A. Anderson et al., "Measuring active-to-sterile neutrino oscillations with neutral current coherent neutrino-nucleus scattering", *Physical Review D* **86** (2012), no. 1, 013004.
- [74] J. A. Formaggio et al., "Sterile neutrinos, coherent scattering, and oscillometry measurements with low-temperature bolometers", *Physical Review D* **85** (2012), no. 1, 013009.
- [75] J. Billard et al., "Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments", *Physical Review D* **89** (2014), no. 2, 023524.
- [76] C. Giunti et al., "Fundamentals of neutrino physics and astrophysics". Oxford university press, 2007.
- [77] C. McCabe, "Astrophysical uncertainties of dark matter direct detection experiments", *Physical Review D* **82** (2010), no. 2, 023530.
- [78] N. W. Evans et al., "Refinement of the standard halo model for dark matter searches in light of the Gaia Sausage", *Physical Review D* **99** (2019), no. 2, 023012.
- [79] L. Necib et al., "Dark Matter in Disequilibrium: The Local Velocity Distribution from SDSS-Gaia", *arXiv:1807.02519* (2018).
- [80] T. Piffl et al., "The RAVE survey: the Galactic escape speed and the mass of the Milky Way", *Astronomy & Astrophysics* **562** (2014) A91.
- [81] Y. Sofue, "Grand Rotation Curve and Dark-Matter Halo in the Milky Way Galaxy", *Publications of the Astronomical Society of Japan* **64** (2012), no. 4, 75.
- [82] R. Catena et al., "A novel determination of the local dark matter density", *Journal of Cosmology and Astroparticle Physics* **2010** (2010), no. 08, 004.
- [83] J. Bovy et al., "On the local dark matter density", *The Astrophysical Journal* **756** (2012), no. 1, 89.
- [84] M. Schumann, "Direct detection of WIMP dark matter: concepts and status", *Journal of Physics G: Nuclear and Particle Physics* **46** (2019), no. 10, 103003.
- [85] R. H. Helm, "Inelastic and elastic scattering of 187-MeV electrons from selected even-even nuclei", *Physical Review* **104** (1956), no. 5, 1466.

- [86] G. Jungman et al., “Supersymmetric dark matter”, *Physics Reports* **267** (1996), no. 5-6, 195–373.
- [87] J. Lewin et al., “Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil”, technical report, SCAN-9603159, (1996).
- [88] P. Toivanen et al., “Large-scale shell-model calculations of elastic and inelastic scattering rates of lightest supersymmetric particles (LSP) on I 127, Xe 129, Xe 131, and Cs 133 nuclei”, *Physical Review C* **79** (2009), no. 4, 044302.
- [89] R. W. Schnee, “Introduction to dark matter experiments”, in *Physics of the Large and the Small: TASI 2009*, pp. 775–829. World Scientific, 2011.
- [90] A. L. Fitzpatrick et al., “The effective field theory of dark matter direct detection”, *Journal of Cosmology and Astroparticle Physics* **2013** (2013), no. 02, 004.
- [91] A. Bueno et al., “Observation of coherent neutrino-nucleus elastic scattering at a beta beam”, *Physical Review D* **74** (2006), no. 3, 033010.
- [92] XENON-Collaboration, “Constraining the spin-dependent WIMP-nucleon cross sections with XENON1T”, *Physical Review Letters* **122** (2019), no. 14, 141301.
- [93] J. Xia et al., “PandaX-II constraints on spin-dependent WIMP-nucleon effective interactions”, *Physics Letters B* **792** (2019) 193–198.
- [94] D. Akerib et al., “Limits on spin-dependent WIMP-nucleon cross section obtained from the complete LUX exposure”, *Physical Review Letters* **118** (2017), no. 25, 251302.
- [95] N. Ackerman et al., “Observation of Two-Neutrino Double-Beta Decay in Xe 136 with the EXO-200 Detector”, *Physical Review Letters* **107** (2011), no. 21, 212501.
