

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Discrimination in Liquid Xenon and Calorimetry in Superfluid Helium for the Direct Detection of Particle Dark Matter

Permalink

<https://escholarship.org/uc/item/1hm8m6zt>

Author

Velan, Vetri

Publication Date

2022

Peer reviewed|Thesis/dissertation

Discrimination in Liquid Xenon and Calorimetry in Superfluid Helium for the
Direct Detection of Particle Dark Matter

by

Vetri Velan

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Daniel N. McKinsey, Chair

Professor Marjorie D. Shapiro

Professor Karl A. van Bibber

Summer 2022

Discrimination in Liquid Xenon and Calorimetry in Superfluid Helium for the
Direct Detection of Particle Dark Matter

Copyright 2022
by
Vetri Velan

Abstract

Discrimination in Liquid Xenon and Calorimetry in Superfluid Helium for the
Direct Detection of Particle Dark Matter

by

Vetri Velan

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Daniel N. McKinsey, Chair

In this dissertation, we present several recent developments in instrumentation for dark matter detectors. The liquid/gas two-phase xenon time projection chamber is well-established as an excellent technology to search for Weakly Interacting Massive Particles (WIMPs). We present analyses of data from the Large Underground Xenon (LUX) and Particle Identification in Xenon at Yale (PIXeY) experiments, in which we study signal vs. background discrimination as a function of detector parameters. This informs design decisions for current and future xenon detectors. Then, we focus on cutting-edge calorimetric technologies geared towards searches at the MeV/c² scale. We present measurements of the superconducting transition of several materials, and we discuss their potential use as transition-edge sensors (TES). Helium as a detection target with TES signal readout is explored, and its sensitivity to dark matter-induced nuclear recoils is calculated.

Contents

Contents	i
1 Introduction	1
2 Foundations of Dark Matter and Direct Detection	4
2.1 Evidence for Particle Dark Matter	5
2.2 Weakly Interacting Massive Particles	9
2.3 Models of Dark Matter at the MeV Scale	10
2.4 The Nuclear Recoil Signal	11
2.4.1 Direct Detection of Nuclear (and Other) Recoils	11
2.4.2 The Dark Matter Signal Model	13
2.5 The Neutrino Fog	19
2.6 Current Status of the Field	23
3 Particle Detection in Liquid Xenon	25
3.1 History of Xenon in Particle Physics	26
3.2 Properties of Liquid Xenon	29
3.3 Energy Partitioning in Liquid Xenon	31
3.3.1 Three Channels of Energy Deposition	31
3.3.2 Electronic Recoils	34
3.3.3 Nuclear Recoils	39
3.3.4 ^{83m}Kr Decays	43
3.4 Two-Phase Liquid/Gas Xenon Time Projection Chambers	46
3.5 Noble Element Simulation Technique (NEST)	51
4 ER/NR Discrimination in the LUX Detector	55
4.1 Backgrounds for Xenon TPCs	56
4.1.1 LZ and G3 Backgrounds	56
4.1.2 LUX Backgrounds	60
4.2 The Large Underground Xenon (LUX) Experiment	61
4.2.1 About the detector	61
4.2.2 Calibrations	64

4.2.3	Electric field variation	65
4.3	Tools for Discrimination	65
4.4	Electronic and Nuclear Recoil Bands	70
4.4.1	Electronic Recoils	70
4.4.1.1	Aside: Statistical Strategy for Fitting $\log_{10}(S2c/S1c)$	78
4.4.2	Nuclear Recoils	78
4.4.3	Variation with g_1	79
4.4.4	Variation with drift field	82
4.5	Leakage and Discrimination Results	85
4.5.1	Charge-to-light discrimination	85
4.5.1.1	Variation with g_1	88
4.5.1.2	Variation with drift field	91
4.5.2	Pulse Shape Discrimination	92
5	Simulations of Liquid Xenon Detectors	103
5.1	A Model of ER and NR Band Skewness	104
5.1.1	ER Skewness	104
5.1.2	NR Skewness	111
5.2	Fluctuations of the ER Band	111
6	ER/NR Discrimination in the PIXeY Detector	115
6.1	The Particle Identification in Xenon at Yale (PIXeY) Experiment . .	116
6.1.1	About the detector	116
6.1.2	Run Details	117
6.2	Data Processing	119
6.3	Detector Calibration with ^{83m}Kr	125
6.3.1	Position-Dependence of S1 and S2	125
6.3.2	Definition of Fiducial Region	129
6.3.3	Determination of g_1 and g_2	129
6.4	Nuclear Recoils	131
6.4.1	NR Data Selection	131
6.4.2	NR Band as a Function of Drift Field	132
6.5	Electronic Recoils	134
6.5.1	ER Data Selection	134
6.5.2	ER Band as a Function of Drift Field	134
6.6	Leakage and Discrimination	139
7	Calibration of Low-Temperature Sensors for the TESSERACT Project	145
7.1	SPICE, HeRALD, and TESSERACT	148
7.2	Experimental Apparatus: Dilution Refrigerator at LBNL	151
7.3	Introduction to TES Dynamics	155

7.4	Measurement of 46 mK Transition Edge Sensors	158
7.4.1	Experimental Setup for TES Measurements	159
7.4.2	A Single TES Measurement	164
7.4.3	TES Measurement Results: Bias Power and Conductance . . .	167
7.4.4	TES Measurement Results: Normal and Parasitic Resistance .	169
7.5	Measurement of T_c and Film Stress for Several Materials	169
7.5.1	Experimental Setup for T_c Measurements	170
7.5.2	T_c as a Function of Chamber Gas	171
7.5.3	Relationship Between T_c and Thin-Film Stress	174
8	Helium as a Dark Matter Direct Detection Target	181
8.1	Detector Concept	182
8.1.1	Energy Partitioning in Superfluid ^4He	182
8.1.2	Sensing Energy Depositions	186
8.1.3	Electronic Recoils, Nuclear Recoils, and Discrimination	188
8.2	Background Model	190
8.3	HeRALD Sensitivity to Nuclear Recoils	194
8.4	HeRALD Sensitivity to Electronic Recoils	200
8.5	The Neutrino Fog in Helium	200
9	Conclusion	202
	Bibliography	205

Acknowledgments

This dissertation is seven years in the making, and I am grateful to everyone who supported me through it. I couldn't possibly list every single person who impacted me, so thank you to those whom I have forgotten below or can't list by name.

My first words of gratitude go to my advisor, Dan. Thank you for many years of support and training. Your breadth of knowledge, passion for experimentation, and ability to identify exciting projects are remarkable and have made me a better physicist. Thank you for encouraging me to apply for fellowships, attend conferences and workshops, and so on, and for allowing me the time to pursue those opportunities. I'm glad that my first semester, Marjorie encouraged me to take 290E with you, where I learned about dark matter detection and decided to dedicate my PhD to it!

Thank you to Marjorie and Karl for reviewing this dissertation. Your feedback has made it a better work. Marjorie, thank you also for being my first-year advisor and for hooding me at graduation.

All my colleagues and collaborators have taught me an immense amount, but I acknowledge four in particular for directly training me. Kevin O was my first research mentor in grad school. Nearly everything I know about xenon TPCs, I learned from him, as well as the evergreen advice of always looking at the next step in my career. Quentin, I've rarely met someone better at physics analysis, even though we disagree on Python vs. C++. Thanks also for keeping your espresso cabinet stocked. Junsong, words cannot describe the support you've given me over several years of sharing an office. Whether that's questions about helium physics, how to use BearBuy, or reviewing my fellowship applications, you've always been willing to take time to help me. I am always inspired by your focus and energy. Toki, thanks for taking in a 6th-year grad student who just needed a hardware project to graduate, and getting me excited about hardware work. Thank you for teaching me a ridiculous amount about TESSs, dilution refrigerators, and SQUIDs, and for writing me a recommendation letter for postdoc applications.

I am grateful to all my LUX, SPICE/HeRALD, and NEST collaborators, and I highlight a few. Matthew and Rick, thank you for providing valuable insight to my LUX discrimination analysis over the course of some three years. Matthew, thank you further for inviting me into NEST. Scott He, you have provided career advice at multiple stages of my grad school tenure, and valuable wisdom on all things helium, both as a postdoc and as a professor. Kevin L, although you were not my official supervisor (that will change soon!), I've always been able to come to you for advice on various aspects of grad school and doing science. Thank you for your unwavering spirit in leading the LBNL group, and for the recommendations to restaurants and hiking spots. Rick, Scott, and Kevin, thank you further for your recommendation letters for various fellowship applications over the years.

At UC Berkeley and LBNL, I had the fortune of working alongside an amazing group of people. Peter, Simon, Aaron, Scott K, Scott Ha, Shilo, Ibles, Hao, Suerfu, Xinran, Kaja, Andy, Reed, Ryan S, Roger, Jose, Ryan G, Yue, Olivia, Ethan, Evan, Brian, Daniel H, Elizabeth, Lucie, Kate, Kelsey, Anisha, and Eve, thank you for many memories.

I have received many other forms of professional support that I am grateful for. During parts of graduate school, I was lucky to be financially supported by scholarships from the Physics Department and CPAD. I am also grateful for the Physics Department administrative staff, who were able to answer all my random questions throughout my time at UC Berkeley. I am further thankful to everyone who helped me become a better teacher and GSI, including the instructors and co-GSIs of 7B. Finally, thank you to all of my colleagues, advisors, and teachers from before I arrived in Berkeley. Even if we regrettably haven't kept in touch, I entered graduate school more prepared because of your impact. In particular, thank you to my undergrad physics research mentors: Sunil, Frank, and Ahmed.

Beyond my professional network, I am endlessly grateful to my social supports. To begin, I acknowledge all of the amazing friends I made at Rutgers and Berkeley. I regret that I have let too many of these friendships lapse, especially in the past few years. Regardless of whether we are still in touch or haven't spoken in a while, I am so thankful for the time we spent together.

To my friends in and around New Jersey, you always make the time to see me when I'm back. I wish I could see you more frequently, but I am grateful for the various dinners, hangouts, and (more recently) weddings that have occurred during my time in grad school. These have reminded me of warm memories and created new ones to cherish. My friends from freshman year in Barr Hall (and nearby dorms, but let's be honest, we were usually in Barr), chemical engineering classes, EGC, OT, and C&S: I've loved every single time we've spent together. More of you need to move out to the West Coast! In Berkeley, I recognize all the friends I've made around campus—graduate students in Physics, students in other departments, residents or mafia-players at the I-House, members of the organizations listed below, and others. I also appreciate all the apartment-mates I had in Berkeley for being gracious and respectful—and especially Ari, Vir, Reed, and Lizzy for the additional friendship.

Thank you to Shiv for many years of friendship. You're my oldest friend that I'm still regularly in touch with; I think our friendship is now old enough to vote. Ajay, thanks for moving to Berkeley and many memories, and for helping me move apartments. Micah, I appreciated the coffee breaks, knowledge of Canadian politics (I now know what a "riding" is), and your frequent presence in the 7B office despite never having a key. Mike R, I enjoyed the many vigorous political arguments and our common taste in TV. You also always reached out, even when I was swamped and hadn't texted in a while, and I am grateful for that. Jessica, we met while stressing about grad school decisions, and we managed to continue stressing seven

years later, but this time about postdoc applications and decisions. I so appreciate having someone who could relate during that process. I still haven't found dark matter, though—oh. Oops. Spoiler alert, I guess.

Surprising no one, I spent grad school involved in many things, often a welcome respite from research. First, I appreciate everyone I met through my science policy journey. Friends, colleagues, and mentors helped me mature into a stronger advocate. I am especially grateful for the friendships of Kathy and Eric. Our conversations about funny tweets and badly-made graphs bring me joy, as did the time we enjoyed pizza and wine in the park near my apartment. Second, I am grateful to all the organizers I worked with in UAW 2865 and Indivisible Berkeley, who taught me countless political skills. Sam and Jaime, in particular, thank you for your friendship; you always make me feel welcome and supported. I fondly remember enjoying pizza and wine in...a park near my apartment. There seems to be a trend emerging. But there was a pandemic, there wasn't much else we could do!

Next, I am immensely grateful for the friends I regularly hung out with at a particular establishment in Berkeley. To Aniket, Carly, and Harrison, especially, thank you for many years of support. Our hangouts were a delight, frequently the highlight of my week. I always appreciated spending time with you, whether through libations, trips to Tahoe, or virtual movie and game nights. Besides the people I lived and worked with, I probably saw you more frequently than I saw anyone else, and I'm so fortunate to have you all in my life.

Shane, Saad, Jake, and Anne, thank you for everything. The over-investment in minor Twitter drama, discussion of the multiple parallel calamities in the world, jokes and memes, help making life decisions, empathy for work/life ordeals, and constant support whenever I felt overwhelmed—I am eternally grateful for all of it. You're the best fake editorial board any publication could ask for. Sharlina, you have always been there for me, even when I was unsure about so much. Thank you for your friendship and your support over many years. Even when months have passed by without talking, it always feels like nothing has changed when we do. My gratitude and appreciation are boundless.

My gratitude also goes to my cousin Ragavi, and her family Suriya and Rakshiv. You have always offered me a home in your home, without hesitation. From letting me stay with you when I moved to Berkeley, to inviting me in whenever I need a break, I am grateful for all of it. I'm so fortunate that you moved to the Bay Area at nearly the same time I did. Rakshiv, as you get older, remember you can accomplish anything you set your mind to.

Finally, and most importantly, thank you to my mom for 29 years of love. From day one (literally), you put all your energy into giving me the best life you could. Through good times and bad times, you have never stopped caring about me. I do not know where I would be with your presence and your support. Words cannot describe how grateful I am.

Chapter 1

Introduction

The direct detection of particle dark matter through its interactions with a terrestrial experiment is one of the greatest challenges in modern particle physics. Some dozens of experiments have been built in the past few decades to discover what dark matter is made of, and the search will continue into the future. These experiments deploy a variety of technologies—liquid noble elements, gaseous noble elements, calorimetry, crystal scintillation, bubble chambers, toroidal magnets, and qubits, just to name a few. For each experiment, scientists need to understand not only how dark matter galactic halo particles interact with the Standard Model, but also the nature of energy deposition in the relevant technology.

This dissertation describes several studies, unified by the pursuit of developing world-class instruments for dark matter direct detection. This pursuit is approached through two complementary avenues: understanding the dynamics of energy deposition in different materials, and enhancing an experiment’s ability to sense and read out a signal.

We begin in Chapter 2 by laying out the evidence for particle dark matter and motivating models with dark matter candidates in the MeV/c^2 – TeV/c^2 mass range. These models can induce nuclear recoils in normal atoms, making nuclear recoil searches the dominant experimental strategy. We introduce the signature and describe the current and future status of direct detection experiments.

The next four chapters are dedicated to studying liquid/gas two-phase xenon time projection chambers, currently the most sensitive technology in searches for Weakly Interacting Massive Particle (WIMP) dark matter. Chapter 3 describes the current understanding of energy deposition and particle detection in liquid xenon. Chapter 4 presents an analysis of nuclear recoil vs. electronic recoil discrimination using data from the Large Underground Xenon (LUX) experiment, studying how discrimination varies based on detector design. Electronic recoils are the dominant background for nuclear recoil searches. Chapter 5 builds on this analysis, using the LUX data to produce detector simulations, which in turn inform us about the causes behind better and worse discrimination. Chapter 6 employs data from the Particle Identification in Xenon at Yale (PIXeY) experiment. PIXeY is not a dark matter search, but it deploys the same technology, and we use its data to study discrimination across a wider range of detector parameters.

The next two chapters are focused on calorimetric detectors, which can probe dark matter with masses below the WIMP scale. In Chapter 7, we introduce the TESSERACT project, a three-target approach to search for sub-GeV dark matter using transition edge sensor (TES) readout. A tungsten film on a silicon substrate is a promising TES material, so the superconducting transition temperature and stress of several films are measured, in order to find the fabrication procedure that minimizes both variables. The bias power of two sensors is also measured. In Chapter 8, we focus on one of the three TESSERACT experiments: the Helium Roton Apparatus for Light Dark Matter (HeRALD). We describe how to identify energy depositions in superfluid ^4He , and the reach of a ^4He dark matter detector

is calculated.

Finally, we conclude in Chapter 9, summarizing our main findings and motivating future work.

Chapter 2

Foundations of Dark Matter and Direct Detection

Non-luminous, or *dark*, matter comprises 26.6% of the energy density of the universe [1]. It interacts gravitationally but not electromagnetically. Dark matter could potentially interact through other unknown channels, and discovering these channels is one of the primary unsolved problems in modern physics.

In this Chapter, we begin by discussing the evidence for dark matter (DM) as a particle. We then describe a few dark matter models, most notably the Weakly Interacting Mass Particle (WIMP), but also a few alternatives at lower masses. We then describe how to identify dark matter particles through nuclear recoil searches, and we briefly describe how neutrino interactions will eventually be an obstacle to this search. We conclude by presenting the current best limits on dark matter interactions.

2.1 Evidence for Particle Dark Matter

The evidence for particle dark matter is abundant. The most well-known evidence is galactic rotation curves. The first demonstration of inconsistencies in these curves was by Fritz Zwicky in 1933. Zwicky measured the radial velocity dispersion of galaxies in the Coma cluster and found large discrepancies from the behaviors predicted by the virial theorem [2]. But it wasn't until decades later that the problem became solidified. In the 1970s, Vera Rubin, Kent Ford, and Norbert Thonnard measured the rotational speed profiles of several galaxies and found, strikingly, that the speeds were roughly constant even at radial distances of 20 kpc (see Fig. 2.1). This was inconsistent with observations of the light emitted from these galaxies, which is a proxy for the visible mass of the galaxies. According to astronomical observations, the mass decreases at high radii, so the rotational velocity (propelled by gravity) should decrease as well. One way to explain this was an extra type of missing matter that extended in a “halo” to higher radii than the visible matter.

This alone is not sufficient to theorize a whole new type of matter, of course. One could alternatively come up with new dynamics for gravity, as has been done most famously by Mordehai Milgrom in his theory of Modified Newtonian Dynamics (MOND) [3]. But other evidence for dark matter exists, too.

Cosmological evidence for dark matter is particularly compelling—specifically, the Cosmic Microwave Background (CMB). The CMB is the remnants of photons created about 370,000 years after the Big Bang (equivalently, at a redshift of $z \approx 1100$) [6]. These photons were generated during recombination, when the universe cooled down sufficiently that protons and electrons could form neutral hydrogen, and photons could propagate through the cosmos relatively unhindered by scattering. Today, the photons are still observable at a temperature of 2.7255 K, corresponding to a maximum frequency of 160 GHz. CMB photons are relatively uniform across the sky, but small-scale fluctuations in temperature and polarization

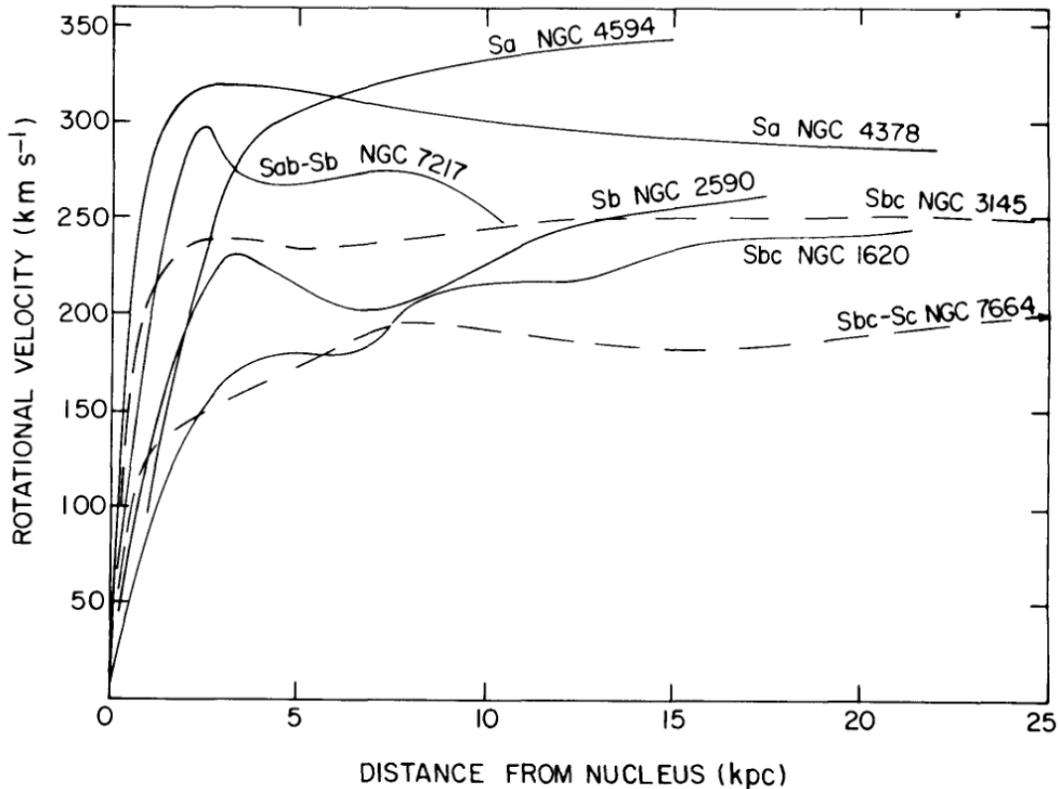


Figure 2.1: Rotation velocity vs. radius for seven galaxies showing flat velocities at high distances from the galactic nucleus, measured by Rubin et al. Copied from [4].

(on the order of 10^{-5}) reveal how these photons were scattered before and during their journey.

At the instant before recombination, the universe was comprised of a plasma of baryons and photons, and potentially dark matter. Quantum fluctuations led to certain regions of the plasma being denser than others. Gravitational attraction and radiation pressure drove oscillations in the plasma, which were frozen (i.e., could not propagate further) when baryons and photons decoupled during recombination. These regions of high and low density are still visible today in the CMB temperature power spectrum, measured through a spherical harmonic decomposition of CMB photon thermal fluctuations across the sky. Planck's measurement of the power spectrum is shown in Fig. 2.2. Higher values of l correspond to smaller scales of density fluctuation in the plasma. The odd-numbered peaks in the power spectrum—first, third, fifth, etc.—correspond to compression, and the even-numbered peaks correspond to rarefaction. Since any matter induces a gravitational potential well, all peaks are affected by the total matter density. Their amplitudes are diminished if the initial matter density was high. Meanwhile, non-baryonic matter only con-

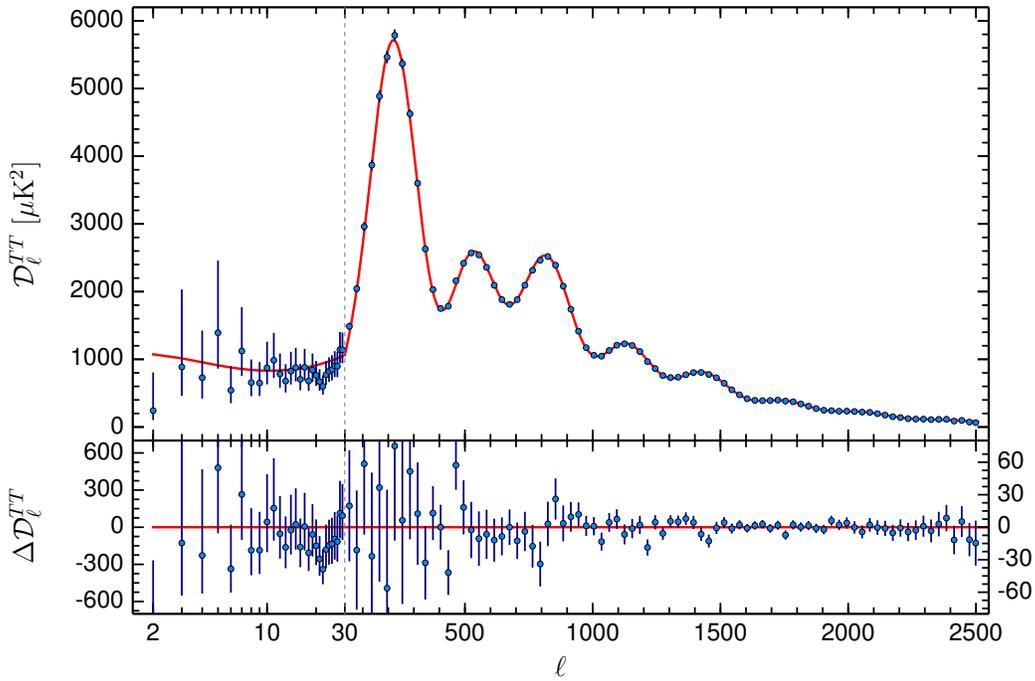


Figure 2.2: Temperature power spectrum of the CMB from Planck: data in blue, fit to Λ CDM model in red. Figure copied from [1].

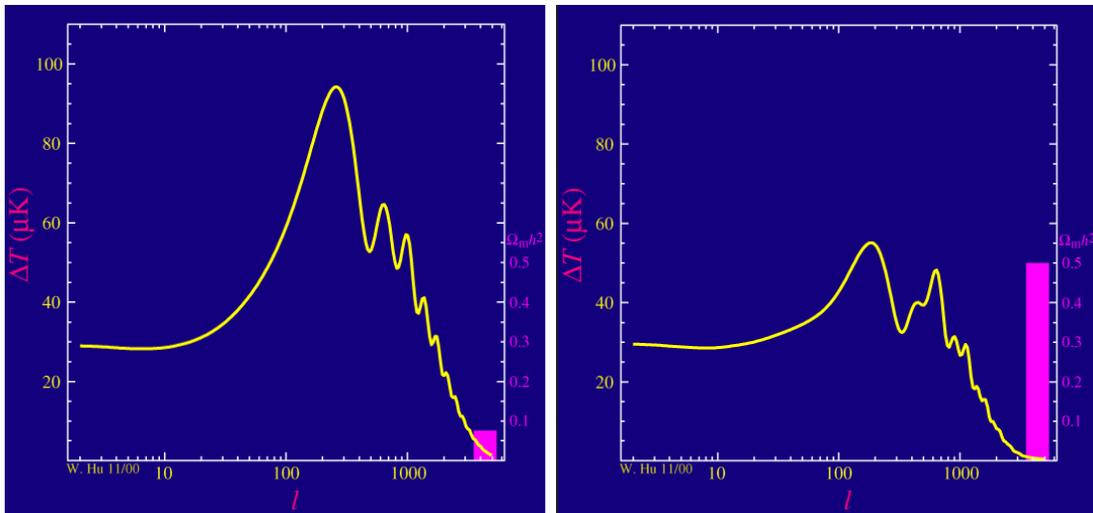


Figure 2.3: The CMB power spectrum in a universe with low (left) and high (right) dark matter density. Here, Ω_m is the total matter density, and the baryonic contribution is kept constant, so only the DM fraction varies. Figures from [5].

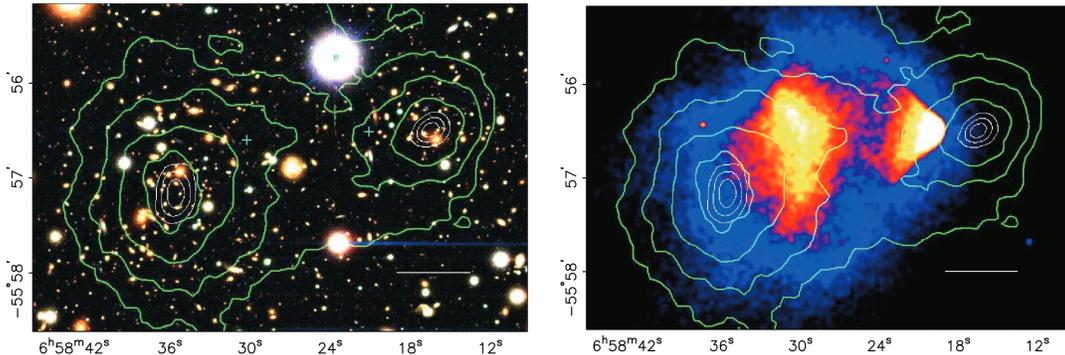


Figure 2.4: Optical (*left*) and X-ray (*right*) images of the galaxy cluster 1E0657-558, commonly known as the “Bullet Cluster.” The green lines in both images show contours of the reconstructed mass from gravitational lensing. Images from [8].

tributed to compression, not rarefaction, in the plasma. The even-numbered peaks are thus governed solely by the baryons. Crucially, the third peak is enhanced to the level of the second peak or higher if the non-baryonic matter contribution was high. The fit to Planck data gives $\Omega_b h^2 = 0.02225$, $\Omega_m h^2 = 0.1428$, and $h = 0.6727$, where Ω_b is the current baryon density, Ω_m is the current matter density, and h is the Hubble constant normalized to 100 km/s/Mpc. Clearly, some 84% of the matter density arises from sources other than baryons.

Graphically, the CMB animations by Wayne Hu are a convenient way to observe the difference between a universe with and without dark matter [7]. We show this in Fig. 2.3, where Ω_b is kept constant while Ω_m varies (in essence, varying the dark matter density). In a high-DM universe, the total amplitude of the power spectrum is suppressed, while the third peak is enhanced relative to the second.

The third primary evidence for particle dark matter is gravitational lensing, when light is deflected as it passes by massive objects, including clusters of dark matter. Many such clusters have been observed through gravitational lensing, but the most famous example is 1E0657-558, commonly known as the Bullet Cluster. The Bullet Cluster is shown in Fig. 2.4 (reproduced from Clowe et al. [8]) and displays the merger of two galaxy clusters. The optical and X-ray signals were detected through the Magellan Telescope and Chandra Observatory, respectively. However, when gravitation lensing was performed, the reconstructed mass profile was inconsistent with these measurements. In general, the bulk of the mass was much further away from the center of the merger than optical and X-ray measurements showed. The conclusion is that most of the mass is not in the photon-emitting baryons, but in a separate source.

Astrophysical evidence also allows us to exclude dark matter as a Massive Astrophysical Compact Halo Object, or MACHO. MACHO theories describe dark matter as composed of non light-emitting objects like black holes, neutron stars,

brown dwarfs, and rogue planets. However, these have been excluded by gravitational microlensing [9, 10], measurements of the cores of dwarf galaxies [11], and constraints on dynamical heating of stars in compact ultra-faint dwarf galaxies [12]. We thus conclude that dark matter must be composed of a stable, fundamental particle.

2.2 Weakly Interacting Massive Particles

The fundamental dark matter particle is not a known particle in the Standard Model (SM). It cannot be electrically charged, allowing us to reject quarks, hadrons, charged leptons, and charged gauge bosons. It must be long-lived, so we can reject the Higgs and Z. Neutrinos are a reasonable hypothesis at first glance, but the constraint on their mass (<1.1 eV [6]) is too tight. They are so light that their velocities are too high to allow large-scale structure formation in the early universe [13]; this is synonymous with the requirement that dark matter should be “cold”¹.

An entirely new particle needs to be introduced into our database. One important parameter is its mass. The constraints on dark matter mass are incredibly loose. It should not be higher than the Planck mass of 10^{19} GeV/ c^2 , at least without introducing GUT-level physics or complex multi-particle objects. It should also not be lower than 10^{-21} eV/ c^2 , because otherwise dark matter would not provide a stable underpinning for galaxies; its de Broglie wavelength would be larger than the galaxy [14]. Thus, in order to discover the particle nature of dark matter, we need to first build models that are theoretically and experimentally compelling.

In this section, we discuss WIMPs: Weakly Interacting Massive Particles. These are hypothetical new additions to the Standard Model that couple to Standard Model particles through the weak interaction. The WIMP mass would be on the order of the W or Z mass, in the tens of GeV/ c^2 to tens of TeV/ c^2 range.

One of the arguments for the WIMP is the so-called “WIMP miracle”. We follow the procedure in [15]. Suppose a new stable massive particle χ exists. Suppose further that χ was in thermal equilibrium with the known SM particles in the early universe, but when the universe cooled, its self-annihilation cross-section dropped sufficiently that the particle essentially stopped annihilating. This particle would still be present today, with some relic density having “frozen out.” One can calculate this relic density:

$$\Omega_\chi h^2 = \frac{3 \times 10^{-27} \text{cm}^3/\text{s}}{\langle \sigma v \rangle}, \quad (2.1)$$

where $\langle \sigma v \rangle$ is the thermally-averaged annihilation cross-section. We know from Planck data that $\Omega_c h^2 = 0.12055$, where Ω_c is the non-baryonic contribution to

¹A “sterile” neutrino—i.e., a fourth-generation heavy neutrino that participates in oscillation but not in the weak interaction—could be an important constituent of dark matter, though.

the universe's matter energy density. So if we want χ to explain dark matter, we should expect $\langle \sigma v \rangle = 2.5 \times 10^{-26} \text{ cm}^3/\text{s}$. But we also know that a particle with weak-scale interactions should have an annihilation cross-section:

$$\langle \sigma v \rangle \approx \left(\frac{\alpha}{100 \text{ GeV}} \right)^2 = \alpha^2 (10^{-21} \text{ cm}^3/\text{s}) \quad (2.2)$$

Therefore, the new particle χ can be both a WIMP and the dominant constituent of dark matter if $\alpha \approx 1/200$, very plausible for a weak interaction. WIMPs are also the favored model of many physicists because they are possible to detect experimentally via nuclear recoils, as we will discuss in detail in Sec. 2.4.

Most of the Z^0 -mediated WIMP parameter space was excluded by early dark matter experiments, but some of this parameter space remains untested [16]. A particle can also be mediated by other processes and still fall under the WIMP umbrella. The Higgs boson (and its hypothetical cousins that appear in many supersymmetric theories) is another particle that can serve as the portal between the Standard Model and dark matter particle(s). The Higgs-mediated DM parameter space has mostly been excluded by recent experiments, similarly to the parameter space for a Z^0 mediator. However, there are still some viable Higgs-mediated models where the WIMP mass is very close to half the Higgs mass or greater than a few TeV/c^2 [17]. Currently and in the near future, most experiments probe models where tree-level interactions are forbidden, and the dominant contribution is at the loop level. This can occur, for example, if dark matter is a fermion and the portal is a pseudoscalar [18]. In such models, the nuclear recoil cross-section is suppressed to levels within the sensitivities of upcoming experiments.

Two dominant constraints exist on the mass of a WIMP. First, Lee and Weinberg calculated that weak-interacting particles could not explain dark matter if the particle mass was $< 2 \text{ GeV}/c^2$, because the annihilation cross-section would be too high to produce the relic density observed today [19]. Second, WIMPs cannot be heavier than about 340 TeV because of unitarity [20].

2.3 Models of Dark Matter at the MeV Scale

Since WIMPs have been unobserved in 30 years of searching, we should also probe alternate dark matter models.

Let us only keep the assumption of thermal production in the early universe, along with the astrophysical requirements of a stable, gravitationally-interacting particle. We remove the assumption that the DM-SM mediator is a known electroweak particle. Then the DM mass can be as low as keV/c^2 , and it can be either a fermion or boson [21]. The mass constraint here comes from the previously-discussed observations of large-scale structure. On the high end, the unitarity bound still restricts the mass to a few hundred TeV/c^2 . The added parameter space is commonly

called *sub-GeV dark matter*. In this dissertation, we also use the term *MeV-scale dark matter*. This is done to avoid confusion with sub-keV/c² models of ultra-light bosonic dark matter, most famously the QCD axion [6]. Furthermore, the experiments we discuss are generally only sensitive to MeV/c²-scale particles, not thermally produced keV/c² particles.

What types of dark matter can exist at the MeV scale? The short answer is: lots. We can come up with a plethora of theories with new MeV-mass particles and corresponding mediators. Probably the most widely-discussed theory is an SM extension with an extra U(1) symmetry, corresponding to a new particle called the “dark photon.” This is kinetically mixed with the SM photon, allowing a portal between the SM and DM particles. The Lagrangian vacuum terms are (copied from [21]):

$$\mathcal{L} \supset -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \frac{1}{4}V_{\mu\nu}V^{\mu\nu} + \frac{\kappa}{2}F_{\mu\nu}V^{\mu\nu} + \frac{1}{2}m_V^2V_\mu V^\mu + eA_\mu J_{EM}^\mu + g_\chi V_\mu J_D^\mu \quad (2.3)$$

Above, $F_{\mu\nu}$, A_μ , J_{EM}^μ , and e are the standard electromagnetic terms. Similarly, $V_{\mu\nu}$, V_μ , J_D^μ , and g_χ are the corresponding terms for the dark photon. The third term represents kinetic mixing, where κ is the strength of mixing. In some theories, the dark photon is the dark matter itself, but more commonly, it is the mediator between the Standard Model and dark matter. As the mediator, the dark photon mass is quite unconstrained, while the dark matter mass is still restricted to the MeV scale.

2.4 The Nuclear Recoil Signal

2.4.1 Direct Detection of Nuclear (and Other) Recoils

We need to test these models experimentally. A few channels are available, but in this dissertation, we focus primarily on *direct detection*. The general principle of direct detection is shown in Fig. 2.5. We build a detector out of some Standard Model particles. These serve as targets for the dark matter in the Milky Way halo. As dark matter particles pass by the detector, they have some probability of interacting with the terrestrial particles, if there is a SM-DM portal. This will cause some transfer of energy, which can be captured by a sufficiently sensitive detector.

For WIMPs and MeV-scale dark matter, we primarily search for nuclear recoils, in which the target is an atomic nucleus. We can use essentially any atom, or even complex molecules or crystals. Technically, the dark matter interacts with individual nucleons, but the nucleus behaves coherently if the momentum transfer is low enough. This constraint is met if the recoil energy is less than about 10 keV, which is always true for low-mass nuclei and low-mass dark matter. The constraint can fail for high-mass nuclei and DM particles, in which case the maximum deposited

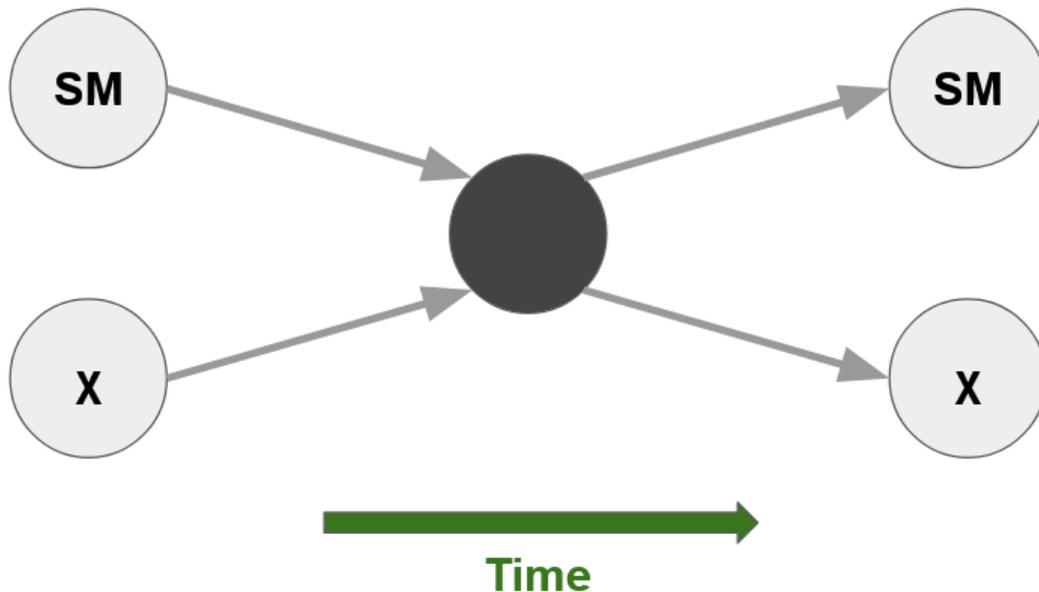


Figure 2.5: Toy Feynman diagram showing the principle of direct detection. A dark matter particle χ and a Standard Model particle collide, interact via some unknown mediator, and scatter. Typically, the SM particle is an atomic nucleus, but it could be an electron, a molecule, or something else.

recoil energy is shifted downwards. In this dissertation, when discussing nuclear recoils, we focus solely on interactions that are independent of the nuclear spin. In addition, the DM couplings to the proton and the neutron are assumed to be identical. Each dark matter model is identified by two parameters: the particle mass and DM-nucleon scattering cross-section. However, one can imagine spin-dependent interactions as well. In fact, an effective field theory approach allows us to consider 16 fundamental interactions [22, 23].

In general, nuclear recoils are an attractive signal because relatively few backgrounds induce nuclear recoils (for example, see [24, 25, 26]). *Backgrounds* are events in an experiment that look like the signal, but actually arise from some other source. The dominant nuclear recoil backgrounds are neutrons and coherent neutrino scattering, neither of which is a dominant issue for experiments in the next 10 years. The latter will be an obstacle in the long-term, as we discuss in Sec. 2.5. On the other hand, electronic recoils create a plethora of events in any dark matter search. Their sources are abundant: β radiation, γ radiation, and neutrinos. Fortunately, we can usually distinguish between electronic and nuclear recoils, based on the fraction of energy deposited into different channels: light, charge, and heat. The exact details depend on the target and detector technology. As a result, the relevant

background from electronic recoils is the events that accidentally contaminate, or *leak* into, the nuclear recoil region of interest. Discriminating between these interactions is critical to the success of any nuclear recoil search. It is also a dominant theme of this dissertation, discussed in every Chapter except Chapter 7.

Some experiments look for other dark matter signatures, such as DM-electron scattering. These searches are fruitful because they allow us to probe much lighter dark matter masses than we could with a heavy nucleus. Dark matter with a keV–TeV mass travels at non-relativistic speeds, so the reason is fundamentally Newtonian. Kinetic energy transfer is optimized when the colliding particles have similar masses—consider two baseballs colliding vs. a basketball and a tennis ball. On the other hand, conducting an electronic recoil search forces the experimenter to accept all the electronic recoil backgrounds. DM-electron searches are briefly discussed in this dissertation, but the focus is on DM-nucleus searches.

2.4.2 The Dark Matter Signal Model

In this subsection, we calculate the energy distribution of dark matter scatters. This is the signal model for all WIMP and MeV-scale DM nuclear recoil searches. Recoil energy is the most important observable; it is only dependent on properties of the target and the dark matter. We primarily follow the calculations of Lewin and Smith [27] but also use the works of Cerdeño and Green [28] and Schnee [29]. Other observables depend on the detector technology, so we delay discussion of those until the relevant Chapters.

To begin, we use basic scattering theory. Consider a collection of particles A , static in the Earth’s reference frame, with volume V_A , mass density ρ_A , number density n_A , and mass $m_A = \rho_A/n_A$. An infinite uniform stream of other particles B is incident on the collection, with mass density ρ_B , number density n_B , and mass $m_B = \rho_B/n_B$. The stream’s velocity is v , and the cross-section for interaction between the particles is σ . Then the rate of collisions, in collisions per unit time, is:

$$R = V_A n_A n_B v \sigma \tag{2.4}$$

The rate per unit mass of A is:

$$\tilde{R} = \frac{R}{V_A \rho_A} = \frac{n_B v \sigma}{m_A} \tag{2.5}$$

Now, let’s introduce dark matter detection. We identify A and B with the nuclear target N and the dark matter particle χ , respectively. We don’t know the dark matter mass, but we do know its mass density through astrophysical measurements. We also know that the dark matter particles in the galactic halo do not have identical velocities; they follow some distribution $f(\vec{v})$. Furthermore, we know that the cross-section is dependent on both the relative velocity and deposited energy E_R .

Therefore, we can calculate the collision rate, per unit target mass, per deposited energy:

$$\frac{dR}{dE_R} = \frac{\rho_\chi}{m_\chi m_N} \int v f(\vec{v}) \frac{d\sigma_N}{dE_R}(E_R, v) d^3\vec{v} \quad (2.6)$$

We are ignorant about the dark matter model and the exact form of $\frac{d\sigma_N}{dE_R}(E_R, v)$. However, in general, for WIMPs and MeV-scale dark matter, we can write:

$$\begin{aligned} \frac{d\sigma_N}{dE_R}(E_R, v) &= \frac{m_N}{2\mu_N^2 v^2} (\sigma_{0N,SI} F_{SI}^2(q^2) + \sigma_{0N,SD} F_{SD}^2(q^2)) \ , \\ \text{where } \mu_N &= \frac{m_N m_\chi}{m_N + m_\chi} \text{ is the reduced mass} \end{aligned} \quad (2.7)$$

and $q = \sqrt{2m_N E_R}$ is the momentum transfer.

We have separated the spin-independent (SI) and spin-dependent (SD) interactions. From here on, we ignore the spin-dependent interactions and drop the *SI* labels. In addition, we have isolated the velocity dependence as $\frac{1}{v^2}$. The form factor $F(q^2)$ describes the coherence of the nucleus, i.e., the ability for the nucleon wave functions to add in-phase. For typical momentum transfers in the range 1–100 MeV/c, full coherence is represented by $F = 1$, and $F \rightarrow 0$ corresponds to decoherence. Lastly, we have isolated the details of the WIMP-nucleus interaction to a single number σ_{0N} .

However, the cross-section σ_{0N} is for the entire nucleus. This is not helpful for comparing experiments with different target nuclei, so we instead seek the dark matter-nucleon cross-section σ_0 . To do this, we rewrite the cross-section in terms of the proton and neutron couplings f_p and f_n :

$$\sigma_{0N} = \frac{4\mu_N^2}{\pi} [Z f_p + (A - Z) f_n]^2 \quad (2.8)$$

We further assume $f_p \approx f_n$, giving us:

$$\sigma_{0N} = \frac{4\mu_N^2}{\pi} A^2 f_p^2 \quad (2.9)$$

In the case that we only have one nucleon:

$$\sigma_0 = \frac{4\mu_p^2}{\pi} f_p^2 \quad (2.10)$$

Here, μ_p is the DM-proton reduced mass (or the DM-neutron reduced mass, which is effectively the same). As a result, we can translate the DM-nucleus cross-section to the DM-nucleon cross section:

$$\sigma_{0N} = A^2 \left(\frac{\mu_N^2}{\mu_p^2} \right) \sigma_0 \quad (2.11)$$

The total nuclear cross-section is proportional to A^2 , as we would expect from coherence. This is a primary reason experimenters often prefer high-mass target nuclei for WIMP searches.

Putting these results back into our rate formula, Eq. 2.6, we find:

$$\frac{dR}{dE_R} = \sigma_0 \left[\frac{\rho_\chi}{2m_\chi \mu_p^2} \right] [A^2 F^2(q^2)] \left[\int \frac{f(\vec{v})}{v} d^3\vec{v} \right] \quad (2.12)$$

During a search, we seek to measure or constrain σ_0 . The first bracketed quantity is set by nature, the second is set by the experimenter, and the third is set by a combination of the two.

The first bracketed term in Eq. 2.12 depends only on the mass of the dark matter particle. We take $\rho_\chi = 0.3 \text{ GeV}/c^2/\text{cm}^3$, motivated by astrophysical measurements [30]. The second bracketed term depends only on the identity of the target nucleus. The form factor F is commonly taken to be the Helm form factor:

$$F(q) = 3 \frac{j_1(qr_N/\hbar)}{qr_N/\hbar} e^{-(qs/\hbar)^2/2} \quad (2.13)$$

$$j_1(x) = \frac{\sin x}{x^2} - \frac{\cos x}{x}$$

Here, r_N is the nuclear radius, s is the nuclear skin thickness, and $j_1(x)$ is the first spherical Bessel function. We use Lewin and Smith's [27] reported fits to various nuclei, which found $s = 0.9 \text{ fm}$ and $r_N = (1.14 \text{ fm}) A^{1/3}$. The Helm form factor is shown graphically in Fig. 2.6. It is unity at low momentum transfer, indicating robust nuclear coherence, but falls off rapidly for $qr_N/\hbar > 1$. Coherence is lost at lower recoil energies for heavier nuclei, e.g. a few keV for ^{131}Xe vs. hundreds of keV for ^4He . We note that technically, the form factor is nonzero even at the highest recoil energies, but its amplitude is dramatically diminished, as seen in the bottom panel.

The final bracketed term in Eq. 2.12 is the velocity phase space, which is based both on the dark matter and the target. Without knowing anything about the dark matter halo, we can determine the integrand limits. If a dark matter particle with kinetic energy $E_\chi = \frac{1}{2}m_\chi v^2$ hits a static nucleus, and the scattering angle is θ , the nuclear recoil energy is:

$$E_R = \frac{1}{2} E_\chi r(1 - \cos \theta) \quad (2.14)$$

$$r = \frac{4m_\chi m_N}{(m_\chi + m_N)^2}$$

For a given recoil energy, the initial velocity of the DM particle could range from v_{min} to ∞ , where $v_{min} = \sqrt{\frac{2E_R}{m_\chi r}}$. Thus, this term (which we dub P) can be written:

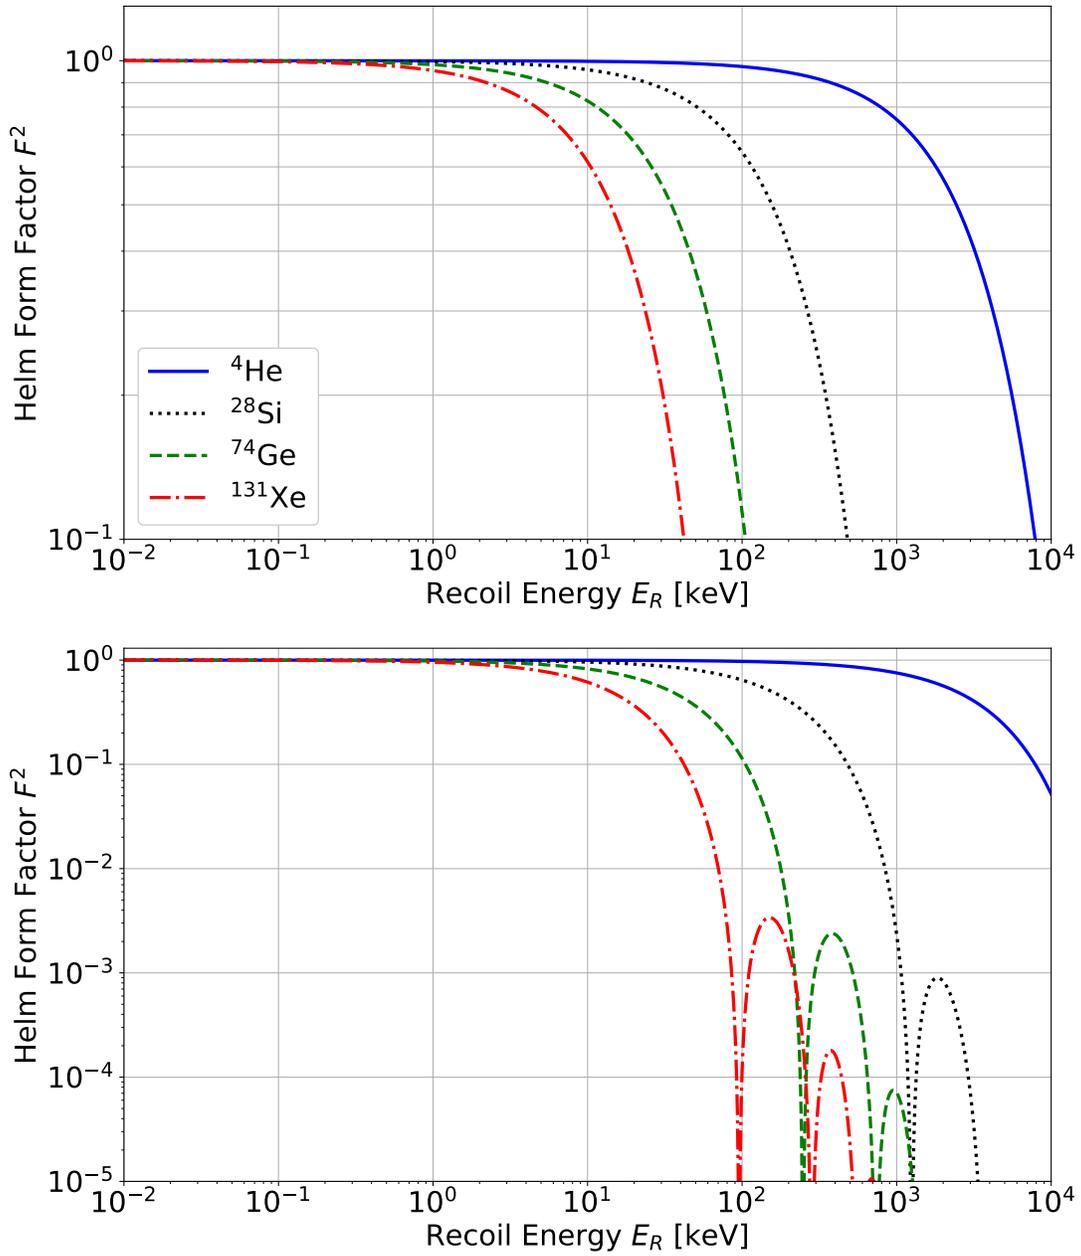


Figure 2.6: The Helm form factor F^2 for various nuclei. Momentum transfer is converted to recoil energy using $q^2 = 2m_N E_R$. The two panels are identical, but the bottom one is zoomed out to show the higher-order bounces of the Bessel function.

$$P \equiv \int_{v_{min}}^{\infty} \frac{f(\vec{v})}{v} d^3\vec{v} \quad (2.15)$$

Now, we can add galactic halo information. The velocity distribution we use is shown below in Eq. 2.16. It is a Maxwell-Boltzmann distribution in the halo's reference frame, with the most probable velocity v_0 . It is truncated above some velocity v_{esc} , the escape velocity of the galaxy, and N_{esc} is the corresponding unitless modification to keep the distribution normalized. The distribution is adjusted to represent velocities in the Earth's reference frame by adding the Earth-halo relative velocity \vec{v}_{lab} . In the Earth's frame, the dark matter halo "wind" is fastest in June and slowest in December.

$$f(\vec{v}) = \frac{1}{N_{esc} (\pi v_0^2)^{3/2}} \exp \left[-\frac{|\vec{v} + \vec{v}_{lab}|^2}{v_0^2} \right] \text{ if } |\vec{v} + \vec{v}_{lab}| < v_{esc} \quad (2.16)$$

= 0 otherwise

We integrate this distribution, assuming that over the course of a year, the angle between \vec{v} and \vec{v}_{lab} varies isotropically, i.e., $\cos \phi$ is uniform. The result is:

$$P = \frac{1}{N_{esc}} \frac{2}{\sqrt{\pi} v_0} \left\{ \frac{\sqrt{\pi}}{4} \frac{v_0}{v_{lab}} \left[\text{erf} \left(\frac{v_{min} + v_{lab}}{v_0} \right) - \text{erf} \left(\frac{v_{min} - v_{lab}}{v_0} \right) \right] - e^{-v_{esc}^2/v_0^2} \right\}$$

$$N_{esc} = \text{erf} \left(\frac{v_{esc}}{v_0} \right) - \frac{2}{\sqrt{\pi}} \frac{v_{esc}}{v_0} e^{-v_{esc}^2/v_0^2} \quad (2.17)$$

The magnitudes of these velocities can be obtained from astrophysics, but there are moderate uncertainties on them. We use the values in [31]. The modal DM velocity is $v_0 = 220$ km/s, the mean Earth velocity is $v_{lab} = 232$ km/s, and the galactic escape velocity is $v_{esc} = 544$ km/s. We plot Eq. 2.17 for several DM masses and target nuclei in Fig. 2.7. Above some recoil energy, P drops to zero because the minimum velocity to induce that recoil is higher than the galactic escape velocity. The higher the DM mass, the higher the cutoff. The trend with nuclear mass is more complicated; as the nuclear mass increases, the DM mass dependence becomes stronger. For example, suppose we change the DM mass from 1000 GeV/c² to 100 GeV/c². This causes the phase space to be reduced by a factor of five for a ¹³¹Xe target, a factor of three for ⁷⁶Ge, a factor of 1.5 for ²⁸Si, and barely changes the phase space for a ⁴He target. This trend is the mathematical form of our statement above: "kinetic energy transfer is optimized when the colliding particles have similar masses." For a 122 GeV/c² Xe nucleus, the difference in its ability to steal kinetic

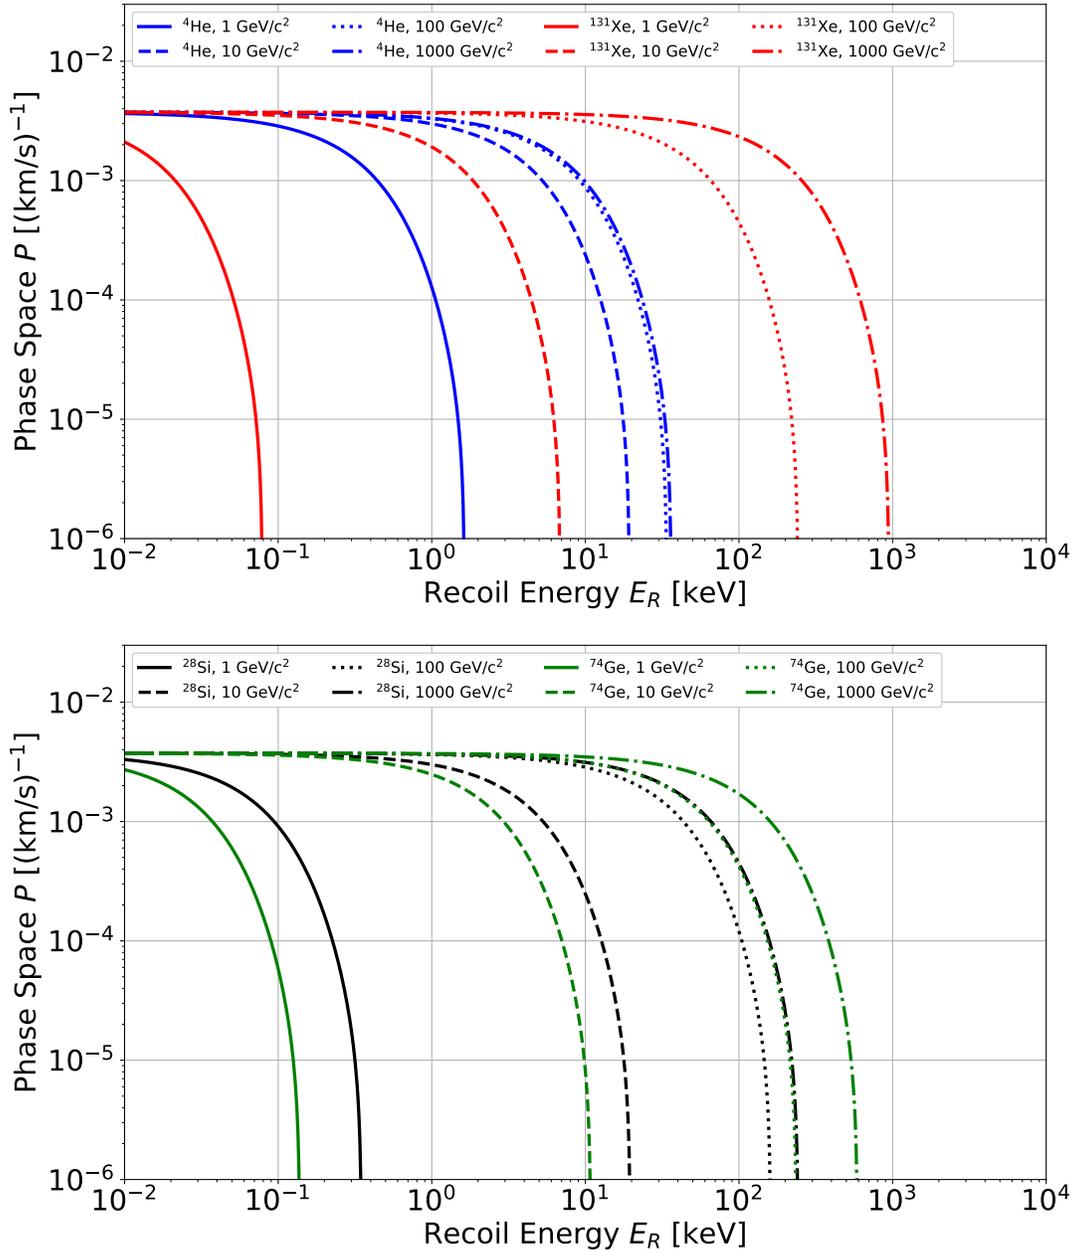


Figure 2.7: The velocity phase space term from Eq. 2.15 as a function of recoil energy, for several DM masses and nuclei.

energy from a $1000 \text{ GeV}/c^2$ or $100 \text{ GeV}/c^2$ particle is substantial. Meanwhile, efficient kinetic transfer is impossible for the ${}^4\text{He}$ nuclei in either situation, so the effect is muted.

Finally, we put it all together. We plot the nuclear recoil rate in Fig. 2.8 for several dark matter masses and target nuclei. The cross-section is kept constant. The rate is typically reported in units of $\text{DRU} = \text{kg}^{-1}\text{day}^{-1}\text{keV}^{-1}$.

We see several trends. First, at low energies, the rate is flat with recoil energy; then, it decreases until hitting some cutoff. This is consistent with we saw earlier in the energy dependence of F and P . Second, for a given DM mass and nucleus, the cutoff can come either from the form factor or from the phase space, whichever occurs at lower recoil energy. For example, a $1 \text{ GeV}/c^2$ particle on ${}^4\text{He}$ can deposit up to 2 keV , set solely by the available phase space. However, a $1 \text{ TeV}/c^2$ particle on ${}^{131}\text{Xe}$ can deposit up to 100 keV , primarily set by the form factor (for the first Bessel bounce); the velocity term allows depositions up to 1000 keV . Third, as DM mass increases, the recoil energy distribution shifts to higher energies, but the low-energy rate decreases. The latter effect is because the DM density in mass per volume is known, so as we increase DM mass, we must decrease number density. Fourth, for the highest of masses, the dark matter mass doesn't affect the recoil energy distribution and only affects the amplitude of the low-energy component. The distribution shape is generally dependent on the reduced DM-nucleus mass μ_N , which is independent of m_χ for $m_\chi \gg m_N$. As a result, changing the mass only changes the number density. Fifth, the trend with nuclear mass is identical to the kinematic efficiency previously discussed; each nucleus is optimized to measure dark matter near its own mass.

This last point can be well-illustrated by Fig. 2.9. We show, for a given nucleus, the relationship between dark matter mass and the maximum energy deposited. As the DM mass increases, it has more kinetic energy, so more energy can be transferred to the nucleus. However, we see crucially that the fraction of energy transferred is not monotonic. Instead, this fraction is maximized near the nuclear mass.

In this dissertation, we will describe direct detection searches for WIMP and MeV-scale dark matter. It is clear from this discussion that to identify dark matter nuclear recoils at the WIMP scale, we should choose heavy nuclei; at the MeV scale, we should choose light nuclei. Indeed, the experiments we will describe use xenon and helium targets, respectively.

2.5 The Neutrino Fog

One of the backgrounds in any nuclear recoil search is from coherent elastic neutrino-nucleus scattering, often abbreviated CEvNS. This is a process in which neutrinos scatter off an entire nucleus, inducing a recoil. Since these events have a nuclear

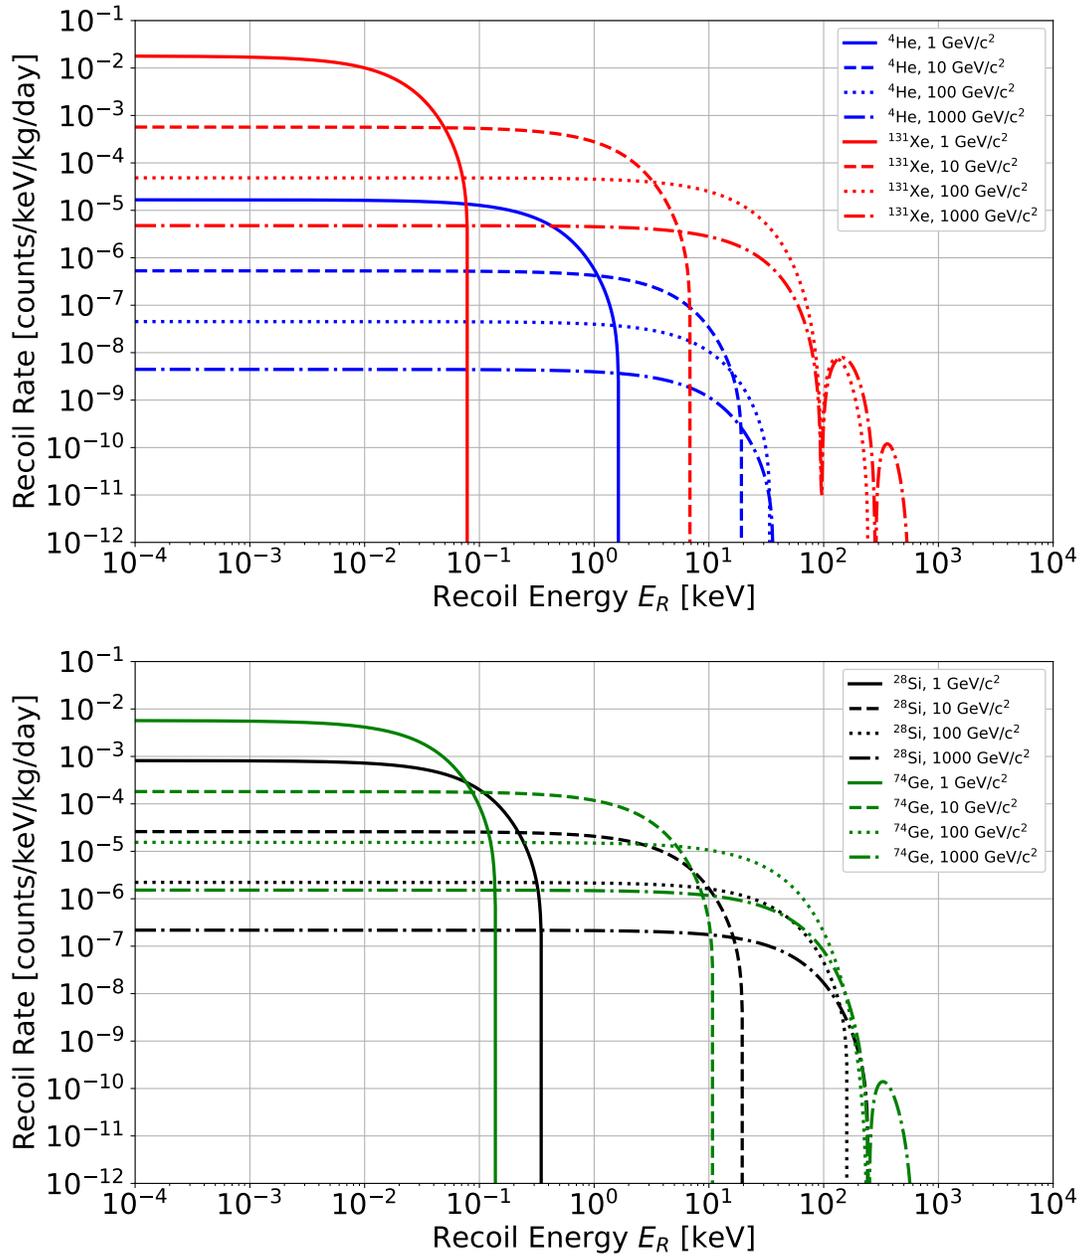


Figure 2.8: The dark matter-nucleus recoil rate as a function of recoil energy, for several DM masses and nuclei. The DM-nucleon scattering cross-section is fixed to 10^{-45} cm^2 .

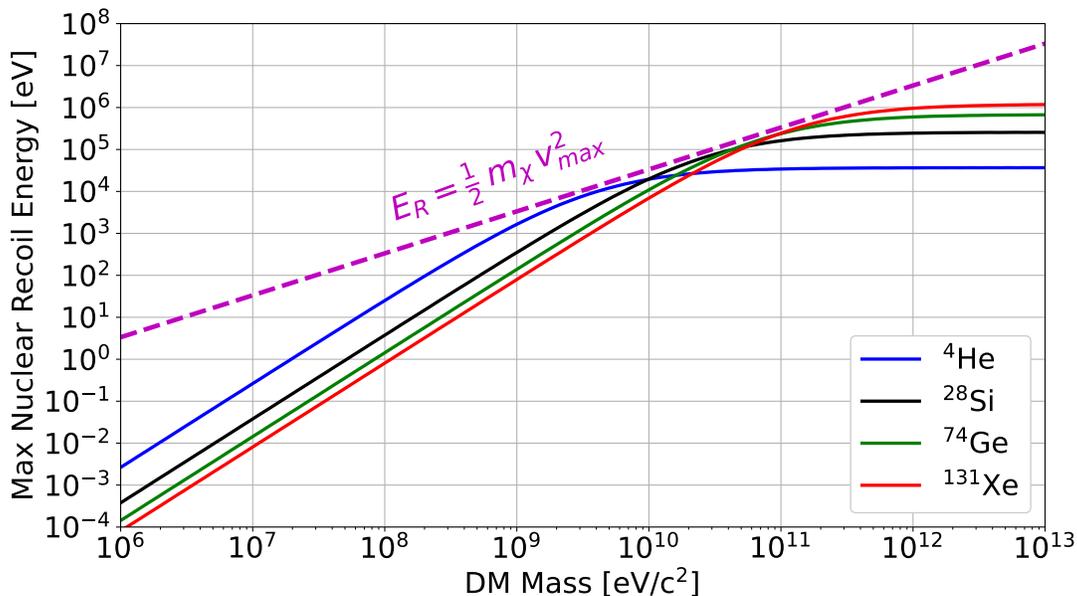


Figure 2.9: The maximum recoil energy deposited for a given nucleus-dark matter combination. We also show in magenta the maximum kinetic energy of a galactic halo particle, using $v_{max} \equiv v_{esc} + v_{lab}$.

recoil signature, they are indistinguishable from a dark matter recoil with the same recoil energy.

Neutrino NR backgrounds are minor for existing experiments. This is demonstrated in Fig. 2.10, where we show the total CEvNS rate on a ^{131}Xe target; the details of this calculation are described in Sec. 8.5. We overlay the recoil rates of several dark matter models, specifically using the models corresponding to the PandaX-4T 90% limit from 2021 [32]. The dark matter rates are noticeably higher, especially above the few-keV energy threshold of xenon-based experiments. This indicates that the CEvNS background is subdominant.

However, as dark matter experiments get larger, they will become sensitive to lower cross-sections. Simultaneously, the neutrino background will become a more pressing issue. This means that below a certain cross-section, the neutrino background will overwhelm any potential DM signal, and it will be impossible to measure these dark matter models. This phenomenon is called the *neutrino fog* or *neutrino floor*. It represents a line in mass vs. cross-section parameter space, below which dark matter models cannot be probed with direct detection techniques.

Some scientists distinguish between “floor” to refer to the actual line in mass vs. cross-section parameter space, and “fog” to refer to the phenomenon of limiting neutrino backgrounds. In this dissertation, we solely use the term “neutrino fog,”

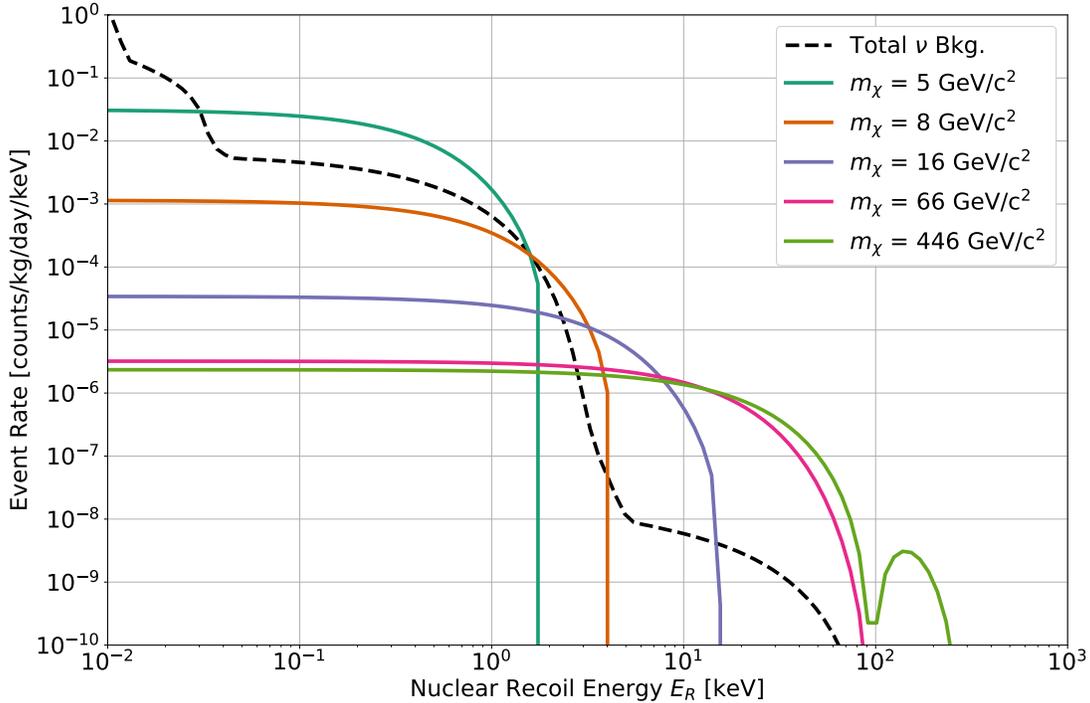


Figure 2.10: The rate of coherent neutrino-nuclear scattering for solar, atmospheric, and supernova neutrinos combined, overlaid with several rates of dark matter scattering. For all curves, the target nucleus is ^{131}Xe . For each DM mass shown, the cross-section used is the 90% exclusion limit of PandaX-4T [32].

because in fact, the line is not a clear floor. First of all, there are multiple definitions of the neutrino fog, and there is not consensus on what it means for a dark matter model to be undetectable; see [33] for a discussion of this. Second, it's important to realize that the neutrino fog is not actually a consequence of the neutrino rate. Dark matter models below the neutrino fog could potentially be probed with a large enough detector. For example, suppose a detector was so enormous that it observed 1000 CEvNS events over some period of time. Then, one could still observe DM particles at the 5σ level if there were at least $5\sqrt{1000} = 158$ events, even ignoring the different energy distributions. The required detector mass for a desired sensitivity σ would scale with $\sigma^{-1/2}$. The true problem that the neutrino fog describes is that the neutrino flux is uncertain at the few percent level. Therefore, for low enough DM cross-sections, it becomes impossible to determine whether an excess of events is due to dark matter or an upward fluctuation in the flux. In the example above, a flux uncertainty of 1.5% would make it impossible to conclude that 158 events of dark matter were observed. The third reason for using the term “fog” is that in fact, we can peer through it, given the right tools. A detector with the ability to measure the

direction of a recoil would be able to push below the neutrino fog. Solar neutrinos are the dominant neutrino background; they originate from the sun, the location of which is well-known. Atmospheric and supernova neutrinos are more diffuse in location but also lower-flux. Meanwhile, in the Earth’s reference frame, the dark matter wind points at us in a specific direction, given our motion through the galaxy. It comes from the location of the star Deneb in the Cygnus constellation [34]. In a directional detector, the different source locations would be a strong discriminant between signal and background.

The neutrino fog is described in more detail in Chapter 8, including how to calculate it. We show it in Fig. 2.11, using a combination of Billard’s calculation in xenon [31] and ours in helium.

2.6 Current Status of the Field

The current status of the field of dark matter direct detection at the keV/c^2 – TeV/c^2 mass scale is shown in Fig. 2.11. The solid black lines show the oldest and newest WIMP-scale limits from the Gotthard and LUX-ZEPLIN experiments [35, 26], respectively. Over the past 30 years, we have enhanced our sensitivity to WIMP-induced nuclear recoils by 12 orders of magnitude. We still have about two orders of magnitude available before reaching the neutrino fog. On the other hand, the MeV-scale parameter space is relatively untouched, and the few exclusions that exist have all been explored in the past few years.

In this dissertation, we will describe two classes of experiments that explore these regions of parameters space. First, we will spend four chapters discussing the dynamics of two-phase liquid/gas xenon time projection chambers—in particular, how they distinguish between nuclear recoils and electronic recoils. This technology is the leading candidate to search for WIMPs. Second, we will use two chapters to describe calorimetric technologies to probe MeV-scale parameter space. We will describe R&D efforts on instrumentation and readout, and then we will focus on superfluid helium as a dark matter nuclear recoil target. Helium is the second-lightest element, making it the prime candidate to search for dark matter at the lowest masses.

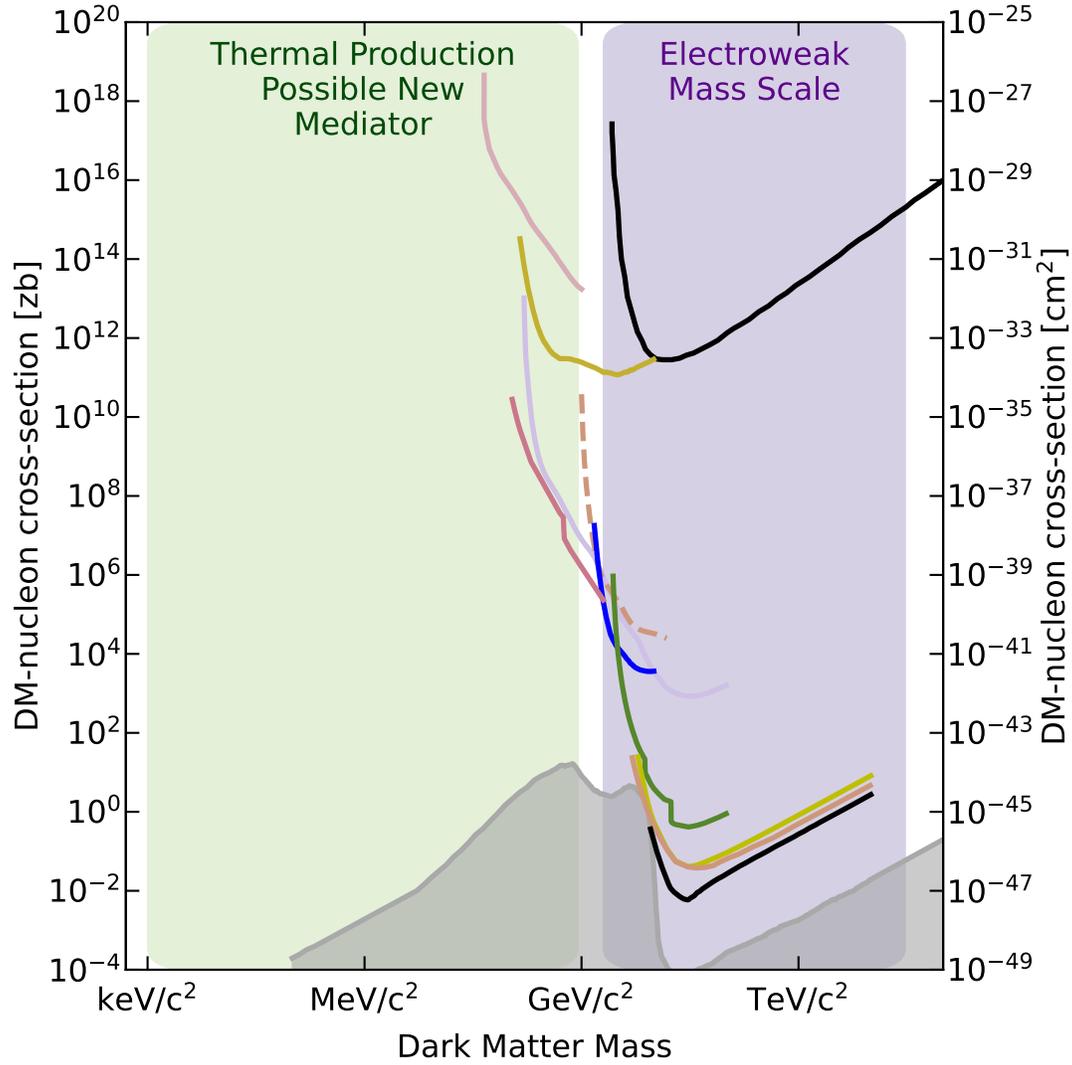


Figure 2.11: Current world-leading limits. The mass scales for electroweak interactions (WIMPs) and a new mediator with thermal production are highlighted. Limit curves are from [35, 36, 37, 38, 39, 40, 41, 42, 43, 32, 26]. Neutrino fog calculation is from [31] and Chapter 8 of this dissertation.

Chapter 3

Particle Detection in Liquid Xenon

In this chapter, xenon is described as a medium for particle detection. The history of xenon and other noble elements for particle detection is summarized. Its properties are explored, with a particular emphasis on the properties that make it a suitable target for dark matter. Energy partitioning between the channels of light, charge, and heat is discussed. We explore dynamics of two-phase liquid/gas time projection chambers, the dominant technology used for dark matter detection. Finally, the Noble Element Simulation Technique, a computational model for simulating interactions in liquid xenon, is introduced.

3.1 History of Xenon in Particle Physics

Liquid noble elements have been used for particle detection since at least 1953 [44], when John Marshall at MIT measured ionization induced by β particles in liquid argon. In the late 1960s and early 1970s, scientists at Lawrence Berkeley National Laboratory and the Space Sciences Laboratory introduced the idea of liquid xenon, and they were able to measure α and γ particles in xenon-filled single-wire proportional chambers and multiple-wire ionization chambers [45]. Over the next several decades, the field of noble element particle detection burgeoned, as physicists around the world characterized liquid argon, xenon, and other noble elements. Two excellent reviews of this history were written by Claude Brassard in 1979 [46] and Tadayoshi Doke in 1993 [47], and the reader is encouraged to explore them.

Meanwhile, the use of xenon detectors for WIMP detection has a history stretching across nearly three decades. The first published discussion of this strategy appears to have been in 1993 [48]. A group of physicists—mostly at Italian universities, but also at CERN and UCLA—proposed the use of liquid xenon scintillation for WIMP detection and built a prototype detector. They simultaneously measured the scintillation and ionization signals from α and γ sources, using the same “S1” and “S2” terminology that is prevalent today to refer to the respective signals. They also found a varying S2/S1 ratio between the particle types; this ratio will be a major subject of this dissertation.

Over the rest of the 1990s, the concept matured, and the ZEPLIN-I project was underway by the turn of the century [49]. This experiment was a UK-based single-phase liquid xenon detector that used scintillation pulse shape to separate potential WIMPs from β and γ backgrounds. In 2005, ZEPLIN-I measurements were used to set the first WIMP nuclear scattering constraints from a xenon experiment [50]. Simultaneously, physicists were advocating for the advantages of two-phase liquid/gas detectors [51, 52, 53, 54].

The two-phase liquid/gas xenon time projection chamber (xenon TPC) is the primary technology discussed in this dissertation. The basic features are displayed in Fig. 3.1. A central liquid xenon chamber serves as the primary target volume, and a small layer of gaseous xenon exists on top of it; these are the two phases in

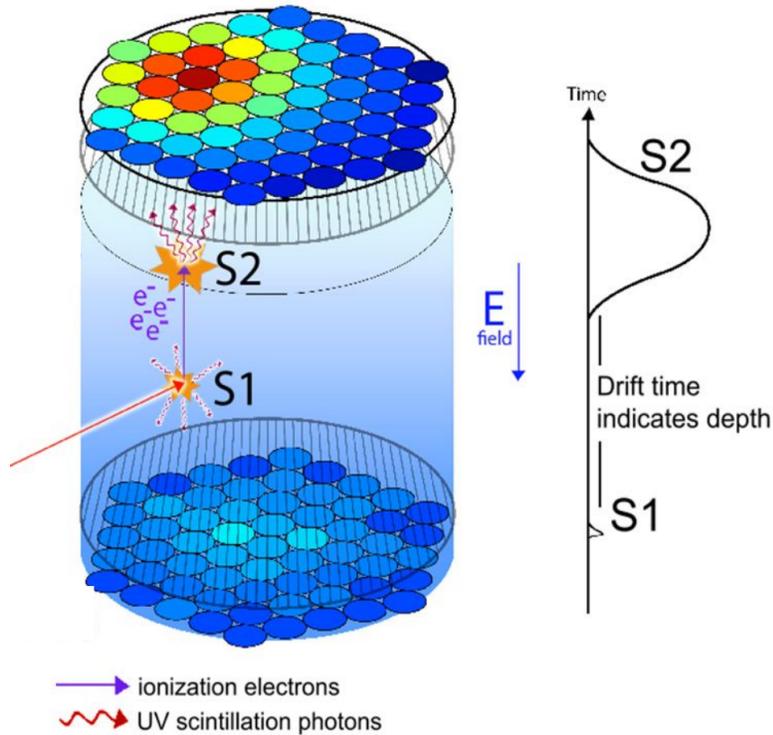


Figure 3.1: A cartoon schematic of the Large Underground Xenon (LUX) detector; it is generally applicable to any TPC. Created by Carlos Faham.

the name. Photomultiplier tubes (PMTs) with single-photon sensitivity are located at the top and bottom of the detector. A TPC for a WIMP search may have up to hundreds of PMTs. A downward-pointing electric field (the “drift field”, typically hundreds of V/cm) is established in the central volume by means of a cathode grid at the bottom of the detector and a gate grid immediately below the liquid surface. A stronger downward-pointing field (the “extraction field”, typically thousands of V/cm) is set up in the gas region between the gate grid and a higher anode grid. Each energy deposition in the liquid xenon above a certain threshold produces two observable signals. First, scintillation photons are detected promptly by the PMTs, giving a signal called S1. Second, ionized electrons drift upwards through the electric field, are extracted into the gas phase, and produce additional scintillation in the gas. The resulting photons are detected by the same PMTs, giving a signal called S2. The combination of S1 and S2 signals allows experimenters to reconstruct the deposited energy and the location of the deposition. The z -position is obtained through the time between the S1 and S2 pulses. The x, y position is obtained through the hit pattern on the top PMTs. Section 3.4 describes the dynamics of this technology in more detail.

In 2007, both the XENON10 and ZEPLIN-II collaborations succeeded at building

xenon TPCs and setting the first WIMP limits from this technology [55, 56]. Since then, the community has dedicated vast resources (on the scale of hundreds of millions of dollars) to studying and exploiting xenon TPCs for WIMP searches. ZEPLIN-III, located at the Boulby mine in the United Kingdom and a successor to ZEPLIN-II, set WIMP nuclear recoil limits in 2009 and 2011 [57, 58]. XENON100, located at Laboratori Nazionali del Gran Sasso (LNGS) in Italy and a successor to XENON10, did the same in 2010, 2011, and 2012 [59, 60, 61]. LUX, located at the Sanford Underground Research Facility (SURF) in South Dakota, U.S.A., continued the effort by releasing constraints in 2013, 2015, and 2016 [62, 63, 64]. In parallel, the PandaX-I experiment at the China Jinping Underground Laboratory set limits on WIMP interactions in 2014 and 2015 [65, 66]. Its successor, PandaX-II, did the same in 2017 [67]. The next experiment in the XENON program, XENON1T, published its results in 2018 [43]. Similarly, the next experiment in the PandaX program, PandaX-4T, set a limit in 2021 [32]. Finally, as of writing, the most recent WIMP limit is from LUX-ZEPLIN (LZ), an experiment located at SURF and run by scientists who were previously associated with the LUX and ZEPLIN projects. The LZ limit is world-leading and was released in July 2022 [32].

