

IN SEARCH OF MAJORANA NEUTRINOS
AND MICRON-SCALE INTERACTIONS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Clarke Alistair Hardy

May 2026

© 2025 by Clarke Alistair Hardy. All Rights Reserved.

Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-3.0 United States License.

<http://creativecommons.org/licenses/by/3.0/us/>

This dissertation is online at: <https://purl.stanford.edu/cc863pp6175>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Giorgio Gratta, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Daniel Akerib

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Savas Dimopoulos

Approved for the Stanford University Committee on Graduate Studies.

Stacey F. Bent, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format.

Abstract

While the Standard Model has been exceptionally successful at describing nature at a fundamental level, there remain key questions that it cannot answer. The discovery of non-zero neutrino masses provides direct evidence of physics beyond the Standard Model, yet the origin of these masses remains unknown. In addition, the Standard Model does not include gravity — a force measured with exquisite precision over macroscopic distances, but largely unexplored at short length scales. In this thesis, I describe experimental efforts on both of these fronts. nEXO is a planned experiment to search for neutrinoless double beta decay, a process that, if observed, would establish neutrinos as Majorana fermions and provide insight into their mass generation mechanism. The experiment will use five tonnes of liquid ^{136}Xe in a cryogenic, single-phase time projection chamber, serving as both the decay source and detection medium. I discuss two projects essential to meeting nEXO's target energy resolution and stringent radioactivity requirements: the development of a novel light response calibration scheme and the design of a low-radioactivity xenon purifier. Following this, I introduce a search for modifications to Newtonian gravity at short distances using optically-levitated microspheres as highly sensitive force sensors. I then present the latest constraints on micron-scale Yukawa interactions from this search and describe a systematic study of the backgrounds limiting the sensitivity.

Acknowledgments

None of the work described in the following pages would have been possible without the continuous support and guidance of my advisor, Giorgio Gratta. His endless curiosity and enthusiasm for physics set an example I will carry with me throughout my own career. I am grateful to Giorgio both for his mentorship and for the enjoyable and productive research environment he cultivated among his students and postdocs.

While working on two separate experiments in Giorgio's group, I had the benefit of learning from excellent scientists from a wide range of research backgrounds. Throughout my time in the liquid xenon lab, I learned a lot from Evan Angelico, Jacopo Dalmasson, Ralph DeVoe, Brian Lenardo, Lin Si, Marie Vidal, and Shuoxing Wu. The long hours we spent together in the lab formed an instrumental part of my education. Brian Lenardo deserves particular credit for his contributions to the work described in Chapter 5, both for conceiving the project and overseeing my own work. The nEXO experiment would not be possible without the dedication of every member of the nEXO Collaboration. In addition to my aforementioned Stanford colleagues, Isaac Arnquist, Dmitry Chernyak, Brian Mong, Jespere Calderone Nzobadila Ondze, Andreas Piepke, Peter Rowson, and Raymond Tsang all made essential contributions to the work reported in Chapters 5 and 6. My correspondence with Paolo Chiggiato of CERN on the topic of Chapter 6 was also a valuable resource.

My work for the microspheres group was done primarily in collaboration with Gautam Venugopalan, and I credit much of my learning to his extensive knowledge of optics and signal processing. Gautam ran the lab and continually improved the experimental hardware while I wrote software and performed the analysis. I am well aware that I had the easier job. On this project, I also benefited from collaboration

with Jacqueline Huang, Chengjie Jia, Kenneth Kohn, Meimei Liu, Lorenzo Magrini, Zhengruilong Wang, and Yuqi Zhu. I joined the microspheres group late in my fourth year with little experience outside of particle physics. It was through continued collaboration with these individuals that I began to understand and appreciate new approaches to experimental physics.

During my time in the Gratta group, I enjoyed countless conversations with many people, including those I did not work with directly. In addition to those already mentioned, this includes Soud Al Kharusi, Chas Blakemore, Chiara Brandenstein, Alexander Fieguth, Joey Howlett, Albert Nazeeri, and Nadav Priel. Our group lunches and coffee chats, on physics and seemingly everything else, made for a fun and memorable graduate school experience.

Finally, I want to thank my parents for their unconditional support and encouragement throughout this process. Their commitment to my success and happiness made this accomplishment possible.

Contents

Abstract	vii
Acknowledgments	ix
1 Introduction	1
2 Neutrino Masses and the Standard Model	3
2.1 The Standard Model of Particle Physics	3
2.1.1 The Fermion Mass Mechanism	5
2.1.2 A Short History of Neutrino Physics	6
2.1.3 The Three-Neutrino Paradigm	7
2.1.4 Other Neutrino Mass Probes	12
2.1.5 Open Questions	15
2.2 Massive Neutrinos	16
2.2.1 The Effective Field Theory Approach	16
2.2.2 Adding Right-Handed Neutrinos	17
2.2.3 The Seesaw Mechanism	17
2.3 From Theory to Experiment	20
2.3.1 Two-Neutrino Double Beta Decay	20
2.3.2 Neutrinoless Double Beta Decay	22
2.3.3 The Decay Rate	23
2.3.4 The Experimental Program	26
2.3.5 Implications of a Discovery	30

3	Gravity and New Physics at the Micron Scale	33
3.1	Theoretical Considerations	33
3.1.1	The Hierarchy Problem	33
3.1.2	Potential Resolutions	34
3.1.3	Large Extra Dimensions	35
3.1.4	Other Theoretical Motivations	37
3.2	Experimental Tests	39
3.2.1	The Landscape of Measurements	40
3.2.2	A New Experimental Technique	42
4	The nEXO Experiment	45
4.1	Xenon Time Projection Chambers	45
4.1.1	Xenon Microphysics	46
4.2	The nEXO Detector	48
4.2.1	Measuring Charge Signals	49
4.2.2	Measuring Light Signals	51
4.3	Three-Parameter Analysis	52
4.3.1	Energy	52
4.3.2	Standoff Distance	55
4.3.3	Event Topology	55
4.4	Backgrounds	56
4.5	Sensitivity	57
5	A Xenon-127 Calibration Scheme for nEXO	61
5.1	The Energy Resolution Model	61
5.2	Position Dependence of the TPC Response	63
5.2.1	Standard Calibration Techniques	65
5.2.2	^{127}Xe as a Calibration Source	66
5.3	Experimental Demonstration	67
5.3.1	Source Production	67
5.3.2	Deployment in a Test TPC	70
5.3.3	Electron Lifetime Calibration	72

5.4	Projected Calibration Performance for nEXO	75
5.4.1	Simulating the Light Response	75
5.4.2	Lightmap Reconstruction Strategies	77
5.4.3	Lightmap Calibration Performance	82
5.5	Conclusions and Prospects for nEXO	84
6	A Purifier for nEXO using Pure Zirconium	87
6.1	Introduction to Xenon Purification	88
6.1.1	Getter Materials	89
6.1.2	Zirconium Alloys	89
6.1.3	High-Purity Zirconium	90
6.2	Design of a Xenon Purifier	91
6.2.1	Getter Pellets	91
6.2.2	Geometry & Pressure Drop	93
6.2.3	Design & Assembly	97
6.3	Measurements of Purification Efficiency	100
6.3.1	The SLAC Xenon Purity Monitor	100
6.3.2	Purification Measurement Procedure	101
6.3.3	Impurity Transport Model	102
6.3.4	Results	103
6.4	Radioactivity of High-Purity Zirconium	105
6.4.1	Radioassay Measurement Campaign	105
6.4.2	Measured Activity	107
6.5	Future Work	111
6.5.1	Dry Distillation for Radium Removal	111
6.5.2	Capacity and Regeneration Measurements	111
7	The Search for Micron-Scale Interactions	113
7.1	Introduction to Optical Tweezers	114
7.1.1	The Force from a Laser	114
7.1.2	Creating an Optical Trap	115
7.1.3	The Equation of Motion	116

7.1.4	Microspheres as Force Sensors	118
7.2	Experimental Setup	118
7.2.1	The Laser	119
7.2.2	Position Sensing	119
7.2.3	The Trapping Region	120
7.2.4	Improvements Since the Previous Result	122
7.3	Overview of the Search	124
7.4	The Statistical Model	125
7.4.1	Likelihood Function Construction	127
7.4.2	Test Statistic Construction	128
7.4.3	Incorporating Systematic Uncertainties	130
7.4.4	Harmonic Selection Criteria	131
7.5	Constraints on Yukawa Interactions	133
7.6	Conclusions	135
8	Background Management Strategies	137
8.1	Techniques for Background Subtraction	137
8.1.1	Physics-Based Model	138
8.1.2	Asynchronous Measurement	139
8.1.3	In Situ Witness	139
8.2	General Features of the Backgrounds	139
8.3	Mechanical Vibrations	141
8.4	Scattering of Stray Laser Light	141
8.4.1	Constructing a Scattered Light Witness Stream	146
8.5	Coupling to the Electric Dipole Moment	159
8.5.1	Electric Dipole Moment of Silica Microspheres	159
8.5.2	Contact Potentials	160
8.5.3	In Situ EDM Background Monitoring	161
8.6	Summary and Current Status	168
	Conclusions	169

Summary for Non-Scientists	173
A $0\nu\beta\beta$ Discovery Probability	179
B The Stanford Liquid Xenon Lab	183
B.1 Overview	183
B.2 Cryogenics	184
B.3 Xenon Filling & Handling	185
B.3.1 Performance	187
B.3.2 Filling & Recovery	189
B.3.3 Recirculation	192
B.4 Slow Controls & Monitoring	193
B.4.1 LabVIEW Code	193
B.4.2 Monitoring Code	194
B.5 Experiments Currently Supported	196
B.5.1 Test TPCs	197
B.5.2 Data Acquisition	198
B.5.3 Light & Charge Sensing	199
B.5.4 LXe Compatibility Test Chambers	200
B.5.5 Calibrations	201
C Microspheres and Measurements	205
D The OptLevAnalysis Package	209
D.1 The <code>data_processing</code> Module	209
D.1.1 The <code>FileData</code> Class	209
D.1.2 The <code>AggregateData</code> Class	212
D.2 Other Modules	213
D.2.1 The <code>plotting</code> Module	213
D.2.2 The <code>stats</code> Module	214
D.2.3 The <code>signals</code> Module	215
D.2.4 The <code>synthetic_data</code> Module	215

D.2.5	External Code	216
E	Exclusion Limit Calculation	219
E.1	Computing the Log-Likelihood Functions	219
E.2	Limits from the Test Statistic	222
F	The Physics of Contact Potentials	225
	References	229

List of Tables

2.1	Field content of the Standard Model.	4
2.2	Global fits to neutrino oscillation parameters.	12
6.1	List of the high-purity zirconium samples that were procured from ALB Materials.	92
6.2	The radioassay measurements conducted by various participating in- stitutions on the high-purity zirconium samples.	106
C.1	Properties of the microspheres used in this work.	205
C.2	Summary of microsphere measurement configurations.	207

List of Figures

2.1	Neutrino flavor oscillation probability.	9
2.2	Neutrino mass ordering scenarios.	13
2.3	Measured values of all fermion masses.	15
2.4	Feynman diagram showing the Weinberg operator.	16
2.5	Feynman diagram for the Seesaw Mechanism.	19
2.6	Mass excess parabolas showing β -decay and $\beta\beta$ -decay.	21
2.7	Electron energy spectra for $2\nu\beta\beta$ and $0\nu\beta\beta$	22
2.8	Feynman diagrams for $2\nu\beta\beta$ and $0\nu\beta\beta$	23
2.9	Feynman diagrams illustrating the Black Box Theorem.	24
2.10	The lobster plot showing the allowed parameter space for the effective Majorana mass.	25
2.11	Q -values and abundances for isotopes which undergo $2\nu\beta\beta$	27
2.12	The sensitivity of a selection of experiments to $0\nu\beta\beta$	28
3.1	Constraints on modifications to the gravitational inverse square law.	41
3.2	Two sample apparatuses for the study of short-range gravity.	43
4.1	Production of scintillation and ionization signals in xenon.	48
4.2	CAD rendering of the nEXO TPC.	50
4.3	Schematic of a nEXO charge-sensing tile.	51
4.4	Three parameters used for signal/background discrimination in nEXO.	53
4.5	Rotated energy axis constructed from ionization and scintillation signals.	54
4.6	The projected sensitivity of nEXO to $0\nu\beta\beta$	58

5.1	^{127}Xe electron capture decay scheme.	67
5.2	Ionization and scintillation signals produced by the decay of ^{127}Xe	68
5.3	Neutron capture cross section and reactor flux spectrum.	69
5.4	Activities of radioisotopes produced by neutron irradiation of the SS cylinder of xenon gas.	71
5.5	Test TPC and charge tile used in the ^{127}Xe calibration demonstration.	72
5.6	Deployment of the ^{127}Xe source in a test detector.	73
5.7	Data from the electron lifetime calibration using a ^{127}Xe source.	74
5.8	Simulated charge and light detected in nEXO.	77
5.9	Sequence of steps in a simulated lightmap calibration scheme.	78
5.10	Lightmap error using a Kernel Smoothing reconstruction method.	83
5.11	Lightmap error in the TPC with increasing calibration dataset sizes.	84
5.12	Lightmap error attainable with a ^{127}Xe calibration scheme.	85
6.1	XPS data collected with a high-purity zirconium sample.	92
6.2	SEM images of St 707 and high-purity zirconium pellets.	94
6.3	Pressure drop across the purifier as a function of diameter.	97
6.4	Design of the custom purifier using high-purity zirconium.	98
6.5	Photos of the custom purifier during and following assembly.	99
6.6	The xenon purity during the run at the SLAC XPM	104
6.7	Xenon purity over time when purifying with the SAES.	105
6.8	Xenon purity over time when purifying with the SCP.	106
6.9	The ^{238}U decay chain.	108
6.10	Radioassay measurements of the high-purity zirconium samples.	109
7.1	Optical forces on a sphere in the geometric optics regime.	116
7.2	Schematic of the optics system used in this work.	121
7.3	Diagram of the trap region.	122
7.4	Time evolution of the force on the microsphere.	123
7.5	Measured force spectra compared to a signal template.	125
7.6	Measurements of the force at selected harmonics.	126
7.7	Systematic errors included in the constraints on Yukawa interactions.	131

7.8	New constraints on Yukawa interactions from this work.	134
8.1	Variation of the backgrounds over time.	140
8.2	Backgrounds caused by mechanical vibrations.	142
8.3	Effect of an aperture and the microsphere-attractor separation on back- grounds.	144
8.4	Fitted signal strengths as a function of microsphere-attractor separation.	145
8.5	Microsphere resonances seen in each of the quadrant photodiode sensors.	148
8.6	Visualization of the position-sensing matrices.	154
8.7	Resonances and the null stream before and after calibration.	155
8.8	Response to driven forces before and after position-sensing calibration.	156
8.9	Backgrounds as measured in x , y , and the null stream.	157
8.10	The technique for measuring the dipole moment of a microsphere. . .	160
8.11	Demonstration of electric dipole interaction background witnessing. .	165
A.1	Effective Majorana mass as a function of Majorana phases.	181
B.1	Photo of the liquid xenon detector test facility.	184
B.2	Piping & instrumentation diagram for the liquid xenon handling system.	185
B.3	Cooling and warming of a detector in the large test stand.	188
B.4	Stability of the xenon cell temperature.	189
B.5	Natural warming of a xenon cell in the large test stand.	190
B.6	Mass flow rate measured during recirculation through the purifier. . .	193
B.7	Functional schematic of the distributed monitoring system.	195
B.8	CAD renderings of three test TPCs.	198
B.9	CAD renderings of the xenon compatibility test chambers.	202
D.1	Flow chart showing the analysis pipeline for the micron-scale interac- tion search.	217
F.1	Schematic illustrating the origin of the contact potential.	227
F.2	Elimination of the contact potential with an applied bias.	228

Chapter 1

Introduction

As its title suggests, this thesis concerns Majorana neutrinos and micron-scale interactions, two topics that may appear unrelated. Each of these topics is the focus of a separate experiment in the domain of fundamental physics. Despite their apparent differences, the two experiments are united by the goal of discovering phenomena that are not explainable under the current theoretical paradigm. Both experiments approach this goal by targeting a known gap in the existing theory.

In the first case, the gap is in the neutrino sector. The Standard Model in its current form does not include masses for the neutrinos, though there are well-motivated ways in which it could be extended to accommodate them. As I will show, the existence of a process called neutrinoless double beta decay can be seen as a natural consequence of this extension. An observation of this process would represent a major breakthrough with far-reaching implications for particle physics and cosmology. Chapter 2 introduces the theory underlying this elusive process, before turning to the landscape of experiments aiming to discover it. Chapter 4 describes nEXO, a search for neutrinoless double beta decay, and outlines many of the experimental considerations that will be key to its eventual success. Two of these are expanded upon in Chapter 5 and Chapter 6, which describe my own work for nEXO on developing a calibration scheme and designing a xenon purifier, respectively.

The second gap in the existing theory is much larger; it is the complete absence of a quantum description of gravity. Gravity has never been directly measured for

separations below fifty micrometers, and many attempts to reconcile gravity with quantum mechanics postulate new interactions that manifest as modifications to the gravitational force at short distances. A selection of these theoretical ideas is introduced in Chapter 3. The latest constraints from the experiment we conducted at Stanford using optically-levitated microspheres to search for interactions of this kind are presented in Chapter 7. As the sensitivity of the search was limited by backgrounds, much of my work was focused on understanding their origin. In Chapter 8, I introduce some diagnostic tools for this purpose.

In the Conclusions, I reiterate the findings of this thesis and look ahead to the future of the projects described within. And for readers without a technical background, or for anyone who does not find discussions of vacuum systems and statistical minutiae as enthralling as I do, the Summary for Non-Scientists provides a condensed overview of this thesis at a more accessible level.

Chapter 2

Neutrino Masses and the Standard Model

For the last quarter-century, it has been understood that the neutrino has a non-zero rest mass. Incorporating this fact into the existing theoretical framework is not trivial, however. In this chapter, that framework is introduced and some relevant findings from neutrino physics are reviewed. The chapter concludes with a description of an experimental signature that would shed light on the nature of neutrino masses.

2.1 The Standard Model of Particle Physics

The Standard Model (SM) of particle physics is a gauge quantum field theory which describes all known fundamental particles and interactions, with the notable exception of gravity [1]. The basic ingredients of the theory are the fields listed in Table 2.1 and the symmetry properties they exhibit. These include fermions comprising matter (expressed as Weyl spinors¹), force-carrying gauge bosons, and a Higgs scalar. The constituent fields can be classified according to their transformation properties under

¹Throughout this chapter, I use the notation convention from Ref. [2], with spinor indices omitted for clarity. All fermion fields are expressed as two-component left-chiral Weyl spinors; the right-chiral components required to build Dirac spinors are obtained via Hermitian conjugation. A bar over a field labels it as an $SU(2)_L$ singlet, while the unbarred fields are $SU(2)_L$ doublets.

the gauge symmetry group,

$$\text{SU}(3)_C \times \text{SU}(2)_L \times \text{U}(1)_Y, \quad (2.1)$$

where C stands for color in the strong interaction, L refers to the weak interaction's parity-violating preference for left-handed chirality, and Y symbolizes weak hypercharge. In addition to the gauge symmetries, the SM respects two U(1) global symmetries: baryon number, B , and lepton number, L , although these symmetries are *accidental*; they are not built into the theory in advance.

Table 2.1: The fields that make up the Standard Model along with their respective symbols and transformation properties under the gauge groups. The index i runs from 1 to 3 and refers to the fermion generation.

Name	Symbol	SU(3) _C	SU(2) _L	U(1) _Y
Electron/muon/tau	\bar{e}_i	1	1	+1
Lepton doublet	$\ell_i \equiv (e_i \ \nu_i)^T$	1	2	-1/2
Down/strange/bottom quark	\bar{d}_i	$\bar{3}$	1	+1/3
Up/charm/top quark	\bar{u}_i	$\bar{3}$	1	-2/3
Quark doublet	$q_i \equiv (u_i \ d_i)^T$	3	2	+1/6
Gluons	G_μ	8	1	0
Weak isospin gauge bosons	W_μ	1	3	0
Weak hypercharge gauge bosons	B_μ	1	1	0
Higgs scalar	Φ	1	2	+1/2

The full Lagrangian of the SM can then be generated according to a simple prescription: include all couplings of the fields that respect Lorentz invariance, gauge invariance, and renormalizability [2]. This is in accordance with the *totalitarian principle*, attributed to Murray Gell-Mann, which says, “that which is not forbidden is compulsory” [3]. Stated differently, any interactions that do not explicitly violate the symmetries of the theory should be expected to occur with some probability. This prescription will be useful later in this chapter when we confront the task of modifying the theory to accommodate massive neutrinos. Before addressing neutrino mass generation, however, a discussion of ordinary fermion masses will set the stage.

2.1.1 The Fermion Mass Mechanism

Fermions obey the Dirac equation; they are described by four-component Dirac spinors $\Psi_D \equiv (\chi \ \xi^\dagger)^T$ with a left-chiral part (χ) and a right-chiral part (ξ^\dagger) [4]. Alternatively, as was first demonstrated by Ettore Majorana in 1937 [5], a Majorana spinor $\Psi_M \equiv (\psi \ \psi^\dagger)^T$ can be constructed from a single Weyl spinor, ψ , and its conjugate, reducing the number of real degrees of freedom from four to two. This describes a particle that is its own antiparticle, and therefore requires that the particle carry no gauge charges. The mass terms in the Dirac Lagrangian for these spinors are

$$\mathcal{L}_D = -m_D(\chi\xi + \chi^\dagger\xi^\dagger) \quad \text{and} \quad \mathcal{L}_M = -m_M(\psi\psi + \psi^\dagger\psi^\dagger), \quad (2.2)$$

where m_D and m_M are the Dirac and Majorana masses, respectively. Neither of these terms is among the SM Lagrangian terms generated according to the prescription defined above, however. The reason is that such terms are not gauge invariant². The problem of giving the fermion fields masses in a gauge invariant way is solved by the Higgs field, Φ . The following terms are allowed by the gauge symmetries of the SM (considering only a single fermion generation for simplicity):

$$\mathcal{L}_{SM} \supset -y_u\Phi q\bar{u} - y_d\Phi^\dagger q\bar{d} - y_e\Phi^\dagger \ell\bar{e} + \text{h.c.} \quad (2.3)$$

These terms represent the Yukawa couplings between Φ and the left- and right-handed³ fermion fields, with $y_{u,d,e}$ determining the strength of the couplings. At high temperatures, the effective Higgs potential has a minimum at $\Phi = 0$, rendering the fermions massless. When the temperature decreases, the potential takes on its famous wine bottle shape, and Φ acquires a non-zero vacuum expectation value (vev) as it settles into the new minimum. In this process, known as Electroweak Symmetry Breaking (EWSB), the $SU(2)_L \times U(1)_Y$ gauge symmetry is spontaneously broken

²To see this, note that the Dirac mass term for the electron $-m_e e\bar{e}$ includes both the $SU(2)_L$ singlet \bar{e} and the $SU(2)_L$ doublet e , so the full term is not a singlet under $SU(2)_L$.

³In this context, left-handed and right-handed refer to the $SU(2)_L$ transformation properties of the fields, with “left-handed” meaning an $SU(2)_L$ doublet, and “right-handed” meaning an $SU(2)_L$ singlet. When referring to chirality of a field, I will instead use the terms “left-chiral” and “right-chiral.”

down to the electromagnetic gauge symmetry, $U(1)_{\text{EM}}$. In the unitary gauge, the Higgs doublet in the broken phase can be expressed as

$$\Phi(x) = \frac{v + H(x)}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad (2.4)$$

where $H(x)$ is a real scalar field and v is the Higgs vev. The Lagrangian terms then become

$$\mathcal{L}_{\text{SM}} \supset -\frac{y_u v}{\sqrt{2}} q\bar{u} - \frac{y_d v}{\sqrt{2}} q\bar{d} - \frac{y_e v}{\sqrt{2}} \ell\bar{e}. \quad (2.5)$$

Finally, the familiar Dirac mass terms have appeared, with $m_u = y_u v/\sqrt{2}$, and similarly for m_d and m_e . The masses are generated by the Yukawa coupling between the left- and right-handed fermions and the Higgs, and they are proportional to the coupling strengths and the Higgs vev. While this explains the origin of the up quark, down quark, and electron masses, the neutrino is left massless. This is because the SM does not include an $SU(2)_L$ singlet counterpart for the neutrino field, ν . Indeed, in the minimal SM, this was by design — at the time it was being developed, neutrinos were taken to be massless. Since then, it has been established that at least two of the three neutrino masses must be non-zero, rendering the minimal SM incomplete.

2.1.2 A Short History of Neutrino Physics

With the general framework having been described, we can now look at how neutrinos fit into the picture. Neutrinos were first postulated by Wolfgang Pauli in 1930 as a way to explain the continuous energy spectrum of electrons from β^- decay, which seemed to violate the law of conservation of energy without a third body [6]. It wasn't until more than two decades later that Cowan and Reines found the first direct evidence of their existence. In a series of two experiments, they observed the capture of antineutrinos from a nuclear reactor on hydrogen, first using a hydrogenated liquid scintillator target, and later with water targets sandwiched between liquid scintillator detectors [7, 8]. With the discovery of parity violation in the weak interaction by Chien-Shiung Wu, after the idea was proposed by Lee and Yang, evidence was mounting to suggest

that neutrinos are left-handed [9, 10]. The Goldhaber experiment conducted in 1957 confirmed that neutrinos (antineutrinos) have negative (positive) helicity [11].

The 1960s ushered in a new approach to the study of neutrino properties with the first detection of solar neutrinos. Neutrinos from the Sun, predominantly those produced in the β^+ decay of ^8B in the solar core, were measured by the Homestake experiment conducted by Raymond Davis Jr. and John Bahcall [12, 13]. The experiment measured the production of ^{37}Ar through neutrino capture on ^{37}Cl in a large chlorine tank deep underground in the Homestake gold mine [14]. Bahcall's detailed modeling of nuclear processes in the solar core gave a predicted neutrino detection rate approximately three times the measured rate [15–17]. This became known as the solar neutrino problem. A solution first proposed by Bruno Pontecorvo suggested that if the neutrino masses were non-zero, neutrinos produced in one flavor could oscillate into other flavors undetectable by the Homestake experiment [18]. By the early 2000s, results from the Sudbury Neutrino Observatory (SNO) [19] and the Super-Kamioka Neutrino Detection Experiment (Super-Kamiokande, or Super-K) [20] had confirmed that at least two of the neutrino masses are non-zero, solving the solar neutrino problem and establishing the existence of physics beyond the minimal SM.

2.1.3 The Three-Neutrino Paradigm

In the current paradigm, three flavor states (ν_e, ν_μ, ν_τ) are related to the three mass states (ν_1, ν_2, ν_3) by a unitary mixing matrix, U_{PMNS} [18, 21, 22]:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \quad (2.6)$$

The matrix U_{PMNS} is known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix. It is parameterized as follows:

$$U_{\text{PMNS}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (2.7)$$

where $s_{ij} \equiv \sin \theta_{ij}$ and $c_{ij} \equiv \cos \theta_{ij}$. The PMNS matrix has four free parameters: the mixing angles, θ_{12} , θ_{13} , and θ_{23} , and a phase, δ_{CP} that describes CP violation in the neutrino sector. The off-diagonal terms in the PMNS matrix mean that the mass eigenstates and flavor eigenstates do not align, and mixing is expected. In the relativistic limit, the energy of a propagating mass eigenstate $|\nu_i\rangle$ is given by⁴

$$E_i = \sqrt{p_i^2 + m_i^2} \approx p_i \left(1 + \frac{m_i^2}{2p_i^2} \right) = p_i + \frac{m_i^2}{2p_i} \approx E + \frac{m_i^2}{2E}, \quad (2.8)$$

where E is the total neutrino energy, and in the last step, the approximation $E \approx p_i$ is used. Applying the time evolution operator with energy E_i to the mass eigenstates in Eq. (2.6) allows a flavor eigenstate $|\nu_\alpha\rangle$ to be described by

$$|\nu_\alpha(t)\rangle = \sum_i U_{\alpha i} \exp \left[-i \left(E + \frac{m_i^2}{2E} \right) t \right] |\nu_i\rangle. \quad (2.9)$$

To illustrate the phenomenology of neutrino oscillation, it is helpful to consider the simpler two-neutrino case, with flavors $|\nu_\alpha\rangle$ and $|\nu_\beta\rangle$, mass eigenstates $|\nu_1\rangle$ and $|\nu_2\rangle$, and a 2×2 mixing matrix with a single angle, θ . One can show that the probability that a neutrino produced in state $|\nu_\alpha\rangle$ will be detected in state $|\nu_\beta\rangle$ at time t is given by

$$P_{\alpha \rightarrow \beta} = |\langle \nu_\beta(t) | \nu_\alpha(0) \rangle|^2 = \sin^2(2\theta) \sin^2 \left(\frac{\Delta m^2 L}{4E} \right), \quad (2.10)$$

where $\Delta m^2 \equiv m_2^2 - m_1^2$, and $t \approx L$ has been used to express the probability in terms of the travel length, L . Evidently, the flavor state oscillates with a frequency set by the difference between the squared masses and the ratio of the propagation distance to

⁴In this chapter and the next, I use natural units with $c = \hbar = 1$.

the neutrino energy. This is true in the three-neutrino case as well, though amplitude modulations and interference between multiple oscillation frequencies also play a role. These features are visible in Fig. 2.1. The dependence on L has permitted a broad range of experiments to target different oscillation parameters by placing detectors at different baselines. A brief survey of these techniques follows.

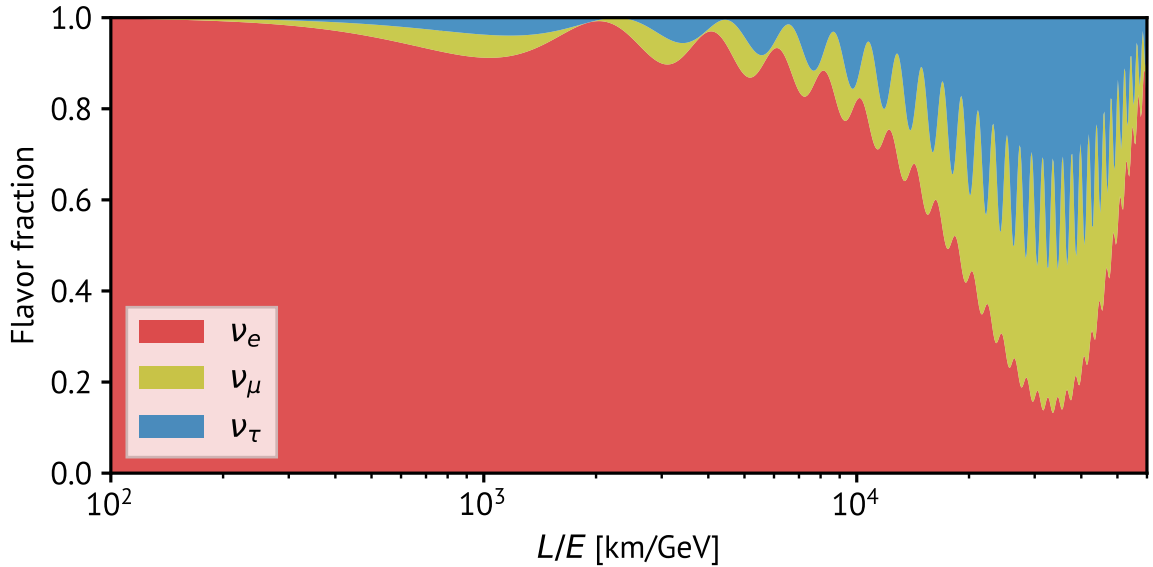


Figure 2.1: Illustration of the phenomenon of neutrino oscillation. Neutrinos produced in one flavor (in this case, electron neutrinos) develop a probability of being measured in a different flavor after some distance of propagation, depending on their energy.

Solar Neutrinos

A number of experiments have measured θ_{12} and Δm_{21}^2 , known as the solar neutrino parameters. Solar neutrinos travel a significant distance through the Sun before exiting. Due to the electron background in the solar interior, the mass states that diagonalize the Hamiltonian are not eigenstates of free propagation. This results in a modification to the oscillation behavior described above and is known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect [23–29]. The admixture of flavor states that constitute eigenstates of the Hamiltonian is different for neutrinos and

antineutrinos, allowing the sign of Δm_{21}^2 to be determined. The solar neutrino flux was first measured by the Homestake experiment [30], and then concurrently by the Gallium Experiment (GALLEX), later called the Gallium Neutrino Observatory (GNO) [31, 32], and the Soviet-American Gallium Experiment (SAGE) [33]. The GALLEX/GNO and SAGE programs used neutrino capture on ${}^{71}\text{Ga}$, with the reaction threshold of 233.2 keV allowing for much lower energy neutrinos to be detected, compared to the ${}^{37}\text{Cl}$ threshold of 814 keV. First evidence of neutrino oscillations in solar neutrinos was published by SNO, which separately measured charged current, neutral current, and neutrino-electron elastic scattering interactions on a heavy water target. The first interaction channel is sensitive only to electron neutrinos, the second equally to all flavors, and the third preferentially to electron neutrinos, but with some sensitivity to the other two as well [34, 19, 35]. Super-K and Borexino have also measured solar neutrinos, the former with sensitivity to both neutral and charged current elastic scattering on electrons in an ultrapure water target, and the latter using a liquid scintillator target [36–38].

Atmospheric Neutrinos

The first direct evidence of neutrino oscillations was observed by Super-K in neutrinos produced by cosmic rays in the upper atmosphere [20]. The atmospheric neutrino oscillation parameters, θ_{23} and $|\Delta m_{32}^2|$, have since been measured with improved precision by Super-K and the IceCube DeepCore experiment [39–41]. IceCube is a cubic-kilometer of ice at the South Pole instrumented with PMTs to measure Cherenkov radiation for through-going charged particles; the denser sub-array of PMTs in IceCube known as DeepCore has sensitivity to GeV-scale atmospheric neutrinos, in contrast to the much larger energy astrophysical neutrinos for which IceCube is optimized. Measurements of neutrinos on an upward trajectory through the Earth allows for a full Earth-diameter to be used as the baseline.

Reactor Neutrinos

Disappearance of electron antineutrinos produced in nuclear reactors over medium (~ 1 km) baselines has enabled direct measurements of θ_{13} and $|\Delta m_{31}^2|$. The Double-Chooz experiment, the Reactor Experiment for Neutrino Oscillation (RENO) and the Daya Bay Reactor Neutrino Experiment have all reported such measurements [42–44]. Reactor neutrinos measured over a longer baseline by Kamioka Liquid Scintillator Antineutrino Detector (KamLAND) have provided further data on Δm_{21}^2 , along with the mixing angles θ_{12} and, with less precision, θ_{13} [45].

Accelerator Neutrinos

Accelerated protons impacting a graphite target produce pions and kaons, which decay predominantly into muons and muon neutrinos. The resulting neutrino beam can be used to study oscillation parameters via the disappearance (appearance) of muon (electron) neutrinos or antineutrinos between a near detector and a far detector. The Main Injector Neutrino Oscillation Search (MINOS) using the Fermilab accelerator, along with its successor, the NuMI Off-axis ν_e Appearance (NO ν A) experiment, and the Tokai to Kamioka (T2K) experiment using the neutrino beam from the J-PARC facility in Tokai, have all produced measurements of the mixing angle θ_{23} and the mass-squared difference $|\Delta m_{32}^2|$ [46–49]. As these experiments use both neutrino and antineutrino beams and are sensitive to differences between them, they have some sensitivity to δ_{CP} . Furthermore, because these beams propagate through the Earth, the MSW effect provides sensitivity to the sign of $|\Delta m_{32}^2|$ as well, though the current sensitivity is not enough to provide definitive evidence in favor of one over the other.

Global Fits to Neutrino Oscillation Data

Joint fits over the many individual datasets can provide a more complete picture of the current status of the three-neutrino paradigm. The NuFIT collaboration [50, 51] provides continually-updated fit parameters as new measurements are reported; results from the latest release are included in Table 2.2. A remaining unknown is the ordering of the mass eigenstates. The scenario in which $\Delta m_{32}^2 > 0$ is known as

the normal ordering (NO) scenario, as it implies that $m_1 < m_2 < m_3$. If $\Delta m_{32}^2 < 0$, then the inverted ordering (IO) scenario is correct and $m_3 < m_1 < m_2$. As there is no *a priori* reason to prefer one scenario over the other, the labels “normal” and “inverted” are somewhat misleading. Fig. 2.2 depicts both scenarios.