- [96] A. Gando et al., “Measurement of the double- β decay half-life of ^{136}Xe with the KamLAND-Zen experiment”, *Physical Review C* **85** (2012), no. 4, 045504.
- [97] J. Albert et al., “Sensitivity and discovery potential of the proposed nEXO experiment to neutrinoless double- β decay”, *Physical Review C* **97** (2018), no. 6, 065503.
- [98] D. Akerib et al., “Projected sensitivity of the LUX-ZEPLIN experiment to the $0\nu\beta\beta$ decay of ^{136}Xe ”, *arXiv:1912.04248* (2019).
- [99] A. Gando et al., “Search for Majorana neutrinos near the inverted mass hierarchy region with KamLAND-Zen”, *Physical Review Letters* **117** (2016), no. 8, 082503.
- [100] J. Albert et al., “Search for neutrinoless double-beta decay with the upgraded EXO-200 detector”, *Physical Review Letters* **120** (2018), no. 7, 072701.

- [101] J. Martín-Albo et al., “Sensitivity of NEXT-100 to neutrinoless double beta decay”, *Journal of High Energy Physics* **2016** (2016), no. 5, 159.
- [102] E. Majorana, “Teoria simmetrica dell’elettrone e del positrone”, *Il Nuovo Cimento (1924-1942)* **14** (1937), no. 4, 171.
- [103] E. Aprile et al., “Observation of two-neutrino double electron capture in ^{124}Xe with XENON1T”, *Nature* **568** (2019), no. 7753, 532–535.
- [104] G. A. Cook, “Argon, Helium, and the Rare Gases: The Elements of the Helium Group”, volume 2. Interscience Publishers, 1961.
- [105] E. Aprile et al., “Liquid xenon detectors for particle physics and astrophysics”, *Reviews of Modern Physics* **82** (2010), no. 3, 2053.
- [106] B. Lenardo et al., “A global analysis of light and charge yields in liquid xenon”, *IEEE Transactions on Nuclear Science* **62** (2015), no. 6, 3387–3396.
- [107] V. Chepel et al., “Liquid noble gas detectors for low energy particle physics”, *Journal of Instrumentation* **8** (2013) R04001, doi:10.1088/1748-0221/8/04/R04001.
- [108] K. Fujii et al., “High-accuracy measurement of the emission spectrum of liquid xenon in the vacuum ultraviolet region”, *Nuclear Instruments and Methods in Physics Research Section A* **795** (2015) 293–297.
- [109] A. Hitachi et al., “Effect of ionization density on the time dependence of luminescence from liquid argon and xenon”, *Physical Review B* **27** (1983), no. 9, 5279.
- [110] C. E. Dahl, “The physics of background discrimination in liquid xenon, and first results from Xenon10 in the hunt for WIMP dark matter”, technical report, (2009).
- [111] J. Lindhard et al., “Integral equations governing radiation effects”, *Mat. Fys. Medd. Dan. Vid. Selsk* **33** (1963), no. 10, 1–42.
- [112] P. Sorensen et al., “Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter”, *Physical Review D* **83** (2011), no. 6, 063501.
- [113] D. Akerib et al., “Low-energy (0.7-74 keV) nuclear recoil calibration of the LUX dark matter experiment using DD neutron scattering kinematics”, *arXiv:1608.05381* (2016).
- [114] E. Aprile et al., “Response of the XENON100 dark matter detector to nuclear recoils”, *Physical Review D* **88** (2013), no. 1, 012006.
- [115] A. Manzur et al., “Scintillation efficiency and ionization yield of liquid xenon for monoenergetic nuclear recoils down to 4 keV”, *Physical Review C* **81** (2010), no. 2, 025808.

- [116] M. Horn et al., “Nuclear recoil scintillation and ionisation yields in liquid xenon from ZEPLIN-III data”, *Physics Letters B* **705** (2011), no. 5, 471–476.