The experiments listed in the previous paragraph have all been primarily designed to search for WIMPs, and the majority have set world-leading constraints on WIMP-induced nuclear recoils. The author has attempted to provide an exhaustive list of xenon WIMP detectors. However, the same experiments have been used for other physics searches. These include searches for MeV/c²-scale particle dark matter [41], axion-like particles [68, 69, 70], mirror dark matter [71], and double electron capture [72]. In addition, a plethora of smaller detectors (typically at the scale of <1 kg) have been built to study the dynamics of xenon TPCs. These include, but are not limited to, MaXe [73, 74], PIXeY [75], XeD [76], neriX [77], and Xurich [78]¹. The previously mentioned WIMP detectors have also been utilized for R&D purposes.

At present, there are three collaborations operating tonne-scale xenon TPCs. This generation of experiments is often called “Generation 2” or simply “G2”, to note its succession beyond experiments at the 100 kg scale—i.e., LUX and XENON100. These are LZ, with a “fiducial” volume (the xenon volume in which we search for nuclear recoils) of 5500 kg [26]; XENONnT, with a fiducial volume of 4370 kg [68, 82]; and PandaX-4T, with a sensitive volume of 2670 kg [32]. All three collaborations have released results—LZ and PandaX-4T have released limits on DM-nucleon

¹The astute reader, and even the unastute reader, will note physicists’ eagerness to name experiments with clever acronyms. I have pondered the natural question, “Why?”, for several years and participated in this process myself. My conclusion is that it is a fact of nature, no more or less fixed than the electron’s mass of 511 keV/c². Other scientists have explored the question in more detail [79, 80]. Yet, it seems physicists are not alone in this proclivity; *The Atlantic* has noted the same tendencies amongst members of the United States Congress [81]. We can only hope that large physics collaborations do not adopt the filibuster.

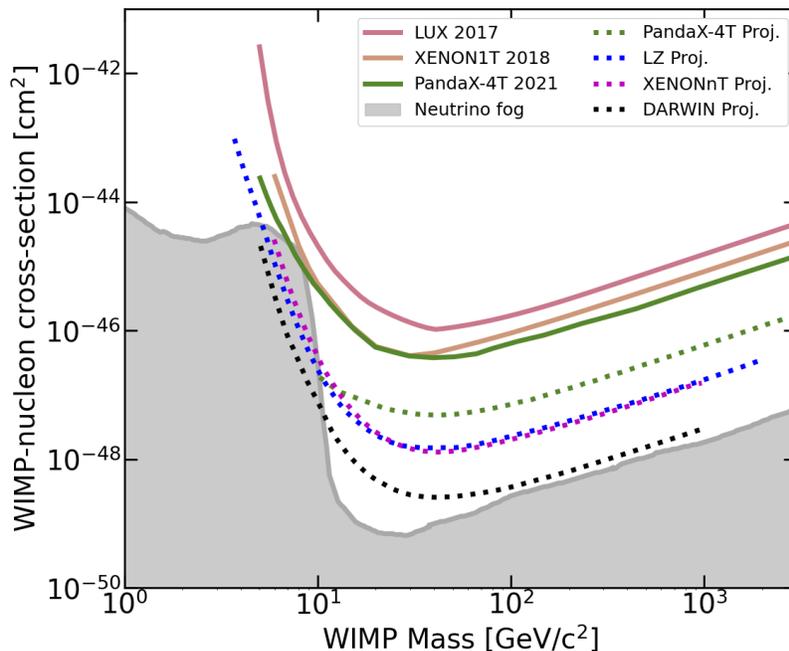


Figure 3.2: The current status of liquid xenon dark matter experiments, represented as 90% exclusions on spin-independent nuclear recoils. The solid lines represent constraints from existing experiments [64, 43, 65], while the dashed lines represent projected sensitivities [85, 86, 82]. The neutrino fog from [31] is also shown.

scattering, while XENONnT has performed a search for electronically interacting solar axions.

These three experiments will make substantial progress in exploring parameter space towards the neutrino fog, but they will not reach the neutrino fog. For this, we need a detector containing on the order of 100 tonnes of liquid xenon. This ambition is often called “DARWIN” after an early proposal [83] or simply “G3” [84].

Figure 3.2 summarizes the current and future landscape of xenon dark matter experiments. This dissertation will touch on all generations of this technology. In Chapter 4, an analysis of data from LUX (a G1 experiment) is presented, which informed the operating electric drift field of LUX-ZEPLIN (a G2 experiment). Then, in Chapter 6, we present an analysis of data from PIXeY, an R&D-scale detector; this analysis provides useful input for the construction of a potential G3 experiment.

3.2 Properties of Liquid Xenon

Xenon is a noble element with 54 protons. It has several naturally-occurring isotopes, and its standard atomic weight is 131.293 amu [90], indicating that the aver-



Figure 3.3: Photographs of xenon, taken with the Xenon Breakdown Apparatus experiment (XeBrA) [87]. (*Left*) Bubbling in liquid xenon. (*Right*) Solid xenon accumulated at the bottom of the chamber. Note that in a properly functioning dark matter detector, neither bubbles nor solid xenon should be present in the liquid xenon chamber. Images screenshotted from videos taken by Lucie Tvrznikova [88, 89].

age xenon atom has about 77 neutrons. The heavy atomic weight is one of xenon’s strengths as a dark matter target, since the interaction rate for spin-independent interactions is proportional to A^2 . Xenon is transparent in all states; Fig. 3.3 shows photos of liquid and solid xenon.

The boiling point of xenon is 165 K. Its density is 2.9 g/cm^3 [91] (technically, liquid xenon’s density depends on its temperature, but xenon’s melting point is only 4 K below its boiling point at 1 atm, so this is a small effect). The high density of xenon is an advantage for a dark matter search. First, it allows us to pack a large number of targets (i.e., nuclei) into a small space. Second, it allows for “self-shielding.” In a liquid xenon chamber, the outer layers of xenon will block most radioactivity from reaching the interior volume. Therefore, by reconstructing the 3D position of an interaction (as we can do in a TPC), we can focus our analysis on the quiet inner region; this process is called “fiducialization.”

The naturally occurring isotopes of xenon are ^{124}Xe , ^{126}Xe , ^{128}Xe , ^{129}Xe , ^{130}Xe , ^{131}Xe , ^{132}Xe , ^{134}Xe , and ^{136}Xe . The corresponding abundances of these isotopes are 0.095%, 0.089%, 1.910%, 26.4%, 4.017%, 21.23%, 26.91%, 10.44%, and 8.86% [92]. This distribution of isotopes presents another advantage for xenon as a dark matter target. Most of these isotopes are radio-stable, so they will not create harmful backgrounds in the interior xenon volume. One notable exception is ^{136}Xe , which decays via two-neutrino (and potentially neutrinoless) double beta decay. The total energy carried by the outgoing electrons ranges from 0 to 2.5 MeV, giving rise to an important background for G2 experiments like LZ [86]. The isotope ^{134}Xe also decays via double beta decay, with an expected endpoint of 825.8 keV. However, both the two-neutrino and neutrinoless modes are rare enough that neither has

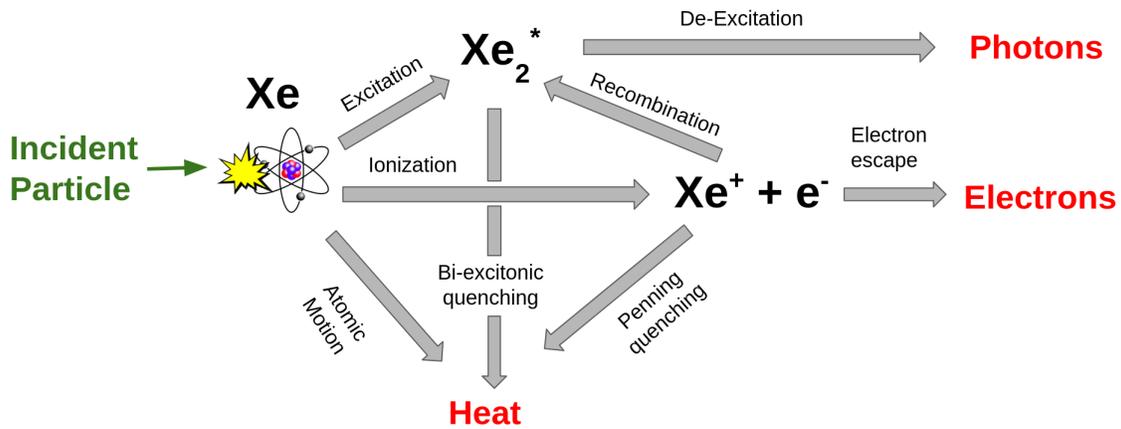


Figure 3.4: A generic picture of energy partitioning in liquid xenon. See text for details.

been observed [93], so this is not an important background. Meanwhile, ^{124}Xe is radioactive via double electron capture, but this process deposits 64.3 keV of energy in the surrounding xenon, which is outside the WIMP search range [72]. In addition to the natural isotopes, there are other isotopes of xenon that can be produced cosmogenically. Most of these decay within a matter of hours or a few days, but ^{127}Xe has a 36.4-day half-life, meaning it can be present for months in the detector. Some of the ^{127}Xe electron-capture decay energies are in the WIMP search region, so it is a background for dark matter experiments like LZ, albeit one with fairly low impact.

3.3 Energy Partitioning in Liquid Xenon

The plethora of experiments listed in Sec. 3.1, as well as others that were not mentioned, have not only searched for WIMP dark matter; they have also sought to understand how energy is deposited in liquid xenon. Here, we describe the current understanding of this process. In this section, we describe energy deposition generically; in the next section, we will explore this phenomena specifically in two-phase time projection chambers.

3.3.1 Three Channels of Energy Deposition

Figure 3.4 presents a generic picture of how energy is deposited in liquid xenon—and, indeed, in liquid noble elements more broadly. We can read it from left to right. An incident particle strikes a xenon atom. To start, we are agnostic to whether the collision was with the xenon nucleus or an atomic electron. The recoiling xenon

atom dissipates its energy (typically, on the order of 1-100 keV in this dissertation) through three channels. First, if the initial interaction is a nuclear recoil, a substantial amount of energy goes into atomic motion, i.e., heat. The xenon nucleus carries a great deal of kinetic energy, and it can collide with other nuclei and induce high velocities in them. Second, regardless of the initial interaction, energy will be transferred into ionized electrons. The initial recoil either produces a high-charge, high-speed xenon nucleus or a high-speed electron. Either way, this charged object will race through the xenon and ionize surrounding atoms. The emergent delta rays can ionize more atoms themselves, and so on. Third, energy will be transferred to excited xenon atoms. This is analogous to the ionization process, except that the atomic electrons simply jump into a higher shell, rather than being totally ionized. Notably, in noble elements, an excited atom will find a neutral atom very quickly (sub-ns, perhaps sub-ps) and dimerize:



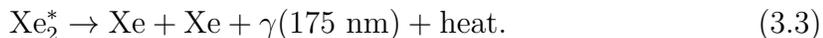
Thus, we can consider the energy as being deposited into Xe_2^* “excimers.” Note that in the literature, the term “excitons” is also used, but this usually refers to excited atoms Xe^* , or to the quanta created by the primary interaction. We will use “excimer” to refer to the molecule, and “exciton(s)” otherwise.

At this stage, we think of the xenon as having some number of excitons N_{ex} and some number of electron/ion pairs N_i . Often, we refer to the ratio $\alpha \equiv N_{ex}/N_i$. We also refer to the total number of quanta N_q :

$$N_q \equiv N_{ex} + N_i. \quad (3.2)$$

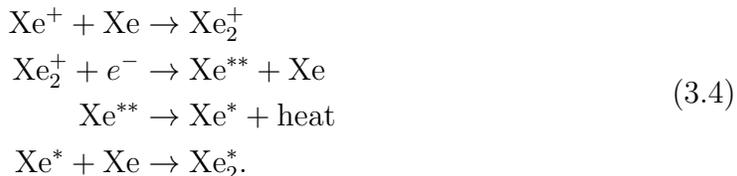
We ignore the initial heat, since it is impossible to measure this channel in a detector at 165 K, but we keep it in mind, as it will be important later.

The excimers de-excite, quenching more energy into the heat channel and producing 175 nm photons [94]:



We see that the dimerization step is critical, because it allows production of light that the xenon itself is transparent to. The de-excitation will occur on one of two timescales, corresponding to the excimer state. Singlet excimers $^1\Sigma_u^+$ de-excite with a time constant of 3.3 ns, while triplet excimers $^3\Sigma_u^+$ de-excite with a time constant of 24 ns [95]. The total decay time distribution follows a sum of two falling exponentials with these time constants, weighted by the initial fraction of singlet and triplet states. These fractions will be discussed later, as they offer a tool for particle identification and discrimination.

Meanwhile, the ionized electrons can have one of two fates. They may simply escape the site of the recoil. This is facilitated by an electric field, as deployed in a TPC. Alternatively, they can undergo recombination, as depicted:



The net process from Eq. 3.4 is simply:



We see that recombination effectively trades an electron/ion pair for an excimer and quenched heat; the excimer then de-excites by Eq. 3.3. If the ionized electron recombines with its own parent ion, this is called “geminate” recombination.

Recombination is typically parameterized through two steps: recombination probability and recombination fluctuations. The phrases are apt. The recombination probability r is the probability that a given ionized electron will recombine with any ion. Recombination fluctuations provide a measure of the statistical variation in recombination. Although it can be defined in many ways, the most common is σ_R , which has units of quanta. For example, suppose $N_i = 100$, $r = 0.4$, and $\sigma_R = 10$. Then the distribution of the number of electrons that recombine will have a mean of 40 electrons and a width of 10 electrons.

After recombination, we are left with the same number of quanta N_q (and additional energy lost to heat), but rearranged into photons and escaped electrons:

$$\begin{aligned}
N_{ph} &= N_{ex} + rN_i \\
N_e &= (1 - r)N_i \\
N_{ph} + N_e &= N_{ex} + N_i
\end{aligned}
\tag{3.6}$$

The reader should carefully note that in this dissertation and the xenon literature, N_e refers to the number of outgoing electrons after recombination, not the total number of electrons produced. In Eq. 3.6, we have made an implicit assumption here that every exciton, whether produced directly or via recombination, produces exactly one outgoing photon. This assumption seems naively true, but it is difficult to confirm experimentally. If this efficiency were non-unity, it would simply be absorbed into the W , g_1 , and g_2 experimental parameters discussed below.

The total energy deposited in a liquid xenon interaction can thus be summarized

$$E = fW(N_{ph} + N_e),\tag{3.7}$$

where $1 - f$ is the fraction of energy lost to heat, and W is the average energy associated with the production of one quantum, either an electron or photon. Interestingly, for purely electronic interactions, W seems to be constant [96]. All of the other steps mentioned so far are dependent on the nature of the interaction and any

applied electric field, as we will discuss soon. However, W has been observed to be independent across a wide range of energies (1 keV to several MeV) and field (10s to 1000s of V/cm). This is in spite of the fact that the ionization and excitation energies are different, and the N_{ph}/N_e ratio varies substantially by energy and field! The constant nature of W indicates that the amount of energy required to create either a photon or electron is the same, after combining the energy lost directly to the quanta and indirectly to heat via de-excitation. It also indicates that the fraction of energy lost to heat is constant for electronic interactions (this fraction has never been measured, however). We can simply absorb the loss of heat into the definition of W . To be explicit, we take $f = 1$ for electronic recoils. For nuclear recoils, $f < 1$. This will be explored in more detail in Sec. 3.3, but typically $f \approx 0.15$ for nuclear recoils.

Although W is empirically constant, its value is presently the subject of controversy. Dahl measured $W = (13.7 \pm 0.2)$ eV in 2009 [76], and this was used as the standard value for a decade. However, in 2019, the EXO-200 collaboration measured $W = (11.5 \pm 0.5_{\text{sys}} \pm 0.1_{\text{stat}})$ eV [97], and in 2021, Baudis et al. measured $W = 11.5^{+0.2}_{-0.3}(\text{sys})$ eV [98]. The discrepancy is not yet understood. One hypothesis is a poorly-calibrated g_2 (the gain associated with the measurement of each electron) by Dahl, but this explanation has not yet been taken as authoritative by the community. The effects of a potentially incorrect W value are complicated. Many experimental parameters are linearly reliant on W , so an incorrect value for W would simply require a rescaling in those parameters. This rescaling would not be noticeable in data. However, in simulations like NEST (see Sec. 3.5), changing W would change many backend dynamics, particularly relating to statistical fluctuations in detected signals. Conversations around the value of W are ongoing, but in this dissertation, we will use $W = 13.7$ eV unless otherwise stated. Much of the work was completed before the EXO-200 measurement, and the vast majority was completed before the Baudis et al. measurement.

3.3.2 Electronic Recoils

As described in Chapter 2, the dominant backgrounds in WIMP searches are electronic recoils (ER). To illustrate the dynamics of ERs, Fig. 3.5 shows how energy is partitioned for a 10 keV recoil. This energy is within the WIMP search region, although somewhat on the higher end of it. ER dynamics are strongly dependent on the applied electric field (in a TPC, this field is used to drift ionized electrons), so we choose a field of 200 V/cm, similar to the 2013 run of LUX. This figure was made using Xenimation, a Github repository created by the dissertation author to easily generate animations of energy partitioning in xenon. The width of each arrow is proportional to the amount of energy going into each channel. Xenimation relies on NEST [99], a simulation package we will describe in Sec. 3.5, to calculate the number and types of quanta produced. One can think of NEST as an average of

world data. The dissertation author is an active member of the NEST Collaboration that produces these models and simulation tools.

Let us consider the 10 keV recoil, step-by-step. Recall that no primary energy is lost to heat, but a constant fraction of energy is lost to heat via de-excitation. This fraction is unknown, but for the sake of the figure, we set it as $1 - (10 \text{ eV} / 13.7 \text{ eV}) = 0.27$, approximating that each exciton and electron/ion pair was created using 10 eV of energy.

The remaining 73% of energy goes into quanta production. For each keV of energy deposited, 73 quanta are produced (the inverse of $W = 13.7 \text{ eV/quantum}$). In electronic recoils, significantly more electron-ion pairs are created than excitons. Many experiments have observed that $N_{ex}/N_i \approx 0.2$ [101, 102, 103], differing from the theoretical value of 0.06. NEST employs a model in which $N_{ex}/N_i = 0.18$ at high energies and goes to zero at the lowest energies; $N_{ex}/N_i = 0.09$ for $E = 10 \text{ keV}$.

Then, we consider recombination. For the 10 keV recoil, $r = 0.59$, so an average of 401 electron-ion pairs recombine out of the 678 pairs created; this is the total of geminate and non-geminate recombination. Although recombination fluctuations are not shown in the figure, their magnitude is modeled as $\sigma_R = 34$ pairs. This is significantly higher than what one would calculate from binomial statistics: $\sigma_R = \sqrt{r(1-r)N_i} = 10$ pairs. At first glance, this is confusing—an ionized electron can either recombine or not, so binomial statistics should be accurate. However, every liquid xenon experiment has observed super-binomial recombination fluctuations. For some examples, the reader can see Refs. [76, 96] and Chapter 5 of this dissertation. This phenomenon is not fully understood as of writing, but work is ongoing to explain it [104].

After recombination, we are left with 277 escaped electrons and 465 excimers. Theory and measurements [95, 105] indicate that most of these are slow-decaying triplets, and we model that only 8% are fast-decaying singlets.

At this stage, we are done with the “microphysics” of electronic recoils; everything else about the interaction is based on the exact nature of the detector. Yet before proceeding to nuclear recoils, it is worth pausing. We have only discussed ERs for a specific field and energy. How do the dynamics vary across a wider range of settings?

See Fig. 3.6 for the answer. The top panels show the light and charge yields (the number of photons and electrons emerging from the interaction site, per unit energy) as a function of recoil energy and electric field. One immediate trend is that as we increase field, the charge yield increases, and the light yield decreases. This is easy to explain; when the electric field increases, electrons are pulled away from the recoil site, and recombination is suppressed. The same effect is seen directly in the lower-left panel. Another trend is that the light yield is maximized at 40–60 photons/keV for energies in the tens of keV. This trend is more complicated (and also more interesting) to explain. It is partly because in NEST, the ratio of N_{ex}/N_i

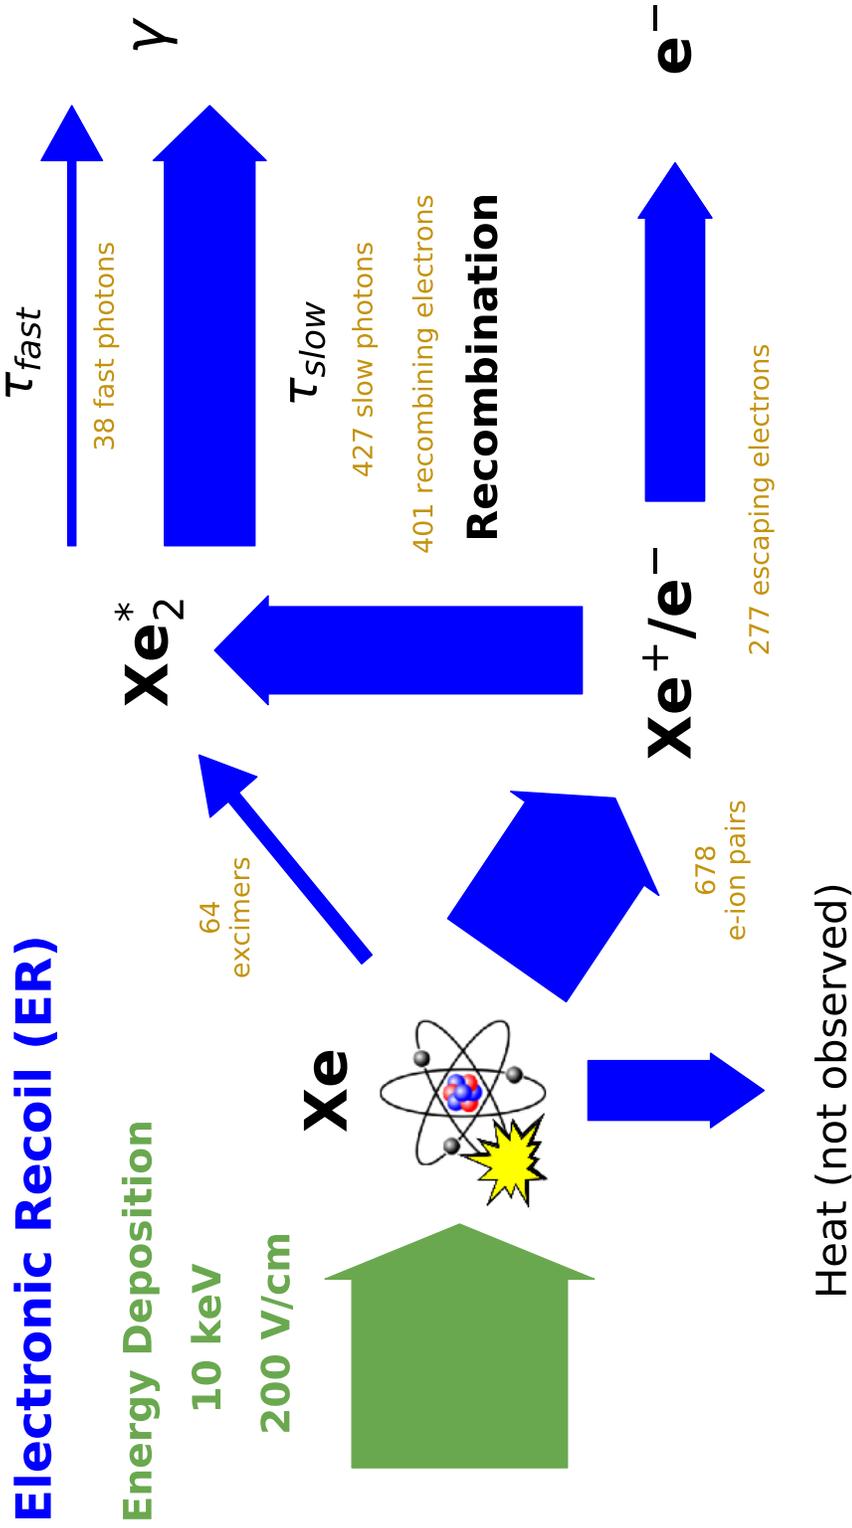


Figure 3.5: The energy partitioning of a 10 keV electronic recoil with an applied field of 200 V/cm. The width of each arrow is proportional to the amount of energy in the corresponding channel. Created with Xenimation v1.0.0 [100], using NESTpy v1.5.0 and NEST v2.3.0. [99].

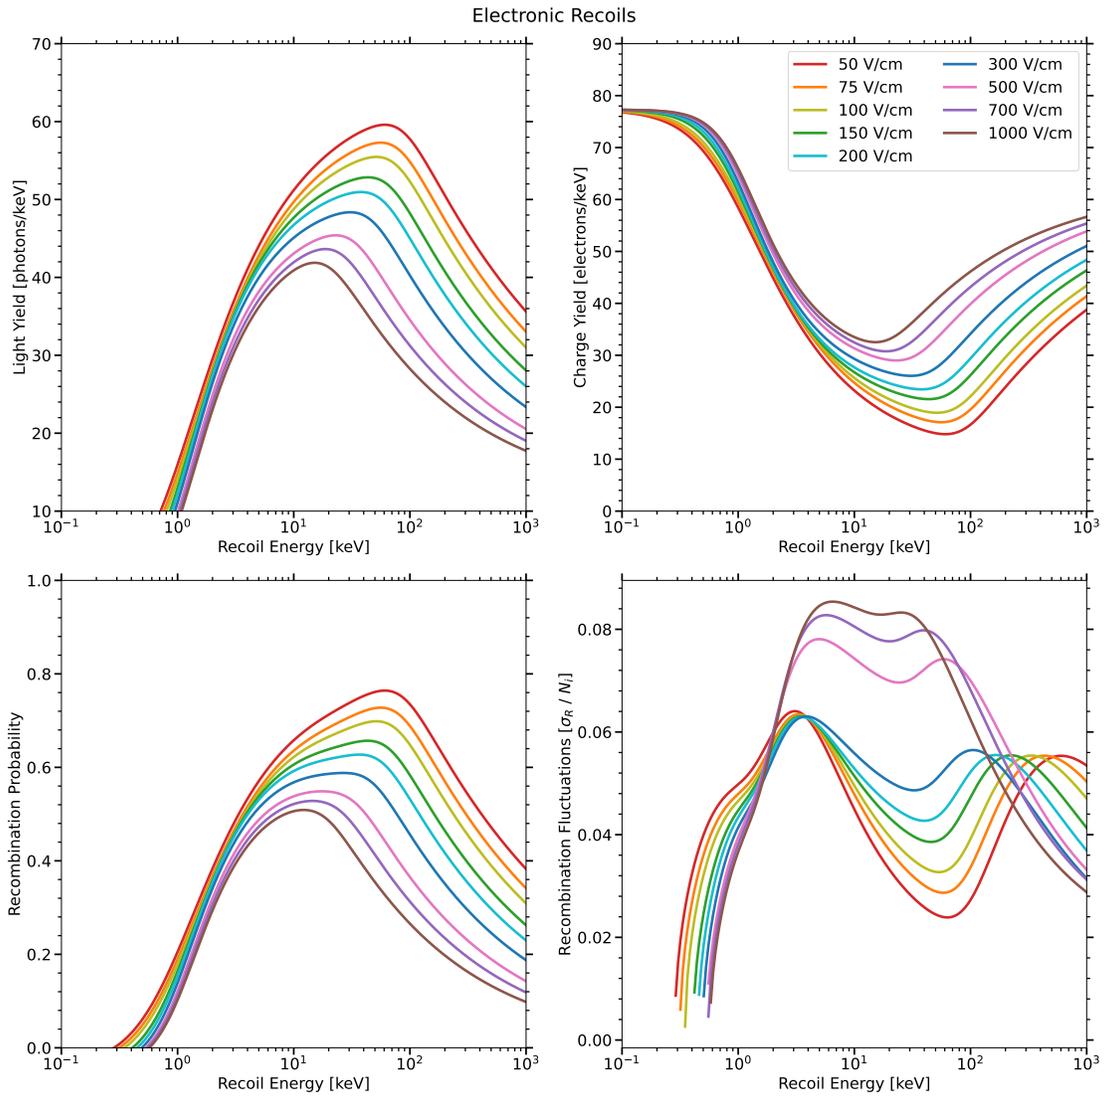


Figure 3.6: The light yield, charge yield, recombination probability, and recombination fluctuations for electronic recoils in liquid xenon at a variety of electric fields and recoil energies. The light and charge yields are post-recombination. Plots made using the beta model from NESTpy v1.5.0 and NEST v2.3.0 [99].

increases with energy below 30 keV. However, the more important contributor is that recombination is maximized at moderate energies, as seen in the lower-left panel. Our current understanding of this maximization is that there are two recombination “regimes.” At high energies, recombination for ERs follows the “Doke-Birks” model, and at low energies, it follows the “Thomas-Imel Box” model.

The Doke-Birks model relates the recombination probability to the ionization density in a track dE/dx . It is based on Birks’ Law, which relates the scintillation yield dL/dx to the ionization density by

$$\frac{dL}{dx} = \frac{A dE/dx}{1 + B dE/dx}. \quad (3.8)$$

Doke et al. extended this law by accounting for geminate recombination with a constant term [106]. Szydagis et al. extended it further by translating the scintillation yield to a probability and noting that at nonzero field, the constant term is best-fit to zero [107]. The full model is

$$r = \frac{k_B \text{LET}}{1 + k_B \text{LET}}, \quad (3.9)$$

where LET is the linear energy transfer, $\text{LET} = (dE/dx)/\rho$, and we have only one free parameter k_B . In plain English, the model says that with higher ionization density, there will be more recombination. Ionization density is negatively correlated with energy, so recombination decreases when recoil energy increases.

At low energies, the Doke-Birks model predicts recombination should be close to 100%, because dE/dx is high. In fact, the opposite is empirically observed; recombination goes to zero at low energy. To explain this, we employ the Thomas-Imel Box (TIB) model in this energy range (typically below 20 or so keV). This model was originally developed by Thomas and Imel in 1987 [108] to explain observations of charge collection vs. field that were incompatible with the geminate-only Onsager theory. It was then successfully modified and deployed by Dahl in 2009 to explain low-energy electronic interactions [76]. The TIB model states that recombination occurs in a box of size a around each ion; essentially, the electrostatic ability of the ion to attract electrons and force them to recombine is smeared on a length scale a . The value of a varies based on the applied electric field but falls in the range of a few hundred to a few thousand nm. The model succeeds if the total size of the ER track is smaller than a , which is true for low-energy depositions. Mathematically, the model is expressed as

$$r = 1 - \ln \frac{1 + \xi}{\xi}, \text{ where} \quad (3.10)$$

$$\xi = \frac{\alpha N_i}{a^2 v 4}.$$

Here, α is the recombination rate per unit volume, and v is the drift velocity, but only the combination of parameters $\frac{\alpha}{a^2v}$ is a free parameter. Its value is typically fit around 0.03. The greater the number of ions and electrons within the box, the greater the probability that a given electron will find an ion and recombine. Thus, recombination should be correlated with the number of ions produced in an interaction, consistent with the observed low-energy behavior.

NEST is consistent with the Doke-Birks and Thomas-Imel models in the high-energy and low-energy limits, as seen in Fig. 3.7.

Another trend seen in Fig. 3.6 is how recombination fluctuations vary with energy and field. Recall that the behaviors shown in the lower-right panel are purely empirical, not based on theory, but they are straightforward to explain. The fluctuations increase when the recombination probability approaches 50% from either direction. This makes sense, since $r = 0.5$ corresponds to the most uncertainty in whether an electron will recombine.

One last ER behavior that is invisible in Fig. 3.6 is a potential discrepancy between β and γ energy depositions. This discrepancy is not universally accepted, but J. Balajthy of the NEST Collaboration noticed that the yields from β and γ radiation were slightly different [109]. This phenomenon is not understood and is somewhat strange, since both interactions create a recoiling electron. One hypothesis is that in a β interaction, the incident electron is still left over and has a noticeable effect on the track. Regardless of the reasoning, NEST contains both a “beta” and “gamma” model. In this dissertation, the beta model is used unless otherwise stated.

3.3.3 Nuclear Recoils

Nuclear recoils (NR) exhibit a few different behaviors from electronic recoils. As we did for ERs, we take a characteristic recoil—again 10 keV at 200 V/cm—and display its specific partitioning in Fig. 3.8. One immediate difference is that a much larger fraction of energy is lost to heat, about 81%. As a consequence, only 139 quanta are created, contrasted with 742 quanta for the ER. Second, the partitioning between excitons and electron-ion pairs is much more equal for NRs than for ERs. This is based on several experimental measurements showing $N_{ex}/N_i \approx 1$ for NRs [76, 111, 112]. Third, the recombination probability is much lower, only 23%, due to the lower number of ions created. Fourth and finally, the ratio between singlet and triplet excimers is different; significantly more excimers are produced in the fast singlet state for NRs, about 23%.

We also can examine the effect of changing energy and field on nuclear recoils, done in Fig. 3.9. Here, there are some clear similarities and differences compared to the same dynamics for ERs. One similarity is that as we increase field, the charge yield increases and the light yield decreases, for the same reason described earlier; recombination is suppressed. However, the reader should also note that the magni-

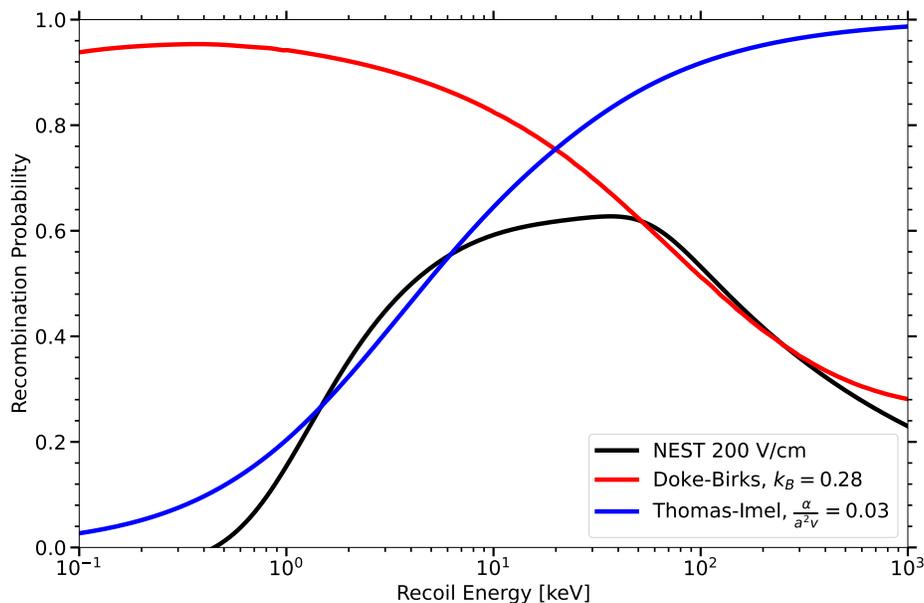


Figure 3.7: A comparison of the NEST v2.3.0 beta model to the Doke-Birks and Thomas-Imel models. The Thomas-Imel model contains one free parameter, set to 0.03. The Doke-Birks model contains one free parameter, set to 0.28, and we use the ESTAR database to obtain LET as a function of energy, using a polynomial extrapolation below 10 keV, the minimum energy in ESTAR [110].

tude of the trend is significantly smaller. This is a well-established fact for nuclear recoils in liquid xenon that every experiment has seen; these interactions are mostly field-independent. We also see that recombination increases monotonically with energy. Nuclear recoil tracks are very small, much smaller than the Thomas-Imel box size. As a result, they do not enter the Doke-Birks regime at all, and recombination is governed only by the TIB model. Regarding fluctuations in recombination, we again see that they are maximized when $r \approx 0.5$, but they are also high at low energies. This is simply because very few electron-ion pairs are created, so the relative size of the fluctuations is high.

There are two ER vs. NR differences that we have not explained: the N_{ex}/N_i ratio and the singlet-to-triplet ratio. Unfortunately, universally accepted explanations for these differences do not exist. However, they are likely linked. Excimers created via primary interactions likely have different singlet and triplet fractions than excimers created via recombination, because of the varying degrees of freedom of angular momentum. Direct excitation almost exclusively creates singlet excimers to conserve angular momentum. When an electron and ion recombine, they are some distance apart and a relative spin flip can generate triplet excimers. An electron can also recombine with a non-parent ion, leading to the same effect. Electronic recoils

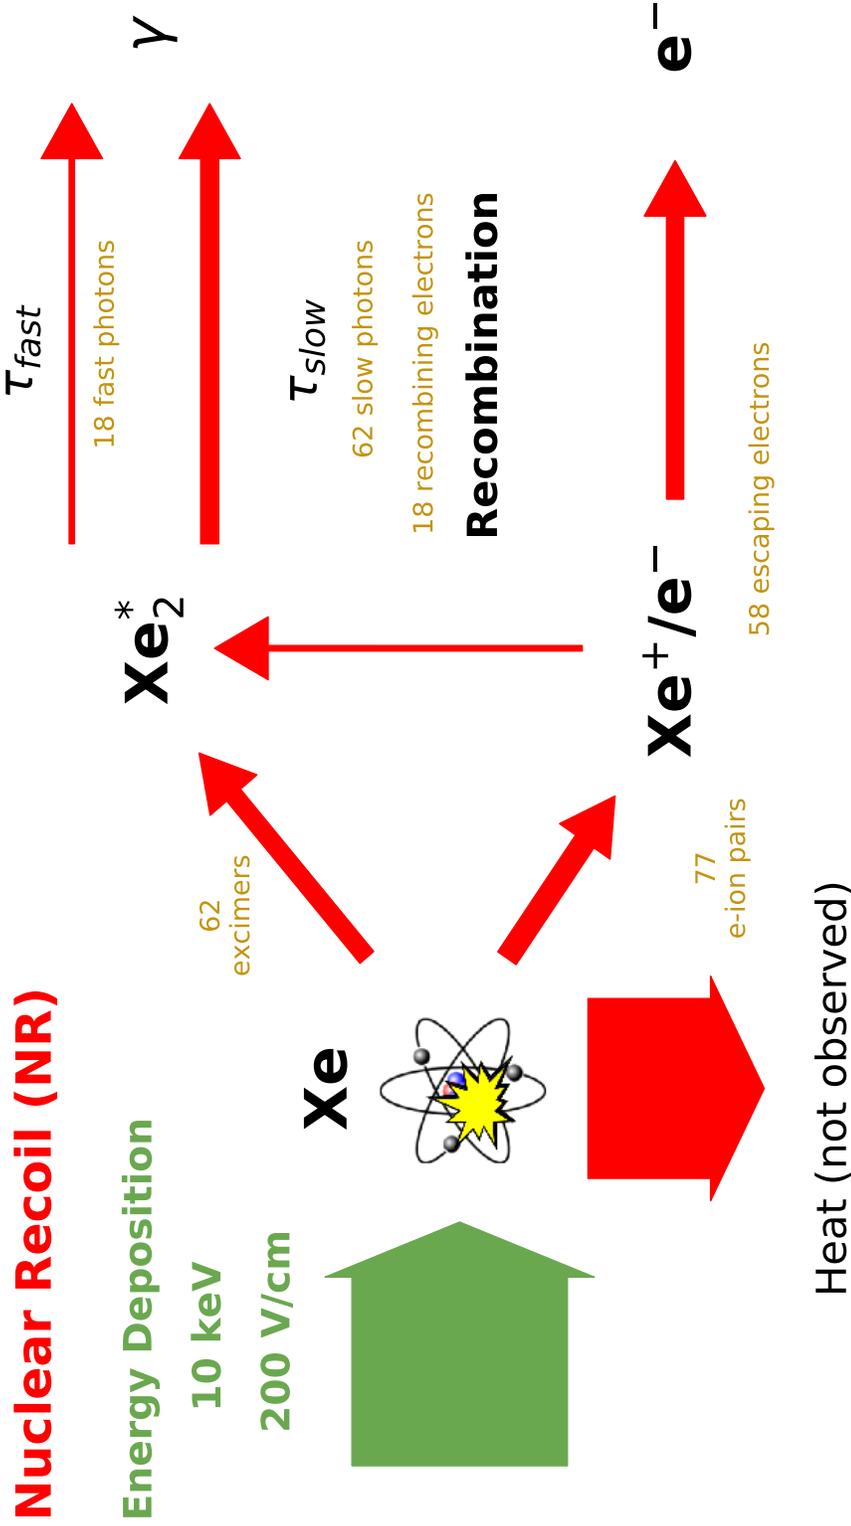


Figure 3.8: The energy partitioning of a 10 keV nuclear recoil with an applied field of 200 V/cm. The width of each arrow is proportional to the amount of energy in the corresponding channel. Created with Xenimation v1.0.0 [100], using NESTpy v1.5.0 and NEST v2.3.0. [99].

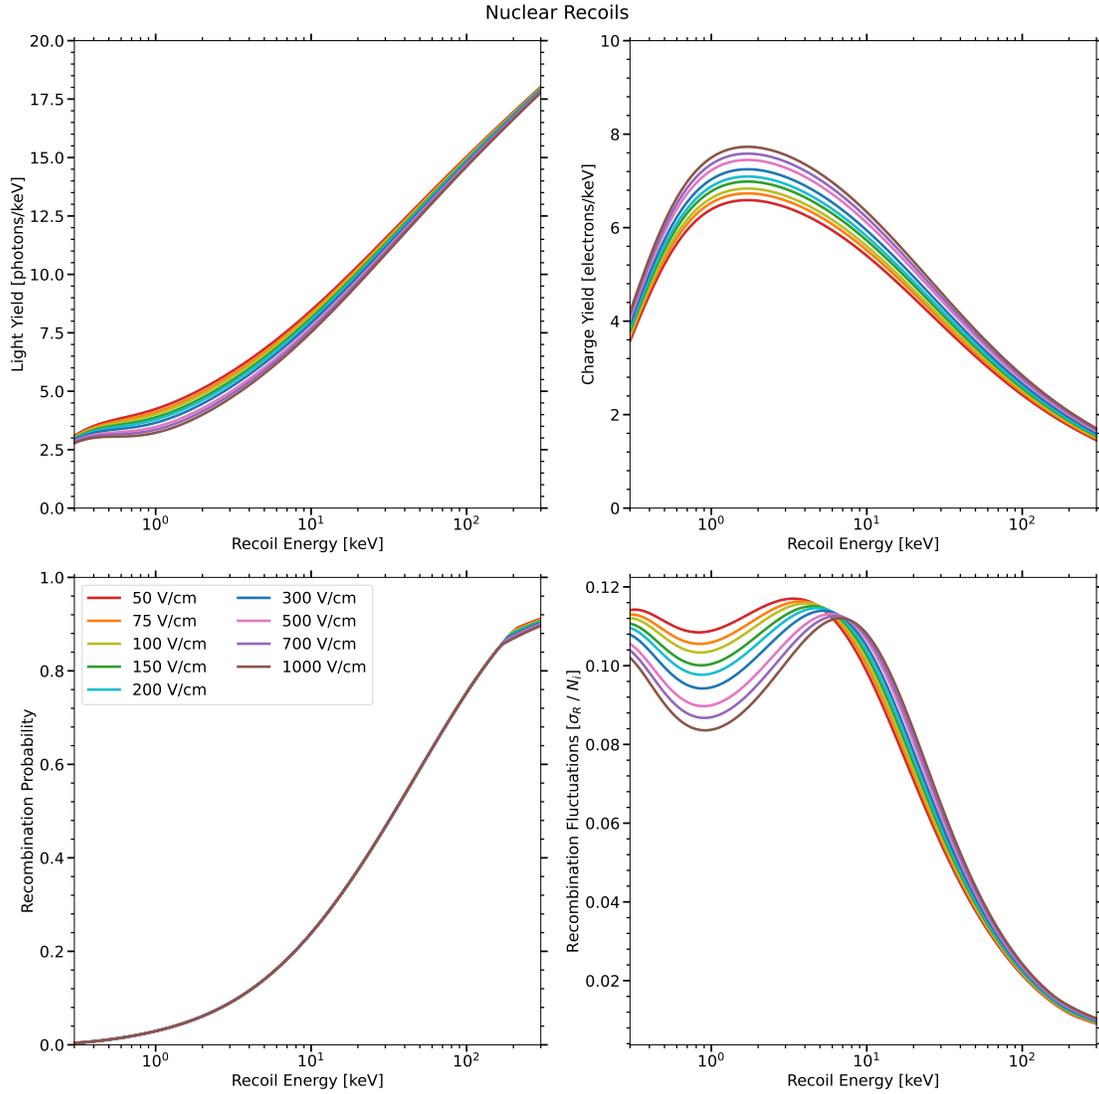


Figure 3.9: The light yield, charge yield, recombination probability, and recombination fluctuations for nuclear recoils in liquid xenon at a variety of electric fields and recoil energies. The light and charge yields are post-recombination. Plots made using the NR model from NESTpy v1.5.0 and NEST v2.3.0 [99].

produce more recombination and thus higher triplet fractions than nuclear recoils. In addition, argon physicists explain the singlet-to-triplet ratio discrepancy (which is much more dramatic in argon) through NRs’ higher ionization density [113]. The same phenomenon should exist in xenon. Remember that we are in the Thomas-Imel regime, so ionization density doesn’t affect total recombination probability, but it affects the fraction of geminate recombination and, accordingly, the creation of singlets and triplets.

Finally, we should discuss more carefully the biggest difference between the types of recoils: energy lost to heat. We have explained that NRs lose significantly more energy to heat, and this fraction varies with energy, but specifically how much heat is lost? This is a central question to any dark matter search. How could we possibly detect a WIMP if we don’t know much observable energy it deposits? Many models have been built to answer this question. The fraction of recoil energy in observable electronic modes (i.e., outgoing photons and electrons), relative to the fraction of energy in these modes for an ER of the same recoil energy, is called L . This is shorthand for “Lindhard factor.” $L = f$ from Eq. 3.7, and by definition $L = 1$ for ER events. One of the models to determine L in nuclear recoils is the “Lindhard model” [114, 115], somewhat unfortunate terminology since it is not the only model, nor even the best one at present. NEST deploys its own model for L that fits the world data better. Another model was developed to be compatible with LUX data; this model will be used frequently in this dissertation, especially in Chapters 4 and 6. All three models are depicted in Fig. 3.10.

3.3.4 ^{83m}Kr Decays

One unique interaction that is a common calibration source in xenon TPCs is the decay of ^{83m}Kr , which is worth discussing in its own section. ^{83m}Kr decays via a two-step process depicted in Fig. 3.11. The first step occurs with a 1.83 hr half-life and releases 32. keV of energy; the second step proceeds shortly after, releasing 9.4 keV with a 154 ns half-life. In both steps, the energy is primarily released via internal conversion electrons, with smaller contributions from Auger electrons, X-rays, and gammas. Early work on understanding ^{83m}Kr interactions in xenon TPCs was done by Kastens et al. [73] and Manalaysay et al. [116].

One may guess that ^{83m}Kr should behave like a standard electronic recoil with 41.5 keV of energy, but this is incorrect. The first step does behave roughly like a 32.1 keV ER. However, the short time between the two steps means that when the second deposition occurs, the excitons and electron-ion pairs from the first deposition are still present. They can interact with the excitons and electron-ion pairs generated by the second step. Thus, the light yield of the 9.4 keV deposition is dependent on the time Δt between the depositions, which is of course exponentially distributed. This is exactly what is seen by various experiments [75, 117, 118]. In theory, the charge yield should also depend on Δt , but experiments typically cannot distinguish

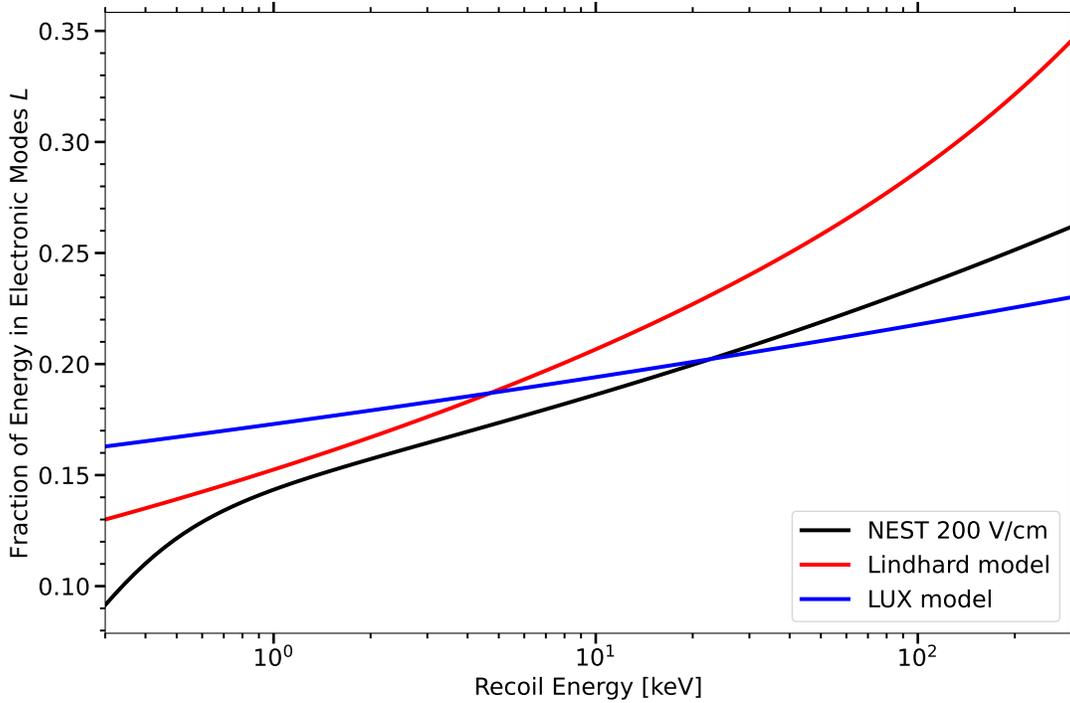


Figure 3.10: The fraction of deposited energy in observable electronic channels (i.e., not lost to heat) for nuclear recoils, also known as the Lindhard factor L , for a few different models.

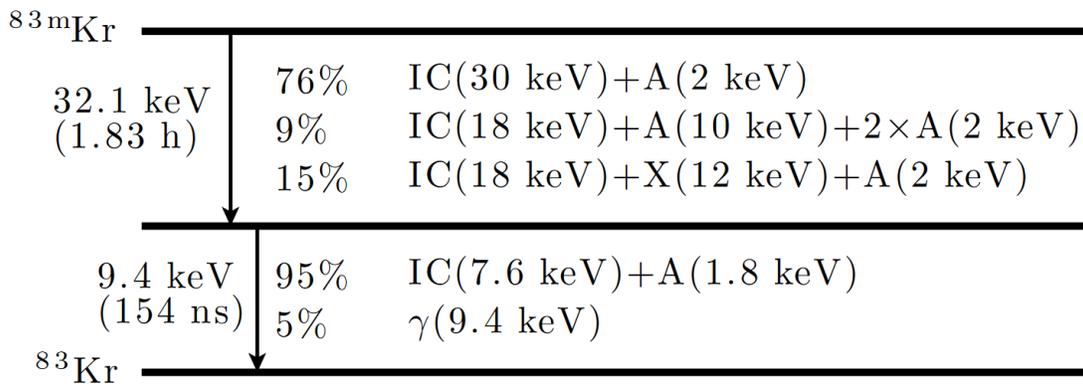


Figure 3.11: The two-step decay mechanism for ^{83m}Kr . Figure taken from [116].

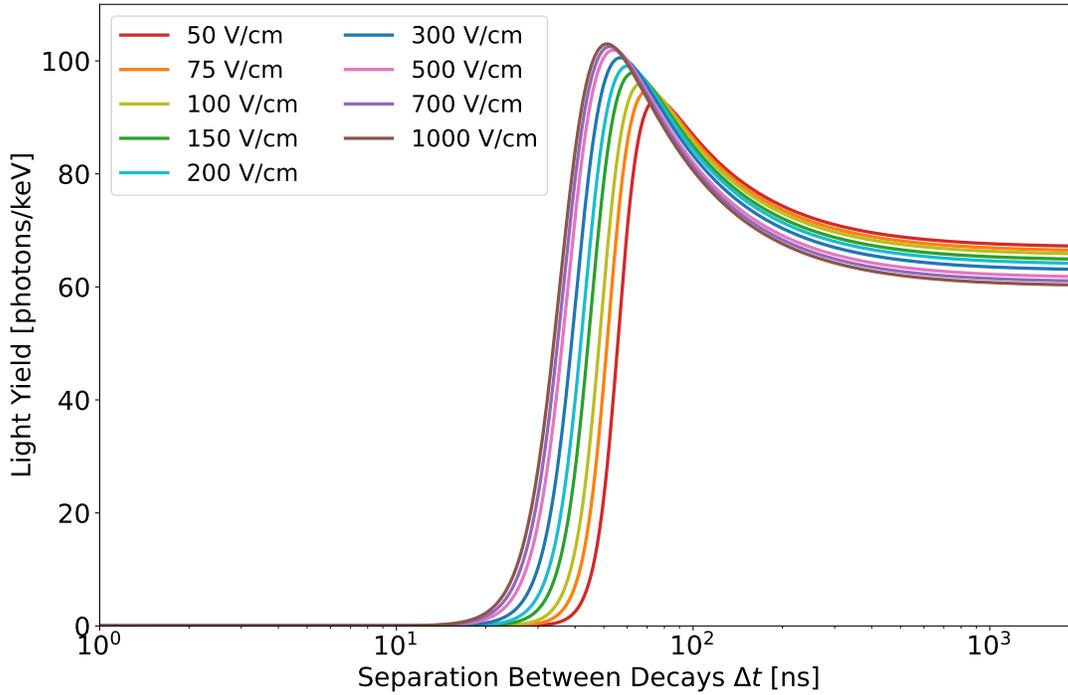


Figure 3.12: The light yield of the second step of the ^{83m}Kr decay, which deposits 9.4 keV, as a function of field and the time separation between the two steps.

charge signals separated by only 100s of ns, so only the combined charge signal from both depositions is observed. If Δt is low enough, the light signals cannot be distinguished either.

The light yield of the 9.4 keV step is shown vs. Δt and field in Fig. 3.12. We used the ^{83m}Kr NEST model to calculate this; note that the dissertation author co-lead the effort to build this model, with G. Rischbieter. Above $\Delta t = 100$ ns, the light yield decreases with field as expected and increases with Δt . Below $\Delta t = 100$ ns, the field trend reverses and the light yield goes to zero, but this should not be taken literally. Xenon TPCs are generally unable to resolve the light signals from the 32.1 and 9.4 keV depositions if $\Delta t < 100$ ns, so data is unavailable in this region. Instead, the model was forced to zero at low Δt in order to match combined 41.5 keV data—data where the two light signals were merged into a single pulse.

Finally, in Fig. 3.13, we compare the ^{83m}Kr yields directly to the two ER models in NEST. As mentioned, the 32.1 keV deposition is basically the same as the beta model. For 9.4 and 41.5 keV, we need to choose a Δt range. We choose a few values for minimum Δt for 9.4 keV depositions, since the pulses must be resolvable; we do the opposite for 41.5 keV.

^{83m}Kr is an incredibly useful calibration tool for liquid xenon TPCs. As a mo-

noenergetic source that fills the detector volume and decays relatively quickly, it allows physicists to measure position-dependence of the light and charge signals, monitor the detector’s performance over time, and measure electric field uniformity.

3.4 Two-Phase Liquid/Gas Xenon Time Projection Chambers

The two-phase liquid/gas xenon time projection chamber (xenon TPC) is the most widespread xenon particle detection technology. Each experiment has unique characteristics, but they all look roughly like the schematics in Figs. 3.1 and 3.14. The former is a cartoon of LUX, and the latter is a more realistic CAD drawing of LZ. A central liquid xenon chamber serves as the primary target volume, and a small layer of gaseous xenon exists on top of it. Photomultiplier tubes (PMTs) are located at the top and bottom of the detector. These convert an incident photon to a photoelectron, which is then amplified by a factor of 10^5 or 10^6 by means of strong electric fields and a dynode chain. An R&D-scale TPC will typically have in the single-digit or low double-digit numbers of PMTs, whereas a TPC for a WIMP search will have hundreds of PMTs. A downward-pointing electric field (the “drift field”, typically hundreds of V/cm) is established in the central volume by means of a cathode grid at the bottom of the detector and a gate grid immediately below the liquid surface. A stronger downward-pointing field (the “extraction field”, typically thousands of V/cm) is set up in the gas region between the gate grid and a higher anode grid. Extra grids are used to shield the PMTs, and field-shaping rings circle the liquid volume. Not portrayed in the figures are panels of solid polytetrafluoroethylene (PTFE), a type of Teflon, surrounding the xenon volume. There are also robust systems for cryogenics, circulation, purification, and shielding surrounding the TPC, but these are not relevant for this dissertation so are ignored².

Consider an energy deposit in the xenon. It produces some number of photons and electrons leaving the recoil site, as described in Sec. 3.3.

The 178 nm VUV photons propagate through the xenon at essentially the speed of light (reduced slightly by the index of refraction, 1.69 [120]) and can be detected by the PMTs. This is called the S1 signal (for “primary scintillation”). Not all photons are detected. An obvious issue is that the photons can travel in any direction and might hit the detector walls; however, they will be reflected by Teflon, which is >95% reflective in liquid xenon [121, 122]. Second, the photons often undergo total internal reflection at the liquid-gas interface, so they are detected in the bottom PMTs more frequently than the top PMTs. Additionally, the photocathode quantum efficiency is not unity—in fact, the best PMTs tend to have $\approx 30\%$ QE. For the photons that do induce a photoelectron, each photoelectron produces

²Interested readers can see the LZ Technical Design Report [119].

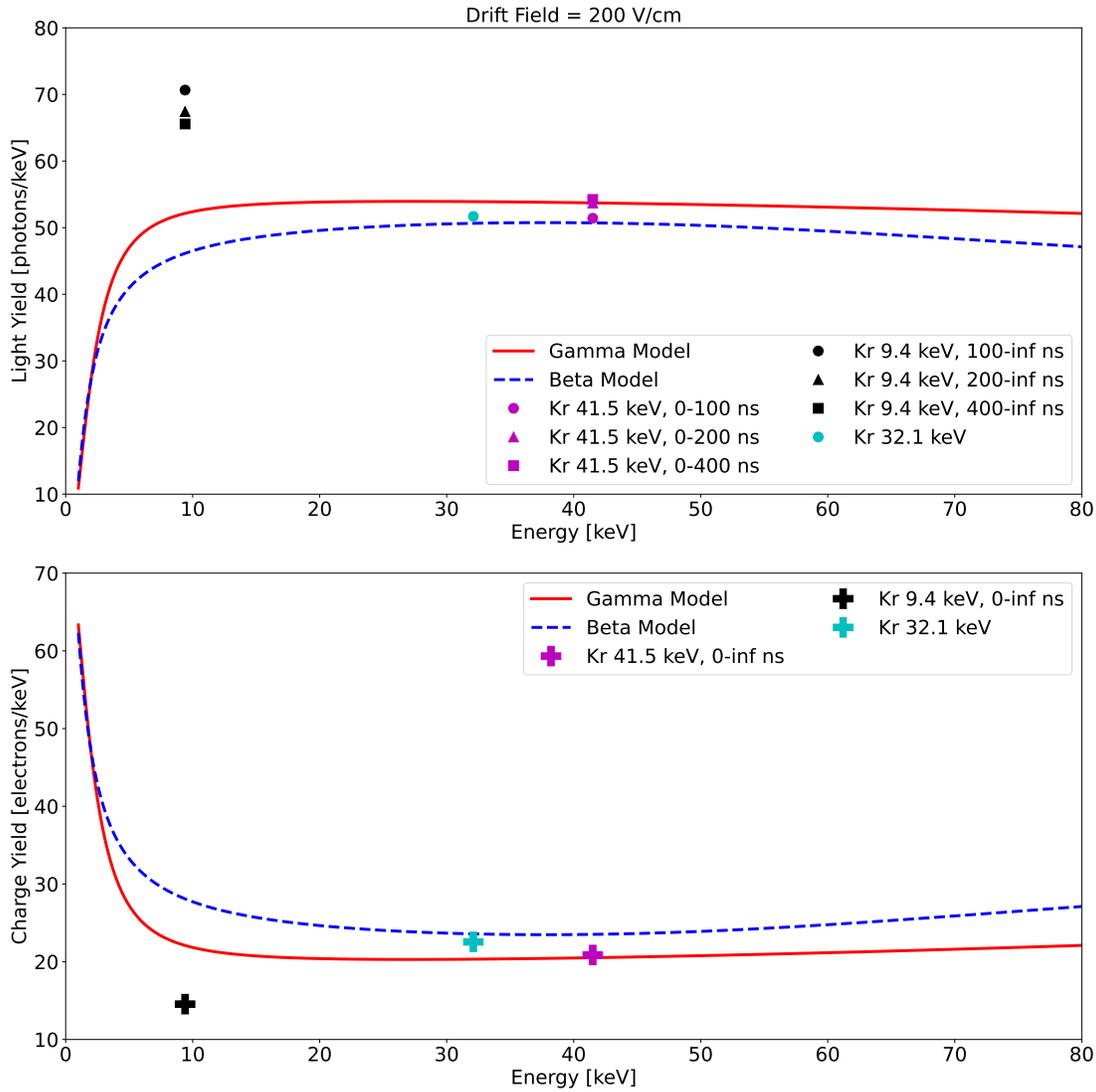


Figure 3.13: A comparison of the NEST ^{83m}Kr model for light and charge yield to the beta ER and gamma ER models. The 9.4 keV light yield depends on the time separation Δt , and we average over different Δt ranges, corresponding to an experiment's ability to resolve the S1a and S1b pulses. The 41.5 keV points show the combined yields for the 32.1 and 9.4 keV depositions, in the low- Δt situation when we cannot resolve the pulses. The charge yield is shown averaged across all possible Δt , since the two S2 pulses are typically merged.

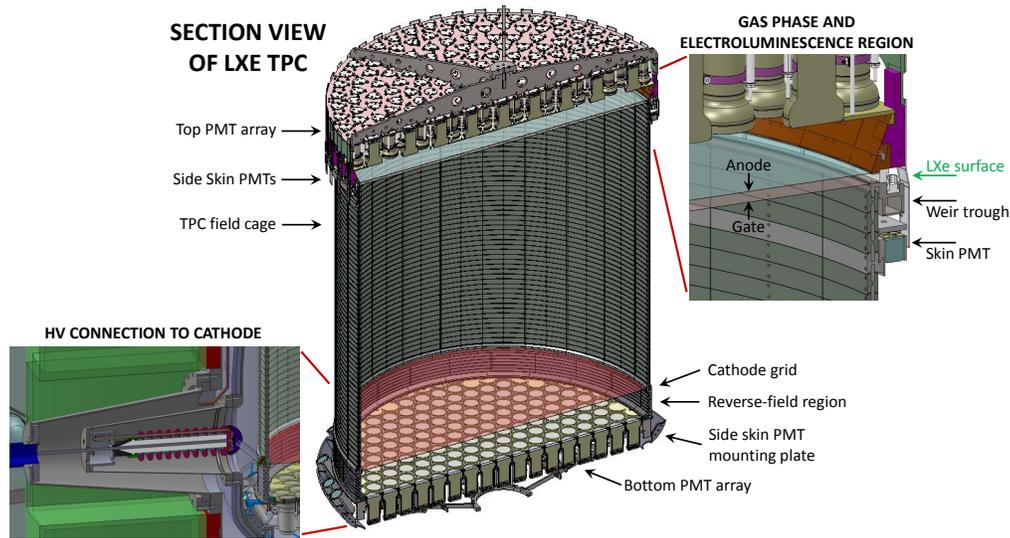


Figure 3.14: Schematic of the LUX-ZEPLIN detector, copied from [119].

a voltage pulse in the PMT output that is detected by some analog-to-digital data acquisition (DAQ) system. The area of the pulse from each photoelectron (in units of [Voltage] \times [Time]) is smeared due to intrinsic PMT resolution, typically with a width of 30–40%. One final complication is that there is a significant probability, typically $\approx 20\%$, that a photon incident on a PMT photocathode produces two photoelectrons [123, 124]. This is called the double photoelectron, or 2-PE, effect. The full S1 pulse corresponding to all photons from a given recoil can be characterized through various quantities, most importantly its area. S1 area is measured in units of photoelectrons [phe] or photons detected [phd], where the latter metric divides out the 2-PE effect. In addition, the S1 signal is typically detected within 10s of ns. For each photon produced by a recoil, the probability it generates 1 phe (or alternatively 1 phd) in any PMT is g_1 (g for gain), typically about 0.1 [125, 26]. This is a low probability! In general, the smallest measurable S1 signal sets the detector’s energy threshold. But it is sufficient for our purposes—for example, the LUX energy threshold was 1.2 keV for electronic recoils and 3.3 keV for nuclear recoils [126].

Meanwhile, the electrons leaving the recoil site drift upwards, propelled by the drift field³. The drift velocity depends on the drift field, but is typically on the scale of 1 mm/ μ s. As a result, they take hundreds of μ s to reach the liquid-gas interface, and significant longitudinal and transverse diffusion occurs during their journey. A significant fraction of electrons do not complete the journey, instead being captured by electronegative impurities, and xenon purity is crucial to minimize this

³There is a small liquid xenon region below the cathode grid, but above the PMTs, where the electrons drift downwards. This is called the reverse field region.

effect. When they reach the liquid surface, most are extracted into the gas phase by the much stronger extraction field. The higher the extraction field, the higher the extraction efficiency. In gas, they collide with xenon atoms and induce scintillation through the same mechanisms outlined for liquid xenon. These scintillation photons, known as S2 for “secondary” (or sometimes proportional) scintillation, are detected by the same PMTs that measure the S1 signal. However, most of the S2 signal is measured in the top array of PMTs. The same smearing effects described in the previous paragraph are relevant here. The S2 signal is much larger than the S1; it is typically detected over the course of $1 \mu\text{s}$, and its area is measured in units of phe or phd. For each electron leaving the recoil site, g_2 gives the average total PMT signal. Values of g_2 vary widely. For example $g_2 = 12.1$ for the first run of LUX [125], and $g_2 = 47.1$ for LZ [26].

The xenon TPC technology offers several advantages. First, we are able to reconstruct the 3D position of each recoil event. We get the z -position through the drift time, i.e. the time between the S1 and S2 pulses. Drift time equals zero at the liquid surface. We get the x, y position through the hit pattern on the top PMTs. In an ideal detector, the electron cloud drifts vertically upwards, so most of the S2 light will be located in the PMTs immediately above the recoil. In a real detector, position reconstruction is more complicated because of field non-uniformity and PMT-to-PMT variation, but these effects can be accounted for [127, 128]. 3D position reconstruction allows us to “fiducialize”, i.e. select events only from the central region of the xenon. Most backgrounds are located towards the edges of the xenon volume, due to self-shielding.

Another advantage is strong energy resolution—not compared to solid-state detectors like high-purity germanium, but compared to single-phase noble detectors and other detectors based purely on scintillation. The largest contributor to energy resolution is recombination fluctuations, which exchange a photon for an electron; since we detect both, we minimize the effects of these fluctuations. For a given recoil, we reconstruct the energy using Eq. 3.7, but modified for a xenon TPC:

$$E_{ee} = W \left(\frac{S1}{g_1} + \frac{S2}{g_2} \right). \quad (3.11)$$

The ee subscript stands for “electronic equivalent,” indicating that this formula can be used to reconstruct the entire energy of an electronic recoil. If we want to measure the energy of a nuclear recoil, we use

$$E_{nr} = \frac{E_{ee}}{L} \quad (3.12)$$

where L is the Lindhard factor and requires use of some model. When energy is reconstructed using either of these two equations, the units are appropriately written as keV_{ee} or keV_{nr}.

We also need to account for the position-variation of the S1 and S2 signals. Two identical recoils in different locations, but exposed to the same electric field, will produce the same N_{ph} and N_e . They will, however, produce different S1 and S2 areas. S1 variation is primarily due to photons reflecting off the liquid-gas interface and the detector walls, so the S1 signal tends to be stronger for recoils near the bottom of the detector. S2 variation is primarily due to the tendency of electrons to be captured by electronegative impurities as they drift upwards (this is called “electron lifetime”). As a result, the S2 signal is stronger for recoils near the top of the detector. Both signals also exhibit xy variation due to geometric effects, but this tends to be smaller than the variation based on z . We account for these variations by using “position corrections,” which are multiplicative factors applied to the S1 and S2 areas. Typically, the S1 area is normalized to the center of the liquid xenon volume, and the S2 area is normalized to $(x = 0, y = 0)$ at the liquid surface. The areas are denoted S1c and S2c after corrections are applied⁴. For example, suppose a given recoil produces an S2 with area 3000 phd. The drift time for this recoil is 100 μ s, and the xenon purity is such that the electron lifetime is 800 μ s. Then $S2c = S2 / \exp(-100/800) = 3400$ phd. Position corrections allow us to strengthen our energy resolution, giving us our final formula for energy reconstruction:

$$E_{ee} = W \left(\frac{S1c}{g_1} + \frac{S2c}{g_2} \right). \quad (3.13)$$

We conclude by showcasing an example of a waveform in the PIXeY detector, which will be the subject of Chapter 6. See Fig. 3.15. This is a 10 keV electronic recoil, the drift field is 407 V/cm, and the gas extraction field is 7980 V/cm. The detector has 7 PMTs each on the top and bottom arrays. The top panel shows the summed waveform across all 14 PMTs in units of phe/sample; 1 sample is 4 ns. The S1 and S2 pulses are highlighted in blue, and other identified pulses are highlighted in red. The drift time is 16 μ s. The second panel zooms in on the seven identified pulses and shows the waveform from each PMT; it also shows the area of each pulse. Visually, the five pulses that are classified as “Else” (neither S1 nor S2) are clearly single electron pulses. These are pulses where an electron is extracted from the liquid into the gas, but this electron is not part of an S2 pulse. They create jagged pulses with areas in the 10s of phe and widths in the high hundreds of ns. Single electrons in xenon TPCs are an interesting topic, but outside the scope of this dissertation, so interested readers are directed to [129, 130]. The middle panel zooms in on the S1 and S2 pulse specifically. For the S1, we see several overlapping photoelectrons within a short time, where each exhibits a sharp rise followed by a slower exponential fall. For the S2, we see a long Gaussian-shaped pulse in each PMT. The S1 and S2 durations are 60 ns and 920 ns, respectively, not shown in the figure. Finally, in the bottom two panels, we show the amount of S1 and S2 light

⁴Some experiments like XENON use the notation cS1 and cS2.

in each PMT. The S1 is mostly in the bottom array, as expected. The S2 light is evenly split between the two arrays, which is normal for PIXeY since it’s so small, but the S2 light is concentrated in the top array for larger TPCs like LUX.

3.5 Noble Element Simulation Technique (NEST)

We have alluded to the Noble Element Simulation Technique (NEST) several times, and it is worth describing this project in detail. The NEST collaboration [131] is a group of scientists from across the liquid noble element detection community, including members from LUX, LZ, XENON, (n)EXO, RED100, COHERENT, DUNE, ICARUS, MicroBooNE, and SBN ⁵. NEST is also the name of the simulation package produced by this collaboration, which models interactions in liquid xenon and argon; the current version of NEST is tagged as v2.3.7 [132]. It is the primary simulation tool for many liquid xenon experiments, including LZ. We will use “NEST” to refer to both the mathematical models and the code used to implement them, as these two are identical. The code is written primarily in C++, but bindings are available to use NEST in Python; there is also compatibility to use NEST with ROOT and GEANT4.

For full details, one should look through the resources available on NEST’s Github repository in [132], including the mathematical form of the models, tools to make benchmark plots, and written documentation. Here, we list the components that are crucial for the TPC analyses in this dissertation.

First, it is important to understand that NEST models are mostly empirical. They are built as averages of world data from liquid, gas, and solid xenon experiments going as far back as the 1970s [133]. Physically-motivated models like Thomas-Imel and Doke-Birks are incorporated when feasible, but they are not prioritized.

NEST works by walking through every step of energy partitioning described in Sec. 3.3, followed by the signal detection mechanisms described in Sec. 3.4. To start, the user sets up a file with all the parameters of their detector. An example is shown in Fig. 3.16. These parameters include the detector size, number of PMTs, PMT resolution and threshold, drift field, g_1 and g_2 , electron lifetime, and any position corrections to be applied to the S1 and S2 signals. Then, the user requests a specific particle type—e.g., ER, NR, alpha, ^{83m}Kr —and energy deposition. NEST models have been constructed to determine the average charge and light yield for the interaction, as well as the average recombination probability and fluctuations. NEST then stochastically simulates the number of excitons and ions produced by the energy deposit, and the number of electrons and photons leaving the recoil site.

⁵The dissertation author is one of these scientists.

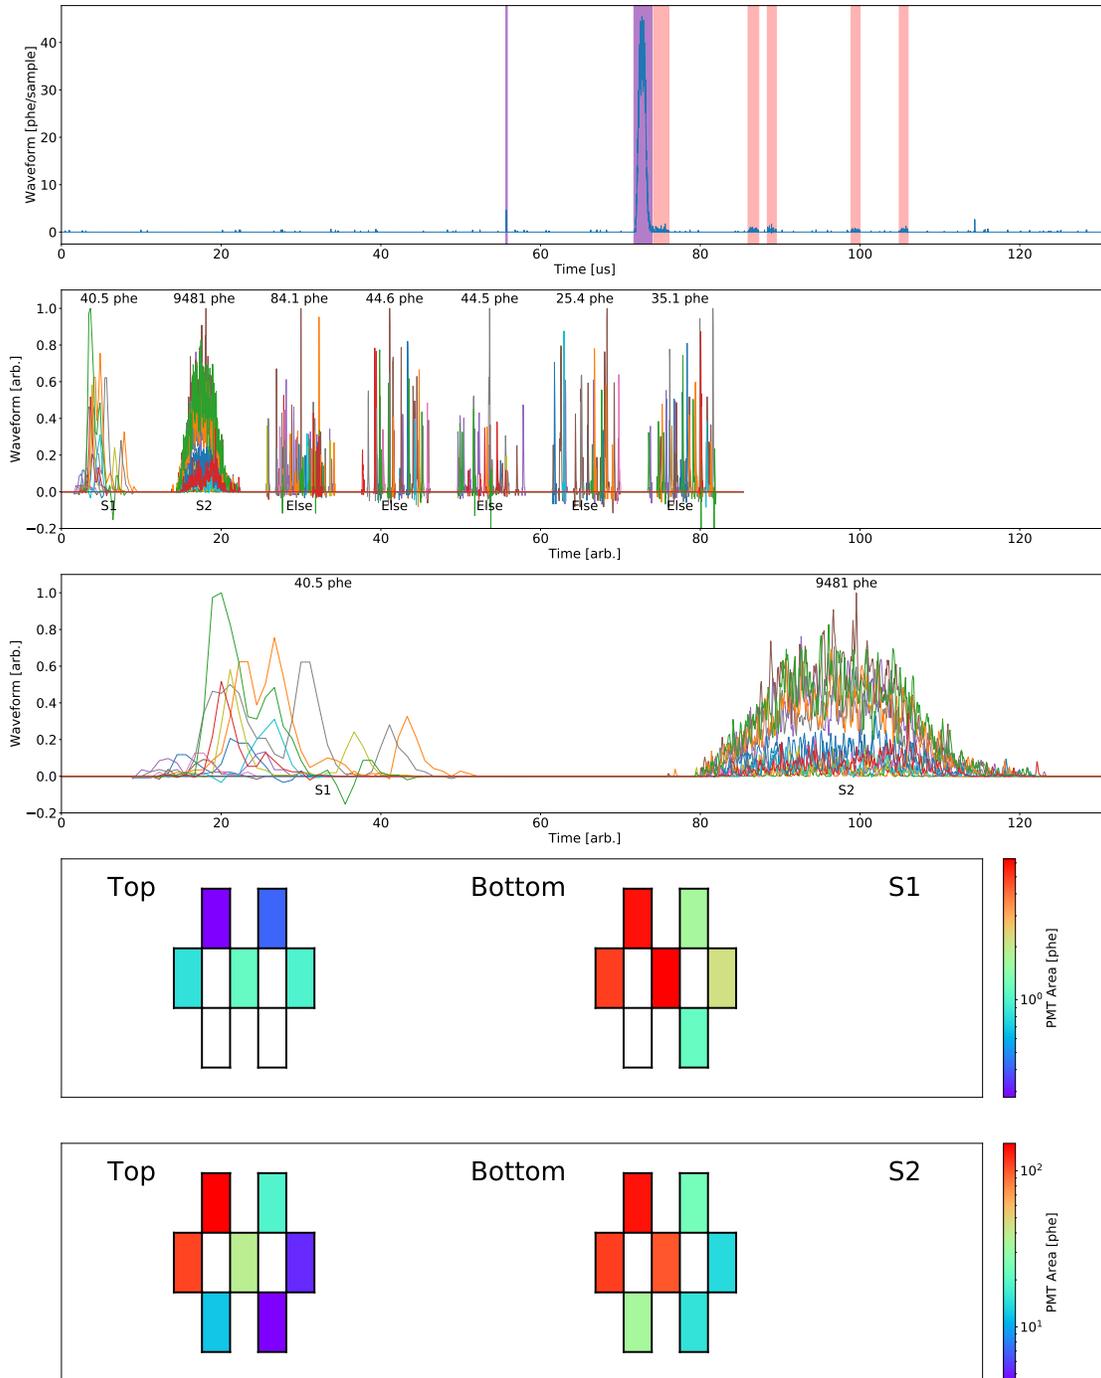


Figure 3.15: An example waveform for a 10 keV electronic recoil event in the PIXeY detector. See text for details.

Finally, the detector response is simulated. For a given number of photons, NEST simulates how many photons are actually detected and the size of the S1 signal in each PMT. It does the same for the S2 signal. If the user simulates many deposits, they can get a distribution of events that should match the distribution from data.

```

27 void Initialization() override {
28     // Primary Scintillation (S1) parameters
29     g1 = 0.1170; // 0.117+/-0.003 WS,0.115+/-0.005 D-D,0.115+/-0.005
30             // CH3T,0.119+/-0.001 LUXSim. UNITS: phd per photon (NOT
31             // photoelectrons!!)
32     sPEres = 0.37; // arXiv:1910.04211. UNITS: phe a.k.a. PE or photoelectrons
33     sPEthr = (0.3 * 1.173) / 0.915; // arXiv:1910.04211. UNITS: phe
34     sPEeff = 1.00; // arXiv:1910.04211. UNITS: fractional
35     noiseBaseline[0] = 0.00; // arXiv:1910.04211 says -0.01. UNITS: phe
36     noiseBaseline[1] = 0.08; // arXiv:1910.04211. UNITS: phe
37     noiseBaseline[2] = 0.; // UNITS: e-'s
38     noiseBaseline[3] = 0.; // UNITS: e-'s
39     P_dphe = 0.173; // arXiv:1910.04211. UNITS: fractional
40
41     coinWind = 100; // 1310.8214. UNITS: ns
42     coinLevel =
43         2; // 1512.03506. UNITS: number of PMTs for coincidence requirement
44     numPMTs = 119; // 122 minus 3 off. UNITS: number of PMTs
45
46     OldW13eV = true; // default true, which means use "classic" W instead of
47             // Baudis / EXO's
48     noiseLinear[0] = 0.0e-2; // 1910.04211 p.12, to match 1610.02076 Fig. 8.
49             // UNITS: fraction NOT %!
50     noiseLinear[1] = 0.0e-2; // 1910.04211 p.12, to match 1610.02076 Fig. 8.
51             // UNITS: fraction NOT %!
52
53     // Ionization and Secondary Scintillation (S2) parameters
54     g1_gas = 0.1; // 0.1 in 1910.04211. UNITS: phd per e-
55     s2Fano = 3.6; // 3.7 in 1910.04211; this matches 1608.05381 better.
56             // Dimensionless
57
58     s2_thr = (150. * 1.173) / 0.915; // 65-194 pe in 1608.05381. UNITS: phe
59     E_gas = 6.25; // 6.55 in 1910.04211. UNITS: kV/cm
60     eLife_us =
61         800.; // p.44 of James Verbus PhD thesis Brown. UNIT: microseconds (us)
62
63     // Thermodynamic Properties
64     // inGas = false; //duh
65     T_Kelvin = 173.; // 1910.04211. UNITS: Temperature in Kelvin
66     p_bar = 1.57; // 1910.04211. UNITS: pressure in bar
67
68     // Data Analysis Parameters and Geometry
69     dtCntr =
70         160.; // p.61 Dobi thesis UMD, 159 in 1708.02566. UNITS: microseconds
71     dt_min = 38.; // 1608.05381. UNITS: microseconds
72     dt_max = 305.; // 1608.05381. UNITS: microseconds
73
74     radius = 200.; // 1512.03506. UNITS: mm
75     radmax = 235.; // 1910.04211. UNITS: mm
76
77     TopDrift = 544.95; // 544.95 in 1910.04211. UNITS: mm
78     anode = 549.2; // 1910.04211 and 549 in 1708.02566. UNITS: mm
79     gate = 539.2; // 1910.04211 and 539 in 1708.02566. UNITS: mm
80     cathode = 55.90; // 55.9-56 in 1910.04211,1708.02566. UNITS: mm
81
82     // 2-D (X & Y) Position Reconstruction
83     PosResExp = 0.015; // arXiv:1710.02752 indirectly. UNITS: mm^-1
84     PosResBase = 70.8364; // 1710.02752 indirectly. UNITS: mm
85 }

```

Figure 3.16: NEST detector parameter file (C++) used for the LUX 2013 run.

Chapter 4

ER/NR Discrimination in the LUX Detector

In Chapter 2, we introduced the WIMP nuclear recoil signal. In Chapter 3, we discussed the differences between electronic and nuclear recoils from an energy partitioning standpoint. Yet we haven't explored how these differences play out in an actual dark matter detector. We do so in this Chapter, using data from the Large Underground Xenon (LUX) detector.

We begin by exploring the dominant backgrounds to WIMP searches in xenon TPCs, which are mostly electronic recoils. This motivates the analysis of ER/NR discrimination and its importance to designing future detectors. We then describe the LUX detector: its design, run time, calibrations, and other characteristics. Next, we briefly summarize, at a high level, how ER/NR discrimination is implemented. Finally, the majority of the chapter is dedicated to an analysis of LUX data to learn how discrimination is affected by drift field, light collection efficiency, and recoil energy.

Most of the work in this Chapter and the next Chapter has been published in Physical Review D [134], with the dissertation author serving as the primary author of the publication.

4.1 Backgrounds for Xenon TPCs

Backgrounds are energy depositions in the xenon that behave like a WIMP signal—they interact only in the xenon, producing a single S1 and S2 signal with a reasonable pulse area and 3D position—but are actually from another source. To run a successful experiment, we need to mitigate backgrounds through two methods: preventing these depositions from occurring in the detector at all, and isolating and rejecting the events that do occur. Our analysis focuses on the latter. As a result, it is important to understand where backgrounds arise from.

The background sources in all xenon TPCs are essentially the same, but their relative importance varies based on the detector. One particularly crucial factor is size, because as detectors get larger, self-shielding becomes more powerful and the impact of external backgrounds is reduced. The discrimination analysis is aimed at informing future detectors, so we primarily focus on backgrounds in LZ, which has a 7-tonne active volume, and larger TPCs. We then briefly touch on backgrounds in LUX, with a 250-kg active volume, to note the differences.

4.1.1 LZ and G3 Backgrounds

To discuss WIMP Search backgrounds in LZ and larger detectors, we rely heavily on the LZ projected sensitivity paper [86]. The LZ results paper is more recent [26], but the sensitivity paper is more detailed in its accounting of backgrounds. Figure 4.1 shows the energy profile of all the ER and NR backgrounds, and Table 4.1 shows the corresponding number of events. Both are taken directly from [86] (with some

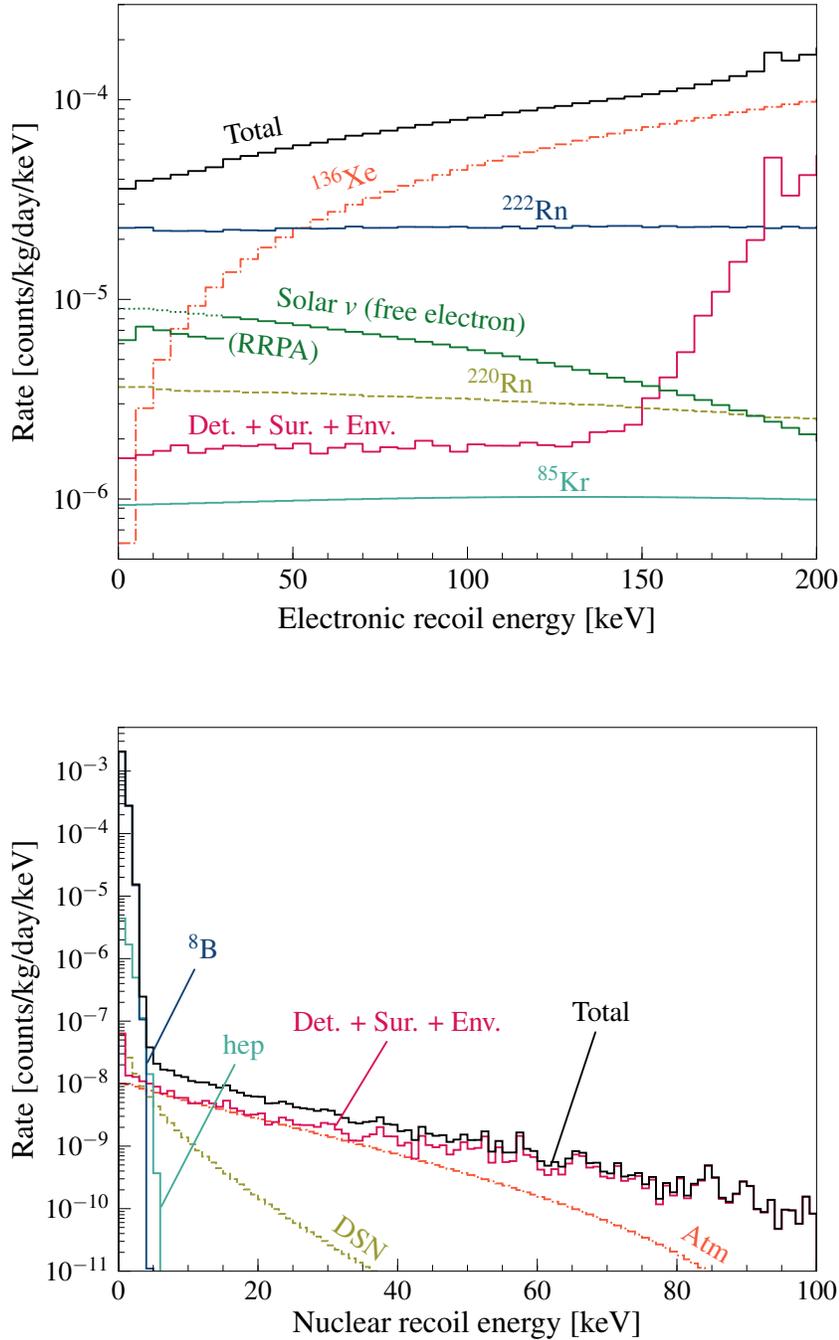


Figure 4.1: Simulated electronic (top) and nuclear (bottom) recoil backgrounds for LZ. Includes single-scatter events in the 5.6-tonne fiducial region. Details of individual backgrounds in text. Copied from Figs. 4 and 5 in [86].