Table 2.2: Best-fit values of the neutrino oscillation parameters along with their 1σ uncertainties, from Ref. [51]. In the fifth row, $\ell = 2$ under the NO scenario and $\ell = 1$ under the IO scenario.

Parameter	NO Fit ($\pm 1\sigma$)	IO Fit ($\pm 1\sigma$)
$\sin^2 \theta_{12}$	$0.308^{+0.012}_{-0.011}$	$0.308^{+0.012}_{-0.011}$
$\sin^2 \theta_{23}$	$0.470^{+0.017}_{-0.013}$	$0.550^{+0.012}_{-0.015}$
$\sin^2 \theta_{13}$	$0.02215^{+0.00056}_{-0.00058}$	$0.02231^{+0.00056}_{-0.00056}$
Δm_{21}^2 [10^{-5} eV ²]	$7.49^{+0.19}_{-0.19}$	$7.49^{+0.19}_{-0.19}$
$ \Delta m_{3\ell}^2 $ [10^{-3} eV ²]	$+2.513^{+0.021}_{-0.019}$	$-2.484^{+0.020}_{-0.020}$
δ_{CP} [°]	212^{+26}_{-41}	274^{+22}_{-25}

2.1.4 Other Neutrino Mass Probes

In addition to the aforementioned oscillation experiments, the precision measurement of β -decay spectra and analysis of astrophysical data can also serve as probes of the neutrino mass, as described below.

Beta Decay Spectra

In β -decay, the shape of the electron kinetic energy spectrum near the Q -value is sensitive to a weighted average of the three neutrino masses, m_β , given by

$$m_\beta = \sqrt{\sum_{i=1}^3 |U_{ei}|^2 m_i^2}. \quad (2.11)$$

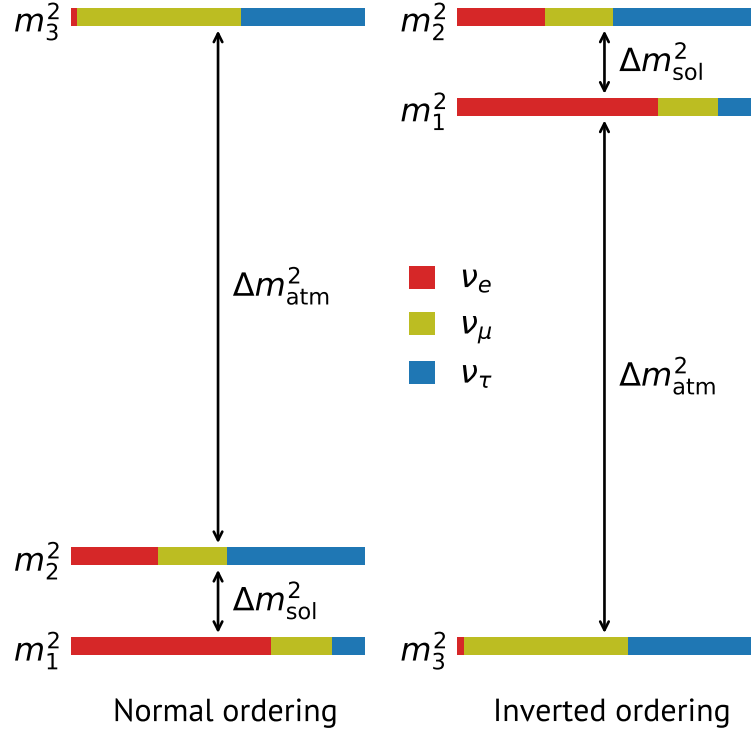


Figure 2.2: The two possible neutrino mass ordering scenarios allowed by current neutrino oscillation data, with the solar neutrino mass squared difference defined as $\Delta m_{\text{sol}}^2 \equiv \Delta m_{21}^2$ and the atmospheric neutrino mass squared difference defined as $\Delta m_{\text{atm}}^2 \equiv \Delta m_{32}^2 \approx \Delta m_{31}^2$.

In principle, this allows for m_β to be determined from precise measurements of the spectral shape near the endpoint, though a direct measurement of m_β is currently beyond the limits of experimental sensitivity. The Karlsruhe Tritium Neutrino (KATRIN) experiment has demonstrated sensitivity to m_β at the level of 0.45 eV in tritium β -decay using a large spectrometer based on Magnetic Adiabatic Collimation with Electrostatic (MAC-E) filters [52, 53]. The planned Project 8 experiment aims to improve on this limit by approximately an order of magnitude using the novel Cyclotron Radiation Emission Spectroscopy (CRES) technique [54, 55].

Cosmology

Measurements of cosmological observables can be used to constrain the sum of neutrino masses, $\Sigma \equiv \sum_i m_i$, as neutrinos leave their imprint at multiple stages throughout the evolution of the Universe [56]. The cosmic microwave background (CMB) power spectrum is a reflection of temperature anisotropies as a function of length scale (expressed as a multipole moment) on the last scattering surface, which in turn depends on scale-dependent density perturbations at the time of recombination. Neutrinos can affect the CMB power spectrum in a number of ways. By contributing to the radiation budget, they delay the epoch of matter-radiation equality (via the early integrated Sachs-Wolf effect), predominantly affecting the first peak. They also affect the expansion history and the angular scale. Before recombination, they inhibit structure formation on length scales smaller than their free streaming scale, resulting in the attenuation of density fluctuations at high multipole moments. After recombination, the suppression of structure formation reduces the gravitational lensing potential; this affects the sharpness of the high-multipole peaks. Baryon acoustic oscillations at recombination are visible today as a preferred separation scale in the distribution of galaxies, reflecting the length scale at the baryon-photon decoupling time. Measurements of this feature in galaxy surveys can break the degeneracies in the CMB fit parameters, improving the constraints on Σ . The latest constraint, combining data from the Planck satellite [57, 58] and the Atacama Cosmology Telescope (ACT) [59, 60] with BAO measurements from the Dark Energy Spectroscopic Instrument (DESI) [61] requires $\Sigma < 0.082$ eV [62, 63]⁵.

These disparate probes all point to an extraordinarily small absolute neutrino mass scale. Indeed, the maximum allowed mass from current data is more than a million times smaller than that of any other fermion (Fig. 2.3). If neutrinos get Dirac masses from the Higgs, this would suggest an unnaturally small Yukawa coupling strength of $y_\nu^j \sim 10^{-12}$.

⁵These results are in tension at the level of $\sim 3\sigma$ with oscillation data, potentially indicating unknown systematics or hinting at new physics.

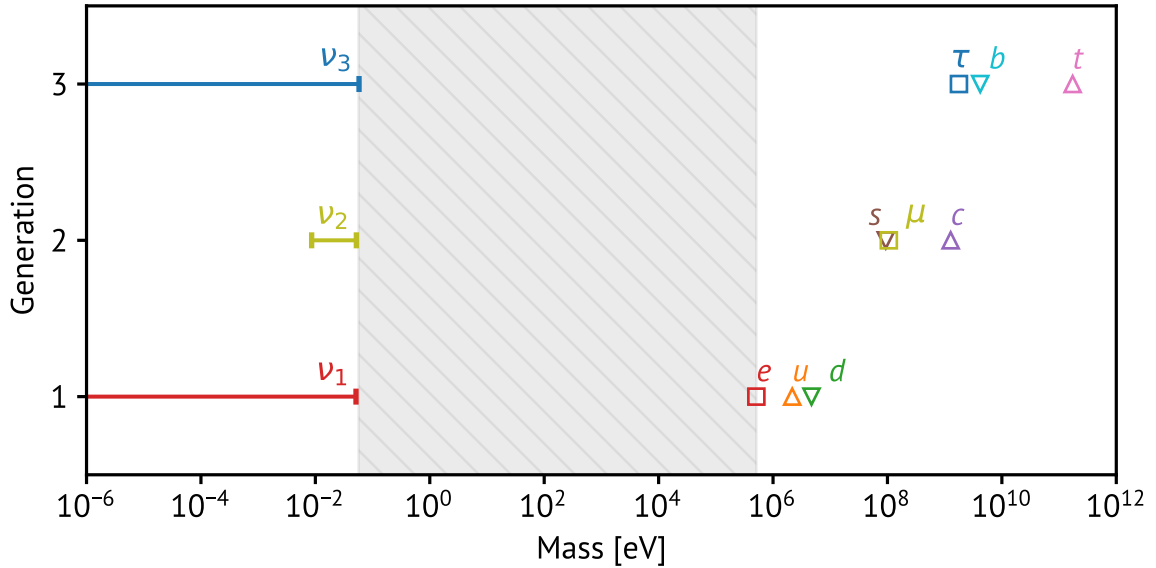


Figure 2.3: Current measured values for the masses of all three generations of fermions. Limits on the neutrino mass are also shown, with both ordering schemes considered.

2.1.5 Open Questions

Despite substantial progress on narrowing in on the neutrino mass parameters, the following unknowns remain:

- The ordering of the masses. The solar neutrino mass splitting (between m_1 and m_2) is much smaller than the atmospheric neutrino mass splitting (between m_2 and m_3). However, sign of the former is known, while that of the latter is not.
- The degree of CP-violation in the neutrino sector, though preliminary evidence suggests a departure from the fully-CP-conserving scenario.
- The absolute neutrino mass scale, though β -spectrum endpoint measurements and limits from cosmology both suggest it is well below the eV scale — six orders of magnitude smaller than the next lightest fermion.
- The origin of the neutrino mass.

The last point will be the focus of the following section.

2.2 Massive Neutrinos

It should now be evident that the SM must be extended to incorporate the new, massive neutrino paradigm. In what follows, two standard approaches to neutrino mass origin are considered. For simplicity, I will ignore flavor mixing and consider only a single generation.

2.2.1 The Effective Field Theory Approach

To take a model-independent approach, we can use the Effective Field Theory (EFT) framework, and consider non-renormalizable operators parameterizing beyond-the-SM (BSM) physics that may give rise to neutrino masses at low energy. Fortunately, there is just such an operator — the coupling of ν to Φ at energy scale Λ_5 , given by

$$\mathcal{L}_{\text{EFT}} \supset -\frac{1}{\Lambda_5} \Phi \Phi \nu \nu, \quad (2.12)$$

is allowed by the gauge symmetries of the SM. Indeed, this is the first non-renormalizable term to be omitted from the SM, as it is the only allowed dimension-5 term. It is known as the Weinberg operator [64]. When the Higgs acquires a vev, this interaction gives the neutrino a Majorana mass $m_\nu \sim v^2/(2\Lambda_5)$. This approach yields a simple neutrino mass mechanism without the introduction of any new particle content.




Figure 2.4: Feynman diagram showing the Weinberg operator, which describes the interaction of two ν and two Φ .

2.2.2 Adding Right-Handed Neutrinos

An alternative approach is to suppose that neutrinos get a mass in the same way as all the other fermions: via a Yukawa coupling joining the left- and right-handed fields. We can therefore introduce a new right-handed neutrino, \bar{N} . It is trivial to determine the gauge charges of \bar{N} , as since right-handed neutrinos have never been observed experimentally, they must be singlets under all three gauge groups. The usual Yukawa coupling term is then allowed by gauge symmetries, and as before, when the Higgs acquires a vev, the neutrino gets a Dirac mass,

$$\mathcal{L}_\nu \supset -y_{\nu i}^j \Phi \ell_i \bar{N}^j \xrightarrow{\text{EWSB}} -\frac{y_{\nu i}^j v}{\sqrt{2}} \ell_i \bar{N}^j. \quad (2.13)$$

Now, recall that the prescription for building the SM is to write down all renormalizable terms that respect the gauge symmetries and Lorentz invariance given the field content. We have now added to the field content of the theory and therefore must ensure we have included all possible terms for the added field. It turns out that there is an additional term that is allowed: a Majorana mass term for \bar{N} , given by

$$\mathcal{L}_\nu \supset -m_M \bar{N} \bar{N}. \quad (2.14)$$

Introducing \bar{N} gives *two* new mass terms: a Dirac mass, $m_D \equiv y_\nu v / \sqrt{2}$, and a Majorana mass, m_M . Moreover, the inclusion of both terms provides a natural reason for the small neutrino mass scale, as described below.

2.2.3 The Seesaw Mechanism

With the introduction of \bar{N} , the neutrino mass portion of the Lagrangian becomes,

$$-\mathcal{L}_\nu = \frac{m_D}{2} \nu \bar{N} + \frac{m_M}{2} \bar{N} \bar{N} + \text{h.c.}, \quad (2.15)$$

with a Dirac mass coupling ν and \bar{N} to one another and a Majorana mass coupling \bar{N} to itself. There is no Majorana mass term for ν , as it is forbidden by gauge

symmetry⁶. These mass terms can be rewritten in matrix form as

$$-\mathcal{L}_\nu = \frac{1}{2} \begin{pmatrix} \nu & \bar{N} \end{pmatrix} \begin{pmatrix} 0 & m_D \\ m_D & m_M \end{pmatrix} \begin{pmatrix} \nu \\ \bar{N} \end{pmatrix} + \text{h.c.} \quad (2.16)$$

The physical mass eigenstates can then be obtained by diagonalizing the mass matrix. The eigenvalues are,

$$m_{R/L} = \frac{m_M \pm \sqrt{m_M^2 + 4m_D^2}}{2}, \quad (2.17)$$

with subscripts R or L denoting right- or left-handed, respectively, for a reason that will become clear shortly. It is natural to take $y_\nu \sim 1$ so that $m_D \sim v \sim 10^2$ GeV and the Dirac mass originates from the electroweak scale. By contrast, m_M is likely to originate from some as-yet-unseen high energy physics. Suppose we take this to be at the scale of grand unification, $m_M \sim m_{\text{GUT}} \sim 10^{16}$ GeV. Since $m_M \gg m_D$, the mass eigenvalues can be approximated by

$$m_L \simeq \frac{m_D^2}{m_M} \quad \text{and} \quad m_R \simeq m_M. \quad (2.18)$$

This is known as the *Type I Seesaw Mechanism*, since one mass eigenvalue goes up as the other goes down [66–68]. Plugging in the numbers from above, we find $m_L \sim 1$ meV. This simple model, built from a minimalistic extension to the SM and using natural estimates of the energy scales involved, produces a state near the true neutrino mass scale. The corresponding eigenstates can be written as,

$$\nu_L \simeq \nu - \frac{m_D}{m_M} \bar{N} \quad \text{and} \quad \nu_R \simeq \bar{N} + \frac{m_D}{m_M} \nu, \quad (2.19)$$

⁶Adding a Higgs triplet to the SM allows for ν to have a Majorana mass. This condition gives rise to what is known as the *Type II Seesaw Mechanism* [65].

meaning that ν_L is mostly the left-handed neutrino, ν , and ν_R is mostly the right-handed neutrino, \bar{N} . The diagonalized neutrino mass term can then be expanded,

$$-\mathcal{L}_\nu = \frac{1}{2} \begin{pmatrix} \nu_L & \nu_R \end{pmatrix} \begin{pmatrix} m_D^2/m_M & 0 \\ 0 & m_M \end{pmatrix} \begin{pmatrix} \nu_L \\ \nu_R \end{pmatrix} + \text{h.c.} \quad (2.20)$$

$$= \frac{m_D^2}{2m_M} \nu_L \nu_L + \frac{m_M}{2} \nu_R \nu_R + \text{h.c.}, \quad (2.21)$$

revealing that both physical mass eigenstates ν_L and ν_R have a Majorana mass only. It is in this sense that the neutrino is a Majorana particle despite the constituent fields having both a Dirac and a Majorana mass term⁷.




Figure 2.5: The interaction resulting from both a Dirac mass term for ν and \bar{N} and a Majorana mass term for \bar{N} . The resulting mass state is much lighter than either m_D or m_M via the Seesaw Mechanism. In the limit where the $m_M \gg m_D$, the three vertices converge to a point, and the diagram in Fig. 2.4 is recovered.

It is worth pointing out that this gives identical phenomenology to the EFT approach: we can identify the cutoff Λ_5 as the Majorana mass m_M . This specific approach is captured by the effective operator, as expected.

Since the two minimalistic extensions to the SM discussed here both generate Majorana masses, it is worth considering how they could be integrated into the current neutrino paradigm. The only change required to the paradigm described above is the

⁷In the limit that $m_M \rightarrow 0$, the two Majorana states are maximally mixed and have the same mass m_D , allowing them to be expressed as a single Dirac spinor, as expected.

addition of two new phases η_1 and η_2 to the neutrino mixing matrix:

$$U_{\text{PMNS}}^{\text{Majorana}} = U_{\text{PMNS}} \cdot \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\eta_1} & 0 \\ 0 & 0 & e^{i\eta_2} \end{pmatrix}. \quad (2.22)$$

The usual mixing phenomenology is unaffected by these additions. The same is true for the β -spectrum endpoint measurements, by the definition of m_β . Since the Majorana phases leave no experimental signature in oscillation data or β -spectrum endpoint measurements, a different experimental probe is required to verify the neutrino's Majorana nature.

2.3 From Theory to Experiment

The claim that “a neutrino is its own antiparticle” leads naturally to the question of why an “antineutrino” has never been observed interacting as a “neutrino” (e.g., through the $\nu_e + n \rightarrow p + e^-$ detection channel). In principle, if neutrinos are Majorana particles, this is indeed allowed. But antineutrinos are almost always ultrarelativistic⁸ with right-handed helicity. This results in a suppression by $m^2/E^2 \sim 10^{-14}$ of the left-chiral component — the one responsible for enabling neutrino-like interactions. This suppression, combined with the inherently small interaction cross section, renders such a detection practically impossible. Fortunately, though, nature has provided an ideal laboratory for studying processes of this kind: the atomic nucleus.

2.3.1 Two-Neutrino Double Beta Decay

β -decay occurs when a nucleus (A, Z) decays to a final state $(A, Z + 1)$. This is only possible when the daughter nucleus has a lower binding energy than the parent. Since nuclear binding energy is lowered by spin coupling between nucleons in the same energy level, even-even nuclei (those with an even number of both protons and

⁸An exception is the cosmic neutrino background — neutrinos which decoupled from thermal equilibrium in the early Universe, and have since been redshifted to $\sim 10^{-5}$ eV. Unfortunately, this makes them effectively impossible to detect with current technology.

neutrons) tend to have lower binding energy than other isobars. In such cases, β -decay can be energetically unfavorable, while the process $(A, Z) \rightarrow (A, Z + 2)$ may not be. This process of two simultaneous β decays is called double beta decay ($\beta\beta$ -decay), or sometimes two-neutrino double beta decay ($2\nu\beta\beta$) to distinguish it from the neutrinoless alternative discussed below. In this process, two neutrons decay to two protons, releasing two electrons and two antineutrinos. It was first described by Maria Goeppert-Mayer in 1935 [69]. It is exceedingly rare, with typical half lives of $\sim 10^{20}$ years [70].

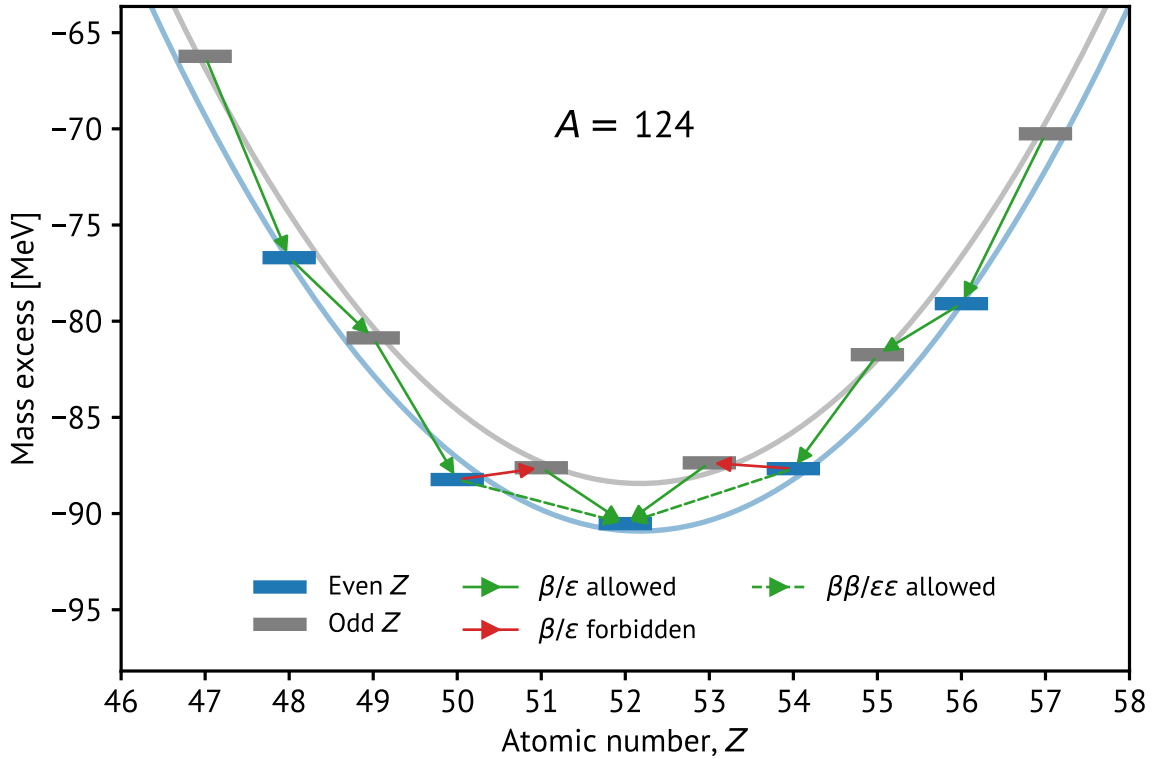


Figure 2.6: Mass excess parabolas for $A = 124$ showing data from Refs. [71, 72]. The horizontal steps represent the measured mass excess and the parabolas are fits. There is a vertical offset between the mass parabolas for nuclei with odd and even Z . β -decay or electron capture (ϵ -decay) can change Z by one unit if it is energetically allowed, i.e. if the Q -value is positive (note that selection rules may still prohibit or suppress these decays). In cases where β -decay or ϵ -decay is energetically prohibited, $\beta\beta$ -decay or double electron capture ($\epsilon\epsilon$ -decay) may still occur.

2.3.2 Neutrinoless Double Beta Decay

$\beta\beta$ -decay is intriguing from the perspective of searching for Majorana neutrinos. If a neutrino is its own antiparticle, then there is a possibility of it being emitted at one vertex and absorbed at the other (Fig. 2.8), leaving the electrons to take all the energy (Fig. 2.7). This hypothesized process is called neutrinoless double beta decay ($0\nu\beta\beta$), first proposed by Wendell Furry in 1939 [73], and it has yet to be observed. The particular Feynman diagram for $0\nu\beta\beta$ shown in Fig. 2.8b is called the





Figure 2.7: Energy of the emitted electrons in $2\nu\beta\beta$ and $0\nu\beta\beta$. In the former, the electrons share the energy with the antineutrinos, and the spectrum is continuous; in the latter, the electrons carry away the full decay energy. E_1 and E_2 are the kinetic energies of the two electrons, while $Q_{\beta\beta}$ is the Q -value of the decay. The relative scaling of the two distributions is arbitrary.

light neutrino exchange model, because the process joining the two β -decay diagrams involves exchange of the light neutrino state ν_L . More generally, though, the process does not need to be mediated by light neutrinos. In fact, it need not involve neutrinos at all⁹. The Schechter-Valle theorem, commonly called the Black Box Theorem, holds

⁹Other mechanisms, including so-called Left-Right Symmetric Models, include right-handed weak currents which could give rise to neutrinoless double beta decay [74, 75].



(a) Two-neutrino double beta decay ($2\nu\beta\beta$) (b) Neutrinoless double beta decay ($0\nu\beta\beta$)

Figure 2.8: Feynman diagrams showing both two-neutrino and neutrinoless double beta decay. In $2\nu\beta\beta$ (left), two neutrons decay to two protons and two electrons, accompanied by two antineutrinos. In $0\nu\beta\beta$ (right), the same process occurs without the neutrinos, resulting in the emission of only two electrons.

that if $0\nu\beta\beta$ occurs, then whatever operator is responsible for it also *induces* a non-zero Majorana mass for the neutrino¹⁰ (Fig. 2.9) [77, 76]. An observation of $0\nu\beta\beta$ would therefore 1) establish definitive evidence of lepton number non-conservation, and 2) prove that neutrinos are Majorana fermions.

2.3.3 The Decay Rate

While $2\nu\beta\beta$ is itself rare, $0\nu\beta\beta$ (if it exists) is even more so. The half life of a nucleus which can undergo $0\nu\beta\beta$ is given by

$$[T_{1/2}^{0\nu}]^{-1} = \frac{\langle m_{\beta\beta} \rangle^2}{m_e^2} G_{0\nu} g_A^4 |\mathcal{M}^{0\nu}|^2, \quad (2.23)$$

where m_e is the electron mass, $\langle m_{\beta\beta} \rangle$ is the effective Majorana mass, $G_{0\nu}$ is the kinematic phase space factor, g_A^4 is the axial-vector coupling constant¹¹, and $\mathcal{M}^{0\nu}$ is

¹⁰This contribution to the Majorana mass is at the level $\sim 10^{-28}$ eV [76], so while it *technically* implies that neutrinos are Majorana particles, it may not have any phenomenological consequences.

¹¹It is often assumed that $g_A \approx 1.27$, as has been observed in free neutron β -decay; the necessity of applying quenching to this value is a topic of ongoing debate [78].




Figure 2.9: The operator which causes $0\nu\beta\beta$ (left) need not involve Majorana neutrinos. By crossing symmetry, however, if such an operator exists, it would contribute an effective Majorana mass for neutrinos at the four-loop level (right), enabling particles to oscillate into antiparticles and vice versa. This result is known as the Black Box Theorem.

the nuclear matrix element. $G_{0\nu}$ and $\mathcal{M}^{0\nu}$ are both isotope dependent; the former can be calculated to high precision [79–81], while the latter is subject to considerable uncertainties. Several techniques have been employed and result in estimates that span more than a factor of 3 for each isotope; see Ref. [82] for a recent review. The effective Majorana mass is defined as

$$\langle m_{\beta\beta} \rangle = \left| \sum_{i=1}^3 U_{ei}^2 m_i \right|. \quad (2.24)$$

The sensitivity of an experiment to $0\nu\beta\beta$ is often reported in terms of the lowest $\langle m_{\beta\beta} \rangle$ value to which it is sensitive. It is common to show the parameter space on a plot of $\langle m_{\beta\beta} \rangle$ as a function of the lightest neutrino mass (m_1 under NO or m_3 under IO) or of the sum of masses, Σ . This is the so-called “lobster” plot. Fig. 2.10 shows both versions of the lobster plot for both ordering scenarios. The yellow shaded areas enclosed by the black curves are regions of the parameter space allowed at 3σ by the measured oscillation parameters, while the color scale shows the Bayesian posterior probability density in the allowed regions (see Appendix A for a description of the calculation). The color scale is typically omitted as it relies on a choice of priors and NME values, though its inclusion here is intended as a reminder that the allowed regions should not be implicitly interpreted as flat probability distributions.

For a portion of the parameter space, the allowed regions under both IO and NO are degenerate. In the IO scenario, the entire $0\nu\beta\beta$ parameter space can be excluded by any experiment with sensitivity to $\langle m_{\beta\beta} \rangle \sim 10$ meV. This is in stark contrast to the NO scenario, where a sensitivity of $\langle m_{\beta\beta} \rangle \sim 1$ meV covers all the parameter space but the tail from $m_1 \sim 1$ meV to 10 meV. This tail arises due to the possibility that η_1 and η_2 conspire to orchestrate a precise cancellation of the three terms in Eq. (2.24).

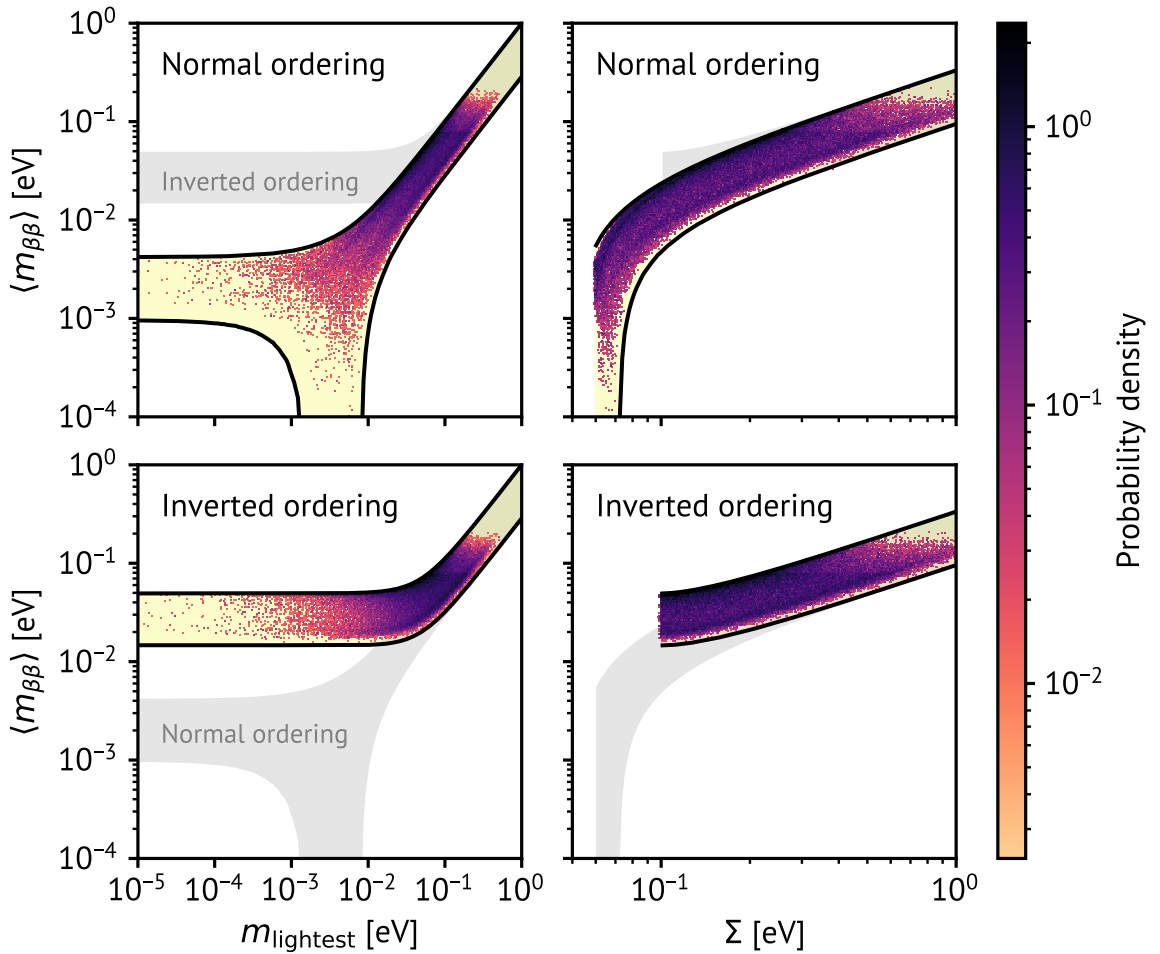


Figure 2.10: The lobster plot showing the allowed parameter space for the effective Majorana mass. The yellow shaded regions bounded by black lines indicate the parameter space allowed at 3σ by a recent global fit to oscillation data. The color scale shows the posterior distribution in this parameter space. For each ordering scenario, the other scenario is shown in light gray as a point of reference.

2.3.4 The Experimental Program

There are 35 isotopes that can undergo $2\nu\beta\beta$ [83], and therefore, in principle, can undergo $0\nu\beta\beta$ as well if $\langle m_{\beta\beta} \rangle > 0$. Which of these to use for an experimental search depends on a number of considerations. Two worth noting here are 1) the decay Q -value, $Q_{\beta\beta}$, and 2) the natural abundance. The first is important because, to a good approximation, the phase space factor (and by consequence, the decay rate) scales with $Q_{\beta\beta}^5$ [84]. There is a more concrete experimental reason to prefer isotopes with high Q -values as well: above a few MeV, the number of potential radioactive backgrounds diminishes quickly with increasing energy. The most problematic backgrounds originate from daughters of the ^{238}U and ^{232}Th decay chains due to trace concentrations of these isotopes in detector materials.

The second consideration, high natural abundance, is also important for experimental feasibility. Maximizing the decay rate requires an instrumented volume that includes as many atoms as possible. For solid targets, this often requires growing crystals to sufficiently large sizes. Isotopic enrichment allows for significant concentrations of the $\beta\beta$ -decaying isotope to be instrumented, though the procurement and enrichment processes can pose significant challenges in themselves. Fig. 2.11 shows the Q -value of the decay as a function of the natural abundance for all isotopes that can undergo $2\nu\beta\beta$. β and γ backgrounds from the aforementioned decay chains are also shown. α decays have been omitted, as they tend to be easier to identify and reject.

Many experiments have searched for $0\nu\beta\beta$ (see Refs. [88–90] for recent reviews); here, I describe a selection of these to highlight the technologies that will be used in future large-scale experiments. The current generation has begun probing the IO band, and next generation searches will be sensitive to nearly all of it. Fig. 2.12 summarizes the sensitivities (both demonstrated and projected) of these experiments to $\langle m_{\beta\beta} \rangle$.

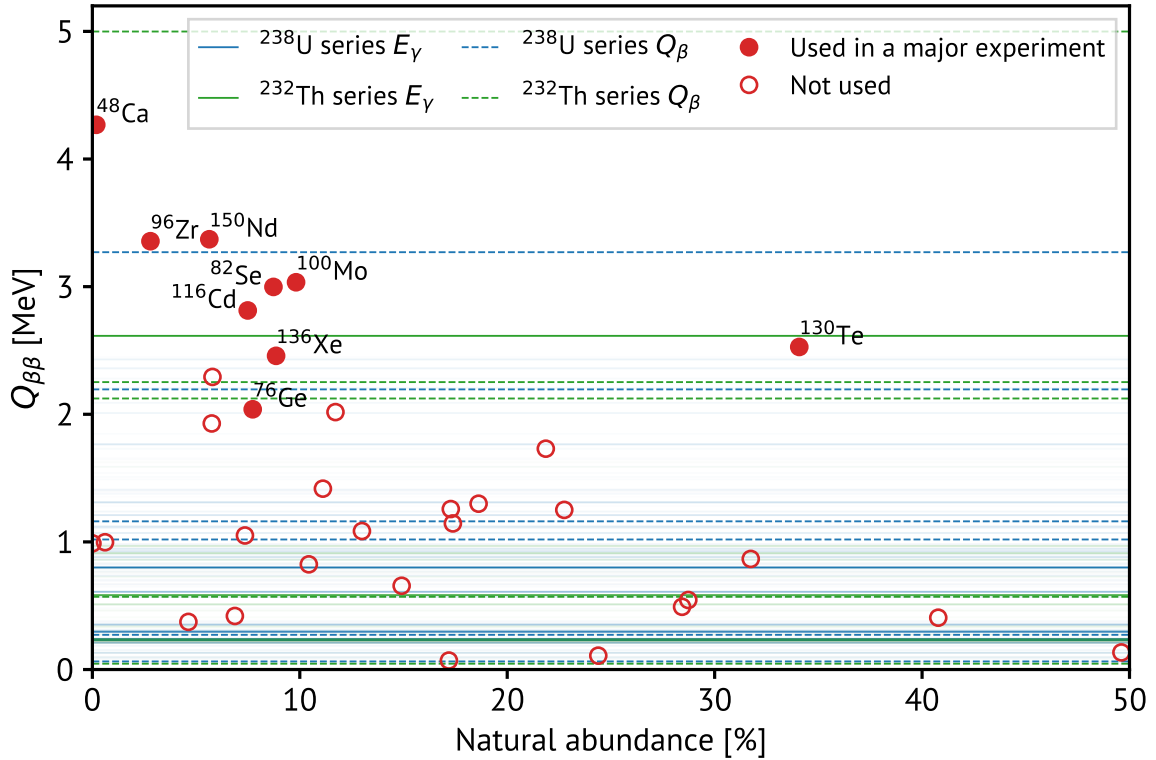


Figure 2.11: Q -values for the double beta decay of the 35 isotopes in which the process is allowed, plotted against their natural abundance [85]. Filled circles indicate isotopes that are currently or have previously been used in major $0\nu\beta\beta$ searches. Also shown are the β -decay Q -values and the γ energies for decays in the ^{238}U and ^{232}Th decay chains [86, 87]. The opacity of the γ lines reflects their relative intensity. Only γ s (β s) with an intensity (branching ratio) of $> 1\%$ are shown.