- [117] E. Aprile et al., “Simultaneous measurement of ionization and scintillation from nuclear recoils in liquid xenon for a dark matter experiment”, *Physical Review Letters* **97** (2006), no. 8, 081302.
- [118] B. Lenardo et al., “Measurement of the ionization yield from nuclear recoils in liquid xenon between 0.3–6 keV with single-ionization-electron sensitivity”, *arXiv:1908.00518* (2019).
- [119] D. Akerib et al., “Tritium calibration of the LUX dark matter experiment”, *Physical Review D* **93** (2016), no. 7, 072009.
- [120] L. Goetzke et al., “Measurement of light and charge yield of low-energy electronic recoils in liquid xenon”, *Physical Review D* **96** (2017), no. 10, 103007.
- [121] D. Akerib et al., “Improved measurements of the β -decay response of liquid xenon with the LUX detector”, *Physical Review D* **100** (2019), no. 2, 022002.
- [122] D. Akerib et al., “Ultralow energy calibration of LUX detector using Xe 127 electron capture”, *Physical Review D* **96** (2017), no. 11, 112011.
- [123] M. Szydagis et al., “NEST: a comprehensive model for scintillation yield in liquid xenon”, *Journal of Instrumentation* **6** (2011), no. 10, P10002.
- [124] M. Szydagis et al. <http://doi.org/10.5281/zenodo.1314669> (2018).
- [125] B. Edwards et al., “Extraction efficiency of drifting electrons in a two-phase xenon time projection chamber”, *Journal of Instrumentation* **13** (2018), no. 01, P01005.
- [126] C. Monteiro et al., “Secondary scintillation yield in pure xenon”, *Journal of Instrumentation* **2** (2007), no. 05, P05001.
- [127] D. Akerib et al., “Improved limits on scattering of weakly interacting massive particles from reanalysis of 2013 LUX data”, *Physical Review Letters* **116** (2016), no. 16, 161301.
- [128] D. Akerib et al., “Discrimination of electronic recoils from nuclear recoils in two-phase xenon time projection chambers”, *arXiv:2004.06304* (2020).
- [129] K. Abe, K. Hiraide, K. Ichimura et al., “Direct dark matter search by annual modulation in XMASS-I”, *Physics Letters B* **759** (2016) 272–276.
- [130] G. Alner, H. Araujo, G. Arnison et al., “First limits on nuclear recoil events from the ZEPLIN I galactic dark matter detector”, *Astroparticle Physics* **23** (2005), no. 5, 444–462.

- [131] D. Akerib et al., “Liquid xenon scintillation measurements and pulse shape discrimination in the LUX dark matter detector”, *Physical Review D* **97** (2018), no. 11, 112002.
- [132] P. Adhikari et al., “The liquid-argon scintillation pulseshape in DEAP-3600”, *arXiv:2001.09855* (2020).
- [133] W. Lippincott et al., “Scintillation time dependence and pulse shape discrimination in liquid argon”, *Physical Review C* **78** (2008), no. 3, 035801.
- [134] D. Akerib et al., “Results from a search for dark matter in the complete LUX exposure”, *Physical Review Letters* **118** (2017), no. 2, 021303.
- [135] D. Akerib et al., “Calibration, event reconstruction, data analysis, and limit calculation for the LUX dark matter experiment”, *Physical Review D* **97** (2018), no. 10, 102008.
- [136] D. S. Akerib et al., “First results from the LUX dark matter experiment at the Sanford Underground Research Facility”, *Physical Review Letters* **112** (2014), no. 9, 091303.
- [137] D. Akerib, H. Araújo, X. Bai et al., “Chromatographic separation of radioactive noble gases from xenon”, *Astroparticle Physics* **97** (2018) 80–87.
- [138] D. Akerib et al., “Radiogenic and muon-induced backgrounds in the LUX dark matter detector”, *Astroparticle Physics* **62** (2015) 33–46.