Table 4.1: Simulated electronic and nuclear backgrounds, integrated across energy, for LZ. Includes single-scatter events in the 5.6-tonne fiducial region, for a 1000-day exposure. Energy range is 1.5–6.5 keV for ERs and 6–30 keV for NRs. Details of individual backgrounds in text. Modified from Table III in [86] to condense details beyond the scope of this dissertation.

Background Source	$^{238}\text{U}_e$	$^{238}\text{U}_l$	$^{232}\text{Th}_e$	$^{232}\text{Th}_l$	^{60}Co	^{40}K	n/yr	ER (cts)	NR (cts)
	mBq/kg								
Detector Components (e.g., PMT, TPC, Cryostat, OD)	47.1	12.83	7.97	7.92	5.99	45.46	8751	9.15	0.07
Surface Contamination (Dust, PTFE plate-out, ^{210}Bi mobility, ion misreconstruction, ^{210}Pb in PTFE)								40.2	0.39
Laboratory and Cosmogenics (Rock walls, muon-induced neutrons, cosmogenic activation)								4.8	0.06
Xenon contaminants									
^{222}Rn (1.8 $\mu\text{Bq/kg}$)								681	-
^{220}Rn (0.09 $\mu\text{Bq/kg}$)								111	-
^{nat}Kr (0.015 ppt g/g)								24.5	-
^{nat}Ar (0.45 ppb g/g)								2.5	-
							subtotal	819	0
Physics									
^{136}Xe $2\nu\beta\beta$								67	-
Solar neutrinos: $pp+^7\text{Be}+^{13}\text{N}$, $^8\text{B}+hep$								191	0.00
Diffuse supernova neutrinos (DSN)								-	0.05
Atmospheric neutrinos (Atm)								-	0.46
							subtotal	258	0.51
Total								1131	1.03

simplification of the table). In total, for a 1000-day exposure, LZ expects to see 1131 ER background events and only about 1 NR event. These numbers are after fiducialization, i.e. the process in which we limit our analysis to the central region of the TPC, which is assumed to have a mass of 5.6 tonnes.

The discrepancy between ERs and NRs is staggering. Essentially all the backgrounds are ERs. As a rule of thumb, the sensitivity tends to scale as X/\sqrt{B} , where X is the exposure, and B is the number of background events observed. This means that the difference between keeping and rejecting ER backgrounds corresponds to a sensitivity improvement of 34x. Fortunately, as we will show, we can indeed reject the vast majority of ER backgrounds based on the partitioning dynamics shown previously. A somewhat conservative discrimination power of 99.8% reduces the LZ ER background to 2.3 events and the total background to 3.3 events, and it improves sensitivity by a factor of 19.

The electronic recoil backgrounds from detector components are generated from trace amounts of ^{238}U , ^{232}Th , ^{60}Co , and ^{40}K in all detector materials: PMTs, TPC construction, and more. These isotopes and their progenies emit gammas. Surface contamination is a source of backgrounds via dust containing the aforementioned

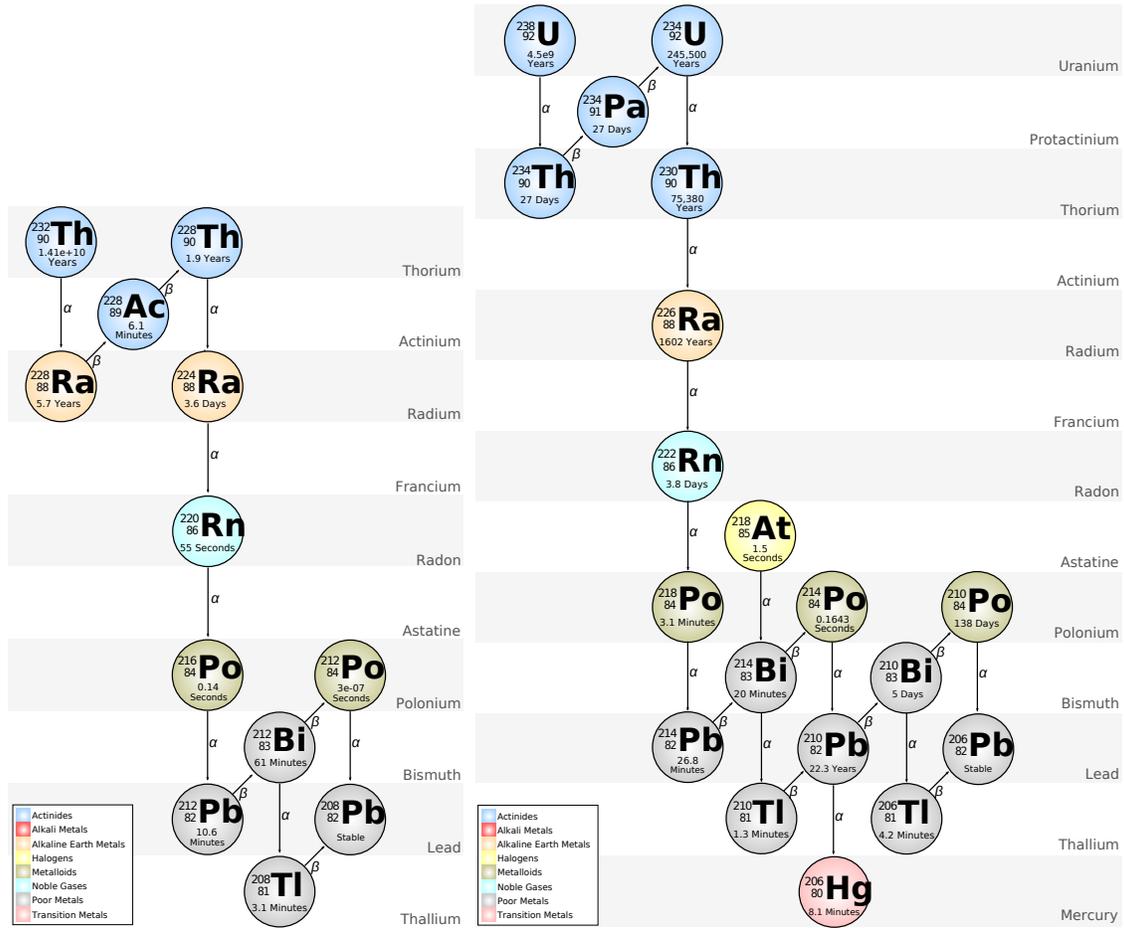


Figure 4.2: Decay chains for ^{232}Th and ^{235}U . Copied from [135].

isotopes and plate-out of beta-emitting ^{222}Rn daughters. LZ also models backgrounds from natural radioactivity in the lab (again, mostly gammas from the the same isotopes) and cosmogenically-activated ^{127}Xe in the fiducial volume. As seen in Table 4.1, these three are minor components of the ER background, only contributing 54 events, or $<5\%$. This is due to fiducialization; most of these backgrounds occur near the walls of the detector. On the other hand, these sources contribute half of the NR background, through neutron recoils and erroneous reconstruction of heavy ions as NRs. These neutrons are generated through (α, n) interactions, fission primarily from ^{238}U , and muons.

Xenon contaminants are the primary background for LZ. Xenon is a noble gas, so it is easy to purify from chemically active elements—for example, through a getter. However, radioactive isotopes from other noble elements are a concern. When they enter the xenon volume, they mix with the xenon, and decays can occur anywhere inside the detector; fiducialization does not make a difference. Two isotopes of note

are the beta emitters ^{85}Kr and ^{39}Ar . They enter the xenon supply because all three elements are components of air, and when air is purified to extract xenon, some argon and krypton remains. But their contributions pale in comparison to ^{220}Rn and ^{222}Rn , which contribute 70% of the ER background. These primarily originate from the same ^{238}U and ^{232}Th isotopes mentioned earlier; see Fig. 4.2 for the full decay chain. Radon produced from those isotopes in detector materials/surfaces can leach into the xenon. Then, radon and its daughters will decay and release various radiation. The most dangerous radon daughter backgrounds are the so-called “naked betas,” which produce single-site electronic recoils with no associated radiation. These include ^{212}Pb , ^{208}Tl , ^{214}Pb , ^{210}Tl , ^{210}Pb , ^{210}Bi , and ^{206}Tl . None of these contaminants create nuclear recoil backgrounds.

The final backgrounds are double beta decay and neutrino backgrounds. The isotope ^{136}Xe decays via the simultaneous emission of two electrons and two antineutrinos (and potentially in a neutrinoless form, but this has yet to be observed [136]). The outgoing electrons carry up to 2.5 MeV of energy, and they induce ERs which can be in the low-energy WIMP region of interest. The solar, atmospheric, and supernova neutrino backgrounds are discussed in Chapters 2 and 8. They can scatter off both electrons and nuclei; in xenon, the lower-energy neutrinos and higher-energy neutrinos are relevant for ER and NR backgrounds, respectively.

As mentioned previously, the relative importance of these backgrounds changes as we build larger, more advanced detectors, so a G3 detector at the 100–1000 tonne scale will be dominated by different backgrounds. LZ is already large enough that detector and laboratory backgrounds are a minor issue, and this should also be true for a G3 experiment. The question of radioactive contaminants is outstanding. LZ has already implemented procedures for removal of Kr and Rn, but in theory, future detectors could try to improve these processes. In addition, removal of ^{136}Xe is certainly possible; EXO-200 purified this isotope for their neutrinoless double beta decay experiment [136]. However, removing the ^{136}Xe is unappealing, because it precludes an experiment from searching for neutrinoless double beta decay. But supposing that Rn, Kr, and ^{136}Xe are removed, this still leaves the neutrino background, which is impossible to eliminate. Neutrinos will be the limiting backgrounds for future TPCs, and as with LZ, they primarily deposit ER depositions but contribute a significant NR background.

4.1.2 LUX Backgrounds

We briefly discuss LUX backgrounds, only because LUX is the detector we use for our analysis. LUX was a much smaller TPC, with only 250 kg of xenon in the instrumented volume, so self-shielding was less effective (though still quite powerful). As a result, radioactive components in the detector materials were the dominant background. Internal contaminants like ^{85}Kr and Rn daughters were present, but smaller fractions of the total ER background compared to LZ. Neutron backgrounds

were also present from the same external sources mentioned for LZ—indeed, the two experiments are located in the same cavern. Backgrounds from ^{136}Xe and neutrinos were at negligible levels (much less than 1 expected event). These are rare processes not visible in the small volume of LUX. Details can be found in [137].

4.2 The Large Underground Xenon (LUX) Experiment

4.2.1 About the detector

The LUX experiment was a two-phase liquid/gas xenon time projection chamber that operated at the 4850' level of the Davis Cavern at the Sanford Underground Research Facility in Lead, South Dakota. A schematic and photo can be seen in Fig. 4.3. It had two primary science runs. The first ran from April to August 2013; we will usually refer to it as WS2013, but the reader should note that it is also called “Run 3” or “Run03.” The second ran from September 2014 to August 2016; we will call it WS2014–16, but is is also referred to by “Run 4” or “Run04.” The active mass was 250 kg of liquid xenon, while the fiducial mass for the dark matter search was about 100 kg. The fiducial cut removed events that occurred less than about 3 cm from the detector wall. An additional 1 cm of gaseous xenon was present above the liquid, converting the ionization response into an optical signal via electroluminescence as described previously. The detector was instrumented with 122 5.6-cm diameter Hamamatsu R8778 photomultiplier tubes (PMTs), with 61 PMTs at the top of the detector (in the gas phase) and 61 at the bottom (immersed in the liquid phase). Furthermore, the detector was instrumented with three wire grids to control the electric field in the liquid and the gas—a cathode at the bottom of the detector, a gate slightly below the liquid level, and an anode in the xenon gas above the liquid level—and two grids in front of the PMT arrays to prevent stray fields from affecting the PMT photocathodes. Full technical details of the experiment’s configuration can be found in [138]. Here, we focus on how signals are converted into usable data for the relevant analysis. This discussion is brief. Most details can be found in [125], but we will be explicit when our analysis choices differ from the choices in [125].

Data from all PMTs was recorded continuously at 100 MHz (1 sample = 10 ns), and the raw voltage waveforms were stored in binary files. The locations of S2 pulses were identified by the hardware trigger and stored in the same binary files. Offline software used the trigger locations to identify event waveforms. Each event is processed to identify up to 10 pulses and classify them as S1, S2, single photoelectron (SPE), single electron (SE), or Else. S1 pulses are required to have hits in at least 2 PMTs to reduce spurious SPE backgrounds; this sets the energy threshold for our analysis.

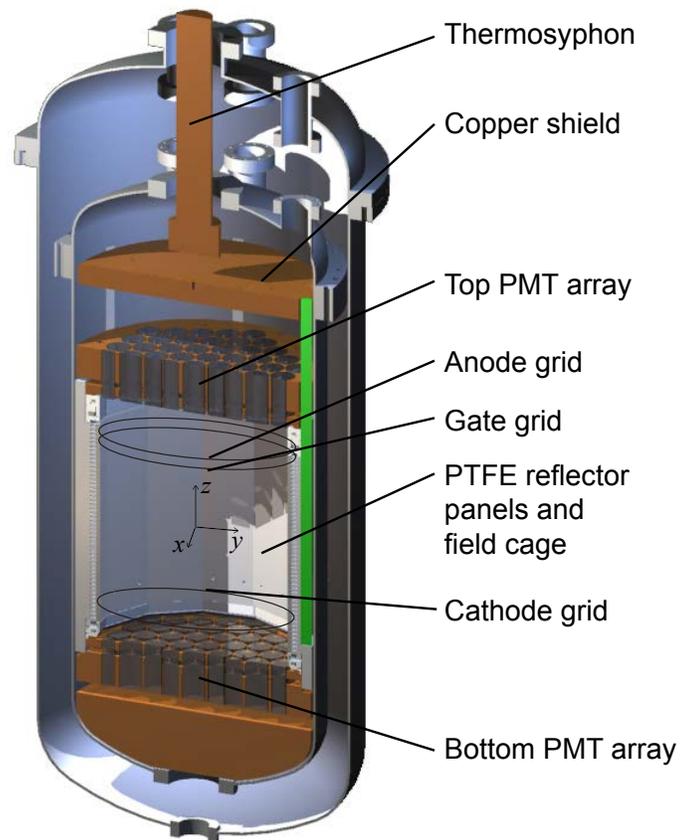


Figure 4.3: CAD schematic of LUX detector, and photo of LUX detector inside the water tank. Copied from Figs. 1 and 4 in [125].

For each pulse, a plethora of some 200 “reduced quantities” (RQs) are calculated and stored in ROOT, Matlab, and binary files. These RQs include information such as start time, pulse area, pulse width, Gaussian fit quality, etc. When the S1 pulse is relatively small, we can measure S1 in two ways: by integrating the full pulse area or by counting the number of photoelectron “spikes” recorded in each PMT. In our ER/NR discrimination analysis, S1 or S1c refers to spike count if the pulse area is less than 80 detected photons, and it refers to integrated pulse area otherwise. This hybrid variable is used because spike counting leads to better discrimination at low energies, but it cannot be reliably determined for large photon statistics at higher energies. Note that when calculating pulse area, we adjust for the 2-PE effect, using a double photoelectron emission probability of about 17%, so the units we use are [photons detected] or [phd]. 3D position was also calculated at the RQ-calculation step and used to correct the S1 and S2 areas. The position corrections themselves were calculated using data from regular ^{83m}Kr injections. The position corrections are dependent on z only in WS2014–16 data and on the full xyz position in WS2013 data. However, the WS2013 corrections are dominantly z -dependent, and when we compare results between the two science runs, we use z -dependent position corrections for WS2013 data.

Finally, for our discrimination analysis, single scatter events are identified. These are events with only one S2 pulse, and only one S1 pulse in the T_{max} preceding the S2, where T_{max} is the maximum drift time. This is in contrast to multiple scatters, a common phenomenon for gamma and neutron interactions, where the incident particle scatters multiple times before thermalizing or leaving the xenon volume. For such events, we only see one S1 pulse (the scatters are so close in time that the separate light signals cannot be resolved), but there are multiple S2 pulses.

As described previously, the energy deposited from electronic and nuclear recoils can be reconstructed by Eqns. 3.13 and 3.12, respectively. When we use these formulas, we refer to the energy as having units of keV_{ee} or keV_{nr} . However, we need a model for the Lindhard factor. We use

$$\begin{aligned}
 L &= [AE_{ee}^{\gamma-1}]^{1/\gamma} = AE_{nr}^{\gamma-1}, \text{ or equivalently} \\
 E_{ee} &= A E_{nr}^{\gamma}, \text{ where} \\
 A &= 0.173 \text{ and } \gamma = 1.05.
 \end{aligned}
 \tag{4.1}$$

We have confirmed that, by using this relationship, we are able to match LUX D-D nuclear recoil calibration data to its theoretical energy spectrum. The reader should note that since $\gamma \approx 1$, Eq. 4.1 is comparable to a linear scaling by a factor of 5–6. This model is similar but not identical to the Lindhard model [114] often used to describe nuclear recoils in liquid xenon. The discrepancy is reasonable because the Lindhard model does not perfectly reproduce the nuclear recoil energy scale across all energies; see e.g., Fig. 15 of [115]. A comparison of these two models can be seen in Fig. 3.10.

4.2.2 Calibrations

LUX underwent several calibration campaigns throughout WS2013 and WS2014–16 to understand the detector’s response to different types of energy depositions. Both runs featured three specific calibration strategies that we focus on here. First, we injected a tritiated methane source into the xenon [139, 140]; this is a molecule that is chemically similar to methane, CH_4 , but with one of the hydrogen atoms replaced by tritium. We usually refer to this calibration source as ^3H , but some figures may refer to it as CH_3T . Tritium is a β^- emitter with a half-life of 12.3 years and an end point of 18.6 keV, making it useful for calibrating low-energy electronic recoils. It also filled the entire detector volume, allowing us to examine effects in different locations. Second, we ran nuclear recoil calibration campaigns by generating 2.5 MeV neutrons from deuterium-deuterium fusion (referred to as a D-D calibration), which deposit up to 74 keV on a xenon nucleus [115, 141]. These were produced by a neutron generator placed outside the xenon volume, and the height of this generator varied during WS2014–16. Third, we regularly (approximately weekly) calibrated the detector with $^{83\text{m}}\text{Kr}$. In addition to these, LUX ran a ^{14}C calibration campaign after the final WIMP search, in August 2016; we injected a $^{14}\text{CH}_4$ methane molecule, which allowed us to calibrate the detector up to 156.5 keVee [140, 109].

Each of these calibration campaigns consisted of several acquisitions—in the case of ^3H and D-D, hundreds of them. To identify which acquisitions to use, we manually scanned the LUX SQL database, called the “LUG.” Acquisitions were excluded if they exhibited any issues. Examples include overlap between calibrations with similar energies (for example, a $^{83\text{m}}\text{Kr}$ injection immediately following a ^{14}C calibration would be contaminated with ^{14}C events around 41.5 keV), circulation outages, a tripped electrode, a tripped PMT, and unusually high pulse rates.

After processing the relevant data files and identifying single-scatter events, we applied event-level quality cuts. Most of these are described in past literature [139, 115, 125, 64, 63]. These include the fiducial cut to select recoils in the central region; cuts on S1 and S2 area to select events in the appropriate energy range; cuts to reject misshapen S1 and S2 pulses; and a cut to remove multiple scatters that are misclassified as single scatters. For the discrimination analysis, we added additional cuts to increase the purity of our dataset; LUX calibration campaigns were abundant enough in statistics that we were able to sacrifice some quantity for quality. For each ER injection, we required events to occur within 3 half-lives after the start of the injection—in this case, the “half-life” is not the isotope half-life, but the removal time of methane from the TPC, about 6 hours. For D-D calibrations, we required events to be in the path of the neutron beam; this was a tighter requirement than the fiducial cut and accounted for the varying beam height during LUX’s runs.

The final datasets for ^{14}C , ^3H , and D-D, with all cuts applied, are showcased in Figs. 4.4, 4.5, and 4.6, respectively. We show data from WS2014–16, but data from WS2013 are consistent. In these figures, we show the distributions of energy, S1, S2,

xy-position, z-position, and drift time. From the energy distributions, we see that the ER distributions are well-fit to the β -decay spectrum. Slight inconsistencies are due to finite energy resolution, unaccounted for in the fit function. The ^{14}C and ^3H distributions both have energy thresholds around 1.2 keVee. Meanwhile, we see that at low energies, the energy distribution of the D-D calibration is well-fit to a 63 GeV WIMP with a 3.5 keVnr threshold. This is relevant because the discrimination we will calculate later depends on the NR energy distribution. As mentioned earlier, the S1 and S2 position corrections are done using z-position only. We see that for ER calibrations, events are uniformly distributed in xyz throughout the detector. On the other hand, for the D-D calibration, event density decreases as we get farther from the beam position. The beam was located near $x = 10$ cm, $y = -25$ cm, and the height varied, leading to multiple peaks in z and drift time.

4.2.3 Electric field variation

In WS2013, the drift field was fairly uniform across the liquid xenon target region at 177 ± 14 V/cm. However, in WS2014–16, the drift field varied significantly throughout the detector from 30 V/cm at the bottom of the fiducial region to 600 V/cm at the top. In [127], the LUX Collaboration hypothesized that the drift field variation was created by net charge buildup within the polytetrafluoroethylene (PTFE) detector walls and that this buildup of charge was induced by the strong VUV fluxes experienced during grid conditioning. A method for converting an event’s 3D position to the electric field at the recoil site was described in that publication. This was a complication for the WIMP search analysis, but it provides us with an opportunity to examine how discrimination is affected by electric field. Figure 4.7 shows the distribution of field in the LUX fiducial volume, as well as the field distribution of events in the calibrations mentioned in Section 4.2.2; the reader may observe the dramatic difference between the two runs. The uncertainty on the electric field magnitude is estimated to be $\sim 10\%$, based on comparisons between light and charge yields in simulation and data [142].

4.3 Tools for Discrimination

As we have convincingly shown, ER/NR discrimination is necessary for a xenon-based dark matter experiment to complete a competitive search. Our discussion in Chapter 3, particularly Figs. 3.5 and 3.5, shows the two strategies for accomplishing this: the ratio of charge to light leaving the recoil site, and the ratio of singlet to triplet excimers produced.

In a WIMP Search analysis, we would discriminate between signal and background by inserting these observable quantities into a statistical test, such as a

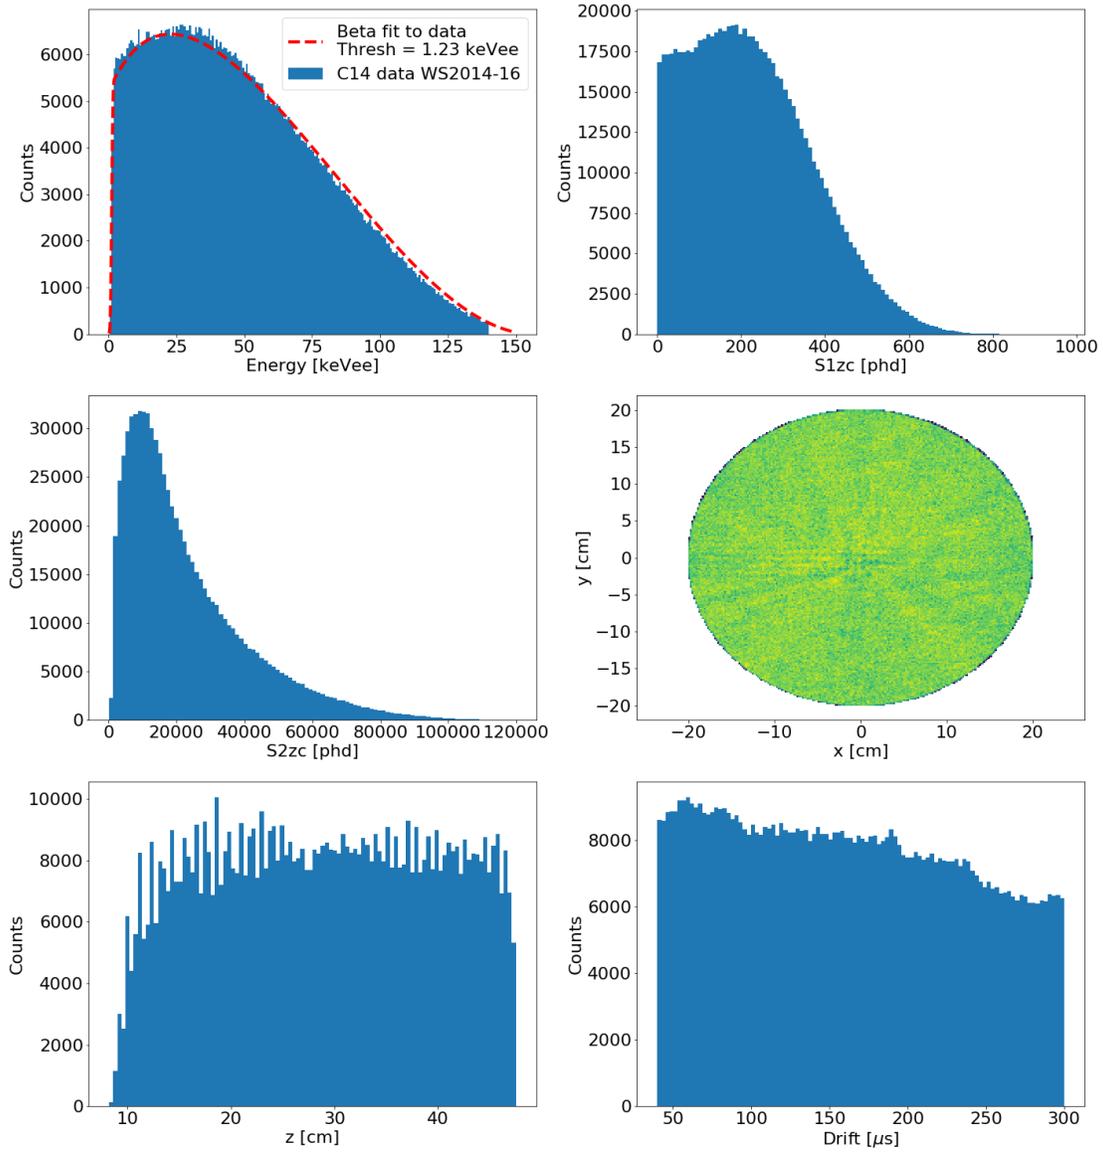


Figure 4.4: Distributions of energy, S1, S2, xy-position, z-position, and drift time for the ^{14}C WS2014-16 calibration. The top-left panel also includes a fit of data to the β decay spectrum with some energy threshold, fit to 1.23 keVee. The S1 and S2 areas are corrected using z-position only.

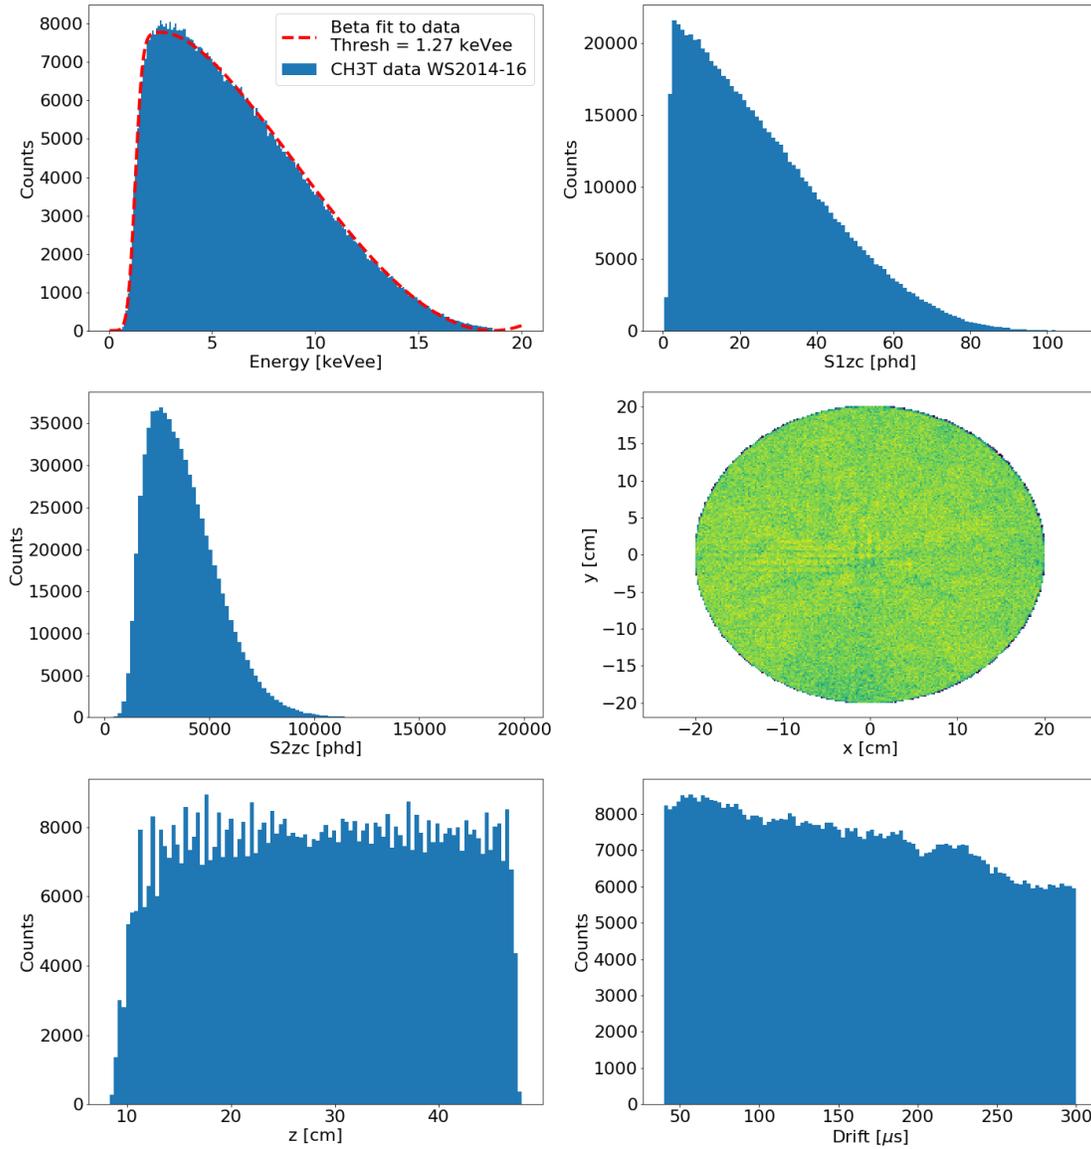


Figure 4.5: Distributions of energy, S1, S2, xy-position, z-position, and drift time for the ^3H WS2014–16 calibration. The top-left panel also includes a fit of data to the β decay spectrum with some energy threshold, fit to 1.27 keVee. The S1 and S2 areas are corrected using z-position only.

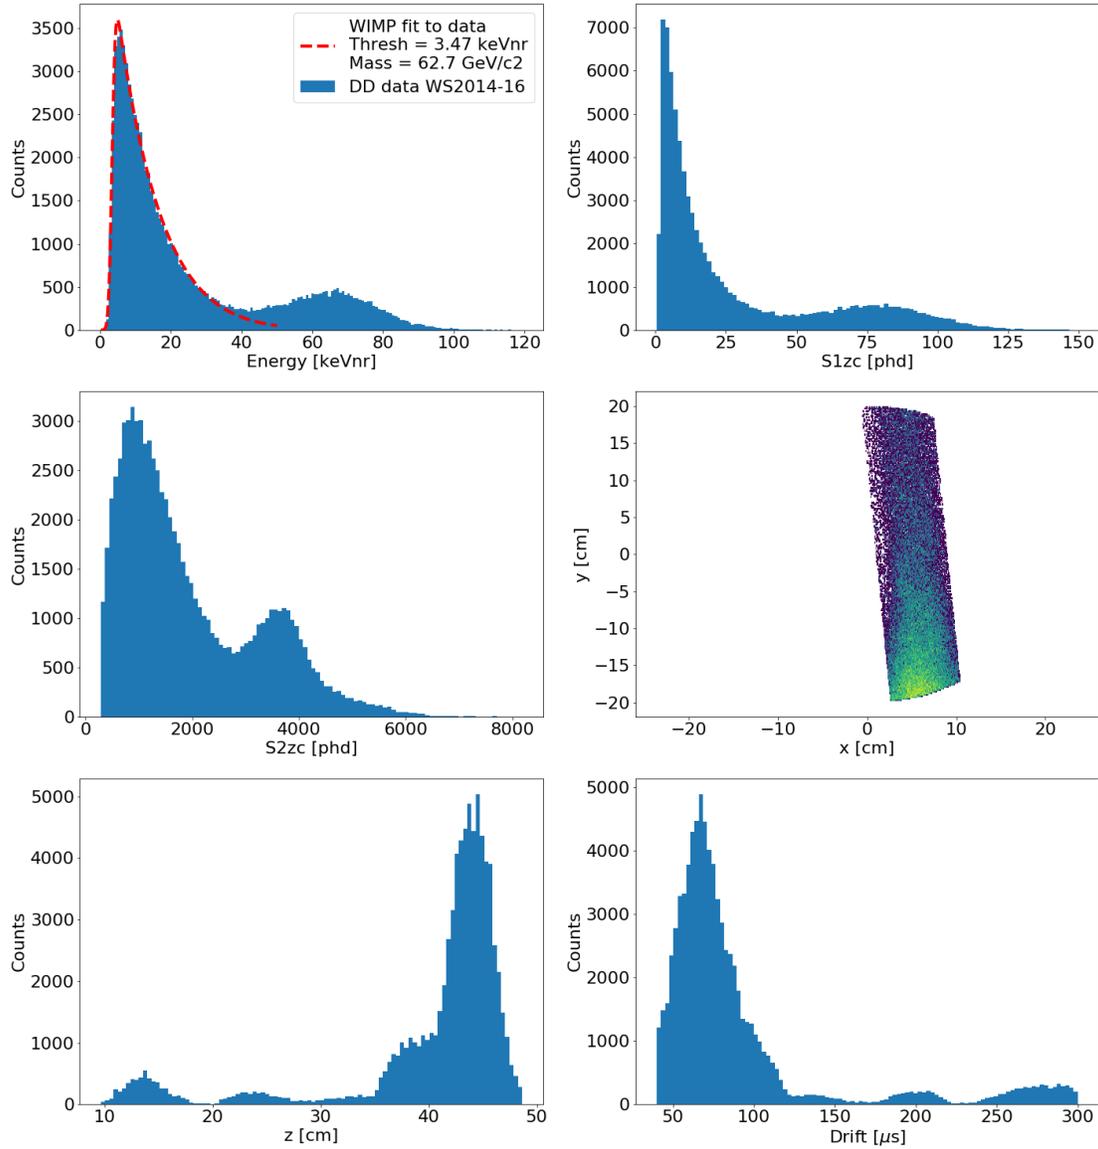


Figure 4.6: Distributions of energy, S1, S2, xy-position, z-position, and drift time for the D-D WS2014-16 calibration. The top-left panel also includes a fit of data to a WIMP recoil spectrum with some energy threshold. The WIMP mass and threshold are fit to 62.7 GeV/c² and 3.47 keVnr, respectively. The S1 and S2 areas are corrected using z-position only.

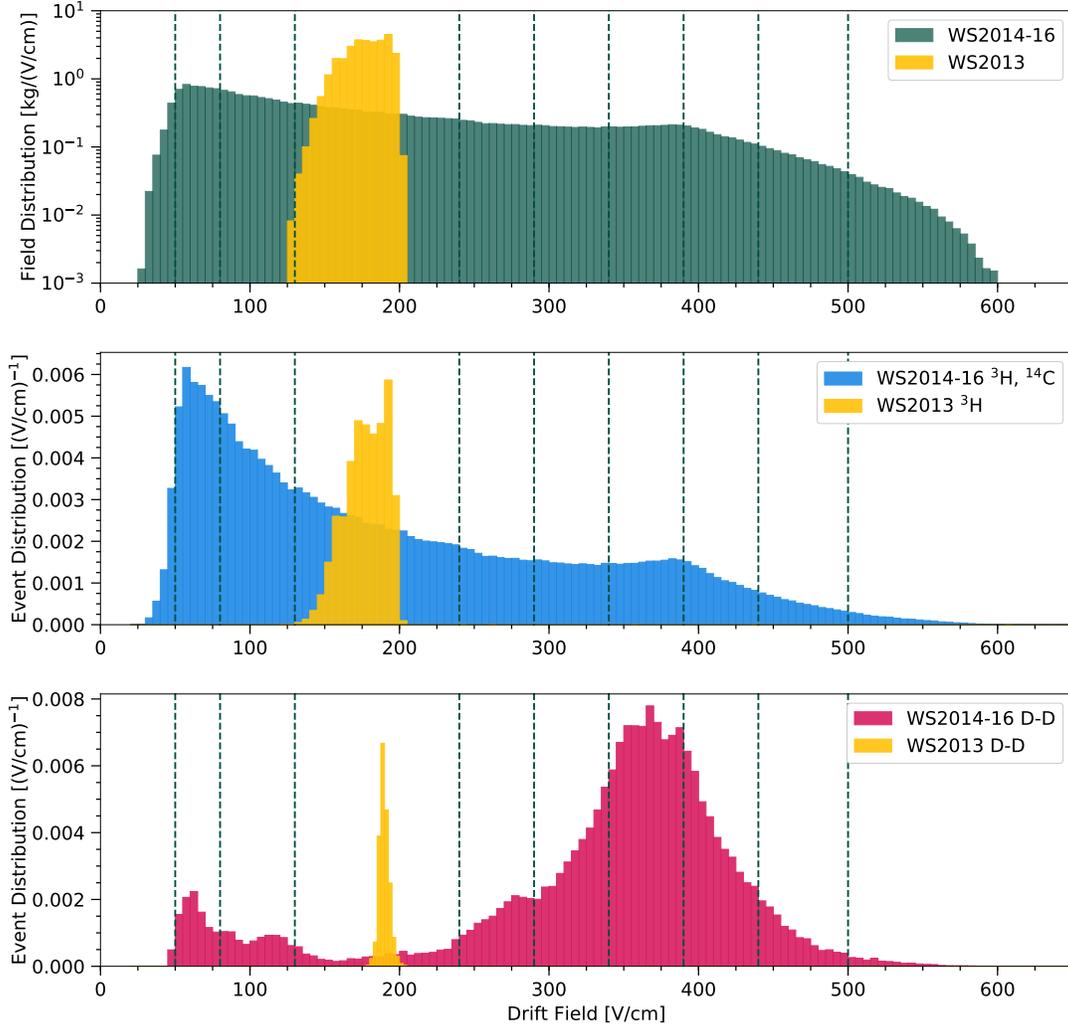


Figure 4.7: The distribution of drift fields in the LUX datasets. (*Top*) The mass distribution of field within the LUX fiducial volume. In this analysis, we define the WS2013 fiducial volume as $r < 20$ cm and $38 < (t_{drift}/\mu s) < 305$, and the WS2014–16 fiducial volume as $r < 20$ cm and $40 < (t_{drift}/\mu s) < 300$. (*Middle*) For each electronic recoil, the field at the recoil site is calculated using the results of [127], and we plot a normalized histogram of the results. The ^3H and ^{14}C datasets are combined for WS2014–16 because they both fill the entire detector volume, and thus have identical distributions. Black dashed lines are used to indicate the field bins used in Section 4.4. The WS2013 and WS2014–16 histograms are normalized separately in order to visualize the data effectively, so the relative heights of the blue and yellow histograms should not be considered an expression of the number of events in each dataset. (*Bottom*) The same as the middle panel, but for nuclear recoils.

profile likelihood ratio. However, we can learn about discrimination through simpler, more transparent techniques.

For each electronic recoil in the dataset, the LUX detector observes a single S1 signal, followed by a single S2 signal. As has been widely observed by liquid xenon experiments [143, 76, 144, 96, 64], one can plot these recoils on axes of $\log_{10}(S2c/S1c)$ vs. S1c to obtain a “band” of events. We will refer to this as the *ER band*, as is common in the literature. The same is true for nuclear recoils and the corresponding *NR band*. Figure 4.8 demonstrates this for our LUX WS2014–16 data. In these axes, we can trade-off between signal acceptance and background rejection as we would for a simple cut-and-count analysis. For example, we can select events below the NR band median, keeping 50% of NR events and admitting a small fraction of ER events. This is, in fact, precisely what we will do in Sec. 4.5.1, but it is important to realize that we can set any acceptance we want, as long as we are willing to accept the corresponding ER background. Note that in recent years, many experiments have been using $\log_{10}(S2c)$ as the vertical axis; see Fig. 1 of [26]. However, we use $\log_{10}(S2c/S1c)$ in this dissertation because it is skew-Gaussian distributed. The form of this distribution and its emergence from underlying processes will be discussed below and in Chapter 5.

The second strategy forces us to rely on pulse shape information. Since the singlet and triplet excimers have different decay times, we can look at the fraction of the S1 pulse in a short time at the beginning of the pulse; a previous LUX analysis [95] showed that 32 ns is the appropriate range. Nuclear recoils will have a higher “prompt fraction” than electronic recoils, shown in Fig. 4.9. Similarly to charge-to-light leakage, we can set some NR acceptance threshold and calculate the amount of ER leakage into the NR region.

Detailed analysis will occur in the remaining portion of this chapter, but we give some brief estimates here. Charge-to-light and pulse shape discrimination tend to give ER leakage fractions on the order of 0.2% and 26%, respectively, for an NR acceptance of 50%. Clearly, pulse shape discrimination is not very powerful on its own. The appropriate strategy is to combine the two variables into a single discriminant, as we will do in Sec. 4.5.2.

4.4 Electronic and Nuclear Recoil Bands

4.4.1 Electronic Recoils

We begin our analysis by calculating the ER band in LUX.

First, we need to account for the irregular energy spectrum of the dataset. The calculation of the ER band depends on the energy distribution of the underlying data. We choose to define the ER band using a flat energy spectrum, a definition that is easy to universalize. It also has the fortunate side effect of yielding a re-

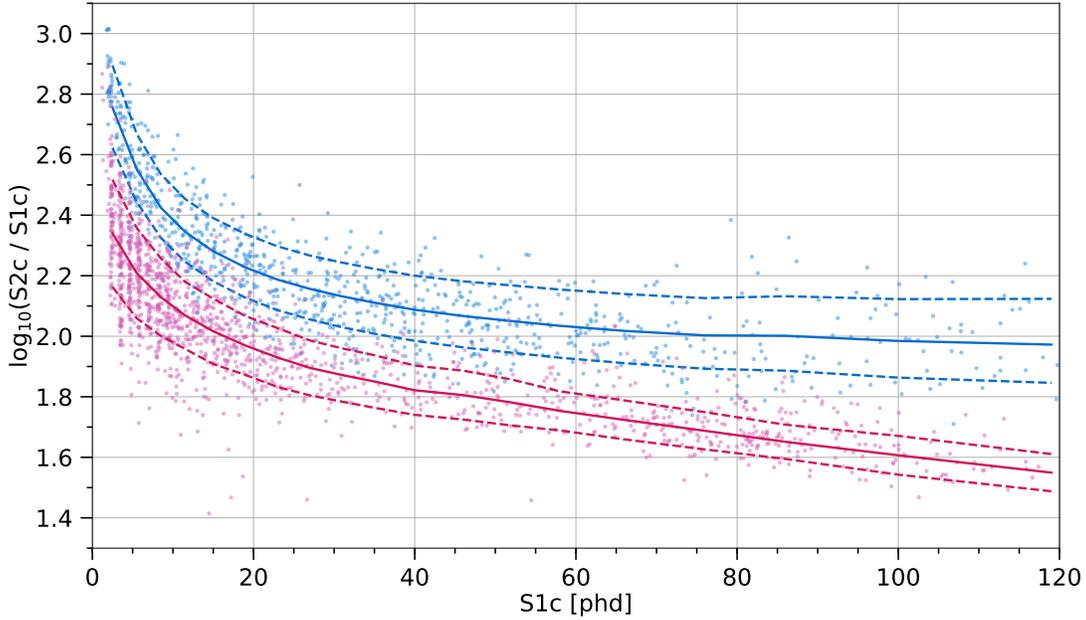


Figure 4.8: A sample of electronic and nuclear recoils, along with the associated bands. A randomly selected 1500 electronic (nuclear) recoils from WS2014–16 are shown in blue (red) dots, the median of the ER (NR) band is shown as a solid blue (red) line, and the 15.9th percentile and 84.1th percentile of the ER (NR) band are shown as dashed blue (red) lines.

sult that is appropriate for the dominant ^{220}Rn and ^{222}Rn backgrounds, which are relatively constant in energy over the range of energies relevant for dark matter direct detection. For each event, a weight is calculated such that the weighted energy distribution is proportional to $f(E)$ in Eq. 4.2, in which E is the recoil energy determined with Eq. 3.13.

$$f(E) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{E - E_\mu}{E_\sigma \sqrt{2}} \right) \right]. \quad (4.2)$$

The parameters E_μ and E_σ are determined by fitting the ^3H and ^{14}C energy distributions to their beta decay spectra multiplied by $f(E)$, as we saw earlier in Figs. 4.4 and 4.5. E_μ and E_σ are fit to about 1 keVee and 0.3 keVee, respectively. Effectively, E_μ is the energy threshold for measuring electronic recoils, and E_σ is the “width” of this threshold. In this way, the energy spectrum of the dataset is transformed into a flat distribution, apart from the threshold behavior at low energy. See Fig. 4.10 for a depiction of this weighting.

We then split the electronic recoil data into small bins of S1c. Within S1c bins, the distribution of $\log_{10}(\text{S2c}/\text{S1c})$ is often assumed to be Gaussian [139, 125, 55],

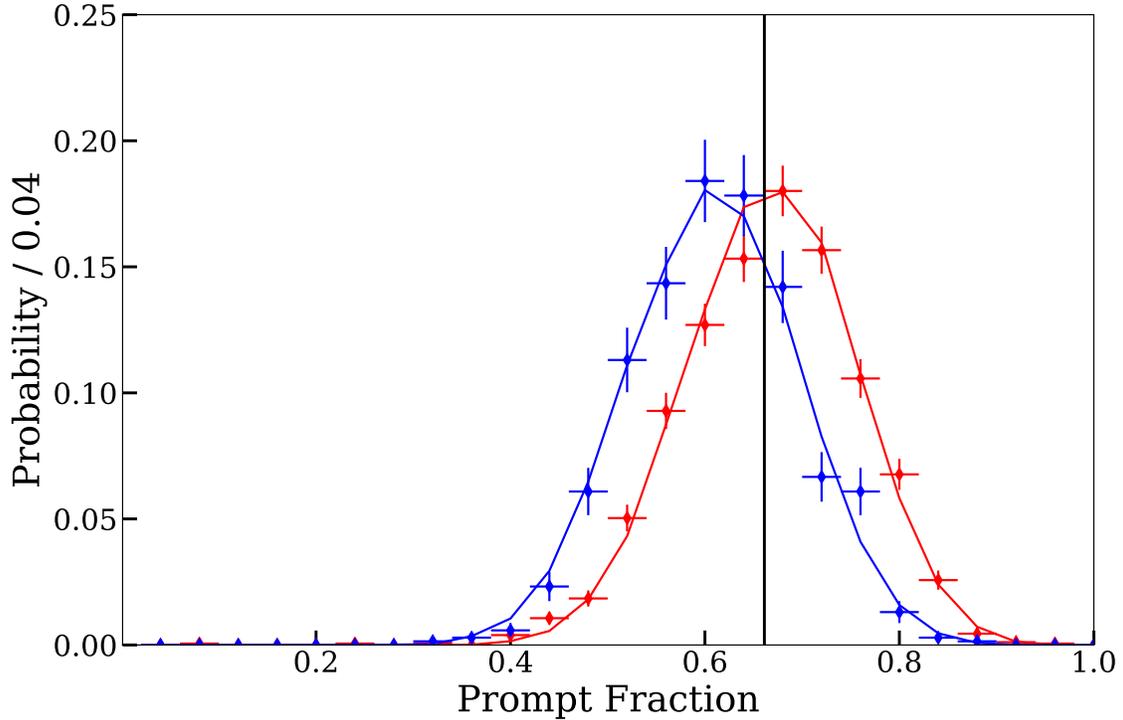


Figure 4.9: Distributions of prompt fraction, i.e. the fraction of light in the first 32 ns of the S1 pulse, from LUX data. NR and ER events are shown in red and blue, respectively. The black line shows the NR median. Copied from [95].

but we observe that a skew-Gaussian distribution is a better fit for the electronic recoil data, as also observed in [57]. A skew-Gaussian distribution follows the probability density function (PDF) in Eq. 4.3. This distribution is similar to a Gaussian distribution, if we identify ξ and ω with the mean and standard deviation. However, the skew-Gaussian distribution is modified by a parameter α , biasing the PDF toward higher values than a Gaussian PDF if $\alpha > 0$ and lower values if $\alpha < 0$. As a result, the mean μ and variance σ^2 of the skew-Gaussian distribution are given by Equations 4.4 and 4.5, respectively [145].

$$f(x) = \frac{1}{\omega\sqrt{2\pi}} e^{-\frac{(x-\xi)^2}{2\omega^2}} \left[1 + \operatorname{erf} \left(\frac{\alpha(x-\xi)}{\omega\sqrt{2}} \right) \right]. \quad (4.3)$$

$$\mu = \xi + \sqrt{\frac{2}{\pi}} \frac{\alpha\omega}{\sqrt{1+\alpha^2}}. \quad (4.4)$$

$$\sigma^2 = \omega^2 \left(1 - \frac{2}{\pi} \frac{\alpha^2}{1+\alpha^2} \right). \quad (4.5)$$

We will refer to α as the skewness parameter, but it is important to note that α does not correspond to the algebraic skewness of the distribution (i.e. the third standard-

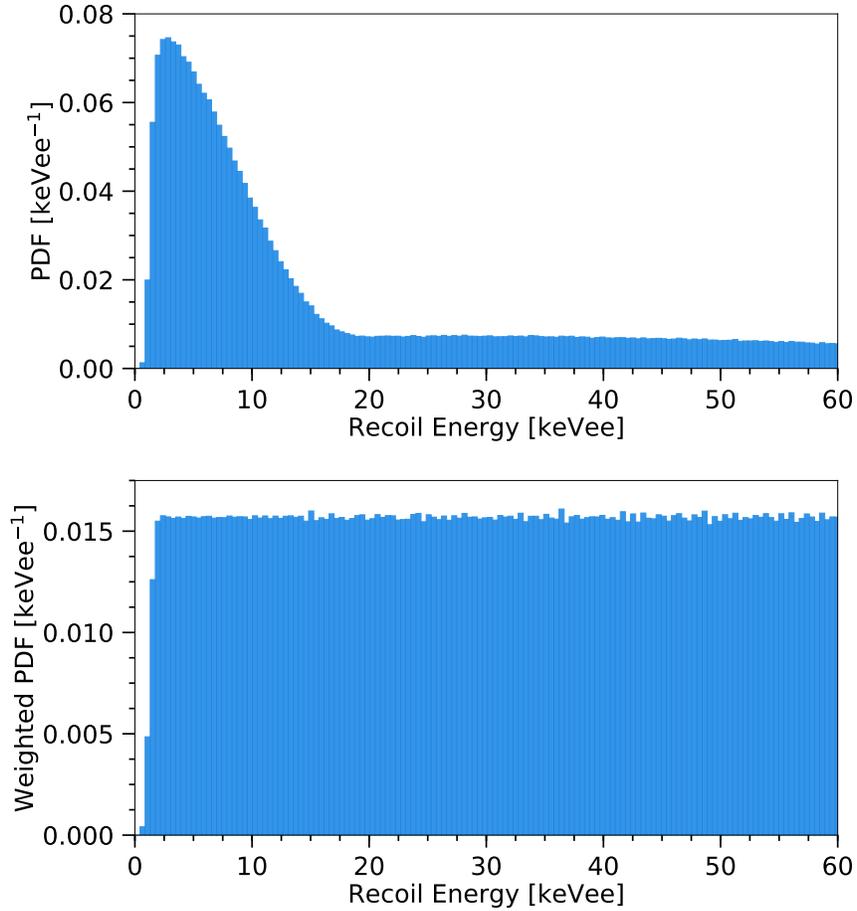


Figure 4.10: (*Top*) The full recoil energy spectrum of the WS2014–16 electronic recoil dataset, including ${}^3\text{H}$ and ${}^{14}\text{C}$ decays. (*Bottom*) The same energy spectrum, but with weights applied such that the spectrum is flat with a threshold at low energy.

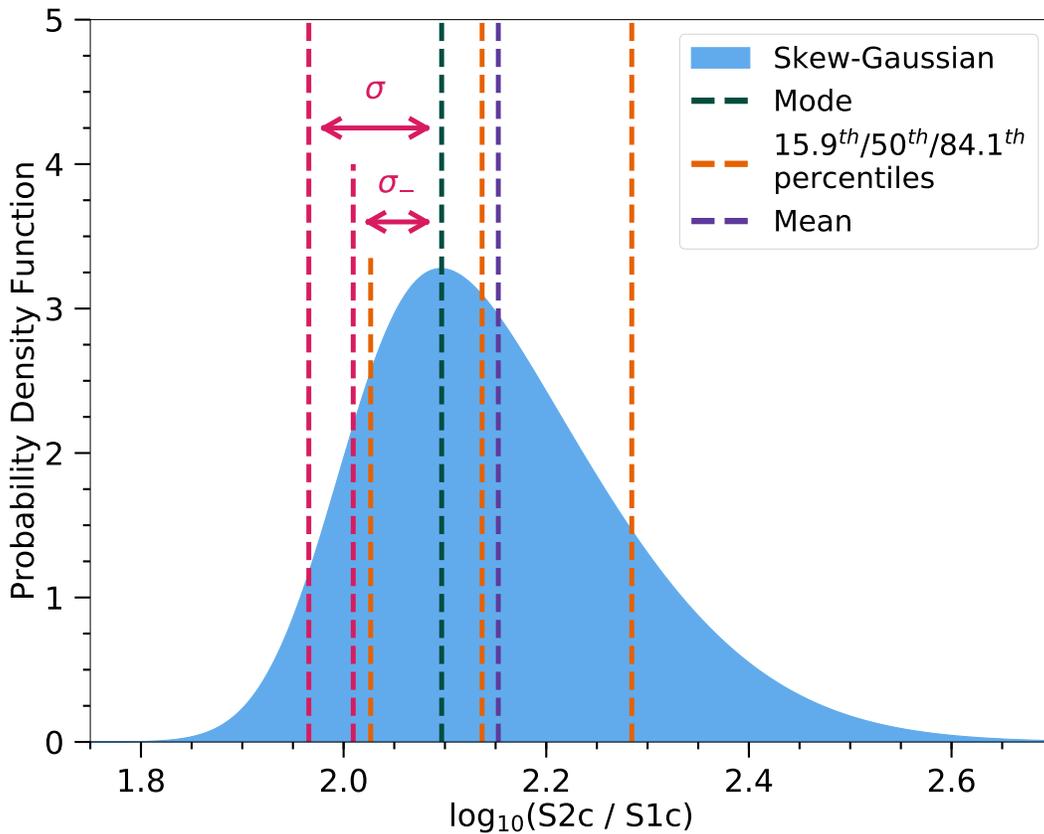


Figure 4.11: A skew-Gaussian distribution with $\xi = 2$, $\omega = 0.2$, and $\alpha_B = 3$. The mode, median, and mean of the distribution are shown, as well as the 15.9th and 84.1th percentiles. We also graphically show the standard deviation σ from Eq. 4.5 and σ_- from Eq. 4.6, relative to the mode of the distribution. In real $\log_{10}(S2c/S1c)$ data, α_B is typically smaller than 3, but a high skewness parameter is shown for ease of viewing.

ized moment). Furthermore, when referencing skew-Gaussian fits to distributions of $\log_{10}(S2c/S1c)$, we denote this parameter as α_B . The subscript “ B ” identifies this quantity as a trait of the ER (or NR) band.

In our energy range, electronic recoil data nearly always display positive skewness; $\alpha_B > 0$. Figure 4.11 shows the effects of positive skewness; the mean is greater than the median, and both are greater than the mode. We emphasize that positive skewness is not a statistical artifact, such as from Poisson statistics in the S1 signal; it seems to be the result of liquid xenon recombination physics, as we will explore in Chapter 5.

In each S1c bin, we fit the weighted histogram of $\log_{10}(S2c/S1c)$ to a skew-

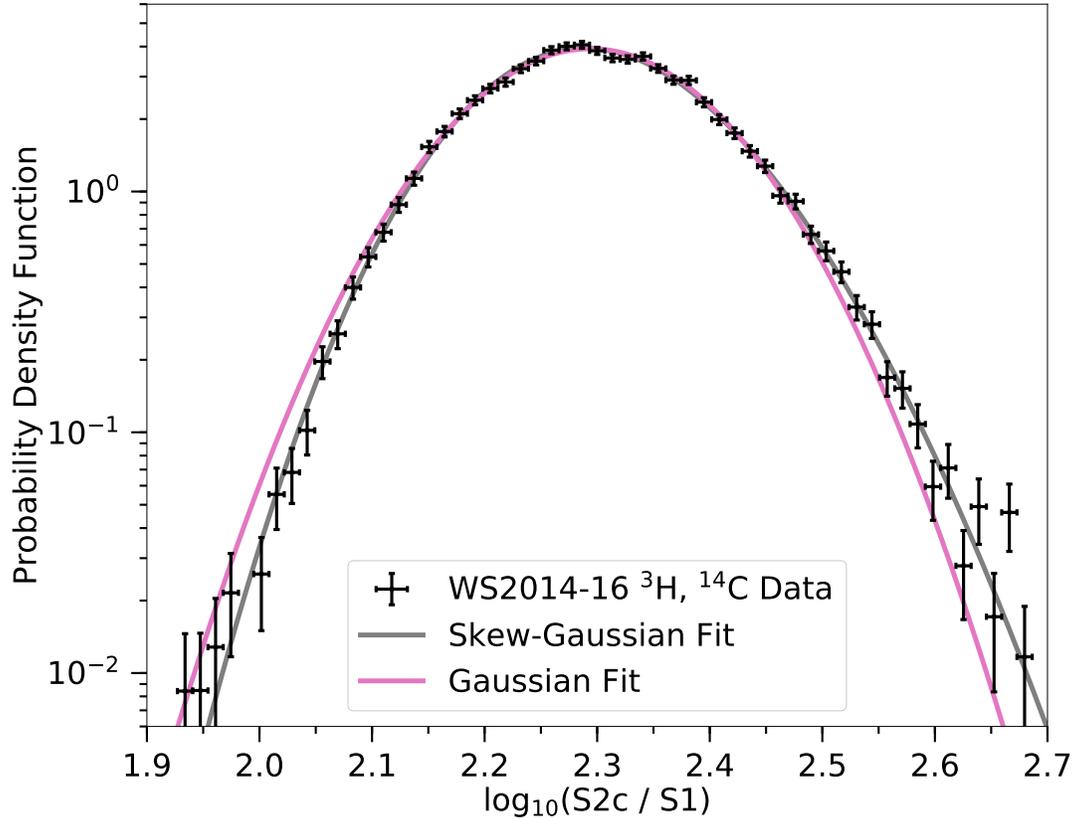


Figure 4.12: An example histogram of $\log_{10}(S2c/S1)$ for electronic recoil data, and associated fits to a skew-Gaussian and Gaussian distribution. The events have an S1 signal between 13 and 16 phd and a drift field between 80 and 130 V/cm; they are weighted based on their energy as described in Section 4.4.1. The best-fit Gaussian parameters are $\mu = 2.2941 \pm 0.0010$ and $\sigma = 0.1018 \pm 0.0010$, and the fit quality is $\chi^2/\text{DOF} = 142.148/51$. Meanwhile, the best-fit skew-Gaussian parameters are $\xi = 2.222 \pm 0.003$, $\omega = 0.128 \pm 0.002$, and $\alpha_B = 1.14 \pm 0.06$, and the fit quality is $\chi^2/\text{DOF} = 44.275/50$. See also Fig. 4.13 for a comparison across energies.

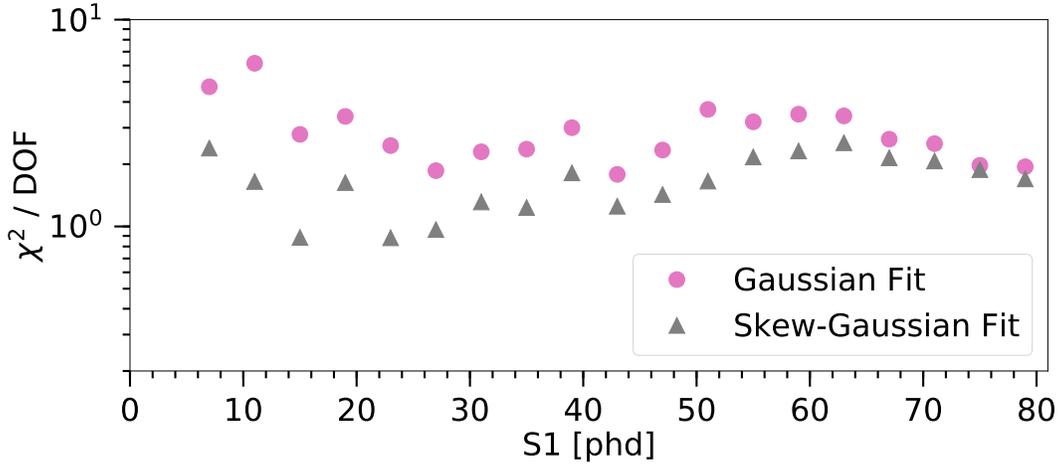


Figure 4.13: A comparison of Skew-Gaussian and Gaussian fit quality χ^2/DOF across S1. Fits are done to WS2014–16 electronic recoil data with drift field between 80 and 130 V/cm. See Fig. 4.12 for an example, using the S1=13–16 phd bin.

Gaussian distribution, using χ^2 minimization. Figure 4.12 shows an example of this fit; note that the skew-Gaussian fit more closely matches the data than the fit to a Gaussian. We show the superiority of the skew-Gaussian fit even more clearly in Fig. 4.13, where we compare χ^2 between the fits across a range of energies. The median of the distribution is easily extracted. The width is defined in two ways. First, the width of the total distribution σ is obtained by using Eq. 4.5. Second, we use Eq. 4.6 to define a quantity that we call σ_- , which is relevant for discrimination. In $\log_{10}(S2c/S1c)$ vs. S1c space, electronic recoils lie above nuclear recoils, so the leakage of electronic recoils into the nuclear recoil region is based only on the lower part of the $\log_{10}(S2c/S1c)$ distribution. Thus, σ_- serves as a measure of the portion of the width due only to downward fluctuations, and it is determined by the condition

$$\int_{m-\sigma_-}^m f(x) dx = 0.68 \int_{-\infty}^m f(x) dx , \quad (4.6)$$

where m is the mode of $f(x)$.

The uncertainties of the skew-Gaussian fit parameters, which are extracted from the χ^2 minimization, are used to estimate the uncertainties of the ER band median and width: $\delta(\text{median}) = \delta\xi$ and $\delta(\sigma_-) = \delta(\sqrt{\sigma^2})$.

The left half of Fig. 4.14 shows an implementation of our procedure—in this case, the ER band from WS2013 data.

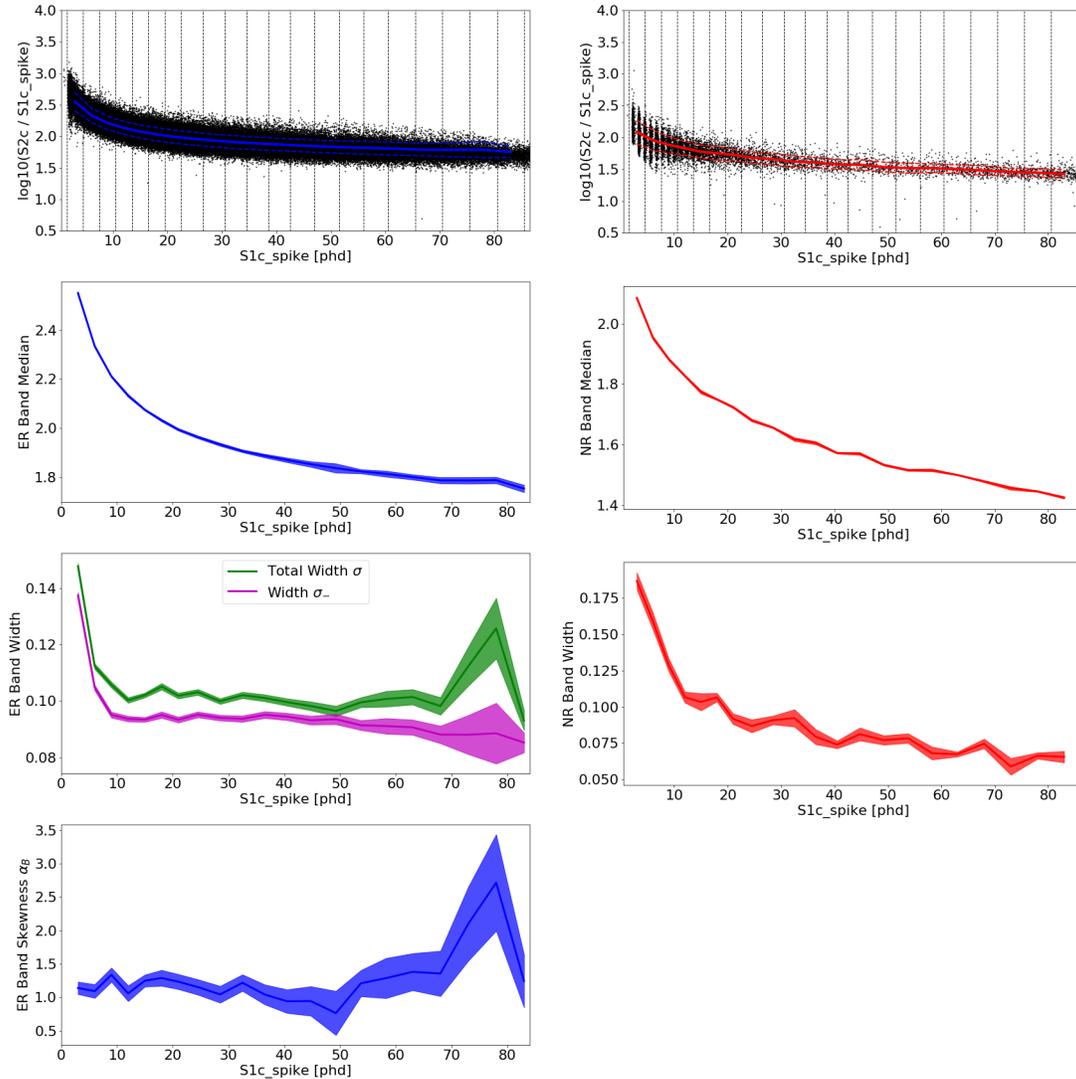


Figure 4.14: The ER (left) and NR (right) bands computed from WS2013 data, using the procedures in Sec. 4.4.1 and 4.4.2, respectively. The top panel shows the raw data in black dots with the measured band median and width overlaid in blue or red. The black dashed lines indicate the S1c bins used. The second-highest panel shows the band median alone, with the shaded region indicating 1σ uncertainty. The third-highest panel shows the band width, where for the ER band, we show both σ from Eq. 4.5 and σ_- from Eq. 4.6. For ER, the bottom panel shows the skewness parameter.

4.4.1.1 Aside: Statistical Strategy for Fitting $\log_{10}(\text{S2c}/\text{S1c})$

Before we move to nuclear recoils, we briefly comment on how the skew-Gaussian and Gaussian fits are performed. In our analysis, we report the results from χ^2 minimization. This is generally not the most accurate strategy, because the strategy assumes the number of events in each bin is a Gaussian-distributed variable. In fact, it is Poisson-distributed. Therefore, fits to histograms should be done via maximum likelihood estimation (MLE) with a Poisson estimator. If the number of events in a substantial number of $\log_{10}(\text{S2c}/\text{S1c})$ bins is small (<10 or so), the two strategies will give different results.

In our situation, the question is mostly academic, because we don't have many bins with few events. However, it is important to note that in fact, MLE with a Poisson estimator is not appropriate for our analysis. The crucial difference is the weighting done on the ER data to simulate a flat-energy spectrum. This weight is correlated with $\log_{10}(\text{S2c}/\text{S1c})$, the same variable we histogram. As a result, the assumption that the number of events in each bin is Poisson distributed is invalid. In fairness, the number of events is not Gaussian-distributed either, but for χ^2 minimization, we simply need to know the standard error on the number of events in a (weighted) bin, which is known. We don't need to know the PDF from which this is drawn, which is unknown. Therefore, χ^2 minimization is the appropriate technique for our fits.

We have confirmed this by fitting the ER data to skew-Gaussian histograms with three methods (named A, B, and C, respectively): χ^2 minimization on a weighted histogram, Poisson-ML optimization on a weighted histogram, and Poisson-ML optimization on an unweighted histogram. As expected, Methods A and B give the same fit parameters. However, the fit uncertainties are dramatically different, with Method B giving a much larger uncertainty. Meanwhile, Methods A and C give different fit parameters, as we would predict for a weight-change, but they give roughly the same uncertainties. This is strong evidence that by weighting the data, the Poisson distribution is no longer appropriate, causing the fit to become weaker.

To summarize, χ^2 minimization of a weighted histogram gives both the correct skew-Gaussian fit parameters and the correct size of the parameter uncertainties. This is relevant to our LUX analysis because it affects the statistical errors on our reported results, but it is even more relevant to our discrimination analysis in Chapter 6, where low-stats bins are common.

4.4.2 Nuclear Recoils

Nuclear recoils can be analyzed similarly to electronic recoils. One modification we make to the procedure outlined in Sec. 4.4.1 is that we eliminate the energy-based event weights. Instead, we use the unweighted D-D calibration data, which has a recoil energy spectrum similar to that of a 60 GeV/ c^2 WIMP. The other adjustment

for nuclear recoils is that in bins of S1c, we assume the distribution of $\log_{10}(S2c/S1c)$ is Gaussian. As will be described in Sec. 5.1, a skew-Gaussian distribution actually fits the NR data better, but we model the NR band as Gaussian for two reasons. First, due to the low statistics of the NR data, the skew-Gaussian fit often fails to converge or gives large errors on the fit parameters. Second, the Gaussian fit reproduces the same median and width as the skew-Gaussian fit, and these parameters have a greater impact on discrimination and sensitivity than the skewness itself. The uncertainties on the NR band median and width are simply the uncertainties on the Gaussian fit. The right half of Fig. 4.14 shows an implementation of this procedure—in this case, the NR band from WS2013 data.

We also note a small source of bias in the NR band calculation. To improve data quality, we have removed events with $S2 < 270$ phd (164 phd) in the WS2014–16 (WS2013) D-D data. In the lowest S1c bin, this removes up to 10% of events. When the Gaussian fit is performed, the best-fit mean and width are higher and lower, respectively, than they would be if the dataset contained events with a smaller S2 signal. The shift in these best-fit parameters is expected to be $< 2\%$, as estimated from simulation. The shift is small but could impact electronic recoil discrimination, as will be described in Sec. 4.5.1. This effect is not relevant for higher S1c bins in the nuclear recoil data and any S1c bins in the electronic recoil data, because all events have S2 signals significantly larger than the analysis threshold.

4.4.3 Variation with g_1

A key detector parameter in two-phase xenon dark matter experiments is the prompt light collection gain g_1 , which is primarily dictated by the detector geometry, the reflectivity of the inner surfaces, and the quantum efficiency of the PMTs. In WS2013, the average value of g_1 was 0.117 [139], while in WS2014–16, it varied from 0.0974 to 0.0994. The time dependence of g_1 could be caused by varying impurity concentration in the Xe bulk, changes in wire grid reflectivity, or decreases in PMT quantum efficiency. We expect g_1 to have a strong impact on discrimination; as more light is collected, the S1 signal will grow in magnitude, and the relative size of S1 fluctuations will decrease. Thus, g_1 should be positively correlated with discrimination power.

This is an effect we can observe in LUX through a novel procedure. For each event, the S1c signal is a sum of the signals in each of LUX’s 122 PMTs (adjusted for position-dependent and PMT-dependent effects). By adding together the pulses in only a fraction of the PMTs, we are able to artificially reduce g_1 . We use $^{83\text{m}}\text{Kr}$ WS2013 calibration data [146] to determine the effective g_1 for a given subset of PMTs. The $^{83\text{m}}\text{Kr}$ decay is a two-step process, emitting 32.1 and 9.4 keV conversion electrons, and the time between the two decay steps Δt is exponentially distributed with a half-life of 154 ns [75]. However, our analysis only uses events in which the two light signals are merged—generally true for events in which $\Delta t < 1200$ ns. We

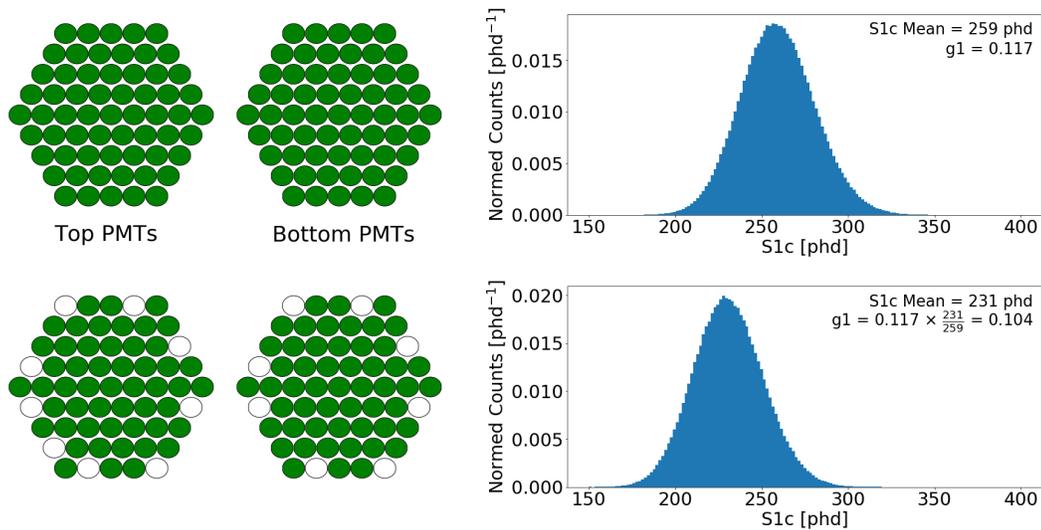


Figure 4.15: A demonstration of our g_1 reduction process. The top panel shows the S1c peak from WS2013 ^{83m}Kr data (180 V/cm drift field) at 259 phd. All PMTs are used, and LUX has previously measured $g_1 = 0.117$. In the bottom panel, we remove 17 PMTs (9 in the top array, 8 in the bottom array) and recalculate the S1c area for each event using the same ^{83m}Kr data, but summing pulses only in the remaining PMTs. We find the new S1c peak at 231 phd. Accordingly, we find that we have reduced g_1 to 0.104.

can thus treat the ^{83m}Kr decay as monoenergetic with a single, field-dependent S1c and S2c peak. The mean value of the ^{83m}Kr S1c peak is reduced when we add the signals in a subset of the 122 PMTs, relative to its value when using the full LUX detector. The reduction in the value of the S1c peak is proportional to the reduction in g_1 . For example, one PMT configuration has 105 PMTs, and when S1c is recalculated for all ^{83m}Kr events using only the signals detected by these 105 PMTs, the average S1c is reduced by 11% relative to adding the signals in all 122 PMTs. Thus, we infer that the effective g_1 obtained by using these 105 PMTs is $g_1 = 0.89 \times 0.117 = 0.104$. A visual demonstration of this is depicted in Fig. 4.15.

We isolate the effect of g_1 on the ER and NR bands by considering only WS2013 data, which have a uniform drift field. First, ten PMT configurations are chosen, and the corresponding g_1 values are calculated. We intentionally choose PMT configurations so the resulting g_1 values are evenly distributed between 50% and 100% of g_1 for the full detector. For each configuration of PMTs, we calculate new S1c values for each event in the WS2013 ^3H and D-D data. The S1 signal is obtained by adding together the signals from only the PMTs in that subset, and this is translated to S1c with the same position-dependent correction factor used in the analysis of all 122 PMTs. Then, we recalculate the ER and NR band.

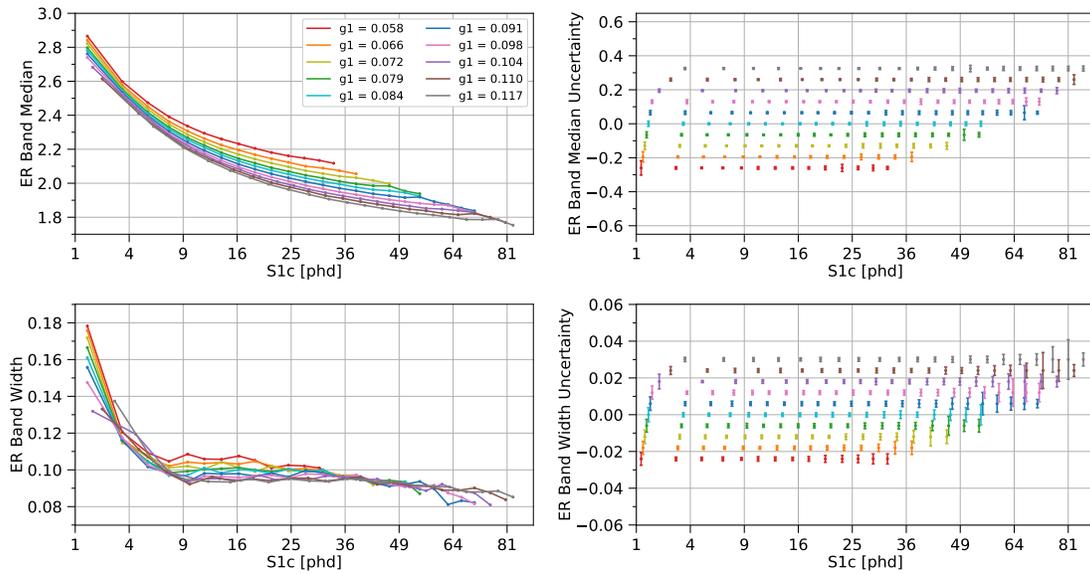


Figure 4.16: The median and width σ of the ER band for several values of g_1 , using WS2013 data. The left plots show measurements, and the right plots display error bars corresponding to these measurements. The S1c axis is proportional to $\sqrt{S1c}$ for ease of visualization. In each row, the y -axes have the same range; the size of the error bars on the right plot can be directly translated to the points on the left plot. For ease of visualization on the right plots, the S1c values are slightly shifted relative to their true value, and the error bars are centered at a different y -value for each g_1 . Note the S1c range varies for each g_1 because as g_1 decreases, the ^3H end point in S1c space decreases.

The results for the ER band are shown in Fig. 4.16. As g_1 increases, the median of the ER band shifts down; this is a fairly straightforward result, because a larger g_1 implies a larger S1c and thus a lower $\log_{10}(S2c/S1c)$. Also, as g_1 increases, the absolute ER band width decreases, particularly for S1 values less than 30 phd. This also matches our expectations, because as the light collection increases, the relative size of the fluctuations in the number of photons detected decreases. Note that the leftmost point for $g_1 = 0.117$ in the bottom panel of Fig. 4.16 appears to be an outlier, showing a different behavior than the other measurements. However, it is not an outlier. Instead, this appearance is due to the changing energy-S1c correspondence as g_1 varies. Above 30 phd, the shrinking of the ER band width with g_1 plateaus, and we can account for this with three explanations. First, since ^3H has an end point in our region of interest, the changing g_1 changes the maximum S1c, which excludes certain curves at high energy. Second, the number of events in each S1c bin decreases as we near the end point, making the error bars larger and reducing our sensitivity to any small differences. Third, as the number of photons

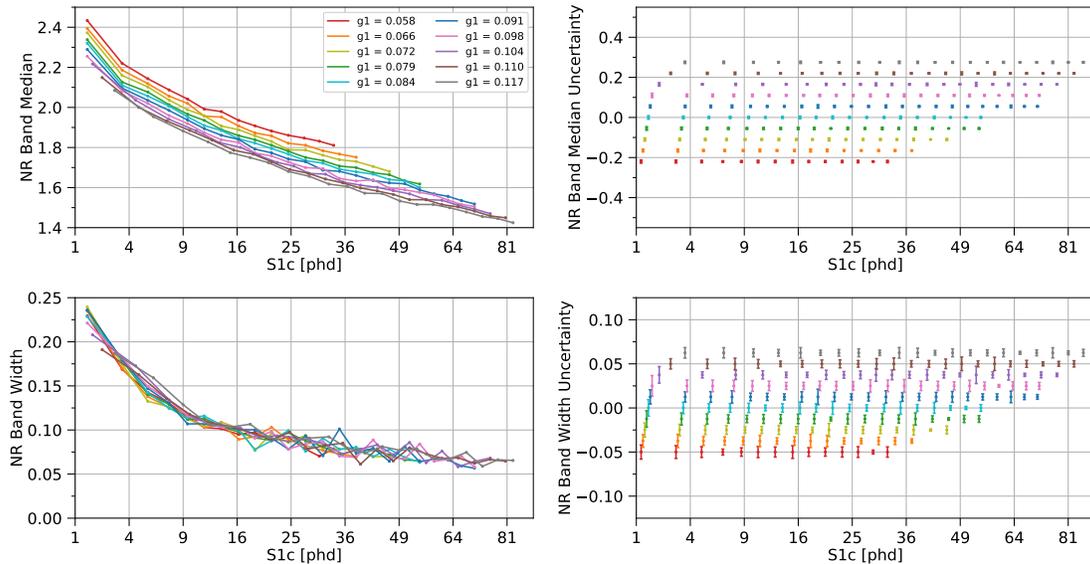


Figure 4.17: The median and width of the NR band for several values of g_1 , using WS2013 data. The left plots show measurements, and the right plots display error bars corresponding to these measurements. The S1c axis is proportional to $\sqrt{S1c}$ for ease of visualization. In each row, the y -axes have the same range; the size of the error bars on the right plot can be directly translated to the points on the left plot. For ease of visualization on the right plots, the S1c values are slightly shifted relative to their true value, and the error bars are centered at a different y -value for each g_1 . Note the S1c range varies for each g_1 because as g_1 decreases, the D-D end point in S1c space decreases.

detected increases, the relative fluctuations in the S1 signal become smaller, and the total ER band width is dominated by other g_1 -independent fluctuations such as recombination.

The variation of the NR band with g_1 , shown in Fig. 4.17, is similar to that of the ER band. It shifts down with g_1 for straightforward reasons; as light collection increases, $\log_{10}(S2c/S1c)$ must decrease. The impact of g_1 on the NR band width is more muted, however.

4.4.4 Variation with drift field

Another crucial detector parameter is the drift field. As described in Section 4.2.3, WS2014–16 saw significant field variation in the liquid xenon volume; we can use this to study the effect of electric field on the ER and NR bands.

First, we separate the electronic recoil and nuclear recoil data into bins based on the field at the recoil site. For WS2014–16 data, the bin boundaries are [50,

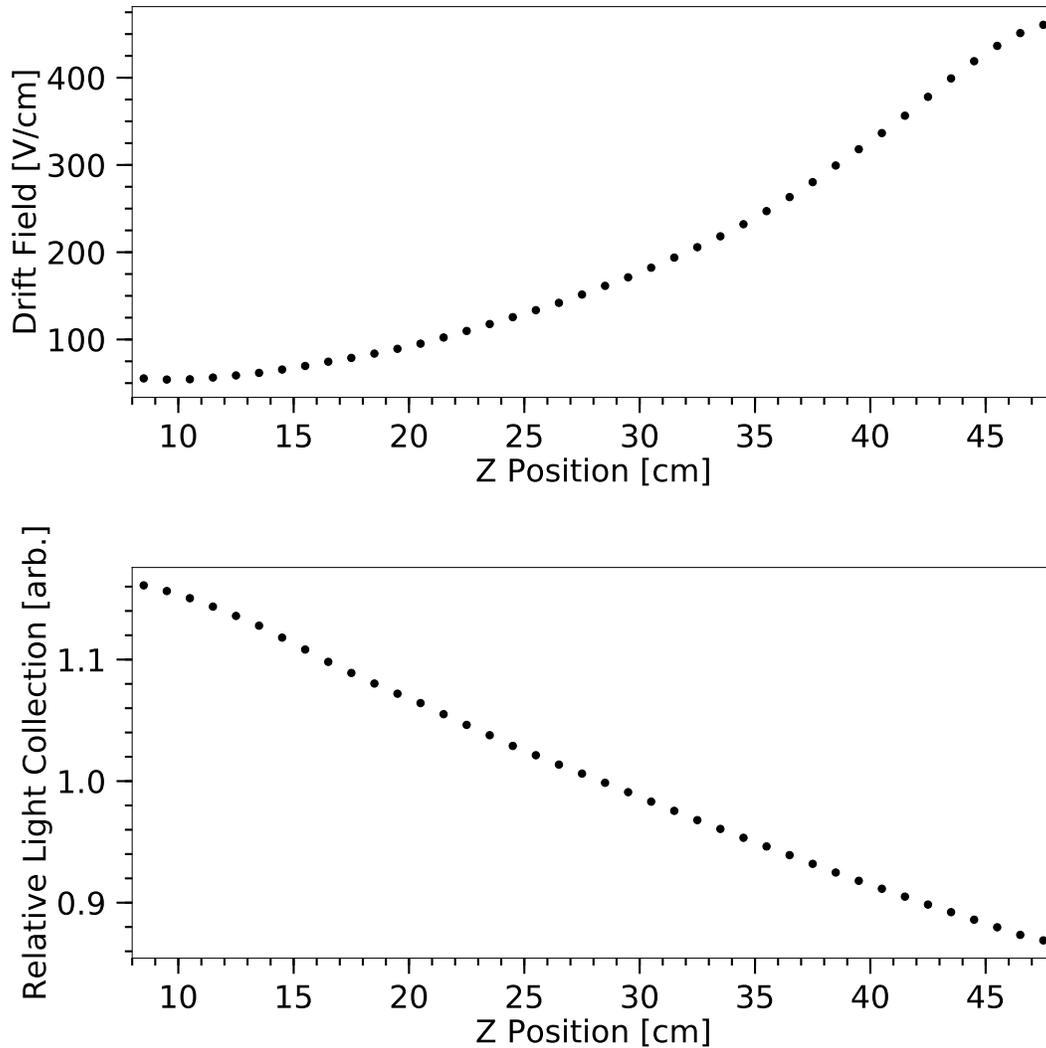


Figure 4.18: The average drift field and relative light collection efficiency vs. z , using WS2014–16 electronic recoil data. The field strength varied throughout the course of the run, but here, we average over the entire two-year period.