Semiconductor Detectors

High-purity germanium (HPGe) detectors, used commonly in γ spectroscopy, can be fabricated from crystals enriched in ^{76}Ge . The growth process can yield extremely pure crystals, and these detectors are known to achieve excellent energy resolution. Current ^{76}Ge -based experiments include the Germanium Detector Array (GERDA) and the MAJORANA DEMONSTRATOR, both of which use point-contact detectors enabling pulse shape discrimination to distinguish signals from multisite and α backgrounds [91, 92]. HPGe detectors must be operated at cryogenic temperatures due to the low band gap. GERDA uses crystals immersed in a liquid argon (LAr) cryostat,




Figure 2.12: The 3σ discovery sensitivity (demonstrated and projected) of a selection of experiments to $0\nu\beta\beta$. The vertical extent of the boxes shows the range of $\langle m_{\beta\beta} \rangle$ corresponding to the reported half life based on different calculations of the NMEs. Color indicates the isotope used in each experiment. Experiments that are currently running or have run in the past are shown as filled boxes, while experiments that are in the planning or construction phases are shown as unfilled boxes. MJD stands for MAJORANA DEMONSTRATOR and KLZ stands for KamLAND-Zen. Data in this plot is taken from Tab. IV of Ref. [89].

enabling scintillation in the LAr to be monitored for background rejection. These two experiments have merged to form the Large Enriched Germanium Experiment for Neutrinoless Double Beta Decay (LEGEND) Collaboration. LEGEND-200 is currently taking data, while the tonne-scale LEGEND-1000 experiment anticipates first data beginning in the early 2030s [93, 94].

Bolometers

Bolometers measure energy depositions in cryogenic (~ 10 mK) crystals via minute increases in temperature. Like HPGe detectors, they benefit from clean detector materials due to the crystal growth process. The Cryogenic Observatory for Underground

Rare Events (CUORE) experiment uses towers of TeO_2 crystals in a dilution refrigerator, instrumented with Neutron Transmutation Doped (NTD) germanium thermistors, to search for the $0\nu\beta\beta$ of ^{130}Te [95]. The CUORE Upgrade with Particle ID (CUPID) experiment plans to use scintillating bolometers consisting of Li_2MoO_4 crystals enriched in ^{100}Mo , instrumented with germanium wafer light sensors, to search for $0\nu\beta\beta$ with improved discrimination of α backgrounds [96]. The Advanced Molybdenum Based Rare Process Experiment (AMoRE) uses CaMoO_4 crystals to search for the $0\nu\beta\beta$ of ^{100}Mo [97]. The crystals are enriched in ^{100}Mo to maximize the signal, but depleted in ^{48}Ca to minimize backgrounds from the high- Q -value $2\nu\beta\beta$. Metallic magnetic calorimeters are used to detect phonons, both from energy depositions in the crystal and from scintillation light absorbed in a germanium wafer coupled to the crystal. AMoRE is currently taking data with half of the eventual target mass.

Scintillator Detectors

Liquid scintillator (LS) can be loaded with a $\beta\beta$ -decaying isotope to form a highly scalable combined decay source and detection medium. The SNO+ experiment is using existing infrastructure from SNO in a multiphase neutrino physics program [98]. In the $0\nu\beta\beta$ phase, tellurium will be loaded into a LS cocktail, enabling the wavelength-shifted scintillation light from $0\nu\beta\beta$ in ^{130}Te to be detected by the surrounding SNO photomultiplier tubes (PMTs) [99]. Similarly, the KamLAND-Zen experiment uses a LS cocktail loaded with xenon enriched in ^{136}Xe . The xenon-loaded LS is contained in a balloon which sits in a mineral oil buffer surrounded by an array of PMTs. KamLAND-Zen currently has the most stringent constraints on $0\nu\beta\beta$ in ^{136}Xe , and plans to increase the target mass from 320 kg to 1000 kg in the second iteration, KamLAND2-Zen, along with other upgrades [100, 101]. $\beta\beta$ -decaying isotopes can be combined with solid scintillators as well. The CANDLES experiment uses scintillating CaF_2 crystals in a LS vessel to search for the $0\nu\beta\beta$ of ^{48}Ca [102]. Scintillation light from the crystals and the LS is read out by PMTs coupled to light pipes in the vessel.

Tracking Calorimeters

Unlike many experiments mentioned so far, tracking calorimeters separate the decay source from the detection medium. The source isotope is contained in a foil, on either side of which is a gas tracking region and then a calorimeter. A magnetic field through the tracking region allows for both electron tracks to be identified and reconstructed, while the calorimeters provide a measurement of the energy. The NEMO program has demonstrated this technology using ^{100}Mo , with smaller quantities of several other isotopes as well [103].

Time Projection Chambers

Chapter 4 describes the physics of time projection chambers (TPCs) in detail, so they are introduced here only briefly. TPCs use the prompt arrival of a scintillation signal, combined with the delayed arrival of a drifting charge signal, to reconstruct the three-dimensional position and energy of events in the active volume. The ratio of scintillation to ionization can be used for particle discrimination. Like LS-based detectors, TPCs for $0\nu\beta\beta$ are homogeneous detectors with a combined decay source and detection medium, a feature which lends itself to scalability. The Neutrino Experiment with a Xenon TPC (NEXT) uses high pressure xenon gas enriched in ^{136}Xe as the source/medium for their $0\nu\beta\beta$ searches [104, 105]. Operating in the gas phase enables the reconstruction of individual electron tracks and the identification of track endpoints, allowing improved background discrimination based on event topology. Liquid-phase detectors benefit from a much higher density of the decay source at the expense of full electron track reconstruction. EXO-200, the flagship experiment of the Enriched Xenon Observatory (EXO) program, used liquid xenon enriched in ^{136}Xe to search for $0\nu\beta\beta$ and reported the first detection of $2\nu\beta\beta$ in that isotope [106, 107]. nEXO, the successor to EXO-200, is the subject of Chapters 4 to 6.

2.3.5 Implications of a Discovery

A discovery of $0\nu\beta\beta$ would have profound implications for particle physics. It would provide a natural explanation for the extraordinarily small neutrino mass scale. It

would hint at the possibility of new physics arising from the grand unification scale — far beyond anything accessible in state-of-the-art accelerator facilities, or even those that could be constructed in the foreseeable future. It would demonstrate lepton number non-conservation via the creation of two leptons. Relatedly, it would address one of the most pressing questions in modern cosmology — why there is more matter than antimatter in the Universe. The creation of two leptons, unaccompanied by the creation of two antileptons, would provide a pathway for the observed matter asymmetry via a process known as leptogenesis¹² [108, 109]. In addition, certain classes of models explain dark matter via heavy right-handed neutrinos of the kind present in the Seesaw Mechanism [110]. A discovery of $0\nu\beta\beta$ would lend support to Seesaw-based models that serve to explain both dark matter and the neutrino mass.

To conclude, this chapter can be summarized as follows:

- SM fermions acquire masses via the Yukawa interaction between the Higgs field and the left- and right-handed fermion fields.
- When the SM was established, neutrinos were thought to be massless, though recent evidence has demonstrated that this is incorrect.
- Unlike other fermions, neutrinos can have both Majorana and Dirac masses, and natural extensions to the SM often give rise to Majorana masses.
- An observation of $0\nu\beta\beta$ would be evidence of a non-zero effective Majorana mass and BSM physics.
- Global fits to oscillation, β -decay, and cosmological data suggest that a first detection of $0\nu\beta\beta$ may be near or within the reach of next generation searches.

After a brief detour into another shortcoming of the SM, I will introduce nEXO, a proposed next-generation experiment with the aim of discovering $0\nu\beta\beta$.

¹²The decay $\bar{N} \rightarrow \ell\Phi$ in the early Universe can lead to a baryon asymmetry, provided that 1) it occurs out of thermal equilibrium, so that the asymmetry is not washed out by the reverse process; 2) CP violation in the Yukawa couplings causes the rate of $\bar{N} \rightarrow \ell\Phi$ to differ from that of $\bar{N} \rightarrow \ell^\dagger\Phi^\dagger$; and 3) an electroweak sphaleron process converts the lepton asymmetry to a baryon asymmetry while preserving $B - L$.

Chapter 3

Gravity and New Physics at the Micron Scale

Gravity is arguably the most tangible of the fundamental forces due to its ubiquity at macroscopic scales. Four centuries of planetary motion data has established the utility of Newton’s theory of universal gravitation at the scale of the Solar System, while precision tests of General Relativity (GR) have validated its applicability in the strong gravity regime [111, 112]. At shorter length scales, however, the nature of gravity remains unexplored. The gravitational force has never been measured at separations of less than $52\ \mu\text{m}$, despite the implicit assumption that GR holds all the way down to the Planck length — a gap of more than 30 orders of magnitude. In this chapter, theoretical motivations for the study of gravity at the micron scale are developed, and the experimental landscape is introduced.

3.1 Theoretical Considerations

3.1.1 The Hierarchy Problem

One feature of gravity that poses a challenge to the SM is its extraordinary weakness. This can be quantified in terms of the energy scale at which gravitational effects become relevant. Gravity can be considered “strong” in the vicinity of the

Schwarzschild radius, $r_s = 2Gm$, of a source mass, m . When the reduced Compton wavelength of the source mass, $\lambda_c = 1/m$, is of the order of r_s , gravitational effects cannot be ignored. This occurs at the Planck scale, M_{Pl} , where

$$M_{\text{Pl}} \equiv \sqrt{\frac{1}{G}} \approx 10^{19} \text{ GeV}. \quad (3.1)$$

Recall that the electroweak scale, given approximately by v , is $\sim 10^2 \text{ GeV}$ — 17 orders of magnitude smaller. This enormous difference of scales is known as the gauge hierarchy problem [113]. Returning to the EFT picture, we can see why this is termed a “problem,” rather than just a noteworthy empirical observation. In the EFT picture, the SM is taken to be an effective descriptor up to some cutoff scale, Λ . The sensitivity of measurable quantities in the low-energy theory to high-energy physics can then be expressed in terms of Λ . Of particular concern is the Higgs mass, m_{H} , which consists of a bare mass term, m_{H}^0 , and quantum corrections, δm_{H} :

$$m_{\text{H}}^2 = (m_{\text{H}}^0)^2 + \delta m_{\text{H}}^2. \quad (3.2)$$

The problem is apparent when considering even a single term in δm_{H}^2 ,

$$\delta m_{\text{H}}^2 = \frac{\Lambda^2}{16\pi^2} (-6y_t^2 + \dots), \quad (3.3)$$

where y_t is the top quark Yukawa coupling. The mass-squared correction term is quadratic in Λ , so one should expect the Higgs mass to be of order Λ as well. The fact that it is 17 orders of magnitude smaller is mathematically allowed, but it requires a delicate cancellation between m_{H}^0 and δm_{H} . While the particular manifestation of the hierarchy problem shown here depends on the regularization scheme, the general problem — sensitivity of the Higgs mass to high-energy physics — does not [114].

3.1.2 Potential Resolutions

There are a number of hypothesized solutions to the hierarchy problem which can broadly be classified into two groups:

1. The existence of a new symmetry which results in cancellations in δm_{H}^2
2. A cutoff at near the electroweak scale rather than the Planck scale

Supersymmetry (SUSY) is a popular solution which falls into the first category. Embedding the SM in a larger symmetry group which associates to each fermion (boson) a bosonic (fermionic) superpartner results in the addition of opposite-sign counterterms to cancel the existing terms. Continuing with the example above, the superpartner of the top quark, the “stop squark”, adds an opposite-sign term to δm_{H}^2 , canceling the contribution from the top quark. In reality, SUSY must be broken since the stop squark is constrained to be significantly heavier than the top quark; nonetheless, the Higgs mass is protected from these quadratically-divergent quantum corrections. Still, the lack of observation of any superpartners from collider searches up to the TeV scale requires some fine-tuning even if the SUSY scale were just beyond experimental reach.

In the second category, composite Higgs models suggest that the Higgs is not fundamental, and the appropriate cutoff scale is at the energy scale at which its constituents could be separated. Another solution in the second category, worthy of mention here due to its relevance to Chapter 3, brings down the cutoff scale by postulating the existence of large extra dimensions.

3.1.3 Large Extra Dimensions

The Arkani-Hamed–Dimopoulos–Dvali (ADD) model is an illustrative example of large extra dimensions (LED) as a solution to the hierarchy problem [115]. The motivation for large extra dimensions can be seen from Gauss’s law: in a theory with D spatial dimensions, the gravitational potential from a point mass scales with radial distance r as $V \sim r^{2-D}$ (for $D > 2$). The gravitational flux is *diluted* by additional dimensions — the existence of extra dimensions could therefore explain the apparent weakness of gravity¹. Evidently, our macroscopic world consists of only three spatial dimensions and one time dimension, so the extra dimensions, if they exist, must

¹This requires that the remaining fundamental interactions remain confined to the 3+1 spacetime dimensions

be hidden. One way for extra dimensions to hide is through compactification. In this scenario, the extra dimensions are curled up in a small, closed manifold with a specific geometry, making them unobservable at macroscopic distances. If there are n extra dimensions all compactified in circles with the same radius, R , then the gravitational potential energy for a configuration of point masses m_1 and m_2 separated by r becomes²:

$$V(r) \sim \begin{cases} \frac{m_1 m_2}{M_D^{n+2}} \frac{1}{r^{n+1}}, & r \ll R, \\ \frac{m_1 m_2}{M_D^{n+2} R^n} \frac{1}{r}, & r \gg R. \end{cases} \quad (3.4)$$

where M_D is the real energy scale of gravity. At macroscopic scales, we get the familiar r^{-1} scaling, while Newton's gravitational constant G can be identified as $M_D^{-n-2} R^{-n}$. At short length scales, the r^{-1} scaling would be replaced with r^{-1-n} , and the "true" gravitational constant M_D^{-n-2} would emerge.

Across all scales, the gravitational force is mediated by the exchange of massless spin-2 bosons called gravitons. The energy-momentum relation for a graviton in extra dimensions can be written as

$$E^2 = \sum_{i=1}^3 p_i^2 + \sum_{i=4}^{3+n} p_i^2, \quad (3.5)$$

where the first sum is over the usual 3-momentum components and the second is over the momentum components in the n extra dimensions. A comparison to the usual energy-momentum relation,

$$E^2 = \sum_{i=1}^3 p_i^2 + m^2, \quad (3.6)$$

reveals that momentum components in the extra dimensions take the place of the mass term and therefore make the graviton appear massive. Moreover, when boundary conditions with period R are applied in the extra dimensions to account for compactification, the resulting modes comprise an infinite spectrum of discrete states,

²There are compactification-dependent prefactors that account for the surface area and volume of the compactified geometry which are omitted here.

known as Kaluza-Klein (KK) modes, with masses of

$$m_j = \frac{j}{R}, \quad (3.7)$$

where j is a positive integer. The large- r (small- r) limit of the potential above can be interpreted as the case in which none (all) of the massive states contribute to the interaction. When $r \sim R$, only the lightest state in each extra dimension contributes, and the potential takes on a Yukawa form [116, 117],

$$V(r) \sim \frac{m_1 m_2}{M_D^{n+2} R^n} \frac{1}{r} (1 + 2n e^{-r/R}) \quad (3.8)$$

where the factor of 2 accounts for the degeneracy of the lightest state in each dimension.

This model can be tested in multiple ways [118]. Collider experiments can search for the production of KK modes or quantum black holes [119–121]. Small-scale experiments can search for deviations from the inverse-square law (ISL), either of the form in Eq. (3.8) at length scales near R , or of the form in Eq. (3.4) at much shorter length scales [122, 123]. Collider searches have already tightly constrained the $n = 2$ scenario; constraints are slightly weaker for larger n [124–133]. Nonetheless, the ADD model is a single illustrative example out of many models which introduce extra dimensions.

3.1.4 Other Theoretical Motivations

The ADD model is a useful starting point for this discussion as a representative example of a larger class of theories. Extra dimensions with different compactifications can give rise to similar short-range interactions. It is possible that the size of extra dimensions fluctuates across spacetime. This would necessitate a massless scalar field known as a “radion” that parameterizes the radii of the extra dimensions. In order to prevent the radion from mediating long-range forces that have not been observed, a stabilization mechanism which gives the radion a mass is required. This is a general feature of extradimensional models, including the ADD model and those discussed

below.

In the Randall-Sundrum (RS-I) model [134], a single extra dimension is introduced, with the SM residing on a (3+1)-dimensional brane (called the TeV brane) at one boundary of the 5-dimensional anti-de Sitter space (AdS) on which gravity propagates. A second (3+1)-dimensional brane (called the Planck brane) sits at the other boundary, separated from the TeV brane by πr_c , with r_c the compactification radius. If the full geometry is warped along the extra dimension, then the gravitational energy scale in the TeV brane can be exponentially suppressed, leading to a naturally large hierarchy. This model includes massive KK modes and a radion which can become massive as a result of a stabilization mechanism [135]. These are typically too heavy to be detected via short-range gravity experiments but could be visible at colliders [136, 137]. A second model by the same authors (RS-II) takes r_c to infinity, giving rise to infinite-volume extra dimensions [138] and again producing massive KK gravitons. Other models based on warped extra dimensions can include ultralight radions that could be detectable in fifth-force experiments [139].

String theory also motivates the existence of extra dimensions and new short-range interactions; in fact, the ADD model can be embedded within string theory [140]. A recent model inspired by the swampland program [141, 142] introduces a single, micron-scale extra dimension and suggests that the resulting tower of KK modes could comprise the observed dark matter [143].

More generally, string theory vacua are known to contain many massless scalar fields known as moduli, whose vevs parameterize the compactification [144]. If the particular theory includes SUSY broken at low energy, the moduli can become massive, albeit extremely light. This makes them appealing targets for short-range force searches. Of these, the dilaton, which parameterizes the strength of gauge couplings, is among the most well understood.

Extra dimensions aside, generic light scalars can result in similar phenomenology and are worthy of a brief mention. A new scalar field ϕ can be added to the SM with any of the following Yukawa couplings,

$$\frac{m_e}{f} \phi \bar{e}e, \quad \frac{m_d}{f} \phi \bar{d}d, \quad \frac{m_u}{f} \phi \bar{u}u, \quad \frac{1}{f} \phi G_{\mu\nu}^a G^{a\mu\nu} \quad (3.9)$$

where f is the mass scale of the coupling and $G_{\mu\nu}^a$ is the gluon field strength tensor. The first three terms are protected by chiral symmetry, so the couplings should be limited to the chiral-symmetry-breaking scale; i.e., that of the fermion masses. The gluon coupling scale, which is of the order of the nucleon mass scale, will then dominate over these. If the mass of the scalar is m_ϕ , the resulting Yukawa interaction will have a characteristic length scale of $\lambda = 1/m_\phi$.

In addition to those mentioned, there are many other theoretical models that give rise to new short-distance physics. See Ref. [122] for a comprehensive review.

3.2 Experimental Tests

Given the breadth of theoretical models testable in short-range gravity experiments, it is impractical for each experiment to report a separate test of each model. It is common instead to report experimental constraints using a single parameterization that captures the broad classes of new physics described previously. The most common is the Yukawa parameterization, which modifies the gravitational potential $V(r)$ with a Yukawa term of strength α and length scale λ :

$$V(r) = -\frac{Gm_1m_2}{r} [1 + \alpha e^{-r/\lambda}]. \quad (3.10)$$

This parameterization is well motivated as it applies to any process mediated by the exchange of a boson with a mass $m_b = 1/\lambda$. This includes the LED scenario in which the lightest KK mode mediates the gravitational force when r is near the compactification length scale. A somewhat less common parameterization is the power law potential,

$$V(r) = -\frac{Gm_1m_2}{r} \left[1 + \alpha_N \left(\frac{r_0}{r} \right)^{N-1} \right] \quad (3.11)$$

which includes a length scale r_0 and a strength α_N for each of $N = 2, 3, 4, \dots$. In addition to describing the LED scenario at separations well below the compactification scale, this model captures higher-order processes mediated by the exchange of multiple bosons. In this thesis I will use the first parameterization (Eq. (3.10)), though I note that certain experimental tests are more suited to tests of the second (Eq. (3.11)).

Regardless, the generic parameterization is simply a placeholder for the more specific models that would have to be tested in the event that a signal is observed.

3.2.1 The Landscape of Measurements

Modifications to the gravitational ISL have been searched for using many techniques covering a large range of length scales. At the angstrom scale, constraints are derived from measurements of neutron-atom scattering, often using noble gas targets [145–150]. Constraints in the nanometer range have been obtained from experiments which measure the Casimir force and report limits on deviations from the expected behavior [151], though the current best constraints come from neutron interferometry [152]. Experiments most sensitive in the micron to millimeter range often use micromechanical cantilevers or torsion pendulums; these are discussed in more detail below. Above the laboratory scale, the strongest constraints come from geophysical or astrophysical data [153]. Laser Lunar Ranging (LLR) is a technique which uses the retroreflectors left by the Apollo and Luna missions on the lunar surface to precisely measure the Earth-Moon distance. Constraints on the anomalous precession of the Moon’s orbit from LLR are the most sensitive probes of the Yukawa modification, with sensitivity better than a part in 10^{10} at $\lambda \sim 10^8$ m [154].

Micromechanical Oscillators

The strongest constraints in the range from $\lambda = 30$ nm to $8\ \mu\text{m}$ were set in Ref. [157] using a microelectromechanical torsion oscillator (MTO) (Fig. 3.2a). The oscillator consisted of a $500\ \mu\text{m} \times 500\ \mu\text{m}$ plate with an axis along its midpoint. A $300\ \mu\text{m}$ diameter sapphire sphere coated with layers of chromium and gold was glued to one end and acted as the test mass. A disk with concentric rings of azimuthally-alternating gold and silicon sections was glued to a circular glass flat and acted as the source mass. To shield the test mass from Casimir backgrounds, a 150 nm gold coating was deposited on the disk. During the measurement, the disk was positioned near the sphere (as close as 200 nm) and rotated at an angular frequency of $\omega = 2\pi f_r/n$, where $f_r \approx 300$ Hz is the resonant frequency of the MTO and n is the number of gold/silicon

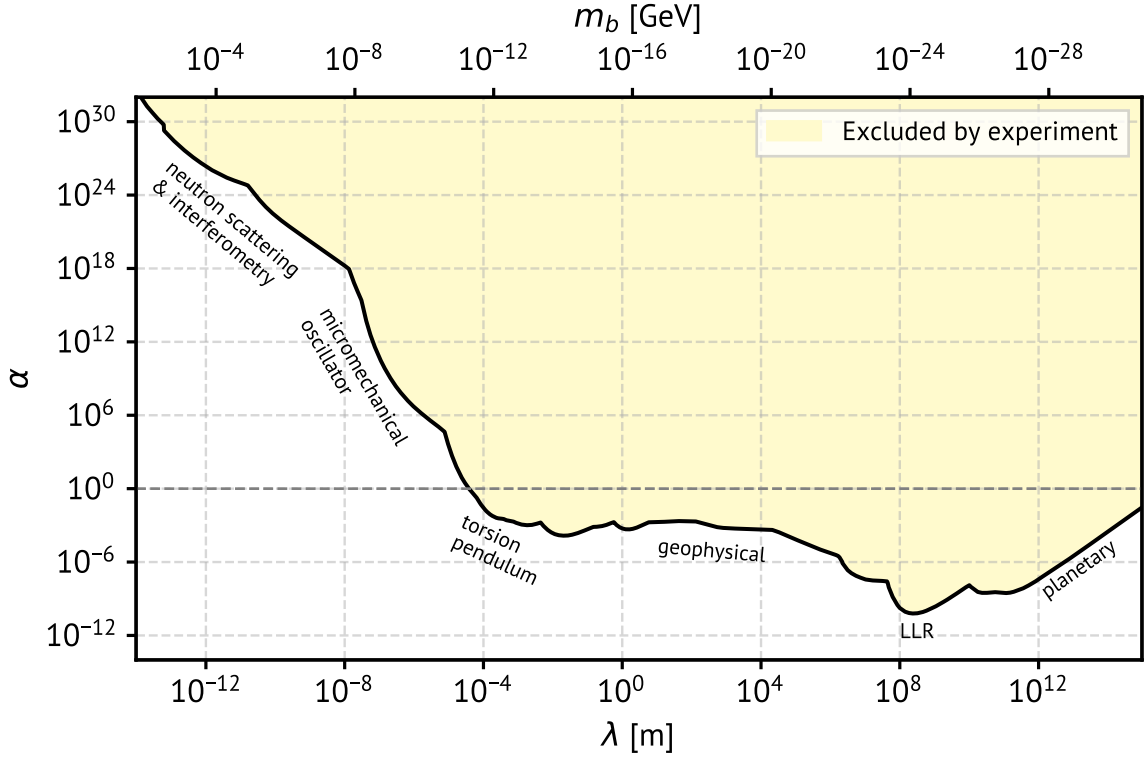


Figure 3.1: Strongest constraints at the 95% CL level on Yukawa modifications to the gravitational ISL across a range of length scales. Constraints come from neutron scattering [145–150], neutron interferometry [152], micromechanical oscillators [155–157], torsion pendulums [158–164], and geophysical/astrophysical measurements [165, 166, 154].

sectors. The resonant response of the MTO was then measured via the capacitance between the oscillator plate and underlying electrodes. In an initial run, a background was observed, though this was later identified as originating from a mechanical step-like wobble in the rotation of the spindle. By repositioning the test mass over the disk, the phase of the background was shifted with respect to that of the hypothetical Yukawa interaction, significantly reducing the in-phase component. This resulted in a measured force that was consistent with zero and enabled strong constraints on Yukawa interactions across a broad range of λ values. A similar technique was employed in Refs. [167, 156] to search in the $\sim 10\ \mu\text{m}$ range. In that experiment, the oscillator consisted of a micromechanical cantilever with a gold test mass mounted

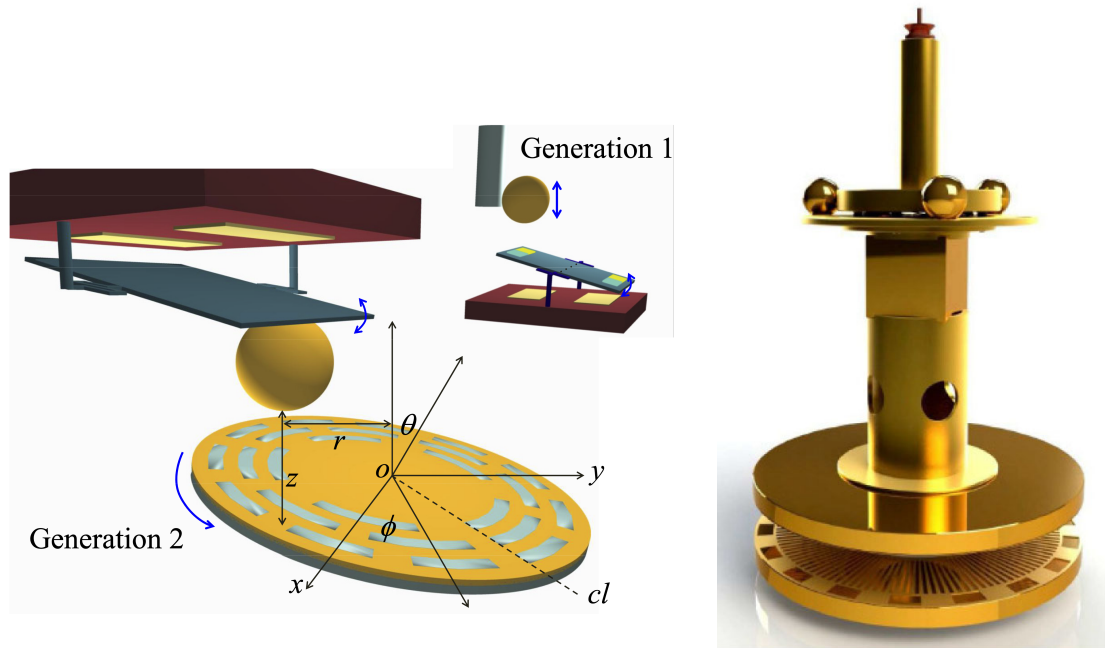
to the end, while an attractor with alternating bars of silicon and gold oscillated horizontally beneath it.

Torsion Pendulums

The first measurement of the gravitational force between two masses in the laboratory was performed by Henry Cavendish in 1797 using a torsion pendulum [168]. While its design has evolved significantly in the years since, the torsion pendulum remains the predominant apparatus for laboratory-scale measurements of the gravitational force. The Eöt-Wash experiment at the University of Washington used a torsion pendulum to measure the gravitational force between two masses with a face-to-face separation reaching as little as $52\ \mu\text{m}$ [164]. This represents the smallest separation at which the gravitational force has been measured. The basic principle on which a torsion pendulum operates is that a test mass suspended from a thin wire can undergo rotational oscillations as a result of minuscule torques. In the Eöt-Wash experiment, the test mass and attractor were gold-coated platinum disks (Fig. 3.2b). Two concentric epoxy-filled hole patterns with 18-fold and 120-fold azimuthal symmetry provided the density modulation. A thin foil was used to shield electromagnetic interactions between the attractor and test mass. Data was collected with the attractor rotating at an angular frequency, ω . The twist angle of the test mass was measured using an autocollimator while the capacitance between the source and test mass provided a measurement of the separation. The torques measured at 18ω , 54ω , and 120ω were then fitted as a function of the separation. Constraints on Yukawa interactions were obtained by repeating the fit with finite λ values. Torsion pendulums have been used in a number of other such measurements [158–163].

3.2.2 A New Experimental Technique

At the scale of a few microns, even state-of-the-art experiments lack the sensitivity to probe gravity directly. This is due to a combination of noise, measurement backgrounds, and the experimental challenge of placing large enough masses sufficiently close together. As these experiments have so far reported only null results, the field



(a) Micromechanical cantilever

(b) Torsion pendulum

Figure 3.2: Left: the micromechanical cantilever used in Ref. [157]. The source mass is a rotating, density-patterned disk. Deflections of the gold-coated sphere are used to measure the force. Right: the torsion pendulum described in Ref. [164]. The density patterned disk below the pendulum rotates and sources the gravitational torque which is then measured from the response of the pendulum. These figures have been adapted from the cited references.

can be said to be *systematics limited*. This is in stark contrast to the prior discussion on neutrino physics; in that field, inadequate statistics limits the current generation of experiments. Given that issues of systematics — not statistics — are limiting progress in the field of short-range gravity experiments, there is immense value in developing novel techniques which complement those that are already more mature. These new techniques, while undoubtedly bringing about their own unique challenges, may bypass existing limitations and open new pathways to improvements. And, in the event that something new is observed, confirmation by multiple experiments employing different techniques will be required for a credible discovery claim. This was a key motivation behind the development of the experiment described in Chapter 7,

one of the few experiments that does not use a mechanical-spring-based oscillator. In that chapter, I will describe the technique in detail, and report updated constraints on micron-scale forces from the latest iteration of the experiment.

In keeping with the theme of this thesis, I have presented the motivations for short-range gravity measurements in terms of the implications that a discovery of something new would carry for particle physics and the SM. There is another important aspect of this work worth mentioning, however. Recently, interest has been growing in experimental techniques that demonstrate quantum mechanical phenomena in the laboratory. These have included cooling nanoscale levitated objects to the motional quantum ground state, developing quantum-limited force sensors, and working towards the preparation of gravitational quantum superpositions of objects [169–173]. Techniques like the one described in Chapter 7 may be of considerable interest to this burgeoning subfield of precision measurement.

Chapter 4

The nEXO Search for Neutrinoless Double Beta Decay

nEXO is a planned experiment to search for the $0\nu\beta\beta$ of ^{136}Xe using a single phase, liquid xenon time projection chamber, with data collection beginning in the 2030s and continuing for 10 years. The nEXO collaboration includes scientists and engineers from 40 participating institutions spread across 9 countries. This chapter introduces the working principle behind the technology employed by nEXO before describing the design of the detector itself. Particular attention is given to those aspects of the experiment that are relevant to later chapters. The chapter concludes with a discussion of nEXO's projected sensitivity to $0\nu\beta\beta$.

4.1 Xenon Time Projection Chambers

A time projection chamber (TPC) is a type of detector that enables three-dimensional reconstruction of a particle trajectory. It typically works by using electric fields to drift electrons liberated via ionization to a plane of charge sensors, thereby *projecting* the trajectory onto the two-dimensional surface. The time elapsed between the prompt detection of scintillation light and the delayed arrival of the drifting charge provides the third dimension. Xenon is an attractive detection medium for TPCs due

to its production of anticorrelated scintillation and ionization signals, allowing the deposited energy to be resolved better than the signals in the either channel alone [174]. The partitioning of energy into charge and light signals is interaction-dependent, and therefore the simultaneous measurement of both channels enables particle identification and discrimination of α backgrounds [175]. This feature has made the xenon TPC a staple of a number of rare event search programs, including searches for dark matter and — since the detection medium itself includes a candidate isotope — for $0\nu\beta\beta$ [176].

Xenon TPCs can use liquid xenon (LXe) or gaseous xenon (GXe) in a “single phase” configuration, or both, in a “dual phase” configuration. The high density of LXe compared to GXe enables LXe TPCs to be more easily scaled to large target masses, though the shorter scattering length prohibits fine-grained trajectory reconstruction, an attractive feature of GXe TPCs. Dual phase TPCs use a strong electric field at the phase boundary to extract the drifting electrons from the LXe target. The light produced via electroluminescence is then detected in the gaseous region by photosensors, resulting in significant amplification of the charge signals. For this reason, dual phase TPCs are the standard in dark matter direct detection programs such as XENON, PandaX, and LZ, where low energy thresholds are critical [177–179]. $0\nu\beta\beta$ experiments, which need excellent energy resolution at much higher energies, often use single-phase TPCs.

4.1.1 Xenon Microphysics

A particle passing through xenon may interact either with a nucleus or a bound electron, resulting in either a nuclear recoil or an electron recoil, respectively. In an electron recoil, the energy deposited by an incident particle is divided entirely between ionization and scintillation. In a nuclear recoil, a portion of the deposited energy is lost to atomic motion which cannot be measured. The process that causes the scintillation and ionization signals is depicted in Fig. 4.1 and unfolds as follows. A population of excited xenon atoms (excitons) is created; each of these then combines with another xenon atom to form an excited dimer (excimer). The relaxation of these

excimers is accompanied by the emission of scintillation light peaked at 175 nm in GXe and 178 nm in LXe [180]. In addition to creating excitons, the primary energy deposition also ionizes a number of atoms, some of which combine with another atom to form an ionized dimer. In a process called recombination, an ionized dimer can capture a free electron to produce a neutral atom and an exciton. The excitons result in an additional scintillation signal, while the electrons which escape recombination constitute the ionization signal. The fraction of electrons which recombine, r , depends on the local electric field, electron mobility, and ion density [181]. The number of scintillation photons, N_{ph} , and the number of ionization electrons, N_{q} , are then given by,

$$N_{\text{ph}} = N_{\text{ex}} + rN_{\text{i}} \quad \text{and} \quad N_{\text{q}} = (1 - r)N_{\text{i}}, \quad (4.1)$$

where N_{ex} (N_{i}) is the number of excitons (ions) created by the energy deposition. The mean energy required to produce a single electron-ion pair, called the W -value, can then be used to reconstruct the energy deposited in the xenon. For electron recoils, where no energy is lost to heat, the total energy is proportional to the sum of the number of excitons and ions, or equivalently, the number of scintillation photons and ionization electrons,

$$E = W(N_{\text{ex}} + N_{\text{i}}) = W(N_{\text{ph}} + N_{\text{q}}). \quad (4.2)$$

A W -value of 13.7 eV has been adopted widely throughout the community, though other measurements [182, 183] suggest a lower value of 11.5 eV. For nuclear recoils, the fraction of the initial energy deposition lost as heat is accounted for by the Lindhard quenching factor \mathcal{L} , with

$$E = \frac{W}{\mathcal{L}}(N_{\text{ex}} + N_{\text{i}}) = \frac{W}{\mathcal{L}}(N_{\text{ph}} + N_{\text{q}}). \quad (4.3)$$

In nEXO, the signal of interest is detected through electron recoils, so Eq. (4.2) will be used. The exciton-ion ratio, $\alpha \equiv N_{\text{ex}}/N_{\text{i}}$, is larger in nuclear recoils than in electron recoils by more than a factor of 5. The resulting difference in the charge-to-light ratio can be used for clean separation of electron recoils from nuclear recoils. It is convenient to think of Eq. (4.2) as defining an energy axis related to the charge

and light axes through a rotation of 45° . In reality, N_q and N_{ph} are estimated from measured quantities subject to statistical fluctuations and measurement uncertainties. A more general definition of the rotated energy axis is,

$$E \propto W (\cos \theta \cdot N_{ph} + \sin \theta \cdot N_q) \quad (4.4)$$

where θ depends on the variances and covariances of the estimated charge and light quanta. Fig. 4.5 show simulated data on this empirically-determined rotated energy axis.