- [139] A. Bolozdynya et al., “Cryogenics for the LUX detector”, *IEEE Transactions on Nuclear Science* **56** (2009), no. 4, 2309–2312.
- [140] D. Akerib et al., “The LUX prototype detector: Heat exchanger development”, *Nuclear Instruments and Methods in Physics Research Section A* **709** (2013) 29–36.
- [141] D. Akerib et al., “Kr 83 m calibration of the 2013 LUX dark matter search”, *Physical Review D* **96** (2017), no. 11, 112009.
- [142] A. Manalaysay, T. M. Undagoitia, A. Askin et al., “Spatially uniform calibration of a liquid xenon detector at low energies using K 83 m”, *Review of Scientific Instruments* **81** (2010), no. 7, 073303.
- [143] J. Menendez et al., “Spin-dependent WIMP scattering off nuclei”, *Physical Review D* **86** (2012), no. 10, 103511.
- [144] D. Akerib et al., “First searches for axions and axionlike particles with the LUX experiment”, *Physical Review Letters* **118** (2017), no. 26, 261301.
- [145] D. Akerib et al., “Results of a search for sub-GeV dark matter using 2013 LUX data”, *Physical Review Letters* **122** (2019), no. 13, 131301.

- [146] D. Akerib et al., “Data acquisition and readout system for the LUX dark matter experiment”, *Nuclear Instruments and Methods in Physics Research Section A* **668** (2012) 1–8.
- [147] D. Akerib et al., “FPGA-based trigger system for the LUX dark matter experiment”, *Nuclear Instruments and Methods in Physics Research Section A* **818** (2016) 57–67.
- [148] D. Akerib et al., “LUX trigger efficiency”, *Nuclear Instruments and Methods in Physics Research Section A* **908** (2018) 401–410.
- [149] D. Akerib et al., “Position reconstruction in LUX”, *Journal of Instrumentation* **13** (2018), no. 02, P02001.
- [150] J. Angle et al., “Search for light dark matter in XENON10 data”, *Physical Review Letters* **107** (2011), no. 5, 051301.
- [151] E. Aprile et al., “Low-mass dark matter search using ionization signals in XENON100”, *Physical Review D* **94** (2016), no. 9, 092001.
- [152] E. Aprile et al., “Light Dark Matter Search with Ionization Signals in XENON1T”, *Physical Review Letters* **123** (2019), no. 25, 251801.
- [153] D. Akerib et al., “Investigation of background electron emission in the LUX detector”, *arXiv:2004.07791* (2020).
- [154] A. Tomás et al., “Study and mitigation of spurious electron emission from cathodic wires in noble liquid time projection chambers”, *Astropart. Phys.* **103** (2018) 49–61.
- [155] P. Agnes et al., “Low-Mass Dark Matter Search with the DarkSide-50 Experiment”, *Physical Review Letters* **121** (2018), no. 8, 081307.
- [156] M. Tanaka et al., “Studies on liquid argon S1 and S2 properties for low mass WIMP search experiments”, in *Journal of Physics: Conference Series*, volume 1468, p. 012052. 2020.
- [157] C. Kouvaris et al., “Probing sub-GeV dark matter with conventional detectors”, *Physical Review Letters* **118** (2017), no. 3, 031803.
- [158] M. Ibe et al., “Migdal effect in dark matter direct detection experiments”, *Journal of High Energy Physics* **2018** (2018), no. 3, 194.
- [159] W. H. Lippincott et al., “Increasing the sensitivity of LXe TPCs to dark matter by doping with helium or neon”, technical report, (2017).
- [160] “Photomultiplier Handbook: Theory, Design, Application”. Burle Industries, 1980.

- [161] D. Akerib, X. Bai, E. Bernard et al., “An ultra-low background PMT for liquid xenon detectors”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **703** (2013) 1–6.