80, 130, 240, 290, 340, 390, 440, 500] V/cm. The bins were chosen to be wide enough such that the number of events in each bin is sufficient for the analysis, but narrow enough to yield precise measurements of field effects; they are overlaid over histograms of the data in Fig. 4.7. For WS2013 data, the data are all collected into a single field bin, leading to nine total field bins.

In the WS2014–16 data, electric field variation is highly correlated with variation in light collection, because they are both dependent on z -position, as seen in

Fig. 4.18. Higher regions (lower drift time) of the LUX detector have higher drift field, but also lower light collection due to total internal reflection at the liquid-gas interface. This causes photons produced near the top of the detector to, on average, pass through more liquid xenon and encounter the PTFE surface more times than photons produced near the bottom of the detector. Thus, we need to adjust the light collection efficiency in order to isolate field dependence. We do this in each field bin through the PMT removal procedure described in Sec. 4.4.3. The adjustment in light collection, relative to the top of the LUX detector, ranges from 0.787 to 1.000 in WS2014–16 and is equal to 0.744 for WS2013. This adjustment effectively accounts for the z -dependent position corrections, and so, in this portion of the analysis, we remove position corrections from the S1 variable.

Within each field bin, we calculate the median and width of the ER and NR bands. For the WS2013 results only, we adjust the band medians so that they are consistent with g_2 in WS2014–16: $g_2 = 12.1$ for WS2013 [139], and the average $g_2 = 19.085$ for WS2014–16. Thus, the WS2013 band medians are shifted up by $\log_{10}(19.085/12.1) = 0.198$.

The results for the ER band are shown in Fig. 4.19. As the drift field increases, the ER band median and width both increase convincingly. The former effect is expected; a plethora of data [76, 143, 140] shows that increasing electric field is correlated with a higher charge signal and smaller light signal, due to lower recombination. The increasing width is a consequence of this—with a lower light signal, the relative size of S1 fluctuations will increase. Crucially, as we will explore later, the width of the ER band is a major factor in discrimination. We note that the outlier width point at 35 phd for the 440–500 V/cm bin is the result of our skew-Gaussian fit converging to a negative skewness, whereas most fits converge to a positive skewness. It is not symptomatic of any trend; in fact, if we consider σ rather than σ_+ , this point is no longer an outlier.

The variation of the NR band with electric field is shown in Fig. 4.20. The behavior of the NR band as we vary electric field is quite different to that of the ER band, indicating fundamental physical differences in these interactions. Primarily, the NR band is substantially less sensitive to electric field than the ER band, a finding that has been seen by others [76]. The median moves up with increased electric field, in a statistically significant but small effect. The width has nearly no discernible variation from the electric field, except that the two highest field bins (390–440 V/cm and 440–500 V/cm) appear to have the largest widths across the entire energy range.

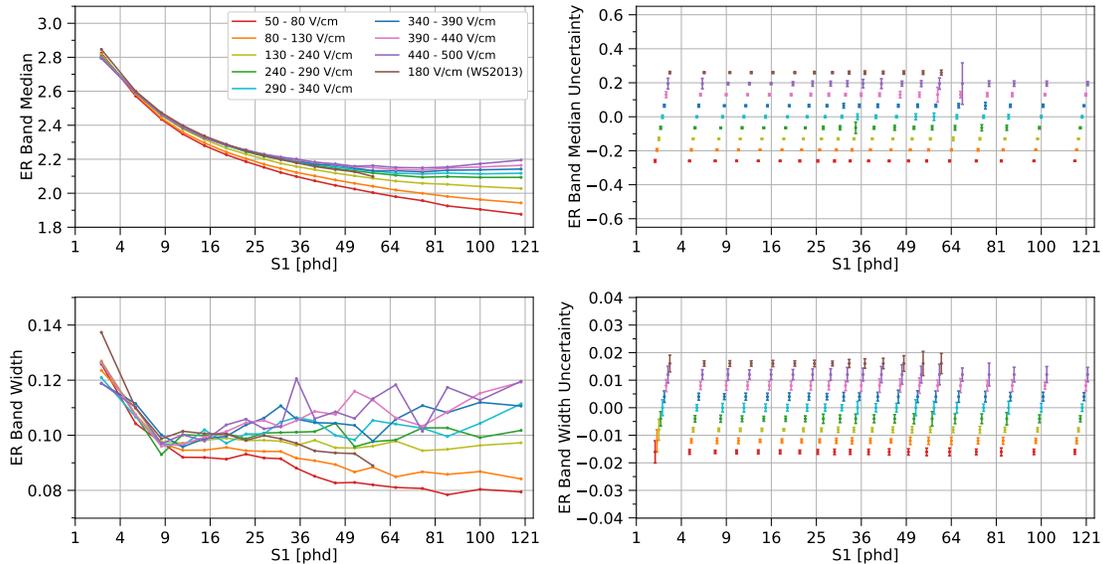


Figure 4.19: The median and width σ of the ER band for several drift fields. The left plots show measurements, and the right plots display error bars corresponding to these measurements. The S1 axis is proportional to $\sqrt{S1}$ for ease of visualization. The ER band for WS2013 is adjusted so g_2 is consistent for the WS2013 and WS2014–16 results. In each row, the y -axes have the same range; the size of the error bars on the right plot can be directly translated to the points on the left plot. For ease of visualization on the right plots, the S1 values are slightly shifted relative to their true value, and the error bars are centered at a different y -value for each field bin.

4.5 Leakage and Discrimination Results

4.5.1 Charge-to-light discrimination

Studying the electronic and nuclear recoil bands separately is informative, but the discrimination power is the critical figure-of-merit for studying how detector parameters affect sensitivity. Figure 4.8 shows charge-to-light discrimination graphically; the electronic recoils lie above nuclear recoils in these axes. We refresh the reader’s memory from Chapter 3 about the causes of this. First, the initial exciton-to-ion ratio varies: it is approximately 1 for nuclear recoils [76, 111, 112] and 0.2 for electronic recoils [101, 102, 103]. Second, recombination varies. Electronic recoils follow the Doke-Birks model [106] at high energies (<10 keVee) [107, 96], in which recombination is based on ionization density; they follow the Thomas-Imel model [108] at lower energies, in which thermal and diffusive effects smear out the track, and recombination can be considered to take place entirely in a small box of size $O(\mu\text{m})$.

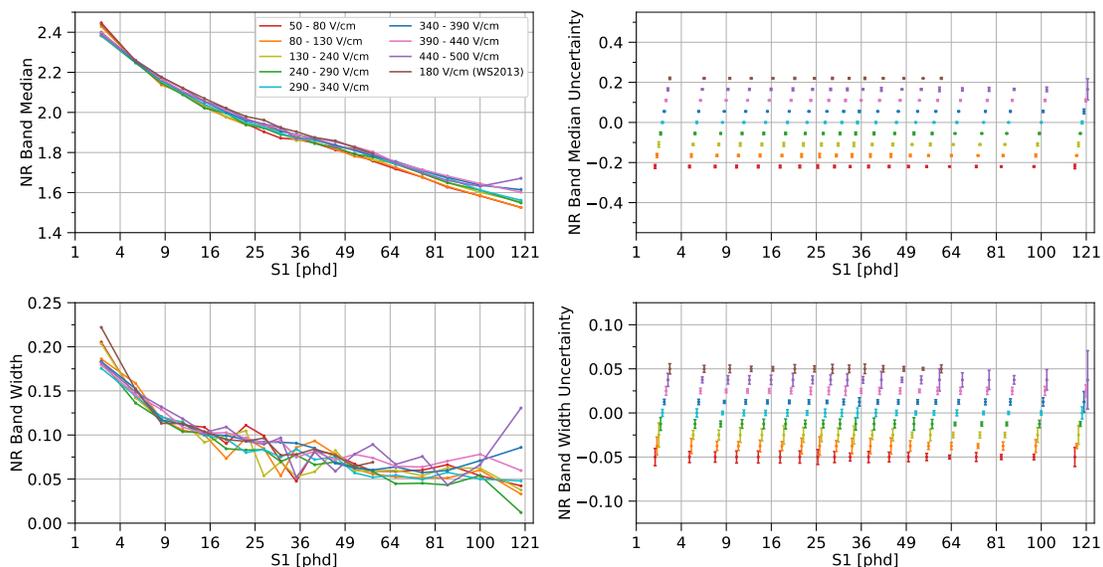


Figure 4.20: The median and width of the NR band for several drift fields. The left plots show measurements, and the right plots display error bars corresponding to these measurements. The S1 axis is proportional to $\sqrt{S1}$ for ease of visualization. The NR band for WS2013 is adjusted so g_2 is consistent for the WS2013 and WS2014–16 results. In each row, the y -axes have the same range; the size of the error bars on the right plot can be directly translated to the points on the left plot. For ease of visualization on the right plots, the S1 values are slightly shifted relative to their true value, and the error bars are centered at a different y -value for each field bin.

Nuclear recoils are governed solely by the Thomas-Imel model at our energies of interest [111]. Thus, at these lowest energies, electronic recoils are disparate from nuclear recoils in their initial exciton-to-ion ratio and the fraction of energy lost to heat.

Within each S1c bin, we can calculate the charge-to-light leakage fraction (or alternatively, its inverse: the discrimination power) at 50% nuclear recoil acceptance in two ways. First, we can count the number of weighted electronic recoils falling below the NR band median. We take the uncertainty on the leakage fraction to be the Poisson error. Second, we can integrate the skew-Gaussian ER distribution below the NR band median. The uncertainty here is found by propagating the errors in the ER band skew-Gaussian fit and the NR band Gaussian fit. The two methods have been confirmed to be consistent with each other, except in the lowest S1c bin where, due to PMT and threshold effects, the distribution of $\log_{10}(S2c/S1c)$ does not match a skew-Gaussian. As an example, see Fig. 4.21, comparing the two methods. The latter method allows us to calculate the leakage fraction even if the

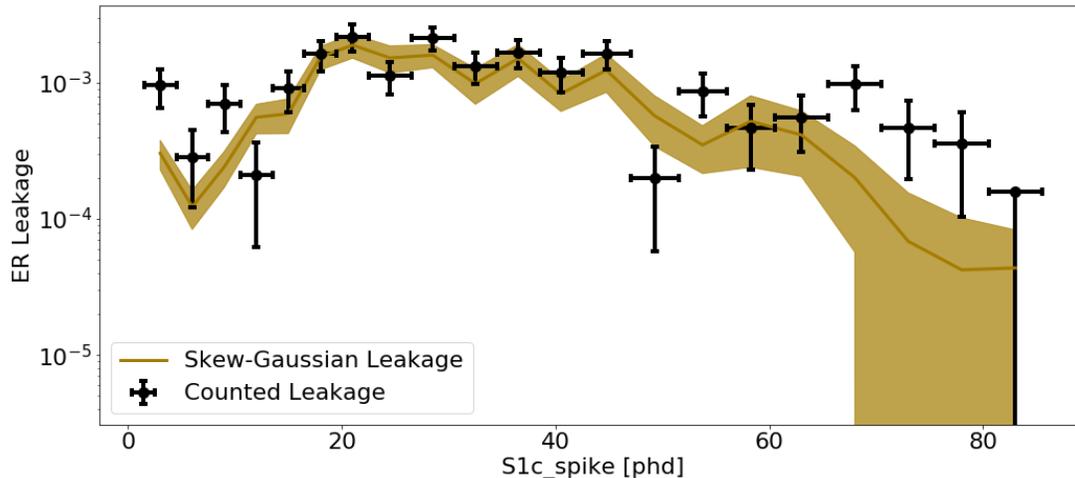


Figure 4.21: The ER leakage in S1c bins for WS2013, using two methods: extrapolating the skew-Gaussian fit below the NR band median, and counting the number of ER events below the NR band median. The two are consistent except in the lowest S1 bin, a finding consistent across our analysis.

number of events in the bin is too low to count the leaked events, so we use it except where specifically mentioned.

In contrast to this consistency between methods, when we use a Gaussian (not skew-Gaussian) fit for the ER band, we get a leakage estimate that is about twice as high! This is due to the positive skewness of the ER band; events are pushed further away from NR parameter space. In other words, the common technique of fitting $\log_{10}(S2c/S1c)$ to a standard Gaussian gives overly pessimistic leakage results.

Before presenting our results, we discuss sources of potential systematic uncertainty on the leakage fraction. First, g_1 and g_2 are uncertain at the 1–3% level; thus, the positions of the ER and NR bands are uncertain at a similar scale. However, this uncertainty will not lead to a systematic error on the leakage fraction, because if the g_1 or g_2 measurement is offset from its true value, the ER and NR bands will move together by the same amount. An error in g_1 could affect the ER band width and thus the electronic recoil leakage fraction, but this effect is insignificant at the level of the uncertainty on g_1 . Second, when we decrease g_1 by using a subset of LUX PMTs, this procedure introduces an extra systematic uncertainty on g_1 . This uncertainty has been calculated and is $<0.1\%$, so it is negligible. Third, the binning of $\log_{10}(S2c/S1c)$ will introduce a bias on the ER skew-Gaussian and NR Gaussian fits. We have experimented with different levels of binning and observed that the leakage fraction is not significantly affected by our choice of binning. The only effect of this choice is whether the ER skew-Gaussian fit converges. Fourth, in the lowest S1c bin only, the NR band median is biased slightly upward due to the finite S2

analysis threshold (see Sec. 4.4.2 for details). This means that the estimated leakage fraction is higher than it would be in a zero-threshold analysis. Using simulations, we have determined that this effect is smaller than the uncertainties on the leakage fraction from statistics and Gaussian-fitting the nuclear recoil data. However, an experiment with a higher S2 threshold could be significantly affected by the shift in the NR band, so caution should be taken if extrapolating our lowest-energy results to such an experiment.

4.5.1.1 Variation with g_1

Calculating the leakage function in S1c bins with g_1 variation gives the results in Fig. 4.22. The most striking effect is that as g_1 increases, the leakage decreases. Furthermore, it shares some features with the bottom of Fig. 4.16, namely that the effect is strongest below 25 phd. This suggests that the improvement in discrimination is due to the shrinking of the ER band width. Above 25 phd, the improvement in discrimination with g_1 is absent or suppressed, but we do not necessarily conclude that g_1 has no effect on discrimination at high energies. Low ^3H statistics at energies near the 18.6-keV end point give rise to large uncertainties on the leakage fractions. As mentioned, the real (counted) leakage does not match the skew-Gaussian leakage in the lowest S1c bin only; the ratio between the two is plotted in Fig. 4.23.

Another way to look at xenon discrimination power is the total leakage in a wide energy range. Using the full set of PMTs and the WS2013 data, we find that the leakage fraction from 0–50 phd, i.e. the WIMP search region used in the 2013 limit [63], is about 0.1%.¹

If we artificially remove PMTs as described in Section 4.4.3, we can still calculate the total leakage, but there is an extra step required due to the ^3H end point. Since the end point is around 85 phd, any setup in which the relative light collection is less than $50/85 = 0.59$ of the full detector will show bizarre behaviors in which the ER band cannot be calculated properly. Thus, we shift the maximum S1c to be proportional to g_1 ; e.g. $S1c_{\text{max}} = 50$ phd for $g_1 = 0.117$, $S1c_{\text{max}} = 25$ phd for $g_1 = 0.0585$, etc. This effectively keeps the maximum energy constant at 9.7 keVee. The results are shown in Fig. 4.24, and they show convincingly that as light collection increases, discrimination improves. The total leakage fraction varies slightly based on the method we use. If we count the weighted number of electronic recoils falling below the NR band median, we generally get a higher leakage than if we use the skew-Gaussian fits; the reverse is true for the lowest g_1 values. This discrepancy is almost entirely due to the discrepancy in the lowest S1c bin.

¹Our measurement of 0.1% is different than the 0.2% reported in [125]. The difference is due to our use of a skew-Gaussian distribution, as well as our energy weighting.

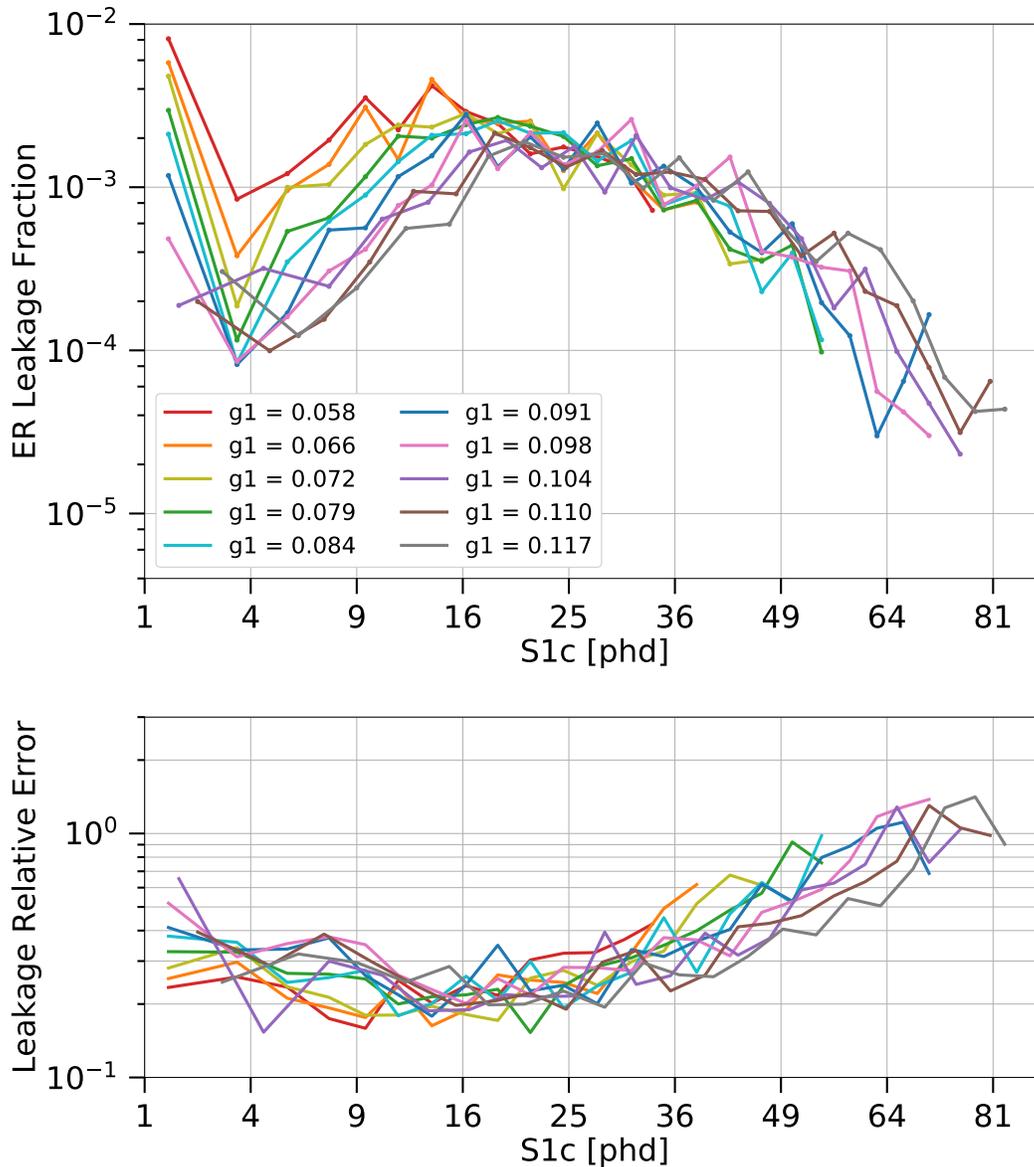


Figure 4.22: (*Top*) The ER leakage fraction for a flat energy spectrum in S1c bins, for various values of g_1 , calculated from a skew-Gaussian extrapolation of the ER band below the NR band median. The S1c axis is proportional to $\sqrt{S1c}$ for ease of visualization. This calculation is consistent with the real counted leakage, except in the lowest S1c bin; see Fig. 4.23 for a comparison of the two calculations in this bin. (*Bottom*) The relative error on these leakage fraction values, defined as: $\text{leakage_fraction_error} / \text{leakage_fraction}$. Note that the leakage relative error can be greater than 1, indicating the leakage fraction is consistent with 0.

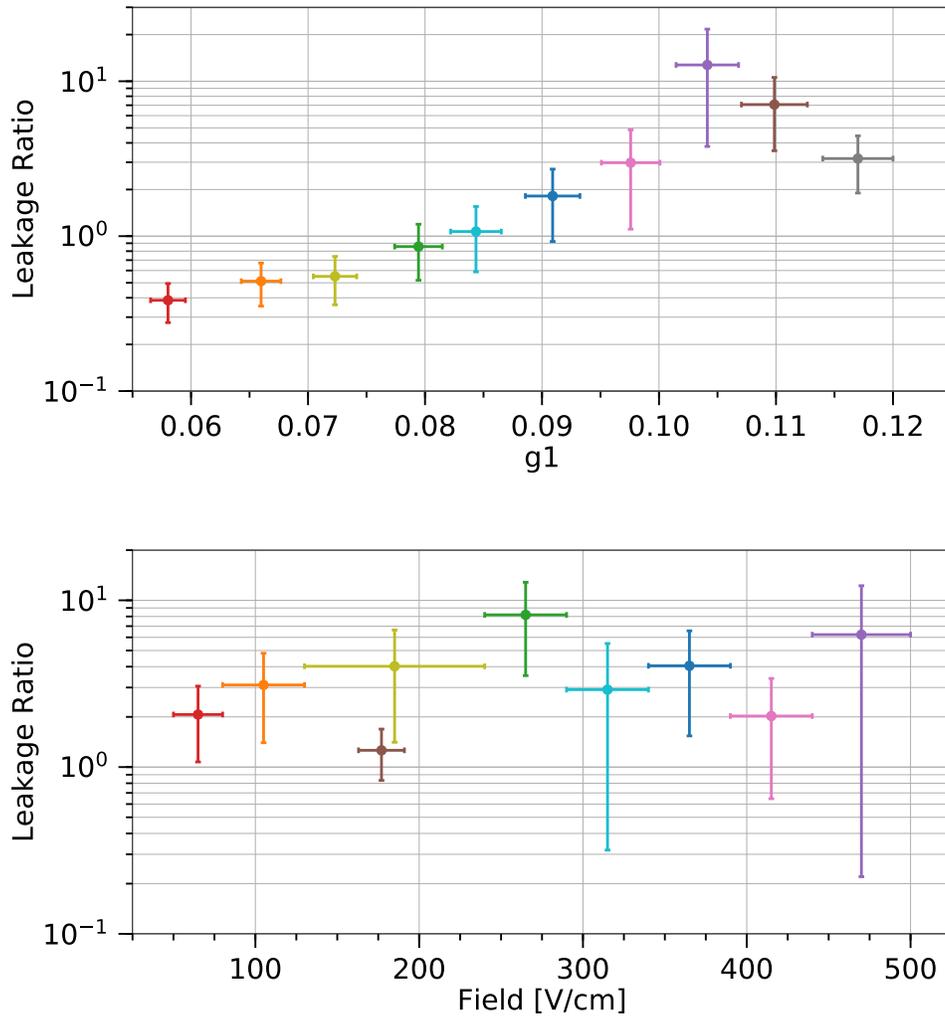


Figure 4.23: The ratio of the real (counted) electronic recoil leakage fraction to the estimated leakage fraction from a skew-Gaussian extrapolation of $\log_{10}(S2c/S1c)$ in the lowest S1c bin. This ratio is shown on the vs. g_1 (top) and vs. field (bottom), using WS2013 and WS2014–16 data, respectively. The color of each data point is degenerate with the g_1 or field value; it is included for consistency with other figures.

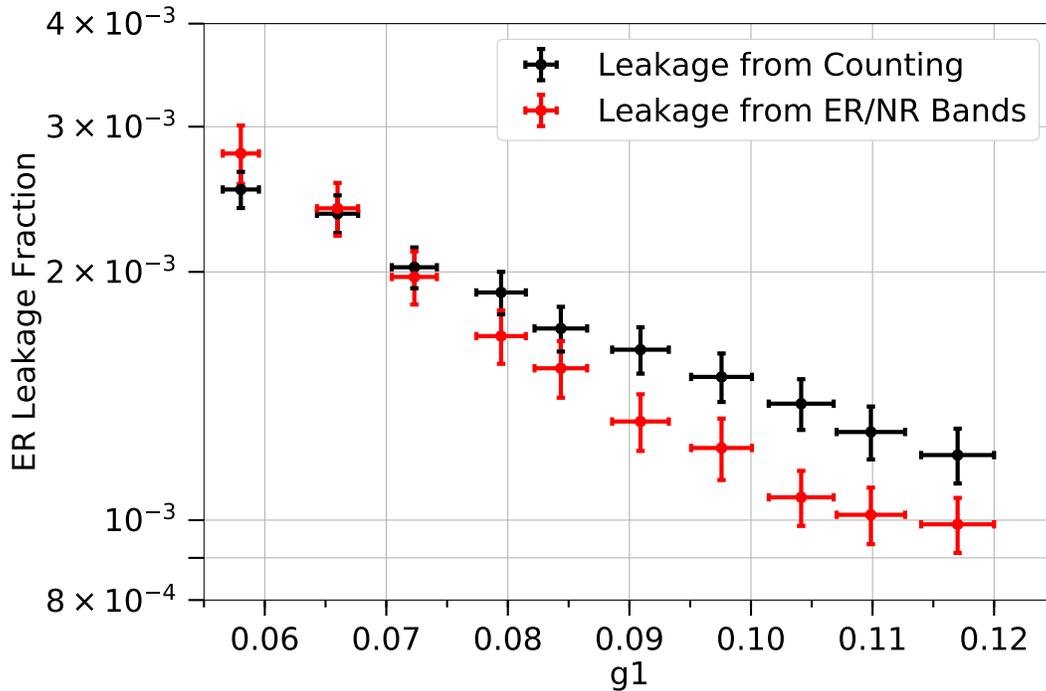


Figure 4.24: The integrated electronic recoil leakage for a flat recoil energy spectrum from 0–9.7 keVee, while varying g_1 in WS2013 data. The max S1c is proportional to 50 photons detected at $g_1 = 0.117$. The leakage is calculated by either counting the number of electronic recoils falling below the NR band (black), or by integrating the electronic recoil skew-Gaussian fits below the NR band (red). The discrepancy between the two methods is explained by a poor fit of the data to a skew-Gaussian distribution in the lowest S1c bin. Statistical errors from Poisson fluctuations are shown.

4.5.1.2 Variation with drift field

Meanwhile, we can also examine the effect of drift field on charge-to-light discrimination, as done in Fig. 4.25. The effect is mostly muted. Drift field does not provide significant variation in the leakage fraction when we look at individual S1 bins. However, we can note some patterns. Across the entire energy range, the lowest field bin of 50–80 V/cm is among the highest leakages for a given S1 bin. Meanwhile, the highest and second-highest fields (390–440 V/cm and 440–500 V/cm, respectively) also often give the highest leakage. Indeed, there seems to be an effect of the leakage reaching a minimum at 240–290 V/cm in several S1 bins.

The WS2013 results are in line with the WS2014–16 results, even though the ER and NR bands separately showed some outlier behavior. A potential explanation for

this latter effect is uncertainties in g_1 , g_2 , and the drift field at the recoil site. The LUX collaboration has previously shown that in order for simulations to correctly mimic data, these quantities need to be slightly adjusted from their measured values [147].

We can also calculate the total leakage up to 80 phd, the maximum pulse area considered in the LZ projected sensitivity [86]. This is done in Fig. 4.26 and shows strong evidence of discrimination being maximized around 300 V/cm. The existence of an optimal drift field in the range accessible to LUX motivated a reduction in the nominal operating field of LZ. The early designs considered a drift field of 600 V/cm [148], while the final design adopts a field of 310 V/cm [86, 119]. (Note, though, that the first scientific run of LZ used a field of 193 V/cm [26].) We compare these results to those from XENON100 [144] at similar g_1 , and we find agreement at the higher fields but a discrepancy at their lowest field of 92 V/cm. However, we emphasize that a direct comparison is impossible, because the two experiments used different S1 thresholds—1 photon detected in LUX and 8 photons detected in XENON100, corresponding to 2 keVnr and 11 keVnr, respectively.

For completeness, we also fit $\log_{10}(S2c/S1)$ in ER data to a standard Gaussian and observe the resulting leakage fraction. This is shown in Fig. 4.27. As mentioned earlier, the leakage is doubled. The improvement ratio is fairly constant, due to the fact that in our energy range, α_B is only weakly dependent on drift field.

4.5.2 Pulse Shape Discrimination

The charge-to-light ratio is undoubtedly the best discriminant in liquid xenon, but under some conditions, its performance can be enhanced with pulse shape information. Xenon excimers are formed in either a singlet or triplet state, and these deexcite on different time scales. The mean lifetime of a singlet excimer is $\tau = 3.27 \pm 0.66$ ns, while that of a triplet excimer is $\tau = 23.97 \pm 0.17$ ns, as measured by the LUX Collaboration [95]. The fraction of excimers produced in each state is found to vary based on the incident particle, with nuclear recoils producing a greater fraction of fast-decaying singlets than electronic recoils. In this paper, we build on the LUX collaboration’s previous analysis of pulse shape discrimination [95]. We explore how our ability to discriminate is dependent on drift field and particle energy.

Figure 4.28 shows an example of how this analysis was conducted. Each event is assigned a *prompt fraction* value, based on the shape of its S1 pulse. The exact calculation is detailed in [95], but in summary: each S1 pulse is decomposed into its detected photon constituents, these detected photons are adjusted based on PMT-specific effects and the location of the recoil, and the fraction of photons within a particular time window is computed. We use a prompt window of -8 to 32 ns and a total S1 window of -14 to 134 ns. For this analysis, we make a key adjustment to the calculation, which is effectively the same g_1 adjustment described in Section 4.4.3. Within each electric field bin, we only consider photons that have hit the PMTs

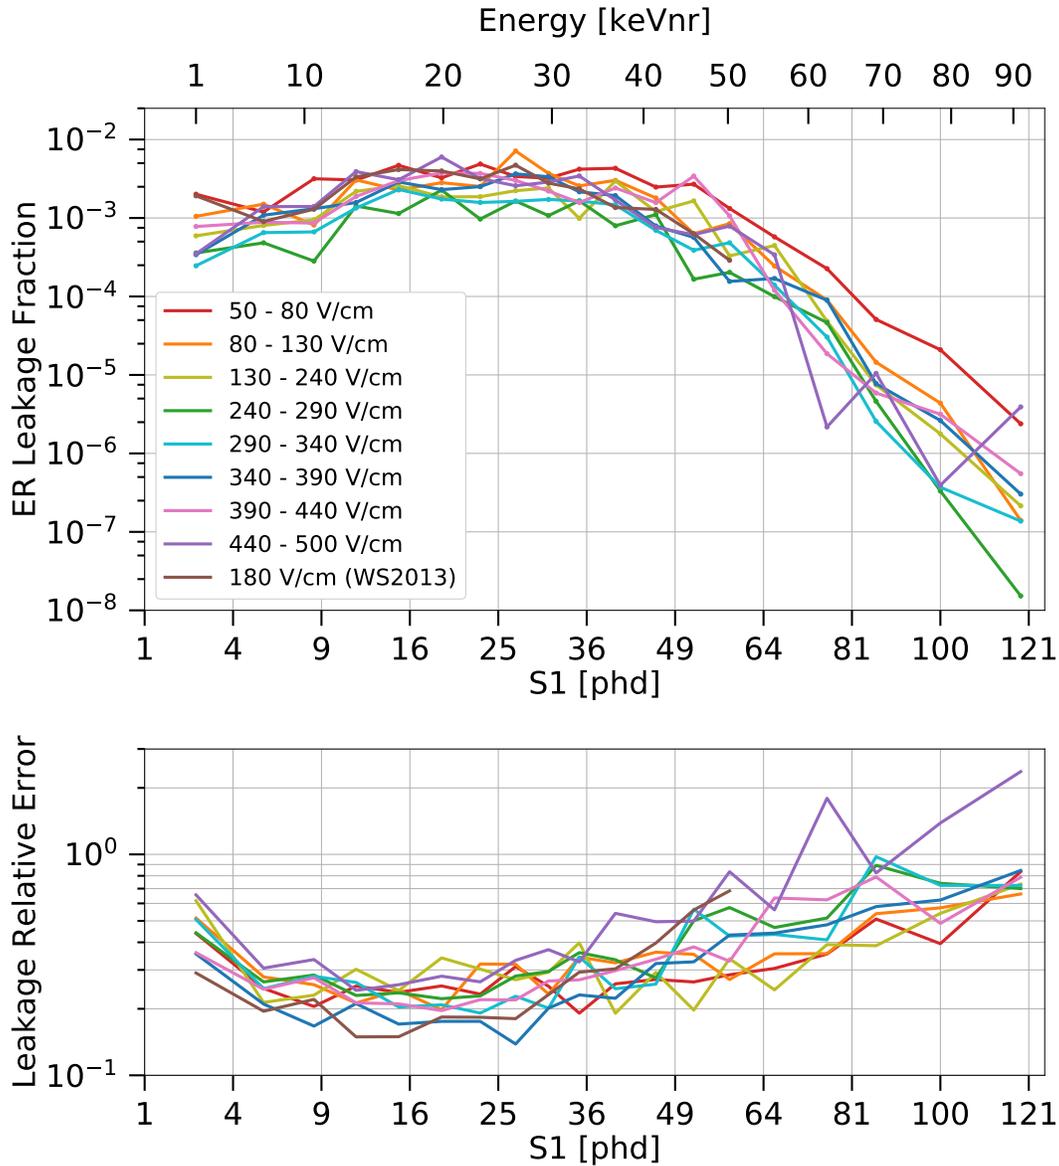


Figure 4.25: (*Top*) The ER leakage fraction for a flat energy spectrum in S1 bins, for various values of drift field, calculated from a skew-Gaussian extrapolation of the ER band below the NR band median. The S1 axis is proportional to $\sqrt{S1}$ for ease of visualization. The equivalent nuclear recoil energy for an S1 is calculated by using the S1 and S2c at the median of the NR band; this varies by field, but not significantly, so we report the energy averaged over the eight field bins. This calculation is consistent with the real counted leakage, except in the lowest S1 bin; see Fig. 4.23 for a comparison of the two calculations in this bin. (*Bottom*) The relative error on these leakage fraction values, defined as in Fig. 4.22.

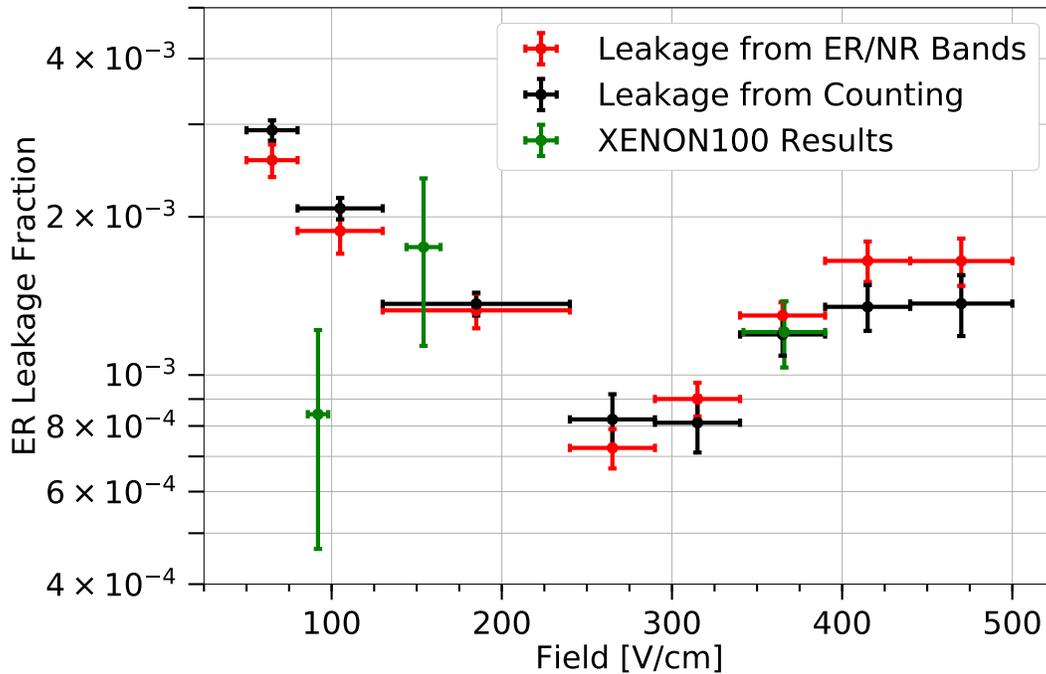


Figure 4.26: The integrated electronic recoil leakage for a flat recoil energy spectrum from 1–80 S1 photons detected (equivalent to 2–65 keVnr), while varying drift field in WS2014–16 data. The leakage is calculated by either counting the number of electronic recoils falling below the NR band (black), or by integrating the electronic recoil skew-Gaussian fits below the NR band (red). Our optimal field over the range examined is ~ 300 V/cm, which is within the expected drift field range of the forthcoming LZ experiment and matches LZ’s design specification of 310 V/cm. However, an exact quantitative prediction of the LZ leakage is impossible because of the higher expected g_1 and g_2 in LZ [86]. Results from XENON100 [144] are shown in green, where we use their leakages at $g_1 = 0.081$ (our results are at $g_1 = 0.087$). The XENON100 leakages correspond to 8–32 photons detected, i.e. 11–34 keVnr.

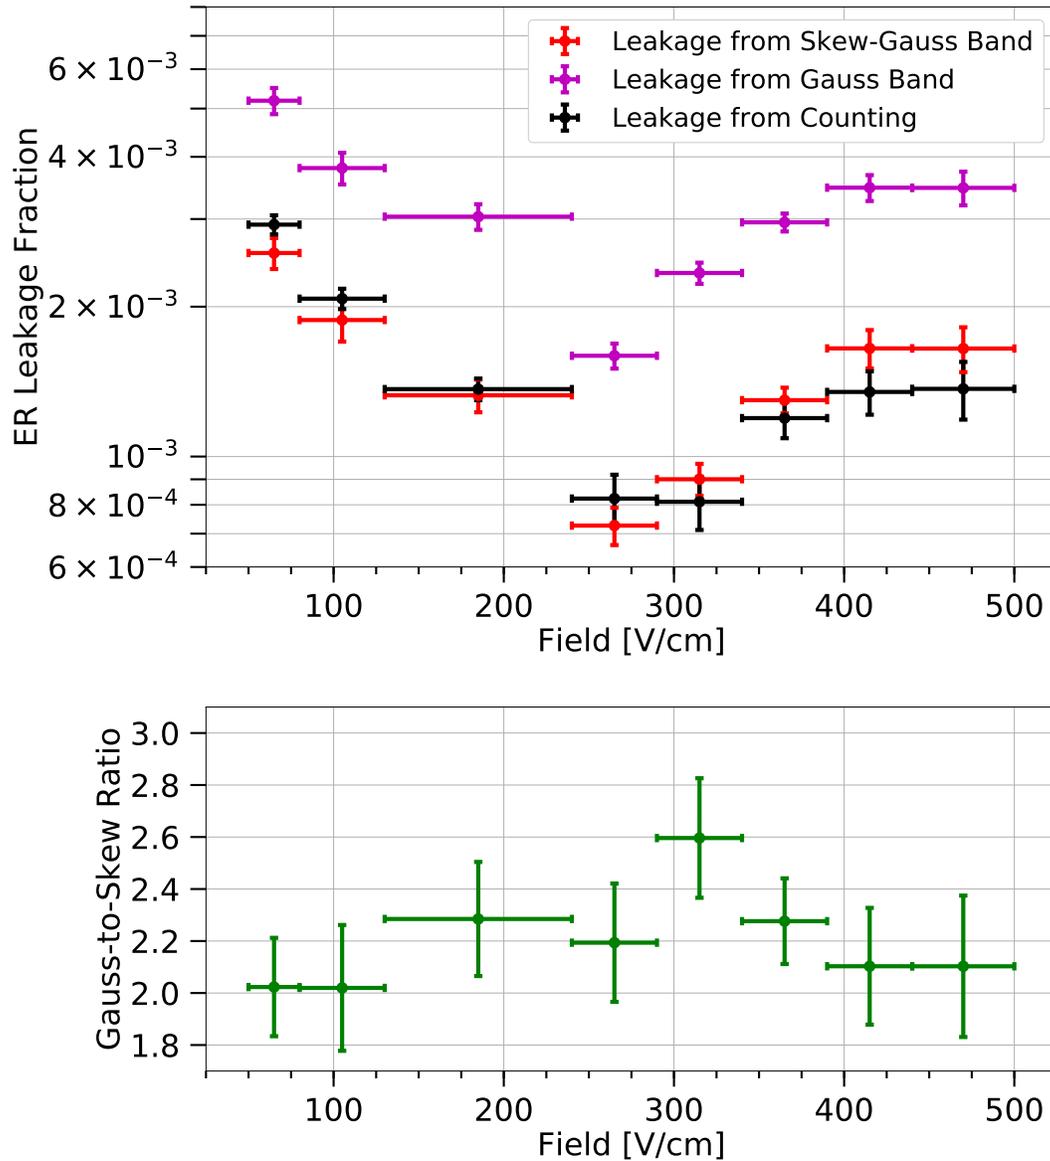


Figure 4.27: A comparison of the electronic recoil leakage fraction using a Gaussian fit to our skew-Gaussian and counting methods shown in Fig. 4.26. The bottom panel shows the ratio of the two fit methods. We see that when we use a Gaussian fit, we over-estimate leakage by a factor of at least 2.

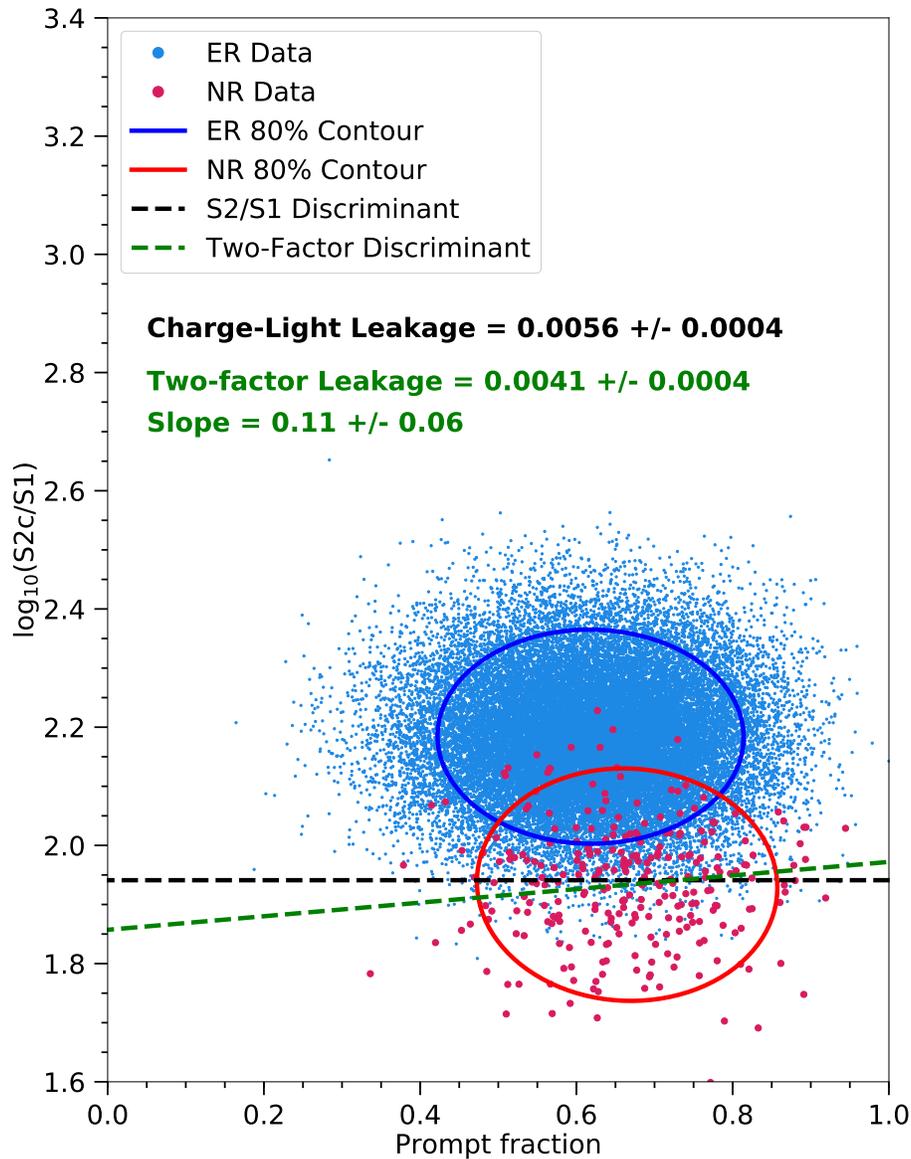


Figure 4.28: An example of how the two-factor leakage is calculated, using data for 80–130 V/cm and 20–30 phd. The electronic recoil and nuclear recoil data are plotted on axes of charge-to-light vs. prompt fraction. Ellipses containing 80% of the data are shown. The black dashed line shows the nuclear recoil median in $\log_{10}(S2c/S1)$ only, and the black text shows the corresponding electronic recoil leakage fraction. The green dashed line shows the optimized discriminating line between the two distributions; the green text shows the resulting electronic recoil leakage, as well as the slope of this line. We note about 27% improvement in the leakage fraction. Further details on this calculation can be found in the text.

used to calculate the ER and NR bands in that bin in order to calculate the prompt fraction. This allows us to adjust for light collection, which we assume accounts for the depth dependence observed in [95]. This fraction is usually between 0.4 and 0.9, but the distribution of prompt fraction for electronic recoils is somewhat lower than the distribution for nuclear recoils (recall Fig. 4.9). As a result, pulse shape serves as a moderately effective discriminant on its own, as also seen by the XMASS experiment [149, 105], the ZEPLIN-I experiment [50], and others [150].

Here, we construct a two-factor discriminant by combining pulse shape with the charge-to-light ratio; this reflects the same strategy as the previous LUX publication and other past analyses [151, 150]. Within each bin of drift field and S1, we consider the prompt fraction and $\log_{10}(S2c/S1)$ in two dimensions. We use maximum likelihood estimation on the ER and NR populations separately to fit the data to a 2D Gaussian distribution. The data are observed to match a 2D Gaussian distribution well except the outermost edges of the electronic recoil data ($<10\%$ of the ER distribution). Then, we choose a line in prompt fraction vs. $\log_{10}(S2c/S1)$ space to discriminate between the two populations. The line is forced to go through the center of the NR 2D Gaussian fit, but the slope is a free parameter; it is determined by minimizing the ER leakage into the NR region. Note one key difference already from [95]: the previous analysis forced this line to pass through the NR median prompt fraction and $\log_{10}(S2c/S1)$, but we find that using the center of the 2D Gaussian gives lower leakage while maintaining 50% NR acceptance. However, for the lowest S1 bin (0–10 phd), the 2D Gaussian fit is poor, because there is an abundance of events with prompt fraction of exactly 0 or 1. If an S1 pulse has only a few photons, there is a significant probability that its prompt fraction is 0 or 1. This fit is so poor that the resulting two-factor leakage ends up being greater than the charge-to-light leakage. As a result, for this bin only, we continue to use the median in both dimensions.

The second addition we make is to use the bootstrap method to determine the slope of the discriminating line and its uncertainty. First, a random selection of N electronic recoil events is chosen with replacement, where N is the total number of electronic recoil events in this field/S1 bin. This means that it is almost certain that some events will be in the bootstrap sample twice or more often. Then, we calculate the optimal slope on this sample, using the procedure described in the previous paragraph. We do this 100 times to get a distribution of slopes (the number of iterations has been chosen to be high enough such that the resulting distribution of slopes is negligibly affected by the pseudo-random number generation). The slope that we use for the final discriminating line of this field/S1 bin is the mean of this distribution, while the error on that slope is given by the standard deviation of this distribution. Finally, we calculate the two-factor leakage by counting the number of (weighted) electronic recoil events falling below the discriminating line. This procedure allows us to obtain an uncertainty on the slope of the discriminating line, and it serves as a safeguard, preventing the calculation from being too dependent

on a single leaked electronic recoil.

The statistical error on the two-factor leakage has two components: the Poisson error on the number of leaked events and the error on the slope of the discriminating line. The total statistical error is not found by adding these in quadrature because they are not independent; the Poisson error is a function of the leakage value, so it is dependent on the discriminating line error. We perform this analysis as follows. Given an S1 and field bin, we calculate the distribution of slopes as described in the previous paragraph. We then draw 100 random slopes, assuming that this distribution is Gaussian with the appropriate mean and standard deviation.² For each slope, we calculate the two-factor leakage and its Poisson error. Then, we randomly choose a leakage from a Gaussian distribution with the two-factor leakage as its mean and the Poisson error as its width. Finally, we take the mean and standard deviation of this 100-sample dataset as the average leakage and its error.

The results are shown in Fig. 4.29, where we plot the ratio of the two-factor leakage to the charge-to-light leakage. A marked improvement in discrimination is observed below 50 phd for the lowest electric fields (50–80 and 80–130 V/cm). The 130–240 V/cm field bin is ambiguous: the WS2014–16 data show improvement for energies between 30–60 phd, but the WS2013 data at 180 V/cm show no improvement over charge-to-light discrimination. For higher electric fields, there does not seem to be a significant reduction in leakage when using the two-factor discriminant. The most likely explanation for this is that higher electric fields are associated with less recombination. Thus, fewer scintillation photons leave the recoil site, and the S1 pulse shape is dominated by the longer triplet decay time for both nuclear and electronic recoils [152]. We also do not observe improvement at higher energies, but this could be due to low statistics; there are plenty of ^{14}C events in the dataset, but the charge-to-light leakage is so robust that virtually none of them falls below the NR band. Although the leakage values appear to be different than the ones reported in [95], this is due to the varying methodology and drift field range. We have confirmed that if we modify our procedure to be identical to the one detailed there, our results are consistent.

We also consider the two-factor leakage across the entire 1–80 phd energy range. Figure 4.30 shows these results, as well as a comparison to the charge-to-light only counted leakage. We see that although there is improvement in discrimination for low fields, the optimal drift field bins are still 240–290 V/cm and 290–340 V/cm. The reader can also find the two-factor leakage in S1 bins in Fig. 4.31, although we emphasize that this is an estimate. The charge-to-light leakage in S1 bins is calculated with a skew-Gaussian extrapolation, whereas the leakage ratio is calculated

²The Gaussian assumption is accurate for the majority of S1/field bins, although there are a few bins where the distribution has a sharp preference for a slope separate from the main peak. In these, a handful of events bias the minimization toward this value, and the use of a Gaussian distribution smooths out this effect.

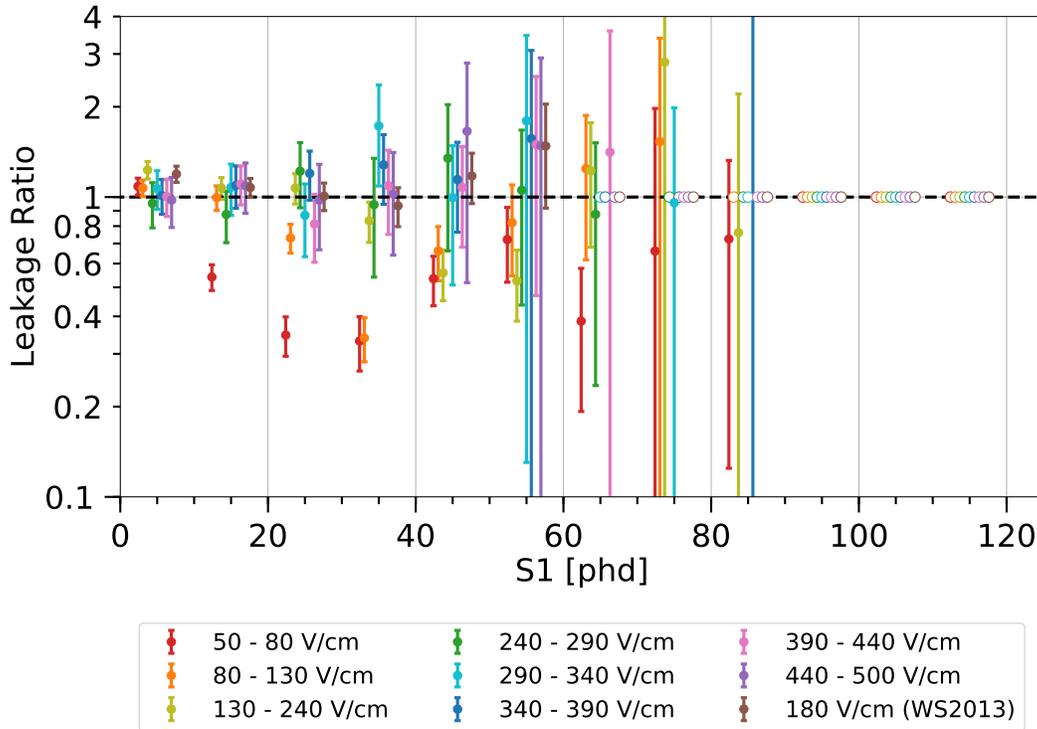


Figure 4.29: The ratio of two-factor leakage to charge-to-light leakage, for various S1 and drift field bins. Error bars are statistical; see text for details. Open circles represent bins for which charge-to-light discrimination alone gives zero electronic recoils falling below the NR band; as a result, it is impossible to calculate the improvement from two-factor discrimination. Leakage ratios with large error bars are made transparent and plotted as dashed lines to draw the eye toward more precise measurements. The plotted S1 values are slightly shifted relative to their true value (by up to 2 phd) for ease of visualization. The true S1 coordinates are 5 phd, 15 phd, 25 phd, etc.

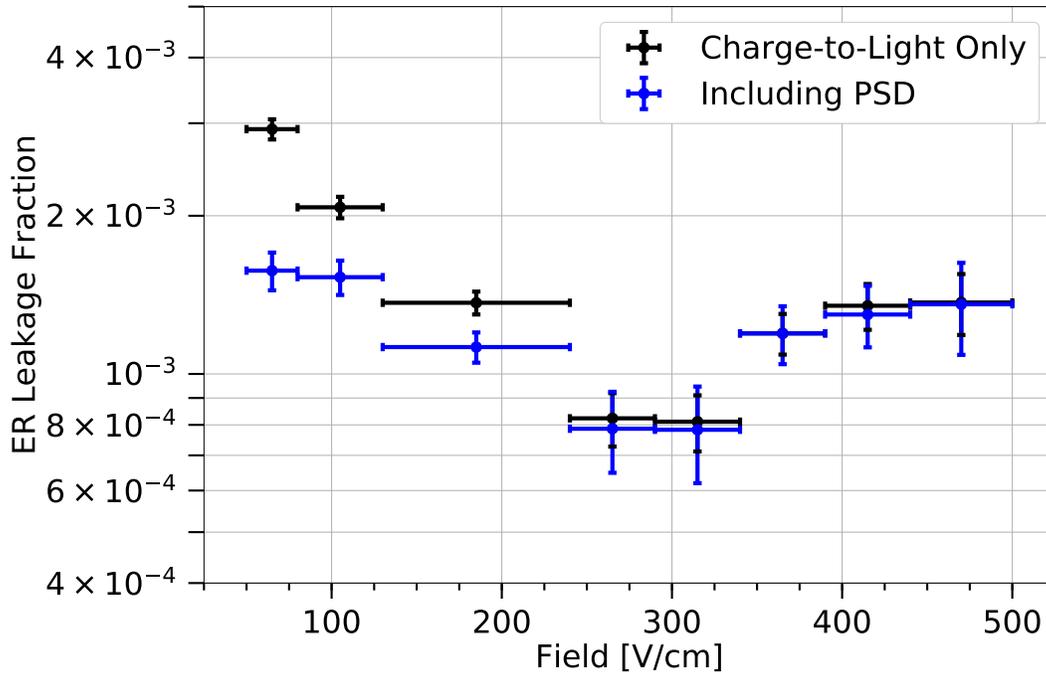


Figure 4.30: The integrated electronic recoil leakage for a flat recoil energy spectrum from 1–80 S1 photons detected (equivalent to 2–65 keVnr), while varying drift field in WS2014–16 data. The leakage is calculated using only the charge-to-light ratio, i.e. $\log_{10}(S2c/S1)$, and using both charge-to-light and pulse-shape discrimination in tandem. Both leakage values are based on the “counting” method described in Fig. 4.26, where we count the number of electronic recoils leaking into the nuclear recoil 50% acceptance region.

by counting electronic recoils in the nuclear recoil acceptance region; thus, it is not exactly consistent to combine the two.

Finally, Fig. 4.32 shows how the slope of the discriminating line varies with electric field and S1. The most striking effect is that the slope is almost always positive, meaning that the ER population is tilted toward higher $\log_{10}(S2c/S1)$ at higher prompt fraction. In addition, there appears to be a weak increase in the slope with energy and no dependence on field.

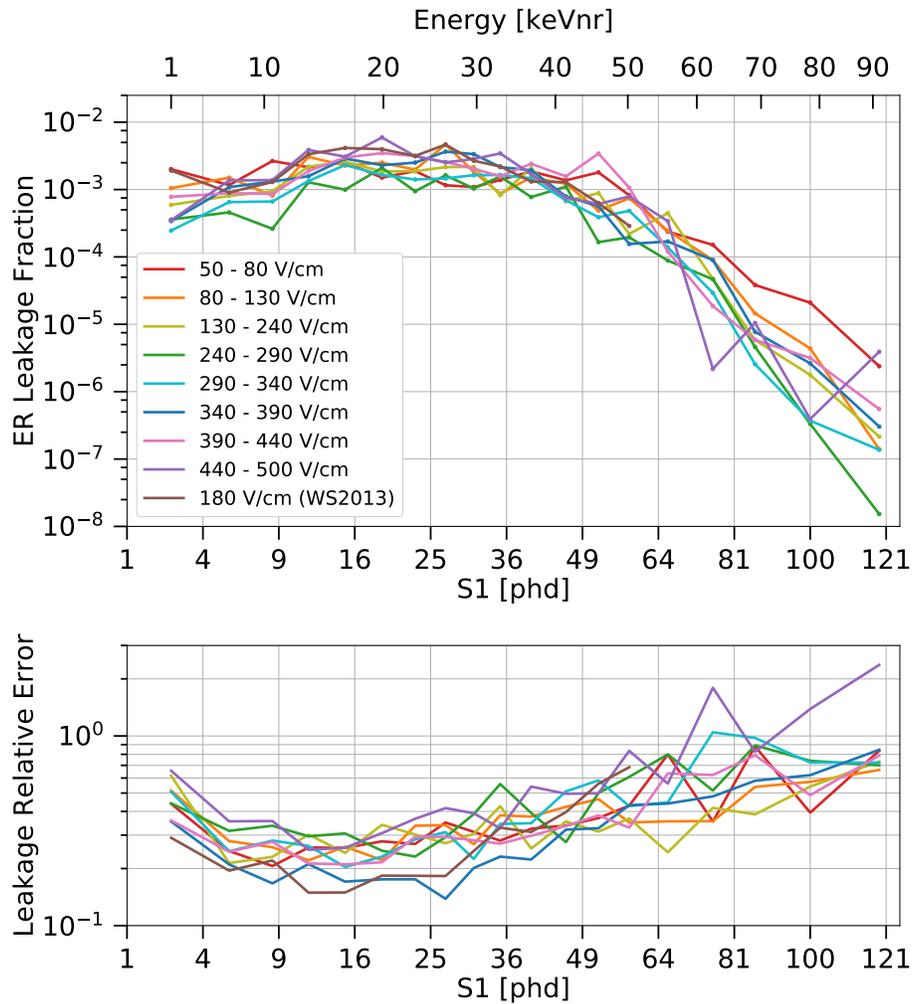


Figure 4.31: (*Top*) An estimate of the electronic recoil two-factor leakage fraction for a flat energy spectrum in S1 bins, for various values of drift field. The leakage fraction is estimated by multiplying the charge-to-light leakage by the ratio of two-factor to charge-to-light leakage (i.e. the results in Figs. 4.25 and 4.29, respectively). This is an estimate for two reasons: first, the two calculations use different S1 bins; second, the charge-to-light leakage is a skew-Gaussian extrapolation, while the ratio is based on counting individual events, so they are not perfectly consistent. The S1 axis is proportional to $\sqrt{S1}$ for ease of visualization. The equivalent nuclear recoil energy for an S1 is calculated by using the S1 and S2c at the median of the NR band; this varies by field, but not significantly, so we report the energy averaged over the eight field bins. (*Bottom*) The relative error on these leakage fraction values, defined as in Fig. 4.22.

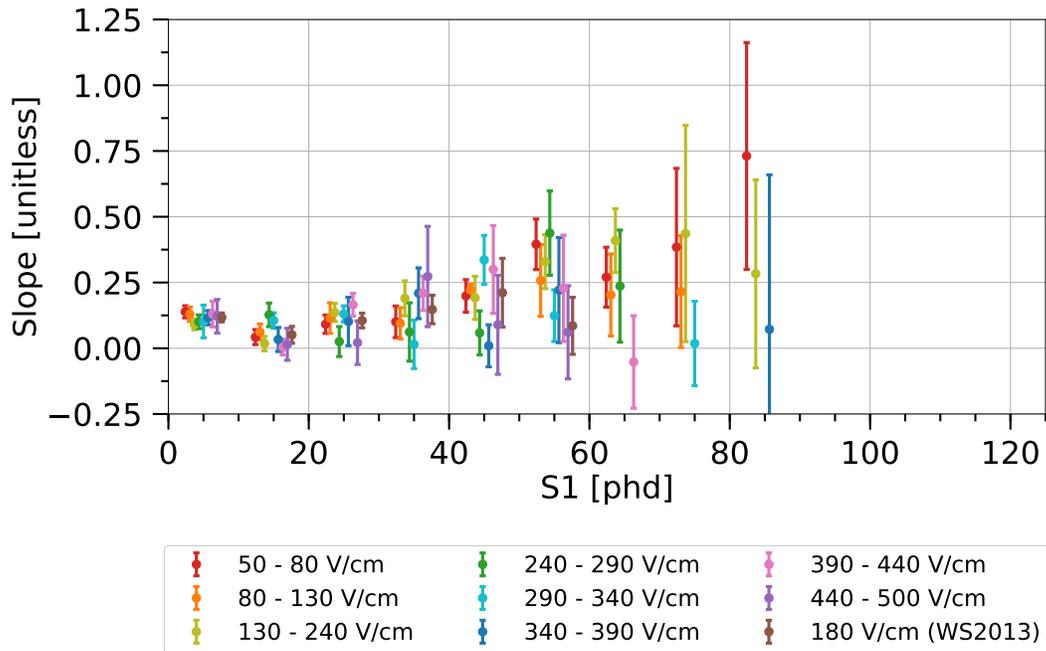


Figure 4.32: The slope of the two-factor discrimination line in $\log_{10}(S2c/S1)$ vs. prompt fraction space, for each S1 and field bin. Missing points represent bins for which charge-to-light discrimination alone gives zero electronic recoils falling below the NR band. The plotted S1 values are slightly shifted relative to their true value (by up to 2 phd) for ease of visualization. The true S1 coordinates are 5 phd, 15 phd, 25 phd, etc.

Chapter 5

Simulations of Liquid Xenon Detectors

We continue our study of energy depositions in liquid xenon by simulating these interactions in the LUX detector. First, we develop a model of band skewness that describes it as a consequence of recombination fluctuations. This is the first-ever model to describe skewness at the microphysical level. We implement this model primarily for electronic recoils, but we also briefly discuss skewness in nuclear recoils. Second, we describe the physical parameters that govern the width of the ER band, one of the primary factors for ER/NR discrimination power.

5.1 A Model of ER and NR Band Skewness

Independent of the LUX analysis described in Chapter 4, skewness of the ER band has been observed by others [57, 153], but no physical motivation for it has emerged.¹ Here, we present one potential explanation by utilizing the Noble Element Simulation Technique, or NEST. We use NEST v2.0.1 [154]. See Sec. 3.5 for a description of this simulation package.

A LUX-specific NEST model, which we will refer to as LUX-NESTv2, has been described in [147]. It has had great success in reproducing the median and width of the ER and NR bands in WS2014–16 data. The only deficiency has been that it fails to correctly reproduce the skewness of the ER and NR bands. Here, we present a model of skewness that can be inserted into NEST and correctly reproduce the data.

5.1.1 ER Skewness

The skewness of the ER band is critical to discrimination and thus to sensitivity in general, so it is equally critical that LUX-NESTv2 models it correctly. In the original version of LUX-NESTv2, if a user simulates the LUX WS2014–16 calibrations of ^3H and ^{14}C , they would arrive at an ER band with (small) negative skewness in the WIMP search region. However, the data clearly show that the ER band has positive skewness in this energy range.

In order to rectify this inconsistency, our solution is to add skewness into LUX-NESTv2 at the level of recombination fluctuations. In LUX-NESTv2, after calculating the quanta produced N_{ions} and N_{excitons} , the code calculates the mean recombination probability r and its variance σ_r^2 ; all of these quantities are deterministic and only based on the particle type, energy, and electric field. It then simulates the number of electrons and photons leaving the recoil site using Eq. 5.1 and Eq. 5.2, respectively.

¹Reference [153] does not directly report skewness. However, they observe that their signal-like mismodeling parameter is fit to a negative value by data. This means that within S1c bins, the S2c distribution is shifted to higher values, an identical effect qualitatively to our observation of positive ER band skewness.

$$N_{\text{electrons}} = G[(1-r)N_{\text{ions}}, \sigma_r^2], \quad (5.1)$$

where $G[\mu, \sigma^2]$ is a randomly generated number from a Gaussian distribution with mean μ and variance σ^2 .

$$N_{\text{photons}} = N_{\text{excitons}} + N_{\text{ions}} - N_{\text{electrons}}. \quad (5.2)$$

However, we update this step such that the number of electrons is drawn from a skew-Gaussian distribution, shown in Eq. 5.3.

$$N_{\text{electrons}} = F[(1-r)N_{\text{ions}} - \xi_c, \frac{1}{\omega_c} \sqrt{\sigma_r^2}, \alpha_R],$$

$$\text{where } \omega_c = \sqrt{1 - \frac{2}{\pi} \frac{\alpha_R^2}{1 + \alpha_R^2}}, \quad (5.3)$$

$$\text{and } \xi_c = \sqrt{\sigma_r^2 \frac{1 - \omega_c^2}{\omega_c^2}}.$$

Here, $F[\xi, \omega, \alpha]$ is a randomly generated number from a skew-Gaussian distribution given by the PDF in Eq. 4.3. This scheme preserves the mean and variance of Eq. 5.1. The number of photons leaving the recoil site is still given by Eq. 5.2. For clarity, we emphasize that there are two skewness parameters that will be frequently referenced: α_R is the skewness parameter in the recombination fluctuations model in Eq. 5.3, while α_B is the skewness parameter of the ER or NR band in $\log_{10}(\text{S2c/S1c})$ space, as described in Section 4.4.1.

If α_R is sufficiently positive, the results of a LUX-NESTv2 simulation will give $\alpha_B > 0$. However, the skewness of the ER band can only be reproduced if α_R varies with energy and field. The model in Eq. 5.4, where E is the total energy deposited by the electronic recoil and F is the drift field at the recoil site, correctly reproduces data with a certain set of parameter values. This model is empirical. We developed it by determining the α_R that reproduces the correct α_B in bins of drift field and S1c. We observed that the α_R required to match the measured α_B behaves differently in the low-energy and high-energy regimes, i.e. above and below E_2 . As a result, we construct a separate model for each energy regime, capturing the energy- and field-dependence of α_R in that regime. The final model is a weighted sum of the two models, in which the weight is an energy-dependent sigmoid function that asymptotically goes to zero and one in the appropriate limits. The transition between the models is field-independent and found to be about 25 keV, which is comparable to the energy at which ER recombination transitions from the Doke-Birks model to the Thomas-Imel Box model (Fig. 3.7).

$$\alpha_R = \frac{1}{1 + e^{(E-E_2)/E_3}} \left[\alpha_0 + c_0 e^{-F/F_0} (1 - e^{-E/E_0}) \right] + \frac{1}{1 + e^{-(E-E_2)/E_3}} \left[c_1 e^{-E/E_1} e^{-\sqrt{F/F_1}} \right]. \quad (5.4)$$

The nine parameters in Eq. 5.4 are not obtained by a rigorous optimization, due to the immense computational power that would be required for a nine-dimensional fit. Instead, we proceed as follows. For each parameter X , we find a value that approximately matches the data. Using this value, we simulate the ^{14}C and ^3H WS2014–16 calibrations, and we calculate the ER bands for six field bins equally spaced between 50 and 500 V/cm. In doing so, we neglect the energy weighting and g_1 adjustments described in Section 4.4.1. Next, we compute the degree to which the simulated ER band skewness is consistent with data by using Eq. 5.5, in which j and k iterate over field and S1c bins, respectively, and δ represents the uncertainty on α_B from the skew-Gaussian fit. By adjusting X slightly and repeating this procedure several times, we obtain a set of points (X_p, χ_p^2) . Finally, we fit a quadratic function to these points. Defining $(\bar{X}, \bar{\chi}^2)$ as the vertex of this parabola, we derive our desired quantities: the estimated value of X is \bar{X} , and the uncertainty on X is the amount δ_X such that $X = \bar{X} \pm \delta_X$ implies $\chi^2 = \bar{\chi}^2 \pm 1$.

$$\chi^2 = \sum_{i \in \{^{14}\text{C}, ^3\text{H}\}} \sum_j \sum_k \left[\frac{(\alpha_{B, \text{Data}} - \alpha_{B, \text{MC}})^2}{\delta_{\text{Data}}^2 + \delta_{\text{MC}}^2} \right]_{i,j,k}. \quad (5.5)$$

We show four examples of this fitting procedure in Fig. 5.1. We intentionally show two high-quality fits and two low-quality fits—the top and bottom panels, respectively. E_2 and E_3 in the bottom panels behave as transition variables, so they are difficult to fit precisely. All nine parameter values determined by this procedure are listed in Table 5.1.

Figure 5.2 shows a plot of Eq. 5.4 for a variety of energies and fields, and Fig. 5.3 shows a comparison of α_B between data and simulation. One observes that the two match well, and that α_B dips below zero at high enough energy. Here, the uncertainty on the skewness is obtained from the fit of the $\log_{10}(\text{S2c/S1c})$ histogram.

We also observe that our skewness model is successful at matching data from other experiments. See Fig. 5.4 for a comparison to ZEPLIN-III data, which reported an average leakage of 1.3×10^{-4} at a 3.8 kV/cm drift field [57, 155, 156]. Furthermore, the authors of [157] used our ER skewness model to accurately simulate ^{37}Ar calibration data in XENON1T.

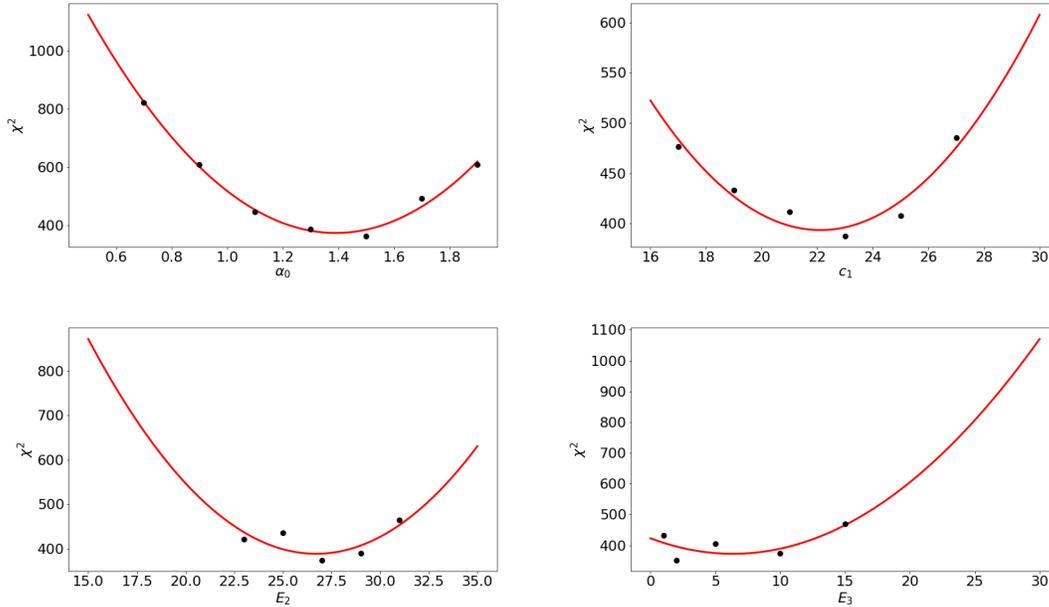


Figure 5.1: Demonstrations of how we find the best-fit values for the α_0 , c_0 , E_2 , and E_3 parameters in Eq. 5.4. For each parameter, we choose several values to test while holding the others constant. We calculate χ^2 using Eq. 5.5 and plot the corresponding black data points. We then fit the points to a parabola, shown by the red line. The best-fit parameter \bar{X} and its uncertainty δ_X are read from the fit.

Table 5.1: The optimal values for the parameters of the electronic recoil skewness model (i.e. Eq. 5.4), based on LUX WS2014–16 ^3H and ^{14}C calibration data.

Parameter	Value \pm Uncertainty	Units
α_0	1.39 ± 0.03	...
c_0	4.0 ± 0.2	...
c_1	22.1 ± 0.5	...
E_0	7.7 ± 0.4	keV
E_1	54 ± 2	keV
E_2	26.7 ± 0.5	keV
E_3	6.4 ± 0.9	keV
F_0	225 ± 12	V/cm
F_1	71 ± 4	V/cm

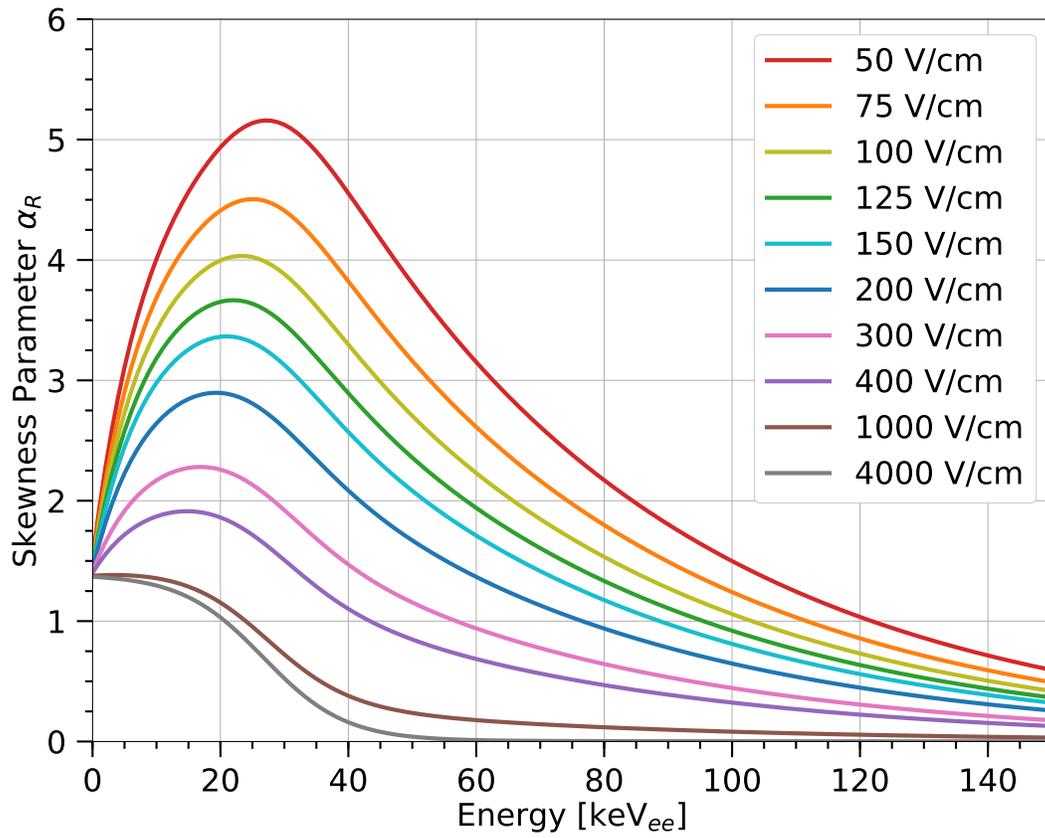


Figure 5.2: The skewness model for recombination fluctuations in Eq. 5.4.

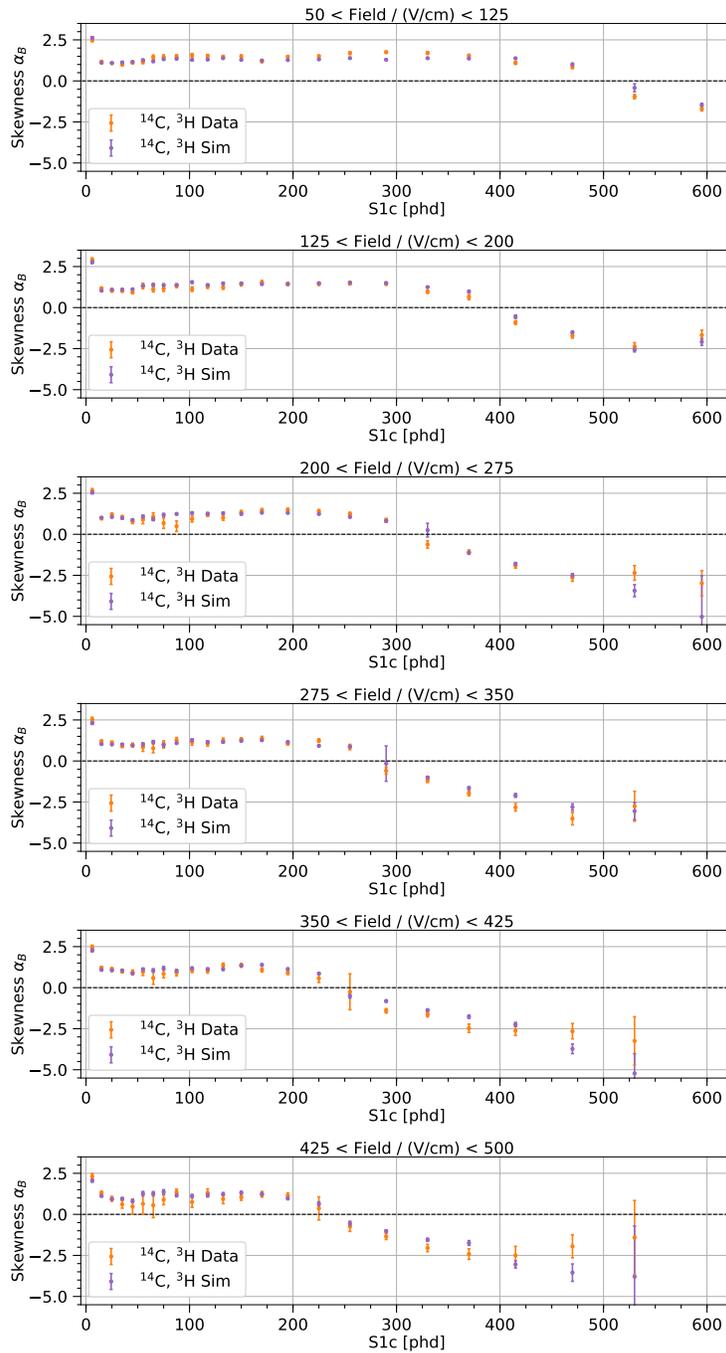


Figure 5.3: A comparison of the skewness of the ER band in WS2014–16 data vs. simulation from LUX-NESTv2, based on our model in Eq. 5.4. Points below 50 phd are from ^3H data and simulation, and points above 50 phd are from ^{14}C data and simulation.

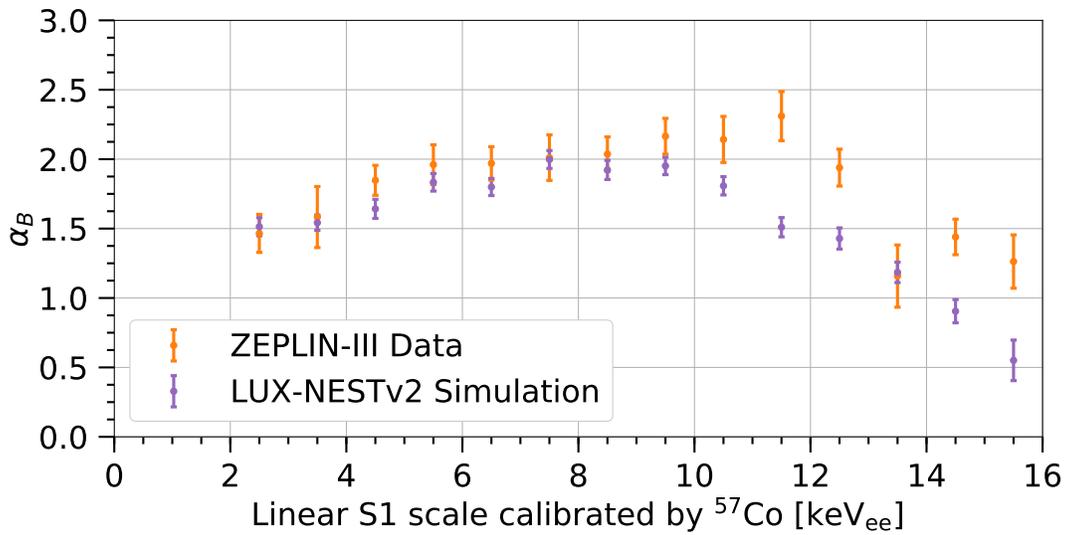


Figure 5.4: A comparison of the skewness of the ER band in ZEPLIN-III data [57, 156] vs. our simulation of the ZEPLIN-III data using LUX-NESTv2 and our model in Eq. 5.4. We use a constant light yield of 1.881 phd/keVee to convert the simulated S1 signal to the energy scale used by ZEPLIN-III. The data points match our simulation within 1 standard deviation below 10 keVee. We do not reproduce the highest-energy data with complete accuracy, but we do observe the same qualitative decrease in skewness between 10 and 16 keVee.

5.1.2 NR Skewness

The NR band exhibits skewness, but it is substantially more difficult to model. There are a few reasons for the difficulty: first, skewness is a third-order effect (as mentioned previously, it is associated with the third standardized moment of the distribution), so correctly measuring it requires a substantial amount of data. This is possible for electronic recoils because in WS2014–16, there are over 1.5 million events. On the other hand, there are only about 80,000 nuclear recoils in the data set, so this dataset is prone to large uncertainties and statistical fluctuations. Second, there is a small number of multiple scatters in the nuclear recoil dataset, because occasionally multiple S2 pulses are so close together that they are classified as a single S2 pulse. We cut these out without significantly reducing the single-scatter acceptance, but a small number do persist, and they have a disproportionately high S2 area. This means that although they have a negligible effect on the NR band median and width, they have a considerable effect on the skewness. Including these multiple scatters, which are prevalent at high energy and high electric field, causes the skew-Gaussian fit to be fit at α_B of 3.0 or above.

To account for this, we remove events at high S2 before binning $\log_{10}(S2c/S1)$ and doing the skew-Gaussian fit, resulting in the data points of Fig. 5.5. The NR band skewness does not affect leakage if it is defined through a cut-and-count procedure, i.e. the fraction of electronic recoils falling below the NR band median. However, most experiments use a profile likelihood ratio or a similar hypothesis test, in which case a positive NR skewness would worsen an experiment’s sensitivity.

The skewness in NR data is still relatively high, even with this change. We simulate recombination fluctuations with Eq. 5.3, but we require $\alpha_R \rightarrow \infty$. To clarify, the skew-Gaussian PDF (Eq. 4.3) is such that as α increases, the PDF tends to “saturate.” This means that for $\alpha > 10$, the PDF does not substantially change; it effectively becomes a unit step function multiplied by a Gaussian. We use $\alpha_R = 20$ in LUX-NESTv2 to simulate nuclear recoils, and the results are shown in Fig. 5.5. The match is moderate; we observe no substantial field or energy dependence.

5.2 Fluctuations of the ER Band

The width of the ER band is crucial to understanding particle discrimination; as the width increases, more electronic recoil events leak below the NR band, and detector sensitivity to dark matter deteriorates. It is therefore an integral part of our analysis to examine the effects of different types of fluctuations on the band width, and especially to see their dependence on drift field and energy.