Figure 4.1: Diagram showing the processes which produce the scintillation and ionization signals resulting from either an electron recoil or a nuclear recoil on a xenon target.

4.2 The nEXO Detector

nEXO will use a single phase LXe TPC with five tonnes of xenon enriched to 90% in ^{136}Xe (natural abundance 8.86%) acting as both the $0\nu\beta\beta$ source and detection

medium. The nEXO TPC will be housed in a cylindrical copper vessel with an inner height and diameter of 127.7 cm. The drift region will be contained between a charge-sensing anode plane at the top and a circular copper cathode at the bottom, with copper rings with an inner diameter of 113.3 cm spaced regularly along the cylinder axis to maintain the electric field uniformity. The cathode will be held at negative high voltage to maintain a 400 V/cm electric field across the 118.3 cm drift length. A CAD rendering of the nEXO TPC is shown in Fig. 4.2. More details on the design can be found in Ref. [184].

nEXO will be installed in the Cryopit at SNOLAB, the laboratory located 2 km underground near Sudbury, Ontario, Canada. At that depth, the cosmic ray muon flux is reduced by a factor of $\sim 5 \times 10^7$ compared to ground level [185]. Surrounding the TPC will be 33,000 kg of Novec 7000, an organic fluid which acts as a thermal bath. A spherical, double-walled cryostat will hold the Novec 7000; the cryostat itself will be suspended in a large water tank instrumented with PMTs to act as both a shield layer and an outer detector (OD) for muon tagging. The TPC, cathode, and field rings will all be electroformed underground at SNOLAB to mitigate intrinsic radioactive contaminants.

4.2.1 Measuring Charge Signals

When ionizations occur in the LXe volume, the electric field causes the positive ions to drift to the cathode and the electrons to the anode. Repeated collisions with xenon atoms along the drift path keep the electron drift speed to a fixed value, $v_d \approx 1.7 \text{ mm}/\mu\text{s}$, which depends on the field strength. As electrons drift much faster than ions, the ions are usually treated as stationary on the timescale of a single event. The drifting cloud of electrons will diffuse in both the longitudinal and transverse directions, enlarging the charge cloud and obscuring any underlying structure in the charge signals. A point-like charge signal produced just above the cathode will have increased in size to $\sim 2 \text{ mm}$ by the time it reaches the anode; energy depositions over multiple sites separated by less than this will appear as a single diffuse charge cloud [186, 187]. Drifting electrons can also be captured by electronegative impurities in




Figure 4.2: CAD rendering showing a section view of the nEXO TPC, with essential components labeled.

the xenon, reducing the amount of charge collected. This reduction can be accounted for via calibration, though it still worsens the energy resolution (see Chapter 5 for further discussion). During operation, the xenon will be continuously boiled off at the TPC outlet, recirculated through a gas-phase xenon purifier, and recondensed at the inlet to remove electronegative impurities. This process is described in more detail in Chapter 6.

Charge signals will be measured using square tiles with side lengths of 10 cm which cover the anode plane. The tiles consist of square-diamond-shaped gold pads deposited on a quartz substrate, as shown in Fig. 4.3 [188]. The pads are connected at the corners to join them into strips along both the x and y directions. Insulation layers between perpendicular gold crossings keep the x and y channels separate. The strip pitch will be 6 mm to ensure adequate position resolution in the $x - y$ plane. To

minimize cabling from the large number of charge channels, the charge signals will be digitized by ASIC chips on the back of each of the tiles. Research is ongoing to determine how best to route the charge signals to the back of the chip, with efforts underway on the development of traces wrapping over the edges and through-quartz vias.




Figure 4.3: Schematic of a nEXO charge-sensing tile with a callout showing the crossings of x and y strips separated by an insulating layer.

4.2.2 Measuring Light Signals

To collect the scintillation signals, the barrel of the TPC will be tiled with staves containing VUV-sensitive silicon photomultipliers (SiPMs). The SiPMs will be grouped into tiles, with 16 readout channels per tile, digitized by an ASIC chip on the reverse side. A staff will be formed from a ten-by-two array of tiles, with 24 staves covering the interior of the TPC. This will result in a total of 7680 light channels. Al+MgF₂ coatings will be applied to the field rings and copper cathode to increase their reflectivity of VUV light and improve the light collection efficiency.

Two candidate manufacturers are under consideration for the SiPMs: Fondazione

Bruno Kessler (FBK) and Hamamatsu Photonics (HPK). Three models (FBK VU-VHD3, HPK VUV4-50, and HPK VUV4-Q-50) have been characterized in terms of their photon detection efficiency, dark count rate, and production of correlated avalanches [189]. Work is ongoing to compare these candidates in terms of external cross-talk and stability over long periods in nEXO-like conditions.

4.3 Three-Parameter Analysis

Energy resolution is an important factor in signal/background discrimination for $0\nu\beta\beta$ experiments. Indeed, many experiments use energy alone as the axis along which signals and backgrounds can be distinguished. However, a distinct advantage of large, monolithic detectors such as nEXO is their ability to leverage spatial and topological information for improved signal/background discrimination capabilities. In addition to energy, nEXO uses the distance from the nearest detector component surface, called the “standoff” distance, and the event topology to separate the signal of interest from backgrounds. The sensitivity is calculated by exploiting the differences in signal and background distributions over these three parameters shown in Fig. 4.4.

4.3.1 Energy

As energy is one of three parameters in which signals differ from backgrounds, energy resolution at the $\beta\beta$ -decay Q -value of $Q_{\beta\beta} = 2458$ keV is critical to achieving optimal signal/background discrimination in nEXO. Many factors contribute to nEXO’s energy resolution; here, I give only a brief overview. In each of the charge and light channels, the following factors contribute to the overall energy resolution:

1. Statistical fluctuations in the number of charge/light quanta produced
2. Statistical fluctuations in the number of charge/light quanta detected
3. Error in the calibrated charge/light detection efficiencies
4. Noise in the charge/light readout electronics

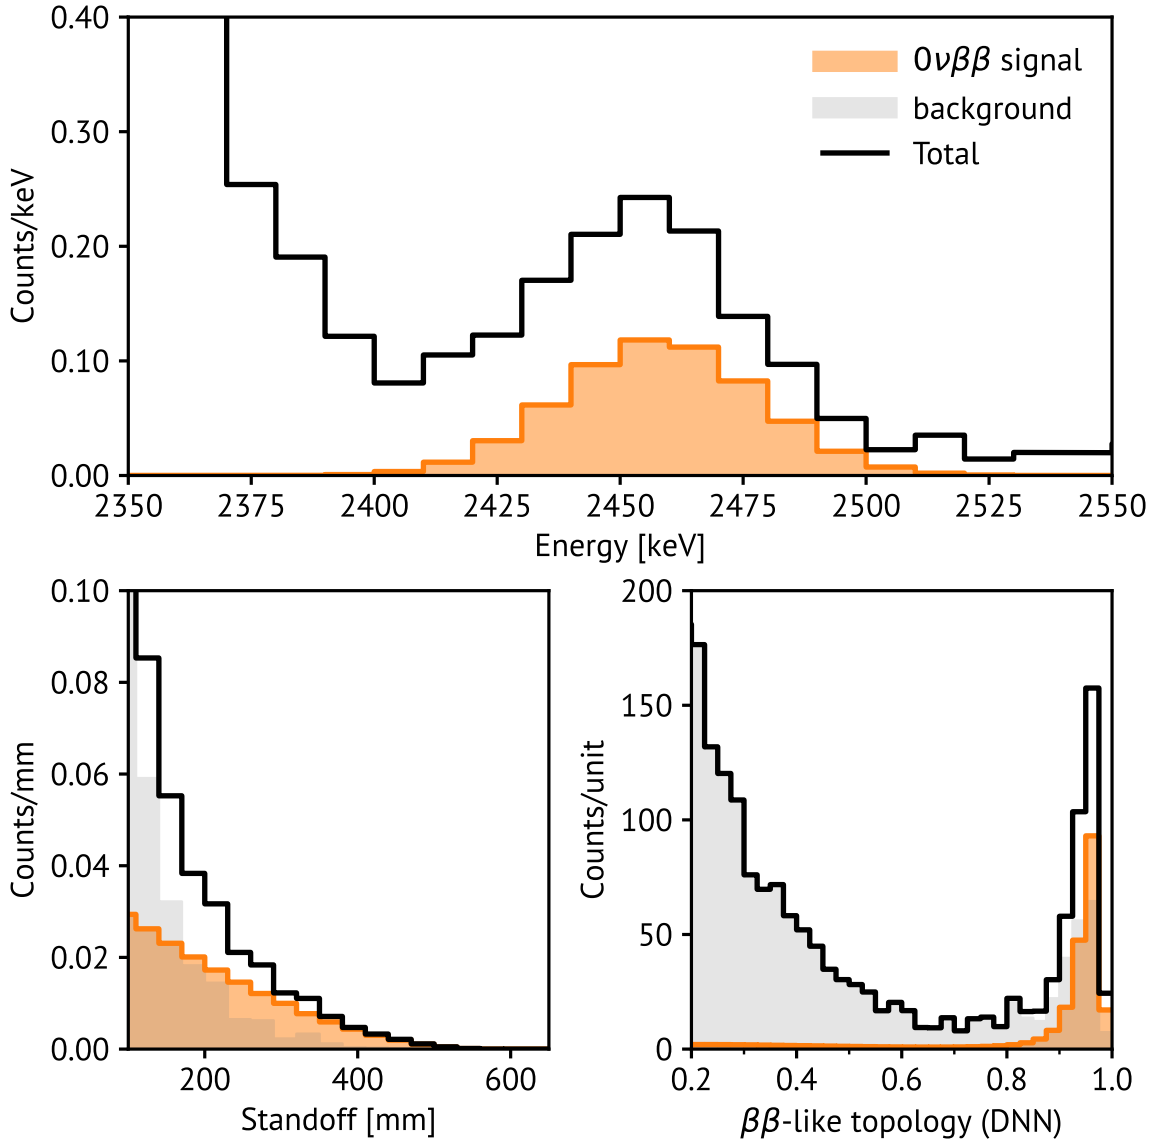


Figure 4.4: Histograms showing simulated $0\nu\beta\beta$ signal with $T_{1/2} = 0.74 \times 10^{28}$ years and background data in all three of the parameters used in the analysis. Each histogram shows data with cuts applied to the other two parameters to select the signal region.

Factor 1 consists of recombination fluctuations, which depends on the covariance between the number of charge and light quanta produced, and an overall Fano factor,

which quantifies fluctuations in the total number of quanta. These have a negligible effect at $Q_{\beta\beta}$. The contribution from factor 2 can be reduced by increasing the charge/light collection efficiency, while that from factor 3 can be reduced with improved calibration techniques. These two factors are discussed in more depth in Chapter 5. Minimizing the contribution from factor 4 involves implementing standard low-noise measurement practices. nEXO projects an energy resolution of 0.8%, including contributions from all of these factors.




Figure 4.5: Distribution of reconstructed charge and light signals for simulated ^{208}Tl decays in the TPC, with the rotated energy axis overlaid in green. The red dashed lines indicate the charge-to-light cuts that select electron recoils and eliminate α events.

4.3.2 Standoff Distance

The attenuation length of a 2.5 MeV γ in LXe is ~ 9 cm [190]. For backgrounds originating from interior detector surfaces, this results in the background rate falling off significantly with increasing standoff distance. In contrast, the $0\nu\beta\beta$ signal rate is independent of the standoff distance. This difference between signal and background rates is more pronounced in detectors with dimensions many times larger than the attenuation length. In nEXO, the inner tonne of the detector (a cylindrical region > 200 mm from the detector walls) is effectively background free, with 0.6 events expected near the Q -value in the full ten-year exposure. The benefits of this strong attenuation go beyond shielding; backgrounds measured in the outer region of the TPC can be used to validate the background model. Moreover, as γ attenuation is extremely well described by simulation, these background measurements can be used to predict the background rate in the inner region.

4.3.3 Event Topology

Events in nEXO can be classified as single-site (SS) or multi-site (MS) depending on whether the individual energy depositions can be spatially resolved. The dominant backgrounds in nEXO are MS events, as γ s tend to Compton scatter multiple times, creating spatially-separated energy depositions which are then detected as separate charge clusters. MeV-scale electrons, on the other hand, have mean ranges of a few millimeters [191] resulting in localized energy depositions and appearing as SS events. Bremsstrahlung from the electrons results in a fraction of $0\nu\beta\beta$ appearing as MS events with a particular event topology. A deep neural net (DNN) based on the ResNet-18 architecture [192] is used for discrimination between signal and background based on event topology. The DNN takes charge waveforms for the x and y strips as input and outputs a DNN score ranging from 0 to 1. It is trained on simulated data, allowing for PDFs for both signal and background events to be produced and used in the sensitivity analysis.

4.4 Backgrounds

The backgrounds which limit sensitivity most severely are those which are similar in the three aforementioned parameters to the signal; i.e., those that have an energy near $Q_{\beta\beta}$, appear at least a few millimeters from detector surfaces, and exhibit $\beta/\beta\beta$ -like topology. These include β -emitters in the LXe volume with Q -values greater than $Q_{\beta\beta}$. One such background source is ^{137}Xe , which undergoes β -decay with a half life of 3.82 minutes and a Q -value of 4162 keV. It is produced when neutrons from muon interactions near the TPC capture on ^{136}Xe . As ^{137}Xe is produced in an excited state and promptly emits a 4026 keV γ , a coincidence trigger between the OD and the TPC can be used to tag ^{137}Xe production and veto the subsequent few minutes of data. The $2\nu\beta\beta$ of ^{136}Xe can also contribute background events in principle, though despite the total activity of 0.2 Bq, the fraction of those with energies near the Q -value is negligible.

Backgrounds from γ transitions with energies within $\sim 1\%$ of $Q_{\beta\beta}$ can also be problematic if not properly mitigated. The β -decay of ^{214}Bi (a daughter of ^{238}U) is of particular concern as it is accompanied by the release of a 2447.9 keV γ . Emission of this γ from detector components in or around the TPC can result in energy depositions in the LXe. For components which are in direct contact with xenon (either the LXe in the TPC or the GXe elsewhere in the xenon handling system), ^{222}Rn produced in the α -decay of ^{226}Ra can emanate into the xenon volume and be transported into the active volume of the TPC. ^{222}Rn that makes it into the LXe volume will subsequently decay, with a fraction of the daughters produced in an ionized state. These charged ^{222}Rn daughters will drift toward the cathode where they will plate out and continue emitting radiation into the LXe. The β -decay of ^{214}Bi is followed by the α -decay of ^{214}Po with a half life of 163.5 μs . With a sufficiently long coincidence window, this “BiPo” decay sequence can be tagged. The tagging efficiency depends strongly on the location of the decay, with near-perfect efficiency in the bulk of the LXe, and $\sim 50\%$ efficiency for decays on the cathode, where the α has equal probability of being emitted into the LXe or the copper. The β decay of ^{208}Tl (a daughter of ^{232}Th) is also of concern due to the accompanying 2614.5 keV γ which can Compton scatter

and deposit energy near $Q_{\beta\beta}$.

To mitigate these backgrounds, all materials must be carefully screened to ensure concentrations of problematic radioisotopes do not exceed stringent limits. In practice, this screening is often done by measuring the activity or concentrations of ^{238}U and ^{232}Th under the assumption that the decay chain is in secular equilibrium. The more relevant activities/concentrations are those of ^{226}Ra and $^{228}\text{Ra}/^{228}\text{Th}$, the nearest long-lived parent radioisotopes to the decays of concern in each decay chain. Radioassay measurements of candidate materials are performed at many of nEXO's participating institutions. Results of these measurements are compiled and stored in an online database to facilitate design studies and background modeling [193].

Radioassay measurements from a number of LXe-based collaborations, along with data from EXO-200, suggest that purifier media constitute a significant source of ^{222}Rn emanation. As the xenon will be continuously recirculated through a gas-phase purifier during data acquisition, much of the ^{222}Rn will be transported into the LXe volume. This has sparked an interest within nEXO in a new design for a low-radioactivity purifier to reduce this source of backgrounds. Chapter 6 discusses the preliminary results of this ongoing effort.

4.5 Sensitivity

Every aspect of nEXO mentioned so far is relevant to the degree that it affects the ultimate sensitivity to $0\nu\beta\beta$. A prior estimate of the sensitivity that nEXO is expected to reach was reported in Ref. [194] based on the multi-parameter analysis and background model described previously. Here, that study is briefly summarized and the results are discussed.

Signal and background data were simulated using the Geant4-based `nexo-offline` simulation framework. The simulation included a complete geometric model of the nEXO TPC, surrounding cryostat, and water tank, with component properties assigned based on shielding properties and radioisotope content. The production of scintillation and ionization signals was calculated with a modified version of the Noble Element Simulation Technique (NEST) code [195] tuned to match the scintillation and

ionization responses measured in EXO-200 [182]. Charge waveforms were simulated for channels near the event, accounting for diffusion, attenuation by electronegative impurities, and electronics noise. The number of scintillation photons detected was sampled from a high-statistics optical simulation of the TPC, with added fluctuations to account for the photon detection efficiency and correlated avalanches. A series of cuts was applied to mimic those present in the analysis of real data, including on event position, reconstructed energy, and ionization to scintillation ratio. The data was then binned in energy, DNN score, and standoff distance to construct PDFs for each component or group of components.

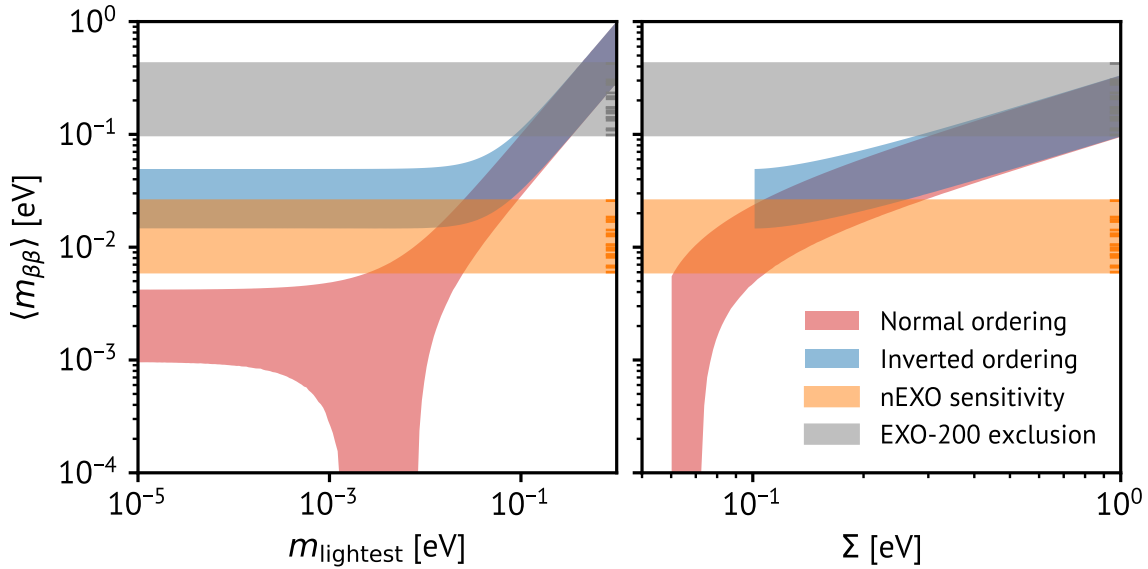


Figure 4.6: The projected 90% CL exclusion sensitivity of nEXO to $0\nu\beta\beta$ (orange band). The 90% CL exclusion limit from the full EXO-200 dataset is shown for comparison (gray band). The bands reflect the uncertainty in where the upper limit is placed originating from different calculations of the NMEs. The ticks on the right side of each band indicate specific calculated values while the bands span their full range.

Toy datasets were produced by sampling the PDFs scaled by weighting factors to account for the relative contributions of each. Fig. 4.4 shows the summed signal and background PDFs in all three parameters. A frequentist statistical analysis based on the profile likelihood ratio method was used to compute the median 90% confidence

level (CL) exclusion sensitivity and 3σ discovery sensitivity from the toy datasets. This method gives a projected 3σ discovery sensitivity of 0.74×10^{28} years and a projected 90% CL exclusion sensitivity of 1.35×10^{28} years after ten years of data collection.

As discussed in Section 2.3.3, there are a range of calculated NMEs for ^{136}Xe ; this is the largest source of uncertainty in translating a reported half life into $\langle m_{\beta\beta} \rangle$. Fig. 4.6 shows a band of $\langle m_{\beta\beta} \rangle$ values corresponding to this range. For all but one of the NMEs, nEXO is projected to cover the entirety of the IO band, and for many, a large portion of the NO band below the degenerate region will be accessible. nEXO is an ambitious, discovery-focused experiment, with discovery potential in both the IO and NO ordering scenarios. Chapter 5 and Chapter 6 focus on addressing two components of the experiment design that will enable nEXO to realize this projected sensitivity when the detector turns on. Those chapters describe a small selection of the R&D activities for nEXO at Stanford; a more complete overview can be found in Appendix B, which provides a comprehensive discussion of the Stanford test facility for LXe TPC development and its many applications.

Chapter 5

Development of a ^{127}Xe Calibration Scheme for nEXO

This chapter has been adapted from Ref. [196], of which I am a primary author. The description that follows places more emphasis on my own contribution — the lightmap calibration scheme — and less on the work led by my collaborators.

5.1 The Energy Resolution Model

Accurate calibration is essential for ensuring that nEXO’s energy resolution target is met. In this section, the quantitative impact of calibration uncertainties on the energy resolution is determined. Both charge and light sensors detect *quanta*, so we can begin by discussing the generic problem before addressing the particularities of light or charge detection. The number of detected quanta, q_{det} , can be modeled as originating from a binomial process with the number of trials given by the quanta produced, q_0 , and the probability given by the collection efficiency at position \vec{x} and time t , $\varepsilon_0(\vec{x}, t)$:

$$q_{\text{det}} \sim \mathcal{B}(q_0, \varepsilon_0). \quad (5.1)$$

Calibration campaigns will be conducted to measure the collection efficiency,

$\varepsilon_c(\vec{x}, t)$, allowing for the estimate of the number of quanta produced, q , to be calculated using¹

$$\hat{q} = \frac{q_{\text{det}}}{\varepsilon_c(\vec{x}, t)}. \quad (5.2)$$

The random variable q has a probability mass function $P(q; q_0, \varepsilon_0, \varepsilon_c)$. To model the energy resolution in nEXO, we require the mean and variance of $P(q; q_0, \varepsilon_0, \varepsilon_c)$ when averaged over the full detector volume, V . It has been shown in Ref. [197] that these quantities are, respectively,

$$\text{E}[P(q; q_0, \varepsilon_0, \varepsilon_c)] = q_0 \text{E} \left[\frac{\varepsilon_0}{\varepsilon_c} \right]_V \quad \text{and} \quad (5.3)$$

$$\text{Var}[P(q; q_0, \varepsilon_0, \varepsilon_c)] = q_0 \text{E} \left[\frac{\varepsilon_0(1 - \varepsilon_0)}{\varepsilon_c^2} \right]_V + q_0^2 \text{Var} \left[\frac{\varepsilon_0}{\varepsilon_c} \right]_V, \quad (5.4)$$

where E and Var denote the expected value and variance of a quantity, respectively.

The first term in Eq. (5.4) is the contribution from the binomial fluctuations in the number of quanta collected, while the second arises purely from errors in the calibration of the collection efficiency. The second term is therefore the one of concern from the perspective of designing and assessing calibration schemes. The variance due to calibration error can be labeled by σ_c^2 and rewritten as

$$\sigma_c^2 \equiv \text{Var} \left[\frac{\varepsilon_0}{\varepsilon_c} \right]_V \approx \text{Var} \left[\frac{\varepsilon_c - \varepsilon_0}{\varepsilon_0} \right]_V, \quad (5.5)$$

where the approximation is obtained with some algebra and the fact that $\varepsilon_c \approx \varepsilon_0$ for a good calibration. With this approximation, the quantity in the square brackets is simply the relative calibration error across the TPC. To isolate the impact of this calibration error, we can set the combined uncertainty in the energy from all other factors to some fixed value, σ_{fixed} , and add the marginal contribution from σ_c according to

$$\frac{\sigma_E}{E} = \frac{\sqrt{\sigma_{\text{fixed}}^2 + q_0^2 \sigma_c^2}}{E}, \quad (5.6)$$

where the quantity on the left-hand side is the total energy resolution and E is the

¹Throughout this chapter, a hat (^) over a variable is used to denote an estimator for that variable.

energy. This can be applied to both the light and charge response calibrations.

5.2 Position Dependence of the TPC Response

The spatial dependence of the ionization signal response arises from the attachment of drifting electrons to electronegative impurities in the xenon. Because charge drifts at a constant speed along the z direction, $v_d = z/t$, charge signals are exponentially attenuated as a function of the drift time, t . This allows for the true charge, Q_0 , to be reconstructed from the detected charge, Q_{det} , as,

$$\hat{Q}_0 = \frac{Q_{\text{det}}}{e^{-t/\tau_e}} \quad (5.7)$$

where τ_e is known as the electron lifetime and represents the timescale over which a free drifting electron will be captured by an impurity in the xenon. The electron lifetime is inversely proportional to the concentration of electronegative impurities in the LXe volume. Calibrating the spatial dependence of the detector response to ionization signals therefore amounts to determining the electron lifetime in the detector. This can be achieved by producing events with a known Q_0 at a range of values of z (equivalently, t) and determining the τ_e that best reproduces the detected charge signals, Q_{det} . Since the concentration of electronegative impurities in the TPC varies on a ~ 1 day timescale, the electron lifetime is time dependent and regular calibrations are required.

The position-dependent response to scintillation signals is more complicated, and depends on the quantum efficiency of the sensors, the geometry of the detector and sensor arrays, and the reflectivity of surfaces within the TPC. LXe is highly transparent to its own scintillation light, with a Rayleigh scattering length of $\gtrsim 30$ cm and an absorption length of > 100 cm [198, 199], though in principle any scattering or attenuation should be accounted for in the position dependence as well. The true scintillation signal, S_0 , can be estimated from the detected scintillation signal, S_{det} , with,

$$\hat{S}_0 = \frac{S_{\text{det}}}{\varepsilon_{\text{QE}} \cdot \varepsilon_{\text{LM}}(x, y, z)} \quad (5.8)$$

where ε_{QE} is the sensor quantum efficiency and ε_{LM} is the position-dependent photon collection efficiency, also known as the “lightmap”. The lightmap describes the fraction of the light produced at position (x, y, z) that will reach sensor surfaces, accounting for the geometry and optical properties of the detector.

To calibrate the lightmap, events with a known S_0 should be produced throughout the full TPC volume. Fitting the three-dimensional lightmap from the corresponding set of Q_{det} values is not straightforward. To see why, it helps to contrast it with the electron lifetime calibration. The electron lifetime is calibrated with the assumption that the impurity concentration is uniform throughout the TPC, while the drift field is applied along the z -axis. This means that the charge signal attenuation anywhere in the TPC can be fully described by the z -coordinate of the event position; moreover, the z -dependence is captured by a single parameter, τ_e . No such symmetry exists for the lightmap. The different optical properties of the cathode and anode along with the regularly-spaced field rings result in a strong z -dependence to the photon transport efficiency (PTE). The solid angle of the SiPMs from the location of the event has a significant dependence on the radial position. There is an approximate azimuthal symmetry which is broken by the pattern of charge tiles on the anode. However, it is nonetheless desirable to have a framework for lightmap reconstruction in three dimensions, allowing for other possible sources of azimuthal asymmetries in the real detector (such as some arbitrary set of non-functioning SiPMs). Strategies for this are discussed in Section 5.4.2. Unlike the electron lifetime, the lightmap is not expected to vary over short timescales, so calibrations can be done with reduced frequency.

The lightmap calibration requirement can be quantified by the lightmap error, σ_{LM} . The target value for nEXO is $\sigma_{\text{LM}} < 0.005$. Potential calibration techniques should therefore be assessed on their ability to meet this metric. By applying Eq. (5.5), σ_{LM} can be written as

$$\sigma_{\text{LM}} = \sqrt{\text{Var} \left[\frac{\varepsilon_{\text{LM}}}{\varepsilon_c} \right]_V} \approx \sqrt{\text{Var} \left[\frac{\varepsilon_c - \varepsilon_{\text{LM}}}{\varepsilon_{\text{LM}}} \right]_V}, \quad (5.9)$$

where ε_c is the calibrated lightmap and ε_{LM} is the true lightmap. Later in the chapter,

this quantity will be used to assess the quality of lightmap calibration schemes in terms of their direct, quantitative impact on the total energy resolution in nEXO.

5.2.1 Standard Calibration Techniques

Both the electron lifetime and the lightmap calibrations require that events of known energy be produced inside the TPC. Two strategies for this have previously been considered and are discussed in turn below.

External γ Sources

One standard technique is to use external γ sources positioned around the TPC. This is the planned technique to calibrate the absolute energy scale in nEXO. Because of this, infrastructure to deploy sources to a series of locations around the TPC is already a part of the design of nEXO, and could be used for regular electron lifetime and lightmap calibration campaigns. Indeed, this is the baseline plan for both the electron lifetime and the lightmap [184]. An additional advantage of using external γ sources is that calibrations can be performed over the full energy range of interest with the choice of appropriate sources. This strategy also does not require introducing anything into the interior of the TPC, and therefore poses no risk of contamination of the xenon, either with sources of background or electronegative impurities.

While this strategy has many advantages, the short attenuation length of γ s in the LXe poses a significant limitation. Achieving a sufficient number of events in the center of the TPC requires accommodating a much higher rate near the walls. To avoid overloading the readout electronics and causing pileup, source activities will have to be sufficiently low and calibration campaigns commensurately long. This increases detector downtime for the calibration campaigns and diminishes the livetime. For this reason, alternative strategies which do not sacrifice livetime are desired.

Internal α Sources

A second option is to inject radioactive isotopes into the LXe volume. This has the advantage of ensuring that events are distributed more uniformly throughout the

TPC, illuminating the full detection volume. This strategy has been used previously with a number of different radioactive isotopes [200–207], including $^{83\text{m}}\text{Kr}$, $^{129\text{m}}\text{Xe}$, $^{131\text{m}}\text{Xe}$, ^{37}Ar , tritiated methane, and ^{220}Rn . These have predominantly been used by experiments searching for nuclear recoils from dark matter particles in the 1–100 keV range, well below the nEXO energy threshold of ~ 200 keV. The Q -value of all but ^{220}Rn fall below this threshold. ^{220}Rn is a possible option for nEXO, but it too has some drawbacks. Calibrations with a ^{220}Rn source use the rate-limiting α decay of ^{212}Pb as a source of 570 keV energy depositions with excellent intrinsic resolution. The 10.6 hour half life is comparable to the recirculation timescale in nEXO, meaning there may not be sufficient time for the source to mix uniformly throughout the xenon volume. Additionally, the β decay of ^{208}Tl overlaps with the $0\nu\beta\beta$ region of interest, resulting in backgrounds and limiting the frequency of these calibrations.

5.2.2 ^{127}Xe as a Calibration Source

^{127}Xe is an attractive alternative to the standard options for a number of reasons, possessing advantages of both options without many of the drawbacks. ^{127}Xe decays via electron capture (EC) to excited states ^{127}I with a Q -value of 662.3 keV — far too low to constitute a background in the $0\nu\beta\beta$ search. The subsequent relaxation to the ground state results in the emission of either 375 keV (branching ratio 0.473) or 203 keV (branching ratio 0.527), accompanied by X-rays or Auger electrons totaling 33 keV (K-shell EC) or 5 keV (L-shell EC) [208]. This results in energy depositions at four energies which, due to the detector resolution, are smeared into two broad peaks in both the scintillation and ionization channels at 236 and 408 keV. The full decay scheme is shown in Fig. 5.1, while Fig. 5.2 shows the simulated charge and light production in nEXO. The 36-day half life ensures that the ^{127}Xe can mix fully into the enriched xenon and will be spread uniformly throughout the TPC.

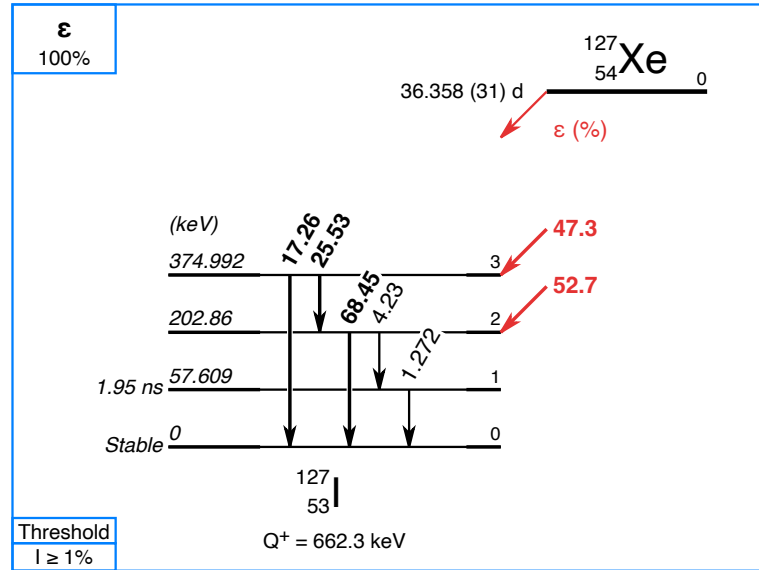


Figure 5.1: Decay scheme for the electron capture decay of ^{127}Xe to excited states of ^{127}I [208]. The red text indicates the branching ratio for each excited daughter state. This figure was produced using the Laraweb tool [209].

5.3 Experimental Demonstration

Before a ^{127}Xe calibration scheme can be adopted for nEXO, its viability must be demonstrated. While a demonstration of the use of ^{127}Xe in nEXO is evidently impossible at present, its use in a small-scale experimental test served to address several practical considerations, from the production of the source to the deployment in a LXe TPC.

5.3.1 Source Production

^{127}Xe is a synthetic isotope. It can be produced by neutron capture on ^{126}Xe , which has an isotopic abundance of $\sim 0.1\%$ in $^{\text{nat}}\text{Xe}$. The neutron capture cross section is shown as a function of energy in Fig. 5.3a. To demonstrate the source production process, a 150 cm^3 316L stainless steel (SS) sample cylinder containing 69 g of $^{\text{nat}}\text{Xe}$ was sent to the McClellan Nuclear Research Center (MNRC), a research reactor operated by UC Davis, for irradiation with neutrons. The cylinder was placed in the Neutron




Figure 5.2: Simulated ionization and scintillation signals produced by the decay of ^{127}Xe in the LXe volume. The panels above and to the right of the scatter plot show one-dimensional histograms of the projections along the scintillation and ionization axes, with the recombination fluctuations causing the decays to appear as two broad peaks in each channel.

Transmutation Doping (NTD) void and subjected to 15 minutes of irradiation with the reactor operating at 250 kW. The neutron flux under nominal operating conditions at 1 MW is shown in Fig. 5.3b. This spectrum was folded with cross sections for isotopes present in both the $^{\text{nat}}\text{Xe}$ and the SS to compute the production of radioisotopes in both. Cross section data was obtained from the ENDF/B-VII.0 library [210] for all reactions except those involving the two metastable isomers $^{129\text{m}}\text{Xe}$ and $^{131\text{m}}\text{Xe}$; for those, the TENDL-2019 nuclear data library was used as it conveniently provides residual cross sections for the desired products [211].