- [162] C. Faham, “Prototype, Surface Commissioning and Photomultiplier Tube Characterization for the Large Underground Xenon (LUX) Direct Dark Matter Search Experiment”. PhD thesis, Brown University, 5, 2014.
- [163] E. Aprile et al., “The XENON dark matter search experiment”, *New Astronomy Reviews* **49** (2005), no. 2, 289–295.
- [164] C. Faham et al., “Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes”, *Journal of Instrumentation* **10** (2015), no. 09, P09010.
- [165] B. López-Paredes et al., “Response of photomultiplier tubes to xenon scintillation light”, *Astroparticle Physics* **102** (2018) 56–66.
- [166] P. Barrow et al., “Qualification tests of the R11410-21 photomultiplier tubes for the XENON1T detector”, *Journal of Instrumentation* **12** (2017), no. 01, P01024.
- [167] E. Aprile et al., “The XENON1T dark matter experiment”, *The European Physical Journal C* **77** (2017), no. 12, 881.
- [168] D. Gibbs et al., “Ageing of materials under intense ultraviolet radiation”, *Metrologia* **32** (1995), no. 6, 601.
- [169] S. Shaw et al., “Ultraviolet characterization of integrating spheres”, *Applied Optics* **46** (2007), no. 22, 5119–5128.
- [170] P. Amaudruz et al., “In-situ characterization of the Hamamatsu R5912-HQE photomultiplier tubes used in the DEAP-3600 experiment”, *Nuclear Instruments and Methods in Physics Research Section A* **922** (2019) 373–384.
- [171] H. Araújo et al., “Radioactivity backgrounds in ZEPLIN–III”, *Astroparticle Physics* **35** (2012), no. 8, 495–502.
- [172] ZEPLIN-III Collaboration, “Single electron emission in two-phase xenon with application to the detection of coherent neutrino-nucleus scattering”, *JHEP* **2011** (Dec, 2011) 115, doi:10.1007/JHEP12(2011)115.
- [173] E. Boulton et al., “Calibration of a two-phase xenon time projection chamber with a ^{37}Ar source”, *Journal of Instrumentation* **12** (2017), no. 08, P08004.
- [174] B. Lenardo, J. Xu, S. Pereverzev et al., “Measurement of the ionization yield from nuclear recoils in liquid xenon between 0.3–6 keV with single-ionization-electron

- sensitivity”, *arXiv preprint arXiv:1908.00518* (2019).
- [175] G. J. Feldman et al., “Unified approach to the classical statistical analysis of small signals”, *Physical Review D* **57** (1998), no. 7, 3873.
- [176] W. A. Rolke et al., “Limits and confidence intervals in the presence of nuisance parameters”, *Nuclear Instruments and Methods in Physics Research Section A* **551** (2005), no. 2-3, 493–503.
- [177] X. Cui et al., “Dark matter results from 54-ton-day exposure of PandaX-II experiment”, *Physical Review Letters* **119** (2017), no. 18, 181302.
- [178] C. Amole et al., “Dark Matter Search Results from the PICO- 60 C 3 F 8 Bubble Chamber”, *Physical Review Letters* **118** (2017), no. 25, 251301.
- [179] R. Agnese et al., “WIMP-Search Results from the Second CDMSlite Run”, *Physical Review Letters* **116** (2016).
- [180] D. Akerib, C. Akerlof, A. Alqahtani et al., “Simulations of Events for the LUX-ZEPLIN (LZ) Dark Matter Experiment”, *arXiv preprint arXiv:2001.09363* (2020).
- [181] S. Agostinelli, J. Allison, K. a. Amako et al., “GEANT4—a simulation toolkit”, *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506** (2003), no. 3, 250–303.
- [182] D. Akerib et al., “Projected WIMP sensitivity of the LUX-ZEPLIN (LZ) dark matter experiment”, *arXiv:1802.06039* (2018).
- [183] G. Cowan, “Statistical data analysis”. Oxford university press, 1998.