LUX-NESTv2 calculates an S1 and S2 signal for each energy deposit, but there are random fluctuations about some mean for these values. We split all these fluctuations into four categories: 1) S1-based fluctuations, including photon detection

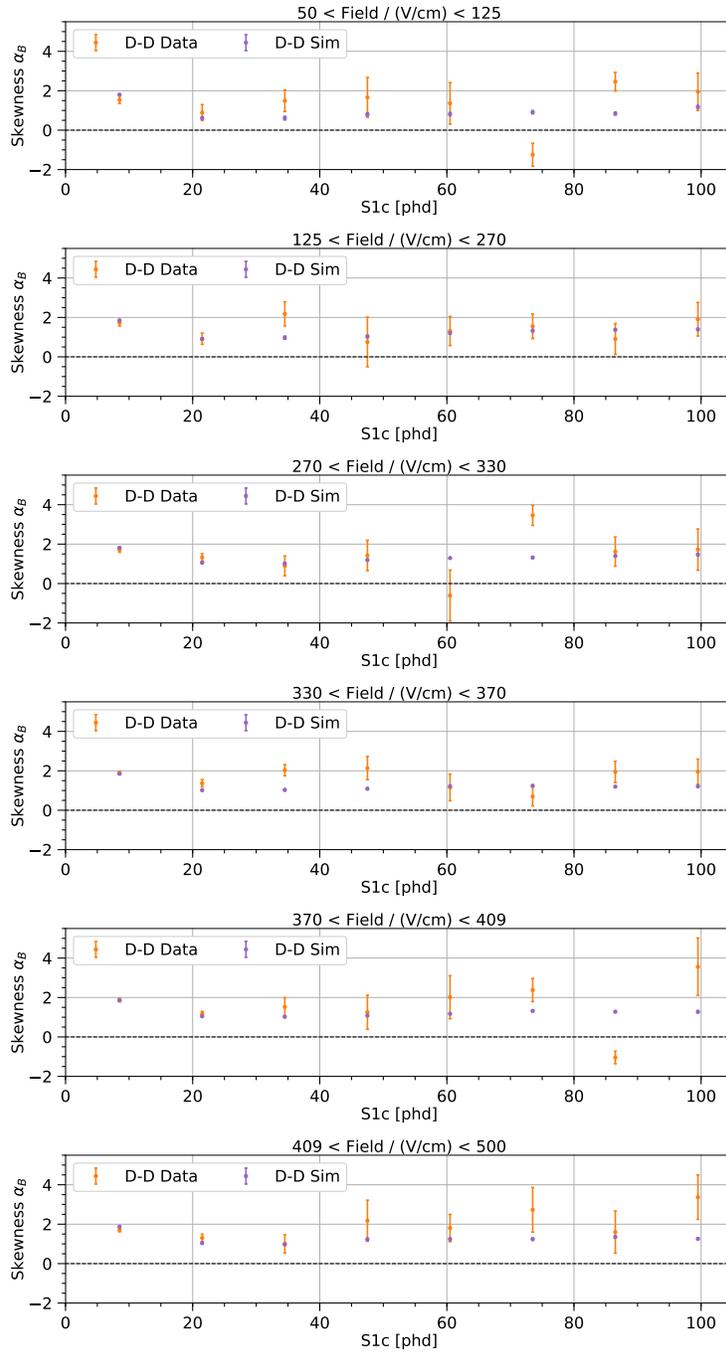


Figure 5.5: A comparison of the skewness of the NR band in WS2014–16 data vs. simulation from LUX-NESTv2, using $\alpha_R = 20$.

efficiency, the double-photoelectron effect, pulse area smearing, PMT coincidence, and position dependence; 2) S2-based fluctuations, including electron extraction efficiency, photon detection efficiency in gas, the double-photoelectron effect, pulse area smearing, and position dependence; 3) recombination fluctuations; and 4) fluctuations in the number of quanta (i.e. excitons and ions) produced for a given energy deposit. For each category, we turn off all other fluctuations in LUX-NESTv2, and we simulate 10 million electronic recoils using a flat energy spectrum, LUX detector-specific parameters, a uniform value of $g_1 = 0.10$, and a uniform drift field. We then calculate the ER band as described in Section 4.4.1, including the skewness model described in Section 5.1.1. We repeat this procedure for electric fields of 180, 500, 1000, and 2000 V/cm. Then, we look specifically at σ_-^2 , the band variance due only to the downward fluctuations. The variance is examined rather than the width because if the fluctuations are independent, adding the variances will give the total variance. The results are shown in Fig. 5.6.

We observe that the fluctuations in the number of quanta are an insignificant portion of the full ER band variance (a few percent at most), but they do grow with field. The S2-based fluctuations contribute to about 5–10% of the full band variance; they are suppressed by both energy and field. The field-dependent suppression of S2-based fluctuations is explained by the fact that a higher electric field is associated with less recombination, so the S2 signal is larger for a given S1 signal. Similarly, an increased energy leads to an increased charge yield (for energies in the Doke-Birks regime) and a suppression of S2-based fluctuations. The S1-based fluctuations are significant at all energies and fields, accounting for 20–30% of the total variance. Their field dependence is weak, but they do get stronger with field, for the same reason that S2-based fluctuations are suppressed by an increased field. Finally, the recombination fluctuations are clearly the strongest contributor to band width, consistent with the findings of [76]. Their field and energy dependence is not easy to summarize quickly, though. At low energies, the recombination fluctuations unambiguously grow with field in this field range. At higher energies, recombination fluctuations begin to shrink with energy in a way that is field-dependent; as a result, the ordering of the fields is not monotonic. For example, looking at just the 2000 V/cm points, recombination fluctuations begin to decrease above ~ 70 phd and continue their downward trend at higher energies. The 2000 V/cm recombination fluctuations are larger than the recombination fluctuations for any other drift field below 70 phd, but they become the smallest at the highest values of S1. One particularly interesting feature is that at very high energies and fields—specifically, the 2000 V/cm simulation above 250 phd, or 110 keVee—the recombination fluctuations become smaller than the S1 fluctuations, which are dominantly from g_1 binomial statistics.

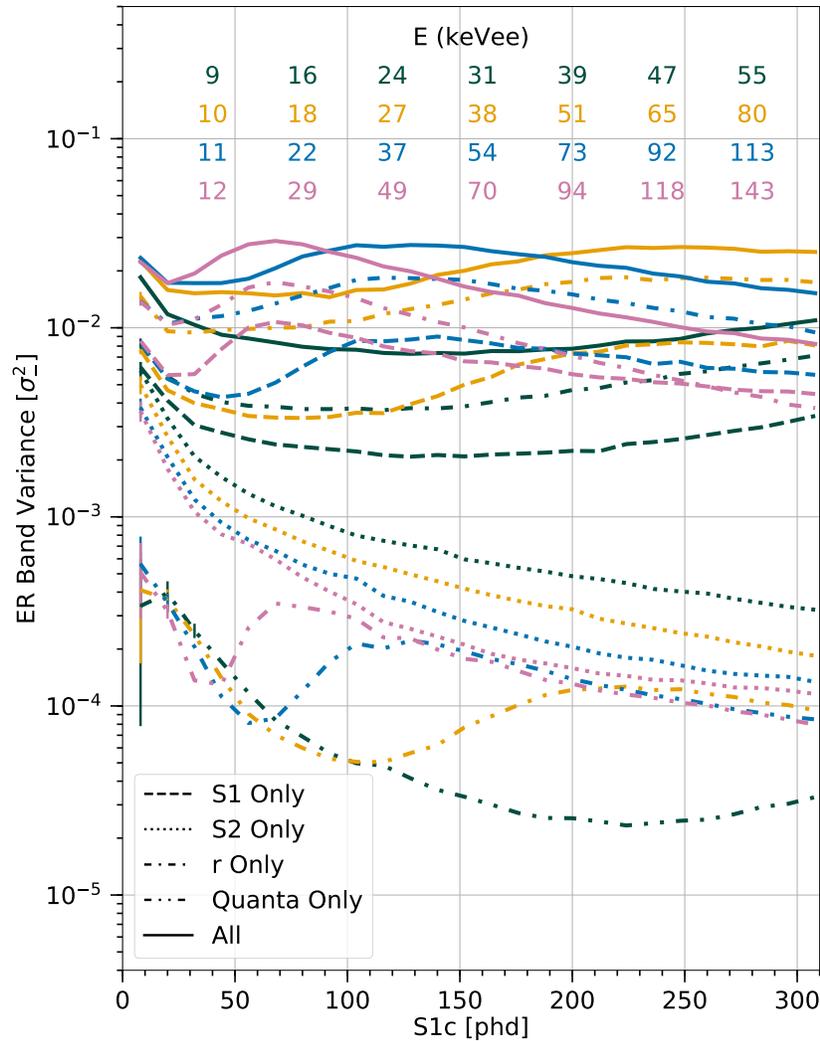


Figure 5.6: The ER band variance σ^2 for different types of fluctuations and fields, based on simulation from LUX-NESTv2. Color represents drift field: green, gold, blue, and magenta represent 180, 500, 1000, and 2000 V/cm, respectively. Line style represents the fluctuations category that is turned on in the simulation; for solid lines, all fluctuations are turned on (the default status for NEST). All points have error bars, but most are too small to be visible, except for the lowest S1c bin. If the fluctuations were uncorrelated, the solid lines would equal the sum of all the other lines for each field, but there are small correlations, so this is not quite true. For a given field, the difference between the solid line and the sum of the other lines is at most 12% except in the lowest S1c bin, where the total variance can be as much as double the sum of the individual component variances. At the top, we show the equivalent energy for S1c in multiples of 40 phd, for each electric field.

Chapter 6

ER/NR Discrimination in the PIXeY Detector

In the final Chapter dealing with liquid xenon detectors, we explore ER/NR discrimination again, but with the Particle Identification in Xenon at Yale (PIXeY) experiment. PIXeY was a small, R&D-scale xenon TPC that operated in the mid-2010s at Yale University to study energy depositions in xenon. As an above-ground detector, it was not capable of setting competitive dark matter limits. However, PIXeY scientists investigated dynamics of the xenon TPC technology in order to inform decisions for large experiments like LUX and LZ.

Here, we effectively redo our charge-to-light discrimination analysis from Chapter 4, using PIXeY data. The advantage of PIXeY is that as a smaller experiment, it is able to achieve much higher drift electric fields than LUX was. Recall that the maximum drift field bin in our LUX analysis was 500 V/cm; in contrast, we have PIXeY data at fields up to 1951 V/cm. Whereas LUX showed us a minimum leakage at 300 V/cm in the 50–500 V/cm range, PIXeY allows us to determine how discrimination changes at much higher fields. As we will describe, the PIXeY data suffer from high background rates (and some hardware issues), so the measured electronic recoil leakage is much higher than observed in LUX and other experiments. Nevertheless, this analysis still provides unique information on the ER and NR bands' relationships to drift field, and thus on the relationship between discrimination and drift field.

The data in this Chapter were collected in 2014–15 by scientists in the Daniel McKinsey group when it was based at Yale. These include, but are not limited to, Ethan Bernard, Elizabeth Boulton, Blair Edwards, Nicole Larsen, Brian Tennyson, Lucie Tvrznikova, and Chris Wahl. The dissertation author was not involved in data collection. However, the data analysis presented was done almost solely by the author. The exceptions are portions of the data processing (Sec. 6.2) where the author based his processing framework off previous work by Drs. Blair Edwards and Quentin Riffard.

6.1 The Particle Identification in Xenon at Yale (PIXeY) Experiment

6.1.1 About the detector

PIXeY was a two-phase liquid/gas xenon time projection chamber, much like LUX. It ran from 2012–2016 at Yale University, during which time multiple calibration campaigns were executed. Sources included ^{22}Na , ^{37}Ar , ^{83m}Kr , D-D neutrons, ^{228}Th , ^{133}Ba , and pulsed LED photons.

A schematic of PIXeY is shown in Fig. 6.1. The detector has seven Hamamatsu R8778 PMTs each in the top and bottom array, arranged hexagonally. The xenon vessel is similarly hexagonal with an 18.4 cm width at its widest point. The gate grid is 5.05 cm above the cathode grid, and the anode grid is another 0.74 cm above

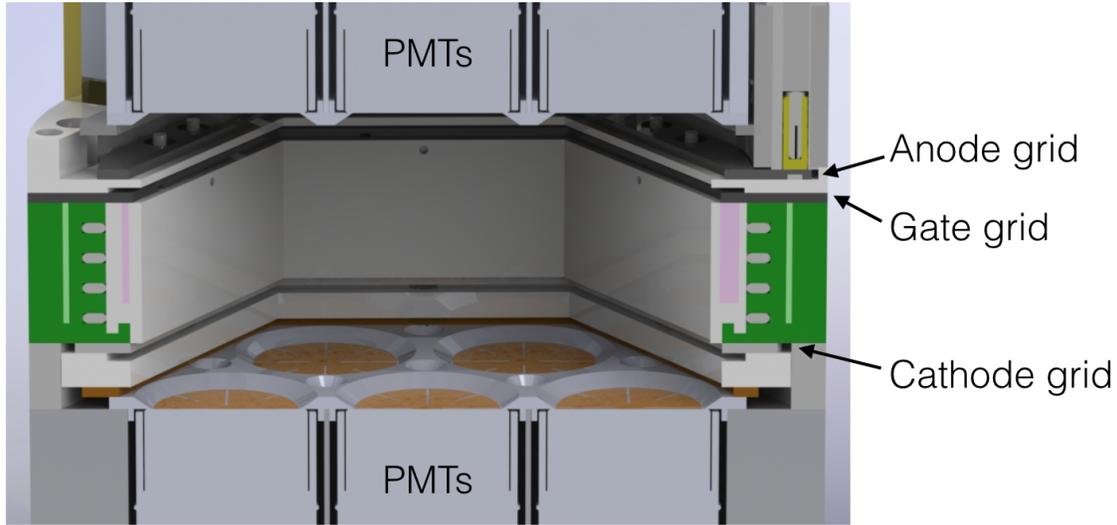


Figure 6.1: A rendering of the PIXeY inner region, with labels for the PMTs and grids. Liquid xenon fills the vessel, apart from the top layer of gas above the gate grid. Figure copied from [158].

that. Data acquisition is triggered on the S2 pulse; the PMT waveforms are sampled at 250 MHz with a 12-bit CAEN V1720 ADC.

6.1.2 Run Details

Although PIXeY ran from 2012–2016, we focus on data taken between October 2014 and April 2015. We focus on data taken with three calibration sources: D-D neutrons to calibrate nuclear recoils, ^{83m}Kr , and ^{22}Na to calibrate electronic recoils. The D-D and ^{83m}Kr sources have been described at length in previous chapters, so we do not repeat that discussion. ^{22}Na decays via β^+ emission to an excited state of ^{22}Ne , which de-excites via emission of a 1275 keV γ . The emitted β^+ quickly annihilates in the ^{22}Na source (which includes casing in addition to the radioactive isotope), and two 511 keV γ photons are emitted. These two photons are emitted in any direction, but always back-to-back. Therefore, we place a separate detector composed of a NaI scintillating crystal and a PMT directly opposite the Xe TPC, and we connect the PMT to our DAQ system. If one 511 keV photon propagates in the direction of the xenon, the other will be incident on the NaI detector. This schematic is shown in Fig. 6.2. We identify decays by searching for an electronic recoil (typically, a Compton scatter) in the xenon and a time-coincident deposition in the NaI.

During our run period, the trigger and veto system were set in several conditions. The trigger time position in a waveform was varied, and the veto was adjusted

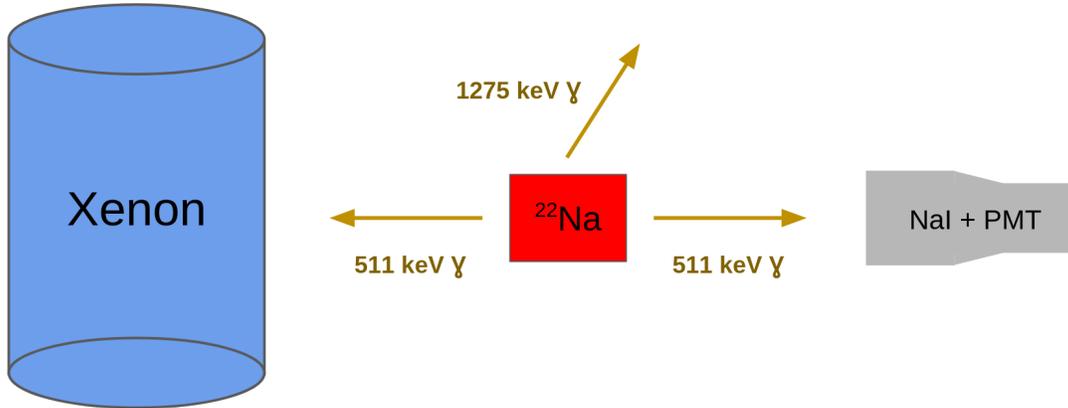


Figure 6.2: A cartoon depicting our detection mechanism for ^{22}Na decays. We detect γ photons in both the xenon volume and an external NaI detector, partnered with a PMT.

to include/exclude events with high single-electron rates or large S2 pulses that could saturate the digitizer. We will show later that the veto variation substantially impacts the energy range of our analysis. For clarity, we divide our data into four sets of acquisitions, based on the veto settings at a given time: November 2014, December 2014, March 2015, and April 2015. These sets correspond to data taken between November 19–21, 2014; November 26, 2014–January 19, 2015; March 13–19, 2015; and March 20–April 29, 2015. Chronologically-later acquisition sets have stricter vetoes.

D-D calibrations are present during the December 2014 and March 2015 sets, but unfortunately, the March 2015 data files contain very few triggered recoil events. The reason for this is unknown, but there are two hypotheses: the neutron flux was set too low, or the veto was set improperly and rejected true nuclear recoils. As a result, we only use D-D data from December 2014. However, there is another complication; there is an issue in the top-center PMT for all data (not just D-D data) taken between December 14–27, 2014. That PMT’s light collection rate is lower than the other PMTs’ by over two orders of magnitude, and lower than its own rate during other periods of time. As (mis-)fortune would have it, all the D-D data was acquired during that two-week period. Our solution is to entirely exclude the top-center PMT from our analysis, for both ER and NR acquisitions. Using $^{83\text{m}}\text{Kr}$ data, we determine that this adjustment reduces g_1 by 3.7% and g_2 by 17.5%. The change in g_2 is irrelevant to discrimination, since the ER and NR bands shift by the same amount. The change in g_1 does have an impact on discrimination, as we

saw in Sec. 4.5.1.1, but a 3.7% decrease should not be significant¹. Furthermore, this modification means that our position reconstruction is less accurate; xy -position is reconstructed further towards the walls of the detector than the true recoil position. We account for this by updating our fiducial cut appropriately.

During this period, the drift and extraction fields are set via the cathode and anode grid voltages, respectively. The cathode voltage is set to seven values between 250 V and 10 kV. The drift field can be estimated by dividing voltage by the drift length, but B. Edwards used COMSOL simulations to determine the actual mean field in the xenon volume. The values in our analysis are 69, 117, 213, 407, 696, 986, and 1951 V/cm. (These are approximated as 50, 100, 200, 400, 700, 1000, and 2000 in some previous PIXeY papers.)

Meanwhile, the anode voltage could be set to 17 values between 3.01 kV and 8.82 kV. Since xenon has a different dielectric constant in liquid and gas, the extraction field is similarly different in liquid and gas; as in LUX, the liquid level is above the gate grid. The extraction field ranges from 2.41–7.08 kV/cm in liquid and 4.45–13.09 kV/cm in gas. The reader can find all 17 values in a previous PIXeY paper on extraction efficiency [159], but in our analysis, we only use data with an anode voltage of 5.38 kV, corresponding to a liquid field of 4.31 kV/cm and gas field of 7.98 kV/cm. This also corresponds to an extraction efficiency of 72.1%.

The two photoelectron emission probability is not measured in PIXeY, but it was measured in the same Hamamatsu R8778 PMTs by Faham et al. to be 17.3% [123]. Therefore, we use this probability to convert between units of [phe] and [phd] (photoelectrons and photons detected).

6.2 Data Processing

All of the data are processed through a multi-stage pipeline to identify the baseline, calculate the single photoelectron area, identify pulses, and classify pulses. This is done event-by-event, where an “event” is the waveform recorded in all 14 PMTs, and the NaI PMT for ²²Na calibrations, after a trigger. The time window for an event is 32768 samples, or 131.072 μ s, and the trigger is somewhere in the middle of the event. Typically, the trigger reacts to the S2 pulse, but sometimes it finds the S1 pulse or something else.

To begin our discussion, we show a fully processed event in Fig. 6.3 (nearly identical to Fig. 3.15). In this case, the event is a 10 keV ER. The waveform is in the top panel, summed across all 14 xenon PMTs, along with seven pulses highlighted. The second panel shows the waveform for the same event in the NaI detector. The third panel shows zoomed-in glimpses of the seven pulses, where each PMT channel is shown in a different color, and the classification of each pulse. The

¹Note: By a 3.7% decrease, we mean that, for example, g_1 changes from 0.090 to 0.087, not from 0.090 to 0.053.

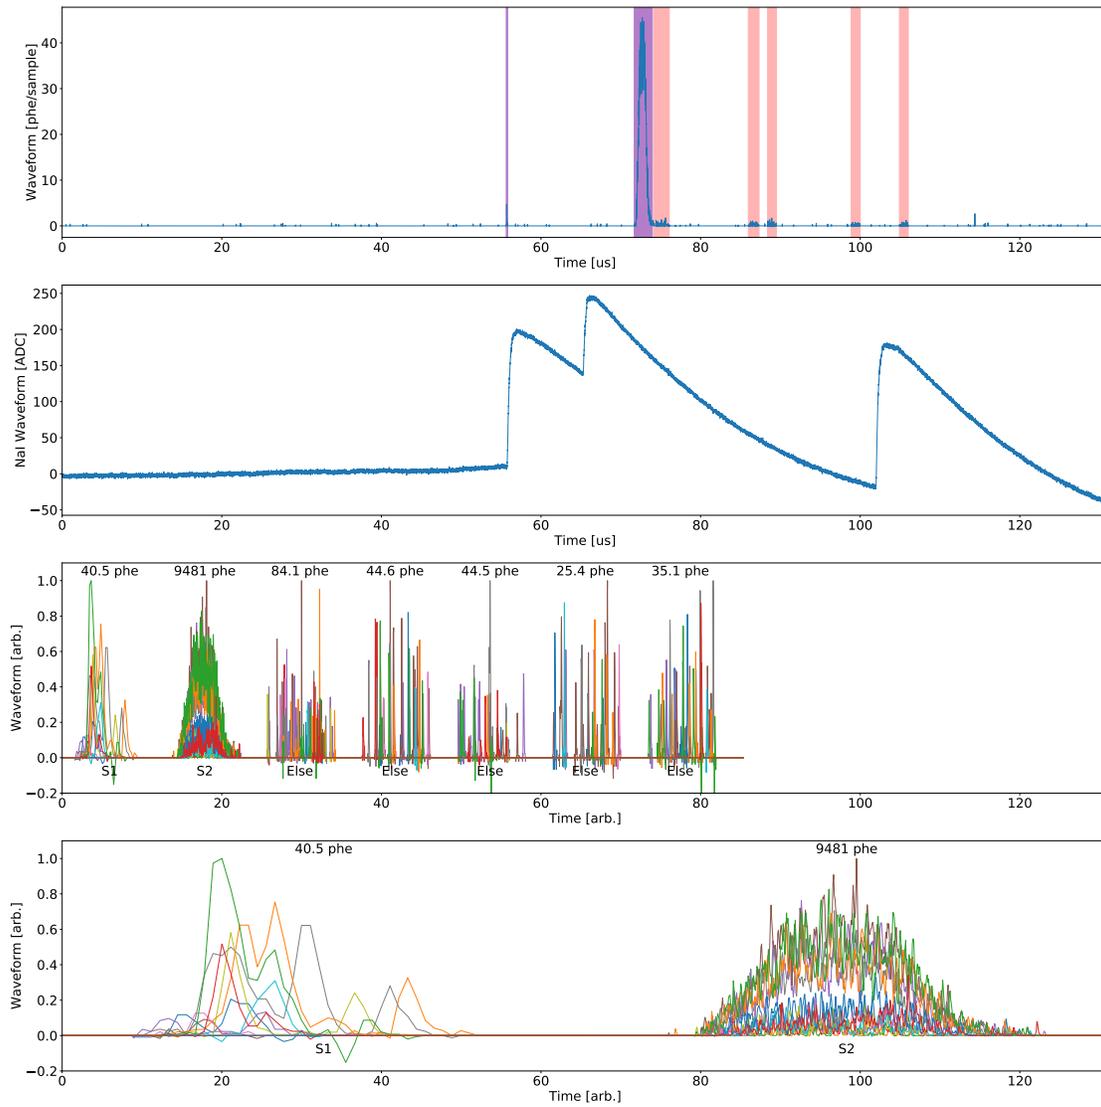


Figure 6.3: An example waveform for a 10 keV electronic recoil event in the PIXeY detector. See text for details.

fourth panel zooms in on the S1 and S2 pulses, specifically. (In Fig. 3.15, the reader can also see the PMT hit pattern for the S1 and S2 pulses.)

The first step of processing is baseline finding. Each PMT waveform starts in units of ADC counts, from 0 to 2^{12} . The flat portion of the waveform is isolated and subtracted from the entire waveform, so that the baseline is at zero. In addition, samples with amplitudes below a certain threshold are set to zero to remove noise, while keeping continuous pulses untouched. Baseline subtraction is also done for the NaI pulses, but the baseline is not perfectly flat (see the region $t < 50\mu s$ in Fig. 6.3. However, we only use the arrival time of the NaI pulse, not the height or area, so a poor baseline is not a concern.

Next, the xenon PMTs and NaI PMT are temporally aligned with each other. Seven of the xenon PMTs are on one ADC digitizer card, and seven are on another, and the two clocks are not perfectly aligned. During operation, the NaI signal was duplicated and sent to the eighth channel of both cards. Thus, we can use the NaI signal to determine the time offset and adjust the xenon waveforms accordingly; the offset is typically between 5–20 samples, centered on 12 samples.

In parallel with the processing pipeline, we determine the single-photoelectron (SPHE) area for each PMT. Equivalently, what is the pulse area in $mV\cdot ns$ or ADC-samples for one photoelectron? The SPHE area ranges between 80–120 $mV\cdot ns$, varying based on PMT and acquisition. Thus, for each acquisition, we calculate the SPHE area for each PMT. We find SPHE pulses using the tools described later in the processing pipeline. Each pulse is fit to a Gaussian, and the area is determined. We create a histogram of the pulse areas, fit it to a skew-Gaussian, and extract the mean, which gives us the conversion between $[mV\cdot ns]$ and $[phe]$. This process is shown in Fig. 6.4, which shows the histogram of pulse areas after SPHE calibration for a single PMT, for an arbitrary ^{83m}Kr data set. We see that the mean single photoelectron area is 1 phe, as forced by our procedure, and the width is 0.3 phe. Across all acquisitions and PMTs, the distribution width ranges between 0.30–0.35 phe.

Returning to the processing pipeline, we divide the waveform by the SPHE area for that acquisition, for each PMT. In doing so, we convert the waveform amplitude from units of ADC counts to phe/sample. The waveform from all 14 PMT channels are added together. By this step, we have arrived at the waveform in the top panel of Fig. 6.3, apart from the highlighted regions.

Next, pulses are identified in the summed waveform. A “pulse” is defined as a continuous time window where some potentially interesting signal exists in the waveform. Pulse finding is done in two stages. First, the algorithm looks for large S2-like pulses. We define an observable $A \times B$, where A is the rolling RMS of the samples in a specified time window, and B is the derivative of the rolling average of the waveform in that same time window. This is a complicated variable, so we show the observable in the second panel of Fig. 6.5, calculated in the vicinity of an S2. A pulse is identified when this observable goes above a certain threshold, crosses zero and goes below a negative threshold, and returns to zero, within a certain

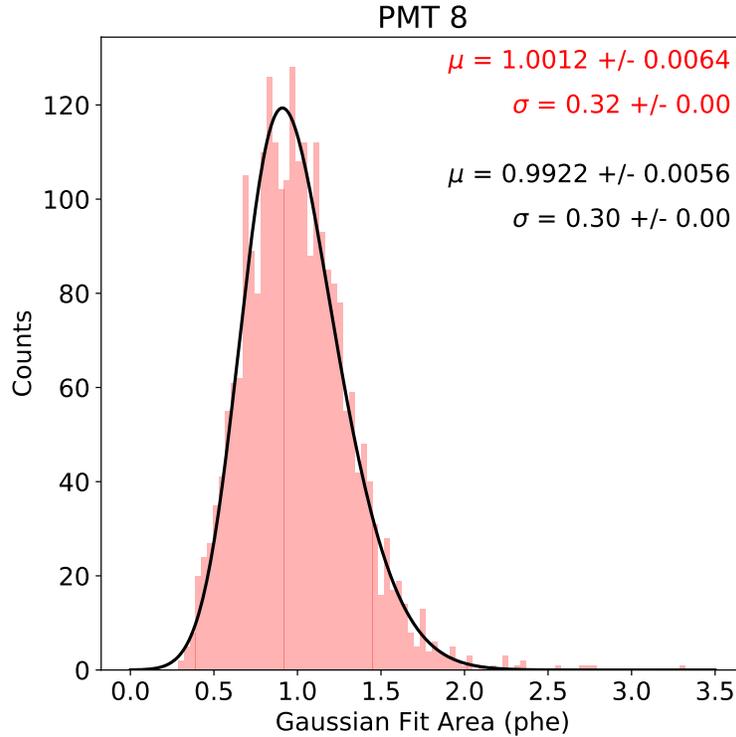


Figure 6.4: A collection of single photoelectron pulses for a single PMT (PMT 8) from a ^{83m}Kr data set, after SPHE area calibration. Each pulse is fit to a Gaussian, and the area is extracted. The plot shows a histogram of these areas, along with a fit to a skew-Gaussian function (in black). The histogram mean and standard deviation are shown in red text, whereas the black text shows the mean and standard deviation from the skew-Gaussian fit.

time. The time window and threshold are set so that only S2 pulses meet these criteria (and occasionally, enormous S1s), and overlapping S2 pulses are identified as such. Second, after excluding any regions identified as pulses by the first step, the algorithm looks directly at the waveform and identifies pulses by determining when it crosses some threshold. The algorithm also merges clusters of small pulses that occur quickly after each other. This accounts for single electrons extracted from the liquid, which look like several S1s in succession. The bottom two panels of Fig. 6.5 highlight pulses identified through the second stage; the three long pulses in the bottom panel are single electrons².

In a given event, we may see several dozens of pulses. Most of these are single

²Note: pulse-finding is the primary segment of this analysis where the author relied on previous work. The two stages were originally developed by Riffard and Edwards, respectively. However, the author modified these procedures to improve their fidelity.

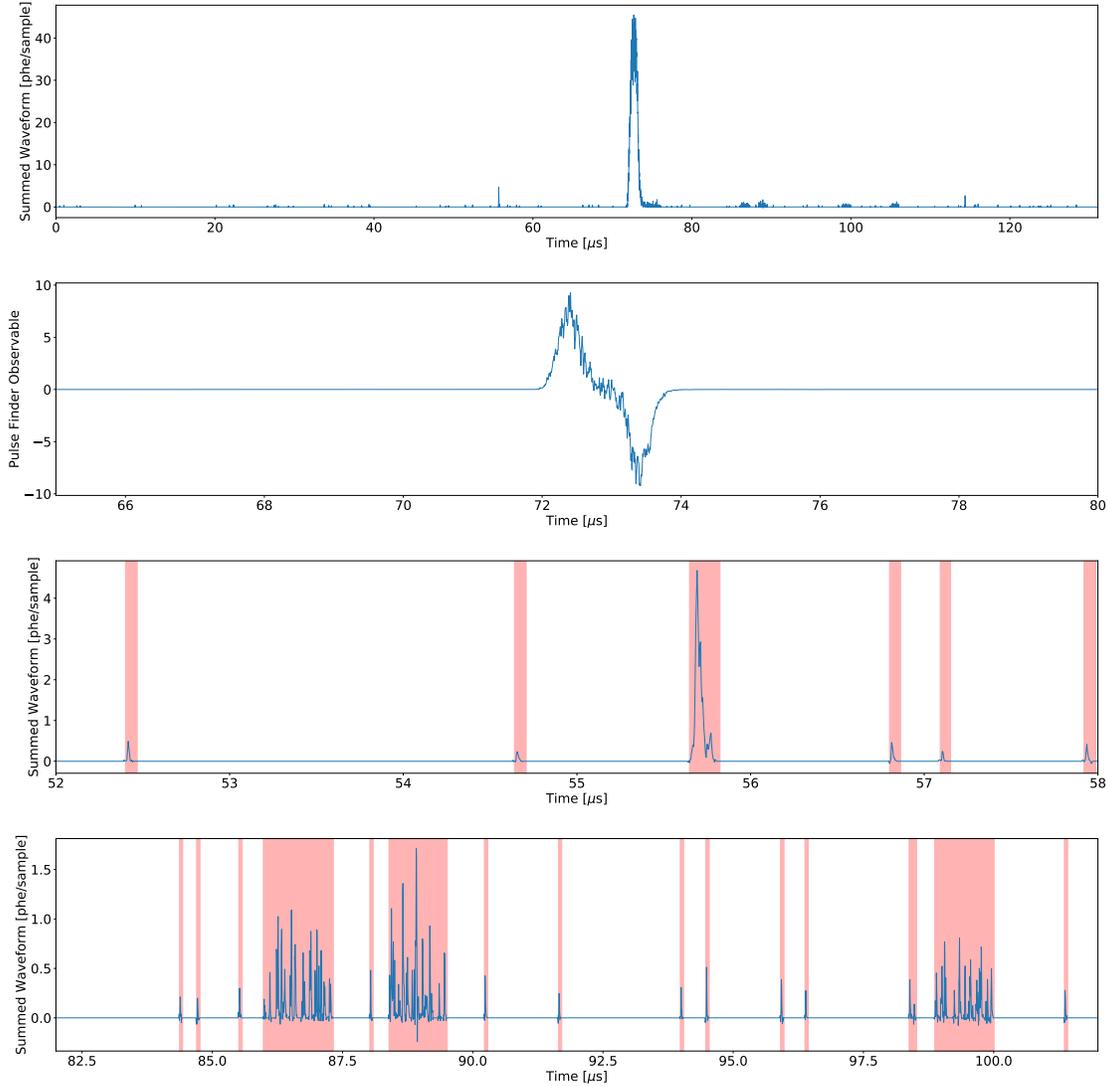


Figure 6.5: Pulse-finding in the same event from Fig. 6.3. The top panel is the waveform summed across all 14 PMTs. The second panel shows the observable used for the first step of pulse-finding, and we zoom in on the region around the S2 pulse (see the x -axis). Pulses are identified when this observable goes above or below a certain threshold. The bottom two panels show zoom-ins of different regions of the waveform, and the identified pulses are highlighted. All the highlighted pulses are found in the second stage of pulse-finding. Apart from the S1 at $55.8 \mu\text{s}$ and the three long single electrons at 86 , 88 , and $99 \mu\text{s}$, none of the highlighted pulses meet the 3-fold coincidence requirement.

photoelectrons, each produced by a single photon hitting a PMT. These are used for the previously-mentioned SPHE area calibration, but otherwise, they are a background. We therefore require each pulse to have a signal in at least three PMTs; this is called “three-fold coincidence.” This removes all single-photon hits, as well as spurious two-photon hits, which are less common but still frequent enough to hinder our analysis. Through simulations, we estimate that 0.1% of event waveforms will contain a spurious two-photon pulse that is coincident with an unrelated NaI pulse, but only 0.01% of waveforms will contain a spurious three-photon, NaI-coincident pulse. The earlier Fig. 6.3 shows the effects of three-fold coincidence; the top and third panels only show identified pulses that meet this requirement. The coincidence requirement also sets the S1 and energy threshold for our analysis. Recall that the SPHE areas follow some distribution for each PMT (e.g., Fig. 6.4); a given photoelectron spike needs to be at least 0.5 phe to produce a detectable pulse with reasonable probability. Therefore, our analysis threshold is an S1 signal of 1.5 phe.

All pulses that meet the 3-fold coincidence requirement are then classified into three categories: S1, S2, and Else. Most of the Else-classified pulses are single electrons. The classification criteria were determined by manually scanning >1000 events and identifying S1 and S2 pulses. Several pulses-level variables (reduced quantities, or RQs, in LUX parlance) were calculated for each pulse. The RQs that best discriminate between S1, S2, and other pulses were chosen. Specifically, these were the pulse width, pulse amplitude, the quality of a Gaussian-fit, and the fraction of the pulse area in the first 90 ns. An S2 is required to have an area equal to at least 4 extracted electrons, about 150 phe. We do not have precise estimates of the efficiency of these classification criteria, but based on manual waveform scanning, it seems to be >90%.

For each pulse, dozens of RQs are computed and stored, including pulse area, pulse width, pulse amplitude, various fits, the fraction of light in a given time window, asymmetry between the top and bottom array, the maximum fraction of light in any PMT, and more. A few of these RQs are also computed and stored for each PMT separately. For the S2 pulses only, the corresponding x and y positions are computed. We do not use the Mercury algorithm [128], as LUX did. Instead, we simply calculate the center-of-mass coordinates, using the coordinates of each PMT and the fraction of S2 light in each PMT.

At this stage, processing is done. However, our analysis only focuses on single scatters, events with one S1 and one S2 pulse, so we need to isolate these events. This is done in two stages. First, for each event, we check for exactly one S2 pulse in the waveform. If so, we scan the previous T samples, where T is the maximum drift time of the TPC; this depends on the drift field, but ranges from 29 μ s to 53 μ s. We search for exactly one S1 pulse. For ^{22}Na calibrations, that S1 pulse must be coincident with a pulse in the NaI. If these criteria are met, the event is labeled a single-scatter. Otherwise, we move to the second stage, which roughly is the same, but in reverse. We search for one S1 pulse first, then scan the next T samples for

one S2 pulse. If we succeed, the event is still labeled a single-scatter. The vast majority of single-scatter events are found in the first stage, but we have confirmed there is no observable difference between the events gathered in either stage. For all single-scatter events, we store the relevant RQs and proceed to the analysis!

6.3 Detector Calibration with ^{83m}Kr

As described in previous chapters, ^{83m}Kr is a useful calibration signal for xenon TPCs, and it was deployed in PIXeY. A complete analysis of ^{83m}Kr yields from PIXeY data can be found in [75]. Here, though, we only discuss its use as a tool to measure g_1 , g_2 , and position dependence of S1 and S2 signals.

6.3.1 Position-Dependence of S1 and S2

Figure 6.6 shows our strategy. For each acquisition, we model the S1- and S2-dependence on radius and z (or drift time) separately. The S2 vs. drift dependence is modeled as exponential, assuming a constant probability of an ionized electron to be captured by an electronegative impurity (i.e., the electron lifetime). The other three relationships are modeled as quadratic. There is no physical motivation for this; a quadratic function fits the data well. We calculate the S1 and S2 position dependencies and invert them to calculate the position corrections. These four corrections are then multiplied into each S1 and S2 signal to give S1c and S2c. The right column of Fig. 6.6 shows the distributions of S1c and S2c; the reader can see that they are flat with r , z , and drift time. Note that we use units of [phd], after accounting for the 17.3% double photo-electron emission probability.

We also need to determine how, and if, these corrections vary as a function of time. Three of the corrections are observed to be independent of time. We can see this in Fig. 6.7, which shows the best-fit quadratic, linear, and constant term of the S1 z -correction vs. date. The terms don't change, beyond their statistical uncertainties. Therefore, we can have just one model of S1 vs. z position corrections for all the data. The same is true for the S1 r -corrections and S2 z -corrections.

However, the electron lifetime does vary significantly from acquisition to acquisition, as shown in Fig. 6.8. We observe that the lifetime is relatively constant before February 2015. During March and April, the lifetime varies substantially, but the lifetime is long enough (the drift time variation is small enough) that the statistical uncertainties are also quite large. Therefore, in our analysis, we define two electron lifetime values: $90.07 \mu\text{s}$ before February 1, 2015, and $289.63 \mu\text{s}$ after. This also corresponds to the acquisition sets we defined earlier: November and December 2014 are “early” data, and March and April 2015 are “late” data.

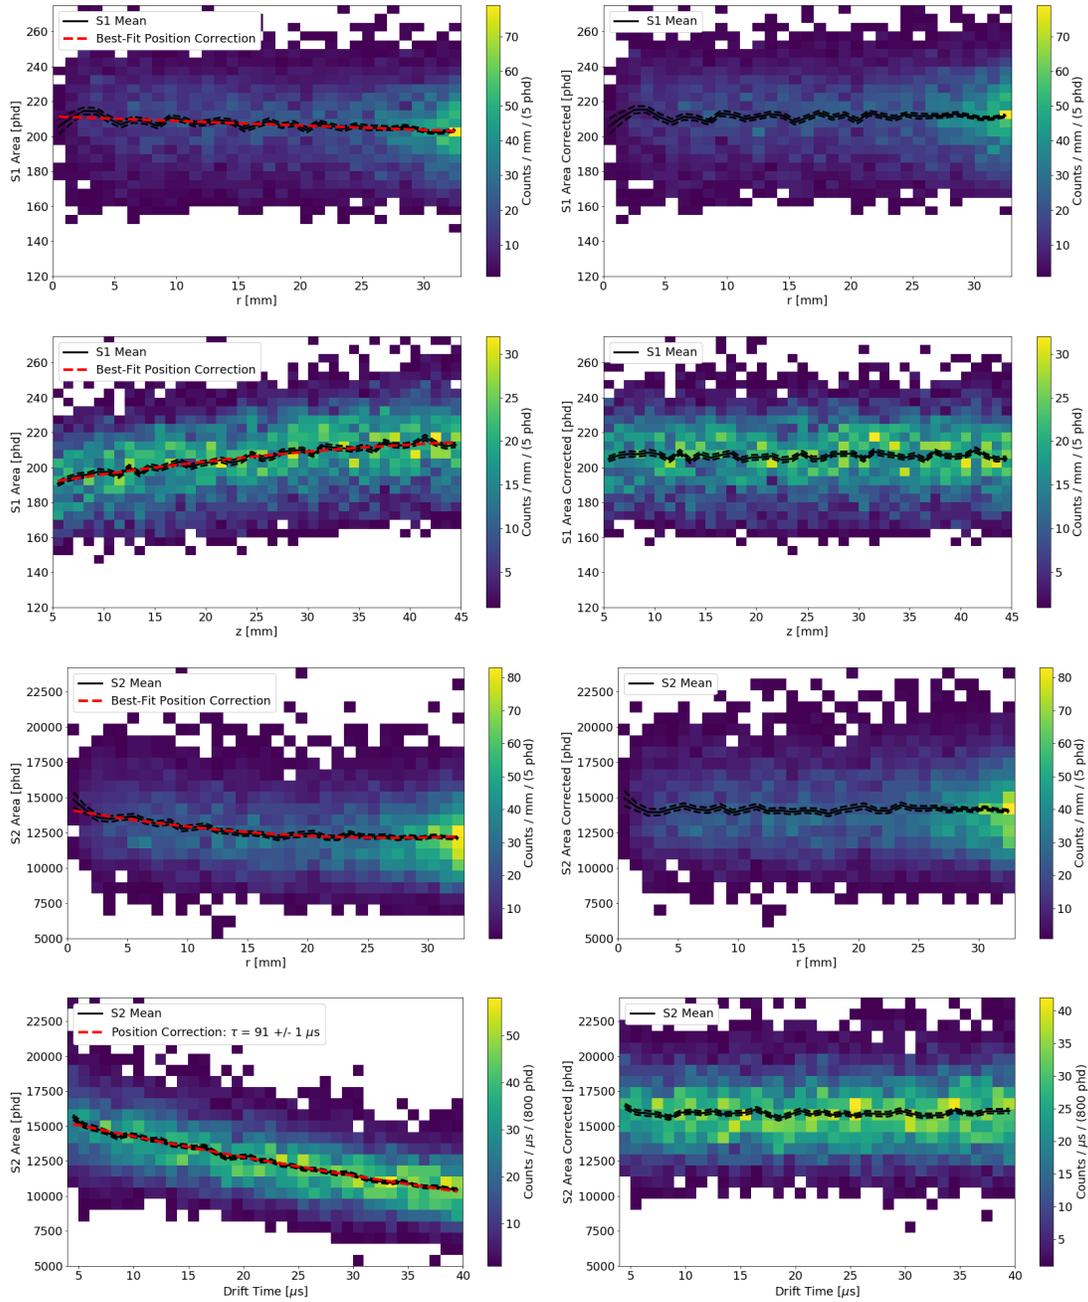


Figure 6.6: The S1 radial, S1 z , S2 radial, and S2 z position corrections, for one ^{83m}Kr acquisition. (*Left*) A 2D histogram of signal size vs. r , z , or drift time. The arithmetic mean is shown in black. We then fit the points to either an exponential function (for S2 vs. drift time) or a quadratic function (for all others), shown in red. These fits are used to calculate position corrections. (*Right*) The same 2D histogram, but with the position corrections applied and the mean re-calculated.

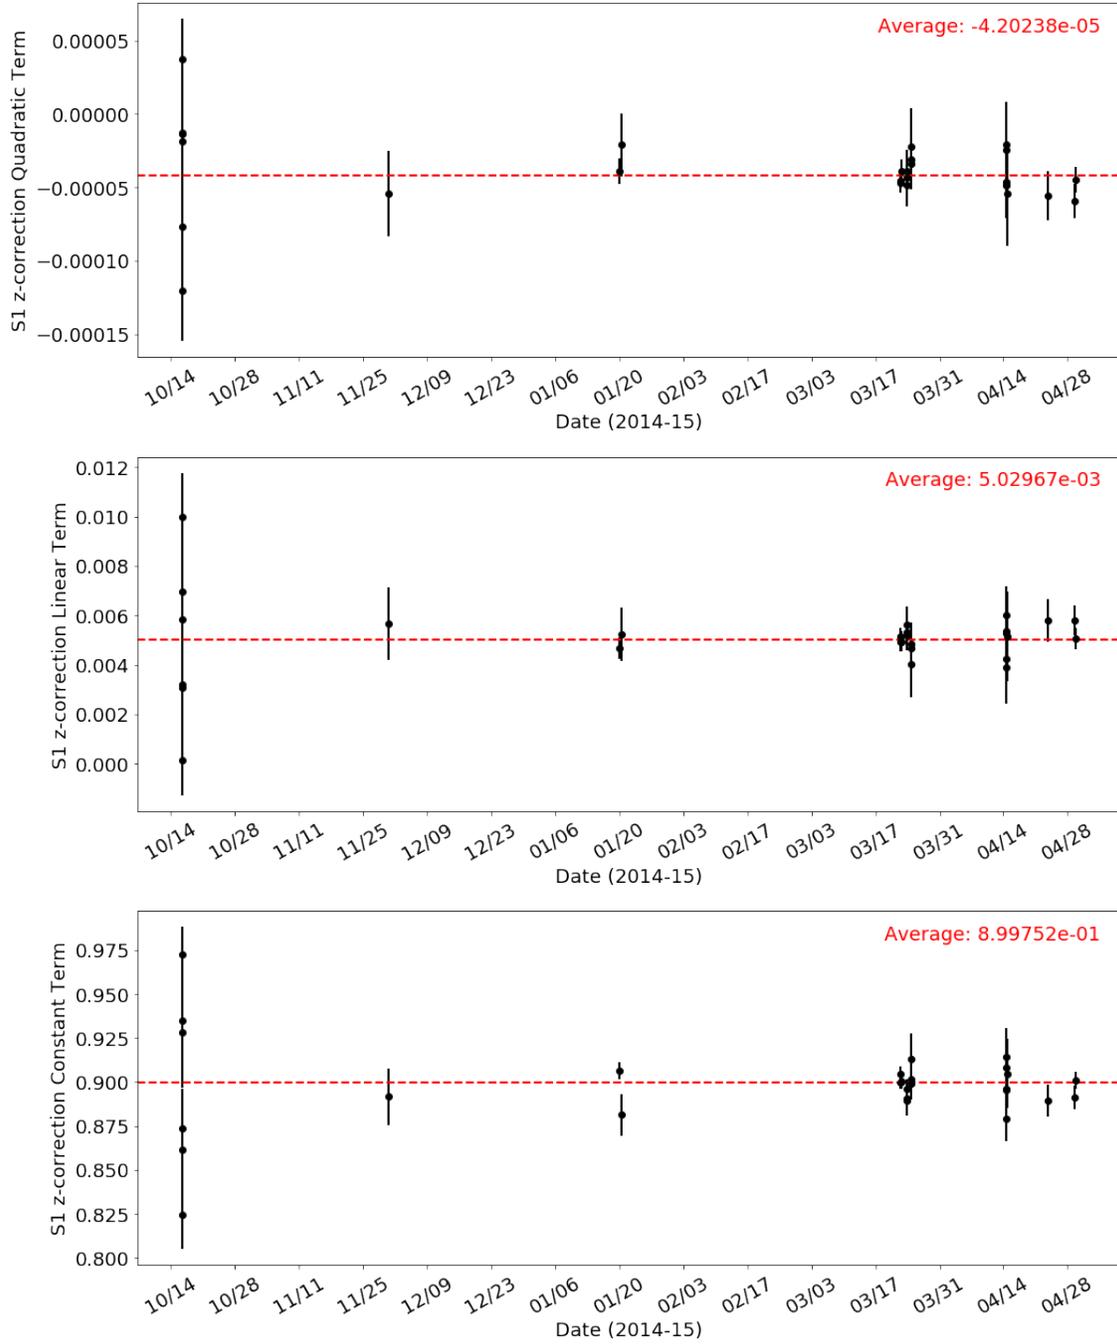


Figure 6.7: The quadratic, linear, and constant terms for the S1 z -correction, vs. date. The error-weighted average of all the values is shown in red.

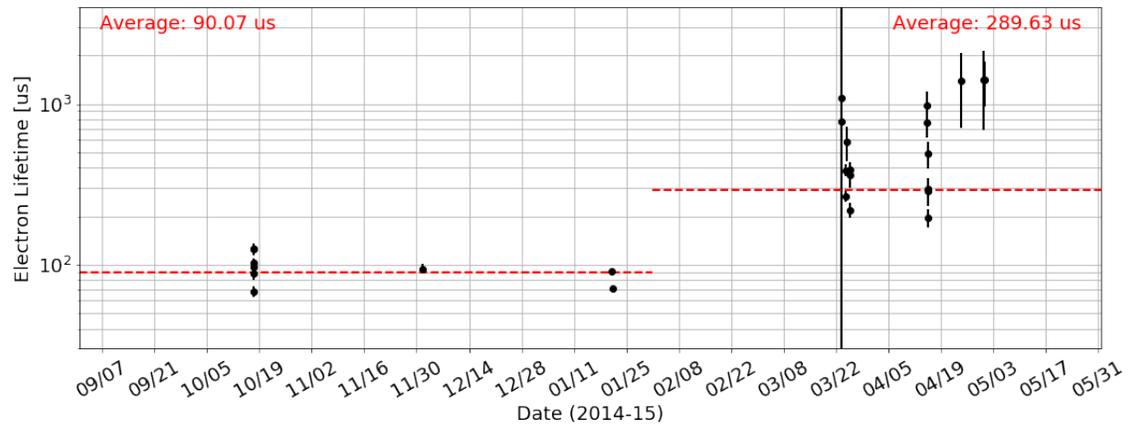


Figure 6.8: The electron lifetime vs. date. The error-weighted average of all the values is shown in red, in two segments: before and after February 1, 2015.

6.3.2 Definition of Fiducial Region

The ^{83m}Kr data can also be used to define the fiducial region, i.e. the region where the analysis is done. We define a cut in both r and z . Drift time is not used because this changes based on drift field.

To define the cut, we look at the distribution of events in 3D position from a single ^{83m}Kr acquisition, as shown in Fig. 6.9. Most events are located near the walls of the detector, but these are often backgrounds. We choose the inner region of the detector, where the sample of events is relatively pure. The fiducial cut is defined visually by excluding event-dense regions corresponding to the inner radius of the hexagonal detector, and the cathode and gate grids.

We define two fiducial cuts, corresponding to the removal vs. retention of the top-center PMT. When this PMT is removed, the position reconstruction is affected; events are reconstructed at 12% larger radii. Both cuts include $4.5 < (z/\text{mm}) < 46$. The radial cut is $r < 34$ mm when the top-center PMT is included and $r < 38$ mm when it is excluded.

6.3.3 Determination of g_1 and g_2

Using the position-corrected S1c and S2c, we are able to measure g_1 and g_2 , using the equation for combined energy. This was shown in an earlier Chapter in Eq. 3.13, but we restate it below.

$$E = W \left(\frac{S1c}{g_1} + \frac{S2c}{g_2} \right). \quad (6.1)$$

We know the energy of the ^{83m}Kr recoils; we only use events with one S1 signal, so the entire 41.5 keV deposition is captured. We also know $W = 13.7$ eV, and S1c and S2c are measured for a variety of drift fields. Therefore, we can measure g_1 and g_2 through a linear fit, as done in the left panel of Fig. 6.10. This figure is often called a ‘‘Doke plot,’’ named after Tadayoshi Doke.

Interestingly, we find the Oct 2014–Jan 2015 and March–April 2015 data lie on different lines in the Doke plot. These lines correspond to $g_1 = 0.0885$ and $g_2 = 27.67$ for the early data and $g_1 = 0.0893$ and $g_2 = 30.45$ for the late data. In fact, g_2 is the main change; g_1 doesn’t change much. We know g_2 is the product of electron extraction efficiency and single electron size, and either of those could have changed during the operation of PIXeY.

The data are more confusing when we look individually by field, as done in the right panel of Fig. 6.10. For a given field, the ^{83m}Kr collected later tend to have higher S1c and lower S2c signals. The (quite substantial) increase in S1c is unexpected since g_1 doesn’t change much, and the decrease in S2c contradicts the increase in g_2 . The Doke plot measurements of g_1 and g_2 cannot be wrong, because

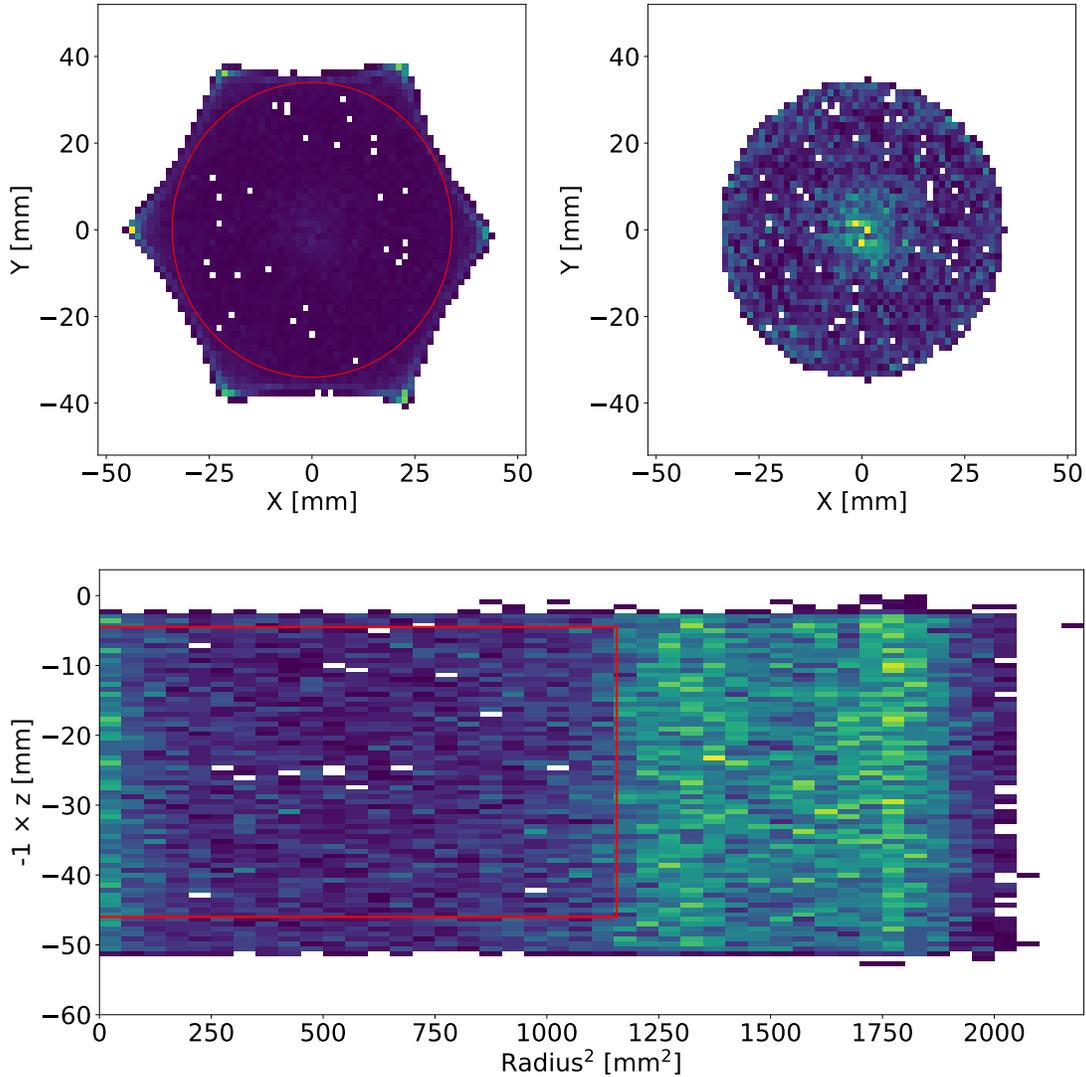


Figure 6.9: Distribution of x , y , z , and r^2 for events from a given ^{83m}Kr acquisition, at top-left and bottom. The fiducial cuts $4.5 < (z/\text{mm}) < 46$ and $r < 34$ mm are highlighted in red, and the top-right shows the event distribution after implementing the cuts. The excess of events at radii of $r^2 = 1300$ mm² and 1800 mm² correspond to the inner and outer radii of the hexagonal detector. The excess of events at $r = 0$ is a side-effect of the position reconstruction algorithm. An advanced algorithm like Mercury [128] would account for the light response function of each PMT, but we do not implement this.

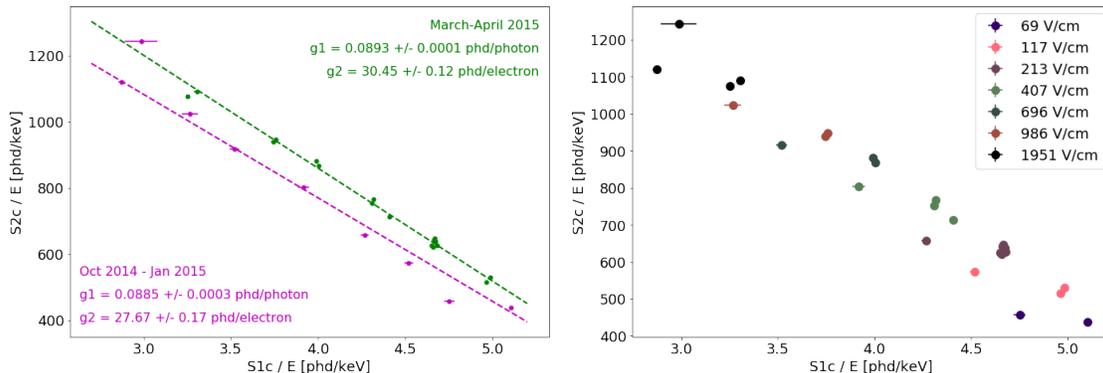


Figure 6.10: The mean S1c and S2c for each acquisition, plotted against each other. (*Left*) We calculate g_1 and g_2 for the early data separately from the late data. (*Right*) The same data, but colored by drift field.

they are just a consequence of quanta conservation. We thus conclude that the drift field must have also changed during the run.

The changing g_2 and drift field between early and late data are unfortunate. They mean that for a given field, the ER band does not stay in the same position for all acquisitions. Consequentially, the leakage measurement is also affected. In Sec. 6.5.2, we confirm a different ER band between acquisitions at the same field. The difference is noticeable, but fortunately, it is small enough to allow us to complete the analysis.

Our final note on this subject is that our measured $\frac{S1c}{g_1} = N_{ph}$ and $\frac{S2c}{g_2} = N_e$ are inconsistent with NEST simulations, either for the early or late acquisitions. This is strange but explainable. As described in 3.3.4, the NEST model includes some unphysical behaviors in order to match both data where the S1a and S1b pulses are merged and data where they are separated. The previous $83m\text{Kr}$ PIXeY analysis [75] was in fact used to write the NEST model, but only the ratio of S1b/S1a was used. Events with a single identified S1 were not included. Therefore, although the inconsistency with NEST deserves further study, it is not unexplainable.

6.4 Nuclear Recoils

6.4.1 NR Data Selection

The data processing and single-scatter event selection described in Sec. 6.2 is still not enough to remove all background events. Thus, quality cuts were developed and implemented on events that had been identified as single-scatters, but might not be recoils of interest. Here, we present these cuts for nuclear recoils.

The most important is the fiducial cut, defined previously. Next, we select

events where the S1 and S2 width fall in an appropriate range. The S2 width cut is a function of both drift field and drift time, since it is a consequence of electron cloud transverse diffusion. Third, we reject events with a high “bad area.” The “total area” of an event is the integrated area, in phd, from the entire waveform. The “good area” is the sum of the S1 and S2 area, and the bad area is the remainder. Consequentially, events with a high bad area fraction tend to have a high rate of single electrons, which can corrupt S1 or S2 identification. Sometimes, these events even have a second S1 or S2 that was not properly identified as such. This cut varies based on the acquisition sets we defined earlier; the stricter the veto settings, the lower the bad area fraction tends to be. The fourth cut is how well the S2 pulse fits a Gaussian function; we demand $r^2 > 0.6$. Fifth, we require the trigger time to be coincident with the S2 pulse, or to be coincident with the S1 pulse if the S1 amplitude is high enough. Sixth and lastly, we require that the S1 light cannot be too concentrated in any one PMT.

The impact of our quality cuts is shown in Fig. 6.11. It shows the NR band from one acquisition on the typical $\log_{10}(S2c/S1c)$ vs. S1c axes. Applying the fiducial cut removes most of the background noise, but our quality cuts improve the purity of the data set even further.

Not shown in Fig. 6.11 are two extra cuts we make when calculating the NR band in the next section. PIXeY was not designed as a low-background experiment, so the background rate is higher than LUX, especially from ER contaminants and wall events. We exclude events above $\mu_{ER} + 0.5\sigma_{ER}$, where these are the ER median and width for a given field. We also exclude events at least 3 standard deviations from the NR median in either direction.

6.4.2 NR Band as a Function of Drift Field

Finally, we compute the NR band, just as we did in Chapter 4 with LUX data. For each drift field, we split the D-D calibration data into bins of S1, create a histogram of $\log_{10}(S2c/S1c)$, fit it to a Gaussian, and extrapolate the median and width. The results are shown in Fig. 6.12. Just as in LUX and other experiments, we see that the NR band is relatively insensitive to drift field. The band width is totally independent of field; the band median shifts slightly upwards with an increased field.

The one exception is the lowest-S1c point for a field of 1951 V/cm, which is at a much higher position (and higher width) than the trend would suggest. Unfortunately, this seems to be an artifact of something in that particular acquisition. That data set, taken on December 17, 2014, is extremely noisy, with more single electrons than almost any other acquisition. A visual scan does confirm that the NR band shifts sharply upward at low S1, and the reason is not understood. As mentioned previously, we also have NR data from March 2015, but it is very sparse. That being said, the March 2015 data with a drift field of 1951 V/cm is quite similar to the 986 V/cm data, consistent with the trend of Fig. 6.12. This suggests a discrepancy

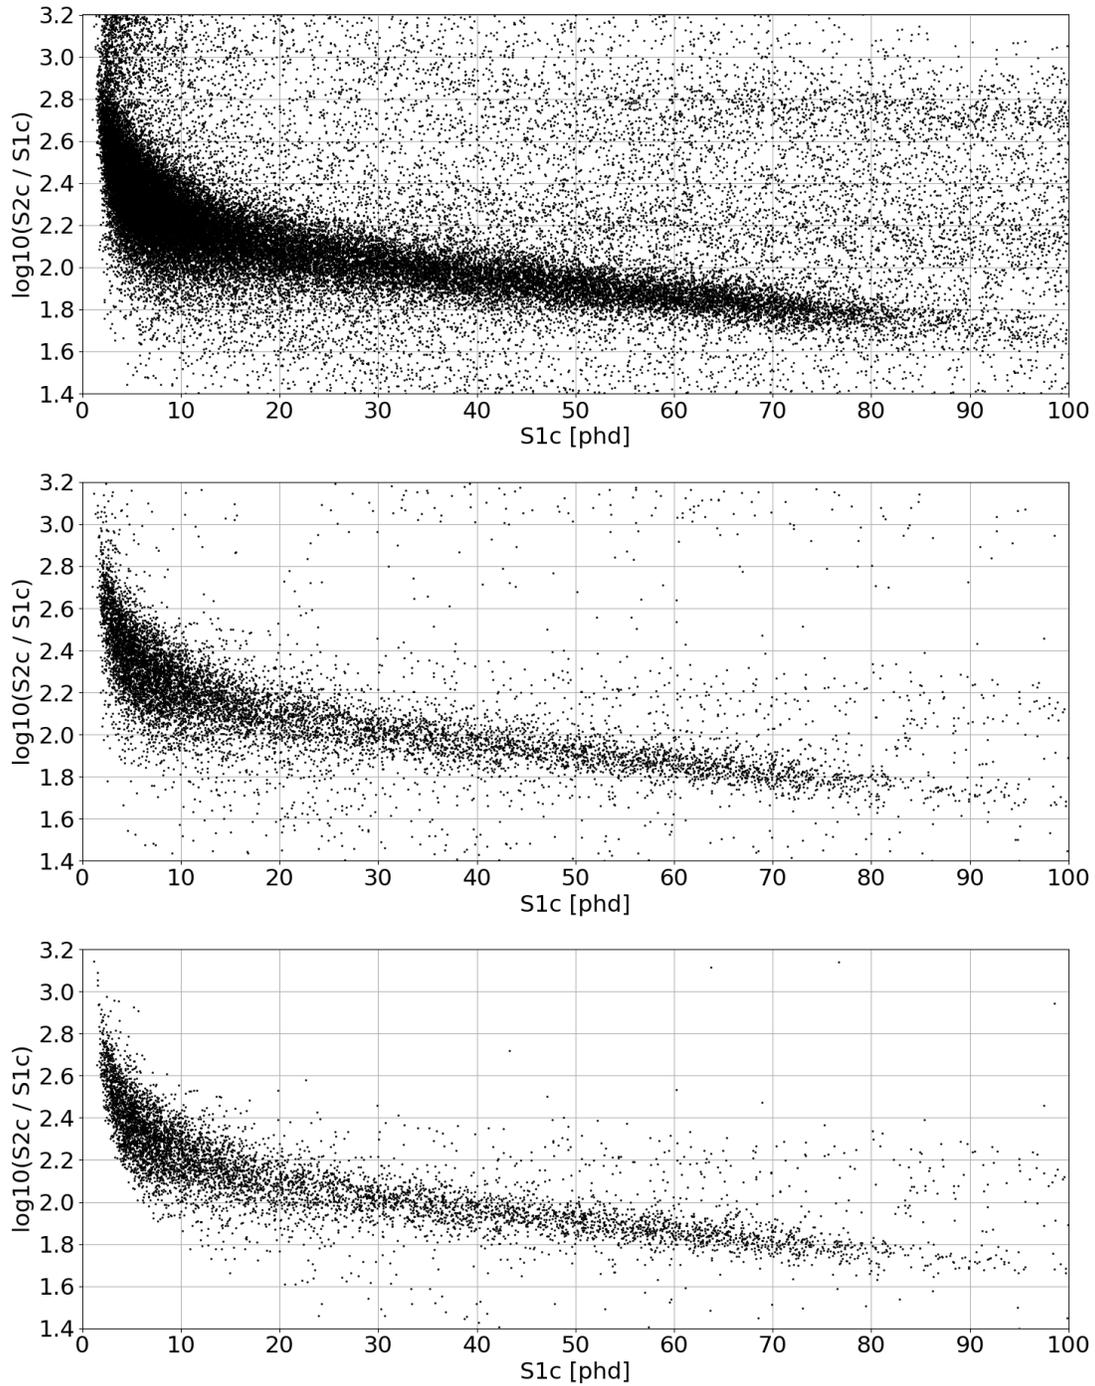


Figure 6.11: The effect of quality cuts on the NR band from one D-D data set. (Top) All single-scatter events. (Middle) Events that pass the fiducial cut. (Bottom) Events that pass the fiducial cut and all additional quality cuts.

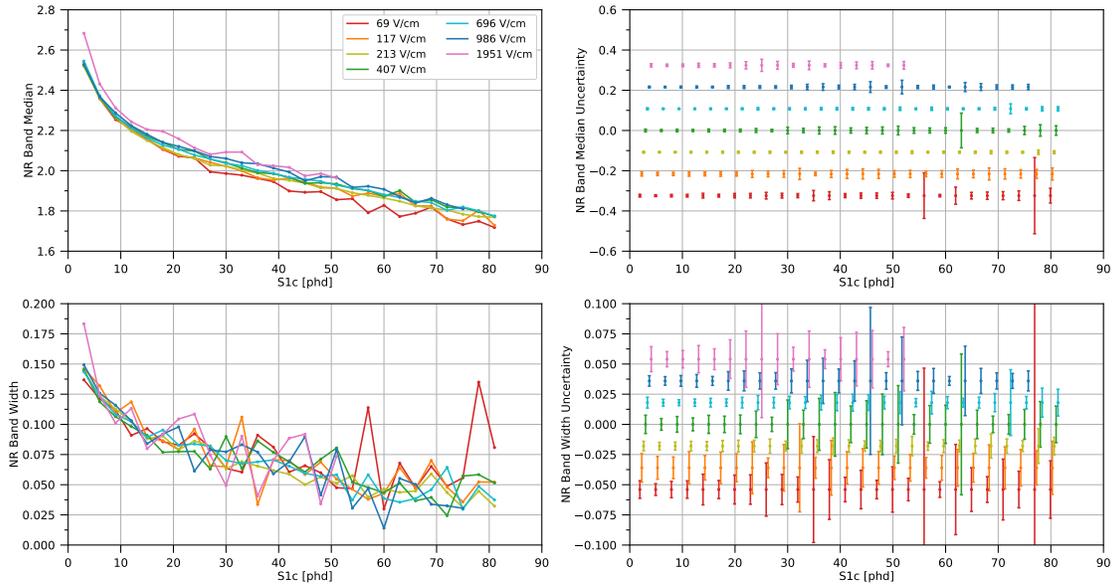


Figure 6.12: The median and width (Gaussian μ and σ of the NR band for several drift fields. The left plots show measurements, and the right plots display error bars corresponding to these measurements. In each row, the y -axes have the same range; the size of the error bars on the right plot can be directly translated to the points on the left plot. For ease of visualization on the right plots, the S1 values are slightly shifted relative to their true value, and the error bars are centered at a different y -value for each field bin.

with the specific acquisition, not necessarily a physical effect. Since we have no other usable data at this field, we keep it, but we remember this feature.

6.5 Electronic Recoils

6.5.1 ER Data Selection

Quality cuts were developed and implemented on the ^{22}Na electronic recoil data, too. They are essentially the same cuts that we used for NR, but with different numerical values. For example, since the S2 pulses are much larger for ER events, we require a better Gaussian fit of $r^2 > 0.8$ and a lower bad area fraction.

6.5.2 ER Band as a Function of Drift Field

The calculation of the ER band in PIXeY data is similar to the LUX calculation, but it is simultaneously simpler and more complicated. It is simpler because we don't use a skew-Gaussian fit. Unfortunately, we have much less PIXeY ER data

than LUX ER data. When we try to fit $\log_{10}(S2c/S1c)$ to a skew-Gaussian, the fit α has huge statistical uncertainties, and it is often fit to a negative value. Therefore, we use a normal Gaussian fit for PIXeY analysis.

The complication is the variation in g_1 , g_2 , drift field, and vetoes between earlier and later data sets. The first three were discussed in Sec. 6.3. The fourth was discussed in Sec. 6.1.2; recall that we split the data into four categories based on the date of acquisition: November 2014, December 2014, March 2015, and April 2015. These roughly correspond to increasingly strict vetoes of single electrons and large S2 pulses. As a result of these variations, the ER band calculated from each set of acquisitions is slightly different.

Let us consider the 407 V/cm data to illustrate this. Look at the top panel of Fig. 6.13, which shows data from November 2014. We see the normal ER band “swoosh,” even seeing the upward turn above ≈ 80 phd where recombination shifts from Thomas-Imel to Doke-Birks. The data goes out to an S1c of 400 phd before starting to taper off, but interestingly, above 200 phd, the data are cut off above some region. We observe that this region aligns perfectly with a line of constant-S2c at a value of $S2c = 120,000$ phd. The reason for this is not entirely understood, but it certainly is related to the veto system, because all other hypotheses have been confirmed false. Manual scans of event waveforms confirm that pulses larger than 120,000 phd do not exist in the data set. The problem is not one of the initial recoil energy distribution, because we clearly see events with higher recoil energy. This is also not an effect of DAQ saturation; large S2 pulses can saturate the 12-bit ADC card, but the area must be higher for that to occur. Thus, we hypothesize that the hardware veto (intentionally or unintentionally) blocked S2 pulses with larger areas during operation.

The consequence of this is that we cannot do our analysis past an S1 of 200 phd or so. For the 400 V/cm data set, this isn’t an issue. After all, our NR data only goes out to 80 phd, as shown in the previous Section. However, this is an issue for the higher-field data, where the ER band is higher, and thus a given S2c cutoff will affect lower S1c signals. Additionally, we see from the three lower panels of Fig. 6.13 that the later data sets have a lower S2c cutoff. The highest-field data sets don’t have November 2014 data, so this will push the maximum S1c even lower. All in all, for each field, our maximum S1c varies, and we will have to account for this when calculating leakage.

Apart from the veto of high-S2 pulses, we also observe that the position of the ER band changes, based on the date of acquisition. We believe this is due to changing g_1 , g_2 , and drift field. See Fig. 6.14 for a demonstration of this. We compare the ER band median and width for all four date sets. The width doesn’t change much, but the median is notably different. December 2014 has the highest ER band median. November 2014 and April 2015 are lower, and March 2015 shows the lowest ER band median. This ordering is true for all drift fields (not all drift fields have November 2014 data, but they all have data from the other three dates). Similarly, the data

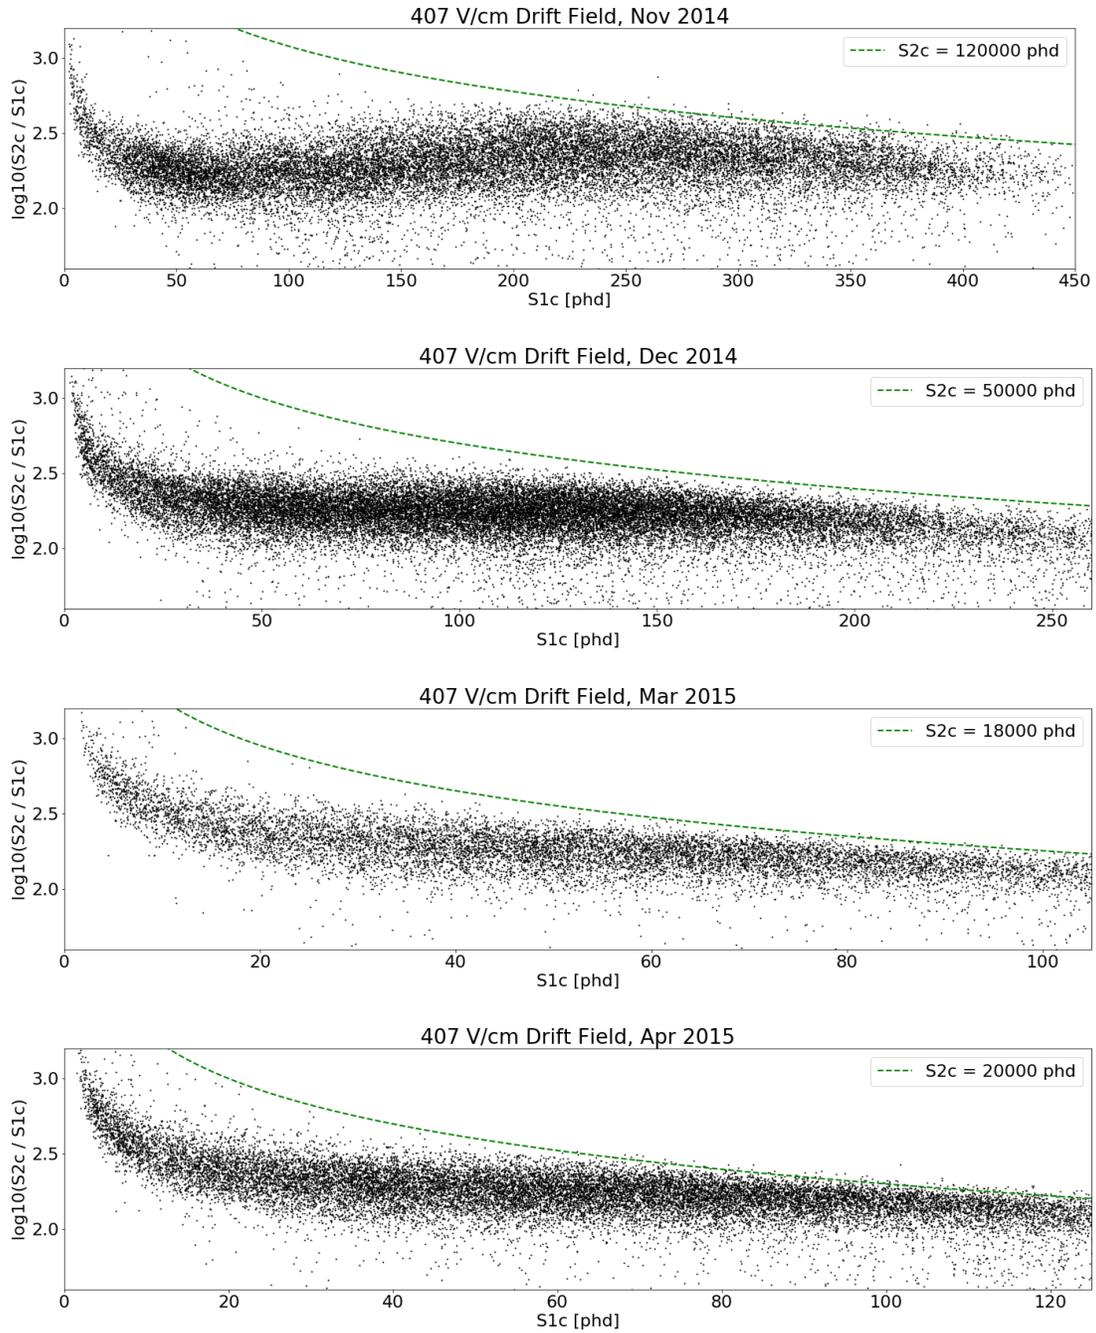


Figure 6.13: Electronic recoil data from ^{22}Na calibrations for a 400 V/cm drift field. Each panel corresponds to a different date set; note the difference in the x -axis. The green line shows a constant- $S2c$ line, above which there are no events because of the hardware veto.

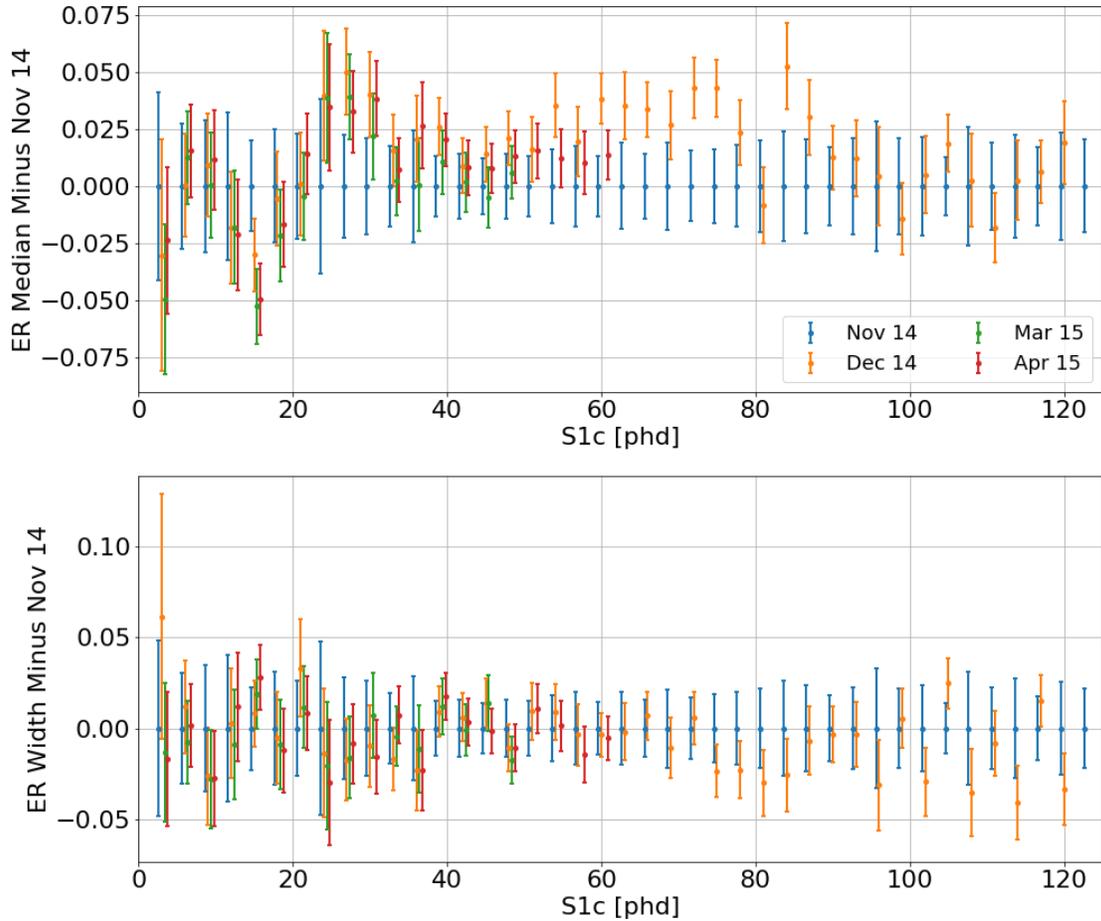


Figure 6.14: A comparison of the ER band median and ER band width between different acquisitions in the 400 V/cm ER data. The reference ER band is Nov 2014.

from all drift fields show little-to-no variation in the band width between data sets.

Regardless, we need to calculate a single ER band for each drift field. We do this by calculating the ER band for each acquisition set, then averaging them together. We treat the median and width separately. For high-S1c bins, we only average the acquisitions where the S2c cutoff has not come into effect.

The ER band median and width for the seven drift fields are shown in Fig. 6.15. Two sets of plots are shown for $S1c < 120$ phd and $S1c > 120$ phd, respectively. We see that the S1c range varies dramatically, from a maximum of 50 phd for the highest drift field to 350 phd for the lowest drift field. This is a consequence of both the decreased S1 signals from high-field interactions and the S2 cutoff we described above.

The behavior of the ER band vs. field is identical to what we saw in LUX.

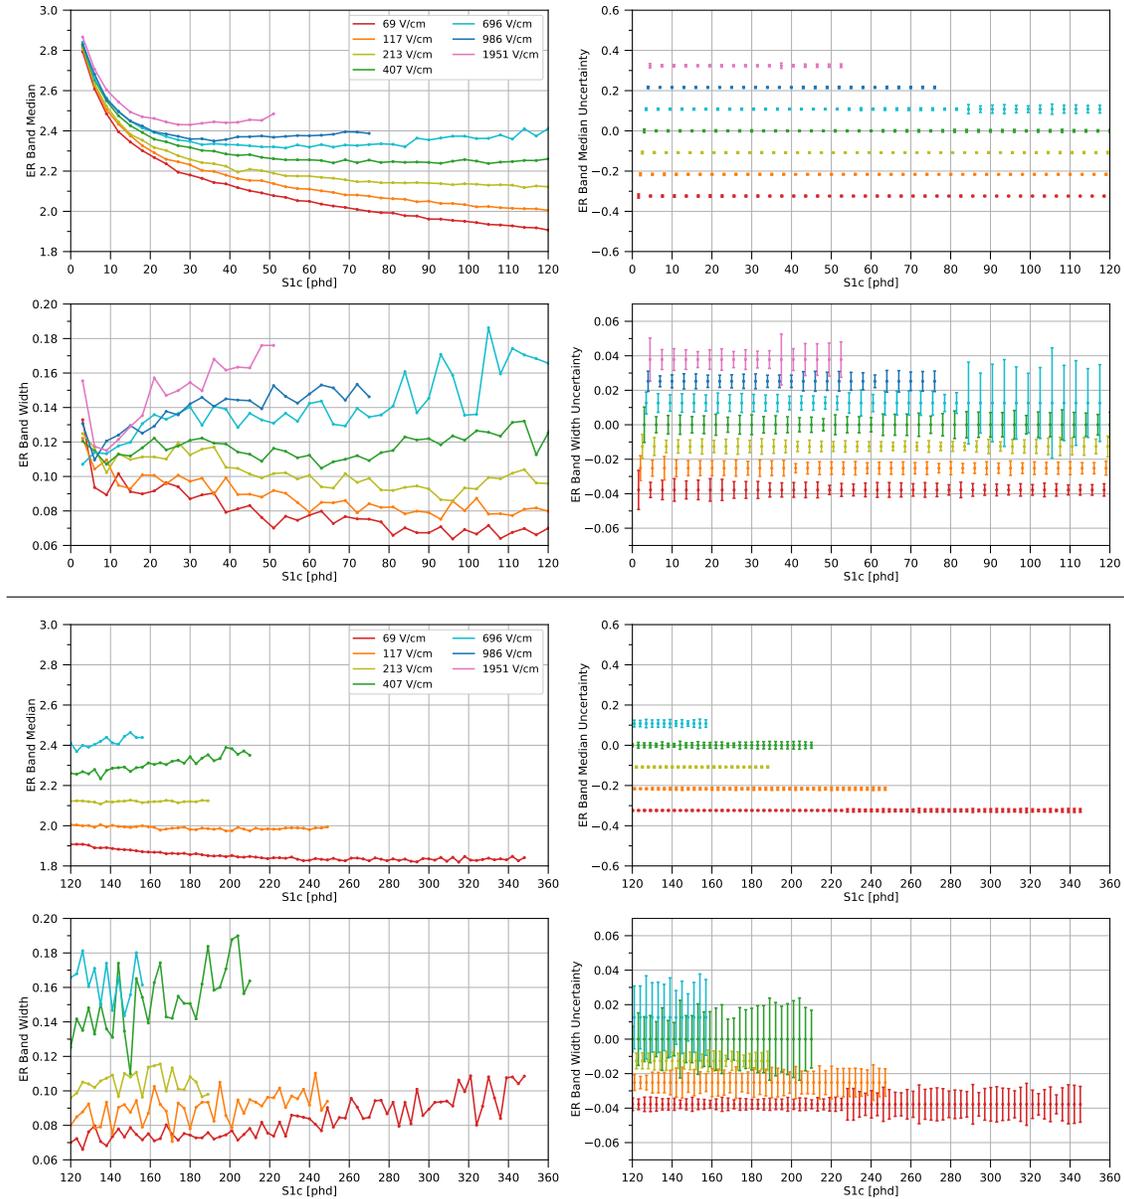


Figure 6.15: The median and width (Gaussian μ and σ) of the ER band for several drift fields. The left plots show measurements, and the right plots display error bars corresponding to these measurements. In each row, the y -axes have the same range; the size of the error bars on the right plot can be directly translated to the points on the left plot. For ease of visualization on the right plots, the S1 values are slightly shifted relative to their true value, and the error bars are centered at a different y -value for each field bin. The upper and lower 2x2 grids represent different ranges in S1c.

The band median and width both shift upward with drift field. In addition, the Thomas-Imel vs. Doke-Birks transition region, identified by the local minimum of $\log_{10}(S2c/S1c)$, shifts to lower S1c for increased drift fields.

6.6 Leakage and Discrimination

Finally, we compute the ER leakage fraction, using Gaussian extrapolation of the ER band below the NR band median. The results are shown in Fig. 6.16, in bins of S1c. The immediate reaction is that the measured leakage is extraordinarily high. For S1c below 40 phd, we see leakages of more than 1%! Our expectation would be an order of magnitude lower, somewhere between 0.1–1%. We observe that the lowest-S1c point for 1951 V/cm is even higher than 10%. This is a consequence of the strange upward turn in the NR band for that field, which we discussed previously.