Following the irradiation, the xenon gas was transferred to a separate sample cylinder, allowing for the activities in the irradiated cylinder and xenon gas to be measured separately. Radioassay measurements of the sample cylinder using HPGe counters were performed at MNRC and Stanford 9 and 25 days after the irradiation,

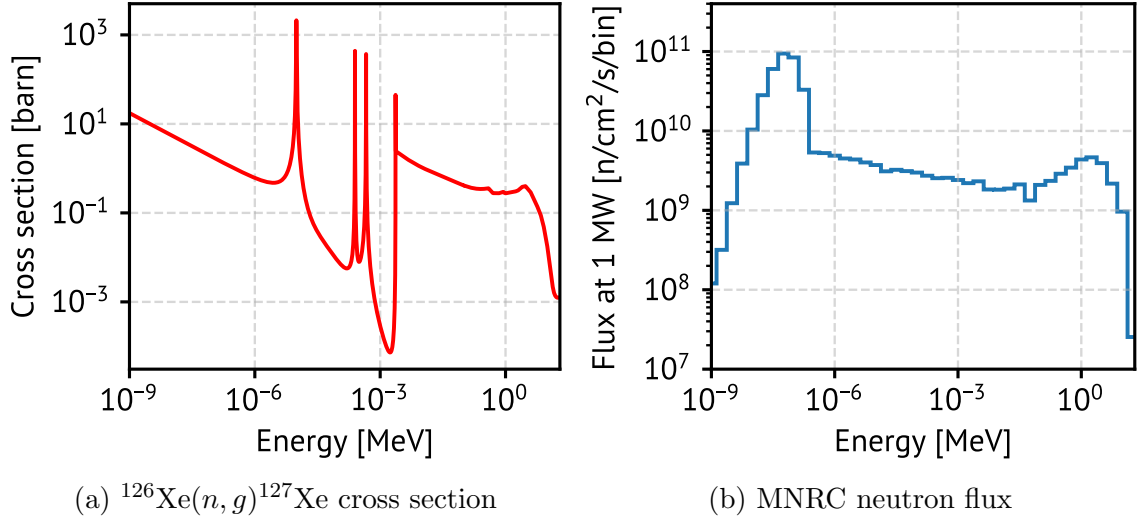


Figure 5.3: Left: neutron capture cross section as a function of energy for ^{126}Xe ; right: Flux spectrum of the MNRC nuclear reactor in the NTD void for the nominal reactor operation of 1 MW.

respectively. Activities were obtained for the long-lived radioisotopes ^{51}Cr , ^{58}Fe , and ^{58}Co ; the first two are produced via thermal neutron capture, while the last results from fast neutron (n, p) reactions. Systematic uncertainties due to the geometry of cylinder/detector setup were estimated to be $\sim 50\%$. These measurements of the activity in the SS allowed for the incident neutron flux to be calibrated by scaling the calculated activity to the measured value, as shown in the left panel of Fig. 5.4. A discrepancy of a factor of ~ 4 was found between the nominal and inferred neutron fluxes; this could be caused by operation of the reactor in transient mode during the irradiation of the sample cylinder. The ratio of activities for the products of thermal neutron and fast neutron reactions were consistent with the expectation from the neutron flux spectrum, however.

The calibrated neutron flux was also used to infer the activities of radioisotopes produced in the xenon (right panel of Fig. 5.4). Of these, ^{127}Xe has the longest half life. After a cool-off period of ~ 100 days, the activities of all other radioisotopes except ^{137}Cs will have dropped to negligible levels, while the ^{137}Cs is expected to be easily removable using standard purification techniques. To verify these predicted activities and ensure that no unexpected radioisotopes were produced, the activities of

radioisotopes within the activated xenon gas were measured using a low-background HPGe counter at the University of Alabama. The low activities required long measurement periods of two weeks (compared to ~ 30 minutes for the SS cylinder). A background measurement was made of the empty cylinder prior to transferring the activated xenon gas. The background spectrum was then subtracted from that obtained during the measurement of the xenon gas in the new cylinder. Five different ^{127}Xe emission lines were observed in the background-subtracted spectrum. The measured activities of the two highest-energy transitions were 6.3 kBq and 6.9 kBq, consistent with the activity of 6.5 kBq predicted from the neutron flux to within $\sim 5\%$. Larger discrepancies were found for the metastable isomers for which thermal neutron cross sections were extrapolated from single measurements with large uncertainties. The most significant discrepancy occurred for ^{137}Cs , which was observed at $\sim 100\times$ lower activity than predicted. This is likely due to attachment of the cesium on the inside of the original sample cylinder, further supporting the hypothesis that the cesium can be removed via purification. A more complete description of the radioassay procedure and results can be found in Ref. [196].

5.3.2 Deployment in a Test TPC

To validate the proposed calibration procedure, the activated xenon was deployed in a test TPC in the Stanford LXe lab (the “Long TPC” described in Appendix B). The TPC consists of a 13.5 cm drift length housed in a 20.3 cm long by 25.4 cm diameter cylindrical SS chamber and holds ~ 30 kg of LXe (Fig. 5.5a). A 400 V/cm drift field is maintained between a SS mesh cathode held at -6 kV at one end of the drift length and a charge-sensing anode plane at the other. Scintillation light is collected on a tile located beneath the cathode mesh containing 24 FBK VUV-HD1 SiPMs paired into 12 channels. Charge collection at the anode is done using the prototype charge tile described in Ref. [188]. The 10 cm by 10 cm tile provides 16 strips with 3 mm pitch in both the x and y directions, though some channels were ganged together as shown in Fig. 5.5b due to the limited number of feedthroughs into the xenon volume. These channels were connected to charge-sensitive preamplifiers located outside the

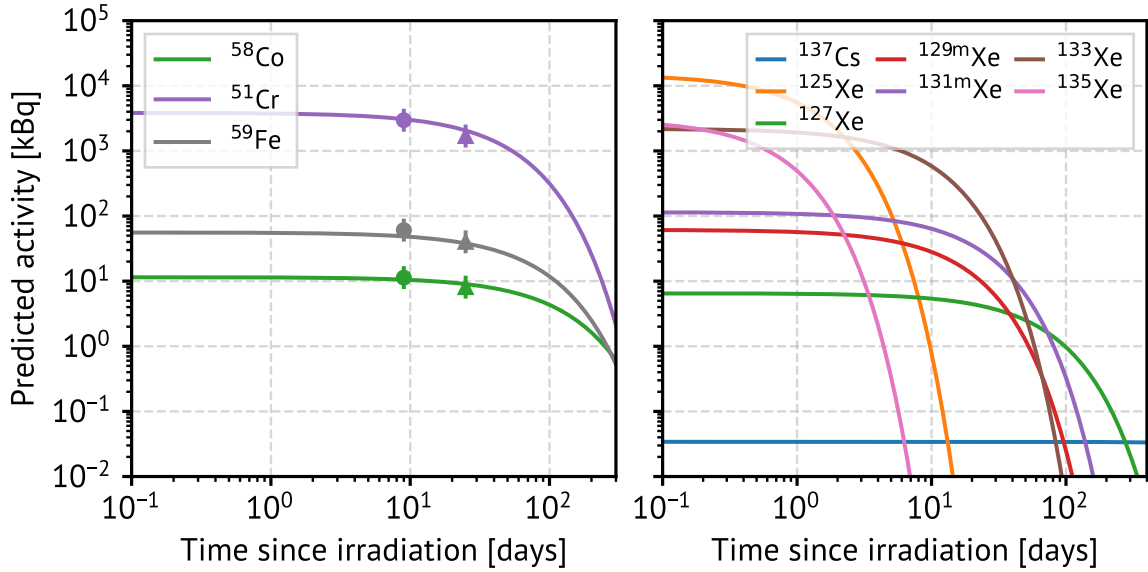


Figure 5.4: Activity of radioisotopes present in the SS cylinder (left) and xenon gas (right) after a 15-minute irradiation in the NTD with the reactor operating at 250 kW. The points with error bars in the left panel indicate radioassay measurements of the SS cylinder used to scale the predicted activity of the radioisotopes plotted.

xenon volume with a design based on Ref. [212]. The light and charge signals were both digitized at 62.5 MS/s using Struck SIS3316 digitizers. Due to the low light collection area provided by the single SiPM tile, the scintillation was used only for triggering, with a three-fold coincidence requirement and the threshold set at the mean single-photon pulse height for each channel.

To introduce the activated xenon into the TPC, the sample cylinder containing the activated xenon was plumbed in to the recirculation manifold, with two valves and an analog pressure gauge in between the cylinder and the main recirculation flow path, as depicted in Fig. 5.6a. The valves were then opened one at a time to first pressurize the intermediate volume, then release the activated xenon into the recirculation path. During this period, the system was undergoing forced recirculation at a flow rate which oscillated between 5 and 30 g/min depending on the stage of the recirculation pump's cycle. Each time xenon was introduced into the intermediate volume, the pressure in the sample cylinder was recorded. Trigger rate measurements were continually acquired throughout the full injection procedure to monitor the

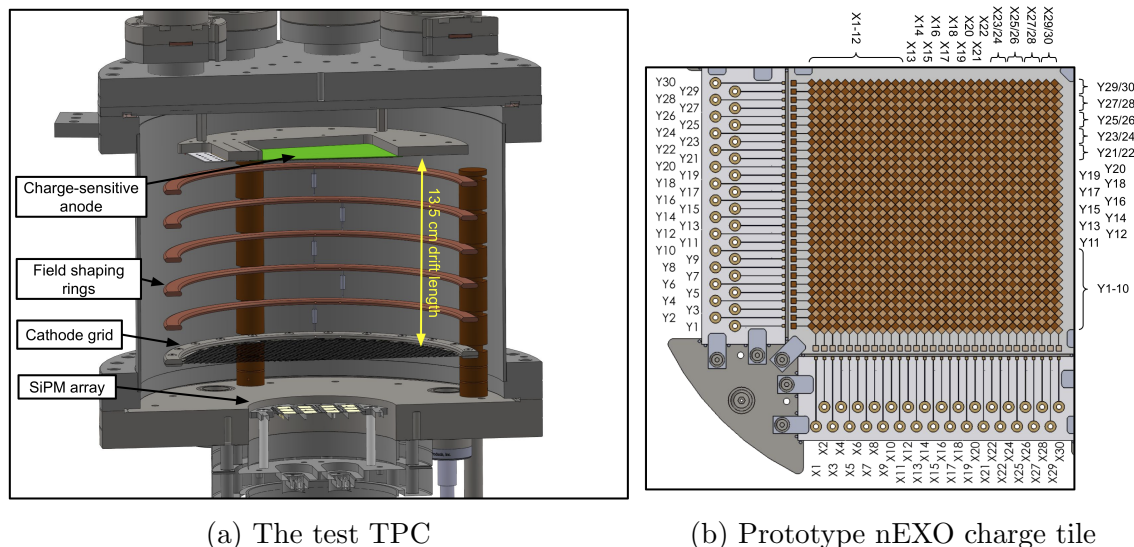


Figure 5.5: Left: CAD model of the test TPC used to demonstrate the deployment of a ^{127}Xe source for calibration of the electron lifetime; right: the prototype charge tile used in the demonstration, with the ganging of multiple strips into charge channels indicated.

rate at which activated xenon made it into the TPC. This injection procedure was carried out twice, with approximately half of the xenon being deployed each time. Data collected after the first injection was used to determine the optimal acquisition settings, while the data following the second injection was used for the electron lifetime measurements reported in Section 5.3.3. The pressure in the sample cylinder and trigger rate in the TPC over the course of the second injection are shown in Fig. 5.6b. The continued increase in trigger rate after the sample cylinder had been emptied suggests some latency in the mixing of the activated xenon into the TPC.

5.3.3 Electron Lifetime Calibration

Following the injection, recirculation continued for a few days to thoroughly mix the ^{127}Xe throughout the LXe. A number of datasets were then collected over the course of two days. Data from four of these datasets was included in the analysis discussed here². The electron lifetime, τ_e , was determined by fitting the charge, Q , as a function

²The specific datasets used are DS16, DS17, DS18, and DS22 of Run 31.

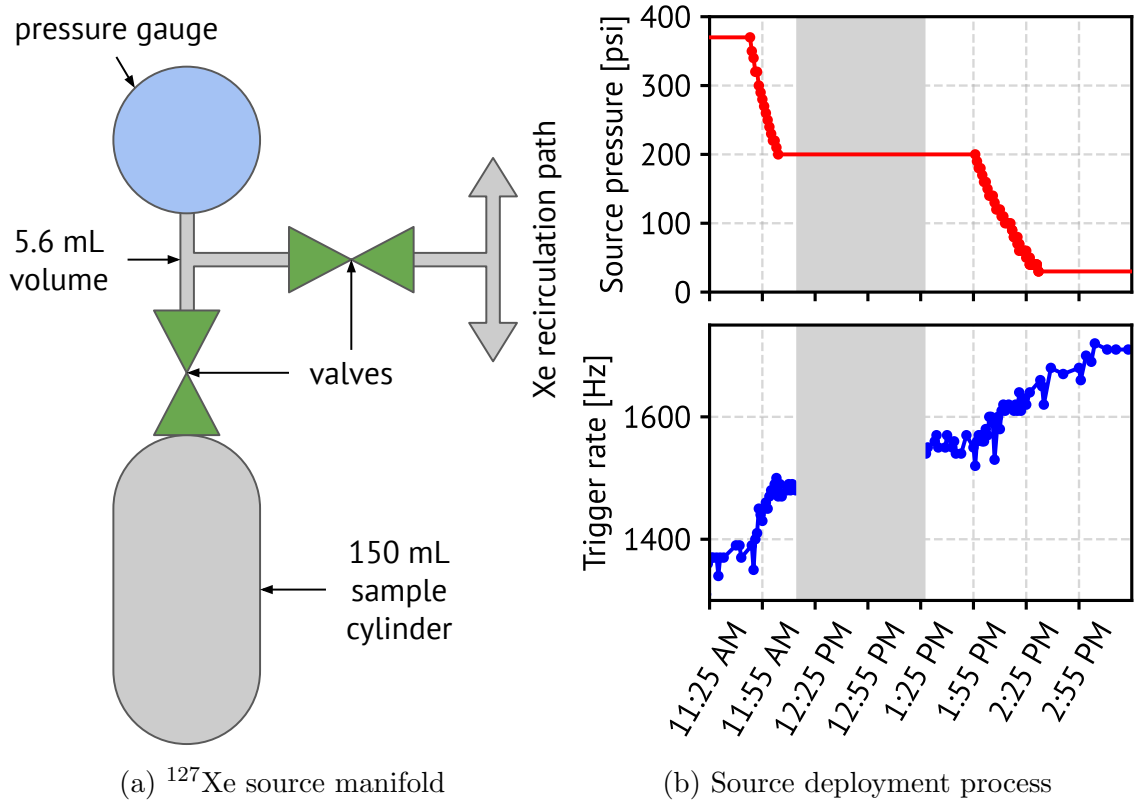


Figure 5.6: Left: schematic showing the source deployment manifold plumbed in to the xenon gas recirculation loop; right: source cylinder pressure and source activity in the TPC as a function of time during the source deployment process.

of the drift time, t , with the function,

$$Q(t) = I(t) \times Q_0 e^{-t/\tau_e} \quad (5.10)$$

where Q_0 is the initial charge produced in the event and $I(t)$ accounts for screening of the charge signal by the positive xenon ions³. For a given event, the total charge was calculated by summing the energies from all charge channels with a pulse height greater than the threshold of $3 \times$ the baseline RMS. The corresponding drift time for each channel was computed from the difference between the 90% rise time and the scintillation trigger time. Cuts were then applied to keep only those events with at

³Equivalently, $I(t)$ accounts for the position-dependent weighting potential near the anode.

least one x and one y channel above threshold and with charge-weighted average x and y coordinates within ± 15 mm of the center of the TPC. Data that survived the cuts was binned by the charge-weighted average drift time for the event, with data from one such bin shown in the histogram in Fig. 5.7a. A joint fit of two Gaussians and an exponentially-decaying background term then allowed for the location of both charge peaks to be determined for each bin.

The resulting charge peaks across all bins are plotted as a function of the drift time in Fig. 5.7b. To extract the electron lifetime, this data was fitted independently with Eq. (5.10) for each of the low and high energy peaks. The function $I(t)$ was computed analytically as described in Ref. [188] by integrating the image charge density resulting from the positive ions, which are treated as stationary on the timescale of an event. The integration was done over two strips, corresponding to the mean number of charge channels with above-threshold signals.

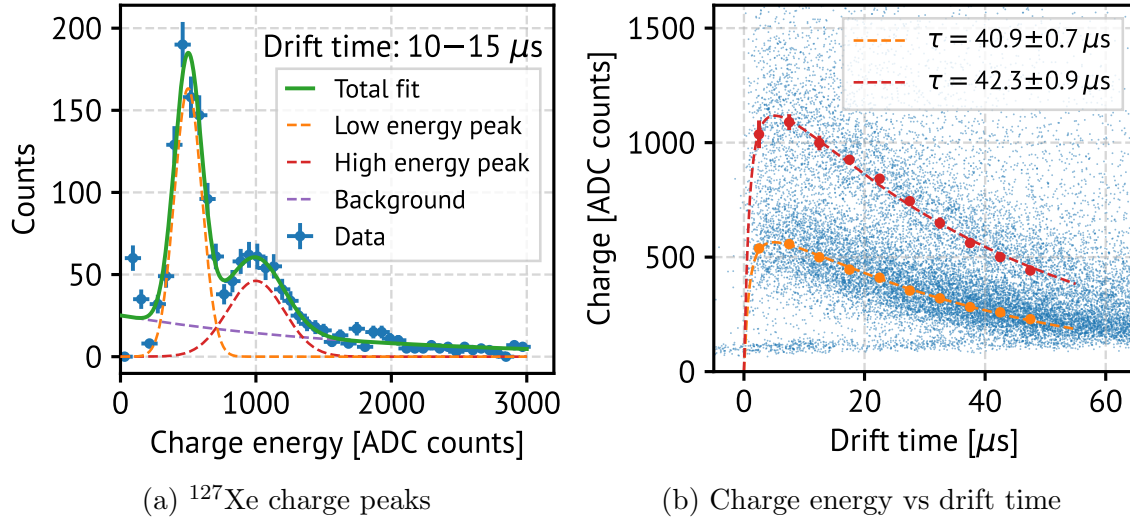


Figure 5.7: Left: plot showing the charge energy in a single drift time bin, with the joint fit and fits to the individual peaks and background spectrum also shown; right: charge energy as a function of drift time for both the low and high energy peaks along with the fits used to extract the electron lifetime.

The electron lifetimes obtained from the low and high energy peaks were $(40.9 \pm 0.7) \mu\text{s}$ and $(42.3 \pm 0.9) \mu\text{s}$, respectively, showing consistency to within $\sim 1.5\sigma$. While the calibration scheme worked as intended, the resulting lifetime indicated poor xenon

purity. This may be explainable by a small segment of PVC-insulated wire which was found during the disassembly of the TPC following the run. Outgassing from the PVC insulation could have worsened the purity, resulting in the poor electron lifetime that was measured.

5.4 Projected Calibration Performance for nEXO

In addition to the experimental demonstration of the calibration procedure in a test TPC, a quantitative projection of the calibration performance in nEXO was obtained through simulations. Simulated calibration datasets were produced using the `nexo-offline` package described in Section 4.5. Calibration events were reconstructed using a custom code specifically designed for this application.

To reconstruct the charge signals, channels with charge exceeding a preliminary $1200e^-$ threshold are selected, and the transverse event position is computed from the charge-weighted average of the channel coordinates. The five nearest channels on each side in both x and y (20 total) are used to compute the energy, enabling channels which see below-threshold charge signals to be included at the expense of added noise. A full description of the electron lifetime reconstruction from simulated charge data is given in Ref. [196]. Here, only the light response calibration is discussed in detail.

5.4.1 Simulating the Light Response

The light response of the detector is modeled with multiple steps, with statistical fluctuations added at each one. For each simulated event, the number of scintillation photons, S_0 , is computed by NEST. As recombination fluctuations dominate the variance in light production, the distribution from which S_0 is drawn can be approximated as a Gaussian with a mean of μ_{S_0} and a variance of $\sigma_{S_0}^2$,

$$S_0 \sim \mathcal{N}(\mu_{S_0}, \sigma_{S_0}^2), \quad (5.11)$$

with the mean and variance determined by the calibration source used. The number of photons collected by a sensor, S_{col} , is sampled from a Binomial distribution,

$$S_{\text{col}} \sim \mathcal{B}(S_0, \varepsilon_{\text{QE}} \cdot \varepsilon_{\text{LM}}(x, y, z)) \quad (5.12)$$

where the number of trials is S_0 , and the probability is the product of the PTE from the true lightmap, ε_{LM} , and the SiPM quantum efficiency, ε_{QE} . The true lightmap was computed with Chroma, a ray tracing code for high energy physics that uses CUDA-enabled GPUs for fast photon transport simulations [213]. A simulated dataset of 5×10^8 photon “bombs” positioned uniformly throughout the TPC volume, each containing 10^4 photons, was binned in r and z (bin widths ~ 0.25 mm) and the PTE in each bin was calculated. Gaussian smearing with a width of 1 bin was then applied to give a smooth, continuous, azimuthally-symmetric lightmap (left panel of Fig. 5.9).

The measured number of scintillation photons is the sum of the number collected by the SiPMs and the number resulting from correlated avalanches. The latter is a Poisson random variable with an expectation given by the former times a factor Λ ,

$$S_{\text{av}} \sim \mathcal{P}(S_{\text{col}} \cdot \Lambda). \quad (5.13)$$

The reconstructed number of photons, correcting for correlated avalanches, is computed with

$$S_{\text{det}} = \frac{S_{\text{col}} + S_{\text{av}}}{1 + \Lambda}. \quad (5.14)$$

The simulation used $\varepsilon_{\text{QE}} = 0.186$ and $\Lambda = 0.2$, the projected values for nEXO. Fig. 5.8a shows the detected charge and light from the simulation prior to any charge or light response calibrations. The absence of a lightmap correction is evident in the vertical smearing resulting from the varying PTE across the detector; this is also visible in the histogram in Fig. 5.8b. The histogram also shows that with a perfect lightmap calibration, the peak width can be corrected back to the true width up to statistical fluctuations. A 20 mm standoff cut is applied to the data to remove events near detector surfaces and an energy selection cut (defined by the red dashed line in Fig. 5.8a) selects only the 408 keV peak since the 236 keV peak is near nEXO’s energy

threshold. 90% of events survive the fiducial cut; 47.2% of those survive the peak selection. This results in 42.5% of the simulated events contributing to the lightmap calibration.

A single calibration dataset consists of a set of event coordinates and the PTE calculated by dividing S_{det} by the expected number of detected photons, $\mu_{S_0} \cdot \varepsilon_{\text{QE}}$ (center panel of Fig. 5.9). This data is used to produce a reconstructed lightmap (right panel of Fig. 5.9) using one of the strategies described in the next section. The accuracy of the lightmap calibration for a given dataset can then be determined by comparing the reconstructed and true lightmaps.

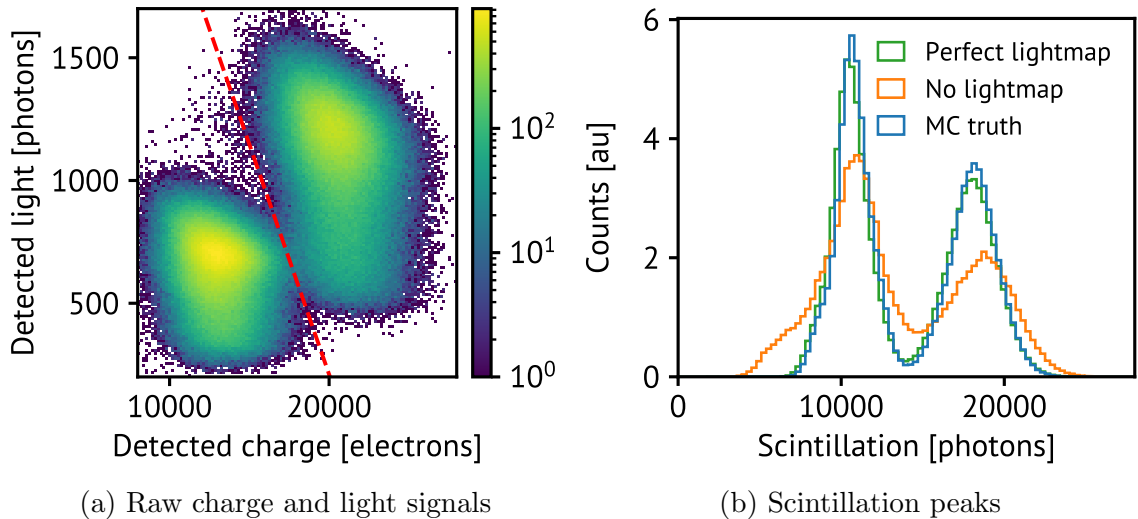


Figure 5.8: Left: simulated charge and light signals detected in nEXO from a ^{127}Xe calibration source prior to the application of any detector response calibration; right: histogram of reconstructed scintillation quanta after correction a constant scaling factor (orange) and a perfect lightmap (green) compared to the true distribution of scintillation quanta. Correction with a perfect lightmap cannot perfectly reconstruct the distribution due to statistical fluctuations in the number of photons detected by the SiPMs.

5.4.2 Lightmap Reconstruction Strategies

Complete three-dimensional lightmap reconstruction with sufficient precision is non-trivial. This is in large part due to the sparsity of calibration events when the dataset

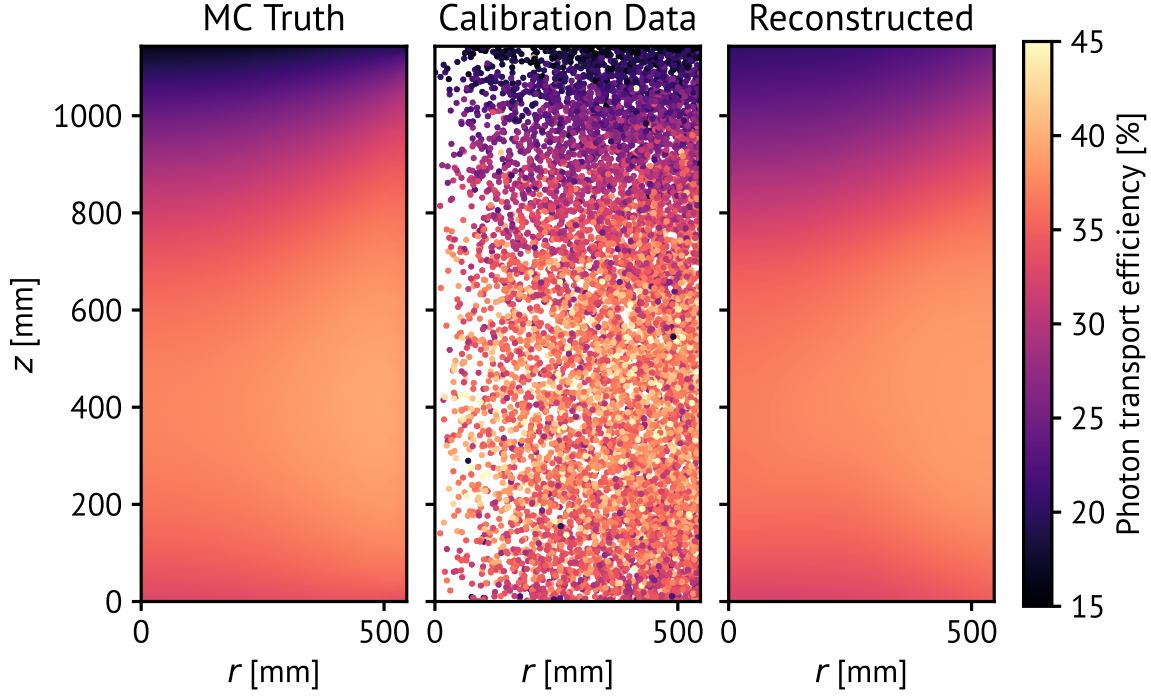


Figure 5.9: The sequence of steps followed during the simulation of a lightmap calibration campaign. The left panel shows the true lightmap from which the scintillation signals are sampled in the simulation. The middle panel shows the reconstructed calibration data with the colors showing the measured PTE. The right panel shows the lightmap as reconstructed using the data in the middle panel.

is diluted into three dimensions, as opposed to just one in the electron lifetime case, as described in Section 5.2. To illustrate, we address the problem of reconstructing the lightmap in a small subvolume of the detector over which the PTE is approximately constant. The PTE in that subvolume can be estimated from the mean of the calculated PTEs for all events, given by,

$$\hat{\varepsilon}_c = \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{S_{\text{det},i}}{\mu_{S_0} \cdot \varepsilon_{\text{QE}}} \quad (5.15)$$

where there are N_c calibration events indexed by i and $S_{\text{det},i}$ is the number of scintillation photons detected in each event. The error on this quantity depends on the variance at each step in the statistical chain described in the previous section. Setting all sampled quantities to their mean values gives a rough estimate of the contribution at each step, indicating that the dominant contribution comes from the scintillation peak width, σ_{S_0} . The uncertainty in ε_c can therefore be approximated by,

$$\sigma_{\hat{\varepsilon}_c} \approx \frac{\hat{\varepsilon}_c}{\sqrt{N_c}} \cdot \frac{\sigma_{S_0}}{\mu_{S_0}}, \quad (5.16)$$

and the required number of calibration events in that subvolume is then given by

$$N_c \approx \left(\frac{\sigma_{S_0}}{S_0} \right)^2 \left(\frac{\hat{\varepsilon}_c}{\sigma_{\hat{\varepsilon}_c}} \right)^2. \quad (5.17)$$

To achieve a total lightmap error of $\sigma_{\text{LM}} < 0.005$, we should aim to have $\sigma_{\hat{\varepsilon}_c}/\hat{\varepsilon}_c < 0.005$ in each subvolume. With a source line width of $\sigma_{S_0}/\mu_{S_0} \approx 0.1$, this implies $N_c \gtrsim 400$ events. What constitutes a single subvolume is not well defined. It should be related to the length scale over which the lightmap changes, yet this varies by position throughout the lightmap, with more rapid variations with position near the cathode and the anode and approximately constant efficiency throughout much of the center of the TPC. The specific requirement on the number of calibration events will depend on the lightmap reconstruction scheme used. Below, some potential schemes are discussed.

Histogram

Consider a calibration campaign in which a ^{127}Xe activity of 2 Bq is maintained for two weeks. This would result in a dataset of $\sim 10^6$ decays in the 408 keV peak. One method of reconstructing the lightmap from the calibration data is to histogram the measured efficiencies in x , y , and z . The total active volume in the TPC is $\sim 1.4 \times 10^9 \text{ mm}^3$, meaning that a cubic voxel containing 400 events it would have a side length of $\sim 80 \text{ mm}$. This is much larger than the spatial scale on which the lightmap varies. Evidently a histogram is not well suited to this problem.

Kernel Smoothing

An alternative strategy is to use a kernel smoothing (KS) algorithm. With this approach, the estimate of the lightmap at location \vec{x} is given by

$$\varepsilon_{\text{KS}}(\vec{x}) = \frac{\sum_{i=0}^{N_t} \frac{S_{\text{det},i}}{\mu_{S_0} \cdot \varepsilon_{\text{QE}}} \cdot w_i(\vec{x})}{\sum_{i=0}^{N_t} w_i(\vec{x})}, \quad (5.18)$$

where N_t is the total number of calibration data events and w_i is the kernel. It is common to use a Gaussian kernel, defined as

$$w_i(\vec{x}) = \exp\left(-\frac{|\vec{x} - \vec{x}_i|^2}{2\sigma_k^2}\right), \quad (5.19)$$

where \vec{x}_i is the position of calibration event i and σ_k is a smoothing length scale which must be chosen based on the data. Due to the Gaussian weighting, only those events with distances less than a few σ_k from \vec{x} will contribute to the sum. If a large number of events are distributed uniformly throughout the detector, the sum over discrete events in Eq. (5.18) can be replaced by an integral over the mean event density, ρ :

$$\sum_{i=0}^{N_t} \exp\left(-\frac{|\vec{x} - \vec{x}_i|^2}{2\sigma_k^2}\right) \approx \int d^3x' \rho \exp\left(-\frac{|\vec{x} - \vec{x}'|^2}{2\sigma_k^2}\right). \quad (5.20)$$

With a uniform event density the calculation is translation-invariant, so for convenience we can shift our coordinate system such that $\vec{x} = 0$. The effective number of events contributing to the lightmap, N_{eff} , can then be calculated as,

$$N_{\text{eff}} = \int d^3x' \rho \exp\left(-\frac{|\vec{x}'|^2}{2\sigma_k^2}\right) = \rho(2\pi\sigma_k^2)^{3/2}. \quad (5.21)$$

As expected, the number of events contributing depends on ρ and σ_k . We can now equate the effective number of events to the required number of calibration events

per subvolume (N_c from Eq. (5.17)) and solve for ρ :

$$\rho = \frac{1}{\sigma_k^3 (2\pi)^{3/2}} \left(\frac{\sigma_{S_0}}{\mu_{S_0}} \right)^2 \left(\frac{\hat{\epsilon}_c}{\sigma_{\hat{\epsilon}_c}} \right)^2. \quad (5.22)$$

Using the same relative source width and lightmap uncertainty as before, we find that the required event density is $\sim 25/\sigma_k^3$. The dataset of 10^6 calibration events gives $\sigma_k \sim 30$ mm, much smaller than the voxel side length in the naive histogram approach.

This is a useful benchmark to which lightmap reconstruction techniques can be compared. However, it is not a computationally feasible way to implement lightmap reconstruction. The lightmap should be a simple function that takes a single argument — a position within the TPC — and returns an efficiency. In the KS method, this efficiency is computed by summing over all calibration events with a weighting that depends on the Euclidean distance between the location of the event and the input position. As this requires computing N_t pairwise distances for each input, it is computationally expensive for large N_t . Moreover, the use of a single parameter σ_k means that the smoothing scale must be constant across the full detector. In reality, there are some regions where the lightmap changes rapidly and a smaller scale is needed, and some where the lightmap varies little and the smaller scale simply adds statistical noise.

Deep Neural Net

For problems such as the one described here, where a complicated function needs to be fitted to a large dataset, a neural net is a natural solution. In particular, it avoids the problem of specifying smoothing length scales, as it is adaptive to the training data. It scales well to large datasets, as after the one-time computational cost of training, it can be sampled rapidly. A deep neural net (DNN) using five hidden layers in a basic multi-layer perceptron architecture was designed for this purpose. The DNN takes as input a vector (x, y, z, x, y, z) of position coordinates within the TPC, and outputs a number from 0 to 1 corresponding to the lightmap value at that point. The three coordinates were repeated in the input vectors to improve extraction of features

with radial dependence. To train the DNN, the reconstructed event positions from simulated calibration datasets were passed as inputs, with the calibration efficiency computed using Eq. (5.15) used as the outputs. During training, half of the calibration events were allocated as validation data. The training loss and validation loss were monitored during training to prevent against overfitting to the training data. For each simulated calibration dataset, the training process was repeated five times and the resulting networks were averaged to produce the final DNN. The architecture and hyperparameter values were chosen to minimize σ_{LM} following a comprehensive study of the dependence of DNN performance on these parameters.

5.4.3 Lightmap Calibration Performance

For both the KS and DNN methods, the performance of the lightmap reconstruction was assessed with calibration datasets of varying sizes. This assessment was done by sampling the reconstructed lightmaps on a three-dimensional grid of points throughout the TPC and computing the relative uncertainty as a function of position. Recall from Section 5.1 that the lightmap error parameter σ_{LM} is the square root of the sample variance of the relevant uncertainty over the full volume. For both reconstruction methods, σ_{LM} was computed for 25 datasets ranging in size from 10^3 to 10^6 events. For each dataset size, the calibration events were sampled with replacement from the full set of 10^7 events.

The left panel of Fig. 5.10 shows the lightmap error achieved with the KS method as a function of the number of calibration events. The lightmap error shows the expected scaling with $\sqrt{\rho}$, or equivalently $\sqrt{N_t}$, as would be expected from Eq. (5.22). The parameter σ_k was chosen at each dataset size by scanning many values and choosing the one that minimizes σ_{LM} ; the optimal values of σ_k are shown in the right panel of Fig. 5.10.

For the DNN, which operates as a black box, it is helpful to visualize the calibration accuracy across the TPC to understand what drives performance. Fig. 5.11 shows the relative lightmap error throughout the TPC for one of these 25 datasets of each size. While the full lightmap reconstruction is done in three dimensions, the

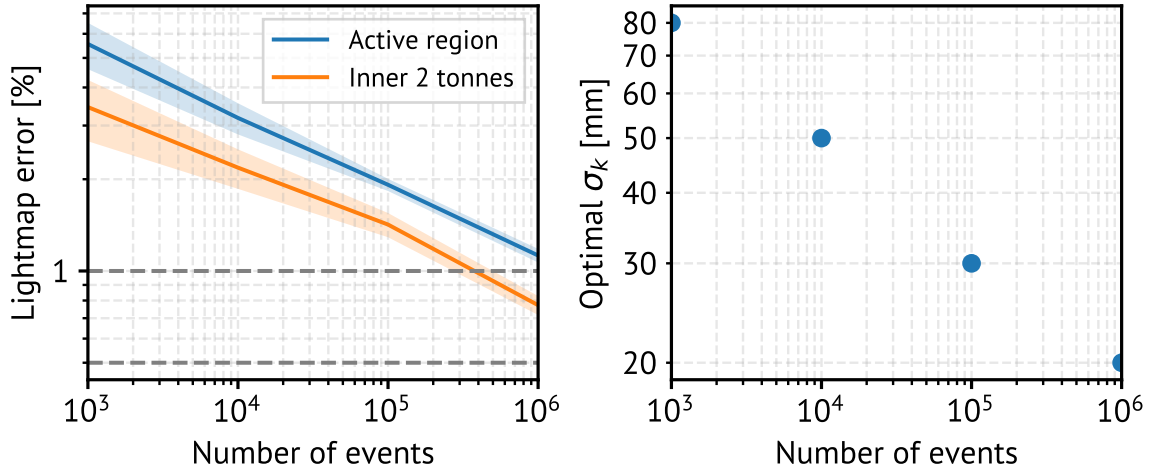


Figure 5.10: Left: lightmap error attainable as a function of the number of calibration event using the Kernel Smoothing reconstruction method; right: the optimal smoothing length, determined by scanning over different parameter values, at each dataset size. The smoothing length on the right at each dataset size is the one used to give the lightmap reconstruction performance shown in the left.

plots show only a single azimuthal slice. The qualitative performance matches expectations, with errors decreasing significantly with increasing dataset size. When calibration data is limited, the DNN tends to produce an overly flat lightmap, resulting in significant overprediction (underprediction) of the PTE in regions where the true PTE is lower (higher) than the average across the TPC. When trained on larger datasets, the DNN is better able to match the steeper gradients in the lightmap, and the errors recede toward the edges of the TPC. The best performance is consistently achieved in the inner two tonnes of the detector; fortunately, this is also the region which drives the physics sensitivity and for which precise calibration is most important.

The mean and $\pm 1\sigma$ range of the σ_{LM} values obtained at each dataset size with the DNN are plotted in Fig. 5.12a. With few calibration events, the performance scales much faster than with the KS method, though it begins to flatten out for larger calibration datasets. As shown in the figure, $\sim 5 \times 10^5$ calibration events are required to meet the nEXO target lightmap error of 0.005 in the inner two tonnes. This can be directly interpreted in terms of the lightmap contribution to the energy resolution by

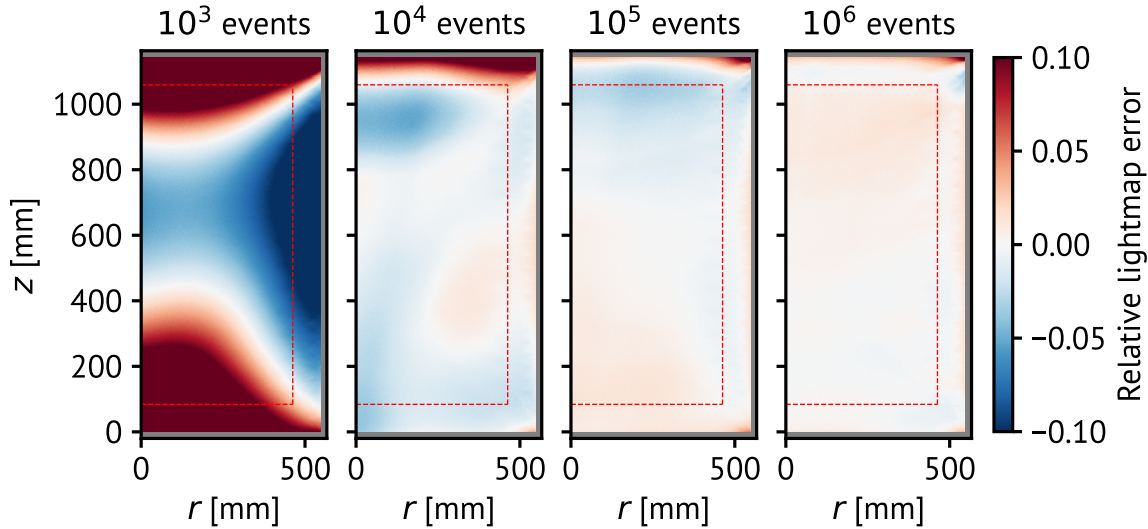


Figure 5.11: The lightmap error as a function of position in an azimuthal slice of the TPC for datasets of four different sizes. The gray region around the edges shows the fiducial cut, while the red dashed line encloses the inner two tonnes of the LXe.

fixing the energy resolution to its nominal value of 0.008 and computing the marginal contribution from the lightmap as a function of the number of events. Fig. 5.12b shows the resulting energy resolution achievable for a range of calibration periods, assuming a constant activity of 1 Bq is maintained in the detector. This study shows that the lightmap error contributes negligibly to the total energy resolution in the inner two tonnes of the detector after only ~ 5 days of calibration data, without any detector downtime. Unlike the default lightmap calibration scheme, which requires interruption of data collection every few months, this calibration could be done on an ongoing basis.

5.5 Conclusions and Prospects for nEXO

In this chapter it was demonstrated that a ^{127}Xe calibration scheme can be used to calibrate both the electron lifetime and lightmap in nEXO. The production of such a source was demonstrated, and it was shown that the expected radioisotope content

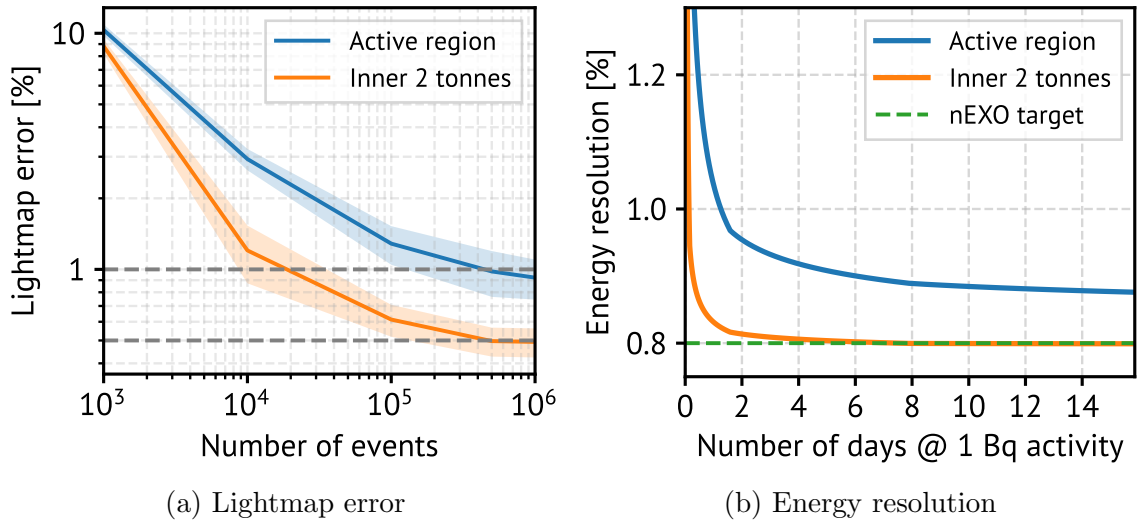


Figure 5.12: Left: the lightmap error parameter as a function of the number of calibration events in nEXO for both the inner two tonnes and the full detector volume; right: the resulting energy resolution as a function of calibration time assuming the nominal nEXO energy resolution model with only the lightmap contribution varied.

matches expectations and does not include any problematic background sources for nEXO. Using this source, nEXO would be able to meet the lightmap error target in less than 6 days if a 1 Bq activity were maintained for the duration of the calibration. While it was not discussed here, a similar analysis was done for the electron lifetime, showing that with the same source activity, it could be calibrated to sufficient precision on a daily basis (see Ref. [196] for details). This work has important implications for nEXO as it removes the requirement for detector downtime during calibration, allowing for more frequent calibrations without a loss of livetime.

Chapter 6

Development of a Xenon Purifier for nEXO Using High-Purity Zirconium

Xenon purification is essential to meeting the energy resolution target in nEXO, as electronegative impurities capture drifting charges and attenuate the charge signals before they are collected. While this attenuation is corrected with the calibrated electron lifetime, binomial fluctuations in the charge collected worsen with more attenuation, so improving the electron lifetime is still critical. The electron lifetime is related to the concentration C_i of each species i of impurity as

$$\tau_e = \frac{1}{\sum_i k_i C_i} = \frac{M_{\text{Xe}}}{\rho_{\text{LXe}} \sum_i k_i x_i} \quad (6.1)$$

where k_i is the rate constant for electron attachment in LXe for species i , and in the second step, the concentration C_i has been replaced with the mole fraction x_i using the LXe density, ρ_{LXe} , and the molar mass of xenon, M_{Xe} . Since the composition of impurities is usually unknown, yet O_2 is typically the dominant contributor to the measured electron lifetime, the oxygen-equivalent concentration, x_{O_2} , is used to

represent all impurities by convention. Rewriting Eq. (6.1) in terms of x_{O_2} gives

$$\tau_e = \frac{M_{\text{Xe}}}{\rho_{\text{LXe}} k_{\text{O}_2} x_{\text{O}_2}} \approx 10 \text{ ms} \times \frac{47 \text{ ppt}}{x_{\text{O}_2}} \quad (6.2)$$

where $k \approx 9.8 \times 10^{10} \text{ M}^{-1}\text{s}^{-1}$ is used for nEXO's planned operating field of 400 V/cm [214]. The steady-state oxygen equivalent concentration in the TPC must therefore be at or below 47 ppt in order to meet the electron lifetime target. At the same time, nEXO's background budget allocates 80 $\mu\text{Bq}/\text{kg}$ of allowable ^{222}Rn activity for the purifier media, so a purifier design must satisfy this constraint as well.

6.1 Introduction to Xenon Purification

In steady state, the impurity concentration in a LXe TPC is limited by the continuous desorption of electronegative impurities from surfaces within the detector. Continuous purification of the xenon is therefore required to achieve and maintain the target purity. A standard purification technique uses continuous forced recirculation of the xenon through a dedicated external purifier. Purification is most commonly done in the gas phase, with a heater at the TPC outlet used to boil off the xenon and a recirculation pump to force it through the purifier, usually a heated bed of zirconium-alloy getter pellets. The purified GXe reenters the TPC through an inlet above the liquid surface where it recondenses. Purification in the liquid phase has also been demonstrated [215], allowing for significantly higher mass flow rates to be achieved at the expense of lower purification efficiency.

Getter purifiers operate by absorbing electronegative impurities on their surface, and the combination of high temperature and large interstitial spacing within the getter matrix enhances the diffusion from the surface into the bulk volume of the material. While this technique has proved to yield exceptional purification efficiencies, the getter pellets are a source of radioactive backgrounds. A sufficiently large mass of getter materials could emanate radon at levels that will limit the physics sensitivity. Indeed, the dominant radon backgrounds in EXO-200 are suspected to have originated from the steady-state population of ^{222}Rn atoms originating from the

getter materials [194]. Furthermore, a diffusion-dominated emanation process could be enhanced at the high temperatures at which typical zirconium-alloy getters operate. The considerations above make clear the need for an effective purifier that does not exceed nEXO's strict background budget. Following a brief introduction to getter materials, the development of a new xenon purifier will be described in detail.

6.1.1 Getter Materials

A getter is a material which exhibits a pumping action by adsorbing gases on its surface to form stable chemical compounds [216, 217]. This process, known as chemisorption, requires that the surface of the material be clean. A clean surface can be obtained in two ways: by evaporating and depositing a film of material in situ, or by heating the material to sufficiently high temperatures as to cause the diffusion of atoms from the surface into the bulk. A material falling into the first category is called an evaporable getter, while one in the second is a non-evaporable getter (NEG). Common getter materials are barium, titanium, and zirconium, with barium and titanium used for evaporable getters and alloys of zirconium used as NEG's [218]. NEG's are preferred where the surface area needed for the required pumping speed is impractically large.

6.1.2 Zirconium Alloys

Zirconium is an excellent gettering agent due to its ability to dissolve large amounts of oxygen [219]. At room temperature, zirconium is capable of adsorbing an oxide layer as large as 1.5 nm to 2 nm in a range of oxidation states [220]. When heated beyond room temperature, rapid diffusion of oxygen from the surface to the bulk is observed, with very little desorption [221, 222]. This makes zirconium an ideal constituent of a gettering material, often in an alloy with aluminum for increased diffusivity at high temperature [216, 217]. However, the activation temperature required for typical zirconium-alloy getters exceeds 700 °C, making in situ activation difficult. The alloy St 707 [223], consisting of 70% zirconium, 24.6% vanadium, and 5.4% iron, was designed with a lower activation temperature of 450 °C for this specific purpose. Its pumping characteristics have been well studied [224–226]. Sintered pellets of

St 707 are used in the commercial SAES Monotorr purifiers that have found use in many xenon-TPC-based experiments due to their excellent purification performance of xenon gas [227].

Despite their superior purification capabilities, zirconium alloys like St 707 are not optimized for use in rare event searches. In particular, the mitigation of radioactive contaminants was not a design consideration, and as a result, purifiers which use St 707 tend to be significant sources of ^{222}Rn , with specific activities varying from $18\ \mu\text{Bq/kg}$ to $293\ \mu\text{Bq/kg}$ [228–231]. As these getters are maintained at $450\ ^\circ\text{C}$ or higher during normal operation, the diffusion of ^{222}Rn through the material could be enhanced, resulting in more transport into the TPC¹. Commercial purifiers are also designed with far more capacity than is required for the steady-state purification of xenon that has already been purified prior to its deployment in the TPC. For low enough impurity concentrations, transport into the getter bulk is not required as the surface alone can capture impurities until it saturates, provided it is periodically regenerated.

6.1.3 High-Purity Zirconium

Evidently nEXO’s specific requirements for a getter are in some ways more stringent, and in other ways more relaxed, than those for commercial getters. This has motivated investigations into alternative purifier designs that meet nEXO’s unique set of design requirements. High-purity zirconium (HPZr) has emerged as a compelling alternative due to its natural gettering properties and its commercial production in large quantities for the nuclear industry.

Indeed, zirconium alloys are commonly used as structural materials for nuclear reactors due to their strong corrosion resistance and favorable mechanical properties [233, 234]. Zirconium is found naturally in the mineral zircon (ZrSiO_4), usually in combination with 1% to 3% hafnium due to their similar chemical properties. Hafnium, in contrast to zirconium, has a high neutron absorption cross section, so

¹The temperature is relevant for diffusion-driven ^{222}Rn transport, which was expected to be the dominant mechanism. More recent evidence suggests that the ^{222}Rn transport through metals may instead be recoil-driven, rendering the high operating temperatures less of a concern [232].

zirconium alloys for nuclear applications are required to have low levels (typically < 100 ppm) of hafnium contamination [235]. While removal of actinides is not a goal of the techniques for separating zirconium and hafnium, it was suspected that lower uranium and thorium concentrations may be a byproduct of the production process of HPZr. The natural gettering properties of zirconium and the potential for obtaining a more radiopure product than in commercial getters were the primary motivations for studying HPZr as an alternative purifier media for nEXO.

6.2 Design of a Xenon Purifier

For a custom xenon purifier to be used in nEXO, it must be capable of handling the nominal recirculation rate of 350 SLPM [184] while maintaining its purification efficiency. It must also meet the ^{222}Rn specific activity design goal of $80 \mu\text{Bq/kg}$, chosen based on the allowable steady-state ^{222}Rn population of 600 atoms and assuming 15 kg of purifier media. A prototype purifier using 1 kg of HPZr pellets was designed as a first step toward this end.

6.2.1 Getter Pellets

HPZr pellets to be used in the purifier were procured in two batches from ALB Materials, a metal supplier which operates out of Nevada. The two batches were ordered six months apart and were produced in separate lots. The first batch consisted of 1 kg of pellets. This batch was used in the prototype purifier after having been sent away for assays and then returned to Stanford. The second batch consisted of another 1 kg of pellets along with dedicated samples for other measurement techniques. Table 6.1 summarizes the samples that were measured while Section 6.4.2 describes the assay results in detail. In addition to the assays, a test was conducted on the regeneration of the gettering surface through baking, and the surface was imaged to understand the potential impacts on radon emanation and active gettering surface area.

Table 6.1: List of the high-purity zirconium samples that were procured from ALB Materials.

Number	Lot	Quantity	Shape	Length	Width	Height
0	ALB-202012	—	—	—	—	—
1	ALB-202102	1 kg	Cylinder	2 mm	2 mm	2 mm
2	ALB-202107	1 kg	Cylinder	2 mm	2.4 mm	2.4 mm
3	ALB-202107	5 pieces	Cylinder	20 mm	2 mm	2 mm
4	ALB-202107	5 pieces	Slab	20 mm	10 mm	2 mm

Activation Test

Sample 4 was sent to CERN to test the surface activation procedure. The test was done by measuring the surface content of the HPZr using X-ray photoelectron spectroscopy (XPS) at a range of temperatures. A HPZr slab was placed in the XPS vacuum chamber (base pressure 4×10^{-10} mbar), then heated to 160 °C, 250 °C, 300 °C, and 350 °C. The sample was kept at each temperature for one hour, after which the XPS data was collected for 25 minutes. The surface composition as a function of temperature is shown in the left panel of Fig. 6.1 while the absolute measurements of carbon and oxygen are shown in the right panel.




Figure 6.1: Left: the surface fraction of four elements as a function of temperature; right: the relative peak area for the C and O 1s peaks as a function of temperature. Data for a standard Ti/Zr/V NEG is shown for comparison.

The XPS data show a decrease in zirconium oxide peaks accompanied by an increase in metallic zirconium peaks, indicating that the activation process is removing adsorbed oxygen from the surface, as intended. The metallic zirconium peak at 350 °C was compared to that of a typical Ti/Zr/V NEG, revealing a similar extent of surface oxygen reduction. However, the decrease in carbon and oxygen on the surface was delayed compared to the typical NEG by ~ 100 °C, suggesting that higher temperatures are required to begin regenerating the surface. A slight increase in carbon content was observed above 250 °C, potentially attributable to uptake from the measurement chamber and the formation of carbides. While this data does show regeneration of the gettering surface at only 350 °C, it indicates a potential reduction in the available gettering surface due to the formation of compounds that cannot be easily removed at that temperature. This is a known phenomenon which can be reduced by keeping other regions of the vacuum system at low temperatures during regeneration bakeouts [236].

Surface Imaging

A pellet from sample 2 and an St 707 pellet from a SAES purifier cartridge were imaged with a scanning electron microscope (SEM) to compare their surfaces. As the St 707 is sintered, it was expected to have a significantly higher surface area and a commensurately higher radon emanation rate for the same bulk concentration. Section 6.2.1 shows an SEM image of each pellet at similar magnifications.

As expected, the HPZr surface looks flat with some surface defects, while the St 707 exhibits structure over a range of length scales, with many deep caverns and extrusions visible. This should limit the impurity capacity of the HPZr accordingly, though makes the capacity more amenable to estimation due to the relatively featureless surface geometry.

6.2.2 Geometry & Pressure Drop

Like the standard SAES Monotorr purifiers, a custom purifier should be constructed using a column of getter pellets through which xenon gas can flow. The fluid dynamics




Figure 6.2: SEM images of an St 707 pellet (left) and a HPZr pellet (right). The HPZr pellet has a comparatively featureless surface with a significantly lower surface area than the St 707 pellet.

of gas flow over a column of pellets is well studied due to the extensive use of this geometry in catalyst bead reactors. Findings from the existing literature can be used in the design of a purifier for nEXO to ensure that the pressure drop across the pellet bed is manageable given the operating pressure and flow rate. The following analysis uses correlations obtained from the analysis of many previous experimental results [237] and CFD simulations [238].

Before beginning the quantitative analysis, some qualitative features of flow over a packed bed of cylindrical pellets are discussed. As we will see, the pressure drop depends on the geometry of the pellets and the tube containing them, along with the Reynolds number, Re , for the flow. In particular, wall effects can significantly increase or decrease the pressure drop depending on Re . At low Re , a boundary layer forms near the walls and the resulting friction increases the total pressure drop. At high Re , the boundary layer is thin and its effect subdominant to that of the increased void fraction near the walls, decreasing the pressure drop. As might be expected, the overall contribution from wall effects diminishes with increasing tube diameter, with the effects becoming negligible when the tube diameter is more than 10 times the pellet diameter.

Another qualitative feature worth mentioning is the significant local variation in

the flow speed over the pellets. The azimuthally-averaged flow speed oscillates by around a factor 2 with distance from the wall with a periodicity on the order of the pellet diameter. Moreover, when the full three-dimensional flow profile is considered, inhomogeneities in the local flow velocity as large as a factor of 8 can arise as a consequence of the random packing of the pellets. While this is unavoidable from the design perspective, it means that areas exposed to high flow should be expected to saturate earlier, and as a result the gettering surface area (and by extension the purification efficiency) may decrease before the purifier becomes fully saturated.

To quantitatively understand the effect of design parameters on the pressure drop, we can use a modified version of the Ergun equation [239] for the friction factor, Ψ_W , as a starting point:

$$\Psi_W = \frac{A_W}{\text{Re}_W} + B_W, \quad (6.3)$$

where A_W and B_W replace the empirically-derived parameters found by Ergun to account for viscous and inertial energy losses, respectively. We consider the specific case of a column of length L and diameter D containing cylindrical pellets whose diameter and height are both equal to d_p . For such a configuration, a combined analysis of thousands of experimental results has yielded

$$A_W = 190 \quad \text{and} \quad B_W = \left(2 \frac{d_p^2}{D^2} + 0.77\right)^2 \quad (6.4)$$

as the best parameters. The modified Reynolds number for flow through a pellet bed, Re_W , is defined as,

$$\text{Re}_W = \frac{\rho u_s d_p}{\mu(1 - \varepsilon)M} \quad (6.5)$$

where ρ , u_s , and μ are the density, superficial velocity, and dynamic viscosity of the fluid, respectively, d_p is the pellet diameter, ε is the bulk porosity of the column, and M is a correction to account for wall effects, defined as

$$M = 1 + \frac{2d_p}{3D(1 - \varepsilon)}. \quad (6.6)$$

The friction factor can be expressed in terms of the pressure drop, ΔP , and the

pellet bed and fluid flow parameters as

$$\Psi_W = \frac{\varepsilon^3}{1 - \varepsilon} \frac{\Delta P d_p}{\rho u_s^2} \frac{1}{L M}, \quad (6.7)$$

allowing for the pressure drop to be calculated directly using

$$\Delta P = \frac{LM\rho u_s^2(1 - \varepsilon)}{d_p \varepsilon^3} \left[190 \frac{\mu(1 - \varepsilon)M}{\rho u_s d_p} + \left(2 \frac{d_p^2}{D^2} + 0.77 \right)^2 \right]. \quad (6.8)$$

When designing a purifier, the amount of getter material to be used should be chosen based on the desired impurity capacity. Supposing a fixed mass M_{Zr} of purifier pellets is specified, D becomes the design parameter which in turn will set L , with $L \propto D^{-2}$. For a fixed mass flow rate during recirculation, the choice of D also determines u_s . Ignoring the secondary dependence of M and B_W on D , we find,

$$\Delta P \propto \begin{cases} D^{-6}, & \text{Re}_W \ll \frac{A_W}{B_W}, \\ D^{-8}, & \text{Re}_W \gg \frac{A_W}{B_W}. \end{cases} \quad (6.9)$$

The strong dependence of the pressure drop on D along with the decrease in wall effects at large D underlines the need for a column with a sufficiently large diameter. At the same time, there are clear drawbacks to using a column that is too short. For one, the length over which an impurity could be captured is shorter, potentially reducing purification efficiency. Additionally, the flow through a short column may be dominated by entrance or exit effects, rendering the prior analysis inapplicable. As a rule of thumb, the minimum column diameter can be set by the allowable pressure drop or by the condition that $D/d_p > 10$, whichever is larger. With that condition satisfied, the column should be made as long as possible. Fig. 6.3 shows the calculated pressure drop as a function of the column diameter for the small-scale prototype purifier (1 kg zirconium purifying at 1 SLPM) and a nEXO-scale purifier (5 kg zirconium purifying at 350 SLPM).




Figure 6.3: The pressure drop across a cylindrical column of zirconium pellets as a function of the column diameter for a small-scale prototype (blue) and a nEXO-scale purifier (orange). The bulk porosity ε is taken to be 0.35 and the shaded regions show the range $0.3 < \varepsilon < 0.4$.

6.2.3 Design & Assembly

Based on the aforementioned design considerations, a Lesker DN35CF 304L SS full nipple is used to house the pellet column. With an inner diameter of 35 mm, this results in $D/d_p \approx 17$ and a column height (for the full 1 kg of pellets) of 24.7 cm. The CF nipple is 91.4 cm long, with the pellets confined to the middle. This allows for the pellets to be baked at the desired temperature of $> 350^\circ\text{C}$ while the flanges at the ends remain near room temperature to avoid annealing the CF knife edges or gaskets. The pellet column is sandwiched between two sintered SS 40 micron filters coupled to positioning rods with custom SS coupling pieces. The positioning rods meet the end flanges of the tube via additional coupling pieces, with a SS spring on one end to compress the pellet bed and minimize motion following assembly.

Baking is done with two 1 kW Tempco band heaters stacked end-to-end around the center of the tube. As SS is a poor thermal conductor, two aluminum clamshell pieces clamp around the tube with the heaters tightened around them to distribute

the heat evenly across the tube and prevent hotspots at the contact points. The complete design is shown in the annotated CAD model in Fig. 6.4.




Figure 6.4: Annotated CAD drawing showing the design of the custom purifier, with certain components called out.

For simplicity and convenient integration with the existing recirculation manifolds in the Stanford LXe lab, 1/2-inch tubing extending from both CF flanges at the end connects to Swagelok bellows-sealed valves with 1/4-inch VCR fittings. This allows the purifier to be easily plumbed into existing LXe purification manifolds. The Stanford Custom Purifier (SCP), as it later became known, was assembled according to this design and mounted in an aluminum T-slot structure.

After assembly, the SCP was plumbed into the recirculation manifold for the large LXe test stand, as shown in the photo on the right of Fig. 6.5. This allowed for it to be pumped out through the manifold during the bakeout. A LabVIEW program was written to control the bakeout, with safety thresholds applied to multiple thermocouples and a required human interaction interval to ensure that the bakeout was not continuing unsupervised. Initial bakeouts were done to monitor the temperature stability and understand the temperature profile along the length of the SCP column.




Figure 6.5: Left: the HPZr pellets being loaded into the CF nipple; right: the fully-assembled custom purifier mounted to a T-slot structure, plumbed into a Stanford LXe TPC test stand, and instrumented with thermocouples to monitor the bakeout during regeneration.

A first purification test was then conducted using a test TPC (the “Long TPC” described in Appendix B) in the lab at Stanford and a ^{220}Rn calibration source. While this did not provide a sufficiently precise measurement of the purity for quantitative conclusions to be drawn as to the purifier’s performance, a > 100 ms electron lifetime was achieved, indicating that the purifier was performing comparably to the SAES in the same system.

6.3 Measurements of Purification Efficiency

6.3.1 The SLAC Xenon Purity Monitor

Following the preliminary test in the Stanford test stand, more quantitative testing was conducted at SLAC using the Xenon Purity Monitor (XPM). The XPM consists of a 10.9 cm LXe drift region between a cathode and anode grid inside a 6.5 cm diameter copper vessel. 266 nm laser pulses are transmitted via a multimode optical fiber into the LXe, where they create photoelectrons on a gold-coated cathode. The cathode, a window grid (6 mm above the cathode and 5 mm below the cathode grid), and the anode (2.5 mm above the anode grid) are all AC-coupled to a low-noise charge preamplifier. Photoelectrons create a rising signal when passing through the window grid and then a falling signal when passing through the anode grid; the difference in amplitude gives a measure of the charge captured by impurities in the drift region. The XPM is controlled by a LabVIEW program which includes a pulse-fitting module to extract the electron lifetime for each measurement. For each measurement, a background-only waveform is produced by averaging a sequence of waveforms collected with a mechanically-actuated shutter blocking the laser just before the input to the fiber. This is then subtracted from an average of waveforms collected with the shutter open, allowing for electromagnetic pickup from the laser firing to be removed. This provides a near-real-time measurement of the xenon purity. A complete description of the SLAC XPM can be found in Ref. [240].

The xenon recirculation manifold includes a SAES Monotorr PS3-MT3 purifier which uses St 707 as the gettering agent. The SCP was plumbed into the recirculation manifold in parallel, allowing for either purifier to be opened to the recirculation path. A magnetically-coupled piston pump [241, 242] with buffer volumes on either side forces recirculation at a tunable flow rate. A heater at the XPM outlet boils off the xenon prior to its recirculation in the gas phase. A mass flow meter and Baratron gauges at the inlet and outlet of the XPM allow the mass flow rate and pressures to be monitored. The electric field in the drift region is maintained at 250 V/cm during regular operation.

6.3.2 Purification Measurement Procedure

In preparation for the purity measurements, the SCP was baked for two hours at 365 °C to regenerate the gettering surface. A relay controlled by an Omega temperature control unit reading a thermocouple on the aluminum surface inside the jaws of one of the heaters was used to switch the heaters on or off to maintain the target temperature. For the duration of the bakeout, the SCP was pumped out through the recirculation manifold.

After filling the XPM with LXe and recirculating through the SAES to achieve a high initial purity, the measurement procedure was followed for both the SAES and the SCP. The procedure was as follows:

1. Dose the xenon with ~ 1 torr of air
2. Start recirculation with both purifiers bypassed to allow the air to mix throughout the xenon
3. Close the bypass valves and open the valves for the purifier under test
4. Continue recirculation until the purity reaches the measurement limit

A short section of tubing separated from a roughing pump by a valve was used as a dosing volume from which air at a known pressure could be metered into the xenon flow path. A TC gauge attached to a tee connector at the inlet of the roughing pump provided a measurement of the pressure in the dosing volume. To fill the dosing volume, the valve was opened with the pump off, putting the volume at atmospheric pressure. The pump was then turned on and the pressure was monitored, allowing for the valve to be closed once the target pressure was reached.

The SAES purifier was tested first. After the dosing step, the purity dropped to an unmeasurably low level, much worse than was predicted given the measured pressure in the dosing volume. Following some investigation, it was discovered that the TC gauge used to set the pressure in the dosing volume was broken, and would read from 1 torr to 1.2 torr even when exposed to atmospheric pressure. This indicated that the amount of air injected into the xenon was likely much larger than intended. Nonetheless, the purity reached ~ 4 ms after a weekend of recirculation through the SAES.

The procedure was then repeated for the SCP with a functional TC gauge. With the correct air pressure in the dosing volume, the purity reached $\sim 100 \mu\text{s}$ following the dosing step. It then rose to $\sim 7 \text{ ms}$ after two days of recirculation. The pressure and mass flow readings were recorded regularly throughout the measurements.

6.3.3 Impurity Transport Model

The evolution of the O_2 -equivalent concentration, x , can be modeled as

$$\frac{dx}{dt} = -\frac{\varepsilon f}{\tau_C} x + \frac{\Lambda}{n}, \quad (6.10)$$

where the first term describes the removal of impurities through purification, and the second, the constant outgassing of impurities into the LXe. τ_C is the time required to recirculate the total xenon mass which depends on the mass flow rate and is typically on the order of a few hours. The rate at which changes in purity propagate throughout the liquid phase is captured by the dimensionless equilibration factor, f . For $f = 0$, the addition/removal of impurities never results in a change in the measured impurity concentration in the liquid phase; for $f = 1$, the change in the measured concentration follows instantaneously. The quantity of interest in this test was the purification efficiency, ε , defined as the fraction of impurities removed in a single pass through the purifier. Λ is the rate of addition of impurities via outgassing and is taken to be constant over the course of a measurement, while n is the number of moles of xenon in the system. The solution to Eq. (6.10) is

$$x(t) = x_0 e^{-\varepsilon f t / \tau_C} + \frac{\Lambda \tau_C}{n \varepsilon f}, \quad (6.11)$$

where x_0 is the initial concentration prior to the start of purification. As f and ε appear only once in Eq. (6.11) as a product, it is impossible to fit each of these independently. Fortunately, the purification performance of the SAES purifier is well understood, with typical impurity removal efficiencies of $> 99.99\%$ [227]. In this analysis, $\varepsilon = 1$ was used for the SAES purifier, allowing f to be determined and used to calculate ε for the SCP.

This model ignores the distribution of impurities between the liquid and gas phases in accordance with Henry’s law. The Henry’s law volatility constant, calculated to be 62.5 for O₂ in xenon [215, 243], is defined as the ratio of the mole fraction in the gas phase to that in the liquid phase at equilibrium. However, the XPM contains 1.5 kg xenon and has a recirculation manifold volume estimated at 13.9 L from rough measurements of the tubing. There is therefore $\gtrsim 100\times$ more O₂ in the liquid phase than the gas phase. With active recirculation, the equilibration timescale is much shorter than τ_C , and therefore the impurities remaining in the gas phase can be ignored.

6.3.4 Results

The electron lifetime over the full XPM run is shown in Fig. 6.6. Each point shown on the plot is an average of ten datapoints collected. The error bars reflect the systematic uncertainty in the measurements due to the resolution of the oscilloscope from which the voltage measurements were made. At very low purity, the anode pulse becomes indistinguishable from noise, while at very high purity, the anode pulse amplitude will be smaller than the cathode pulse amplitude by less than the noise. The measurement noise therefore sets effective lower and upper limits on the measurable electron lifetime, respectively.

To account for this, for each measurement, 100 cathode and anode voltages were sampled from a Gaussian distribution centered at the measured value with a standard deviation of 0.1 mV. These sampled voltages were then used to compute a distribution of electron lifetime measurements. The central values and error bars were taken from the median and 16th and 94th percentiles of each distribution.

The impurity transport model was fitted to each of the two purification periods independently, allowing for the products εf to be determined. Fig. 6.7 shows the results for the SAES and Fig. 6.8 shows the results for the SCP. For both cases, the goodness of fit was computed with a modified χ^2 where the uncertainty for each datapoint was taken as the upper (lower) 1σ uncertainty for those points falling above (below) the median. The goodness of fit is shown in the blue panel on both plots. In




Figure 6.6: The xenon purity as measured by the SLAC XPM over the course of the run. Different stages of the run are shown by the dashed lines and shading as indicated in the legend.

the SAES data, the equivalent O_2 concentration starts at a much higher point than for the SCP data due to the inadvertent injection of more air than intended. This could also explain the higher outgassing rate in the SAES data, as the interior surfaces of the XPM system likely adsorbed a layer of air which was then slowly released during the purification phase.

The equilibration factor obtained from the SAES purification stage was $f = 0.82 \pm 0.02$, indicating that changes in purity propagate quickly throughout the LXe volume relative to the recirculation timescale. Combining this with the SCP purification data, we find that $\varepsilon = 0.83$ for the SCP. While this indicates less efficient purification than the SAES, the worse efficiency can be easily compensated by a slight increase in the mass flow rate (corresponding to a reduction in τ_C).

It is important to note that this purification efficiency was determined with the SCP operating at a lower flow rate and inlet pressure than the target in nEXO. More work is required to understand the impact of the flow rate on the purification performance and to determine whether a purifier with a design similar to the SCP will be capable of operating under nEXO's nominal conditions. If purification efficiency




Figure 6.7: The oxygen-equivalent concentration as a function of time when purifying with the SAES. The parameters obtained when Eq. (6.11) was fitted to the data are shown in the upper right.

decreases with a higher mass flow rate, multiple SCP-like purifiers could be used in parallel to ensure the flow rate through any one is within the optimal operating window.

6.4 Radioactivity of High-Purity Zirconium

6.4.1 Radioassay Measurement Campaign

While the aforementioned design work and purification tests were being conducted, in parallel, the samples procured from ALB Materials were being subjected to an extensive campaign of radioassay measurements. The chemical composition of the HPZr was measured via glow discharge mass spectrometry (GDMS) at the National Research Council of Canada (NRC) and inductively-coupled plasma mass spectrometry (ICPMS) at Pacific Northwest National Laboratory (PNNL). A report from a GDMS analysis performed by Perfect Analysis Technology (PAT) on a previous lot was also provided by ALB Materials. The University of Alabama and Laboratori




Figure 6.8: The oxygen-equivalent concentration as a function of time when purifying with the SCP. The parameters obtained when Eq. (6.11) was fitted to the data are shown in the upper right.

Nazionali del Gran Sasso (LNGS) independently counted samples of pellets with low-background germanium counters. Radon emanation was measured by the Laurentian University nEXO group at SNOLAB and by the nEXO group at SLAC. A summary of the measurements, with references both to the sample number and the identifier used on the nEXO Materials Database, is given in Table 6.2.