The effect is also visible (perhaps more strikingly) in Fig. 6.17, where we show the total leakage between 1.5–52.5 phd. The integrated leakage is over 1% for most fields, except for 407 and 696 V/cm, where it is just under 1%. We emphasize that this is an incredibly high value. A dark matter experiment with 1% ER/NR discrimination would not be able to perform a competitive search. Fortunately, PIXeY is not a dark matter search, so we can investigate further.

The reason for this high leakage is not understood. However, one way to dig into the puzzle is to compare our results to LUX measurements and NEST simulations. The LUX WS2014–16 data are at roughly the same g_1 (for LUX, $g_1 = 0.087$), and the four lowest PIXeY fields are represented in the LUX data. As for NEST, we built a simulation file with the appropriate parameters for the PIXeY detector: g_1 , g_2 , drift field, extraction field, single electron size, and PMT resolution and threshold. The total integrated leakage from LUX and NEST are compared to PIXeY in Fig. 6.17, confirming the substantial discrepancy. We make two remarks here. First, recall that the LUX leakage was calculated with both skew-Gaussian and Gaussian extrapolations (see Fig. 4.27 in Chapter 4). We show both, but the proper comparison is the much higher Gaussian leakage, since that is the PIXeY procedure. Second, the LUX and NEST leakages are over a range of 1–80 phd. Our PIXeY analysis only goes up to 52.5 phd, and we are unable to extend it to higher S1c for high drift fields. We know that leakage decreases with higher energy, so this could naively explain the discrepancy.

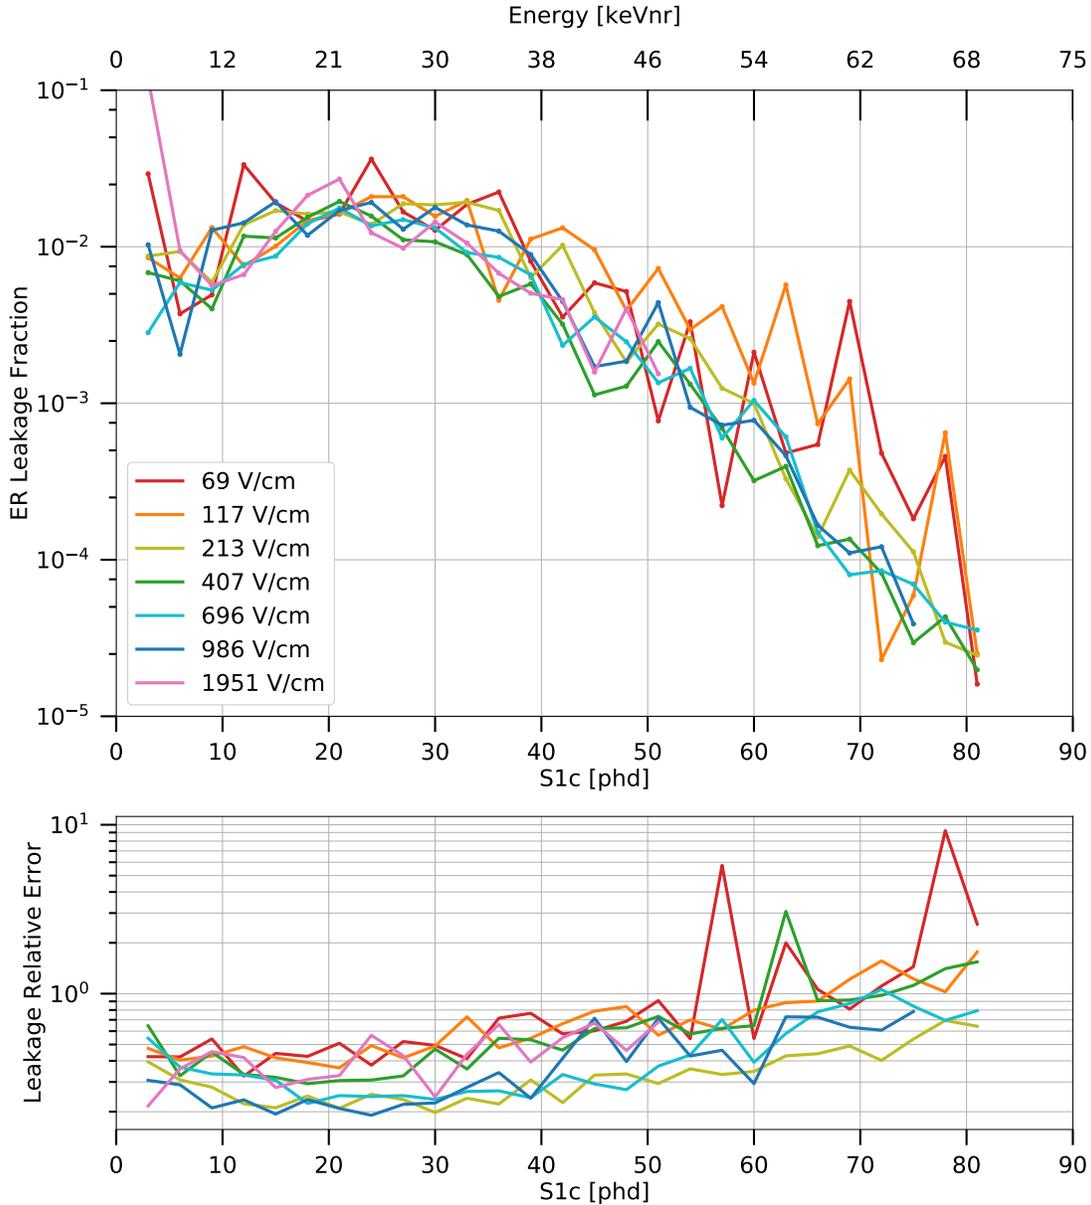


Figure 6.16: (Top) The ER leakage fraction for a flat energy spectrum in S1 bins, for various values of drift field, calculated from a Gaussian extrapolation of the ER band below the NR band median. The equivalent nuclear recoil energy for an S1 is calculated by using the S1c and S2c at the median of the NR band; this varies by field, but not significantly, so we report the energy averaged over the eight field bins. (Bottom) The relative error on these leakage fraction values, defined as in Fig. 4.22.

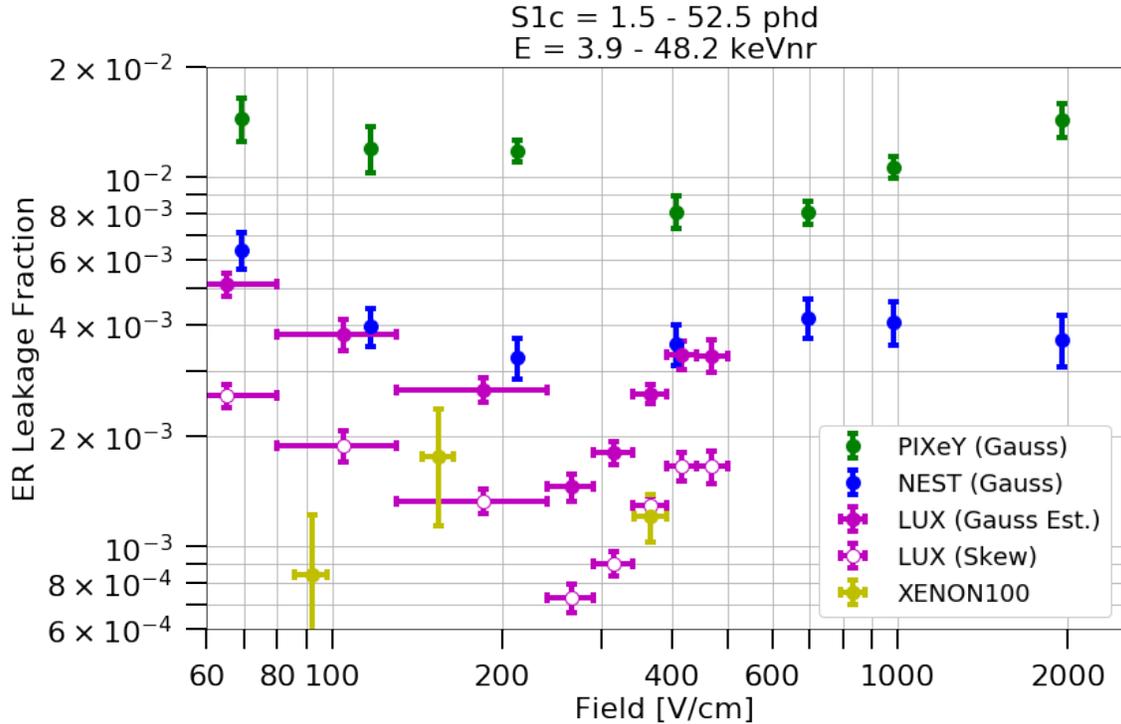


Figure 6.17: The integrated electronic recoil leakage for a flat recoil energy spectrum from $S1c = 1.5\text{--}52.5$ phd (equivalent to $3.9\text{--}48.2$ keVnr), while varying drift field in PIXeY data (green). Results from LUX WS2014–16 are shown in magenta, where we show both the skew-Gaussian and Gaussian extrapolations in open and closed circles. Results from NEST simulations are shown in blue. For both LUX and NEST, leakage is measured in the range $S1c = 1\text{--}80$ phd, equivalent to $2\text{--}65$ keVnr. Results from XENON100 [144] are shown in yellow. We use XENON100 leakages at $g_1 = 0.081$, corresponding to $11\text{--}34$ keVnr.

However, the $S1c$ range is not the reason. Figure 6.18 shows a comparison of PIXeY measurements to LUX and NEST for a 117 V/cm drift field, in $S1c$ bins. The ER and NR band medians, the difference between the two, the ER band width, and the ER leakage fraction are shown. It is representative of the trends for all seven drift fields. We see that the PIXeY leakage is much higher than the NEST simulation and the LUX measurements, even when split up by $S1c$. We also observe that the ER band width is the same as PIXeY and LUX would suggest. The crucial cause of our worse leakage is that the ER and NR bands are closer together. In PIXeY, the ER and NR bands are both lower than LUX and NEST (probably because of a mis-modeled g_1 or g_2), but the ER band is substantially more deviant. As a result, the difference between the band medians is smaller in PIXeY, and the leakage is higher.

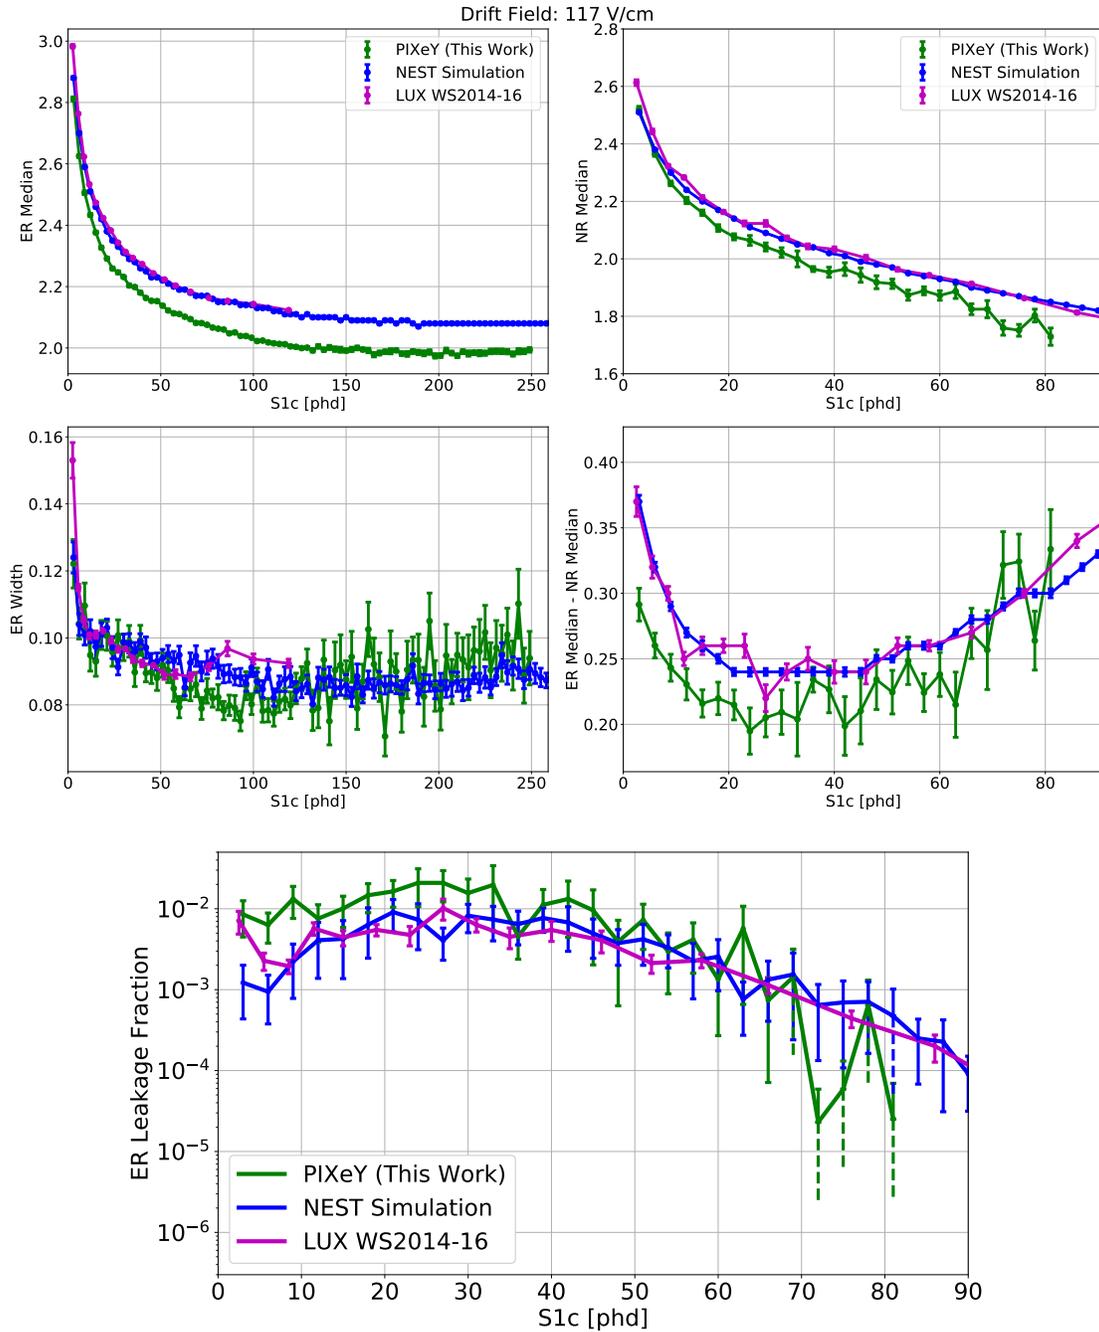


Figure 6.18: Comparison of the ER band median, NR band median, difference in the ER and NR band medians, ER band width, and ER Gaussian-extrapolated leakage fraction between PIXeY measurements at 117 V/cm drift field, a NEST simulation of the PIXeY detector, and LUX measurements in the field range 80–130 V/cm.

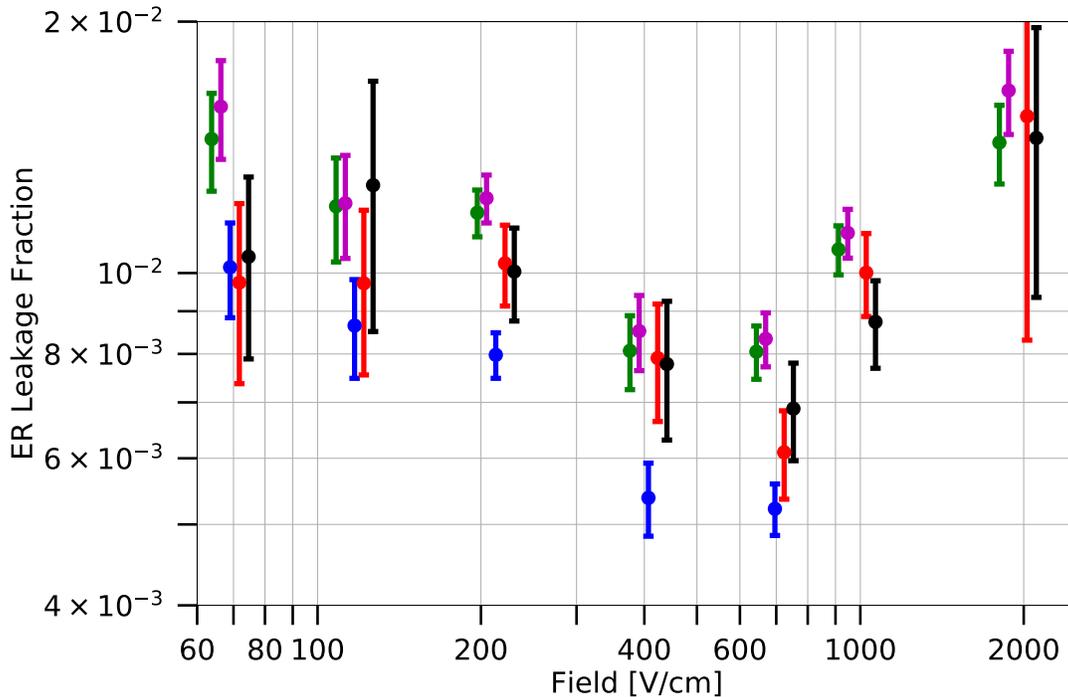


Figure 6.19: PIXEY measurements of the integrated ER leakage fraction, for several energy ranges and fiducial cuts.

Green: $r < 38$ mm, $1.5 < (S1c/phd) < 52.5$, $3.9 < (E/keVnr) < 48.2$;

Magenta: $r < 38$ mm, $4.5 < (S1c/phd) < 52.5$, $6.6 < (E/keVnr) < 48.2$;

Blue: $r < 38$ mm, $1.5 < (S1c/phd) < 79.5$, $3.9 < (E/keVnr) < 67.9$;

Red: $r < 34$ mm, $1.5 < (S1c/phd) < 52.5$, $3.9 < (E/keVnr) < 48.2$;

Black: $r < 32$ mm, $1.5 < (S1c/phd) < 52.5$, $3.9 < (E/keVnr) < 48.2$.

Although we have a technical explanation, we still do not understand the fundamental cause of the poor discrimination. Backgrounds in the calibration data are the most likely culprit, but if there was a background distribution on top of the ER or NR data, the histograms of $\log_{10}(S2c/S1c)$ would not look Gaussian at all; they would be bimodal. However, we observe that the histograms are quite Gaussian.

Our ignorance notwithstanding, the trends we see in band medians and widths vs. field, as well as the minimal discrepancy between different acquisitions of the same field, are quite robust. Thus, we expect that our measured *trends* of leakage vs. field should be meaningful. We showed the integrated leakage in Fig. 6.17, but we show it again in Fig. 6.19, using a few fiducial cuts and energy ranges. This is to confirm that our choice of these parameters is not inappropriately altering the results. Indeed, the leakage doesn't change much, based on our choices.

Thus, we can arrive at some final conclusions for our analysis. The leakage varies

as a function of drift field, and it is minimized at moderate fields, just like LUX saw. However, the leakage-minimizing field appears to be higher in PIXeY data, between 400 and 700 V/cm. Given the larger statistical uncertainties in the PIXeY analysis, though, this is not as firm a statement as it was in LUX. Probably the more important finding is the increase in leakage at the highest fields. If a 100-tonne xenon TPC is built, one of the design questions will be its dimensions. In general, it is better to build a “right cylinder”—a detector where the height equals the diameter—to maximize self-shielding. This is the structure of LUX and LZ. However, if the leakage went down substantially at higher fields, it might be worth building a short, wide detector. That way, experimentalists could achieve a much higher drift field with a given cathode voltage. Our findings show that this scheme will likely not be advantageous, and the right cylinder is the best strategy.

Chapter 7

Calibration of Low-Temperature Sensors for the TESSERACT Project

We now dramatically switch our focus to the search for low-mass dark matter: particles with masses $\mathcal{O}(1 \text{ MeV}/c^2\text{--}10 \text{ GeV}/c^2)$.

Such particles cannot be WIMPs in the strictest sense because of the Lee-Weinberg limit [19], but we consider models where their interactions with the standard model are similar to WIMPs. More precisely, we consider sub-GeV galactic halo particles that exhibit spin-independent interactions of the form $\frac{d\sigma}{dE_R} \propto v^{-2}$. These are the “MeV-scale dark matter” models described in Chapter 2. We can keep all of the experimental strategies we have discussed to identify nuclear recoils, but the target material needs to be changed.

Xenon is unable to detect the sub-keV energy depositions associated with MeV-scale dark matter, because these depositions do not create enough photons to generate an S1 signal. Recall that we found a 50% efficiency threshold of 3.5 keVnr from the LUX D-D calibration (Chapter 4), which required events to have both an S1 and S2 pulse. There are strategies to search for low-mass dark matter using only the S2 pulse [41, 160] or using ER signals [126, 39], but these suffer from high background rates.

On the other hand, we can try an entirely different strategy: calorimetry. In a calorimeter, one can measure the total energy deposited in a material through its corresponding temperature increase. Specifically, in this dissertation, we will discuss cryogenic transition-edge sensors (TES), which can detect photon and phonon absorptions. A TES is a superconducting material cooled down to its transition temperature and kept inside its transition. Energy deposits create temperature changes, which become resistance changes and are thus electrically measurable. For example, the most-common scheme is a voltage-biased TES, in which the resistance change is transformed into a current pulse, which is in turn read out through a DC SQUID. The general electrothermal principles of a TES are shown in Fig. 7.1. We will discuss these in much more detail in Sec. 7.3.

This Chapter presents work done to characterize materials for their use as transition-edge sensors for TESSERACT. We begin by describing the TESSERACT project, which is implemented by the SPICE/HeRALD collaboration and aims to build TES-based dark matter detectors with three target materials. Next, we describe the dilution refrigerator in which our measurements were performed, located at Lawrence Berkeley National Laboratory (LBNL). We then give a theoretical overview of TES dynamics and present an analysis of 46-mK devices, in which we measure their bias power and thermal conductance. Finally, we present measurements of the superconductivity of various materials, allowing us to gauge their potential use as TESs. In this section, we also explore the relationship between superconducting transition temperature and film stress.

Before we begin, a few notes on proper attribution of work. All of the analysis in this Chapter was done by the dissertation author. However, the measurements were taken in an experimental setup shared by the author, LBNL staff scientist Aritoki Suzuki, LBNL postdoc Xinran Li, and LBNL postdoc Kaja Rotermund. As a result,

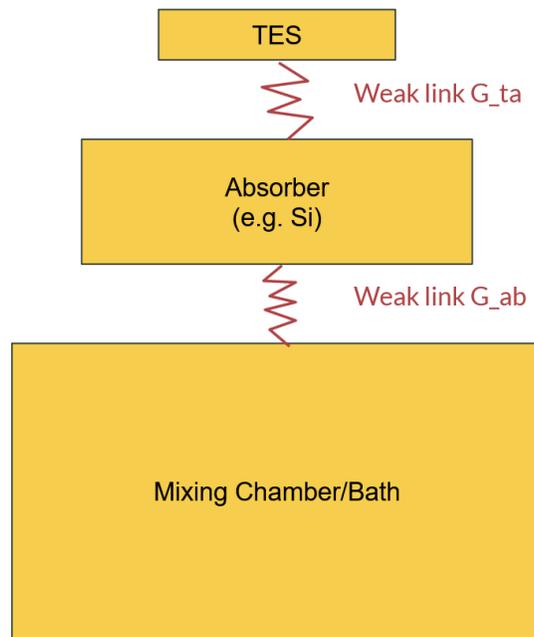
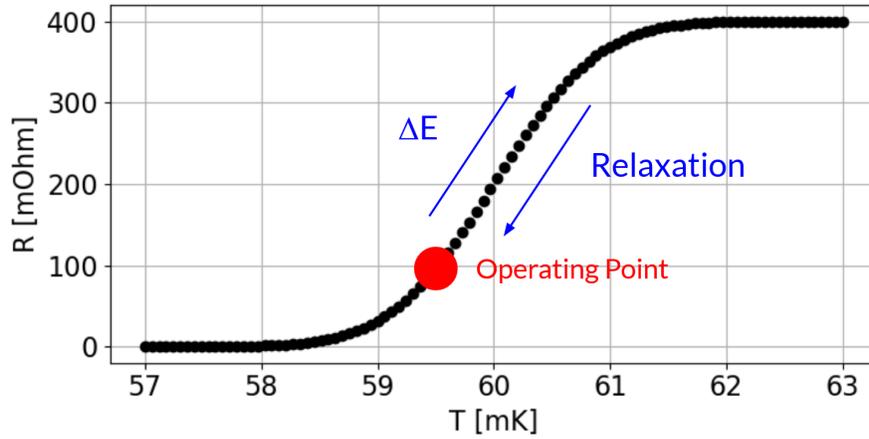


Figure 7.1: A high-level schematic of how a TES works. (*Top*) An example of resistance vs. temperature for some material with a superconducting transition (not real data). The material is kept inside the transition at the Operating Point. If the temperature increases from an energy deposit, the resistance will increase. If the TES is voltage-biased (the common scheme), the Joule power V^2/R will decrease, bringing the temperature back to the Operating point. (*Bottom*) A thermal diagram of the system. The TES is thermally linked to an absorber, which is linked to a thermal bath, e.g., the mixing chamber of a dilution refrigerator. In some situations, the TES itself is the absorber, but this is uncommon in DM experiments.

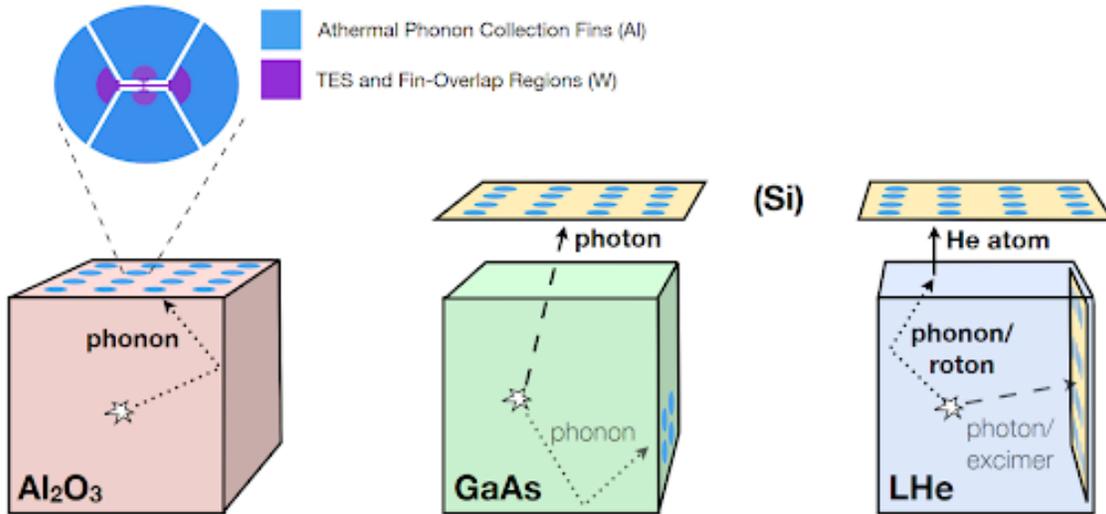


Figure 7.2: A high-level schematic of the TESSERACT strategy. When energy is deposited in three target materials, they produce light and heat signals, which are detected by transition edge sensors. The detectors and sensors are at a temperature of tens of mK. Figure copied from [161].

setting up the experiment and taking measurements were done by a combination of all four scientists. In addition, the device designs and fabrications were done by the Pyle group at UC Berkeley and the Mahapatra group at Texas A&M. The measurements of film stress were done solely by the author.

7.1 SPICE, HeRALD, and TESSERACT

TESSERACT is the “Transition Edge Sensors with Sub-eV Resolution And Cryogenic Targets” project. Its aim is to detect sub-WIMP-mass particle dark matter using three target materials: polar sapphire (Al_2O_3) crystals, gallium arsenide (GaAs) crystals, and superfluid helium. The detectors using sapphire and GaAs are called SPICE (Sub-eV Polar Interactions Cryogenic Experiment), and the helium experiment is called HeRALD (Helium Roton Apparatus for the Exploration of Light Dark Matter). Accordingly, the collaboration of scientists (some 20–30 of us) that runs the project is named the “SPICE/HeRALD Collaboration.” Although the target materials are different, the idea is that TESSERACT serves as a unified framework for delivering calorimeters, detectors, shielding and cryogenics.

Figure 7.2 shows a cartoon of the TESSERACT strategy. Each detector is made of $\mathcal{O}(0.1\text{--}1\text{ kg})$ of the corresponding target material. One advantage of TESSERACT is that we can “stage” our detectors. In other words, we can deploy larger targets with relative ease, so that we can use early runs to understand our technology

and resolve obstacles, and late runs to focus on physics searches. Each interaction deposits energy in three channels, just as we discussed in Chapter 3: light, heat, and charge. The energy partitioning is not the same as xenon¹, but the three channels are the same. The SPICE and HeRALD experiments will not employ electric fields, so we will not be able to read out the ionization signal. At zero field, recombination should be nearly complete, so our loss of observable energy should be minimal. More importantly, the lack of electric fields allows us to lose the backgrounds that electric fields cause through spurious electrons and dark counts. The heat signal manifests as phonons. For the GaAs and sapphire detectors, these phonons are collected by athermal aluminum fins on the surface of the crystal and then thermalize into small transition-edge sensors. The TES temperature increase is sensed as an electrical signal, as we will discuss momentarily. Meanwhile, in superfluid helium, phonons propagate ballistically through the volume and reflect off the surfaces until eventually, they hit the helium-vacuum interface. When this occurs, they have some probability of evaporating a single helium atom off the surface, which is adsorbed onto an upper detector (e.g., silicon). The corresponding energy deposition is measured by TESs. Finally, the GaAs and helium detectors will also be able to measure energy in the excitation channel. Both can measure the photons generated from de-excitation, while the helium detector can additionally measure long-lived ($\tau_{1/2} = 13$ s) excimers. This energy is measured by the same silicon substrates, collection fins, and TESs.

The sensitivity of the TESSERACT experiments to several dark matter interaction models is shown in Fig. 7.3. We see that the three materials are complementary across all four interaction types—nuclear recoils, electronic recoils via a heavy mediator, electronic recoils via an ultralight mediator, and absorption. Some materials are better at probing certain interactions than others (e.g., helium and sapphire are best at searching for nuclear recoils), but in combination, the TESSERACT project has the potential to be world-leading for all the interactions.

In pursuit of this agenda, the author undertook several experiments to determine the best materials to serve as TESs. We have three goals: to achieve the lowest possible superconducting transition temperature (which we note T_c), the lowest film stress, and reliability in our fabrication procedure. Low T_c is important to reduce noise, therefore increasing energy resolution and energy threshold, which allows us to expand sensitivity in both the low-mass and low-cross-section directions. Low T_c also leads to a lower heat capacity and a stronger thermal response for a given energy deposition. Meanwhile, low film stress is crucial to reduce microfracture backgrounds, in which cracks in the absorber material lead to sudden releases of energy through athermal phonons. Such a release is associated with the same pulse shape as a true energy deposition. This background has been observed by CRESST for W

¹Helium and xenon are somewhat similar in their energy deposition behaviors, though, both being noble elements.

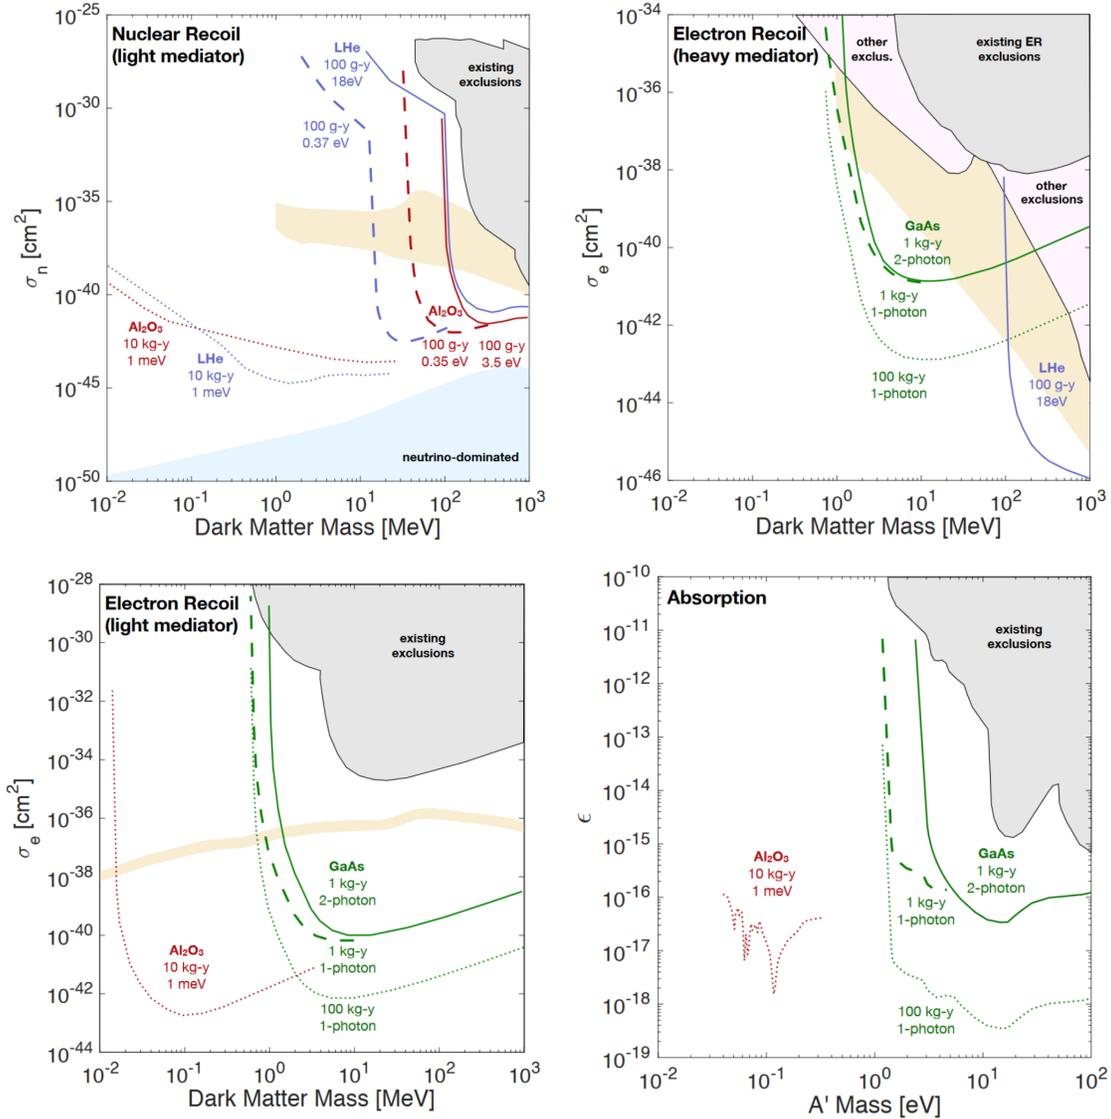


Figure 7.3: Sensitivity of the three TESSERACT constituent experiments to dark matter interactions. We show spin-independent nuclear recoils, electron recoils via heavy and light mediators, and dark matter absorption. The gray regions are excluded parameter space, and the orange regions show models of interest. The solid lines, dashed lines, and dotted lines show the predictions for shovel-ready technology, the expected baseline technology, and a “blue-sky” calculation for the best possible sensitivity. Figure copied from [161].

films on sapphire crystals [162]. Finally, the need for reliability is self-explanatory. These experiments are described in the remaining portion of the chapter, starting with the setup of the apparatus.

7.2 Experimental Apparatus: Dilution Refrigerator at LBNL

All measurements were done in a BlueFors LD-400 cryogen-free dilution refrigerator located at Lawrence Berkeley National Laboratory. The refrigerator is able to cool samples down to <7 mK; it provides a cooling power of $16 \mu\text{W}$ at 20 mK and $500 \mu\text{W}$ at 100 mK. The refrigerator is located inside a $10 \times 10 \times 9$ ft³ EMI shield room that provides 100 dB attenuation to electromagnetic waves in the frequency range 14 kHz to 10 GHz. Power is provided to the shield room only through three noise-reducing filter banks. All other signals entering the shield room (such as a computer monitor) are provided through fiber optic cables. The refrigerator is installed on an active vibration suppression mount, although active vibration was not used for the measurements in this dissertation.

We will describe the operation of a dilution refrigerator briefly; for a more detailed description, readers are directed to copious amounts of literature online, e.g. [163, 164]. First, the cryostat is evacuated of air, pumping down to a pressure of $\approx 10^{-6}$ mbar. Then, the cryostat is cooled down to about 4 K. This can be done either through a liquid helium jacket (a “wet” refrigerator) or through a pulse tube cooler (a “dry” or “cryogen-free” refrigerator). Dry refrigerators are much easier to operate, but the pulse tube adds a substantial amount of vibrations, which can create noise in the sensors. Our apparatus is a dry refrigerator².

The final and most critical step is to use the dilution cycle, which takes advantage of the heat of mixing of ³He and ⁴He to cool the refrigerator to several mK. This process is shown in the phase diagram in Fig. 7.4. Pure ⁴He transitions to a superfluid state when cooled below 2.17 K. A mixture of ³He/⁴He may similarly transition if the ³He concentration is low enough, but the transition temperature decreases with increasing concentration. Furthermore, when cooled below a certain temperature (typically about 0.8 K, depending on concentration), the mixture separates into a ³He-rich phase and a ³He-poor phase. The “concentrated” phase contains close to 100% ³He, whereas the “dilute” phase contains at least 6.4% ³He. The dilute phase is denser than the concentrated phase, so it sinks to the bottom. This is done in the aptly-named “mixing chamber” of a dilution refrigerator. Additionally, ³He has higher enthalpy in the dilute phase than in the concentrated phase. Therefore, to

²The refrigerator has the capability to stay cold for 1–2 hours without the pulse tube by relying on a small vessel of condensed liquid helium, but we don’t use this capability for the measurements in this dissertation.

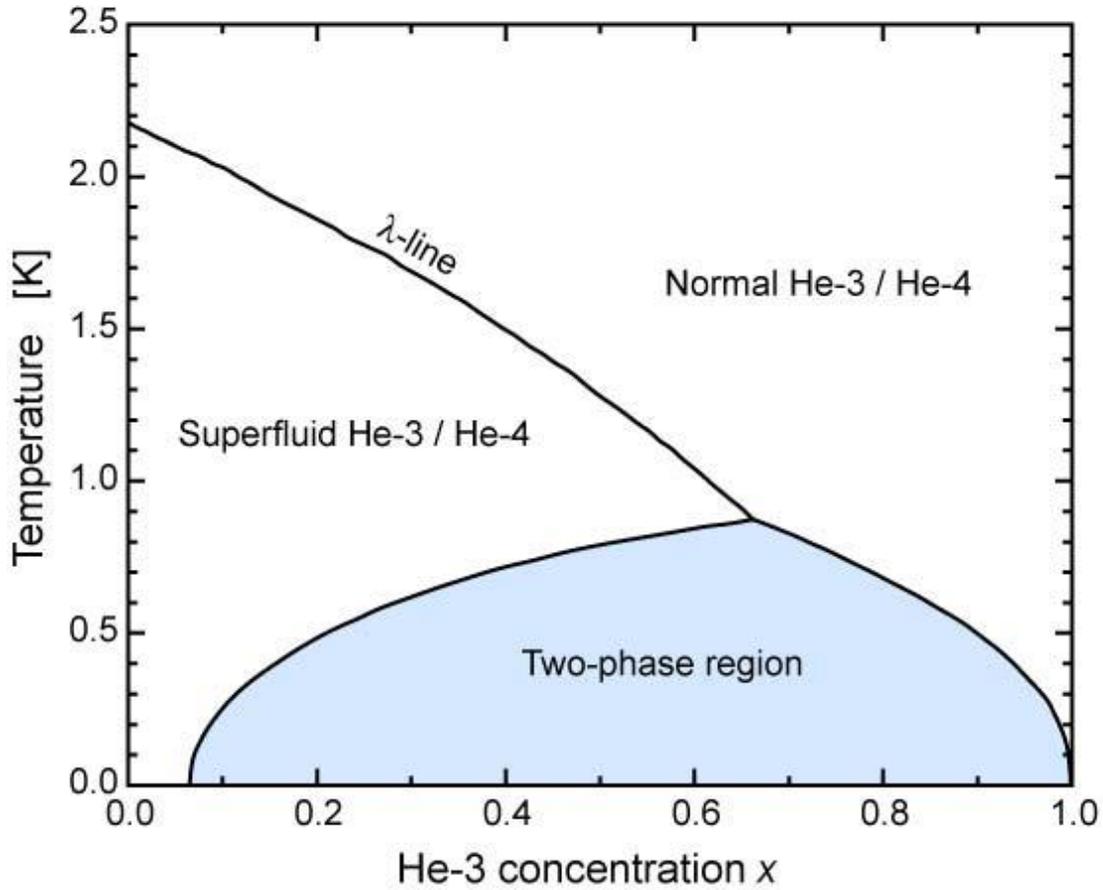


Figure 7.4: The phase diagram of a mixture of ^3H and ^4H . The “two-phase region” is thermodynamically forbidden. Figure copied from [165].

keep the refrigerator cold, we push ^3He into the upper concentrated phase, pull it out of the lower dilute phase. and connect these in a loop. As ^3He migrates across the boundary, it requires energy, which comes from the refrigerator and cools down the system. Anything in physical contact with the mixing chamber will also be cooled. In our case, the mixing chamber is connected to a flange that experiments can be mounted on, with some 900 cm^2 of space. The cooling power of a dilution refrigerator is given by

$$\dot{Q} = \dot{n}_3 (95T_{mc}^2 - 11T_{in}^2) \approx 84 \dot{n}_3 T_{mc}^2 \quad (7.1)$$

where T_{mc} is the mixing chamber temperature, T_{in} is the input temperature of the ^3He , and \dot{n}_3 is the ^3He flow rate. In the second half of Eq. 7.1, we approximate $T_{mc} \approx T_{in}$, which is usually valid. Typically, we operate at $\dot{n}_3 \approx 0.7\text{ mmol/s}$. The

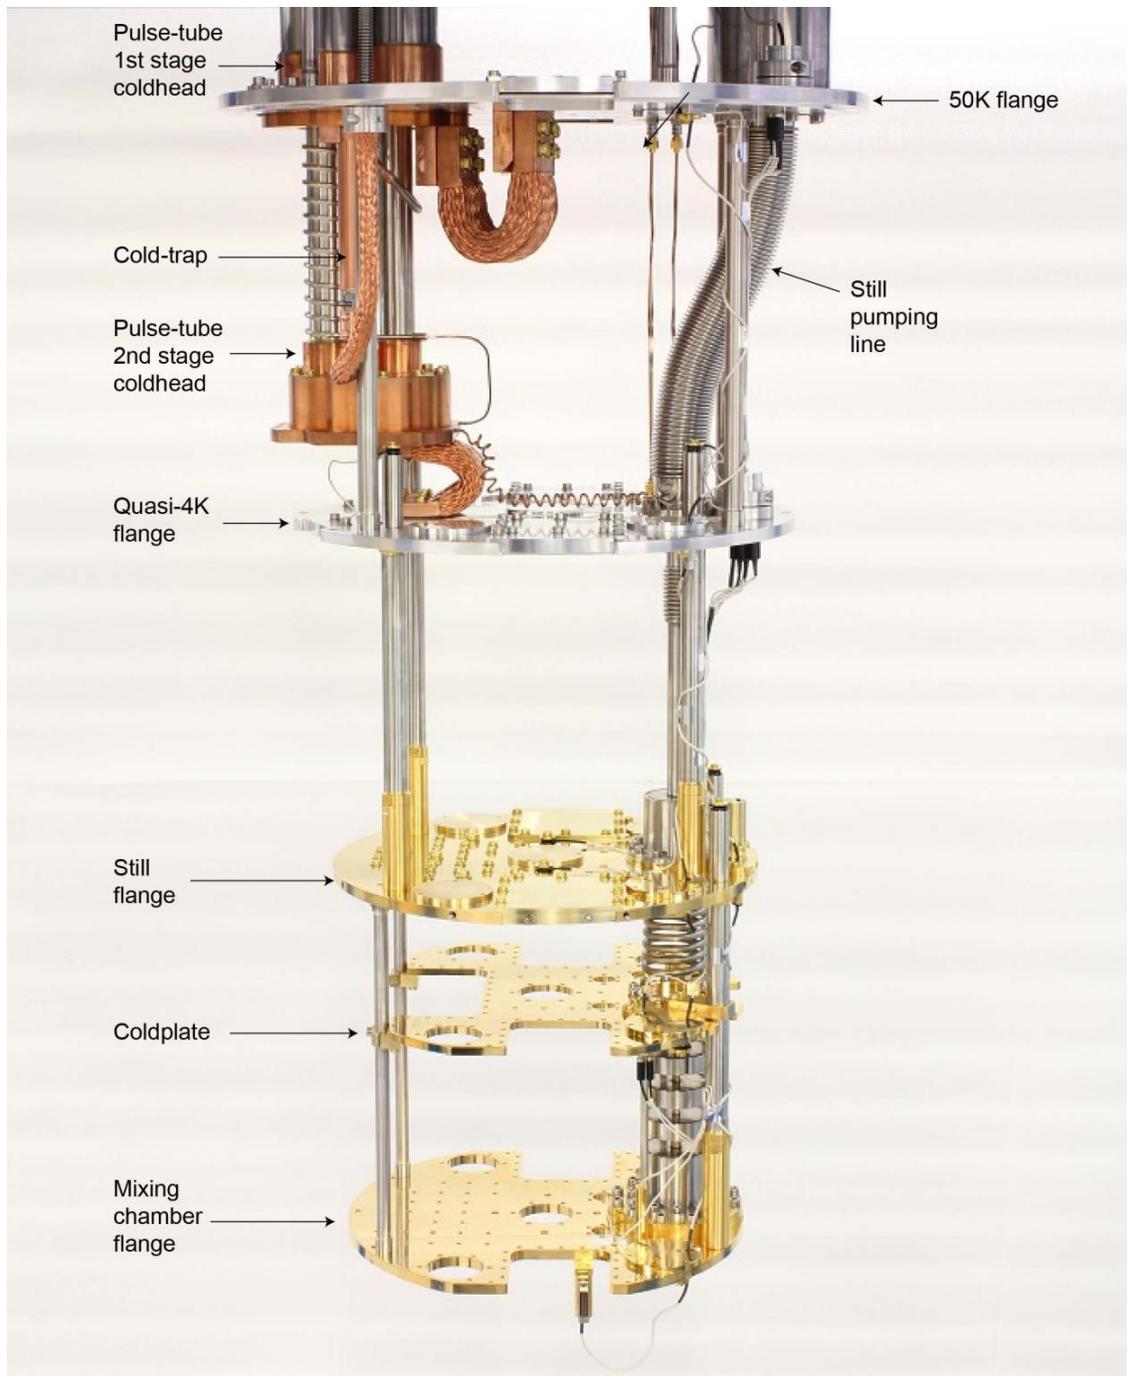


Figure 7.5: Generic photo of the BlueFors LD-400 refrigerator. Figure copied from [165].

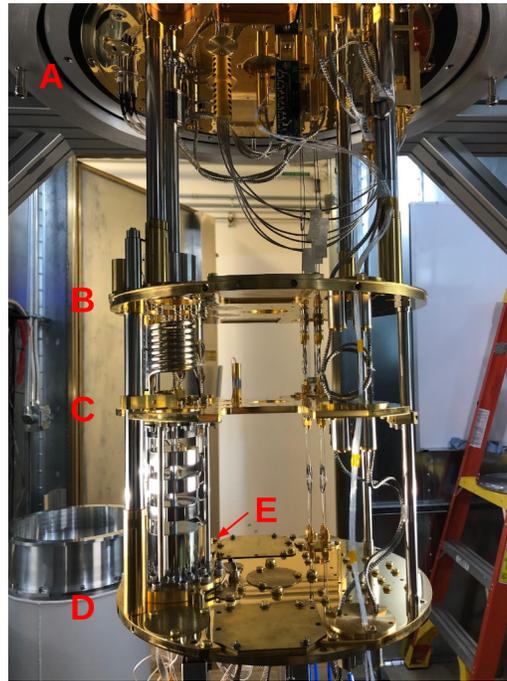


Figure 7.6: Photos of the LBNL refrigerator. Top is open, back view; bottom is closed, front view. **A**: 4K Flange, **B**: Still Flange, **C**: Cold Plate Flange, **D**: Mixing Chamber Flange, **E**: Mixing Chamber, **F**: Closed Fridge, **G**: Anti-Vibration Mount, **H**: Room-Temp Inputs, **I**: Filtered Power Banks.

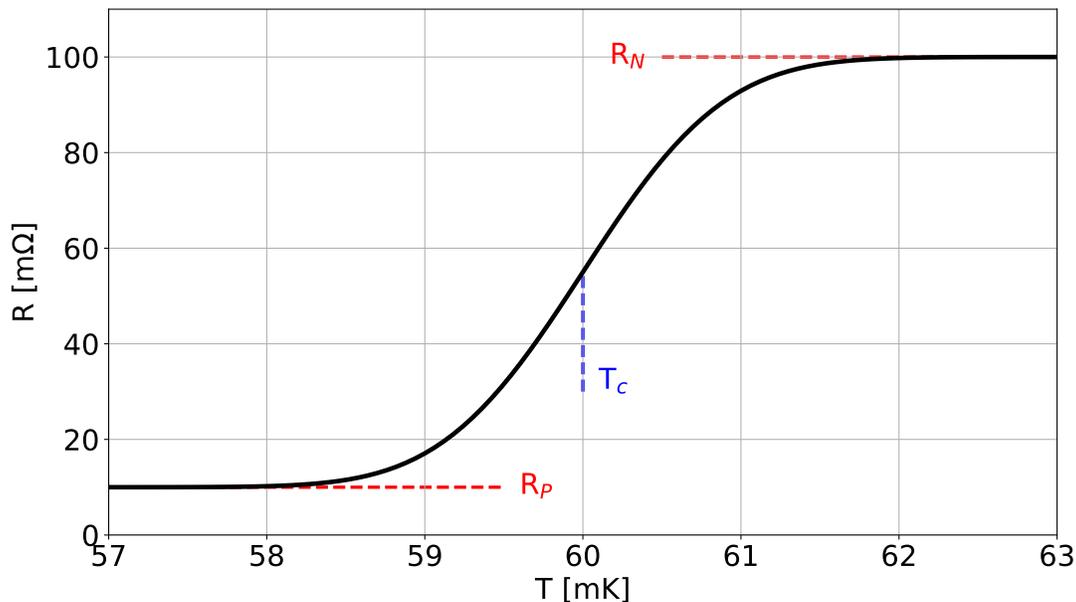


Figure 7.7: An example of resistance vs. temperature for some material with a superconducting transition at temperature T_c (not real data). The normal resistance R_N and parasitic resistance R_P are highlighted. The effect is exaggerated for visualization; typically, the ratio R_P/R_N is much smaller.

lowest T_{mc} we can achieve is <7 mK; we can go as high as several hundreds of mK, but we usually do not go above 100 mK for dark matter sensor measurements.

Figures 7.5 and 7.6 show photos of a generic BlueFors LD-400 refrigerator and our refrigerator at LBNL, respectively. We will discuss the specific experimental setups for our TES and T_c measurements in the relevant sections.

7.3 Introduction to TES Dynamics

Before reporting our TES measurements, it is necessary to discuss the parameters that characterize a transition-edge sensor, and the dynamics that drive their operation. We primarily follow the introductory text by Irwin and Hilton [166], summarizing their key results. In general, we don't derive equations; the interested reader can read Irwin and Hilton.

The TES resistance vs. temperature relationship is shown in Fig. 7.7. This is similar to the top panel of Fig. 7.1, but we highlight a few different characteristics. At temperatures significantly above the transition temperature T_c , the resistance is a constant value R_N , which we call the “normal resistance.”³ At temperatures

³Technically, the normal resistance varies slightly with temperature, just like all materials, but

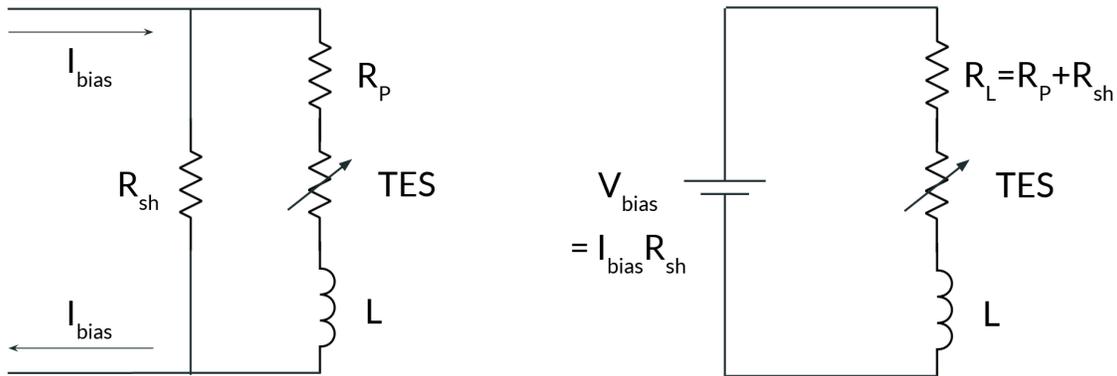


Figure 7.8: (*Left*) The current-biased TES circuit: a shunt resistor in parallel with the combination of a TES, parasitic resistance, and an inductor. (*Right*) The Thevenin equivalent of the same circuit: a voltage bias applied to a load resistance (equal to the shunt plus parasitic resistance), the TES, and the inductor.

significantly below T_c , the resistance is not actually zero; it is typically measured at some trace “parasitic resistance” R_p . The parasitic resistance does not necessarily come from the TES itself; it can come from the wiring or surface connections. A given experiment can live anywhere on the R-T curve. We denote the steady-state values as R_0 and T_0 , and the values at a given time $R(t)$ and $T(t)$.

The generic TES circuit is shown in Fig. 7.8. As discussed earlier, we can achieve negative electrothermal feedback by voltage-biasing the TES. Typically, this is done by current-biasing a TES-shunt pair in parallel, as shown in the left panel. However, the Thevenin circuit, shown at right, indicates that this is equivalent to a voltage-biased TES with bias $V_{bias} = I_{bias} R_{sh}$. We consider the TES and parasitic resistance to be separate resistors in series, so that the TES resistance is exactly zero at low temperatures. The load resistance in the Thevenin circuit is $R_L = R_p + R_{sh}$. When the TES temperature increases from an energy deposition, the current through the TES changes. This, in turn, produces a changing magnetic flux in an inductor L , which can be read out by a sensor that is disconnected from the TES circuit. Flux readout is most commonly done with a Superconducting QUantum Interference Device, or SQUID. We do not describe the operation of a SQUID; it is sufficient for the reader to know that a SQUID can measure small changes in magnetic flux. However, the interested reader can avail themselves of many resources, such as Fagaly’s 2006 review [167]. There are two categories of SQUIDs, RF and DC; in this dissertation, we only discuss DC SQUIDs.

During operation, TES dynamics are governed by two differential equations for the electrical and thermal responses. The electrical equation is shown below, where I and R are the instantaneous TES current and TES resistance, respectively:

this variation is typically visible at the K scale, rather than the mK scale.

$$L \frac{dI}{dt} = V_{bias} - I[R_L + R(T, I)] \quad (7.2)$$

Meanwhile, the thermal equation is below, where T is the instantaneous TES current and C is the combined TES-absorber heat capacity:

$$\begin{aligned} C \frac{dT}{dt} &= -P_{bath} + P_J + P_{steady} \\ &= -K(T^n - T_{bath}^n) + I^2 R(T, I) + P_{steady} \end{aligned} \quad (7.3)$$

Here, we have modeled the power flowing from the TES to the bath as a power-law, where the exponent n depends on the details of the surface boundary. The conductance is $G \equiv \frac{dP_{bath}}{dT} = nKT^{n-1}$. The Joule power dissipated in the TES is the normal I^2R . We have added a term P_{steady} , which denotes a steady-state signal power. This is useful for TESs operated as bolometers—for example, to measure the constant power of the CMB—but it is not relevant for calorimetry, so we set $P_{steady} = 0$.

Note that the electrical and thermal equations are coupled via the Joule power heating and via the temperature-dependence of the TES resistance. This is known as *electrothermal feedback*. Furthermore, note that we write the TES resistance as $R(T, I)$, a function of both the temperature and current. In general, this is true. The resistance can vary with both. As a consequence, we define the temperature sensitivity α and the current sensitivity β :

$$\alpha = \frac{\partial \log R}{\partial \log T} = \frac{T_0}{R_0} \frac{\partial R}{\partial T} \quad (7.4)$$

$$\beta = \frac{\partial \log R}{\partial \log I} = \frac{I_0}{R_0} \frac{\partial R}{\partial I} \quad (7.5)$$

Experimenters typically try to produce sensors with $\beta \approx 0$, but some nonzero β is always present. That being said, $\beta = 0$ is often an acceptable first-order approximation, and it is an approximation we make in this dissertation.

The low-frequency open loop gain is defined below, with I_0 and α_0 noting the steady-state values of I and α :

$$\mathcal{L} \equiv \frac{I_0^2 R_0 \alpha_0}{G T_0} \quad (7.6)$$

For high gain, the fall time of a current pulse can be approximated as $\tau_- \approx \frac{C/G}{\mathcal{L}}$. We remark that C/G is simply the time scale associated with conductive thermal dissipation in the absence of any feedback (e.g., the time it takes an object in one's hand to reach one's body temperature).

Finally, we plot a few sample pulses in Fig. 7.9. We simulate energy deposits of 1–20 eV in a TES and show the electrical and thermal responses. Some reasonable

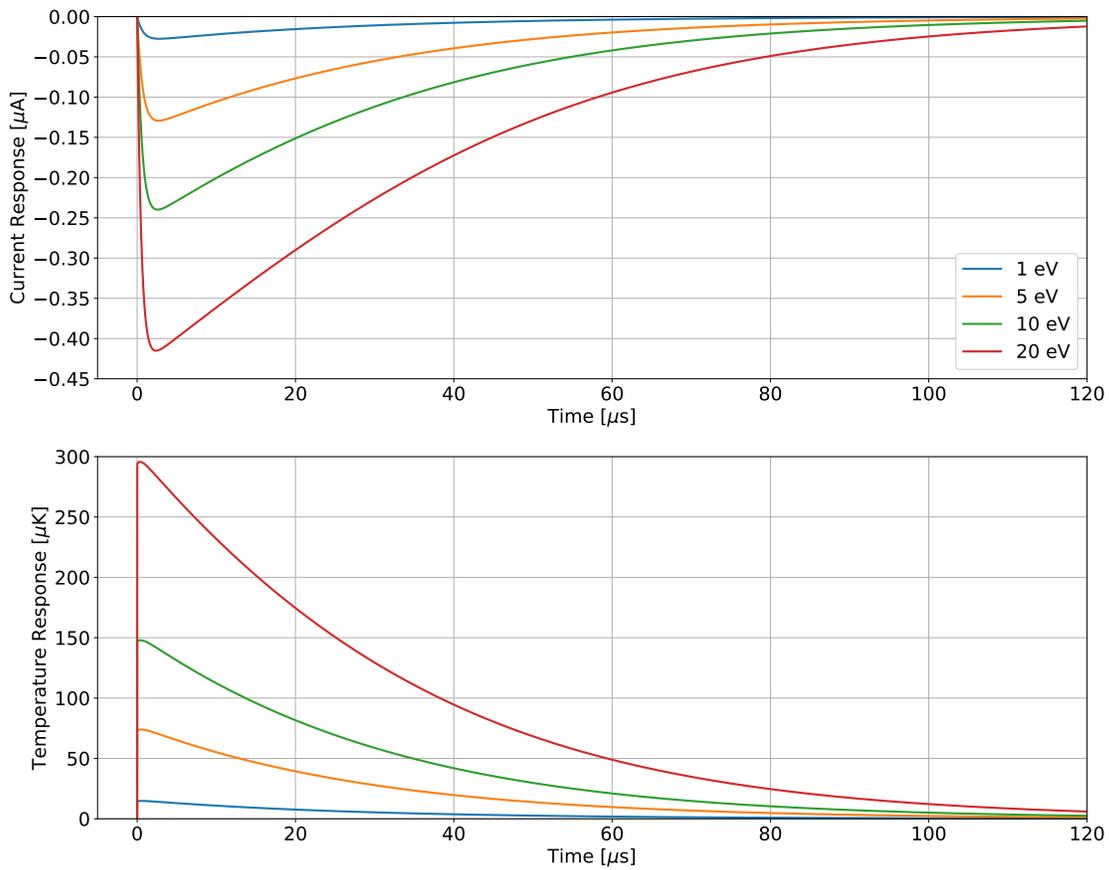


Figure 7.9: Simulated current and temperature pulses in a TES, for several energy depositions. The TES, absorber, and circuit parameters used are: $V_{bias} = 0.18 \mu\text{V}$, $I_0 = 1.3 \mu\text{A}$, $R_N = 400 \text{ m}\Omega$, $R_P = 1.5 \text{ m}\Omega$, $R_{sh} = 0.5 \text{ m}\Omega$, $L = 100 \text{ nH}$, $T_c = 60 \text{ mK}$, $G = 20 \text{ pW/K}$, $n = 5$, $C = 11 \text{ fJ/K}$. We assume the TES is operated at a temperature T_0 such that $R_0 = \frac{1}{3}R_N$, corresponding to $\alpha = 91$ and $\mathcal{L} = 18$.

parameters are chosen, listed in the figure caption. In these pulses, the rise time is a few μs , and the fall time is some tens of μs .

7.4 Measurement of 46 mK Transition Edge Sensors

In this section, measurements of two TESs are presented. They are both thin 40 nm tungsten films on silicon substrates with $T_c \approx 46 \text{ mK}$. They have rectangular areas with dimensions of $800 \mu\text{m} \times 200 \mu\text{m}$ and $400 \mu\text{m} \times 100 \mu\text{m}$, respectively. They were designed by the Pyle group at the University of California, Berkeley, and

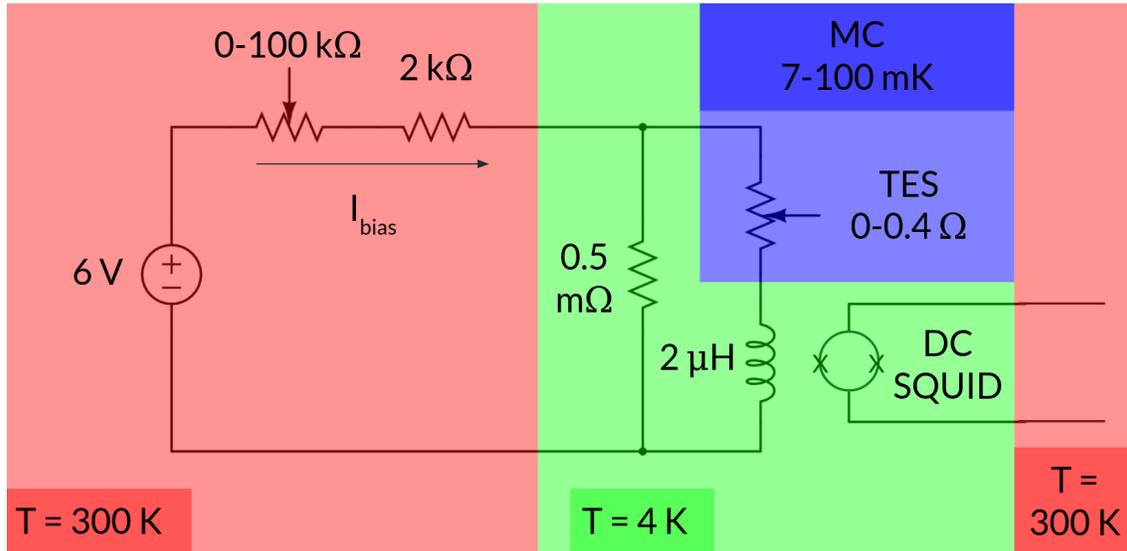


Figure 7.10: Circuit diagram used in our experiment. The shaded colors represent the temperature of each section of the setup. See also photos in Figs. 7.11, 7.12, and 7.13.

fabricated by the Mahapatra group at Texas A&M University. We measure their normal resistance, parasitic resistance, bias power as a function of temperature, and thermal conductance. Measurements were taken in August 2021.

7.4.1 Experimental Setup for TES Measurements

Our experimental setup used to characterize these TESs is shown in cartoon form in 7.10. Actual photos are depicted in Figs. 7.11, 7.12, and 7.13. We have six independent TES and SQUID channels, which are shown.

The room-temperature portions of the apparatus set up a current bias going into the refrigerator, using a 6 V battery, a 2 k Ω resistor, and a 0–100 k Ω 10-turn potentiometer. The current bias can be controlled with the potentiometer to values between 60 μ A and 3000 μ A. The 2 k Ω resistor is used to read the current bias; we can measure the voltage across the resistor using the BNC output in Fig. 7.11. Alternatively, we can use the same BNC connection to input a signal; for example, we may want to add a waveform on top of the DC current bias.

At the 4 K stage, the current bias is transformed into a voltage bias by means of a 0.5 m Ω shunt resistor. The voltage bias can be set between 0.03 μ V and 1.5 μ V, again depending on the potentiometer. We use Model 50 DC SQUIDS from Quantum Design, mounted at the 4 K stage. The DC SQUID and 2000 nH input inductor⁴

⁴Note that a 2000 nH inductance is far too high for a $R_N = 0.4\Omega$ dark matter sensor; the TES will suffer from unstable electrothermal oscillations. However, this setup is sufficient for

are assembled in the same can, but they are not in physical contact, as seen in Fig. 7.10. The SQUID measures the magnetic flux generated by the inductor, which is proportional to the current through the TES. The SQUID output voltage is passed through separate cabling back to the room temperature controller and readout. A printed circuit board has been designed and built to handle the connections between the bias resistor, TES, SQUID, and inductor, shown in Fig. 7.12.

Finally, the TES bias line goes down to the actual sensor at the mixing chamber stage. Our sensor holder, a copper “boat” shown in Fig. 7.13, is mounted on the mixing chamber flange after being encased in Cryoperm to shield the TES against magnetic fields. High magnetic fields can prevent materials from transitioning into the superconducting state. Our chip contains two duplicates of four TESs (eight sensors total). All TESs are 40 nm thick tungsten films on silicon substrates. They have areas of $800\ \mu\text{m} \times 200\ \mu\text{m}$, $400\ \mu\text{m} \times 100\ \mu\text{m}$, $200\ \mu\text{m} \times 50\ \mu\text{m}$, and $100\ \mu\text{m} \times 25\ \mu\text{m}$. Four sensors are fabricated on the chip with no extra support; the other four are identical, but they are surrounded by etched coils to provide a strong thermal link to the chip. In this dissertation, we measure the two largest TESs with coils. The bias power of a TES scales with volume, and we were unable to transition the smaller TESs because their bias power was smaller than the parasitic noise power⁵. The TES chip is glued onto the boat using rubber cement, which holds its adhesion at cryogenic temperatures. In the absence of any current, the sensors should be at roughly the same temperature as the flange, but the holder contains a separate thermometer to record accurate readings.

Absent from the circuit diagram, but present in the setup, is data acquisition equipment. We can probe the analog voltage across the $2\ \text{k}\Omega$ resistor and the analog voltage output of the SQUID. These are both digitized at a 50 kHz sampling rate with a National Instruments 9229 module. Data are saved to a PC with LabView software, which averages continuous samples over some time window, usually 0.1 sec.

characterization.

⁵We have since made noise reduction improvements to our experimental setup, allowing the two smaller TESs to transition. However, the dissertation author was minimally involved in these efforts, so they are not included.

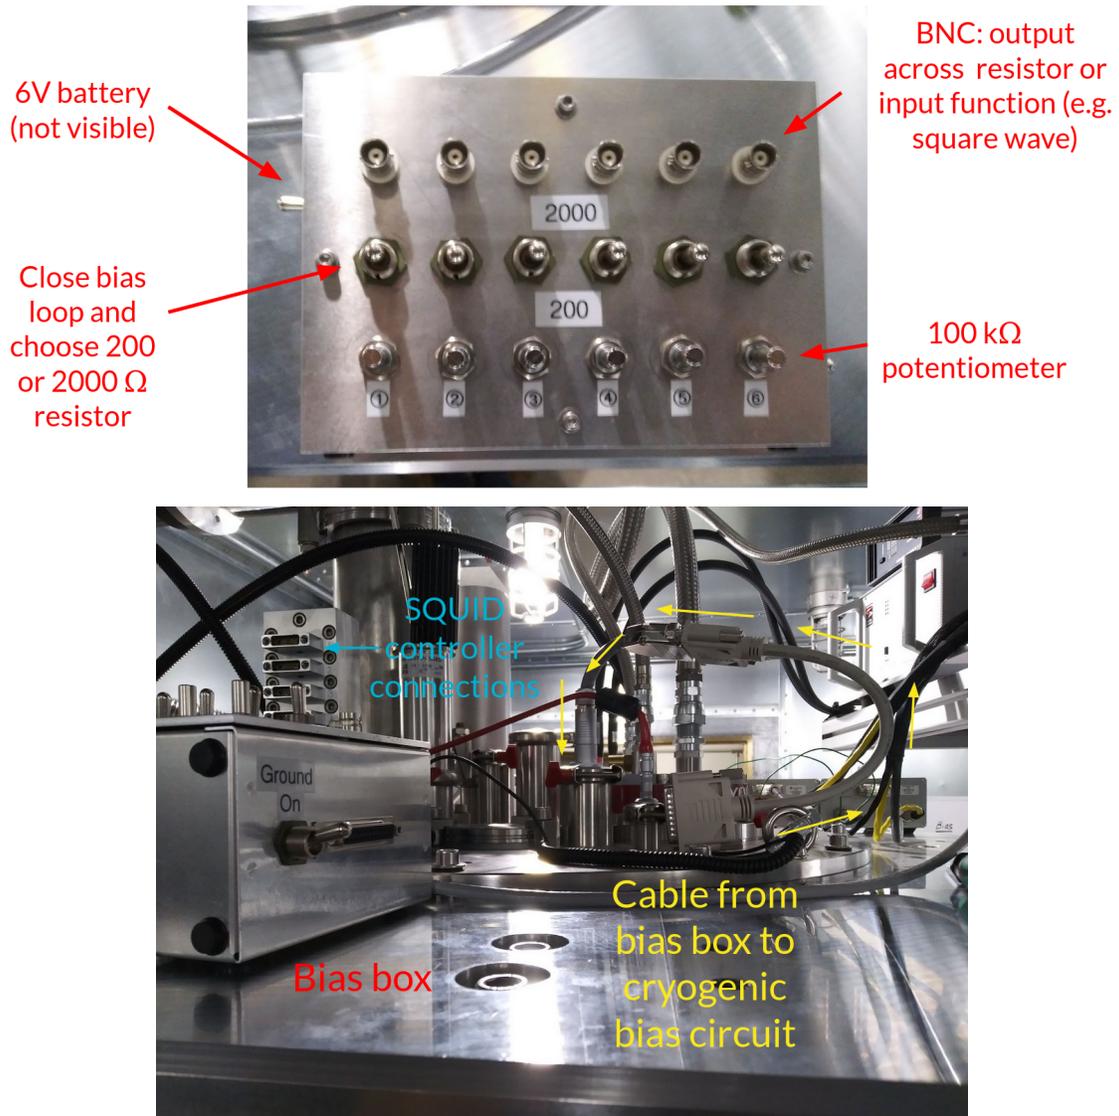


Figure 7.11: Photos of the bias box located on the platform above the refrigerator and the electronics going into the refrigerator. Connector for room-temp SQUID controller is disconnected in image. See also diagram in Fig. 7.10 and photos in Figs. 7.12 and 7.13.

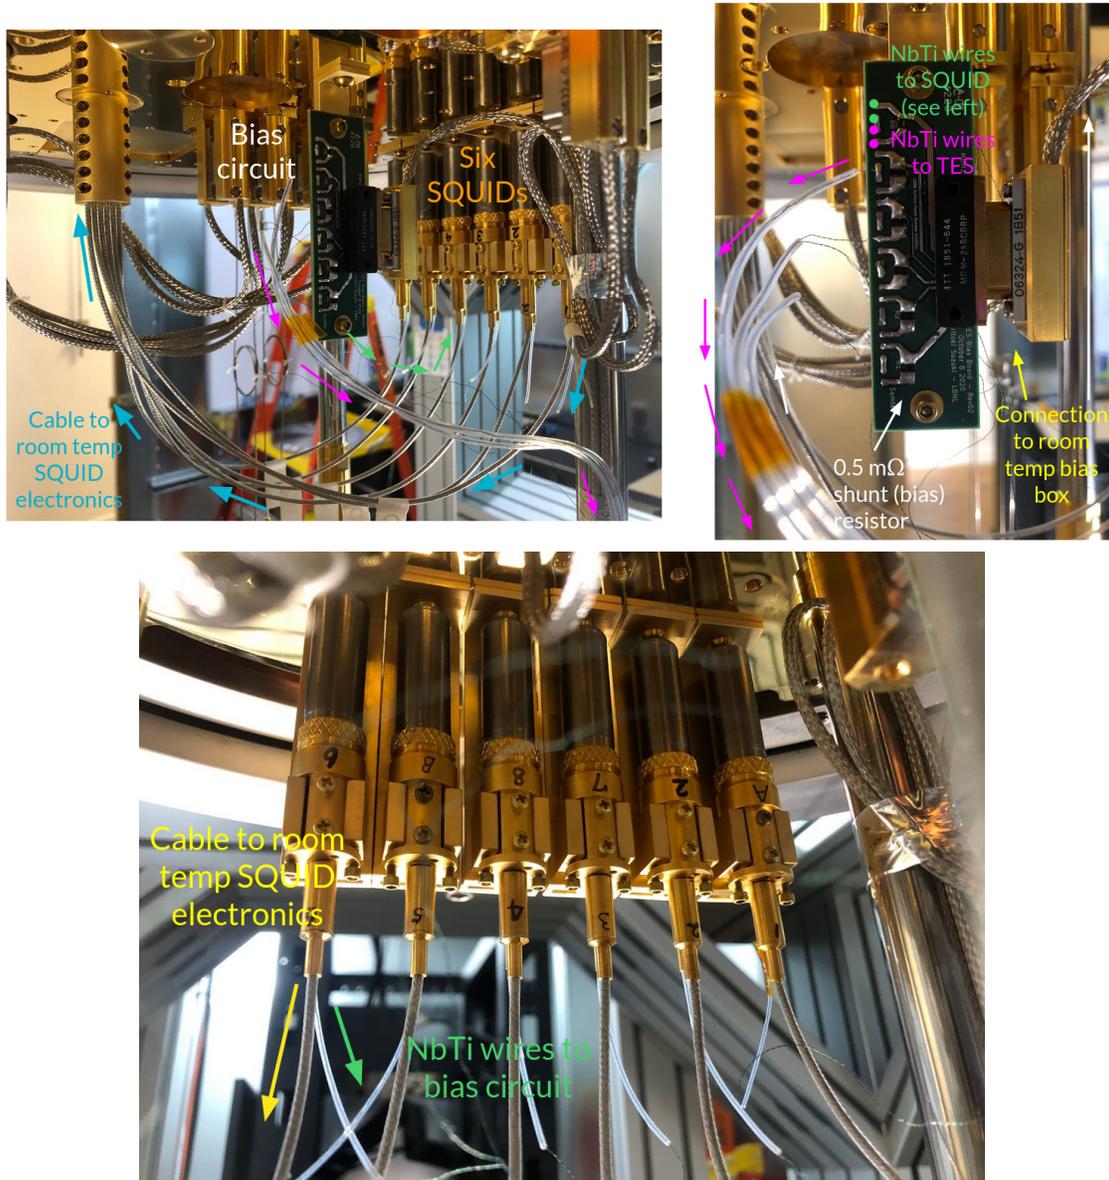


Figure 7.12: Photos of the components of our TES circuit that are mounted to the 4 K flange (flange is the plate at the top of the photos). Bias circuit and resistor are shown, as well as their connections to room-temperature electronics and the TES. DC SQUIDs are also shown; the input inductor is inside the SQUID can. See also diagram in Fig. 7.10 and photos in Figs. 7.11 and 7.13.

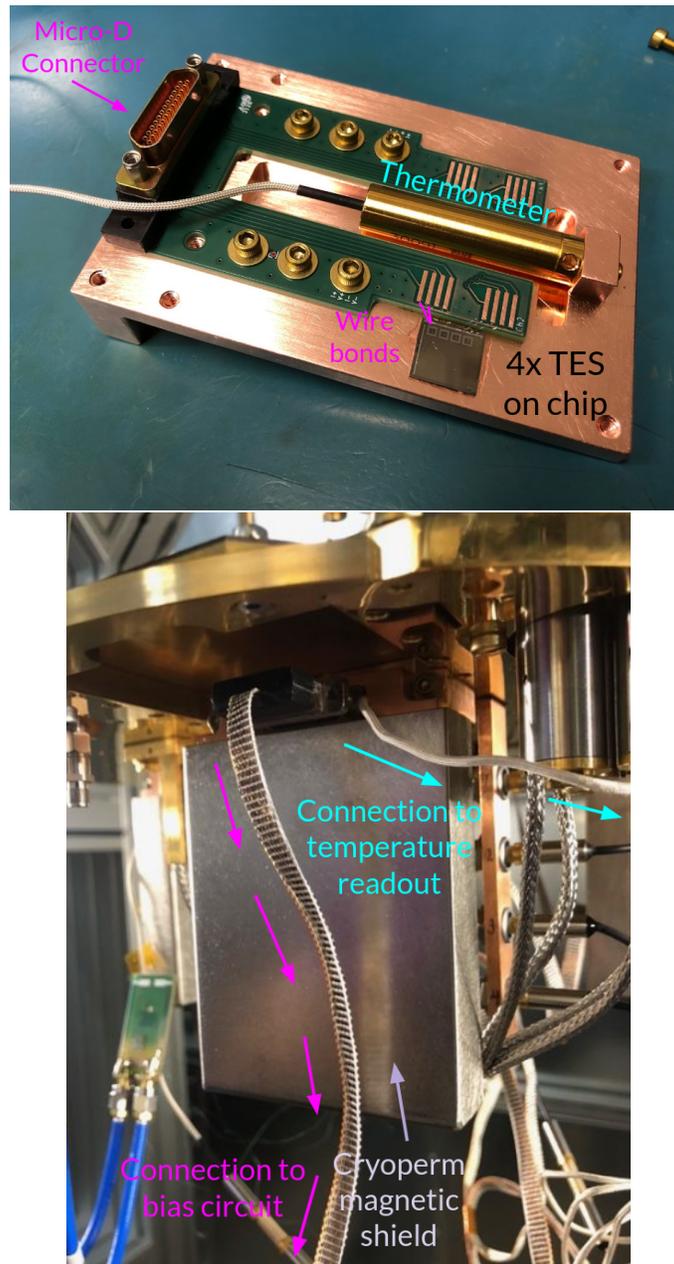


Figure 7.13: Photos of the TES holder, a copper boat mounted to the mixing chamber flange (flange is the plate at the top of the bottom photo). (*Top*) The boat is open. We can see the sample thermometer and the printed circuit board that connects the bias line to the TES chip. Each of the small squares on the edge of the chip is a coil used to strengthen the thermal connection to the silicon substrate; the TESs are inside the squares, but they are too small to be seen without a microscope. (*Bottom*) The boat is closed, encased in Cryoperm for magnetic shielding. See also diagram in Fig. 7.10 and photos in Figs. 7.11 and 7.12.

7.4.2 A Single TES Measurement

To describe our measurement procedure, it is useful to follow an example. We choose the $800\ \mu\text{m} \times 200\ \mu\text{m}$ TES. We know that both TESs have a transition temperature of $T_c \approx 46\ \text{mK}$, so we cool the mixing chamber below this temperature. In this example, the mixing chamber temperature is $34\ \text{mK}$. The mixing chamber flange, which the TES is mounted to, serves as a thermal bath. In the absence of any TES current, the TES will also be at $34\ \text{mK}$ and therefore in the superconducting state. However, we voltage bias the TES sufficiently high that the Joule power V^2/R raises its temperature above its transition point, and thus it behaves like a normal resistor. Recall that the maximum bias in our setup is $1.5\ \mu\text{V}$, which is more than sufficient to keep the TES normal⁶. We also start the data acquisition software, recording the voltage across the $2\ \text{k}\Omega$ resistor and the SQUID output.

We then lower the voltage bias using the potentiometer, decreasing the TES temperature and bringing it into transition. Figure 7.14 shows the results of this procedure. The top-left panel shows the voltage across the $2\ \text{k}\Omega$ resistor vs. time. We start turning the potentiometer at $t = 360$ samples and slowly decrease the bias until $t = 630$. Then, we go in the opposite direction and increase the bias until we return to the original voltage at $t = 870$. The top-right panel shows the corresponding SQUID output voltage vs. time. Between $t = 360$ and $t = 410$, it decreases, but then it increases until $t = 630$, the same time that we hit the minimum voltage bias. Between $t = 630$ and $t = 870$, while we are increasing the voltage bias, the SQUID voltage traces its earlier path in reverse.

However, we see a sharp change in the output voltage at $t = 720$. This is a phenomenon in the SQUID, not the TES. Recall that the current in a SQUID loop will adjust itself so that the enclosed magnetic flux is $\Phi = Z \frac{h}{2e}$, where Z is any integer. During normal operation, this integer stays the same, but our observation at $t = 720$ corresponds to the SQUID jumping from Z to $Z + 1$. We operate the SQUID in linear feedback mode (called “flux-locked loop”) to prevent this, but it sometimes occurs accidentally, usually due to a twitching finger while rotating the potentiometer. Flux jumping also occurs when the measured current is beyond the range of the SQUID, but that is not a possibility based on the current range we see. Regardless, this is an imperfection in the measurement procedure. It could be fixed at the hardware level by e.g., use of a function generator to control the voltage bias. We can also adjust for the flux jump at the analysis level; in the top-right panel, if we subtracted off the voltage jump, then the measurements after $t = 720$ would be a reflection of the earlier path of the SQUID output. We do neither in this analysis⁷ and simply ignore data with a flux jump. When flux jumps occur, we either take the measurement again or (as we will show) we use only the data from

⁶We bias the TES before cooling down, so that it stays in the normal state.

⁷A function generator has been implemented for future measurements, but those measurements are not included in this dissertation.

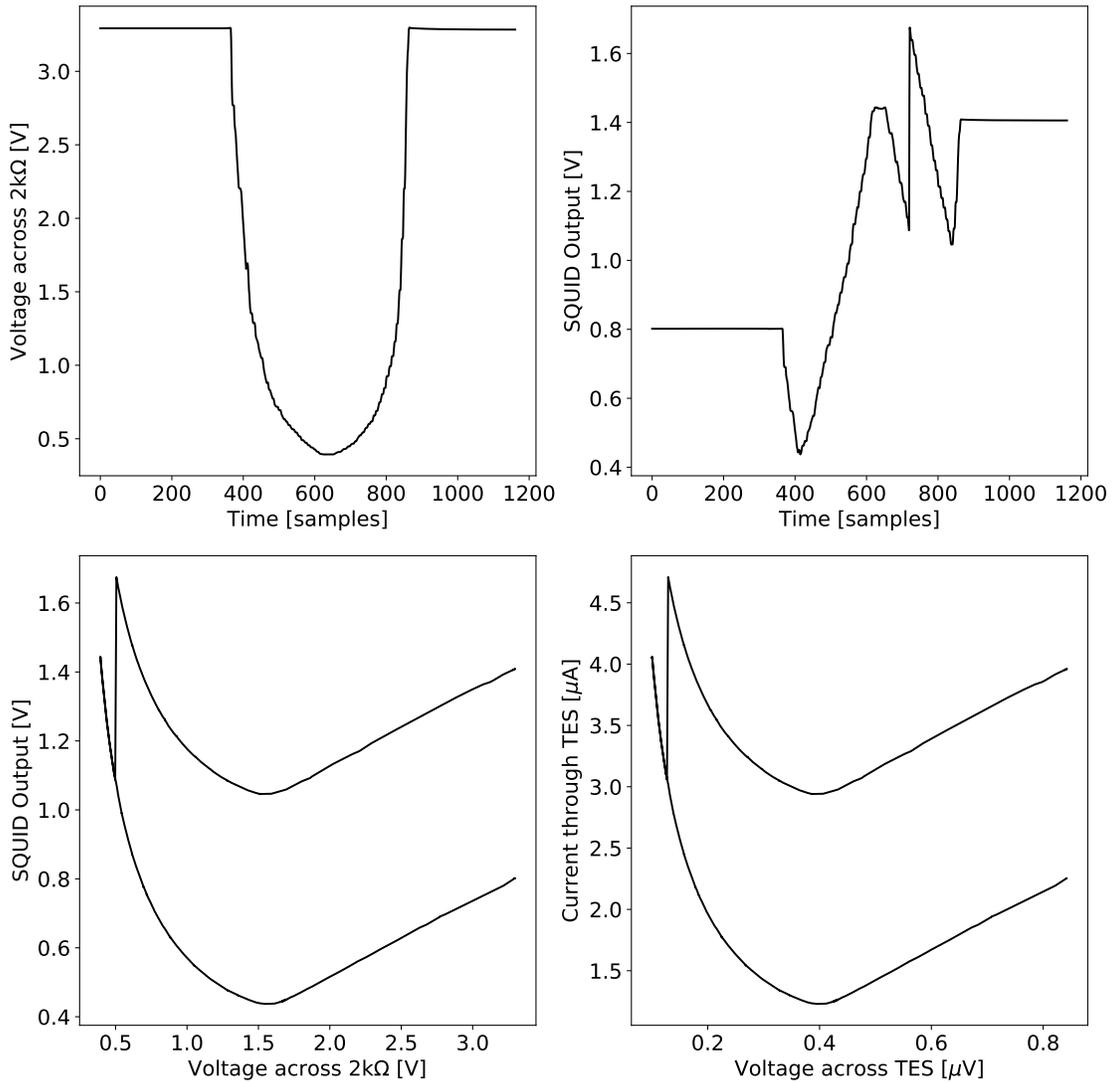


Figure 7.14: Example of a single TES measurement. See text for discussion. (*Top panels*) Voltage across the 2 kΩ resistor and DC SQUID output, vs. time. One sample is equal to 0.1 second. (*Bottom-left*) These two measurements vs. each other. (*Bottom-right*) The 2 kΩ resistor voltage, converted to the voltage across the TES, vs. the SQUID output, converted to the current through the TES.