Table 6.2: The radioassay measurements conducted by various participating institutions on the high-purity zirconium samples.

Sample	Technique	Institution	MaterialsDB Entry
0	GDMS	Perfect Analysis Technology	—
3	GDMS	NRC	R-177.1.1
1	ICPMS	PNNL	R-164.1.2
1	Ge counting	University of Alabama	R-164.1.1
2	Ge counting	LNGS	R-164.2.1
2	Rn emanation	SNOLAB/Laurentian	R-164.2.2
2	Rn emanation	SLAC	R-164.2.3

6.4.2 Measured Activity

The most problematic backgrounds for nEXO are those likely to produce events with energies near the $0\nu\beta\beta$ Q -value of 2.458 MeV, such as the β -decay of ^{214}Bi to ^{214}Po which can be accompanied by the release of a 2.448 MeV γ . ^{214}Bi is a daughter of the ^{238}U decay chain, shown in Fig. 6.9. The ^{222}Rn activity is of particular relevance as this gaseous radionuclide can emanate from detector surfaces into the LXe or GXe volume. Under the assumption that the decay chain is in secular equilibrium, measurements of the activity of any ^{238}U daughters will allow for the ^{222}Rn activity to be inferred; however, it is possible that broken equilibrium of the decay chain could result from various production processes. The suite of radioassay measurements performed on the HPZr samples allow for both a direct measurement of the ^{222}Rn emanation rate and an understanding of the degree to which the decay chain may be out of equilibrium. A summary of the radioassay measurements is shown visually in Fig. 6.10 and the measurement techniques are discussed in turn.

Elemental Composition

The sensitivity of the ICPMS and GDMS measurements was limited to 1 ppb. A measurement of ^{238}U at this concentration corresponds to a ^{222}Rn activity of 12 mBq/kg — well above the acceptable level for nEXO. Any measurable traces of ^{238}U would therefore preclude the HPZr from use in nEXO if the equilibrium of the decay chain were preserved. The ICPMS and recent GDMS measurements both report a measurable concentration of ^{238}U at just above the measurement sensitivity. The results are in mild tension at the level of a factor 2, though the measurements in question were done using different lots of the HPZr, and near the sensitivity limits of the instruments.

Germanium Counting

The first germanium counting measurement from the University of Alabama established the broken equilibrium of the decay chain through tension in the inferred ^{238}U concentration with the ICPMS measurement. A limit was placed on the ^{226}Ra activity




Figure 6.9: The ^{238}U decay chain. Elemental analysis techniques measure concentrations of uranium. Assuming the concentration of uranium is approximately that of ^{238}U (natural abundance 99.3%), and that the decay chain is in secular equilibrium (activity of all daughters is set by the ^{238}U activity), the activity of ^{222}Rn and the problematic ^{214}Bi decay can be inferred.

of 8.37 mBq/kg, corresponding to a ^{238}U concentration of < 678 ppt for a decay chain in equilibrium, while a uranium concentration of (2000 ± 400) ppt was measured using




Figure 6.10: Results from the radioassay measurement campaign. The left axis shows the uranium concentration and the right shows the $^{226}\text{Ra}/^{222}\text{Rn}$ activity. The axes are aligned using the assumption of secular equilibrium of the decay chain and the points are color-coded based on the axis on which they are plotted. The non-alignment of the two sets of points on the vertical scale indicates that the secular equilibrium of the decay chain has been broken. The solid orange line indicates nEXO’s target activity for purifier media while the shaded band shows the range of St 707 activities reported in previous assays.

ICPMS. The broken equilibrium works in nEXO’s favor, with lower concentrations of the more problematic ^{226}Ra than expected from the measured uranium concentration. A more sensitive measurement could allow for the degree of broken equilibrium to be established.

The germanium counting also revealed a troubling inconsistency with the GDMS report from PAT. ^{176}Lu , which has a half life of 3.78×10^{10} years and a natural abundance of 2.5%, was observed with a specific activity of (86.3 ± 9.1) mBq/kg. This corresponds to a concentration (1.65 ± 0.18) ppm, significantly higher than the reported limit of < 5 ppb from the PAT GDMS, the only other assay which reported a lutetium concentration. While lutetium is not itself a concern for nEXO, the discrepancy in the measurement suggests inconsistent production between lots of the HPZr. This is of significant importance, as the radioassay measurements reported here are

only useful insofar as they generalize to future lots that may be used for nEXO.

The germanium counting measurements from LNGS are more sensitive to ^{226}Ra and report a measured specific activity of (0.53 ± 0.50) mBq/kg. This activity corresponds to a concentration of (43 ± 40) ppt uranium, more than 20 times lower than the concentration measured with GDMS. This activity is well above the ^{222}Rn target of $80 \mu\text{Bq/kg}$. The degree to which this is a problem for nEXO depends on what fraction of the ^{222}Rn produced in these decays ends up emanating into the flowing xenon and making its way to the TPC. An understanding of this can also inform both the dynamics of the decay and radon transport, and the distribution of radium in the HPZr pellets, discussed below.

Radon Emanation

Two separate radon emanation measurements were conducted using the same set of pellets a few months apart, both using an electrostatic counter (ESC). The first, done at SNOLAB, reported a ^{222}Rn emanation activity of (0.95 ± 0.27) mBq/kg, while the second at SLAC measured (0.373 ± 0.089) mBq/kg. The source of the discrepancy is not understood. Both measurements exceed the nEXO target. But if the data allow for something about the mechanism by which radon is released from the zirconium matrix to be inferred, then strategies for mitigating the radon emanation may become clear. We can calculate the ratio of the ^{222}Rn activity to that of the ^{226}Ra ; this ratio gives the fraction of radium atoms which make it from the solid zirconium matrix into the flowing gas. This gives 1.79 ± 1.77 and 0.685 ± 0.704 using the SNOLAB and SLAC measurements, respectively. The uncertainties are large enough that it is impossible to draw strong conclusions from this data. But the fact that both measurements are close to 1 indicates that a higher fraction of radon atoms escape from the zirconium matrix than might be expected, suggesting that the majority of the radium may be near the surface of the metal. If so, this could allow for its removal, a possibility discussed in the next section.

6.5 Future Work

This chapter presented the design of a purifier for nEXO which uses HPZr instead of St 707 as the getter material. An in situ surface regeneration technique was demonstrated, and the purification efficiency of the purifier was measured to be comparable, though slightly inferior, to that of the standard SAES Monotorr purifiers. Despite these promising results, the unacceptably high radon emanation rate prohibits this design from being adopted in nEXO as-is. However, based on the radon emanation and radium decay measurements collected throughout the extensive radioassay campaign, there is reason for optimism.

6.5.1 Dry Distillation for Radium Removal

If the radium in the HPZr pellets is indeed concentrated on the surface of the metal, as preliminary evidence suggests, then it may be removable using dry distillation. This technique exploits the difference in the vapor pressures of radium and zirconium, enabling the substantially more volatile radium to be baked off, leaving a more radiopure zirconium surface behind.

This possibility is currently being explored by the nEXO group at PNNL. Small-scale tests done by heating a single zirconium pellet to 1700 °C in an electrothermal vaporizer have shown that barium, a chemical surrogate for radium with a much higher concentration, can be removed using this process. A full-scale test using the 500 g of pellets originally sent for radon counting at SLAC is expected in the near future. Following the bake, a second radon counting measurement will quantify the success of this technique.

6.5.2 Capacity and Regeneration Measurements

Supposing the radium removal procedure is found to be successful, a realistic purification scheme for nEXO will be needed. This will involve determining the quantity of pellets and number of purifiers required. With multiple purifiers, one could be

operational while the other(s) are being regenerated, allowing for nonstop purification during the time that nEXO is running. Designing such a scheme will require a measurement of the impurity capacity of each purifier and the degree to which the surface can be regenerated following saturation. Measurements of this kind could be carried out at the SLAC XPM using a purifier containing fewer pellets and a larger impurity dose to prevent the saturation from taking prohibitively long. Dosing with pure oxygen rather than air would eliminate other systematic effects, including that of nitrogen absorption on the getter surface.

Chapter 7

The Search for Micron-Scale Interactions using Optically-Levitated Microspheres

This chapter presents a recent search for non-Newtonian interactions at the micron scale, parametrized as modifications to the ISL, using optically levitated microspheres. Part of this chapter has been adapted from Ref. [244], of which I am a primary author. As was discussed in Chapter 3, the most stringent constraints on modifications to the ISL of this type have so far been established by measuring responses of mechanical oscillators. These techniques search for variations of one component of a temporally-modulated force vector at one or more harmonics of the modulation frequency. In contrast, the technique presented here uses the time-dependent behavior of the full 3-dimensional force vector to search for new interactions. By conducting the search for a potential signal across multiple spatial dimensions and harmonics, the unique spectral fingerprint of a potential interaction can be exploited, especially to constrain and confirm a possible discovery. The chapter begins with an introduction to optical tweezers — the experimental platform at the heart of this work — with a particular focus on their application at the size scales relevant to what follows.

7.1 Introduction to Optical Tweezers

Pioneered by Arthur Ashkin in the 1970s [245], optical tweezers have found a multitude of applications in physics and biology due to their capability of confining and manipulating nanometer- to micron-sized objects, including atoms and molecules, dielectric particles, and living cells [246–251]. The principle underlying optical tweezers is conservation of momentum; the interaction between a photon and a trapped object causes a change in the photon's path which transfers momentum to the trapped object. This section shows how this effect can be used to confine objects near the focus of a laser beam for their use as force sensors. For the discussion, we assume the laser beam has a single TEM₀₀ (Gaussian) mode.

7.1.1 The Force from a Laser

The force imparted by a laser beam on an object can be resolved into a scattering component along the beam axis and a gradient component in the perpendicular direction. The appropriate mathematical description of these components depends on the relative length scales of the wavelength of the trapping beam, λ_b , and the diameter of the object, d . In the Rayleigh regime, where $d \ll \lambda_b$, the object can be modeled as a point dipole in an electric field oscillating perpendicular to the beam axis. The dipole radiation produced in response creates a net force along the beam axis and accounts for the scattering term. As the intensity of a Gaussian beam falls off radially, the dipole also interacts with the electric field gradient and experiences a force whose time-averaged direction is toward the beam axis.

If instead $d \gg \lambda_b$, the force components are best understood using geometric optics. As a simple example, consider a dielectric sphere with a refractive index higher than that of the surrounding medium that is positioned slightly off the beam axis, as shown in Fig. 7.1. By solving the Fresnel equations for rays entering and exiting the sphere, the momentum imparted by the reflected and refracted rays can be determined; the result is a force vector with components along and toward the beam axis [252]. The experiment described in this chapter uses spheres with $d \approx 10\lambda_b$, near the edge of the regime in which geometric optics can be applied. An exact solution

applicable to spheres of arbitrary size can be obtained through Mie theory in the form of an infinite series [253, 254], though it is not necessary for this application given that calibrations can be performed using known forces.

7.1.2 Creating an Optical Trap

Independent of the length scale, the radial component of the force points toward the beam axis, and the condition for radial confinement is satisfied. For small displacements, the restoring force is approximately linear and well described by Hooke's law. Confinement along the direction of the beam axis is more complicated. Provided there is an appreciable intensity gradient over the length scale of the trapped object, a similar argument from geometric optics can be used to show that there is indeed a confining force [255]. This condition is met when the numerical aperture (NA) is sufficiently high, and the Rayleigh range, z_R , is of order d . For the application described in this chapter, however, it is necessary to bring a comparatively large object in close proximity to the microsphere. As this would obstruct the diverging portion of a tightly-focused trapping beam, a low-NA beam is used instead, and there is negligible optical confinement along the beam axis.

Instead, equilibrium in the axial direction can be established by orienting the trap vertically, with the beam propagating upward [256]. Radiation pressure then balances the gravitational force, with the precise equilibrium position set by the laser power. The required power is minimal in the vicinity of the trap focus and increases with distance from it in both directions. Unlike the radial directions, this does not create a harmonic trap; instead, active feedback is used to adjust the trap beam power in response to small displacements. With a proportional gain feedback scheme, a harmonic trap can be established in the vertical direction for small displacements in regions where the power curve is locally monotonic¹.

¹Depending on the particulars of the trap and sphere, the required power as a function of z can have minima on either side of the trap focus with a local maximum in between [257, 258]. Regardless of the sphere's location on this curve, linear feedback can still be used to stabilize its position provided that the sign of the feedback signal is chosen appropriately and the sphere does not cross an extremum.




Figure 7.1: Optical forces on a sphere in the geometric optics regime and under the paraxial approximation. Two rays are shown, along with the forces imparted by the transmitted and reflected portions at each interface (labeled \vec{F}_T and \vec{F}_R , respectively, at one of the interfaces). As the left ray is closest to the beam axis (gray dashed line), the intensity is higher, and the resulting forces are larger. The sum of the optical forces can be resolved into a scattering component, \vec{F}_{scat} , and a gradient component, \vec{F}_{grad} . The scattering component balances the gravitational force, \vec{F}_g , while the gradient component gives the net restoring force toward the beam axis.

7.1.3 The Equation of Motion

A sphere of mass m in an optical trap with a spring constant k will have a natural frequency of $\omega_0 = \sqrt{m/k}$. We take the z axis to be pointing up, in the direction of the beam. Interactions with the surrounding medium can be modeled as damping described by the translational damping coefficient, β_{tr} . At high vacuum, the viscous damping is caused by collisions with gas molecules in the molecular flow regime. The

translational damping coefficient in such circumstances can be calculated to be

$$\beta_{\text{tr}} = \pi R^2 \sqrt{\frac{128}{9\pi}} \left(1 + \frac{\pi}{8}\right) \frac{p}{v_T} \quad \text{with} \quad v_T = \sqrt{\frac{k_B T}{m_0}} \quad (7.1)$$

for a sphere of radius R in a gas at pressure p and temperature T , where v_T is the thermal velocity of a gas molecule of mass m_0 and k_B is Boltzmann's constant [259]. When used for force sensing, the response of the sphere to a driving force of amplitude F at an angular frequency ω is desired. The equation of motion in x is then simply that of a driven, damped harmonic oscillator,

$$m\ddot{x} + \beta_{\text{tr}}\dot{x} + m\omega_0^2 x = F \cos(\omega t) \quad (7.2)$$

which can be solved with the Ansatz

$$x(t) = A \cos(\omega t - \phi). \quad (7.3)$$

Plugging this into Eq. (7.2) and simplifying allows for the amplitude, A , and phase, ϕ , of the sphere's response to be determined:

$$A = \frac{F}{\sqrt{m^2 (\omega_0^2 - \omega^2)^2 + \beta_{\text{tr}}^2 \omega^2}} \quad \text{and} \quad \phi = \arctan \left[\frac{\beta_{\text{tr}} \omega}{m (\omega_0^2 - \omega^2)} \right]. \quad (7.4)$$

Near the resonance, one can show via Taylor expansion that the amplitude response takes the form of a Lorentzian function with a linewidth of β_{tr}/m . At frequencies far below the resonance, $\omega \ll \omega_0$, the response is flat with an amplitude of $F/(m\omega_0^2) = F/k$. In the absence of external driving forces, the dominant force on the sphere arises from collisions with gas molecules. By the fluctuation-dissipation theorem, the resulting force noise, S_F , can be expressed in terms of the damping coefficient as $S_F = 4k_B T \beta_{\text{tr}}$. This sets the limit on force sensing at low frequencies. At such low pressures, however, gas damping is insufficient to prevent perturbations near the resonance from knocking the sphere out of the trap. Derivative gain (velocity damping) feedback is therefore used to damp motion at the resonance.

7.1.4 Microspheres as Force Sensors

I have described a configuration in which a harmonic optical trap can be created in three dimensions: two using the beam profile itself, and the third using active feedback on the beam power. Like with a mechanical spring, the task of measuring a force reduces to that of measuring the resulting displacement. In the z direction, the way to do this is straightforward: a vertical displacement, Δz , causes a change in the optical path length of reflected light of $2\Delta z$ and is therefore encoded as a phase change, $\Delta\phi = 4\pi\Delta z/\lambda_b$. In x and y , we saw that the deflection of photons causes, by conservation of momentum, a displacement of the microsphere. If a fraction ξ of the photons in a trapping beam of power P_{trap} are deflected by an average angle $\Delta\theta$, then conservation of momentum in the radial direction gives

$$F_{\text{grad}} \approx \frac{\xi P_{\text{trap}} \Delta\theta}{c} \quad (7.5)$$

as the gradient component of the force. Measuring this deflection of the transmitted light allows for the corresponding displacements of the sphere to be inferred. This is typically done by using a split photodiode, which allows the deflection to be measured as the measured power simultaneously decreases in one segment and increases in an adjacent segment. In the Mie regime, the dependence of the spatial profile of the transmitted light on the sphere's position is nontrivial. In the linear regime, in which small displacements cause small deflections, however, a complete understanding of this relationship is unnecessary, as the sensor response can be calibrated using known forces applied to the sphere. The experimental realization of this force sensing scheme is described in detail in the following section.

7.2 Experimental Setup

The experimental platform described in this work uses optically levitated silica microspheres [260], with nominal diameters of either $d = 7.56 \mu\text{m}$ or $d = 9.98 \mu\text{m}$ trapped using vertically-oriented optical tweezers. An early generation of the setup is described in Ref. [261], while an updated diagram of the optical layout, showing the

essential components in the setup used to produce the result described in this chapter, is shown in Fig. 7.2.

7.2.1 The Laser

A 1064 nm laser diode (Innolume LD-1064-DBR-150) provides the seed for a linearly-polarized, single-mode fiber amplifier (IPG Photonics YAR-2-1064-LP-SF). The fiber amplifier output is split into a “trap beam” and a “reference beam” using a fiber-coupled polarization-maintaining beam splitter. Two acousto-optic modulators (AOMs) are used to shift the trap and reference beam frequencies and create a 125 kHz frequency difference between them. At the input optics (left half of Fig. 7.2), the fiber output is collimated ($f = 15.3$ mm) into free space, the residual elliptical polarization is removed, and the beam is expanded through a telescope ($f_1 = 50$ mm, $f_2 = 200$ mm) before it enters the vacuum chamber. Inside the vacuum chamber, it is focused to a waist of ~ 3.5 μm using an off-axis parabolic mirror (NA ~ 0.1) to create the optical trap.

7.2.2 Position Sensing

The light scattered off the microsphere in the forward direction is re-collimated by a second off-axis parabolic mirror. On the output optics side (right half of Fig. 7.2), a reducing telescope ($f_1 = 200$ mm, $f_2 = 50$ mm) focuses the beam through a 50 μm aperture to spatially select light from the microsphere while blocking stray light from elsewhere. The retroreflected light exits the trap through the input viewport and passes through the Faraday rotator in the opposite direction to the incoming light, allowing it to be separated based on its orthogonal polarization.

The forward scattered light and retroreflected light are separately combined with the reference beams to perform interferometric (heterodyne) measurements. The 125 kHz beat signal produced by each combined beam encodes the changes in the phase and amplitude of the trap beam originating from its interactions with the microsphere. A quadrant photodiode (QPD, Hamamatsu S5980) is used to measure deflections of the forward-scattered beam while the z photodiode (z PD, Thorlabs

DET100A2) measures the retroreflected beam. The power measured on these sensors is sampled at 500 kS/s and digitally demodulated, allowing the amplitude and phase of the trap beam to be reconstructed. The difference between the summed QPD amplitudes in the top and bottom halves of the sensor provides y , while that of the left and right halves gives x . The phase of the retroreflected light gives z .

The optical setup also includes a second QPD (the “DC QPD”, Thorlabs PDQ80A) to make a non-interferometric measurement of the x and y positions of the microsphere, along with a 640×480 pixel camera (Allied Vision Mako U-029B) currently being studied as a candidate position sensor. A photodiode (the “spin PD”, Thorlabs PDAPC4) measures the orthogonal (S-polarized) light exiting the trap. This can be used to reconstruct the rotation of the microsphere’s optic axis due to its intrinsic birefringence. Some light is also picked off and measured by two other photodiodes (the “input PD”, Thorlabs PDA20CS2, and “transmitted PD”, Thorlabs PDAPC4) before and after passing through the vacuum chamber. Feedback on the input PD signal is used to stabilize the trap beam intensity.

7.2.3 The Trapping Region

The geometry of the trapping region is illustrated in Figure 7.3. The microsphere is trapped near the focus, allowing the closest distance of approach of a source mass (“attractor”) with minimal interaction with the trapping beam. A piezo deflector outside the input viewport is used to apply velocity damping feedback on the x and y positions of the microsphere to stabilize it and allow operation at high vacuum ($\sim 10^{-7}$ hPa). The trap is surrounded by six electrodes (three pairs, one for each Cartesian axis, separated by ~ 8 mm). The electrodes are spray-coated with AquaDAG [262] to reduce reflections of stray laser light. The monopole electric charge on the microsphere can be controlled with single-electron resolution using a UV flash lamp [263, 264], allowing for a bias on the electrodes to exert a known force on the microsphere, thereby enabling the calibration of the optical position readout to force.

The attractor has a density pattern to source interactions coupling to mass [265]. Gold and silicon are used as the two materials generating a contrast in density, ρ , with




Figure 7.2: Simplified schematic showing the optics system used in this work. All components are shown from a top-down view except the cutaway showing the parabolic mirrors, electrode cube, and microsphere. Some beam-steering mirrors have been omitted. BS refers to a 50:50 beamsplitter; other splitting ratios are labeled explicitly. The cameras depicted here are assumed to include the required focusing optics.

$\rho_{\text{Au}} \sim 8\rho_{\text{Si}}$. By scanning the position of the attractor along y in close proximity to the microsphere at a frequency $f_0 = 3$ Hz, the microsphere is subject to a characteristic force with a distinctive spectral fingerprint in each Cartesian axis that is unlikely to be mimicked by backgrounds. Electromagnetic forces constitute an important source of background, and while the microsphere can be made to have zero overall charge, dipole and higher order moments persist and can couple to time-varying electric field gradients as the attractor is scanned. In order to mitigate this effect, a stationary “shield” (a microfabricated silicon fence structure, conformally coated with ~ 100 nm gold) is positioned between the microsphere and the scanning attractor

using a separate nanopositioning stage. The “top camera” (Allied Vision Mako U-029B) and the “side cameras” (Allied Vision Mako U-029B and U-130B) are used for real-time imaging of the trap region, allowing for the attractor and shield to be positioned precisely for the measurement.




Figure 7.3: Diagram of the experimental setup also defining the coordinate system used and providing the most important physical dimensions. A microsphere is held at the focus of the trapping beam while the density-patterned attractor oscillates by $170 \mu\text{m}_{\text{pp}}$ at $f_0 = 3 \text{ Hz}$ behind the stationary shield. The inset shows a side view with the microsphere diameter and the face-to-face separation between the microsphere and attractor identified. The Platinum Black coating on the attractor, which hides the underlying density pattern, has been omitted in this diagram.

7.2.4 Improvements Since the Previous Result

The first constraints on modifications to gravity using a levitated force sensor were reported in Ref. [266]. The present iteration of the experiment incorporates several improvements over the first. The amount of stray light scattering off the moving attractor and resulting in a background on the detectors was reduced $\gtrsim 100$ -fold. This




Figure 7.4: Time evolution of the force exerted on a $10\ \mu\text{m}$ microsphere by a hypothetical Yukawa interaction with the attractor ($\alpha = 10^6$, $\lambda = 10\ \mu\text{m}$). The z component of the force is not shown as it is only present when a vertical offset is introduced between the microsphere and the attractor. The secondary y axis shows the position of the center of the attractor as it is driven through one full cycle.

was primarily achieved by adding the aperture at the focus of the downstream telescope, and by coating the attractor with a $\sim 3\ \mu\text{m}$ -thick layer of Platinum Black [267], with a reflectivity at $1064\ \text{nm}$ that is $\lesssim 1\%$ that of the original gold. In addition, various parts of the vacuum chamber were stiffened to mitigate the effect of mechanical vibrations. Three accelerometers (Wilcoxon 731A) were added close to the vacuum chamber on a 3-axis mounting cube and a series of microphones (Adafruit MAX9814 Electret Microphone Amplifier) were placed around the lab to sense the acoustic environment near the setup. Optimal (Wiener) filters [268, 269] were then constructed to subtract environmental disturbances coherently sensed by the microsphere and the environmental sensors. Lastly, the noise in the force measurement was improved by over a factor of 10 along the x and y axes, and $\sim 5\times$ along the z -axis. This improvement was achieved primarily with better stabilization of the laser intensity delivered to the trapping region (relative intensity noise, $\text{RIN} \lesssim 5 \times 10^{-7}/\sqrt{\text{Hz}}$). With these upgrades, the forward-scatter detection system has a displacement sensitivity

of $\sim 10^{-11} \text{ m}/\sqrt{\text{Hz}}$, which translates to a force sensitivity of $\sim 10^{-17} \text{ N}/\sqrt{\text{Hz}}$ for the typical trap spring constant of $\sim 10^{-6} \text{ N/m}$.

7.3 Overview of the Search

The results presented in this chapter were obtained from measurements of three microspheres, one with nominal diameter $7.56 \mu\text{m}$ and the other two $9.98 \mu\text{m}$. Their masses were measured in situ, following the procedure described in Ref. [270]. A detailed list of the experimental conditions under which data was collected is available in Appendix C, and the analysis pipeline is described in Appendix D. The search for new interactions is done by determining the value of the signal strength, α , that provides the best fit between signal templates and the measured forces across multiple dimensions and harmonics.

Signal templates are computed for various values of the Yukawa length scale, λ , using a finite-element model of the geometry of the attractor and microsphere over a spatial grid of the trap volume [271, 272]. An example template in the time domain is shown in Figure 7.4. The path traversed by the attractor during each measurement is recorded and used to construct the signal template for that particular trajectory. Both positive and negative values of α are considered, enabling a search for attractive and repulsive modifications to the Newtonian gravitational potential. The data presented here was collected with a negligible offset between the attractor and microsphere centers of mass in the z direction, maximizing the signal along the x and y directions, at the expense of that in z . Fig. 7.5 shows the spectrum of apparent forces measured at the first 13 harmonics in both x and y for a representative dataset, along with the expected signal for a particular choice of Yukawa interaction parameters.

While some power is present above the noise level at harmonics of the 3 Hz scanning frequency, the relative distribution of the excess power does not match the Yukawa signal template and therefore the experiment found no evidence of a new interaction. To better illustrate some features of the background, Figure 7.6 shows both the magnitude and phase of the effective forces measured for three harmonics with the

largest power above the noise. The magnitudes and phases expected for true attractive and repulsive interactions parametrized by Eq. (3.10) are also shown, allowing for a direct comparison between the measurements and the signal template. Three sources of background have been identified and are discussed in detail in Chapter 8: mechanical vibrations, electromagnetic interactions, and scattering of stray light. All three of them substantially improved with respect to the previously reported results.




Figure 7.5: Force spectra in x and y measured by the microsphere for a single dataset, with the templates for a Yukawa interaction with $\alpha = 10^9$ and $\lambda = 10 \mu\text{m}$ shown for comparison.

7.4 The Statistical Model

Since the measurements reflect features that are not consistent with the hypothesized interactions, the data is used to place constraints on α for a range of λ . Before proceeding, it must be stressed that a rigorous statistical methodology for setting limits in the presence of unmodeled, bipolar backgrounds does not exist. Without a predictive model of the backgrounds, the possibility that they are obscuring a real signal

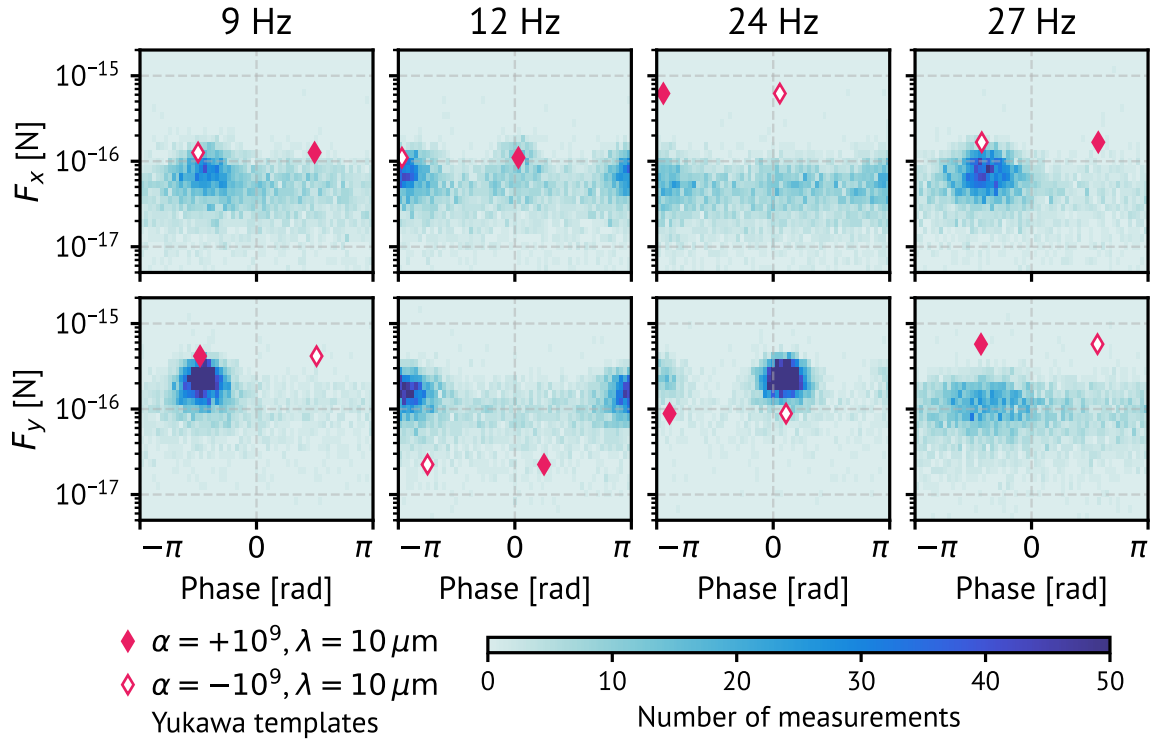


Figure 7.6: Measurements of the x and y components of the force vector at the three harmonics with the largest observed power. The filled (unfilled) diamonds show the expected signal magnitudes and phases for an attractive (repulsive) Yukawa interaction with α and λ shown in the legend. The measurements are not consistent with either signal template across all harmonics.

cannot be ruled out. This is not a far-fetched idea, nor does it require any particular conspiracy among measurements in different harmonics. If nothing but noise were measured, the probability that opposite-sign backgrounds had perfectly canceled a true signal in all harmonics *would* be implausibly low. But given that backgrounds are known to be present, mimicking an attractive force in some harmonics and a repulsive force in others, if a real signal were hiding in the data, the backgrounds would necessarily add to it in some harmonics and subtract from it in others. Depending on which of these dominated, the fitted signal strength would be either larger or smaller than the true value.

If the backgrounds exhibited no preferred polarity, we should expect the latter

scenario roughly half of the time. This would mean that the signal had at least partially been canceled by the backgrounds. The complex spectral fingerprint of a real signal, while a powerful confirmation tool in the absence of backgrounds, is not useful if the signal is overpowered by the backgrounds, even if the signal itself is well above the noise. I point this out to convey that we are forced to do statistics on a problem that is not amenable to standard statistical techniques. Consequently, the strategies below are intended as temporary workarounds rather than comprehensive solutions. The actual solution to this problem is obvious but difficult to implement: either eliminate or model the backgrounds.

7.4.1 Likelihood Function Construction

The statistical analysis is performed using a frequentist framework. Each dataset consists of measurements of the complex amplitude of each Cartesian component of the force at N_h harmonics for each of the N_f 10-second integrations. The complex amplitude is given by the discrete Fourier transform (DFT) over a 10-second integration, scaled by an appropriate normalization factor. For each dataset there is a corresponding signal template, also consisting of a complex amplitude for all components, harmonics, and integrations. This gives a total of $3 \times N_h \times N_f$ measurements and templates. A likelihood function can then be constructed for each, describing the probability of obtaining the measured complex amplitude F_{ijk} given the scaled template $\alpha \tau_{ijk}(\lambda)$, where i , j , and k index the file, harmonic, and force component, respectively. The likelihood is defined as

$$\mathcal{L}_{ijk}(\alpha|\lambda) = \frac{1}{2\pi\sigma_{ijk}^2} \times \exp \left\{ -\frac{[\Re(F_{ijk} - \alpha \tau_{ijk}(\lambda))]^2}{2\sigma_{ijk}^2} \right\} \times \exp \left\{ -\frac{[\Im(F_{ijk} - \alpha \tau_{ijk}(\lambda))]^2}{2\sigma_{ijk}^2} \right\}, \quad (7.6)$$

where σ_{ijk} is the RMS noise as measured with the attractor stationary. One such measurement is done in every tenth integration to account for potential changes in the noise environment over long timescales. The value of σ_{ijk} used in each likelihood

function is the measurement at harmonic j for force component k from the next noise-only file following integration i .

To exploit both amplitude and phase information, the likelihood function includes both the real and imaginary parts of the complex amplitude at each frequency as Gaussian distributions² with means set by the signal template. Assuming statistical independence among measurements, a global likelihood function can be constructed from the product of the individual likelihood functions for each measurement included,

$$\mathcal{L}(\alpha|\lambda) = \prod_i \prod_j \prod_k \mathcal{L}_{ijk}(\alpha|\lambda), \quad (7.7)$$

where the products are over the chosen subsets of files, harmonics, and components selected for the particular statistical test to be conducted. We define $\hat{\alpha}$ as the Maximum Likelihood Estimator (MLE) for α , or the value for which $\mathcal{L}(\alpha|\lambda)$ is maximal. This can be computed for any arbitrary subset of measurements to understand how each subset contributes to the global fit.

7.4.2 Test Statistic Construction

To set a confidence interval (CI) on α for a particular λ , a likelihood ratio test is used [273]. The likelihood ratio is defined as

$$\Lambda = \frac{\mathcal{L}(\alpha|\lambda)}{\mathcal{L}(\hat{\alpha}|\lambda)}. \quad (7.8)$$

In the numerator, α is set to a particular value (which defines the null hypothesis) to be tested for compatibility with the data. In the denominator, α is fixed to its best-fit value, ensuring that $\Lambda \leq 1$, with Λ decreasing toward 0 as α gets further from $\hat{\alpha}$. The likelihood ratio therefore provides a metric for compatibility between the null hypothesis and the data. If the distribution of Λ under the null hypothesis is known, this compatibility can be quantified, and CIs on α can be obtained.

For this purpose, it is convenient to use the quantity $-2 \log \Lambda$, as Wilks' theorem

²This is in contrast to the magnitude and phase, which follow Rayleigh and uniform distributions, respectively.

then gives an analytic expression for its distribution under the null hypothesis [274]. For the likelihood function defined by Eqs. (7.6) and (7.7), this quantity takes a simple parabolic form,

$$-2 \log \Lambda = \frac{(\alpha - \hat{\alpha})^2}{\sigma_{\text{stat}}^2}, \quad (7.9)$$

where $\hat{\alpha}$ is the global MLE for α over all measurements included, and σ_{stat} describes the statistical uncertainty on α . In this case, both $\hat{\alpha}$ and σ_{stat} are calculable directly from the data, with no parameter estimation required (see Appendix E for the calculation). To account for systematic uncertainties, we can define a modified likelihood ratio, Λ' , such that

$$-2 \log \Lambda' = \frac{(\alpha - \hat{\alpha})^2}{\sigma_{\hat{\alpha}}^2}, \quad (7.10)$$

where $\sigma_{\hat{\alpha}}^2 = \sigma_{\text{stat}}^2 + \sigma_{\text{sys}}^2$ is the total variance including both statistical and systematic uncertainties³.