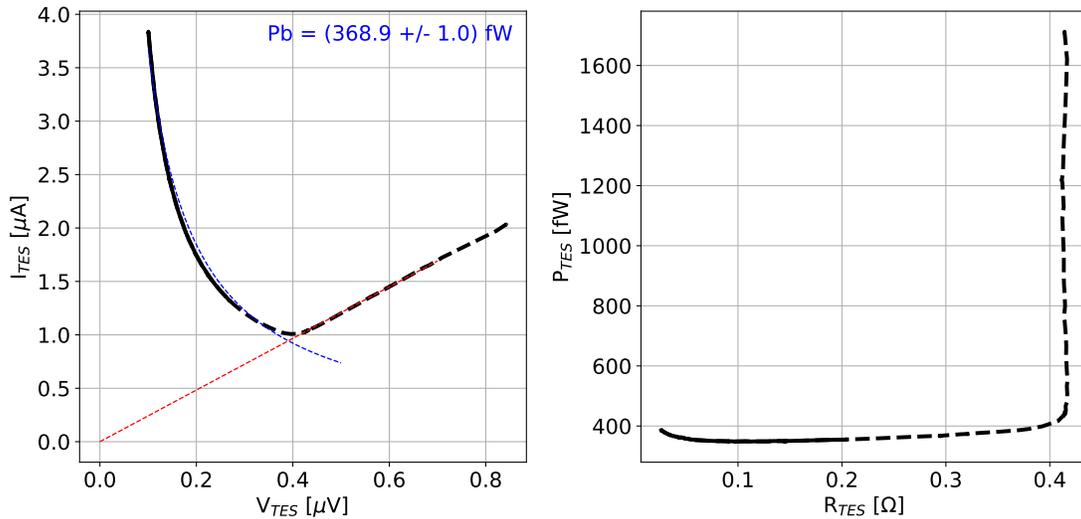


Figure 7.15: (*Left*) Calculation of bias power for the $800 \mu\text{m} \times 200 \mu\text{m}$ TES at a 34 mK bath temperature (same as Fig. 7.14). Data is shown in black. The solid black line represents the region of data used for the inverse fit, shown in blue, whereas the dashed black line represents other data. A linear fit is shown in red. (*Right*) The bias power $P_{bias} = I_{TES}V_{TES}$ vs. TES resistance $R_{TES} = V_{TES}/I_{TES}$ during the measurement.

the increasing-bias or decreasing-bias path. As a result, we don't lose any data from flux jumps.

The two lower panels of Fig. 7.14 show the measured voltages, plotted against each other. Time increases from lower-right to top-right. In the lower-right panel, we have converted the voltage across the $2 \text{ k}\Omega$ resistor to the voltage across the TES, using the circuit diagram in Fig. 7.10. The SQUID voltage has been converted to the current through the TES using a previous calibration of the SQUID. We see that the TES begins by acting like a normal resistor, as we expect since it is above its transition temperature. The plot is linear: $I_{TES} = V_{TES}/R_N$. This corresponds to the decreasing SQUID output voltage in the upper-right panel. However, as we decrease the TES voltage bias (and thus the temperature) even further, the current begins to increase. This indicates that the TES has begun to transition. During transition, the TES resistance varies, but the bias power is constant. Therefore, we observe an inverse relationship: $I_{TES} = P_{bias}/V_{TES}$. If we lowered the voltage bias even further, the TES would go fully superconducting, and we would see another linear relationship based on the parasitic resistance. However, as discussed above, we instead brought the TES back to its normal state.

Now, we can measure the actual bias power. This is shown in Fig. 7.15. We have totally removed the data from the second half of our acquisition, when the bias

voltage was increasing and we saw a flux jump, but we still have sufficient data. First, we use the linear portion of the curve to add a constant offset to the SQUID voltage. In a DC SQUID, zero voltage out does not necessarily correspond to zero measured current, but we can force this by offsetting the voltage. Second, we fit the transition regime to the function $I_{TES} = P_{bias}/V_{TES}$ and obtain a measurement for P_{bias} . In this case, we see $P_{bias} = (368.9 \pm 1.0)$ fW. We also show the power in the TES vs. the TES resistance, which clearly shows the sharp transition between a constant-resistance regime and a constant-power regime.

7.4.3 TES Measurement Results: Bias Power and Conductance

We follow this procedure for both TESs and for several bath temperatures. The results are shown in Fig. 7.16.

During the transition, the TES stays roughly at a constant temperature. It is in a thermal quasi-equilibrium; the power lost to the flange via thermal conduction is equal to the Joule power dissipated I^2 . In general, the conduction power and thermal conductance are given by Eqs. 7.7 and 7.8, respectively.

$$P = K(T_c^n - T_{bath}^n) \quad (7.7)$$

$$G = \frac{dP}{dT_c} = nKT_c^{n-1} \quad (7.8)$$

Here, K is a constant that we can measure or fit, and n is some real number based on the type of interface. Recall from Fig. 7.1 that there are technically two thermal connections. For an ideal TES-absorber interface, we expect $n = 5$, consistent with the thermal impedance between electrons and phonons in small-volume metals with high power densities [166]. For an ideal absorber-bath interface, we expect $n = 4$, as calculated in [168] for interfaces between dissimilar materials with mismatched elastic constants. We did not have the capability to measure the conductances of the two interfaces separately, so our measurements are for the entire TES-bath connection:

$$G^{-1} = G_{ta}^{-1} + G_{ab}^{-1} \quad (7.9)$$

The TES-absorber thermal link is usually weaker than the absorber-bath link (i.e., $G_{ta} < G_{ab}$), so it should dominate. As a result, we expect n to be closer to 5 than to 4. For both TESs, we fit P_{bias} vs. the mixing chamber (bath) temperature and extract K , G , and n . We also redo the fit, but fixing $n = 5$.

Both transition-edge sensors have essentially the same transition temperature of $T_c = 46.68$ mK (using the fit where n is free). We also get $n = 4.52$ for the larger TES and $n = 4.36$ for the smaller TES. The values of n are a bit smaller than we expect,

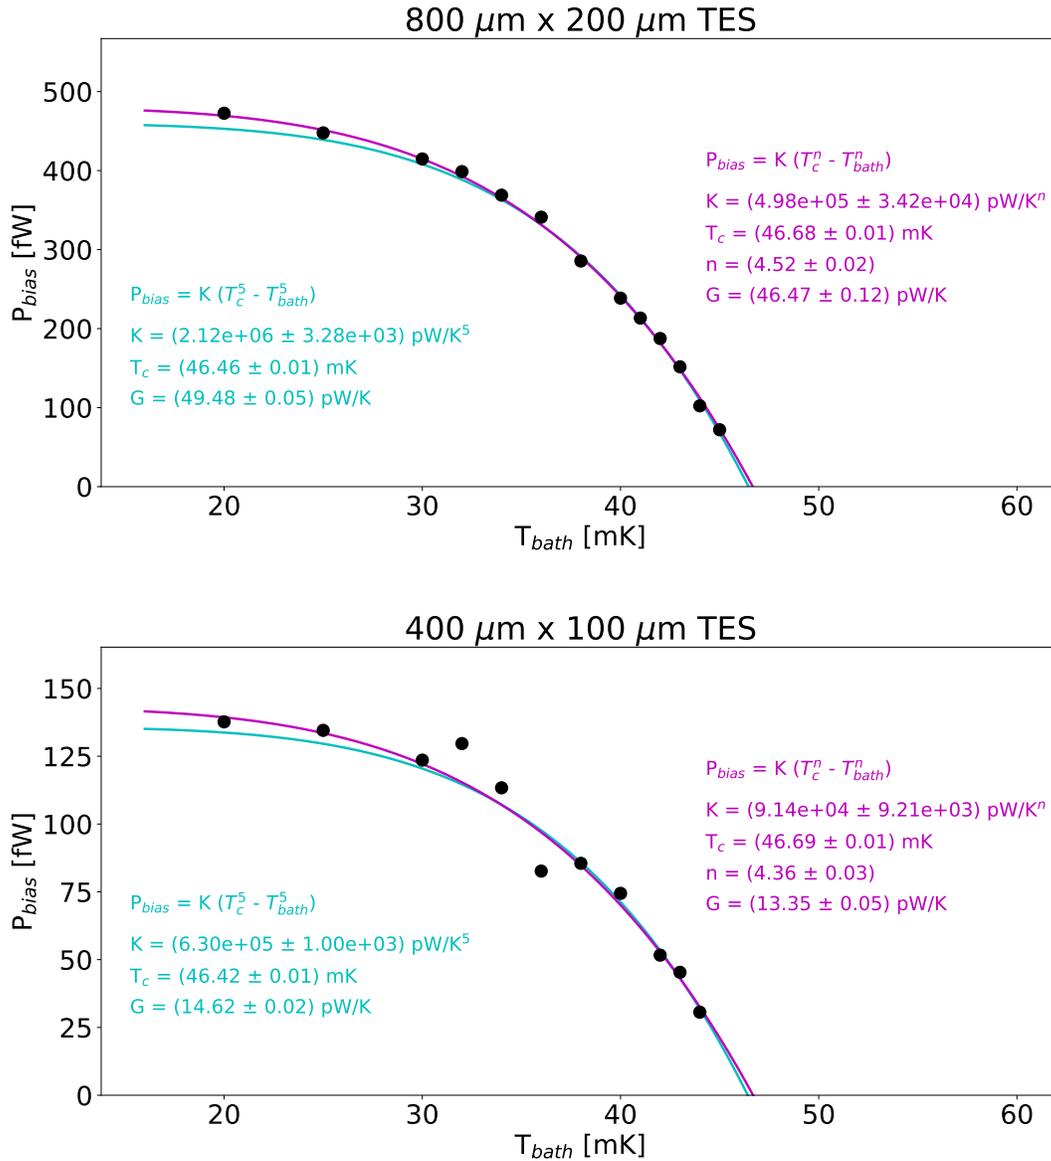


Figure 7.16: Measured bias power vs. the temperature of the mixing chamber flange (i.e., the bath) for each TES. We fit these to Eq. 7.7, both allowing n to float (magenta) and fixing $n = 5$ (cyan).

Table 7.1: Measured normal and parasitic resistances of our two TESs.

	800 μm x 200 μm TES	400 μm x 100 μm TES
R_N [m Ω]	393.83 ± 0.03	392.50 ± 0.01
R_P [m Ω]	4.8830 ± 0.0004	6.1539 ± 0.0005

suggesting irregular phonon transmission at one or both interfaces, perhaps due to imperfect contact. Another indication of irregular transmission is that in the bottom panel, the P_{bias} vs. T_{bath} data are poorly fit to a polynomial function; the statistical error bars on each point are smaller than the points themselves. Meanwhile, we measure $G = 46.47$ pW/K and $G = 13.35$ pW/K for the larger and smaller TES, respectively. The thermal conductance should scale with volume, so we expect a factor of 4 between them. Instead, we see a factor of 3.48 between the measured values of G . We do not have a confirmed explanation for this discrepancy. Some hypotheses include: inaccurate TES fabrication, improper mounting, and parasitic power.

7.4.4 TES Measurement Results: Normal and Parasitic Resistance

We conclude by reporting the normal and parasitic resistance of each TES. These are measured in the same procedure outlined in Sec. 7.4.2; we adjust the voltage bias and measure the current through the TES with the SQUID. To measure the normal resistance, we keep the mixing chamber well above 46 mK, so that the TES does not transition. For the parasitic resistance measurement, the mixing chamber should be well below 46 mK, and the voltage bias should be small enough that the TES stays in the superconducting state. In both measurements, the $I - V$ curve is linear, and we use the slope to measure the resistance. The results are shown in Table 7.1. They are consistent with our expectations of a normal resistance of around 400 m Ω and parasitic resistance of a few m Ω .

7.5 Measurement of T_c and Film Stress for Several Materials

Independently of our TES measurements, we also measured the transition temperature and film stress of several materials to gauge their suitability as TESs. Our goals were to find the lowest-possible T_c and film stress to enhance our experimental sensitivity and reduce backgrounds.

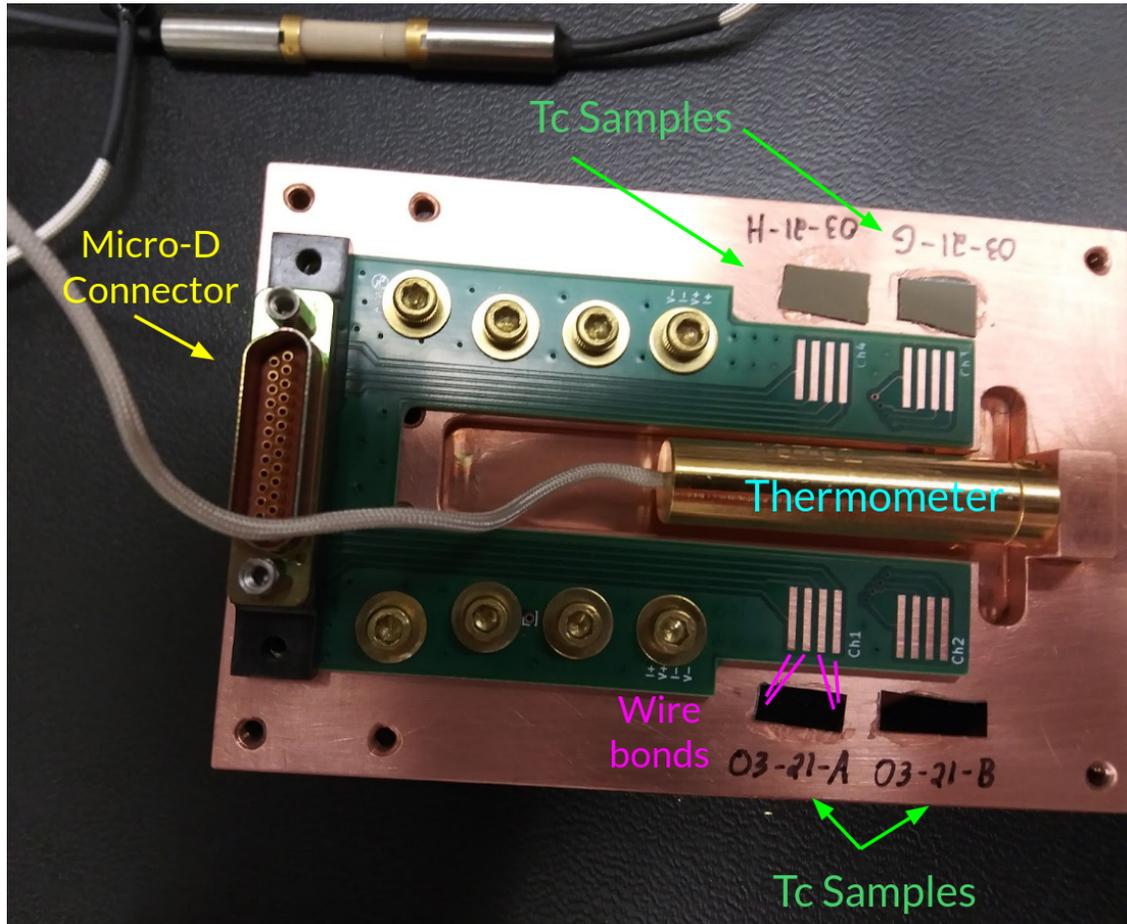


Figure 7.17: The copper holder for our T_c sample measurements. Four samples are mounted on this holder, and the holder is mounted to the mixing chamber flange in the same way as Fig. 7.13. The labels “03-21-A”, etc. are just the sample names.

All our samples are W films on Si substrates, deposited by Mark Platt at Texas A&M University, but the film deposition procedure varies. Most were deposited with a sputtering system built by Semiconductor Engineering Group Inc (SEGI). A few were deposited with a sputtering system manufactured by AJA International. Both systems were modified by M. Platt to accommodate the necessary wafers. Various deposition parameters were varied, such as the deposition time, the chamber gas, the gas flow rate, the sputtering power, and more.

7.5.1 Experimental Setup for T_c Measurements

To measure T_c , we use a Lakeshore Model 372 AC Resistance Bridge, which performs a 4-point resistance measurement. An example of our setup is shown in Fig. 7.17,

where we have mounted four samples to a single copper boat. In this case, a “sample” is simply a small piece of a given material. The reader will notice this is the same structure as the holder for our TES measurement (Fig. 7.13). However, rather than being connected to a bias circuit, the Micro-D connector links directly to the Lakeshore bridge. Each sample is connected to the printed circuit board via four wire bonds, corresponding to the I+, I-, V+, and V- terminals, and it is glued down with rubber cement. Additionally, the holder contains a ruthenium oxide (RuOx) thermometer. This thermometer is connected to a separate Lakeshore bridge; the bridge measures resistance, but the thermometer is calibrated so we can measure temperatures down to 7 mK.

To measure the transition of a given W film, we can sweep the temperature of the mixing chamber flange. BlueFors allows us to do this cleanly with a heater mounted on the flange and a PID controller. Meanwhile, we connect the Lakeshore bridge output to the previously-mentioned NI 9229 digitizer to record data. In this way, we get measurements of the instantaneous temperature of the sample holder (which should be similar to the flange temperature, but not necessarily identical) and the resistance of the sample.

An example measurement is shown in Fig. 7.18. In this case, the sample transitions at 49.12 mK. The transition is rapid, with a total width of about 2 mK. We can also measure the temperature sensitivity α , defined earlier in Eq. 7.4. It depends on the exact point in the transition, so as a reference, we define it at the point where the resistance is one-third its maximum value. We find $\alpha = 87.5$.

7.5.2 T_c as a Function of Chamber Gas

Our primary finding using only T_c measurement data is the effect of the gas used in the sputtering chamber on the transition temperature. Specifically, we find that using xenon decreases the T_c , relative to argon.

In order to make this observation, our TAMU collaborators fabricated tungsten films using two sputtering recipes on the SEGI (a recipe is a combination of flow rate, pressure, power, deposition time, and other parameters). These recipes had been developed previously for use with argon; they were found to have T_c on the order of 120 mK and 50 mK. The primary differences were in the power, gas pressure, and gas flow rate; the lower- T_c recipe has lower power, higher pressure, and higher flow rate. For each recipe, several samples were produced, with varying amounts of xenon and argon gas. The total flow rate was kept constant, but the xenon concentration varied between 0% and 100%.

The results are shown for the two recipes in Fig. 7.19. We see very clearly that as Xe concentration increases, T_c decreases, as does α . For the high-xenon samples in the low- T_c recipe, the T_c is so low that at 7 mK, the samples don't transition at all. We also see that the trend is incredibly sharp, going from 48 mK to <7 mK in the low- T_c recipe when we increase the xenon concentration from 30% to 50%. Note

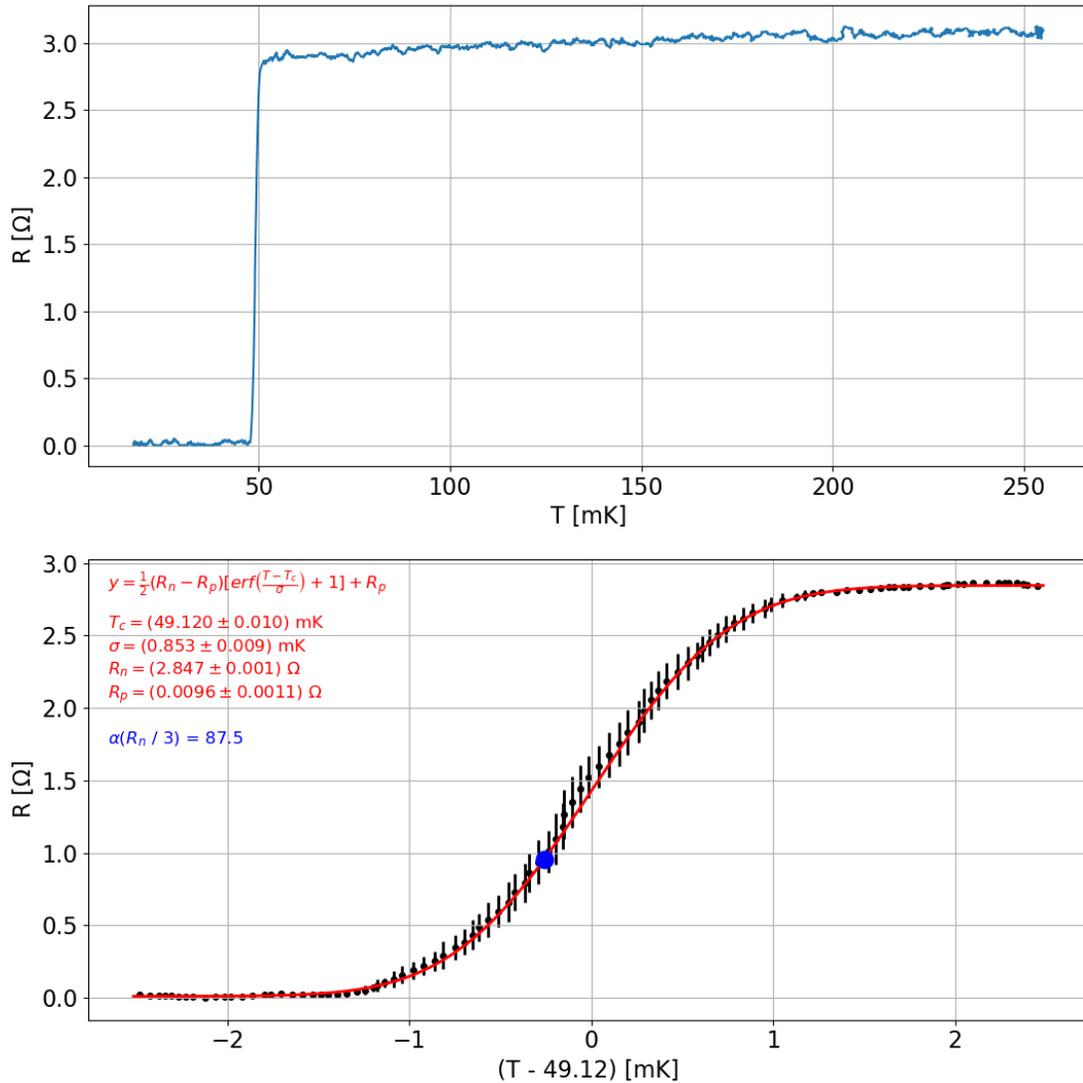


Figure 7.18: An example of a T_c measurement. (*Top*) The resistance vs. temperature measurements for a single T_c sample. (*Bottom*) The same measurement, but zoomed in on the transition region. We fit the transition to an error function, shown in red. The transition temperature is fit to $T_c = 49.12$ mK. The sensitivity of the sample α is shown at the temperature where the resistance is one-third its maximum value (the blue dot).

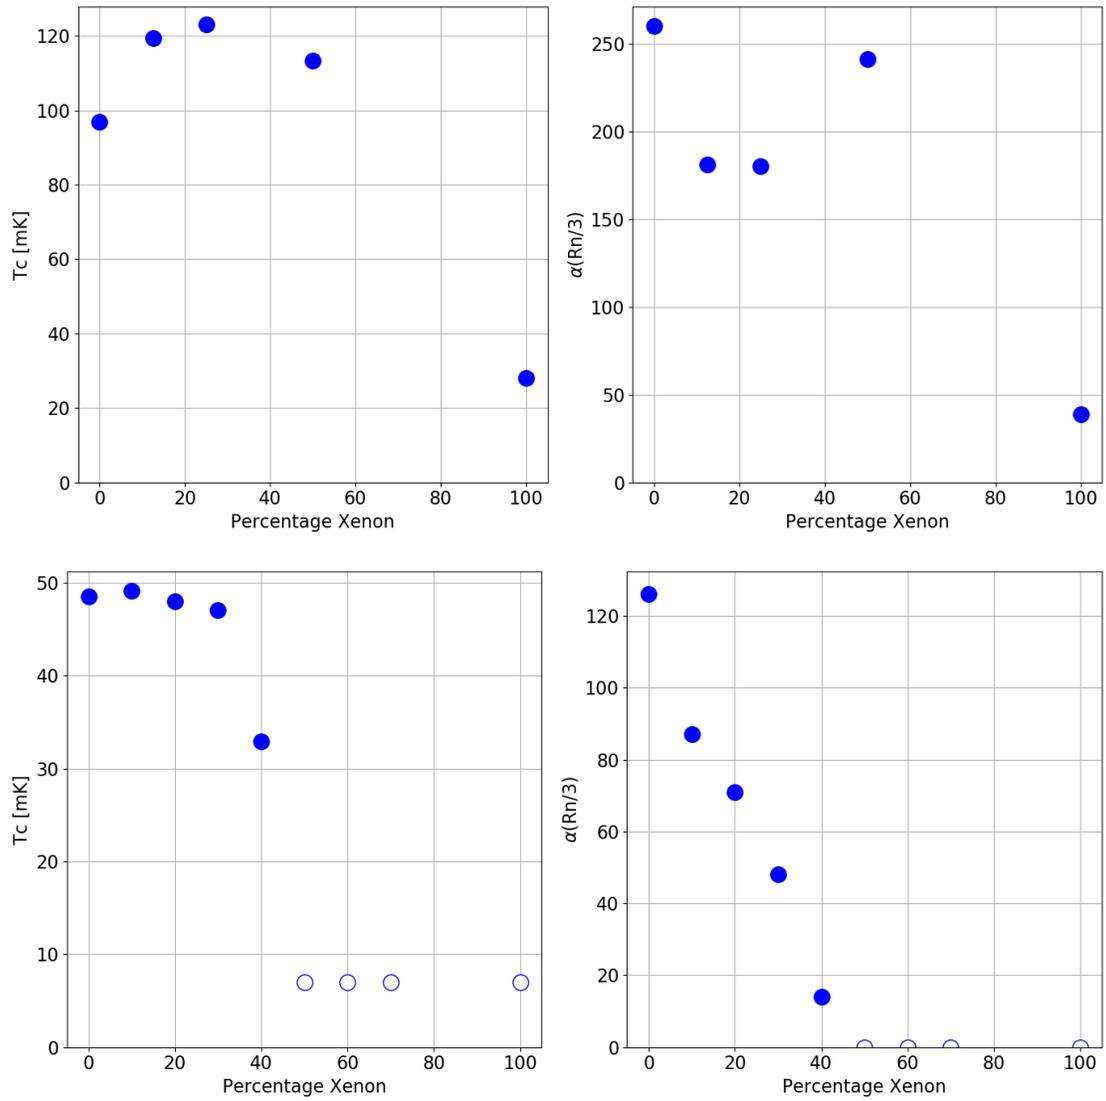


Figure 7.19: Measured T_c and α vs. xenon concentration for the “high- T_c ” (top) and “low- T_c ” (bottom) recipe. Open circles represent samples that did not transition at 7 mK.

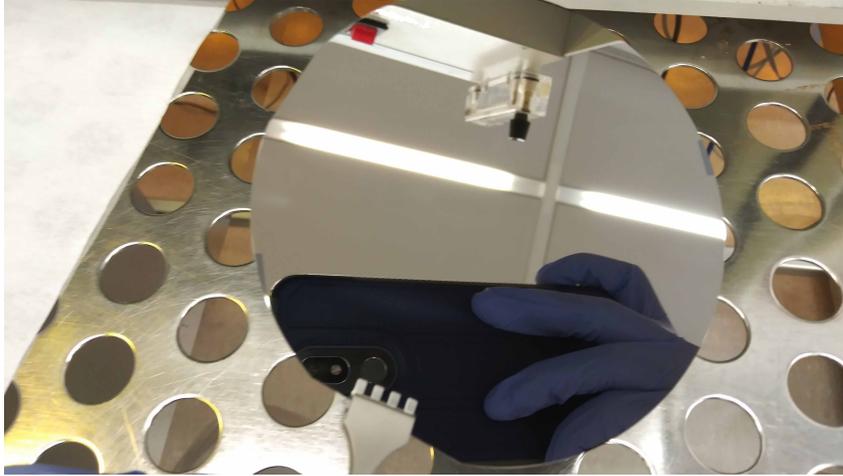


Figure 7.20: Example of a 4-inch diameter silicon wafer. The wafer is bare; no film is deposited. The polished side is facing up, as can be seen from its reflectivity.

that the trend seems to slightly fail for the 0% Xe point in the high- T_c plot, but that sample was from an earlier dataset. That sample was fabricated in November 2020, whereas the others were fabricated in May 2021. Ideally, the same recipe should produce identical T_c , but we do see variability of this magnitude in other datasets.

This is an incredibly important finding. It is clear that to achieve low T_c , the use of xenon gas instead of argon is crucial. The decreasing α is undesirable, since high α provides greater sensitivity to low-energy interactions, but additional work can be done to try to optimize this. Our results are complementary to Bouziane et al. 2005, [169], who found that film stress is higher for W films on Si substrates when the sputtering gas is xenon, compared to argon. Bouziane et al. also found strong dependencies of crystal structure on the sputtering gas. Tungsten exhibits two crystalline phases—the α phase has $T_c = 15$ mK, and the β phase has T_c between 1–4 K [170]. Films sputtered with xenon were associated with higher fractions of α -phase tungsten and larger grain sizes than films sputtered with argon. We cannot necessarily conclude that stress is the mediator between low T_c and chamber gas, though. Instead, we must directly investigate the relationship between T_c and film stress.

7.5.3 Relationship Between T_c and Thin-Film Stress

The relationship between the stress in our films and the measured superconducting transition temperature is not as clear-cut as the effect of sputtering chamber gas. From the results discussed below, it seems that higher tensile stress is correlated with lower T_c , but more work will be needed to confirm this.

To measure film stress, it is not sufficient to use the $\mathcal{O}(\text{cm}^2)$ samples that were

fabricated for only- T_c measurements. Instead, we used 4-inch diameter silicon wafers, as shown in Fig. 7.20. The wafers are Si prime-quality, single-side polished, $\langle 100 \rangle$ oriented, $525 \mu\text{m}$ thickness, and doped with 10^{15} cm^{-3} boron—equivalent to a resistivity of 10–20 Ωcm . The resistivity was chosen to be high enough that only a negligible current flows through the silicon, and the vast majority of the current flows through the tungsten film. Measurements of eight wafers are presented in this Section.

Before any W was deposited, we measured the curvature of the wafer using a Tencor Flexus FLX-2320 (abbreviated as Flexus), housed at the UC Berkeley Marvell Nanofabrication Laboratory. The Flexus uses a laser interferometer to scan the surface of the wafer and measure its deflection from perfect flatness. The deflection can be upwards or downwards. Two example measurements are shown in Fig. 7.21, where we show the deflection vs. distance along the wafer diameter. The deflection is quite small—0.90 and $0.33 \mu\text{m}$, respectively—so the wafer’s curvature is not perfectly circular. In fact, in the bottom photo, we have intentionally shown a wafer whose curvature is quite irregular, like a potato chip. Nevertheless, the Flexus software is able to fit the data and extract a radius of curvature: 977 m and -5411 m, where the sign indicates the direction of curvature.

Then, the wafers were sent to our TAMU collaborators, and a different W film recipe was deposited on each. The film thickness was targeted at 40 nm but ranges from 27 to 68 nm. The W film and Si substrate apply stress to each other, primarily because of mismatched lattices. This stress must be compressive in one material and tensile in the other, with equal magnitude; however, either direction is possible. In this Section, when we describe the “stress,” we refer to the stress in the W film.

The film stress can be measured by repeating the Flexus measurements. Two examples are shown in Fig. 7.22, although not the same wafers we showed pre-film. After deposition, the wafers typically have a larger deflection (either upwards or downwards), and the curvature is closer to circular. However, this is not a fixed rule; the wafer in the bottom photo has a moderately irregular curvature, with an unusually flat center. Regardless, the radius of curvature is extracted by the Flexus software. Finally, to get the film stress σ , we use the Stoney equation below [171], where E_s is the substrate Young’s modulus (equal to 180.5 GPa for Si), t_f is the film thickness, t_s is the substrate thickness, and r_0 and r are the initial and final radii of curvature.

$$\sigma = \frac{E_s t_s^2}{6 t_f} \left(\frac{1}{r} - \frac{1}{r_0} \right) \quad (7.10)$$

The stress measurements are done at room temperature, but we actually need to know the film stress at T_c . As we cool down the wafer, the Si substrate and W film both shrink via thermal contraction. Tungsten has a larger coefficient of thermal expansion, so it tends to shrink more than the silicon; in order to keep the substrate and film in contact across their entire interface, the tungsten must be stretched by

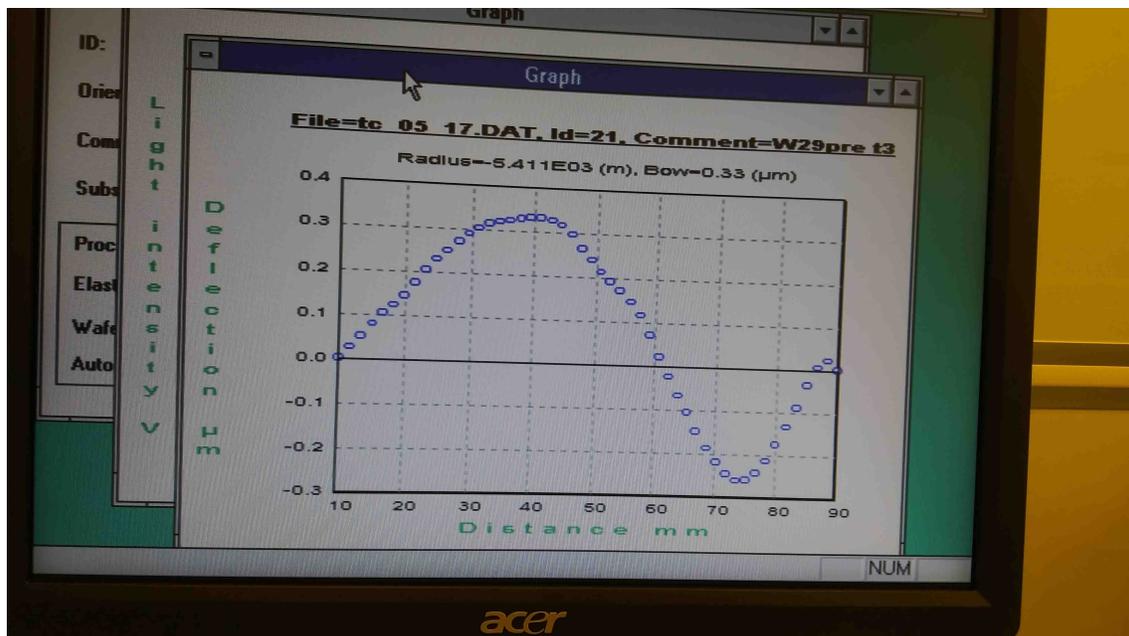
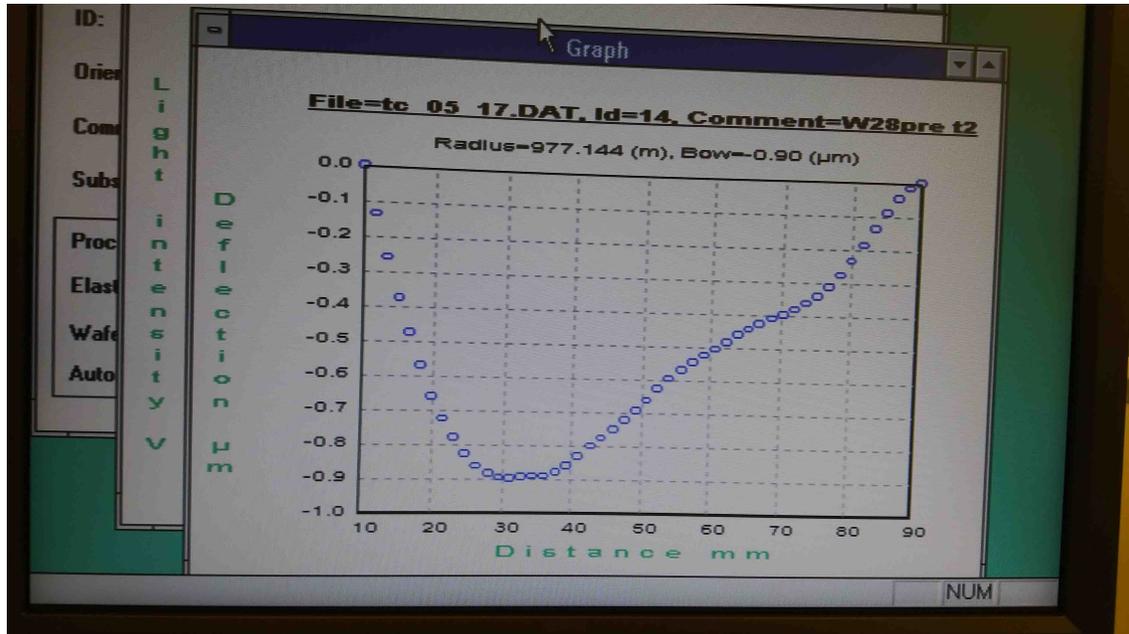


Figure 7.21: Two Flexus measurements of wafer deflection, before any film is deposited.

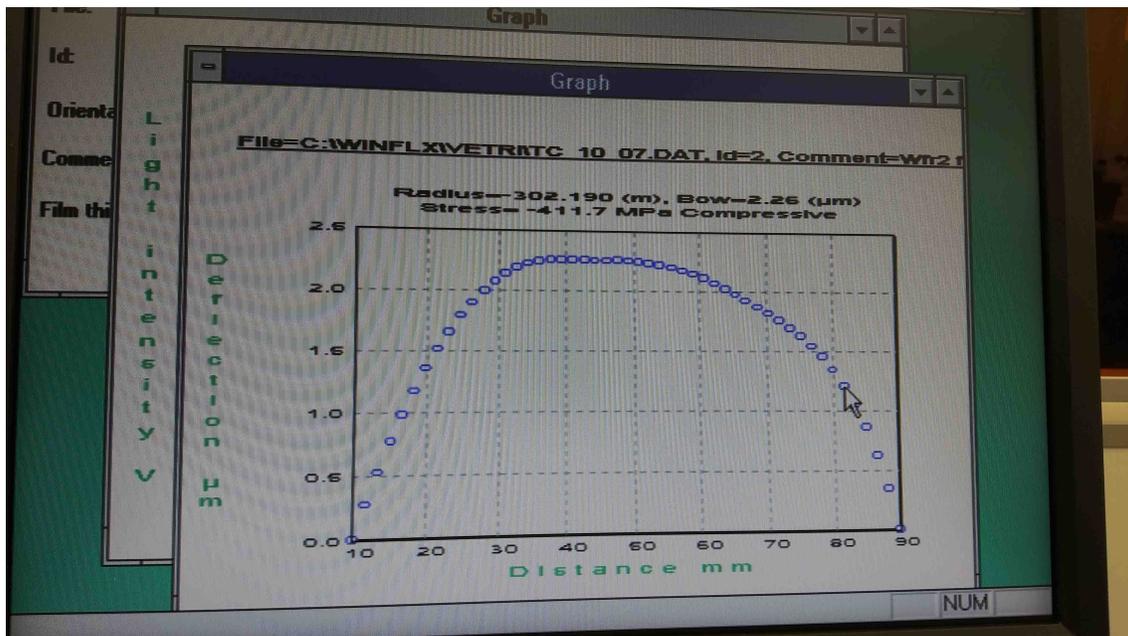
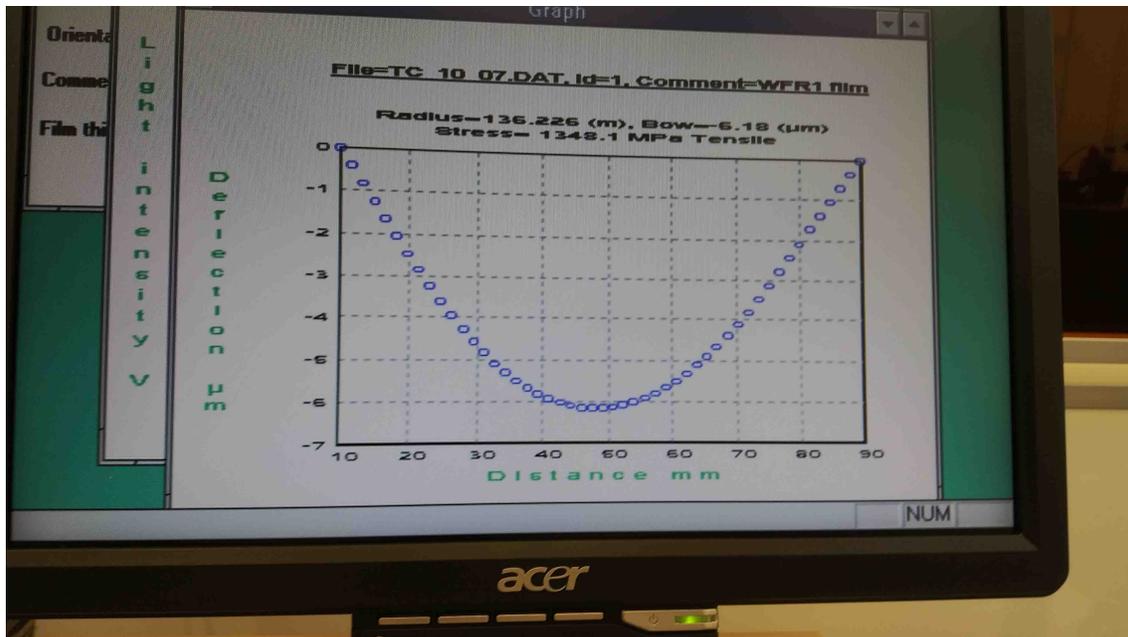


Figure 7.22: Two Flexus measurements of wafer deflection, after W film was deposited. These are not the same wafers from Fig. 7.21. The stress values reported on the screen are approximately correct, but the measurements presented in this Section are more accurate.

Table 7.2: Tungsten film stress for eight wafers. Positive values indicate tensile stress, and negative values indicate compressive stress. Measured at 300 K, but an extra 400 MPa tensile stress is added to account for thermal contraction. The sputtering machine and chamber gas are also noted.

Wafer Label	Machine	Gas	W Film Stress [MPa]
A	SEGI	Ar	2321 ± 4
B	SEGI	Ar	-201 ± 12
C	SEGI	Ar	2021 ± 2
D	SEGI	Xe	4040 ± 40
E	SEGI	60 Ar/40 Xe	1600 ± 11
F	AJA	Ar	-1910 ± 30
G	SEGI	Xe	1850 ± 70
H	SEGI	Xe	1250 ± 130

the silicon. In other words, the cooldown creates an induced tensile stress in the film. We can calculate the induced thermal stress:

$$\sigma_{th} = \frac{E_W}{1 - \nu_W} \int_{300 \text{ K}}^{0 \text{ K}} [\alpha_W(T) - \alpha_{Si}(T)] dT \quad (7.11)$$

Here, E_W is the tungsten Young's modulus, ν_W is the tungsten Poisson's ratio, and α_{Si} and α_W are the temperature-dependent coefficients of thermal expansion of silicon and tungsten. We set the upper integrand limit as 0 K for ease of notation, but in tungsten and silicon (as in most materials), thermal contraction stops below several tens of K.

At room temperature, $E_W = 405 \text{ GPa}$ and $\nu_W = 0.28$ [172]. Unfortunately, these parameters have not been measured at cryogenic temperatures, but through surveying other materials like aluminum, stainless steel, and titanium, we typically observe a 10–20% decrease in E and a 5–10% increase in ν upon cooldown [173]. In addition, $\int_{300 \text{ K}}^{0 \text{ K}} \alpha_W dT = 85.8 \times 10^{-5}$ and $\int_{300 \text{ K}}^{0 \text{ K}} \alpha_{Si} dT = 21.5 \times 10^{-5}$ [174, 175]. Thus, we assume $\sigma_{th} = (400 \pm 50) \text{ MPa}$, where the uncertainty is a rough estimate for our ignorance of E_W and ν_W at 0 K.

With these tools, we can calculate the film stress at 0 K. This is done for eight wafers and shown in Table 7.2. The error bars are determined by doing the Flexus measurements multiple times and seeing the range of radius of curvature. The variation arises from slightly different alignments of the wafer in the Flexus, an issue that is exacerbated by non-circular curvature. Repeated measurements were only done for post-film wafers, but pre-film wafers usually have a significantly higher

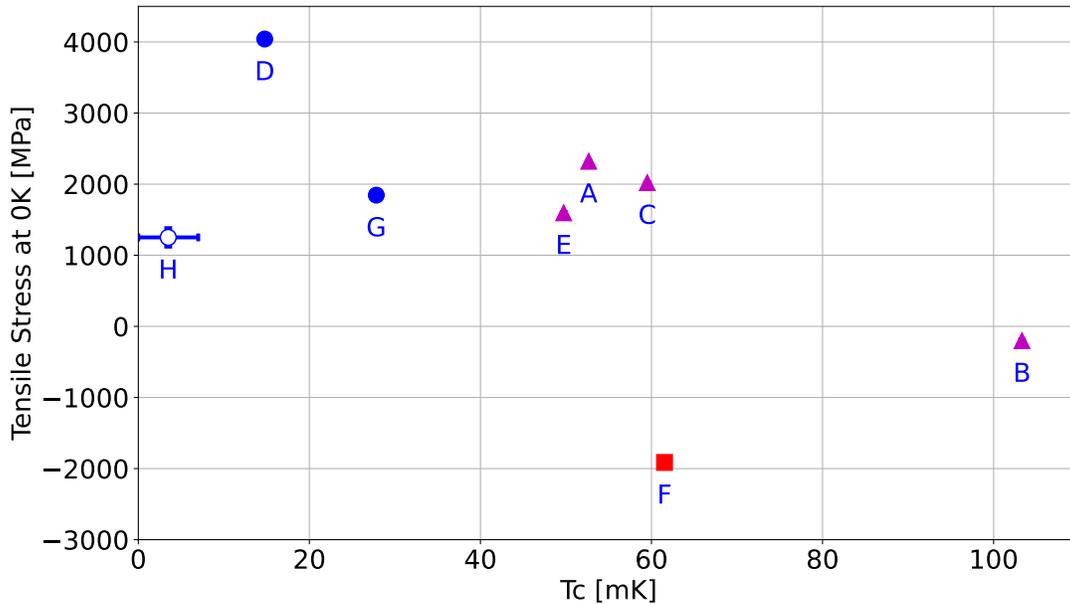


Figure 7.23: Measured stress vs. measured transition temperature for eight wafers. Red squares indicate films deposited with the AJA machine, using Ar gas in the chamber. Magenta triangles indicate films deposited with the SEGI machine, using Ar or a mixture of Ar/Xe. Blue circles indicate films deposited with the SEGI machine, using only Xe. The blue open circle is a wafer that did not transition; we show it as having $T_c < 7$ mK, but it possibly can't transition at all. Under each data point, we note the label of each wafer.

radius of curvature, so their uncertainty is lower-impact. We observe that the total uncertainty on the measured stress is small.

Finally, we cleaved the wafers with a diamond scribe pen and cut off a small $\mathcal{O}(\text{cm}^2)$ piece of each for T_c measurements. The T_c measurements were done identically to Sec. 7.5.1. In Fig. 7.23, we show these measurements vs. the stress measurements. The red squares, blue circles, and magenta triangles are used to note the sputtering machine and chamber gas. Seven of the eight wafers transitioned; Wafer H did not. Our dilution refrigerator goes as low as 7 mK, so it's likely that the transition temperature of Wafer H is < 7 mK, but it's also possible that the wafer cannot transition at all.

The relationship between transition temperature and film stress is weak. There appears to be a negative correlation, but the correlation is not well-established. For example, Sample G has a similar stress as Samples A, C, and E, but a much lower T_c . Meanwhile, Sample F has a similar T_c as A, C, and E, but a much lower film stress. As mentioned earlier in this Chapter, our goals are low T_c and near-zero stress, so

if the relationship is truly negative, it would be impossible to accomplish both goals simultaneously. However, more data is clearly needed to study this relationship, and to establish whether we can develop a low- T_c , zero-stress film.

Chapter 8

Helium as a Dark Matter Direct Detection Target

In this chapter, we focus on the superfluid helium-4 component of TESSERACT. This detector is named the Helium Roton Apparatus for Light Dark Matter, or HeRALD.

HeRALD will be used to perform a search for nuclear recoils on helium. It is primarily sensitive to MeV-scale dark matter; the technology is technically capable of measuring recoils from WIMPs, but the relevant parameter space has already been excluded by other experiments, including LUX. We discuss the HeRALD technology and its advantages, and we explore its sensitivity to dark matter interactions. We also calculate the neutrino fog in helium.

The work done in this chapter was done collaboratively by the dissertation author, Scott Hertel, Andreas Biekert, Junsong Lin, and Daniel McKinsey. It was published in Physical Review D [176]. We give disproportionate focus, though, to the work done by the author: the calculations of experimental reach and the neutrino fog.

8.1 Detector Concept

Helium-4 has several intrinsic advantages for MeV-scale searches. One is its low mass. As mentioned in Sec. 2.4, the kinetic energy transfer is optimized when the nuclear mass and the dark matter mass are similar. A $4 \text{ GeV}/c^2$ nucleus is much more suitable for searches below the Lee-Weinberg limit than a heavy nucleus such as xenon (about $130 \text{ GeV}/c^2$). Helium does not exhibit a high A^2 cross-section enhancement, but as we saw earlier in Fig. 2.11, there is barely any parameter space explored at the sub-GeV scale. Therefore, we don't need to achieve cross-sections as low as LZ to achieve world-leading sensitivity. Another intrinsic advantage is its fluid nature down to absolute zero temperature. This means that the technology is scalable; if we want a bigger detector, we just need a larger vessel and more helium, rather than e.g., building multiple towers of solid-state crystals. Finally, helium is incredibly easy to purify, since it has the lowest boiling point (4.2 K) of any element. It has only two isotopes, ^3He and ^4He , neither of which is radioactive. Our HeRALD concept considers a detector of ultrapure ^4He .

8.1.1 Energy Partitioning in Superfluid ^4He

To go beyond the physical properties of helium, we need to consider the energy partitioning, as we did for xenon. For xenon, the partitioning was shown in Fig. 3.4, whereas for helium, we show it in Fig. 8.1. The reader will note similarities between the two. This is to be expected, since both are noble elements. There are three channels of energy deposition: excitation, ionization, and motion. However, there are notable differences, as well. Since we plan to operate HeRALD in a zero-field

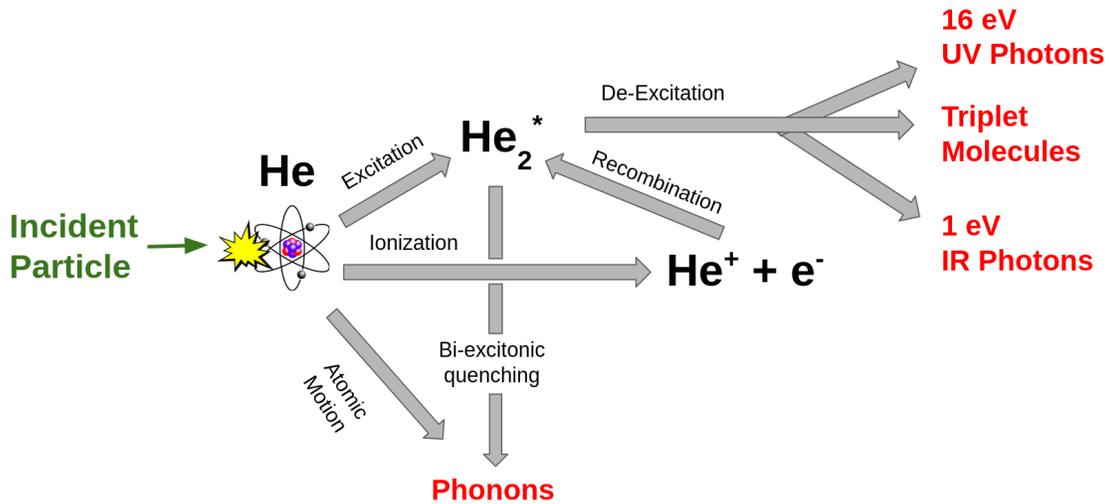


Figure 8.1: A generic picture of energy partitioning in superfluid helium.

environment, electron-ion recombination should be mostly complete. In other words, the ionization signal is effectively absent.

The excitation channel can exhibit the same dynamics as xenon; the excitons de-excite, emitting photons in the process. For helium, the photons have a wavelength of 80 nm, or equivalently an energy of 16 eV. However, this is only true for the singlets $^1\Sigma_u^+$. The triplets $^3\Sigma_u^+$ technically can decay the same way, but their lifetime is incredibly long: 13 s [177]. As a result, we consider triplet excitons to be long-lived molecules, propagating through the surrounding helium. Some energy also goes into infrared ≈ 1 eV photons.

Finally, the heat channel is dramatically different than xenon, since ^4He is a superfluid at our temperatures¹. Superfluid helium acts like a Bose-Einstein condensate with zero viscosity, since ^4He nuclei are bosons. Energy deposited as atomic motion thus manifests itself in quasiparticles—collective motions of the superfluid. As in other systems, these can be called *phonons*, but in superfluid ^4He , certain quasiparticles are also called *rotons*. The top panel of Fig. 8.2 shows the dispersion relation of these quasiparticles, modeled as a high-degree polynomial fit to measured data from [180]. Below a momentum of 2.2 keV/c or above 3.8 keV/c, the energy increases with momentum, as seen in most systems. However, in between these two momenta, the energy actually decreases with increasing momentum. As a result, the dispersion curve is commonly split into three regions. In order of increasing momentum, these are phonons, R- rotons, and R+ rotons.

These quasiparticles propagate ballistically through the helium. Some sponta-

¹The superfluid transition occurs at 2.17 K [179], and recall from Chapter 7 that we are considering sensors with transitions below 100 mK, so the helium needs to be colder than that.

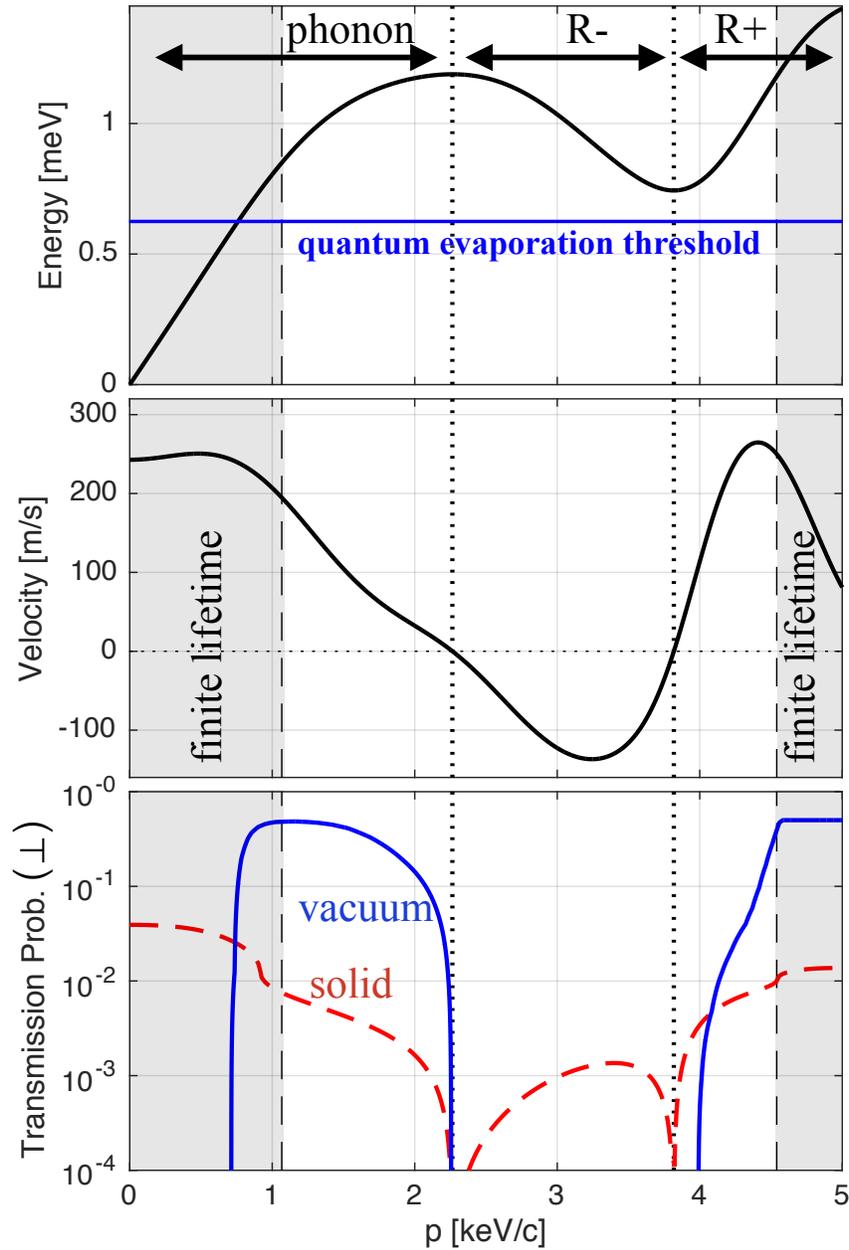


Figure 8.2: (*Top*) The energy vs. momentum dispersion of quasiparticles in ${}^4\text{He}$, along with labels for the phonon, R- roton, and R+ roton momentum regimes. The gray shaded region highlights energies where quasiparticles can spontaneously down-convert to two lower-energy quasiparticles. (*Middle*) The group velocity of the quasiparticles. (*Bottom*) When encountering a solid or vacuum interface at zero degrees, the probability of transmission or evaporation, respectively. Same as top row of Fig. 8.3, for the zero-degree case. Figure made by S. Hertel, published in [176].

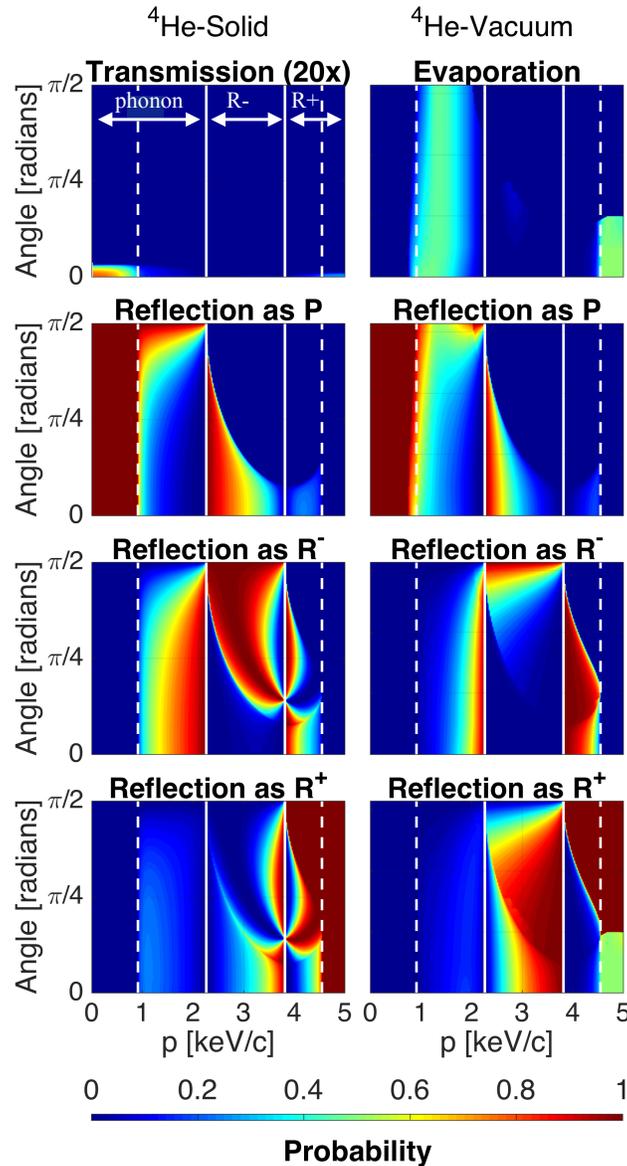


Figure 8.3: (*Left*) Probability of a quasiparticle being transmitted at a solid interface or reflected as any of the three quasiparticle flavors, given the initial momentum and incidence angle. The solid white lines show the momentum cutoffs for each flavor. The dashed white lines show the minimum and maximum momentum for quasiparticle conversion; a quasiparticle can only change flavor if energy is conserved. See Fig. 8.2 for the energy-momentum relation. The transmission probability is close to 0%, so it is multiplied by a factor of 20 for ease of visibility. (*Right*) The same, but for a vacuum interface. Transmission corresponds to quantum evaporation. The quantum evaporation probability from [178] has been halved to better match experimental data. Figure made by S. Hertel, published in [176].

neously decay to two lower-energy quasiparticles, but this is forbidden for quasiparticles with momenta between 0.83–4.54 keV/c. The quasiparticles that are able to decay do so at a rate of $A\epsilon^5$, where $A = 7.12 \times 10^5 \text{ s}^{-1}$ and ϵ is the energy in Kelvin with $k_B = 1$ [181, 182].

Suppose a quasiparticle is incident on an interface between the helium and vacuum. Then, with some probability, the quasiparticle can release a single helium atom from the superfluid if the quasiparticle energy is at least 0.62 meV. This process is called “quantum evaporation.” If the energy is insufficient, or simply if quantum evaporation doesn’t occur because of the non-unit probability, the quasiparticle will be reflected. However, it can actually be reflected as a different quasiparticle flavor, if energy can be conserved. For example, a 1.5 keV/c phonon can be reflected as a phonon, an R+ roton, or R- roton with the same energy, about 1.1 meV. On the other hand, a 5 keV/c R+ roton must be reflected as an R+ roton or evaporate a helium atom. It cannot be reflected as an R- roton or phonon because the dispersion curve does not allow these quasiparticle flavors to have the same energy, about 1.5 meV. The probabilities for these various reflection and evaporation processes are shown in Fig. 8.3. They were calculated using the descriptions in [183, 184], reducing the quantum evaporation probability by a factor of two for consistency with data. We do not provide calculation details in this dissertation; interested readers can read the linked sources or the discussion in [176].

Now suppose a quasiparticle is incident on a helium-solid interface. The same dynamics exist, except that the numeric probabilities are different, shown on the left side of Fig. 8.3. Also, instead of quantum evaporation, phonons can be transmitted into the solid. For the helium-solid interface, the probability of transmission is incredibly low because of the enormous Kapitza resistance—less than 5% for the optimal momentum and incidence angle, and much less than 1% for other situations.

As we will discuss in the next section, we measure the quantum evaporation signal, so phonon transmission at a solid interface is a loss mechanism. Another loss mechanism is quasiparticle down-conversion. Low-momentum phonons and high-momentum R+ rotons can decay, as discussed above, but now we see that medium-momentum quasiparticles can also be converted to lower-momentum flavors through reflection. Those quasiparticles can then decay if their momentum is low enough, or be converted to even lower momentum. During phonon decay, the total energy is conserved, but if the energy decreases below the 0.62 meV quantum evaporation threshold, we will not be able to observe these phonons.

8.1.2 Sensing Energy Depositions

Figure 8.1 shows HeRALD’s general concept of sensing the energy deposited by an interaction. As discussed above, energy depositions create photons, long-lived triplet molecules, and quasiparticles.

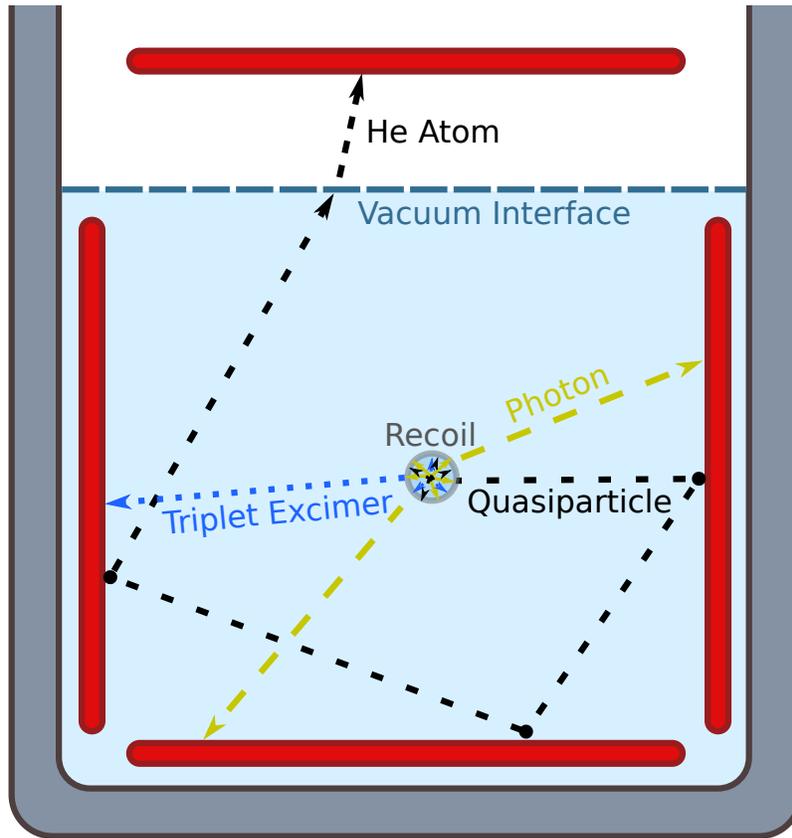


Figure 8.4: HeRALD detector concept. The shaded light blue region is superfluid ^4He , the red panels represent calorimeters instrumented with TES's, and the gray box is a vessel for the helium. Figure made by A. Biekert, published in [176].

The photons (of all wavelengths) are absorbed by the silicon calorimeters, and the corresponding energy is read out by the TES's. The photon absorption occurs within ns or less of the recoil, so the signal shape is dominated by the TES fall time, typically several tens of μs .

The triplets travel ballistically at speeds of $\mathcal{O}(\text{m/s})$ until they are quenched on the silicon calorimeters. This signal occurs some tens to hundreds of ms after the recoil for a $\mathcal{O}(10 \text{ cm})$ detector. The quenching process is complex, and we don't describe it in detail, but we point the interested reader to discussions in [176, 185].

The quasiparticles bounce around the detector, reflecting at the solid surfaces with $>99\%$ efficiency. Eventually, most quasiparticles will hit the vacuum interface and quantum evaporate helium atoms; others will down-convert or decay to sub-evaporation energies, as discussed earlier. We expect this process to take tens to hundreds of ms. The evaporated helium atoms will be adsorbed onto the top calorimeter. The energy deposited in the calorimeter is not the quasiparticle energy less the evaporation energy, as might be naively expected. Instead, the ^4He

atom is attracted to the calorimeter by van der Waals forces. These forces can be controlled by applying a layer of high-attraction material to the silicon. For example, fluorographene has a theoretical adsorption energy of 42.9 meV [186]. A 1 meV phonon would deposit $1 - 0.62 + 42.9 = 43.38$ meV in the calorimeter, a substantial improvement on the atom's energy immediately after evaporation: $1 - 0.62 = 0.38$ meV.

8.1.3 Electronic Recoils, Nuclear Recoils, and Discrimination

As in xenon, nuclear and electronic interactions create different partitioning between the energy channels. The mean fraction of energy going into each channel—quasiparticles as a whole, singlet excitons, triplet excitons, and infrared—is shown in the top panels of Fig. 8.5 for nuclear and electronic recoils. Nuclear recoils have a higher fraction of energy in quasiparticles across all recoil energies, and the singlet and triplet fractions are peaked in certain regimes. The Lindhard model [114] is used to determine nuclear quenching. Meanwhile, electronic recoils tend to have a roughly equal distribution among channels, regardless of the total recoil energy. Helium has a particularly high minimum excitation energy of 19.8 eV and ionization energy of 24.6 eV. Therefore, below 19.8 eV, all energy manifests as quasiparticles.

To understand discrimination between NRs and ERs, we also need to estimate the fluctuations around these mean fractions. First, for a given energy, Poisson statistics are used on the singlet, triplet, and IR means to determine a stochastic number of quanta produced. The remainder of the energy is assigned to quasiparticles. Next, detector efficiencies are modeled with binomial statistics: 95% detection efficiency for singlets and IR photons, $5/6 = 83\%$ for triplets, and 5% for quasiparticles. The quasiparticle detection efficiency is lower than suggested by the previous discussion on reflection and transmission, but we are extremely conservative and somewhat consistent with experimental data [187]. Finally, calorimeter fluctuations are added to each signal channel at the level of 0.5 eV Gaussian noise. All these fluctuations are simulated with a Monte Carlo, shown in the bottom panels of Fig. 8.5.

The ER leakage fraction can be estimated by grouping all energy channels into two categories: excitation and quasiparticles. Then, the discriminant is simply the ratio of quasiparticles to excitation. We can construct ER and NR regions, similar to the ER and NR bands in xenon, and calculate the fraction of ER events falling into the NR region. The ER leakage at 50% NR acceptance is shown in Fig. 8.6. We emphasize that below 19.8 eV, discrimination is impossible because no excitations are created, and all energy is in the form of quasiparticles.

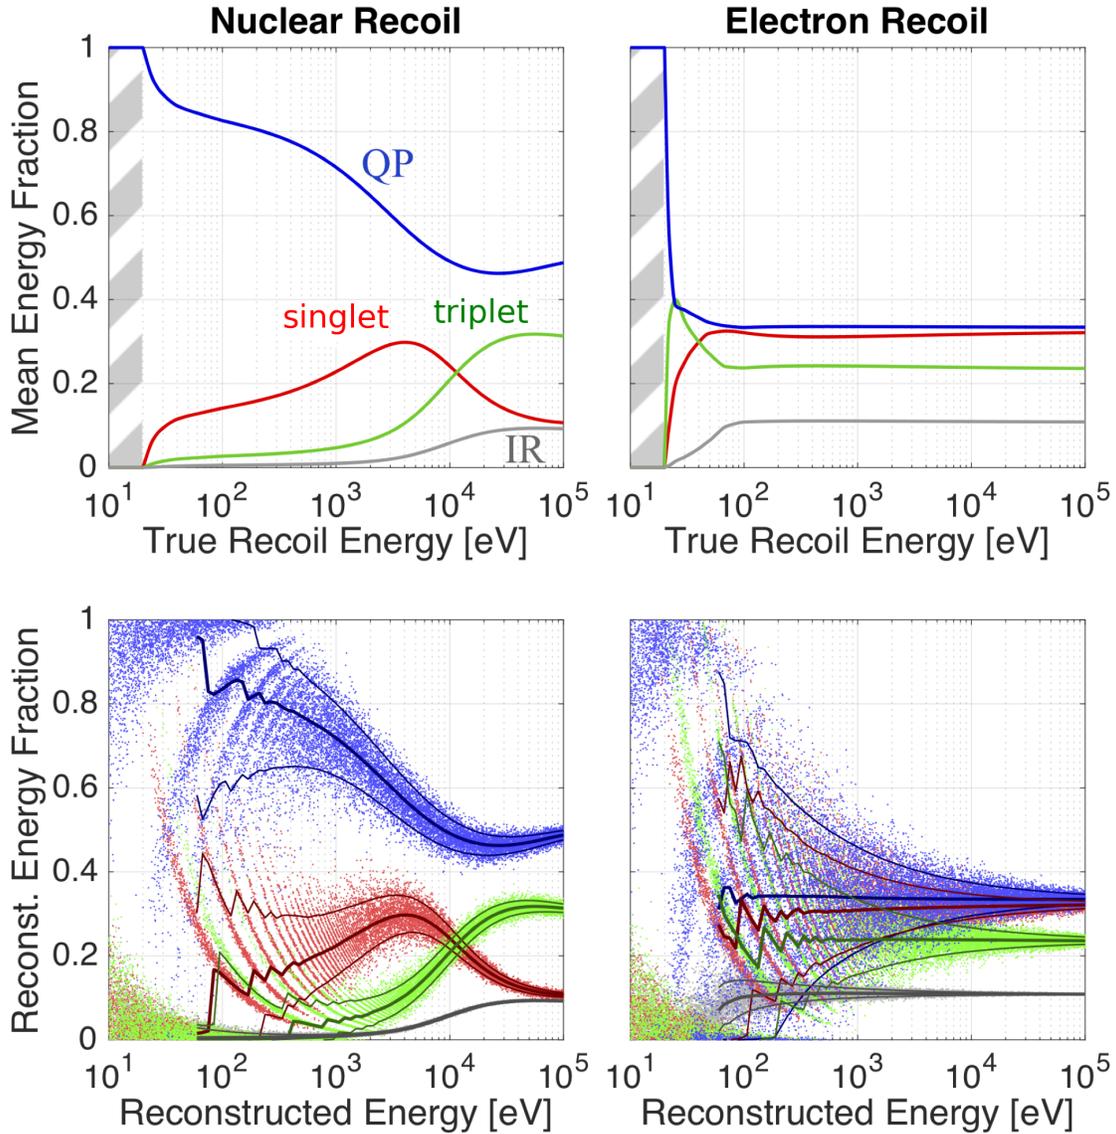


Figure 8.5: (*Top*) Mean fraction of recoil energy going into each channel for nuclear and electronic recoils. The gray shaded region indicates the minimum helium excitation energy of 19.8 eV; for recoil energies below this value, only quasiparticles are produced. (*Bottom*) Simulated recoils, based on these means and several sources of fluctuation; see text for details. Energy is reconstructed based on the simulated signal distribution. The solid lines show 10%, 50%, and 90% contours. Figure made by A. Biekert and S. Hertel, published in [176]. Labels for “singlet” and “triplet” are backwards in [176], but corrected here.

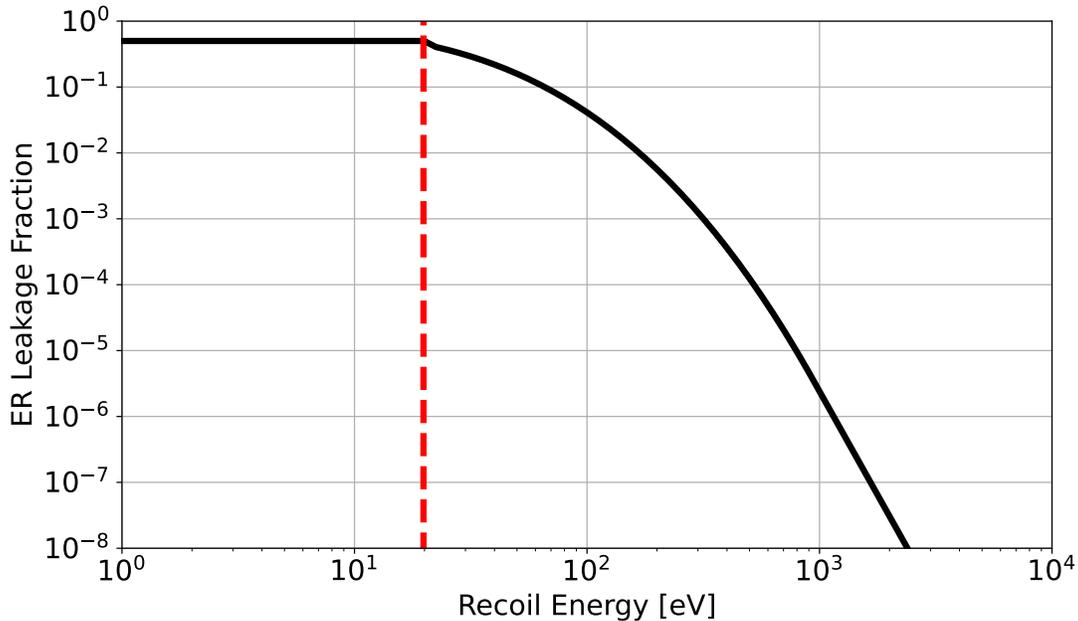


Figure 8.6: The estimated leakage fraction of electronic recoils, at 50% nuclear recoil acceptance. The red dashed line indicates the minimum excitation energy of helium 19.8 eV, below which discrimination is impossible.

8.2 Background Model

To estimate the sensitivity of HeRALD to DM-induced nuclear recoils, we need a background model.

The first background we consider is coherent elastic neutrino-nucleus scattering, often abbreviated as CEvNS. This is a process, long-theorized [188] but first observed in 2017 [189], where neutrinos can scatter off a nucleus through a Z^0 mediator. Similar to WIMP interactions, coherence requires that the momentum transfer be smaller than the nuclear size divided by \hbar . As a nuclear recoil signal, this background is not removable by discrimination.

We follow Billard et. al's treatment of CEvNS backgrounds [31], considering three neutrino sources. The most abundant source of neutrinos is the Sun. The Sun produces neutrinos primarily through the pp chain. Neutrinos are produced as byproducts of deuteron generation through $p-p$ or $p-e-p$ fusion, and then again through later branches: ${}^7\text{Be}$ electron capture, ${}^8\text{B}$ positron decay, or $He-p$ fusion. A smaller number of neutrinos are generated through the CNO cycle, during positron decay of ${}^{13}\text{N}$, ${}^{15}\text{O}$, or ${}^{17}\text{F}$. All of these neutrinos are electron-flavor with lepton number +1 (i.e., not antineutrinos), but they oscillate during the $\sim 10^8$ km journey to the Earth, and only 55% are electron-flavor by the time they are incident on a

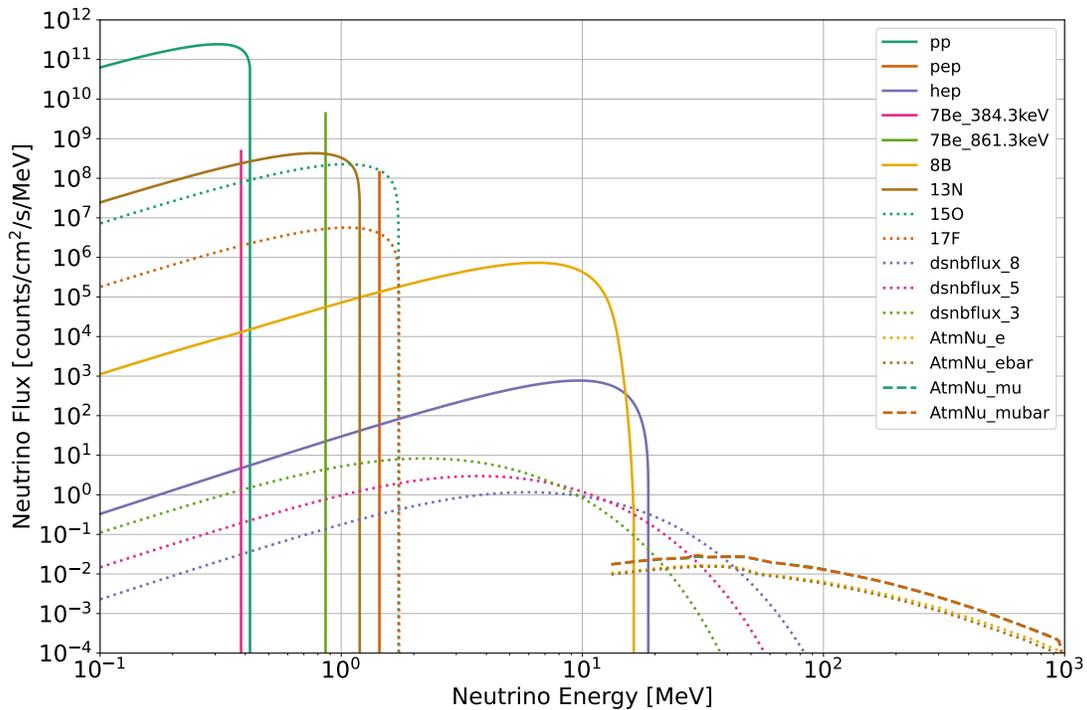


Figure 8.7: The fluxes of all background neutrino species we consider: solar, diffuse supernova, and atmospheric. Some species are monoenergetic, and others have an energy distribution. Fluxes taken from [31].

detector. However, the neutral current interaction is independent of the neutrino flavor or charge, so we consider the total neutrino flux.

Our second source of neutrinos is a diffuse background of neutrinos from all supernovas ever occurring in the observable universe. The third is atmospheric neutrinos. These are produced when cosmic rays collide with atoms in the Earth’s atmosphere, producing showers of pions, kaons, and muons, which decay weakly and emit neutrinos.

The flux of all three of these sources are shown in Fig. 8.7. Supernova neutrinos are broken into three fluxes based on flavor: $T_{\nu_e} = 3$ MeV, $T_{\bar{\nu}_e} = 5$ MeV, $T_{\nu_{other}} = 8$ MeV. Atmospheric neutrinos are split based on flavor and charge. In this Chapter, solar neutrinos are the primary background we are concerned about; in fact, even within that category, the dominant background is from the lowest-energy pp neutrinos. For xenon-based experiments, atmospheric neutrinos and the highest-energy solar neutrinos from ^8B are more problematic. Note we have neglected two neutrino sources: reactor neutrinos, which have a minor flux as long as the experiment is not located near a nuclear reactor, and the thermal Cosmic Neutrino Background, which are too low-energy to produce observable recoils.

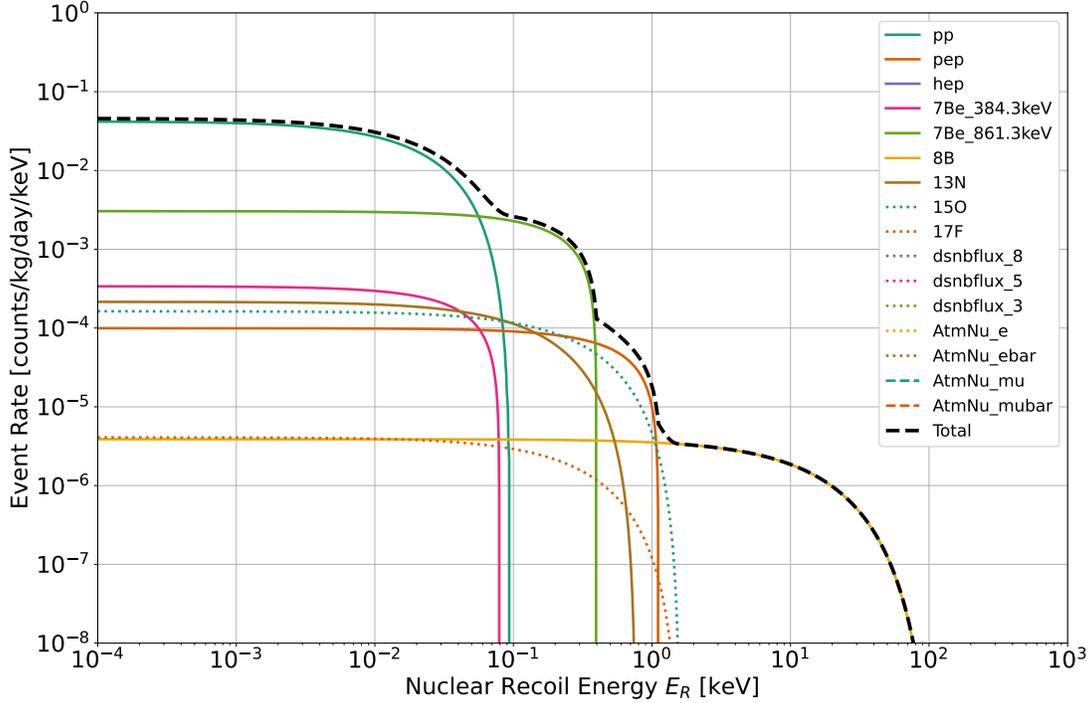


Figure 8.8: The rate of coherent elastic neutrino scattering on ${}^4\text{He}$ nuclei, for each neutrino species separately and summed. Below 10^{-4} keV, the rate is flat with energy.

From Billard et al. [31], the cross-section for a neutrino of energy E_ν to produce a recoil E_R is given below:

$$\frac{d\sigma(E_\nu, E_R)}{dE_R} = \frac{G_F^2}{4\pi} Q_w^2 m_N \left(1 - \frac{m_N E_R}{2E_\nu^2}\right) F^2(q^2) \quad (8.1)$$

$$Q_w = N - (1 - 4\sin^2\theta_w)Z$$

Here, G_F is the Fermi constant, m_N is the nuclear mass, θ_w is the weak mixing angle, and N and Z are the number of neutrons and protons in the nucleus. We use the same Helm nuclear form factor that we used for WIMP recoils. Then, the rate of interactions is given by integrating the neutrino flux dN/dE_ν .

$$\frac{dR}{dE_R} = \mathcal{N} \int_{E_\nu^{min}} \frac{dN}{dE_\nu} \frac{d\sigma(E_\nu, E_R)}{dE_R} dE_\nu \quad (8.2)$$

Here, \mathcal{N} is the density of target nuclei per unit mass, and $E_\nu^{min} = \sqrt{\frac{m_N E_R}{2}}$ is the minimum neutrino energy to produce a E_R recoil. The recoil rate for a ${}^4\text{He}$ target nucleus is plotted in Fig. 8.8.

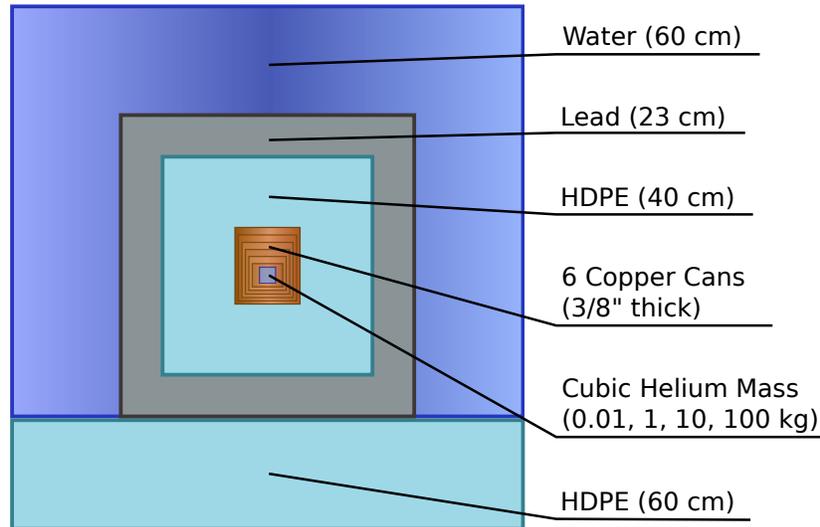


Figure 8.9: Detector geometry simulated to determine radiogenic γ backgrounds. Each component is cylindrically symmetric except for the ${}^4\text{He}$ volume, which is cubic. Four ${}^4\text{He}$ volumes were considered, ranging from 0.01–100 kg; the copper can dimensions were updated accordingly. Figure made by A. Biekert, published in [176].

The other backgrounds in our model are all radiogenic gammas. Unlike in xenon, where γ backgrounds induce only electronic interactions, HeRALD is also susceptible to coherent γ backgrounds. These are very low-energy interactions where a γ scatters off an atomic nucleus, but given the thresholds we can achieve with HeRALD technology, they are relevant. In fact, below 19.8 eV, the minimum energy to excite a helium atom, all recoils are necessarily nuclear.

Simulations are employed to calculate the expected level of γ backgrounds in HeRALD². Gamma rates are determined from the ${}^{238}\text{U}$, ${}^{232}\text{Th}$, ${}^{40}\text{K}$, ${}^{60}\text{Co}$, and ${}^{137}\text{Cs}$ decay chains, using detector impurity concentrations from SuperCDMS SNOLAB [24]. The detector geometry, in turn, is simulated as shown in Fig. 8.9. From innermost to outermost, the shielding is copper, high density polyethylene (HDPE), lead, water, and more HDPE. Using GEANT4, γ 's are simulated and propagated, and the flux in the helium volume is recorded. Four ${}^4\text{He}$ volumes with masses of 0.01, 1, 10, and 100 kg are simulated, representing successive stages of HeRALD. The flux per unit mass of γ photons in the detector volume decreases with increased ${}^4\text{He}$ mass. Finally, the background rate is calculated using Eq. 8.2, replacing the neutrino energy with the gamma energy, and replacing the neutrino-nucleus cross-section with the sum of the gamma-nucleus and gamma-electron cross-sections. The

²The γ and neutron simulations and background calculations presented in this Section were done by A. Biekert.

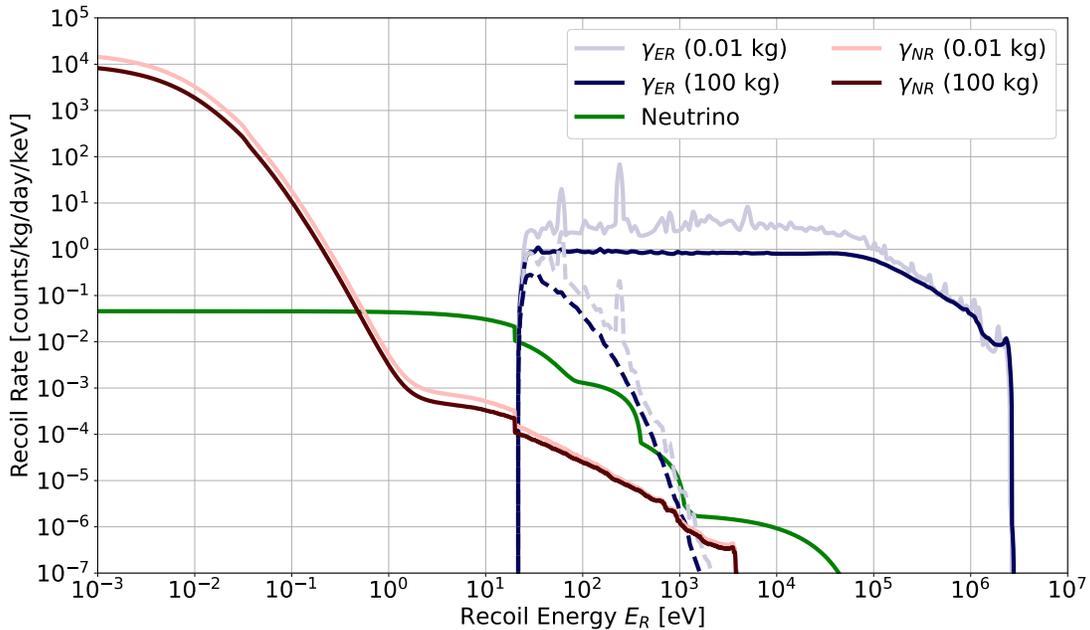


Figure 8.10: Predicted background rates in HeRALD from neutrino NRs, and from radiogenic γ ERs and NRs. We assume detector masses of 0.1 and 100 kg for the γ backgrounds. ER background rates are shown before (solid) and after (dashed) discrimination. NR background rates are multiplied by 0.5 above 19.8 eV, the discrimination threshold, to indicate 50% acceptance.

γ -e interactions considered are Compton scattering and photoabsorption, and the coherent γ -N interactions considered are Rayleigh, nuclear Thomson, and Delbrück scattering.

We do not consider neutron backgrounds in our model. This background was estimated using LZ studies of the Sanford Underground Research Facility, but the rate was found to be negligible.

The background rates vs. recoil energy from neutrinos and γ 's are shown in Fig. 8.10. Above the minimum excitation threshold of 19.8 meV, we show the ER rates before and after discrimination, and we show the NR rates at 50% acceptance.

8.3 HeRALD Sensitivity to Nuclear Recoils

Finally, we calculate projected sensitivities of the HeRALD detector to dark matter nuclear recoils, using a profile likelihood ratio (PLR) statistical test. We use the signal model from Sec. 2.4.2 for a pure ${}^4\text{He}$ detector and the background model from Sec. 8.2.

We summarize the main steps in calculating the projected 90% sensitivity below; the interested reader can also reference Cowan et al. [190]. This procedure works for any statistical test.

1. Model the detector with relevant parameters—for us, the ${}^4\text{He}$ mass, run time, and energy threshold. Also choose the observable quantities—for us, number of events observed and recoil energy of each.
2. Run $\mathcal{O}(1000)$ Monte Carlo simulations of the detector in which no dark matter is present, and only backgrounds are present. Record the observable(s) for each interaction.
3. For each dark matter mass m_χ of interest:
 - a) For each DM-nucleon cross-section $\sigma_{\chi-n}$ of interest:
 - i. Run $\mathcal{O}(1000)$ Monte Carlo simulations of the detector in which the given dark matter model is present. Record the observable(s) for each interaction.
 - ii. Calculate the test statistic q for each DM + background simulation. Note q is a function of the null hypothesis H_0 , which is the existence of DM model $(m_\chi, \sigma_{\chi-n})$.
 - iii. Calculate q for each background-only simulation.
 - iv. Determine q_{med} , the median value of q from the background-only simulations.
 - v. Compare q_{med} to the distribution of q from the DM + background simulations. Calculate the fraction of q that are higher than q_{med} . This is the p-value p .
 - b) Scan the $(\sigma_{\chi-n}, p)$ pairs. Interpolate to find $\sigma_{\chi-n}$ such that $p = 0.1$. Call it σ_{90} ; this is the 90% sensitivity at m_χ .
4. The pairs (m_χ, σ_{90}) are the projected sensitivity.

The PLR likelihood function that we use is

$$\mathcal{L}(m_\chi, \sigma_{\chi-n}) = \frac{e^{-(\mu_\chi + \sum_j \mu_j)}}{N!} \times \prod_{i=1}^N \left[\mu_\chi f_\chi(E_{r_i}) + \sum_j \mu_j f_j(E_{r_i}) \right], \quad (8.3)$$

where i iterates over observed events, j iterates over different background species, μ_χ (μ_j) is the expected number of signal (background) events, N is the total number of observed events, and f_χ (f_j) is the signal (background) recoil energy probability distribution function. The corresponding test statistic is

$$q = -2 \ln \frac{\mathcal{L}(m_\chi, \sigma_{\chi-n})}{\mathcal{L}(m_\chi, \hat{\sigma})}. \quad (8.4)$$

The numerator represents the null hypothesis H_0 , a particular DM model $(m_\chi, \sigma_{\chi-n})$. The denominator represents the alternate hypothesis H_1 , in which m_χ is fixed, but $\sigma_{\chi-n}$ is whatever value maximizes \mathcal{L} .

The detector is simulated as a pure ^4He target with 100% efficiency to measure all recoil energies above some threshold. We consider four generations of experiments with threshold-mass-runtime combinations of (40 eV, 10 g, 100 days), (10 eV, 1 kg, 1 year), (0.1 eV, 10 kg, 1 year), and (1 meV, 100 kg, 1 year). The first generation experiment is “shovel ready” in that it combines several already-demonstrated technologies with no required new R&D: a calorimeter of 3.5 eV baseline resolution [191], a 5% efficiency of converting recoil energy to evaporation [187], and a $9\times$ adsorption gain of He atoms on Si [187]. Energy threshold can be improved with advancements in TES sensing, our pursuit in Chapter 7. The fourth-generation experiment represents a theoretical limit for this technology, when we can measure individual quantum evaporated He atoms. The resulting sensitivities are shown in Fig. 8.11.

Yet these four projections are not the whole story. We can extend our reach to lower-mass dark matter using recently developed techniques. One of these is nuclear bremsstrahlung, a rare occurrence in which a photon is emitted from the nucleus during a collision. It was demonstrated in [197] that the energy carried by the photon can be significantly greater than the nuclear recoil energy.

$$E_{R,max} = 4 \left(\frac{m_\chi}{m_N} \right) E_{\gamma,max} \quad (8.5)$$

The authors of [197] also demonstrated how to calculate the interaction rate for this process. We show this rate in helium in Fig. 8.12 for a 10 MeV/ c^2 particle, comparing it to elastic scattering with the same mass and cross-section³. It is clear to see that although the rate is highly suppressed, the energy range is substantially higher. For a detector threshold of 40 eV, a 10 MeV/ c^2 dark matter particle would be unobservable with elastic scattering, but it would be visible using the bremsstrahlung signal.

Therefore, we also calculate the sensitivity of HeRALD to these interactions. We use the same PLR procedure, but with one crucial difference: although we are still probing nuclear recoils, our signal appears as a photon, in the electronic recoil region. As a result, we are forced to sacrifice ER/NR discrimination. Figure 8.11 shows the results as upper extensions to the four projections (the 4th-generation bremsstrahlung curve is not shown because it is cut off by the Earth shielding curve). Although bremsstrahlung interactions are even rarer than DM elastic recoils and

³Calculations in helium done by J. Lin.

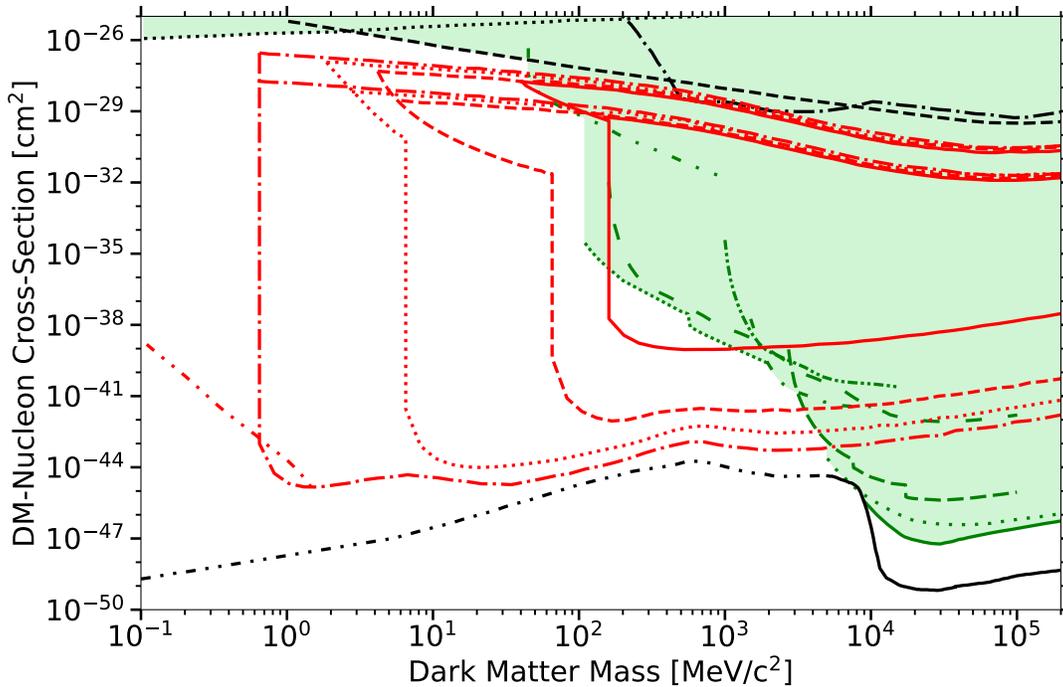


Figure 8.11: Projected HeRALD sensitivity to DM-nucleon spin-independent interactions with 90% confidence, through detection of elastic scattering and bremsstrahlung emission. Four combinations of exposure and energy threshold have been investigated: 1 kg-day with 40 eV (solid red), 1 kg-yr with 10 eV (dashed red), 10 kg-yr with 0.1 eV (dotted red), and 100 kg-yr with 1 meV (dash-dotted red). The sensitivity added by bremsstrahlung emission is visible in the extensions of the curves to lower mass at high cross section ($\sigma_{\text{SI}} > 10^{-32} \text{ cm}^2$). Two maximum cross-section constraints from Earth shielding are shown for each energy threshold, corresponding to depths of 100 m and 1478 m. The dash-double-dotted red curve includes off-shell phonon processes, scaling from Knapen et al. [192] to a 100 kg-yr exposure, assuming a massive mediator and 1 meV energy threshold. The black lines at the bottom of the figure are the neutrino fog; the solid line is from Billard et al. in xenon [31], and the dash-double-dotted line is from this Chapter in helium. The green lines correspond to constraints from various direct detection experiments: solid, LZ [26]; dense dashed, XENON1T S2-Only [41]; sparse dashed CRESST [36]; dense dotted XENON1T Migdal [39]; loose dotted PandaX-4T [32]; dot-dashed CDMSlite [38]; dense dot-dot-dashed DAMIC [37]; sparse dot-dot-dashed EDELWEISS Migdal [42]. The black lines at the top of the figure correspond to constraints from astronomy: dashed, galactic gas cooling [193]; dotted, CMB [194, 195]; dash-dotted, XQC [196]. Excluded parameter space is shaded light green.

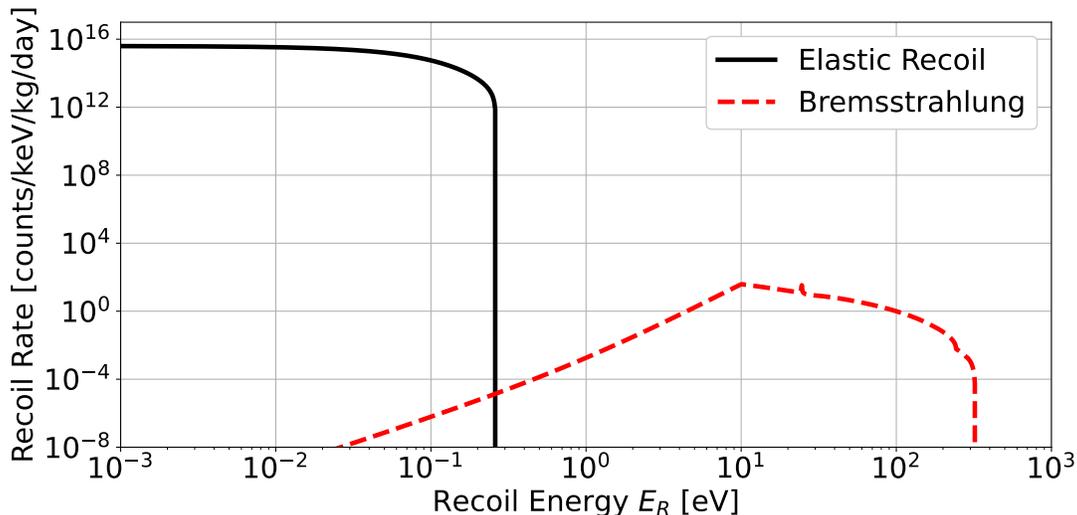


Figure 8.12: The recoil rate for $10 \text{ MeV}/c^2$ dark matter on ${}^4\text{He}$ nuclei, assuming a 10^{-30} cm^2 cross-section, for elastic recoils and nuclear bremsstrahlung.

appear in the high-background ER region, we still gain access to a lot of parameter space by using them. We remark that LUX and XENON1T have successfully used this approach in xenon to extend their reach to lower masses [126, 39].

Another projection shown in Fig. 8.11 is the sensitivity to keV/ c^2 -scale dark matter that couples directly to quasiparticle excitations. Knapen et al. showed that the production of multiple high-momentum excitations in superfluid helium allows experiments to probe DM at these masses, even if an experiment’s energy threshold is higher than the energy deposited by elastic nuclear recoils [192]. We show a version of Knapen et al.’s calculation, scaled to our energy threshold and exposure.

Finally, we have constraints from the opposite direction—the maximum cross-sections we can probe. Any successful version of HeRALD must be built deep underground to protect against pileup from cosmic muons. If the DM-nucleon interaction has a sufficiently high cross section, the DM flux and kinetic energy are degraded as DM particles propagate through the Earth. This upper limit was determined through a random walk simulation, based on the procedure in [198]⁴. Details are available in [176], but briefly: DM particles are simulated at the surface of the Earth with initial downward velocities; they travel through the Earth, scattering off terrestrial nuclei with some probability; and they either escape or reach the underground site with some speed. The upper $\sigma_{\chi-n}$ is chosen when the average DM speed at the detector is equal to $\sqrt{2E_{th}/m_\chi}$ for a detector threshold E_{th} . This procedure was done for two depths—100 m and 1478 m—for the four energy thresh-

⁴Calculation done by J. Lin.

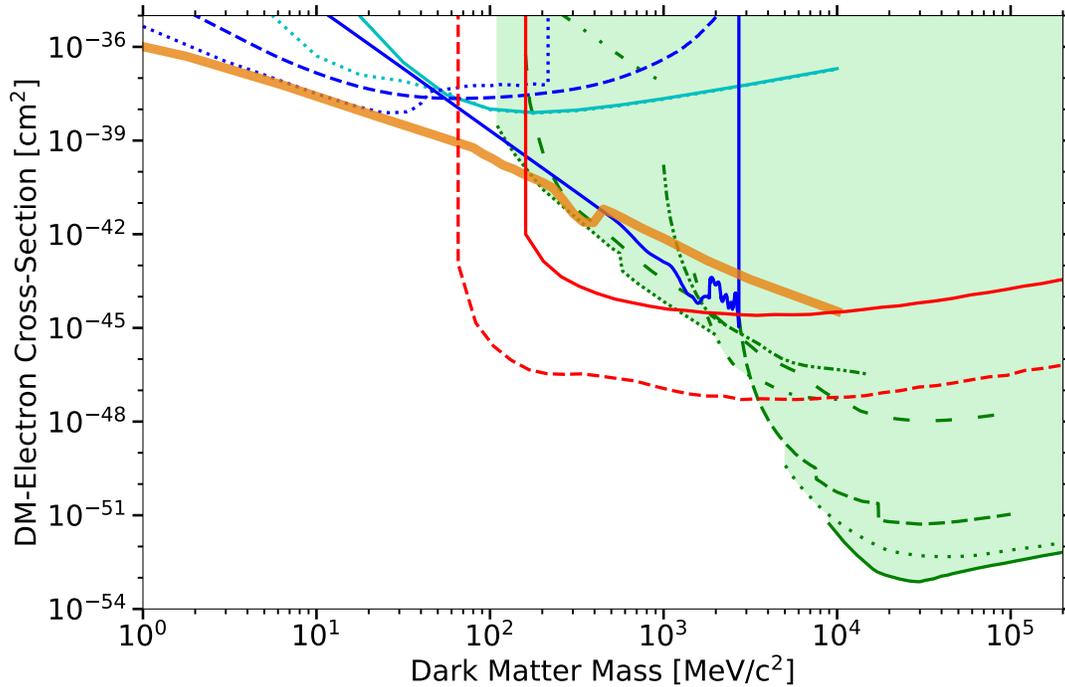


Figure 8.13: Projected HeRALD sensitivity to DM-electron scattering via a heavy dark photon, converting $\bar{\sigma}_n$ sensitivity to $\bar{\sigma}_e$ sensitivity via Eq. 8.6. Two combinations of exposure and energy threshold are displayed: 1 kg-day with 40 eV threshold (solid red) and 1 kg-yr with 10 eV threshold (dashed red). The orange band shows the parameter space for which a complex scalar freeze-out gives the correct relic abundance [199]. Existing constraints from nuclear recoil experiments are shown as green lines, having been similarly converted into the DM-electron space. The line styles are the same as in Fig. 8.11. Constraints from BaBar (solid blue), E137 (dashed blue), and LSND (dotted blue) are shown, converted into the $\bar{\sigma}_e$ space in [199, 200]. Finally, direct constraints on the DM-electron cross section from XENON10 (dotted cyan) and XENON100 (solid cyan) are shown [200, 201]. Excluded parameter space is shaded light green.

olds corresponding to our generations of HeRALD. They are depicted in Fig. 8.11, showing that these cross-sections are incredibly high and not a pressing concern for HeRALD’s reach.

8.4 HeRALD Sensitivity to Electronic Recoils

In the case of a heavy dark photon mediator ($F_{DM} = 1$), the DM-nucleon and DM-electron scattering cross sections are related:

$$\frac{\bar{\sigma}_e}{\sigma_{\chi-n}} = \left(\frac{A}{Z}\right)^2 \left(\frac{\mu_{\chi-e}}{\mu_{\chi-n}}\right)^2, \quad (8.6)$$

where A and Z are the atomic mass number and atomic number of the target nucleus, and $\mu_{\chi-e}$ ($\mu_{\chi-n}$) is the reduced mass between the DM and an electron (nucleon). Thus, we can translate our projected sensitivities into DM-electron space and compare to existing constraints on dark photon interactions. This is done in Fig. 8.13, with current nuclear recoil constraints translated into the $\bar{\sigma}_e$ plane using Eq. 8.6. Note that we have not translated the sensitivities for our third and fourth generation experiments into $\bar{\sigma}_e$ parameter space. In the case of a dark photon-mediated nuclear recoil in helium, the photon propagator is modified by in-medium effects. These effects are negligible for recoil energies greater than 10 eV [202], legitimizing the translation of our first and second generation sensitivities. However, if they are significant for lower recoil energies, Equation 8.6 will not hold. Further work needs to be done to determine whether the other sensitivities can be similarly translated, based on a detailed calculation of the in-helium photon propagator.

8.5 The Neutrino Fog in Helium

Finally, we return to the “neutrino fog,” introduced in Chapter 2. This is the region in parameter space when CEvNS becomes a limiting background for direct detection. As mentioned, the neutrino fog is fundamentally about our ignorance in the flux of neutrino species, not only the absolute rate of neutrino-induced nuclear recoils. We now calculate the neutrino fog in helium using a technique similar to [31]. Recall that for helium targets, solar neutrinos are the primary obstacle; atmospheric and supernova neutrinos are not a dark matter background, so we ignore them.

To start, the coherent elastic scattering of solar neutrinos on helium nuclei is considered to be the only background in a hypothetical detector. A new PLR test statistic is defined by extending the likelihood function to include uncertainties in neutrino flux as “nuisance parameters.”

$$\begin{aligned} \mathcal{L}(\sigma_{\chi-n}, \phi) &= \frac{e^{-(\mu_\chi + \sum_j \mu_j)}}{N!} \\ &\times \prod_{i=1}^N \left[\mu_\chi f_\chi(E_{r_i}) + \sum_j \mu_j f_j(E_{r_i}) \right] \times \prod_j \mathcal{G}_j(\phi_j, \Delta\phi_j). \end{aligned} \quad (8.7)$$

Here, \mathcal{G}_j is a Gaussian distribution centered on the mean value of the flux of neutrino species j , and with standard deviation given by the uncertainty in that flux. Nuisance parameters give us a way to express the uncertainty in the parameters of our model. The neutrino species we consider and their associated uncertainties are: pp (1%), pep (1.7%), ${}^7\text{Be}$ (10.5%), ${}^8\text{B}$ (8.8%), hep (15.5%), and CNO (30%). These values are a combination of theoretical [203, 204], experimental [205], and estimated [31] uncertainties.

CEvNS becomes a limiting, irreducible background for dark matter experiments when the exposure is high enough that flux uncertainties rival Poisson fluctuations. For recoil energies less than about 1 keV, the dominant neutrino species are pp and ${}^7\text{Be}$. The geometric mean of their uncertainties is 3.2%, indicating that the solar neutrino background becomes significant at an exposure corresponding to about 980 expected recoils. We thus define the neutrino fog as the projected sensitivity of a ${}^4\text{He}$ detector with 1.6 tonne-yr exposure, for which the expected number of CEvNS events is 1000. The threshold is set arbitrarily low.

This calculation of the neutrino fog is shown in Fig. 8.11, in the black dot-dot-dashed line. We see that the neutrino fog falls away at the lowest masses, a consequence of the fact that the pp recoil energy spectrum is flat at the lowest recoil energies (see Fig. 8.8). The bumps around 500 MeV/ c^2 are a consequence of the fact that the DM recoil energy spectrum at these masses looks quite similar to the pp and ${}^7\text{Be}$ spectra. The reader can also observe that the neutrino fog in helium lines up smoothly with the neutrino fog in xenon calculated by Billard et al. [31]. This is a coincidence. Different elements aren't required to have the same neutrino fog—in fact, above 10s of GeV/ c^2 , the neutrino fog in helium is a much higher cross-section than in xenon. Regardless, the neutrino fog in any target gives us an expectation for our ultimate reach.

Chapter 9

Conclusion

When I began my PhD program in August 2015, the nature of dark matter was one of the primary questions facing particle physicists. Today, in August 2022, this is still true.

Yet, physicists have made great achievements in that time. Over a dozen direct detection experiments have been operated, producing ever-stronger constraints on dark matter models. The 2013 run of LUX was the world-leading limit then, rejecting WIMP spin-independent scattering with cross-sections above $6.6 \times 10^{-46} \text{ cm}^2$ at a mass of $33 \text{ GeV}/c^2$ [62]. As of July 2022, LUX-ZEPLIN has inherited the title, rejecting cross-sections of above $5.9 \times 10^{-48} \text{ cm}^2$ at a mass of $30 \text{ GeV}/c^2$ [26]. In 2017, the COHERENT collaboration measured coherent elastic neutrino-nucleus scattering [189]—long known to be a dark matter background, but confirmed for the first time. The pursuit for sub-GeV dark matter kicked into high gear, with several cutting-edge technologies being deployed to operate dark matter searches: advanced calorimetry [36, 38, 40, 42], charge coupled devices [37], and new signatures in technologies designed for GeV–TeV-scale WIMPs [39, 126].

The work in this dissertation should be seen as a piece of the same journey. The analyses of ER/NR discrimination in xenon have already motivated the 300 V/cm design operating field of LZ. They will be crucial to building a G3-scale detector by giving quantitative predictions for the tradeoff between drift field, g_1 , cost, construction materials, and other design considerations. This will allow us to reach the neutrino fog at the WIMP mass scale. Meanwhile, TESSERACT fits cleanly into the sub-GeV agenda of the past several years. Our measurements of various tungsten films have already proven the need to use a xenon chamber gas instead of argon when fabricating TES's, in order to enhance sensitivity.

Much work remains to be done. LZ and its peers, XENONnT and PandaX-4T, will run for several years to achieve the best possible sensitivity to spin-independent and spin-dependent nuclear recoils, electron recoils, and other interactions. Discussions are ongoing to decide whether to build a G3-scale (100-tonne) xenon TPC that could finally reach the neutrino fog. There are some outstanding questions about energy deposition in xenon: What is the true value of W ? Why are recombination fluctuations non-binomial, and why are they skewed? Do β and γ particles of the same energy really produce different light and charge yields? However, the most important questions are about design, and the current 10-tonne scale experiments will inform those.

The TESSERACT (and calorimetry more broadly) agenda is rife with unanswered questions. Our measurements did not establish the ideal fabrication procedure for tungsten films, nor did they even establish tungsten as the superior element. Significant efforts will need to be made (and are being made) to enhance TES and calorimetry technology. We need to understand the dynamics of energy partitioning in our three target materials; work is already underway for this project in helium [206, 207]. Future work will require measuring pulse shapes, quantum evaporation efficiency, quasiparticle propagation and loss, and more. Eventually, the gallium

arsenide, sapphire, and helium dark matter experiments will have to be designed and built.

Perhaps a first-year PhD student reading this will graduate in a world where dark matter has been detected. Perhaps that student is a freshman undergrad. Perhaps they are a freshman high schooler, although I certainly hope not. Regardless, the search for dark matter will continue until we finally set our eyes on the truth.

Bibliography

- [1] P. A. R. Ade et al. “Planck 2015 results. XIII. Cosmological parameters”. In: *Astron. Astrophys.* 594 (2016), A13. DOI: 10.1051/0004-6361/201525830. arXiv: 1502.01589 [astro-ph.CO].
- [2] F. Zwicky. “Die Rotverschiebung von extragalaktischen Nebeln”. In: *Helv. Phys. Acta* 6 (1933), pp. 110–127. DOI: 10.1007/s10714-008-0707-4.
- [3] M. Milgrom. “A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis.” In: *apj* 270 (July 1983), pp. 365–370. DOI: 10.1086/161130.
- [4] Vera C. Rubin, W. Kent Ford Jr., and Norbert Thonnard. “Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa through Sc”. In: *Astrophys. J. Lett.* 225 (1978), pp. L107–L111. DOI: 10.1086/182804.
- [5] Wayne Hu. *Dark Matter Density*. 2001. URL: <http://background.uchicago.edu/~whu/intermediate/driving2.html> (visited on 06/20/2022).
- [6] P. A. Zyla et al. “Review of Particle Physics”. In: *PTEP* 2020.8 (2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.
- [7] Wayne Hu. *Intermediate Guide to the Acoustic Peaks and Polarization*. 2001. URL: <http://background.uchicago.edu/~whu/intermediate/intermediate.html> (visited on 06/20/2022).
- [8] Douglas Clowe et al. “A Direct Empirical Proof of the Existence of Dark Matter”. In: *The Astrophysical Journal* 648.2 (Aug. 2006), pp. L109–L113. DOI: 10.1086/508162. URL: <https://doi.org/10.1086/508162>.
- [9] C. Alcock et al. “MACHO Project Limits on Black Hole Dark Matter in the 1–30 [ITAL]M[/ITAL][TINF]⊙[/TINF] Range”. In: *The Astrophysical Journal* 550.2 (Apr. 2001), pp. L169–L172. DOI: 10.1086/319636. URL: <https://doi.org/10.1086/319636>.

- [10] Tisserand, P. et al. “Limits on the Macho content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds ***”. In: *A&A* 469.2 (2007), pp. 387–404. DOI: 10.1051/0004-6361:20066017. URL: <https://doi.org/10.1051/0004-6361:20066017>.
- [11] Xavier Hernández et al. “Scalar field “mini-MACHOs”: A new explanation for galactic dark matter”. In: *Phys. Rev. D* 70 (4 Aug. 2004), p. 043537. DOI: 10.1103/PhysRevD.70.043537. URL: <https://link.aps.org/doi/10.1103/PhysRevD.70.043537>.
- [12] Timothy D. Brandt. “CONSTRAINTS ON MACHO DARK MATTER FROM COMPACT STELLAR SYSTEMS IN ULTRA-FAINT DWARF GALAXIES”. In: *The Astrophysical Journal* 824.2 (June 2016), p. L31. DOI: 10.3847/2041-8205/824/2/131. URL: <https://doi.org/10.3847/2041-8205/824/2/131>.
- [13] S. D. M. White, C. S. Frenk, and M. Davis. “Clustering in a neutrino-dominated universe”. In: *The Astrophysical Journal* 274 (Nov. 1983), pp. L1–L5. DOI: 10.1086/184139.
- [14] E. O. Nadler et al. “Constraints on Dark Matter Properties from Observations of Milky Way Satellite Galaxies”. In: *Phys. Rev. Lett.* 126 (9 Mar. 2021), p. 091101. DOI: 10.1103/PhysRevLett.126.091101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.126.091101>.
- [15] Gerard Jungman, Marc Kamionkowski, and Kim Griest. “Supersymmetric dark matter”. In: *Phys. Rept.* 267 (1996), pp. 195–373. DOI: 10.1016/0370-1573(95)00058-5. arXiv: hep-ph/9506380.
- [16] Clifford Cheung et al. “Prospects and Blind Spots for Neutralino Dark Matter”. In: *JHEP* 05 (2013), p. 100. DOI: 10.1007/JHEP05(2013)100. arXiv: 1211.4873 [hep-ph].
- [17] Miguel Escudero et al. “Toward (Finally!) Ruling Out Z and Higgs Mediated Dark Matter Models”. In: *JCAP* 12 (2016), p. 029. DOI: 10.1088/1475-7516/2016/12/029. arXiv: 1609.09079 [hep-ph].
- [18] Tomohiro Abe, Motoko Fujiwara, and Junji Hisano. “Loop corrections to dark matter direct detection in a pseudoscalar mediator dark matter model”. In: *JHEP* 02 (2019), p. 028. DOI: 10.1007/JHEP02(2019)028. arXiv: 1810.01039 [hep-ph].
- [19] Benjamin W. Lee and Steven Weinberg. “Cosmological Lower Bound on Heavy Neutrino Masses”. In: *Phys. Rev. Lett.* 39 (1977). Ed. by M. A. Srednicki, pp. 165–168. DOI: 10.1103/PhysRevLett.39.165.
- [20] Kim Griest and Marc Kamionkowski. “Unitarity Limits on the Mass and Radius of Dark Matter Particles”. In: *Phys. Rev. Lett.* 64 (1990), p. 615. DOI: 10.1103/PhysRevLett.64.615.

- [21] Tongyan Lin. “Dark matter models and direct detection”. In: *PoS TASI2018* (2019), p. 009. DOI: 10.22323/1.333.0009.
- [22] D. S. Akerib et al. “Effective field theory analysis of the first LUX dark matter search”. In: *Phys. Rev. D* 103.12 (2021), p. 122005. DOI: 10.1103/PhysRevD.103.122005. arXiv: 2003.11141 [astro-ph.CO].
- [23] D. S. Akerib et al. “Constraints on effective field theory couplings using 311.2 days of LUX data”. In: *Phys. Rev. D* 104.6 (2021), p. 062005. DOI: 10.1103/PhysRevD.104.062005. arXiv: 2102.06998 [astro-ph.CO].
- [24] R. Agnese et al. “Projected sensitivity of the SuperCDMS SNOLAB experiment”. In: *Phys. Rev. D* 95 (8 Apr. 2017), p. 082002. DOI: 10.1103/PhysRevD.95.082002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.95.082002>.
- [25] P. Agnes et al. “First results from the DarkSide-50 dark matter experiment at Laboratori Nazionali del Gran Sasso”. In: *Physics Letters B* 743 (Apr. 2015), pp. 456–466. ISSN: 0370-2693. DOI: 10.1016/J.PHYSLETB.2015.03.012. URL: <https://www.sciencedirect.com/science/article/pii/S0370269315001756>.
- [26] J. Aalbers et al. “First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment”. In: (July 2022). arXiv: 2207.03764 [hep-ex].
- [27] J. D. Lewin and P. F. Smith. “Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil”. In: *Astropart. Phys.* 6 (1996), pp. 87–112. DOI: 10.1016/S0927-6505(96)00047-3.
- [28] David G. Cerdeno and Anne M. Green. “Direct detection of WIMPs”. In: (Feb. 2010), pp. 347–369. DOI: 10.1017/CB09780511770739.018. arXiv: 1002.1912 [astro-ph.CO].
- [29] Richard W. Schnee. “Introduction to dark matter experiments”. In: *Theoretical Advanced Study Institute in Elementary Particle Physics: Physics of the Large and the Small*. 2011, pp. 775–829. DOI: 10.1142/9789814327183_0014. arXiv: 1101.5205 [astro-ph.CO].
- [30] Lawrence M. Widrow, Brent Pym, and John Dubinski. “Dynamical Blueprints for Galaxies”. In: *Astrophys. J.* 679 (2008), pp. 1239–1259. DOI: 10.1086/587636. arXiv: 0801.3414 [astro-ph].
- [31] J. Billard, L. Strigari, and E. Figueroa-Feliciano. “Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments”. In: *Phys. Rev. D* 89.2 (2014), p. 023524. DOI: 10.1103/PhysRevD.89.023524. arXiv: 1307.5458 [hep-ph].

- [32] Yue Meng et al. “Dark Matter Search Results from the PandaX-4T Commissioning Run”. In: *Phys. Rev. Lett.* 127.26 (2021), p. 261802. DOI: 10.1103/PhysRevLett.127.261802. arXiv: 2107.13438 [hep-ex].
- [33] D. S. Akerib et al. “Snowmass2021 Cosmic Frontier Dark Matter Direct Detection to the Neutrino Fog”. In: *2022 Snowmass Summer Study*. Mar. 2022. arXiv: 2203.08084 [hep-ex].
- [34] S. E. Vahsen et al. “CYGNUS: Feasibility of a nuclear recoil observatory with directional sensitivity to dark matter and neutrinos”. In: (Aug. 2020). arXiv: 2008.12587 [physics.ins-det].
- [35] D. Reusser et al. “Limits on cold dark matter from the Gotthard Ge experiment”. In: *Phys. Lett. B* 255 (1991), pp. 143–145. DOI: 10.1016/0370-2693(91)91155-0.
- [36] A. H. Abdelhameed et al. “First results from the CRESST-III low-mass dark matter program”. In: *Phys. Rev. D* 100.10 (2019), p. 102002. DOI: 10.1103/PhysRevD.100.102002. arXiv: 1904.00498 [astro-ph.CO].
- [37] A. Aguilar-Arevalo et al. “Results on low-mass weakly interacting massive particles from a 11 kg-day target exposure of DAMIC at SNOLAB”. In: *Phys. Rev. Lett.* 125 (2020), p. 241803. DOI: 10.1103/PhysRevLett.125.241803. arXiv: 2007.15622 [astro-ph.CO].
- [38] R. Agnese et al. “Search for Low-Mass Dark Matter with CDMSlite Using a Profile Likelihood Fit”. In: *Phys. Rev. D* 99.6 (2019), p. 062001. DOI: 10.1103/PhysRevD.99.062001. arXiv: 1808.09098 [astro-ph.CO].
- [39] E. Aprile et al. “Search for Light Dark Matter Interactions Enhanced by the Migdal Effect or Bremsstrahlung in XENON1T”. In: *Physical Review Letters* 123.24 (Dec. 2019), p. 241803. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.123.241803. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.241803>.
- [40] G. Angloher et al. “Results on MeV-scale dark matter from a gram-scale cryogenic calorimeter operated above ground”. In: *Eur. Phys. J. C* 77.9 (2017), p. 637. DOI: 10.1140/epjc/s10052-017-5223-9. arXiv: 1707.06749 [astro-ph.CO].
- [41] E. Aprile et al. “Light Dark Matter Search with Ionization Signals in XENON1T”. In: *Physical Review Letters* 123.25 (Dec. 2019), p. 251801. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.123.251801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.251801>.
- [42] E. Armengaud et al. “Searching for low-mass dark matter particles with a massive Ge bolometer operated above-ground”. In: *Phys. Rev. D* 99.8 (2019), p. 082003. DOI: 10.1103/PhysRevD.99.082003. arXiv: 1901.03588 [astro-ph.GA].

- [43] E. Aprile et al. “Dark Matter Search Results from a One Ton-Year Exposure of XENON1T”. In: *Physical Review Letters* 121.11 (Sept. 2018), p. 111302. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.121.111302. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.111302>.
- [44] John H. Marshall. “A Liquid Argon Ionization Chamber Measurement of the Shape of the Beta-Ray Spectrum of K^{40} ”. In: *Phys. Rev.* 91 (4 Aug. 1953), pp. 905–909. DOI: 10.1103/PhysRev.91.905. URL: <https://link.aps.org/doi/10.1103/PhysRev.91.905>.
- [45] Richard A. Muller et al. “Liquid-Filled Proportional Counter”. In: *Phys. Rev. Lett.* 27 (8 Aug. 1971), pp. 532–535. DOI: 10.1103/PhysRevLett.27.532. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.27.532>.
- [46] Claude Brassard. “Liquid ionization detectors”. In: *Nuclear Instruments and Methods* 162.1 (1979), pp. 29–47. ISSN: 0029-554X. DOI: [https://doi.org/10.1016/0029-554X\(79\)90705-5](https://doi.org/10.1016/0029-554X(79)90705-5). URL: <https://www.sciencedirect.com/science/article/pii/0029554X79907055>.
- [47] Tadayoshi Doke. “A historical view on the R&D for liquid rare gas detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 327.1 (1993), pp. 113–118. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/0168-9002\(93\)91423-K](https://doi.org/10.1016/0168-9002(93)91423-K). URL: <https://www.sciencedirect.com/science/article/pii/016890029391423K>.
- [48] P. Benetti et al. “Detection of energy deposition down to the keV region using liquid xenon scintillation”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 327.1 (1993), pp. 203–206. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/0168-9002\(93\)91442-P](https://doi.org/10.1016/0168-9002(93)91442-P). URL: <https://www.sciencedirect.com/science/article/pii/016890029391442P>.
- [49] I. IVANIOUCHENKOV et al. “ZEPLIN-I: A SINGLE PHASE LIQUID XE DETECTOR FOR DARK MATTER SEARCH”. In: *The Identification of Dark Matter*. 2001, pp. 433–437. DOI: 10.1142/9789812811363_0055. URL: https://www.worldscientific.com/doi/abs/10.1142/9789812811363_0055.
- [50] G.J. Alner et al. “First limits on nuclear recoil events from the ZEPLIN I galactic dark matter detector”. In: *Astroparticle Physics* 23.5 (July 2005), pp. 444–462. ISSN: 0927-6505. DOI: 10.1016/j.astropartphys.2005.02.004. URL: <https://www.sciencedirect.com/science/article/pii/S0927650505000289>.

- [51] David Cline et al. “Measurement of Scintillation Amplification in a Xenon Detector with a CsI-Luminescence Plate”. In: *Sources and Detection of Dark Matter and Dark Energy in the Universe*. Ed. by David B. Cline. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 485–492. ISBN: 978-3-662-04587-9.
- [52] E. APRILE et al. “XENON: A 1 TONNE LIQUID XENON EXPERIMENT FOR A SENSITIVE DARK MATTER SEARCH”. In: *Technique and Application of Xenon Detectors*. 2003, pp. 165–178. DOI: 10.1142/9789812705075_0015. URL: https://www.worldscientific.com/doi/abs/10.1142/9789812705075_0015.
- [53] E. Aprile et al. “The XENON dark matter search experiment”. In: *New Astronomy Reviews* 49.2 (2005). Sources and Detection of Dark Matter and Dark Energy in the Universe, pp. 289–295. ISSN: 1387-6473. DOI: <https://doi.org/10.1016/j.newar.2005.01.035>. URL: <https://www.sciencedirect.com/science/article/pii/S1387647305000473>.
- [54] M. Yamashita et al. “Double phase (liquid/gas) xenon scintillation detector for WIMPs direct search”. In: *Astroparticle Physics* 20.1 (2003), pp. 79–84. ISSN: 0927-6505. DOI: [https://doi.org/10.1016/S0927-6505\(03\)00174-9](https://doi.org/10.1016/S0927-6505(03)00174-9). URL: <https://www.sciencedirect.com/science/article/pii/S0927650503001749>.
- [55] G.J. Alner et al. “First limits on WIMP nuclear recoil signals in ZEPLIN-II: A two-phase xenon detector for dark matter detection”. In: *Astroparticle Physics* 28.3 (2007), pp. 287–302. ISSN: 0927-6505. DOI: <https://doi.org/10.1016/j.astropartphys.2007.06.002>. URL: <https://www.sciencedirect.com/science/article/pii/S0927650507000928>.
- [56] J. Angle et al. “First Results from the XENON10 Dark Matter Experiment at the Gran Sasso National Laboratory”. In: *Phys. Rev. Lett.* 100 (2 Jan. 2008), p. 021303. DOI: 10.1103/PhysRevLett.100.021303. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.100.021303>.
- [57] V. N. Lebedenko et al. “Results from the first science run of the ZEPLIN-III dark matter search experiment”. In: *Physical Review D* 80.5 (Sept. 2009), p. 052010. ISSN: 1550-7998. DOI: 10.1103/PhysRevD.80.052010. URL: <https://link.aps.org/doi/10.1103/PhysRevD.80.052010>.
- [58] D. Yu. Akimov et al. “WIMP-nucleon cross-section results from the second science run of ZEPLIN-III”. In: *Phys. Lett. B* 709 (2012), pp. 14–20. DOI: 10.1016/j.physletb.2012.01.064. arXiv: 1110.4769 [astro-ph.CO].
- [59] E. Aprile et al. “First Dark Matter Results from the XENON100 Experiment”. In: *Phys. Rev. Lett.* 105 (2010), p. 131302. DOI: 10.1103/PhysRevLett.105.131302. arXiv: 1005.0380 [astro-ph.CO].

- [60] E. Aprile et al. “Dark Matter Results from 100 Live Days of XENON100 Data”. In: *Phys. Rev. Lett.* 107 (2011), p. 131302. DOI: 10 . 1103 / PhysRevLett.107.131302. arXiv: 1104.2549 [astro-ph.CO].
- [61] E. Aprile et al. “Dark Matter Results from 225 Live Days of XENON100 Data”. In: *Phys. Rev. Lett.* 109 (2012), p. 181301. DOI: 10 . 1103 / PhysRevLett.109.181301. arXiv: 1207.5988 [astro-ph.CO].
- [62] D. S. Akerib et al. “First Results from the LUX Dark Matter Experiment at the Sanford Underground Research Facility”. In: *Physical Review Letters* 112.9 (Mar. 2014), p. 091303. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.112.091303. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.112.091303>.
- [63] D. S. Akerib et al. “Improved Limits on Scattering of Weakly Interacting Massive Particles from Reanalysis of 2013 LUX Data”. In: *Physical Review Letters* 116.16 (Apr. 2016), p. 161301. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.116.161301. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.161301>.
- [64] D. S. Akerib et al. “Results from a Search for Dark Matter in the Complete LUX Exposure”. In: *Physical Review Letters* 118.2 (Jan. 2017), p. 021303. ISSN: 0031-9007. DOI: 10 . 1103 / PhysRevLett . 118 . 021303. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.021303>.
- [65] Mengjiao Xiao et al. “First dark matter search results from the PandaX-I experiment”. In: *Sci. China Phys. Mech. Astron.* 57 (2014), pp. 2024–2030. DOI: 10.1007/s11433-014-5598-7. arXiv: 1408.5114 [hep-ex].
- [66] Xiang Xiao et al. “Low-mass dark matter search results from full exposure of the PandaX-I experiment”. In: *Phys. Rev. D* 92.5 (2015), p. 052004. DOI: 10.1103/PhysRevD.92.052004. arXiv: 1505.00771 [hep-ex].
- [67] Xiangyi Cui et al. “Dark Matter Results from 54-Ton-Day Exposure of PandaX-II Experiment”. In: *Physical Review Letters* 119.18 (Oct. 2017), p. 181302. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.119.181302. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.181302>.
- [68] E. Aprile et al. “Search for New Physics in Electronic Recoil Data from XENONnT”. In: (July 2022). arXiv: 2207.11330 [hep-ex].
- [69] E. Aprile et al. “Excess electronic recoil events in XENON1T”. In: *Phys. Rev. D* 102.7 (2020), p. 072004. DOI: 10 . 1103 / PhysRevD . 102 . 072004. arXiv: 2006.09721 [hep-ex].
- [70] D. S. Akerib et al. “First Searches for Axions and Axionlike Particles with the LUX Experiment”. In: *Phys. Rev. Lett.* 118.26 (2017), p. 261301. DOI: 10.1103/PhysRevLett.118.261301. arXiv: 1704.02297 [astro-ph.CO].

- [71] D. S. Akerib et al. “First direct detection constraint on mirror dark matter kinetic mixing using LUX 2013 data”. In: *Phys. Rev. D* 101.1 (2020), p. 012003. DOI: 10.1103/PhysRevD.101.012003. arXiv: 1908.03479 [hep-ex].
- [72] E. Aprile et al. “Observation of two-neutrino double electron capture in ^{124}Xe with XENON1T”. In: *Nature* 568.7753 (2019), pp. 532–535. DOI: 10.1038/s41586-019-1124-4. arXiv: 1904.11002 [nucl-ex].
- [73] L. W. Kastens et al. “Calibration of a Liquid Xenon Detector with Kr-83m”. In: *Phys. Rev. C* 80 (2009), p. 045809. DOI: 10.1103/PhysRevC.80.045809. arXiv: 0905.1766 [physics.ins-det].
- [74] A. Manzur et al. “Scintillation efficiency and ionization yield of liquid xenon for mono-energetic nuclear recoils down to 4 keV”. In: *Phys. Rev. C* 81 (2010), p. 025808. DOI: 10.1103/PhysRevC.81.025808. arXiv: 0909.1063 [physics.ins-det].
- [75] A.G. Singh et al. “Analysis of 83m Kr prompt scintillation signals in the PIXeY detector”. In: *Journal of Instrumentation* 15.01 (Jan. 2020), P01023–P01023. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/01/P01023. URL: <https://iopscience.iop.org/article/10.1088/1748-0221/15/01/P01023>.
- [76] Carl Eric Dahl. “The physics of background discrimination in liquid xenon, and first results from Xenon10 in the hunt for WIMP dark matter.” PhD thesis. Princeton University, 2009. URL: <https://inspirehep.net/record/1374815>.
- [77] L. W. Goetzke et al. “Measurement of light and charge yield of low-energy electronic recoils in liquid xenon”. In: *Phys. Rev. D* 96.10 (2017), p. 103007. DOI: 10.1103/PhysRevD.96.103007. arXiv: 1611.10322 [astro-ph.IM].
- [78] Laura Baudis et al. “A Dual-phase Xenon TPC for Scintillation and Ionisation Yield Measurements in Liquid Xenon”. In: *Eur. Phys. J. C* 78.5 (2018), p. 351. DOI: 10.1140/epjc/s10052-018-5801-5. arXiv: 1712.08607 [astro-ph.IM].
- [79] Robert Scherrer. *The Wacky Names that Scientists Give Their Experiments*. Cosmic Yarns. Aug. 2015. URL: <http://www.cosmicyarns.com/2015/08/the-wacky-names-that-scientists-give.html> (visited on 06/04/2022).
- [80] Janelle Shane. *Physics acronyms: there was an attempt*. AI Weirdness. Apr. 2019. URL: <https://www.aiweirdness.com/physics-acronyms-there-was-an-attempt-19-04-12/> (visited on 06/04/2022).
- [81] Philip Bump. *All the Silly Legislative Acronyms Congress Came Up with This Year*. The Atlantic. Aug. 2013. URL: <https://www.theatlantic.com/politics/archive/2013/08/congress-acronyms-reins/312565/> (visited on 06/04/2022).

- [82] E. Aprile et al. “Projected WIMP Sensitivity of the XENONnT Dark Matter Experiment”. In: *JCAP* 11 (2020), p. 031. DOI: 10.1088/1475-7516/2020/11/031.
- [83] J. Aalbers et al. “DARWIN: towards the ultimate dark matter detector”. In: *JCAP* 11 (2016), p. 017. DOI: 10.1088/1475-7516/2016/11/017. arXiv: 1606.07001 [astro-ph.IM].
- [84] J. Aalbers et al. “A Next-Generation Liquid Xenon Observatory for Dark Matter and Neutrino Physics”. In: (Mar. 2022). arXiv: 2203.02309 [physics.ins-det].
- [85] HongGuang Zhang et al. “Dark matter direct search sensitivity of the PandaX-4T experiment”. In: *Science China Physics, Mechanics & Astronomy* 62.3 (Mar. 2019), p. 31011. ISSN: 1674-7348. DOI: 10.1007/s11433-018-9259-0. URL: <http://link.springer.com/10.1007/s11433-018-9259-0>.
- [86] D. S. Akerib et al. “Projected WIMP sensitivity of the LUX-ZEPLIN dark matter experiment”. In: *Physical Review D* 101.5 (Mar. 2020), p. 052002. ISSN: 2470-0010. DOI: 10.1103/PhysRevD.101.052002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.052002>.
- [87] L. Tvrznikova et al. “Direct comparison of high voltage breakdown measurements in liquid argon and liquid xenon”. In: *JINST* 14.12 (2019), P12018. DOI: 10.1088/1748-0221/14/12/P12018. arXiv: 1908.06888 [physics.ins-det].
- [88] Lucie Tvrznikova. *Liquid xenon bubbles*. Youtube. July 2018. URL: <https://www.youtube.com/watch?v=Zy9r8q1wmYc> (visited on 01/10/2022).
- [89] Lucie Tvrznikova. *Xice*. Youtube. Sept. 2018. URL: <https://www.youtube.com/watch?v=vFYziAxx95w> (visited on 01/10/2022).
- [90] Juris Meija et al. “Atomic weights of the elements 2013 (IUPAC Technical Report)”. In: *Pure and Applied Chemistry* 88.3 (2016), pp. 265–291. DOI: doi:10.1515/pac-2015-0305. URL: <https://doi.org/10.1515/pac-2015-0305>.
- [91] Marc Schumann. “Dark Matter Search with liquid Noble Gases”. In: (June 2012). arXiv: 1206.2169 [astro-ph.IM].
- [92] Sujoy Mukhopadhyay. “Xenon Isotopes”. In: *Encyclopedia of Geochemistry: A Comprehensive Reference Source on the Chemistry of the Earth*. Ed. by William M. White. Cham: Springer International Publishing, 2018, pp. 1500–1508. ISBN: 978-3-319-39312-4. DOI: 10.1007/978-3-319-39312-4_203. URL: https://doi.org/10.1007/978-3-319-39312-4_203.
- [93] J. B. Albert et al. “Searches for double beta decay of ^{134}Xe with EXO-200”. In: *Phys. Rev. D* 96.9 (2017), p. 092001. DOI: 10.1103/PhysRevD.96.092001. arXiv: 1704.05042 [hep-ex].

- [94] Keiko Fujii et al. “High-accuracy measurement of the emission spectrum of liquid xenon in the vacuum ultraviolet region”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 795 (2015), pp. 293–297. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2015.05.065>. URL: <https://www.sciencedirect.com/science/article/pii/S016890021500724X>.
- [95] D. S. Akerib et al. “Liquid xenon scintillation measurements and pulse shape discrimination in the LUX dark matter detector”. In: *Physical Review D* 97.11 (July 2018), p. 112002. ISSN: 2470-0010. DOI: 10.1103/PhysRevD.97.112002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.97.112002>.
- [96] D. S. Akerib et al. “Signal yields, energy resolution, and recombination fluctuations in liquid xenon”. In: *Physical Review D* 95.1 (Jan. 2017), p. 012008. ISSN: 2470-0010. DOI: 10.1103/PhysRevD.95.012008. URL: <https://link.aps.org/doi/10.1103/PhysRevD.95.012008>.
- [97] G. Anton et al. “Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment”. In: *Phys. Rev. C* 101.6 (2020), p. 065501. DOI: 10.1103/PhysRevC.101.065501. arXiv: 1908.04128 [physics.ins-det].
- [98] Laura Baudis, Patricia Sanchez-Lucas, and Kevin Thieme. “A measurement of the mean electronic excitation energy of liquid xenon”. In: *Eur. Phys. J. C* 81.12 (2021), p. 1060. DOI: 10.1140/epjc/s10052-021-09834-x. arXiv: 2109.07151 [physics.ins-det].
- [99] M. Szydagus et al. *Noble Element Simulation Technique*. Version v2.3.0. Nov. 2021. DOI: 10.5281/zenodo.5676553. URL: <https://doi.org/10.5281/zenodo.5676553>.
- [100] Vetri Velan. *vvelan/Xenimation: v1.0.0*. Version v1.0. June 2022. DOI: 10.5281/zenodo.6615730. URL: <https://doi.org/10.5281/zenodo.6615730>.
- [101] Tadayoshi Doke et al. “Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles”. In: *Japanese Journal of Applied Physics* 41.Part 1, No. 3A (Mar. 2002), pp. 1538–1545. ISSN: 0021-4922. DOI: 10.1143/JJAP.41.1538. URL: <http://stacks.iop.org/1347-4065/41/1538>.
- [102] E. Aprile et al. “Observation of anticorrelation between scintillation and ionization for MeV gamma rays in liquid xenon”. In: *Physical Review B* 76.1 (July 2007), p. 014115. ISSN: 1098-0121. DOI: 10.1103/PhysRevB.76.014115. URL: <https://link.aps.org/doi/10.1103/PhysRevB.76.014115>.

- [103] Qing Lin et al. “Scintillation and ionization responses of liquid xenon to low energy electronic and nuclear recoils at drift fields from 236 V / cm to 3.93 kV / cm”. In: *Physical Review D* 92.3 (Aug. 2015), p. 032005. ISSN: 1550-7998. DOI: 10.1103/PhysRevD.92.032005. URL: <https://link.aps.org/doi/10.1103/PhysRevD.92.032005>.
- [104] O. Piazza, V. Velan, and D. McKinsey. *A first-principles approach to electron-ion recombination in liquid xenon*. LIDINE: LIght Detection In Noble Elements. 2021. URL: <https://indico.physics.ucsd.edu/event/1/contributions/33/>.
- [105] K. Ueshima et al. “Scintillation-only based pulse shape discrimination for nuclear and electron recoils in liquid xenon”. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2011). ISSN: 01689002. DOI: 10.1016/j.nima.2011.09.011. arXiv: 1106.2209.
- [106] Tadayoshi Doke et al. “Let dependence of scintillation yields in liquid argon”. In: *Nuclear Inst. and Methods in Physics Research, A* (1988). ISSN: 01689002. DOI: 10.1016/0168-9002(88)90892-3.
- [107] M Szydagis et al. “NEST: a comprehensive model for scintillation yield in liquid xenon”. In: *Journal of Instrumentation* 6.10 (Oct. 2011), P10002–P10002. ISSN: 1748-0221. DOI: 10.1088/1748-0221/6/10/P10002. URL: <http://stacks.iop.org/1748-0221/6/i=10/a=P10002?key=crossref.5f8848204896d7ad9ca1ddbcbaac32d9>.
- [108] J. Thomas and D. A. Imel. “Recombination of electron-ion pairs in liquid argon and liquid xenon”. In: *Physical Review A* 36.2 (July 1987), pp. 614–616. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.36.614. URL: <https://link.aps.org/doi/10.1103/PhysRevA.36.614>.
- [109] Jon Balajthy. “Purity Monitoring Techniques and Electronic Energy Deposition Properties in Liquid Xenon Time Projection Chambers”. PhD thesis. University of Maryland, 2018. URL: <https://drum.lib.umd.edu/handle/1903/20824>.
- [110] *ESTAR: CSDA Range for Xenon*. National Institute of Standards and Technology. URL: <https://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html> (visited on 06/06/2022).
- [111] Brian Lenardo et al. “A Global Analysis of Light and Charge Yields in Liquid Xenon”. In: *IEEE Transactions on Nuclear Science* 62.6 (Dec. 2015), pp. 3387–3396. ISSN: 0018-9499. DOI: 10.1109/TNS.2015.2481322. URL: <http://ieeexplore.ieee.org/document/7317823/>.

- [112] Peter Sorensen and Carl Eric Dahl. “Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter”. In: *Physical Review D* 83.6 (Mar. 2011), p. 063501. ISSN: 1550-7998. DOI: 10.1103/PhysRevD.83.063501. URL: <https://link.aps.org/doi/10.1103/PhysRevD.83.063501>.
- [113] P Peiffer et al. “Pulse shape analysis of scintillation signals from pure and xenon-doped liquid argon for radioactive background identification”. In: *Journal of Instrumentation* 3.08 (Aug. 2008), P08007–P08007. ISSN: 1748-0221. DOI: 10.1088/1748-0221/3/08/P08007. URL: <http://stacks.iop.org/1748-0221/3/i=08/a=P08007?key=crossref.c3b7007b577d2afc1e20d58d284e54b2>.
- [114] J. Lindhard et al. “Integral Equations Governing Radiation Effects”. In: *Mat. Fys. Medd. Dan. Vid. Selsk.* 33.10 (1963). URL: <http://gymarkiv.sdu.dk/MFM/kdvs/mfm%2030-39/mfm-33-10.pdf>.
- [115] D. S. Akerib et al. “Low-energy (0.7-74 keV) nuclear recoil calibration of the LUX dark matter experiment using D-D neutron scattering kinematics”. In: (Aug. 2016). arXiv: 1608.05381. URL: <https://arxiv.org/abs/1608.05381>.
- [116] A. Manalaysay et al. “Spatially uniform calibration of a liquid xenon detector at low energies using 83m-Kr”. In: *Rev. Sci. Instrum.* 81 (2010), p. 073303. DOI: 10.1063/1.3436636. arXiv: 0908.0616 [astro-ph.IM].
- [117] E. Aprile et al. “Measurement of the scintillation yield of low-energy electrons in liquid xenon”. In: *Physical Review D* 86.11 (Dec. 2012), p. 112004. ISSN: 1550-7998. DOI: 10.1103/PhysRevD.86.112004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.86.112004>.
- [118] L. Baudis et al. “Response of liquid xenon to Compton electrons down to 1.5 keV”. In: *Physical Review D* 87.11 (July 2013), p. 115015. ISSN: 1550-7998. DOI: 10.1103/PhysRevD.87.115015. URL: <https://link.aps.org/doi/10.1103/PhysRevD.87.115015>.
- [119] B. J. Mount et al. “LUX-ZEPLIN (LZ) Technical Design Report”. In: (Mar. 2017). arXiv: 1703.09144. URL: <http://arxiv.org/abs/1703.09144>.
- [120] V. N. Solovov et al. “Measurement of the refractive index and attenuation length of liquid xenon for its scintillation light”. In: *Nucl. Instrum. Meth. A* 516 (2004), pp. 462–474. DOI: 10.1016/j.nima.2003.08.117. arXiv: physics/0307044.
- [121] F. Neves et al. “Measurement of the absolute reflectance of polytetrafluoroethylene (PTFE) immersed in liquid xenon”. In: *JINST* 12.01 (2017), P01017. DOI: 10.1088/1748-0221/12/01/P01017. arXiv: 1612.07965 [physics.ins-det].

- [122] S. Kravitz et al. “Measurements of Angle-Resolved Reflectivity of PTFE in Liquid Xenon with IBEX”. In: *Eur. Phys. J. C* 80.3 (2020), p. 262. DOI: 10.1140/epjc/s10052-020-7800-6. arXiv: 1909.08730 [physics.ins-det].
- [123] C.H. Faham et al. “Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes”. In: *Journal of Instrumentation* 10.09 (Sept. 2015), P09010–P09010. ISSN: 1748-0221. DOI: 10.1088/1748-0221/10/09/P09010. URL: <http://stacks.iop.org/1748-0221/10/i=09/a=P09010?key=crossref.b436ccc6e15acae238afeff0ecbd6d97>.
- [124] D. S. Akerib et al. “Extending light WIMP searches to single scintillation photons in LUX”. In: *Physical Review D* 101.4 (Feb. 2020), p. 042001. ISSN: 2470-0010. DOI: 10.1103/PhysRevD.101.042001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.042001>.
- [125] D. S. Akerib et al. “Calibration, event reconstruction, data analysis, and limit calculation for the LUX dark matter experiment”. In: *Physical Review D* 97.10 (May 2018), p. 102008. ISSN: 2470-0010. DOI: 10.1103/PhysRevD.97.102008. URL: <https://link.aps.org/doi/10.1103/PhysRevD.97.102008>.
- [126] D. S. Akerib et al. “Results of a Search for Sub-GeV Dark Matter Using 2013 LUX Data”. In: *Physical Review Letters* 122.13 (Apr. 2019), p. 131301. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.122.131301. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.131301>.
- [127] D.S. Akerib et al. “3D modeling of electric fields in the LUX detector”. In: *Journal of Instrumentation* 12.11 (Nov. 2017), P11022–P11022. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/11/P11022. URL: <http://stacks.iop.org/1748-0221/12/i=11/a=P11022?key=crossref.c93e348558314091a48936d33c09e70b>.
- [128] V.N. Solovov et al. “Position Reconstruction in a Dual Phase Xenon Scintillation Detector”. In: *IEEE Trans. Nucl. Sci.* 59 (2012). Ed. by Ken Label, pp. 3286–3293. DOI: 10.1109/TNS.2012.2221742. arXiv: 1112.1481 [physics.ins-det].
- [129] E. Bodnia et al. “The electric field dependence of single electron emission in the PIXeY two-phase xenon detector”. In: *JINST* 16.12 (2021), P12015. DOI: 10.1088/1748-0221/16/12/P12015. arXiv: 2101.03686 [physics.ins-det].
- [130] D. S. Akerib et al. “Investigation of background electron emission in the LUX detector”. In: *Phys. Rev. D* 102.9 (2020), p. 092004. DOI: 10.1103/PhysRevD.102.092004. arXiv: 2004.07791 [physics.ins-det].
- [131] NEST Collaboration. *Noble Element Simulation Technique*. URL: <https://nest.physics.ucdavis.edu/> (visited on 06/04/2022).

- [132] M. Szydagi et al. *Noble Element Simulation Technique*. Version v2.3.7. May 2022. DOI: 10.5281/zenodo.6534007. URL: <https://doi.org/10.5281/zenodo.6534007>.
- [133] T. Doke et al. “Estimation of Fano factors in liquid argon, krypton, xenon and xenon-doped liquid argon”. In: *Nucl. Instrum. Meth.* 134 (1976), pp. 353–357. DOI: 10.1016/0029-554X(76)90292-5.
- [134] D. S. Akerib et al. “Discrimination of electronic recoils from nuclear recoils in two-phase xenon time projection chambers”. In: *Phys. Rev. D* 102.11 (2020), p. 112002. DOI: 10.1103/PhysRevD.102.112002. arXiv: 2004.06304 [physics.ins-det].
- [135] Chloe Reynolds. *Decay Chains*. Nuclear Forensic Search Project. URL: <https://metadata.berkeley.edu/nuclear-forensics/Decay%5C%20Chains.html> (visited on 06/04/2022).
- [136] G. Anton et al. “Search for Neutrinoless Double- β Decay with the Complete EXO-200 Dataset”. In: *Phys. Rev. Lett.* 123.16 (2019), p. 161802. DOI: 10.1103/PhysRevLett.123.161802. arXiv: 1906.02723 [hep-ex].
- [137] D.S. Akerib et al. “Radiogenic and muon-induced backgrounds in the LUX dark matter detector”. In: *Astroparticle Physics* 62 (Mar. 2015), pp. 33–46. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2014.07.009. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650514001054>.
- [138] D. S. Akerib et al. “The Large Underground Xenon (LUX) experiment”. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2013). ISSN: 01689002. DOI: 10.1016/j.nima.2012.11.135.
- [139] D. S. Akerib et al. “Tritium calibration of the LUX dark matter experiment”. In: *Physical Review D* 93.7 (Apr. 2016), p. 072009. ISSN: 2470-0010. URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.072009>.
- [140] D. S. Akerib et al. “Improved measurements of the β -decay response of liquid xenon with the LUX detector”. In: *Physical Review D* 100.2 (July 2019), p. 022002. ISSN: 2470-0010. DOI: 10.1103/PhysRevD.100.022002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.022002>.
- [141] J.R. Verbus et al. “Proposed low-energy absolute calibration of nuclear recoils in a dual-phase noble element TPC using D-D neutron scattering kinematics”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 851 (Apr. 2017), pp. 68–81. ISSN: 0168-9002. DOI: 10.1016/J.NIMA.2017.01.053. URL: <https://www.sciencedirect.com/science/article/pii/S0168900217301158?via%7B%5C%7D3Dihub>.

- [142] Lucie Tvrznikova. “Sub-GeV Dark Matter Searches and Electric Field Studies for the LUX and LZ Experiments”. PhD thesis. Yale University, Apr. 2019. arXiv: 1904.08979. URL: <http://arxiv.org/abs/1904.08979>.
- [143] E. Aprile et al. “Simultaneous Measurement of Ionization and Scintillation from Nuclear Recoils in Liquid Xenon for a Dark Matter Experiment”. In: *Physical Review Letters* 97.8 (Aug. 2006), p. 081302. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.97.081302. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.97.081302>.
- [144] E. Aprile et al. “Signal yields of keV electronic recoils and their discrimination from nuclear recoils in liquid xenon”. In: *Physical Review D* 97.9 (May 2018), p. 092007. ISSN: 2470-0010. DOI: 10.1103/PhysRevD.97.092007. URL: <https://link.aps.org/doi/10.1103/PhysRevD.97.092007>.
- [145] A. Azzalini. “A Class of Distributions Which Includes the Normal Ones”. In: *Scandinavian Journal of Statistics* 12.2 (1985), pp. 171–178. ISSN: 03036898, 14679469. URL: <http://www.jstor.org/stable/4615982>.
- [146] D. S. Akerib et al. “Kr 83 m calibration of the 2013 LUX dark matter search”. In: *Physical Review D* 96.11 (Dec. 2017), p. 112009. ISSN: 2470-0010. DOI: 10.1103/PhysRevD.96.112009. URL: <https://link.aps.org/doi/10.1103/PhysRevD.96.112009>.
- [147] D.S. Akerib et al. “Improved modeling of β electronic recoils in liquid xenon using LUX calibration data”. In: *Journal of Instrumentation* 15.02 (Feb. 2020), T02007–T02007. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/02/T02007. URL: <https://iopscience.iop.org/article/10.1088/1748-0221/15/02/T02007>.
- [148] D. S. Akerib et al. “LUX-ZEPLIN (LZ) Conceptual Design Report”. In: (Sept. 2015). arXiv: 1509.02910. URL: <http://arxiv.org/abs/1509.02910>.
- [149] K. Abe et al. “A measurement of the scintillation decay time constant of nuclear recoils in liquid xenon with the XMASS-I detector”. In: *Journal of Instrumentation* 13.12 (Dec. 2018), P12032–P12032. ISSN: 1748-0221. DOI: 10.1088/1748-0221/13/12/P12032. URL: <http://stacks.iop.org/1748-0221/13/i=12/a=P12032?key=crossref.83f083b7c8b9480f36ccfd78bb23bce9>.
- [150] J. Kwong et al. “Scintillation pulse shape discrimination in a two-phase xenon time projection chamber”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 612.2 (Jan. 2010), pp. 328–333. ISSN: 01689002. DOI: 10.1016/j.nima.2009.10.106. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900209020579>.

- [151] D.Yu. Akimov et al. “Limits on inelastic dark matter from ZEPLIN-III”. In: *Physics Letters B* 692.3 (Aug. 2010), pp. 180–183. ISSN: 0370-2693. DOI: 10.1016/J.PHYSLETB.2010.07.042. URL: <https://www.sciencedirect.com/science/article/pii/S0370269310008944>.
- [152] J Mock et al. “Modeling pulse characteristics in Xenon with NEST”. In: *Journal of Instrumentation* 9.04 (Apr. 2014), T04002–T04002. ISSN: 1748-0221. DOI: 10.1088/1748-0221/9/04/T04002. URL: <http://stacks.iop.org/1748-0221/9/i=04/a=T04002?key=crossref.c6c8f99ce50f4a3f5e1e2a54b90a03d5>.
- [153] E. Aprile et al. “XENON1T dark matter data analysis: Signal and background models and statistical inference”. In: *Physical Review D* 99.11 (July 2019), p. 112009. ISSN: 2470-0010. DOI: 10.1103/PhysRevD.99.112009. URL: <https://link.aps.org/doi/10.1103/PhysRevD.99.112009>.
- [154] Matthew Szydagis et al. *Noble Element Simulation Technique v2.0.1*. Zenodo:3357973. Aug. 2019. DOI: 10.5281/zenodo.3357973. URL: <https://zenodo.org/record/3357973>.
- [155] M. Horn et al. “Nuclear recoil scintillation and ionisation yields in liquid xenon from ZEPLIN-III data”. In: *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* (2011). ISSN: 03702693. DOI: 10.1016/j.physletb.2011.10.038.
- [156] Henrique Araújo. “Revised performance parameters of the ZEPLIN-III dark matter experiment”. In: (July 2020). arXiv: 2007.01683. URL: <https://arxiv.org/abs/2007.01683>.
- [157] M. Szydagis et al. “Investigating the XENON1T Low-Energy Electronic Recoil Excess Using NEST”. In: (July 2020). arXiv: 2007.00528. URL: <https://arxiv.org/abs/2007.00528>.
- [158] E. M. Boulton et al. “Calibration of a two-phase xenon time projection chamber with a ^{37}Ar source”. In: *JINST* 12.08 (2017), P08004. DOI: 10.1088/1748-0221/12/08/P08004. arXiv: 1705.08958 [physics.ins-det].
- [159] B. N. V. Edwards et al. “Extraction efficiency of drifting electrons in a two-phase xenon time projection chamber”. In: *JINST* 13.01 (2018), P01005. DOI: 10.1088/1748-0221/13/01/P01005. arXiv: 1710.11032 [physics.ins-det].
- [160] D. S. Akerib et al. “Improving sensitivity to low-mass dark matter in LUX using a novel electrode background mitigation technique”. In: *Phys. Rev. D* 104.1 (2021), p. 012011. DOI: 10.1103/PhysRevD.104.012011. arXiv: 2011.09602 [hep-ex].

- [161] D. McKinsey et al. *The TESSERACT Dark Matter Project*. Snowmass2021 - Letter of Interest. 2020. URL: https://www.snowmass21.org/docs/files/summaries/CF/SNOWMASS21-CF1_CF2-IF1_IF8-120.pdf.
- [162] J. Åström et al. “Fracture Processes Observed with A Cryogenic Detector”. In: *Phys. Lett. A* 356 (2006), pp. 262–266. DOI: 10.1016/j.physleta.2006.03.059. arXiv: physics/0504151.
- [163] J. Barrier. *Principle of the dilution refrigerator*. 2019. URL: <https://julienbarrier.eu/blog/2019/08/14/dilution-fridge.html> (visited on 06/14/2022).
- [164] G. Batey and G. Teleberg. *Principles of dilution refrigeration: A brief technology guide*. Oxford Instruments. 2015. URL: https://nanoscience.oxinst.com/assets/uploads/NanoScience/Brochures/Principles%5C%20of%5C%20dilution%5C%20refrigeration_Sept15.pdf (visited on 06/14/2022).
- [165] *BF-LD-Series User Manual, Version 1.5.0*. BlueFors Cryogenics. Helsinki, Finland, 2016.
- [166] K.D. Irwin and G.C. Hilton. “Transition-Edge Sensors”. In: *Cryogenic Particle Detection*. Ed. by Christian Enss. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 63–150. ISBN: 978-3-540-31478-3. DOI: 10.1007/10933596_3. URL: https://doi.org/10.1007/10933596_3.
- [167] R. L. Fagaly. “Superconducting quantum interference device instruments and applications”. In: *Review of Scientific Instruments* 77.10 (2006), p. 101101. DOI: 10.1063/1.2354545. eprint: <https://doi.org/10.1063/1.2354545>. URL: <https://doi.org/10.1063/1.2354545>.
- [168] W. A. Little. “THE TRANSPORT OF HEAT BETWEEN DISSIMILAR SOLIDS AT LOW TEMPERATURES”. In: *Canadian Journal of Physics* 37.3 (1959), pp. 334–349. DOI: 10.1139/p59-037. eprint: <https://doi.org/10.1139/p59-037>. URL: <https://doi.org/10.1139/p59-037>.
- [169] Khalid Bouziane, M. Mamor, and F. Meyer. “DC magnetron sputtered tungsten: W film properties and electrical properties of W/Si Schottky diodes”. In: *Applied Physics A* 81 (June 2005), pp. 209–215. DOI: 10.1007/s00339-004-2558-5.
- [170] A.E. Lita et al. “Tuning of tungsten thin film superconducting transition temperature for fabrication of photon number resolving detectors”. In: *IEEE Transactions on Applied Superconductivity* 15.2 (2005), pp. 3528–3531. DOI: 10.1109/TASC.2005.849033.

- [171] Ying Yong Hu and Wei Min Huang. “Thermal Stress Analysis and Characterization of Thermomechanical Properties of Thin Films on an Elastic Substrate”. In: *Handbook of Manufacturing Engineering and Technology*. Ed. by Andrew Nee. London: Springer London, 2013, pp. 1–71. ISBN: 978-1-4471-4976-7. DOI: 10.1007/978-1-4471-4976-7_51-1. URL: https://doi.org/10.1007/978-1-4471-4976-7_51-1.
- [172] E. Grünwald et al. “Young’s Modulus and Poisson’s Ratio Characterization of Tungsten Thin Films Via Laser Ultrasound”. In: *Materials Today: Proceedings* 2.8 (2015). nanoFIS 2014 - Functional Integrated nanoSystems, pp. 4289–4294. ISSN: 2214-7853. DOI: <https://doi.org/10.1016/j.matpr.2015.09.015>. URL: <https://www.sciencedirect.com/science/article/pii/S2214785315007993>.
- [173] F. Schwartzberg and M. Knight. *Cryogenic Materials Data Handbook, Volume 1*. Technical Documentary Report AFML-TDR-64-280. Air Force Materials Laboratory, July 1970. URL: https://www.jlab.org/sites/default/files/magnet-group/materials/cryogenic_materials_data_handbook1.pdf.
- [174] R. J. Corruccini and J. J. Gniewek. *Thermal Expansion of Technical Solids at Low Temperatures: A Compilation From the Literature*. Monograph 29. National Bureau of Standards, May 1961. URL: <https://nvlpubs.nist.gov/nistpubs/Legacy/MONO/nbsmonograph29.pdf>.
- [175] Thomas Middelmann et al. “Thermal expansion coefficient of single-crystal silicon from 7 K to 293 K”. In: *Phys. Rev. B* 92 (17 Nov. 2015), p. 174113. DOI: 10.1103/PhysRevB.92.174113. URL: <https://link.aps.org/doi/10.1103/PhysRevB.92.174113>.
- [176] S. A. Hertel et al. “Direct detection of sub-GeV dark matter using a superfluid ^4He target”. In: *Phys. Rev. D* 100.9 (2019), p. 092007. DOI: 10.1103/PhysRevD.100.092007. arXiv: 1810.06283 [physics.ins-det].
- [177] D. N. McKinsey et al. “Radiative decay of the metastable $\text{He}_2(a^3\Sigma_u^+)$ molecule in liquid helium”. In: *Phys. Rev. A* 59 (1 Jan. 1999), pp. 200–204. DOI: 10.1103/PhysRevA.59.200. URL: <https://link.aps.org/doi/10.1103/PhysRevA.59.200>.
- [178] M. B. Sobnack, J. C. Inkson, and J. C. H. Fung. “Quasiparticle scattering at helium surfaces: A microscopic theory”. In: *Phys. Rev. B* 60 (5 Aug. 1999), pp. 3465–3475. DOI: 10.1103/PhysRevB.60.3465. URL: <https://link.aps.org/doi/10.1103/PhysRevB.60.3465>.

- [179] M.J. Buckingham and W.M. Fairbank. “Chapter III The Nature of the λ -Transition in Liquid Helium”. In: ed. by C.J. Gorter. Vol. 3. Progress in Low Temperature Physics. Elsevier, 1961, pp. 80–112. DOI: [https://doi.org/10.1016/S0079-6417\(08\)60134-1](https://doi.org/10.1016/S0079-6417(08)60134-1). URL: <https://www.sciencedirect.com/science/article/pii/S0079641708601341>.
- [180] R. J. Donnelly, J. A. Donnelly, and R. N. Hills. “Specific heat and dispersion curve for helium II”. In: *J. Low Temp. Phys.* 44.5 (Sept. 1981), pp. 471–489. ISSN: 1573-7357. DOI: [10.1007/BF00117839](https://doi.org/10.1007/BF00117839). URL: <https://doi.org/10.1007/BF00117839>.
- [181] Humphrey J. Maris. “Phonon-phonon interactions in liquid helium”. In: *Rev. Mod. Phys.* 49 (2 Apr. 1977), pp. 341–359. DOI: [10.1103/RevModPhys.49.341](https://doi.org/10.1103/RevModPhys.49.341). URL: <https://link.aps.org/doi/10.1103/RevModPhys.49.341>.
- [182] V. Narayanamurti and R. C. Dynes. “Roton propagation and phonon-roton scattering in He II”. In: *Phys. Rev. B* 13 (7 Apr. 1976), pp. 2898–2909. DOI: [10.1103/PhysRevB.13.2898](https://doi.org/10.1103/PhysRevB.13.2898). URL: <https://link.aps.org/doi/10.1103/PhysRevB.13.2898>.
- [183] M. B. Sobnack, J. C. Inkson, and J. C. H. Fung. “Quasiparticle scattering at helium surfaces: A microscopic theory”. In: *Phys. Rev. B* 60 (5 Aug. 1999), pp. 3465–3475. DOI: [10.1103/PhysRevB.60.3465](https://doi.org/10.1103/PhysRevB.60.3465). URL: <https://link.aps.org/doi/10.1103/PhysRevB.60.3465>.
- [184] I. V. Tanatarov et al. “Interactions of Phonons and Rotons with Interfaces in Superfluid Helium”. In: *J. Low Temp. Phys.* 159.5 (June 2010), pp. 549–575. ISSN: 1573-7357. DOI: [10.1007/s10909-010-0165-8](https://doi.org/10.1007/s10909-010-0165-8). URL: <https://doi.org/10.1007/s10909-010-0165-8>.
- [185] F.W. Carter et al. “Calorimetric Observation of Single He₂^{*} Excimers in a 100-mK He Bath”. In: *Journal of Low Temperature Physics* 186.3 (2017), pp. 183–196. ISSN: 1573-7357. DOI: [10.1007/s10909-016-1666-x](https://doi.org/10.1007/s10909-016-1666-x). URL: <http://dx.doi.org/10.1007/s10909-016-1666-x>.
- [186] M. Nava et al. “Superfluid State of 4He on Graphane and Graphene-Fluoride: Anisotropic Roton States”. In: *Journal of Low Temperature Physics* 171.5 (June 2013), pp. 699–710. ISSN: 1573-7357. DOI: [10.1007/s10909-012-0770-9](https://doi.org/10.1007/s10909-012-0770-9). URL: <https://doi.org/10.1007/s10909-012-0770-9>.
- [187] Simon R. Bandler. “Detection of charged particles in superfluid helium”. In: (Dec. 1994). URL: <https://www.osti.gov/biblio/10115891>.
- [188] Daniel Z. Freedman. “Coherent Neutrino Nucleus Scattering as a Probe of the Weak Neutral Current”. In: *Phys. Rev. D* 9 (1974), pp. 1389–1392. DOI: [10.1103/PhysRevD.9.1389](https://doi.org/10.1103/PhysRevD.9.1389).

- [189] D. Akimov et al. “Observation of Coherent Elastic Neutrino-Nucleus Scattering”. In: *Science* 357.6356 (2017), pp. 1123–1126. DOI: 10.1126/science.aao0990. arXiv: 1708.01294 [nucl-ex].
- [190] Glen Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *Eur. Phys. J. C* 71 (2011). [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)], p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [physics.data-an].
- [191] Matt Pyle. *Searching for 10meV-1GeV Dark Matter with Athermal Phonon Detectors*. <https://indico.fnal.gov/event/18104/session/22/contribution/154/material/slides/0.pptx>. New Technologies for Discovery IV: The 2018 CPAD Instrumentation Frontier Workshop. 2018.
- [192] Simon Knapen, Tongyan Lin, and Kathryn M. Zurek. “Light dark matter in superfluid helium: Detection with multi-excitation production”. In: *Phys. Rev. D* 95 (5 Mar. 2017), p. 056019. DOI: 10.1103/PhysRevD.95.056019. URL: <https://link.aps.org/doi/10.1103/PhysRevD.95.056019>.
- [193] Amit Bhoonah et al. “Calorimetric Dark Matter Detection With Galactic Center Gas Clouds”. In: *Phys. Rev. Lett.* 121.13 (2018), p. 131101. DOI: 10.1103/PhysRevLett.121.131101. arXiv: 1806.06857 [hep-ph].
- [194] Vera Gluscevic and Kimberly K. Boddy. “Constraints on Scattering of keV–TeV Dark Matter with Protons in the Early Universe”. In: *Phys. Rev. Lett.* 121.8 (2018), p. 081301. DOI: 10.1103/PhysRevLett.121.081301. arXiv: 1712.07133 [astro-ph.CO].
- [195] Kimberly K. Boddy and Vera Gluscevic. “First Cosmological Constraint on the Effective Theory of Dark Matter-Proton Interactions”. In: *Phys. Rev. D* 98.8 (2018), p. 083510. DOI: 10.1103/PhysRevD.98.083510. arXiv: 1801.08609 [astro-ph.CO].
- [196] Adrienne L. Erickcek et al. “Constraints on the Interactions between Dark Matter and Baryons from the X-ray Quantum Calorimetry Experiment”. In: *Phys. Rev. D* 76 (2007), p. 042007. DOI: 10.1103/PhysRevD.76.042007. arXiv: 0704.0794 [astro-ph].
- [197] Chris Kouvaris and Josef Pradler. “Probing sub-GeV Dark Matter with conventional detectors”. In: *Phys. Rev. Lett.* 118.3 (2017), p. 031803. DOI: 10.1103/PhysRevLett.118.031803. arXiv: 1607.01789 [hep-ph].
- [198] Timon Emken, Chris Kouvaris, and Ian M. Shoemaker. “Terrestrial effects on dark matter-electron scattering experiments”. In: *Phys. Rev. D* 96 (1 July 2017), p. 015018. DOI: 10.1103/PhysRevD.96.015018. URL: <https://link.aps.org/doi/10.1103/PhysRevD.96.015018>.

- [199] R. Essig et al. “Direct detection of sub-GeV dark matter with semiconductor targets”. In: *J. High Energy Phys.* 2016.5 (2016), p. 46. ISSN: 1029-8479. DOI: 10.1007/JHEP05(2016)046. URL: [http://dx.doi.org/10.1007/JHEP05\(2016\)046](http://dx.doi.org/10.1007/JHEP05(2016)046).
- [200] Rouven Essig, Tomer Volansky, and Tien-Tien Yu. “New constraints and prospects for sub-GeV dark matter scattering off electrons in xenon”. In: *Phys. Rev. D* 96 (4 Aug. 2017), p. 043017. DOI: 10.1103/PhysRevD.96.043017. URL: <https://link.aps.org/doi/10.1103/PhysRevD.96.043017>.
- [201] Rouven Essig et al. “First Direct Detection Limits on Sub-GeV Dark Matter from XENON10”. In: *Phys. Rev. Lett.* 109 (2 July 2012), p. 021301. DOI: 10.1103/PhysRevLett.109.021301. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.109.021301>.
- [202] Yonit Hochberg et al. “Detecting Superlight Dark Matter with Fermi-Degenerate Materials”. In: *JHEP* 08 (2016), p. 057. DOI: 10.1007/JHEP08(2016)057. arXiv: 1512.04533 [hep-ph].
- [203] J. N. Bahcall and A. M. Serenelli. “New Solar Opacities, Abundances, Helioseismology, and Neutrino Fluxes”. In: *The Astrophysical Journal* 621 (Mar. 2005), pp. L85–L88. URL: <http://stacks.iop.org/1538-4357/621/i=1/a=L85>.
- [204] John N. Bahcall and Aldo M. Serenelli. “How do Uncertainties in the Surface Chemical Composition of the Sun Affect the Predicted Solar Neutrino Fluxes?” In: *The Astrophysical Journal* 626 (June 2005), pp. 530–542. URL: <http://iopscience.iop.org/article/10.1086/429883>.
- [205] B. Aharmim et al. “Measurement of the ν_e and total ^8B solar neutrino fluxes with the Sudbury Neutrino Observatory phase-III data set”. In: *Phys. Rev. C* 87 (1 Jan. 2013), p. 015502. DOI: 10.1103/PhysRevC.87.015502. URL: <https://link.aps.org/doi/10.1103/PhysRevC.87.015502>.
- [206] A. Biekert et al. “Nuclear recoil scintillation linearity of a high pressure ^4He gas detector”. In: *JINST* 14.10 (2019), P10028. DOI: 10.1088/1748-0221/14/10/P10028. arXiv: 1907.03985 [physics.ins-det].
- [207] A. Biekert et al. “Scintillation yield from electronic and nuclear recoils in superfluid He^4 ”. In: *Phys. Rev. D* 105.9 (2022), p. 092005. DOI: 10.1103/PhysRevD.105.092005. arXiv: 2108.02176 [physics.ins-det].