There are multiple ways in which a test statistic can be constructed from Λ' , depending on the allowed range of α and whether the test statistic is to be used for discovery or exclusion. Thorough descriptions of many can be found in Ref. [275]. For the search presented here, we use the exclusion-only test statistic q_{α} , defined as

$$q_{\alpha} = \begin{cases} -2 \log \Lambda', & |\alpha| \geq |\hat{\alpha}|, \\ 0, & |\alpha| < |\hat{\alpha}|, \end{cases} \quad (7.11)$$

which prohibits a discovery claim by construction, as the hypothesis that $|\alpha| < |\hat{\alpha}|$ can never be rejected. To set an exclusion limit on α at the 95% CL, one finds the value of α such that q_{α} takes on its value at the 95th percentile of its distribution

³This is a convenient approximation; a rigorous treatment of the systematic uncertainties would include the relevant terms as a set of nuisance parameters in the likelihood function and profile over them across the range of α values. We do not use this approach for a key reason: the observed backgrounds do not vary predictably or consistently with the parameters identified as sources of systematic uncertainty. This means that rather than finding the true values of the nuisance parameters, the likelihood maximization routine would instead simply converge to the parameter values that best absorb the apparent forces in the fit, without reflecting a real discrepancy between the true and measured values of the physical quantity. This is a failure of the statistical model, and it would not be an issue if a predictive background model were included.

under the null hypothesis. This value can be calculated from the analytic form of the distribution of q_α provided by Wilks' theorem. The critical value of α , which is denoted α_{95} , is given by

$$\alpha_{95} = \hat{\alpha} + \text{sgn}(\hat{\alpha}) \cdot 1.64 \sigma_{\hat{\alpha}}. \quad (7.12)$$

Computing this quantity as a function of λ for the selected measurements gives the desired exclusion curve.

7.4.3 Incorporating Systematic Uncertainties

The systematic uncertainties considered are shown in Fig. 7.7. For each parameter, the signal template is recomputed at both of the $\pm 1\sigma$ values, and the best-fit α value over the given selection of files, harmonics, and components is determined. This is repeated for all values of λ to get the fractional error on α at each. These fractional errors are added in quadrature across all values of λ to produce the total uncertainty curve shown in Fig. 7.7. The computed values of σ_{sys} have the effect in Eq. (7.10) of broadening the Λ' parabolas as a function of λ . At small λ , the systematic uncertainty is dominated by uncertainty in the microsphere's position (predominantly in x), while at large λ , the dominant uncertainty comes from the thickness of the attractor.

I have not yet addressed which of the measurements should be included in the calculation of the test statistic. It may seem sensible to use all of them, though this simply adds noise (or backgrounds) to the data even in channels that shouldn't see much or any signal power to begin with. There may be some harmonics that see so much power as to be obviously in statistical tension with the rest, so including them just worsens the fit. But to leverage a key advantage of this technique — the ability to use a complete spectral fingerprint to differentiate signals from backgrounds — a sufficient number of harmonics must be included. The advantages and disadvantages of multiple strategies are discussed below. Even in the absence of a bulletproof statistical methodology, the competing frameworks can be compared based on a number of criteria to ensure that the chosen method is as responsible as possible.

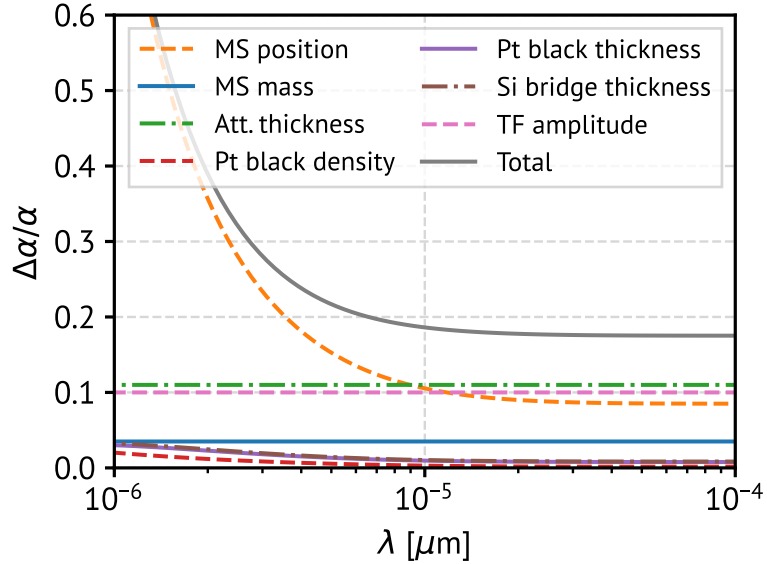


Figure 7.7: Systematic errors included in the constraint set on α . The microsphere position uncertainties in x , y , and z are computed separately, though as the resulting uncertainties in y and z are negligible compared to in x , they are combined in the plot.

7.4.4 Harmonic Selection Criteria

In most force components at most harmonics, the measured backgrounds are either in-phase with or opposite-phase to the attractive Yukawa template. This means that when multiple harmonics are included, the harmonics consistent with $\alpha > 0$ will compete with those consistent with $\alpha < 0$, bringing the global best-fit α value down to a small fraction of the background magnitude in any one component. The exact value of $\hat{\alpha}$ depends on the particular selection of harmonics used. Indeed, switching out one harmonic for another in the fit can even change the sign of $\hat{\alpha}$. This fact underscores the need for a well-justified selection criterion. A sound selection criterion should satisfy the following requirements:

- The reported statistical coverage should be accurate. If a 95% CL upper limit is reported, then a real signal at that level should only be excluded 5% of the time. With unmodeled, time-varying backgrounds in the data, guaranteeing accurate coverage is strictly impossible, though clear bias in favor of undercoverage or

overcoverage should be avoided.

- The harmonics chosen should include the majority of the available signal power. By design, the attractor geometry results in signal power being distributed among many harmonics. This is advantageous due to the signal identification capability it provides, though it means that any one harmonic does not receive much of the total.
- Limits should be placed on both attractive and repulsive interactions. By construction, q_α provides a single-sided limit when computed for any given set of measurements, with the sign of α_{95} given by that of $\hat{\alpha}$. It is trivial to then set conservative limits on $\alpha > 0$ and $\alpha < 0$ separately, by dividing harmonics into those with $\hat{\alpha} > 0$ and those with $\hat{\alpha} < 0$, and computing q_α for each. This naive strategy overcovers significantly, as it eliminates the competition between backgrounds of different signs in different harmonics.

The following three strategies are assessed on their ability to meet these requirements.

Choosing the “Best” Harmonic

In the first analysis reported from this experiment [266], q_α was computed independently for each of the six harmonics considered, and the results were summed together to produce a single test statistic as a function of α . The limit was then set based on this modified test statistic. For most datasets, this is equivalent to including only the harmonic with the smallest value of $\hat{\alpha}$, since that will be the one in which q_α departs from 0 and crosses the critical threshold at the lowest value. For datasets which include bipolar backgrounds, this test statistic construction inherently results in undercoverage. This is because if a true signal were present, the backgrounds would add to the signal in some harmonics and subtract from it in others. The test statistic would then pick out the harmonic with the smallest $|\hat{\alpha}|$, which would preferentially be one of those in which the background had canceled a portion of the signal. Other features of the data, such as the spectral fingerprint or the time-dependence of the

measurements, can be used to make a convincing case that backgrounds are present, or even dominant. But these arguments have no bearing on *where exactly* the limit should be set. Systematic undercoverage like this can be expected to occur whenever the choice of harmonics depends on the measured values of $\hat{\alpha}$ in each one.

Signal Power

In this approach, the signal power in all harmonics is ordered from greatest to least. Harmonics are then included in order until the signal power exceeds some predefined threshold. This is advantageous in that the analysis is not punished by the inclusion of measurements in channels insensitive to a signal. The fraction of signal power to be included is arbitrary. This criterion gives a one-sided limit on α of the same sign as $\hat{\alpha}$ for the harmonics included.

Goodness of Fit

An alternative selection criterion is based on the goodness of fit over some number of harmonics. In this approach, the maximum likelihood is computed for every combination of some specified number of harmonics, and the combination which gives the maximum value for each sign of $\hat{\alpha}$ is used. This can be construed as doing separate statistical tests for $\alpha > 0$ and $\alpha < 0$, where the subset of data included in each is that which looks the most like the signal under test. It therefore has the advantage of constraining both attractive and repulsive interactions. The number of harmonics to include is arbitrary.

7.5 Constraints on Yukawa Interactions

The constraints reported in Ref. [244] and in this chapter were computed from six harmonics chosen by goodness of fit. Limits on α for new attractive and repulsive interactions for λ ranging from 1 μm to 100 μm were computed for each of the datasets considered. The strongest constraints among these are reported in Fig. Fig. 7.8. These new constraints represent an improvement of a factor of ~ 50 over previous

results for length scales of $\lambda > 10 \mu\text{m}$, and $\gtrsim 100$ at $\lambda = 2 \mu\text{m}$. Moreover, they are the first to use reconstruction of multiple spatial components of the force vector as a function of time, hence exploiting the unique signature of a new force and providing a robust pathway for discovery. For example, a z offset in either direction (c.f. the negligible offset used to collect data analyzed in this work) could be used to verify that a signal with the appropriate signature is transferred to this degree of freedom in case of a discovery.




Figure 7.8: Constraints on the $\alpha - \lambda$ parameter space obtained in this work and the prior 2021 result [266]. The yellow shaded region indicates the region of parameter space excluded by other experiments [157, 164, 161, 156].

The sensitivity of this search was limited by backgrounds primarily due to stray laser light scattering off the attractor, creating modulations of the optical power on the QPD at harmonics of the attractor drive frequency. While substantial progress has been made in reducing this and other backgrounds, potential avenues for further improvement have been identified. Work is underway to implement a low-noise, high dynamic range and frame rate 100-pixel sensor, to be used for offline discrimination between true microsphere motion and scattered light effects. An alternative attractor geometry, consisting of a $30 \mu\text{m}$ thick, 1.5 mm diameter rotating disk with azimuthal density modulations [276], has also been prepared for potential use in the

next experimental run. Better free space optics upstream of the trap is expected to produce more stable conditions and backgrounds, specifically addressing slow drifts of the trap focus which appear correlated with small variations in the refractive index of air driven by atmospheric pressure fluctuations. Beyond these immediate steps, a completely enclosed attractor should prevent any direct coupling between its motion and the beam, vastly attenuating optical backgrounds. Design of such an attractor system is a technical challenge to which future engineering efforts will be dedicated.

7.6 Conclusions

In this chapter, I have described the use of optically levitated microspheres as force sensors capable of probing new Yukawa interactions that couple to mass at micron-scale separations. In addition to complementing existing techniques, this platform uses the fully-reconstructed force vector to search for the complex temporal signature that would arise from such interactions, a feature that may become crucial should future work result in a discovery. Techniques to probe the main sources of background limiting the experiment, as well as mitigation strategies, have been identified, and are discussed in detail in Chapter 8. Apart from the direct physics goal of the experiment, this work may be relevant to other uses of levitated optomechanics for fundamental physics [171], including searches for dark matter candidates [277, 278], sterile neutrino emission in weak decays [279], or tests of the neutrality of matter [280]. In addition, this work demonstrates the operation of a setup with microscopic objects levitated a few micrometers away from mechanical structures, while maintaining stable conditions, exploring backgrounds, and performing precision metrology. These are among the challenges that will need to be addressed towards experimental tests of the quantum nature of gravity [173, 172, 281].

Chapter 8

Background Management Strategies for the Micron-Scale Interaction Search

The results presented in the previous chapter, while a significant improvement over the first search for micron-scale interactions with levitated microspheres, were still limited by backgrounds. Techniques for understanding and mitigating these backgrounds will be the focus of this chapter.

8.1 Techniques for Background Subtraction

Before discussing the specific backgrounds present in the search for micron-scale interactions, I will introduce some high-level strategies for background subtraction. I will then discuss particular features of the observed backgrounds with emphasis on which of the subtraction strategies they are most and least amenable to. To frame the general problem, I assume that:

- The experimenter uses a number of signal channels to make a set of measurements. The measurement in each channel is assumed to consist of signal and background in some unknown proportions.

- The backgrounds depend on a set of physical parameters. Some of these parameters may be measurable, or under the experimenter's control, while others may not. The background contribution to the measurement in each signal channel is taken to be a function of the background parameters.
- There are also background channels in which an isolated measurement of the backgrounds can be made simultaneous with regular data collection. The signal couples negligibly into the background channels.
- There is some relationship that defines how the backgrounds appear in the background channels versus the signal channels. If a transformation describing this relationship can be determined, the background channels can be used to infer the background contributions to measurements in the signal channels.

Given these assumptions, the experimenter's challenge is to isolate the signal's contribution to the total measurement by subtracting the background. For the background contribution to be subtracted from a measurement, it must be quantitatively understood. Below are three general approaches to this problem.

8.1.1 Physics-Based Model

If the origin of the backgrounds is well understood, then a background model can be constructed to predict the backgrounds measured in a particular experimental configuration. The model may be based on first principles, simulations, fits to experimental data, or some combination of these. The clear advantage to this approach is the ability to make generalizable predictions. If the backgrounds are known to depend on a finite set of parameters, then a background model should contain a functional dependence on each of these parameters, ensuring the flexibility to use it in more than just a single configuration of the experiment. A basic requirement for an accurate background model is the stability of the backgrounds with respect to any of the model parameters. If the predicted background depends on an experimental parameter more sensitively than that parameter can be measured, the background model has little utility in predicting the backgrounds accurately.

8.1.2 Asynchronous Measurement

If the backgrounds cannot be modeled, perhaps because they exhibit an unknown or nontrivial dependence on the experimental parameters, they may still be measurable. This must be done with the signal “turned off” so that the measurements in the signal channels are direct measurements of the backgrounds. Repeating the measurement with the signal “turned on” then allows for the measured backgrounds to be subtracted. Identifying an experimental configuration in which the signal is eliminated while the backgrounds remain unchanged is often impossible.

8.1.3 In Situ Witness

If a model or direct measurement of the background in the signal channels is not possible, one may still be able to construct a separate data stream which is sensitive to the backgrounds but not the signal. This will enable a moment-by-moment measurement of the backgrounds in this second set of channels. If the transformation relating the background channels to the signal channels is known, it can be used to compute the background contributions to the signal channels, allowing them to be subtracted. This method is advantageous in that it does not require an understanding of the exact mechanism causing the background. Construction of the background-only data streams can be done in multiple ways, such as by finding signal-free linear combinations of sensors (discussed in Section 8.4.1) or by modulating a parameter to create signal-free sidebands of the modulation frequency (discussed in Section 8.5.3).

8.2 General Features of the Backgrounds

Having developed the calculation of $\hat{\alpha}$ for a collection of datasets in Chapter 7, it can now be used to understand the backgrounds in terms of their direct impact on the physics sensitivity. By computing $\hat{\alpha}$ for each force component and harmonic in consecutive subsets of files throughout a measurement, the time evolution of the backgrounds can be determined. This reveals another key feature of the backgrounds: unlike an expected signal, they vary in time, and the time variation of separate

harmonics appears independent, as is visible in Fig. 8.1. This makes modeling the backgrounds as a function of experimental parameters difficult, since each harmonic appears to have a different time dependence. More difficult still, none of the measured environmental parameters show any clear correlations with any of the background measurements.




Figure 8.1: Time evolution of the fitted signal strengths in a number of harmonics for both the x and y force components. Different harmonics appear to evolve independently, with no clear correlations with one another or with any other measured parameters.

Three classes of backgrounds have been identified and will be discussed in turn: mechanical vibrations, stray light scattering off the attractor, and electromagnetic coupling to the dipole moment of the microsphere. There is strong evidence that the second is currently the dominant source; this will be presented and discussed.

8.3 Mechanical Vibrations

Mechanical coupling of the oscillating attractor stage to any optics in the vacuum chamber could cause angular jitter of the trapping beam. This is not easily modeled, but it can be measured. By oscillating the attractor at its nominal frequency and amplitude, but retracted along x away from the trap focus, the signal is removed while the backgrounds remain. Provided it is sufficiently far away, no optical or electromagnetic coupling between the microsphere or trapping beam and the attractor is possible. Data has been collected in this configuration on multiple occasions, with the attractor typically ~ 4 mm from the trap focus. The spectra of apparent forces are shown in Fig. 8.2 for one such dataset taken after the intensity stabilization substantially improved the force sensitivity.

In the dataset shown, the backgrounds appear out to ~ 30 Hz predominantly in x , but they are not in phase with the expected Yukawa interaction. It is possible that in this particular dataset, the attractor stage was retracted by too much, causing it to interfere with cables behind the stack, and that in a typical dataset these backgrounds would not be present. Nonetheless, mechanical vibrations are a subdominant source of backgrounds in the current configuration.

8.4 Scattering of Stray Laser Light

Backgrounds are observed when the attractor oscillates about its nominal position even without a microsphere in the trap, indicating that they do not originate from real forces on the microsphere. While this limits the range of possible sources, they can further be identified as originating from stray laser light scattering off the attractor based on three features: the spatial origin, the dependence on the laser power incident on the attractor, and the dependence on the reflectivity of the attractor.

Spatial Origin

The QPD is used to image the region of the trap surrounding the focus. While the central region of this image contains light transmitted by the microsphere, the




Figure 8.2: Spectra showing the backgrounds caused by mechanical vibrations coupled through the attractor stage and into the trap beam. The top two panels show the force spectra both with and without the attractor oscillating. The bottom panel shows the phase of the measured forces for the case in which the attractor was oscillating. The phase is measured relative to the signal template, indicating that the backgrounds do not appear with the phase expected for a signal.

outer regions could include light originating from other sources. In particular, the top surfaces of the attractor and shield wall could be imaged incidentally due to their proximity to the trap focus. As the attractor oscillates at 3 Hz during a measurement, light reflected from its surface would be modulated at this frequency and higher harmonics to produce the backgrounds. To reduce these backgrounds, a spatial filter can be used to select only light sourced by the microsphere’s motion. The first version of such a spatial filtering system consisted of a circular Aektar-coated [282] aperture placed at the focus of the output telescope to occlude the outer regions of the beam. Apertures of 50 μm and 100 μm were tested, with the 50 μm aperture demonstrating

the most significant reduction of backgrounds. More sophisticated methods of spatial filtering could include the use of a spatial light modulator (SLM) to imprint a mask on the reference beam to selectively interfere with the forward-scattered trapping beam, or a multi-pixel sensor to allow for spatial selection to be done at a post-processing step. Another option is to use the heterodyne phase in addition to the amplitude for position reconstruction, allowing for light originating from the microsphere to be selected based on its optical path length, though this requires further experimental validation.

Dependence on Incident Power

The amount of light scattering from the attractor evidently scales with the amount of light incident on it. This in turn depends on the intensity of laser light at the location of the reflective surfaces. Due to the laser's Gaussian intensity profile, retracting the attractor from the beam center should decrease the amount of light scattered from the attractor surface. While this trend is expected to hold over multi-wavelength separations, it is possible that different behavior could be observed on wavelength or sub-wavelength distances. The shield blocks any direct path by which upward-propagating laser light could reflect off a surface of the attractor, but more complicated scattering paths are still possible. One such possibility is that light diffracts around the top of the shield wall and reflects off either the front face or top surface of the attractor. If this were the mechanism responsible for scattered light backgrounds, then interference fringes may be visible as the shield-attractor separation is increased. Regardless of whether this fringing is observed, the overall trend in decreasing backgrounds with increasing shield-attractor separation should hold if the backgrounds originate from scattered light.

Fig. 8.3 shows the average prominence of the background peaks over eleven harmonics with no aperture and with each of the apertures considered. As expected, the backgrounds fall off sharply with distance as the attractor cuts into less of the trap beam. At close separations, the aperture makes a significant difference and reduces the backgrounds by up to a factor of ~ 10 . This data was collected without a microsphere in the trap, providing additional evidence that backgrounds are present which

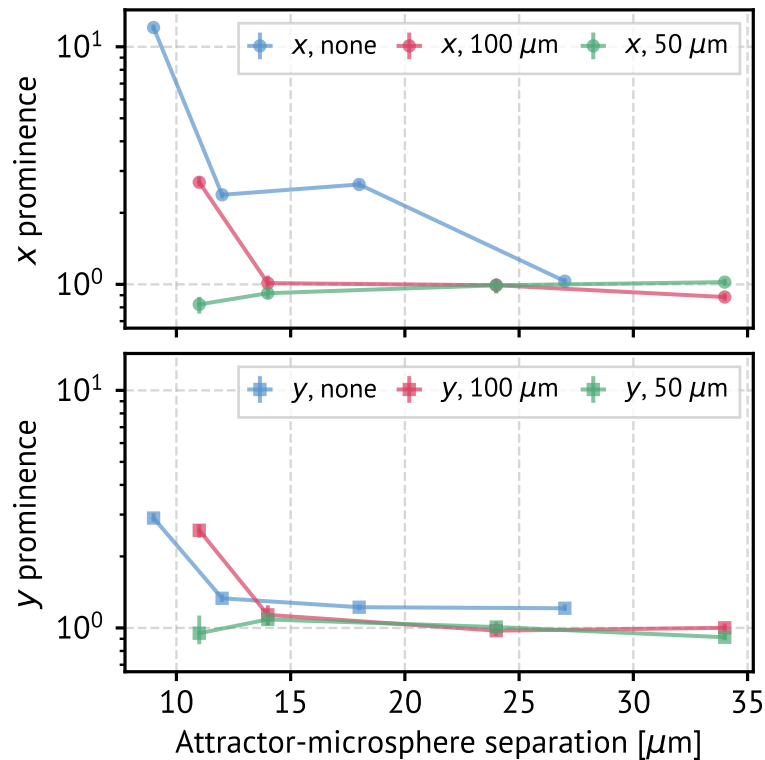


Figure 8.3: Effect of different apertures and the microsphere-attraction separation on the backgrounds. The background prominence is averaged over the fourth through the fourteenth harmonic, as the noise is too high for a meaningful measurement at the first three. For the first three points in each series, the shield was moved with the attractor to maintain a 1 μm separation, while for the last point, the shield was kept in place while the attractor was retracted. The data was collected without a microsphere in the trap, prior to the blackening of the attractor described in Section 8.4.

are not caused by real forces.

Given that the scattered light backgrounds decrease with separation, one might wonder how the scaling with distance compares to that of the expected signal strength. If the two scalings are sufficiently different, there may be an optimal separation at which the backgrounds are minimal compared to a signal. To determine whether this is the case, the distribution of $|\hat{\alpha}|$ was computed for all harmonics and force components at a series of separations, this time with a microsphere in the trap. The absolute value allows backgrounds that appear both attractive and repulsive to be included. Fig. 8.4 shows this data, collected with the 50 μm aperture installed.




Figure 8.4: The range of fitted signal strengths at a series of separations between the attractor and the microsphere. The horizontal markers indicate the median while the vertical bars show the $\pm 1\sigma$ ranges of the computed $\hat{\alpha}$ values over all harmonics and force components.

The spread of $|\hat{\alpha}|$ values at each separation is significant, covering nearly two orders of magnitude. There is a slight increase in the median $|\hat{\alpha}|$ with distance, suggesting that the backgrounds fall off more slowly than the signal strength and that a closer approach is preferred. The combination of a $50 \mu\text{m}$ aperture and the closest possible approach therefore maximizes the ratio of signal strength to backgrounds.

Dependence on Reflectivity

Reducing the reflectivity of the attractor should also reduce the reflected power which makes up the backgrounds. Coating the attractor with a $\sim 3 \mu\text{m}$ layer of Platinum Black, which has a reflectivity $\lesssim 1\%$ that of the underlying gold, did indeed reduce the backgrounds significantly. The coating procedure and resulting reflectivity are described in detail in Ref. [267].

With a microsphere in the trap, the scattered light backgrounds are visible at harmonics of the attractor drive extending beyond 100 Hz, with no drop-off in amplitude. This suggests that they are produced through reflections off structures with a high spatial frequency, potentially including rough surfaces or sharp edges on the attractor. This is consistent with the time variation, which would only be significant in the event that small drifts of experimental parameters (laser power, position of the trap focus, position of the attractor, etc.) caused large and unpredictable changes in the background spectrum. Additionally, considerable effort spent using ray tracing and finite difference time domain (FDTD) software failed to yield any numerical predictions that matched the data, further suggesting that modeling or simulating the backgrounds is untenable. As the presence of the sphere in the trap enhances scattering and alters the light environment in the trap, a measurement of the background spectrum under the same conditions — yet without a signal present — is also not possible. The existence of the backgrounds is fundamentally unseparable from having a sphere in the trap in proximity to the attractor, the same configuration under which a potential signal would be present. For this reason, we turn instead to constructing a background witness data stream.

8.4.1 Constructing a Scattered Light Witness Stream

The x (y) motion of the microsphere is naively calculated by taking the difference between power incident on the left and right (top and bottom) halves of the QPD and dividing through by the total light detected. If the four quadrants are numbered one through four¹ and each quadrant i measures an amplitude q_i , then the x and y

¹By convention, we number the quadrants in the order top right, bottom right, top left, bottom left, as shown later in Fig. 8.6.

positions can simply be calculated using

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \end{pmatrix} \begin{pmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{pmatrix} \quad (8.1)$$

up to some force calibration constants. This assumes that microsphere motion is captured through deflections of the beam centroid on the QPD, and that the x and y axes at the trap focus are aligned with the axes of the QPD. This is a vast oversimplification which ignores the complexities of Mie scattering and the subtle effects stemming from the unique optical properties of each individual microsphere. A data-driven, model-agnostic position detection scheme is preferable to this.

Calibrating to Microsphere Eigenmodes

Stepping back, the task we are faced with is using four sensors (the four QPD quadrants) to measure two physical quantities (x and y). We don't need to know anything about how the microsphere motion interacts with the laser light to produce its signature on the QPD, but we do assume that the coupling is linear. To construct the measurement streams for x and y from the four quadrants, we first require some calibration data representing known motion. The x and y resonances provide such motion, as these both have resonant frequencies at which the motion should be predominantly along the given axis. In practice, feedback is required to keep the microsphere in the trap at high vacuum, so the motion at the resonance includes motion caused by the feedback signal. Instead, we can use data collected with the trap at 1 mbar to 2 mbar with the feedback off, leaving the resonances broadened by gas damping but otherwise undistorted.

Fig. 8.5 shows an example dataset collected under these conditions, with the amplitude spectral density (ASD) of each quadrant plotted independently in the top panel. The two resonances are visible in all four quadrants. The bottom panel shows the measured phase in each quadrant, plotted relative to the phase of the first




Figure 8.5: Amplitude and phase of the responses of each of the four quadrant photodiode sensors to free motion of the microsphere. The resonant frequency in y is lower than that in x . The clear phase relationships between quadrants at both resonances allows for the responses to x and y motion to be determined.

quadrant. A clear phase relationship between quadrants is visible for each of the two resonances: at the y resonance, quadrant 1 is in phase with quadrant 3, while at the x resonance, it is in phase with quadrant 2, as expected.

This calibration data can be used to determine the calibrated position-sensing matrix \mathbf{C}_{xy} which reconstructs the microsphere motion in x and y from the sensor responses. Take the amplitude of x motion at the x resonance and y motion at the y resonance to be 1. Scaling of these quantities into physically meaningful units will be done at a later stage. Assume that the resonances overlap such that the y motion at the x resonant frequency is ϵ_{yx} and the x motion at the y resonant frequency is ϵ_{xy} , with $\epsilon_{yx}, \epsilon_{xy} \ll 1$. Define the complex amplitude measured by quadrant i at the resonant frequency for axis j to be q_{ij} . Then we know that the transformation \mathbf{C}_{xy}

must satisfy

$$\begin{pmatrix} 1 \\ \epsilon_{yx} \end{pmatrix} = \mathbf{C}_{xy} \begin{pmatrix} q_{1x} \\ q_{2x} \\ q_{3x} \\ q_{4x} \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} \epsilon_{yx} \\ 1 \end{pmatrix} = \mathbf{C}_{xy} \begin{pmatrix} q_{1y} \\ q_{2y} \\ q_{3y} \\ q_{4y} \end{pmatrix}. \quad (8.2)$$

We can combine the column vectors into matrices \mathbf{M} and \mathbf{Q} , with

$$\mathbf{M} = \begin{pmatrix} 1 & \epsilon_{xy} \\ \epsilon_{yx} & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{Q} = \begin{pmatrix} q_{1x} & q_{1y} \\ q_{2x} & q_{2y} \\ q_{3x} & q_{3y} \\ q_{4x} & q_{4y} \end{pmatrix}, \quad (8.3)$$

allowing \mathbf{C}_{xy} to be defined by the product

$$\mathbf{M} = \mathbf{C}_{xy} \mathbf{Q}. \quad (8.4)$$

The system described by this equation is underdetermined and infinitely many solutions exist for \mathbf{C}_{xy} . This is because the elements of \mathbf{C}_{xy} can be arbitrarily large while still delicately canceling one another when summed to satisfy Eq. (8.4). To prevent the transformation from amplifying measurement noise, we use the minimum-norm solution for \mathbf{C}_{xy} . The minimum-norm solution is unique; it is given by

$$\mathbf{C}_{xy} = \mathbf{M} \mathbf{Q}^+, \quad (8.5)$$

where \mathbf{Q}^+ is the pseudoinverse of \mathbf{Q} . The pseudoinverse is a generalization of the matrix inverse to matrices that need not be square. One can show that the pseudoinverse can be calculated using

$$\mathbf{Q}^+ = (\mathbf{Q}^\top \mathbf{Q})^{-1} \mathbf{Q}^\top, \quad (8.6)$$

and the full calibration matrix then becomes

$$\mathbf{C}_{xy} = \mathbf{M} (\mathbf{Q}^\top \mathbf{Q})^{-1} \mathbf{Q}^\top. \quad (8.7)$$

The transformation \mathbf{C}_{xy} takes the length-4 vector containing the complex amplitudes of each quadrant, \mathbf{q} , to the length-2 position vector, \mathbf{x} :

$$\mathbf{x} = \mathbf{C}_{xy}\mathbf{q}. \quad (8.8)$$

\mathbf{C}_{xy} is a 2×4 matrix which contains the empirically-determined coefficients by which each quadrant amplitude should be multiplied to optimally reconstruct the microsphere's position.

So far, what I have described is a method to better calibrate the QPD data streams to the motion of the microsphere. But with this comes a method to construct a witness data stream for scattered light backgrounds. As only two linear combinations of the sensor amplitudes are required to reconstruct the microsphere's position, there are two remaining linearly-independent combinations. One of these is the sum of all quadrants. In practice, at each time step every quadrant amplitude is normalized by the sum to reduce noise from intensity fluctuations, so this combination remains fixed at 1. This leaves three remaining degrees of freedom: x , y , and a linear combination orthogonal to the $x-y$ plane. Unlike motion of the microsphere, stray light scattering off the attractor will in general not be confined to the motional plane, and should therefore be visible in the orthogonal linear combination. I refer to this combination as the “null stream” since it represents the null space of \mathbf{C}_{xy} and should be insensitive to motion. The null stream, \mathbf{n} , and the total light detected, \mathbf{t} , can be added to \mathbf{C}_{xy} to make the full 4×4 calibration matrix, \mathbf{C} :

$$\mathbf{C} = \begin{pmatrix} \mathbf{C}_{xy} \\ \mathbf{n}^\top \\ \mathbf{t}^\top \end{pmatrix}_{4 \times 4} \quad \text{with} \quad \mathbf{C}_{xy}\mathbf{n} = 0, \quad \mathbf{n}^\top\mathbf{t} = 0. \quad (8.9)$$

As a point of comparison, it helps to determine what this matrix would look like under the naive position calibration scheme. The first two rows are given by Eq. (8.2). By inspection, one can find that the two remaining linearly-independent combinations

to be added are

$$\begin{pmatrix} \mathbf{n}^\top \\ \mathbf{t}^\top \end{pmatrix} = \begin{pmatrix} 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix}. \quad (8.10)$$

The null vector corresponds to opposite diagonals being excited in phase, as would occur if the beam were stretched along one diagonal and compressed along the other. Combining these with Eq. (8.2) allows us to define the “naive matrix” \mathbf{N} , a useful reference point when interpreting the data-driven calibration matrices:

$$\mathbf{N} \equiv \begin{pmatrix} 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix}. \quad (8.11)$$

To facilitate direct comparisons between the calibrated and uncalibrated matrices, the rows of \mathbf{C} can be scaled such that they all have norm 2. If \mathbf{N} is a relatively good approximation of the optimal calibration matrix, then the elements of \mathbf{C} should be $\approx \pm 1$. Having explained the mathematical underpinnings of this method, I will now describe its implementation in a realistic calibration scheme.

Implementation of the Calibration Scheme

As mentioned, the calibration procedure starts with data collected in the range of 1 mbar to 2 mbar with feedback off so the resonances are visible and undistorted. These resonances can then be fit with Lorentzian functions so that the resonant frequencies can be precisely determined. The cross-coupling used to populate the off-diagonal elements of matrix \mathbf{M} can also be estimated from these fits. The matrix \mathbf{Q} is then constructed using the real part of the complex amplitudes of all four quadrants at each of the resonant frequencies, as defined in Eq. (8.3). Since the quadrants are predominantly either in phase with or opposite phase to one another at the resonances, throwing away the imaginary part should not result in a significant fraction of power being discarded, though this fraction is calculated each time as a basic data quality check.

\mathbf{C}_{xy} is then calculated from \mathbf{Q} and \mathbf{M} using Eq. (8.7). A row of ones is then added to this 2×4 matrix to enforce that the null vector is orthogonal to both the $x-y$ plane and the sum of the quadrants, and the null space of the resulting 3×4 matrix (defined by a single null vector) is computed using singular value decomposition (SVD). The complete calibration matrix is then formed by stacking \mathbf{C}_{xy} , the transposed null vector, and the row of ones². The scaling to ensure each row has a norm of 2 is then applied.

Two additional checks are then performed. First, the angle between the x and y rows of \mathbf{C} is computed. This angle, in sensor space, *not* physical space, gives an indication of how well decoupled the x and y motion can be. Finally, to ensure that the calibration data contain only two motional degrees of freedom as expected, an expanded, 4×4 version of the matrix \mathbf{Q} is constructed, with the two additional columns consisting of complex amplitudes from two additional frequencies. The singular values of this 4×4 matrix are then computed to ensure that it is approximately rank 2. Across multiple microspheres, the first two singular values tend to be $\sim 10^{-4}$, while the next largest is $\sim 10^{-8}$ to 10^{-7} . The smallest is zero to within the limits of numerical precision, as is to be expected given the removal of one degree of freedom by normalizing all quadrant amplitudes by the sum over all quadrants. The fact that the third singular value tends to be $\sim 1000\times$ smaller than the other two confirms that microsphere motion is limited to two physical modes and that other motion (spinning, breathing modes, etc.) do not show up in the position data.

The calibrated matrix is used to compute the calibrated position and null data streams for all data collected with that microsphere. This includes the data used to fit the transfer functions and to compute the Wiener filters. As the null stream does not contain measurements of a real force, it is impossible to assign it meaningful physical units. In the procedure outlined above for determining the calibration matrix, the rows are scaled such that they all have norm 2, and as a result the raw data measured in the null stream will have the same units – scaled ADC counts — as in x or y . For

²A row of ones is evidently redundant when the quadrant amplitudes are normalized by their sum at each timestep, as the corresponding data stream will always remain at 1. However, it is included nonetheless to make the transformation matrix explicitly full rank, even when the data include only three degrees of freedom.

ease of interpretation, the null stream is then scaled to force units using the same calibration factor as for x . This enables direct comparisons between the measurements in the null stream and forces in x , and approximate comparisons to forces in y .

Performance on Real Data

The plots in this section demonstrate the performance of this calibration technique using the calibration data shown in Fig. 8.5. A visual comparison of the calibrated position sensing matrix \mathbf{C} to the naive matrix \mathbf{N} is shown in Fig. 8.6. In each of the visualized matrices, the four numbered regions show the layout of the QPD quadrants, while the colored areas in each quadrant show the weight that the quadrant is assigned in the calculation of each quantity. Hatching indicates that the weight is negative. In the naive matrix (left), the magnitudes of the weights in the calculation of each quantity are equal for all quadrants, and the areas are the same. The calibrated matrix (right) has unequal areas, representing the fact that some quantities rely more on one quadrant than another for optimal reconstruction.

Both the naive and the calibrated position-sensing matrices are applied to reconstruct x , y , and the null stream in the calibration dataset of free resonant motion. The results are shown in Fig. 8.7. Using the naive matrix, the null stream contains significant (20%) cross-coupled resonant motion. This is eliminated when the calibrated matrix is used.

Another way to check the quality of the calibration is to quantify the degree of cross coupling between x , y , and the null stream in the transfer function calibration data. This data is collected with the microsphere charged and driven with a frequency comb (1 Hz spacing from 1 Hz to 700 Hz) in either x or y . When driven along either axis, for well-calibrated position sensing, the response along the other axis and in the null stream should be minimal. Fig. 8.8 shows this data using both the naive and the calibrated position sensing matrices. With proper calibration, the response in x to a driving force in y drops from $\sim 30\%$ to $\sim 3\%$. In the null stream, the $\sim 20\%$ cross coupling in both axes drops to the noise floor in x and $\sim 4\%$ in y . While the null stream is not perfectly motion free, it is serviceable as a background witness; if the null stream measures an apparent force more than $\sim 5\%$ of that seen in y , then the

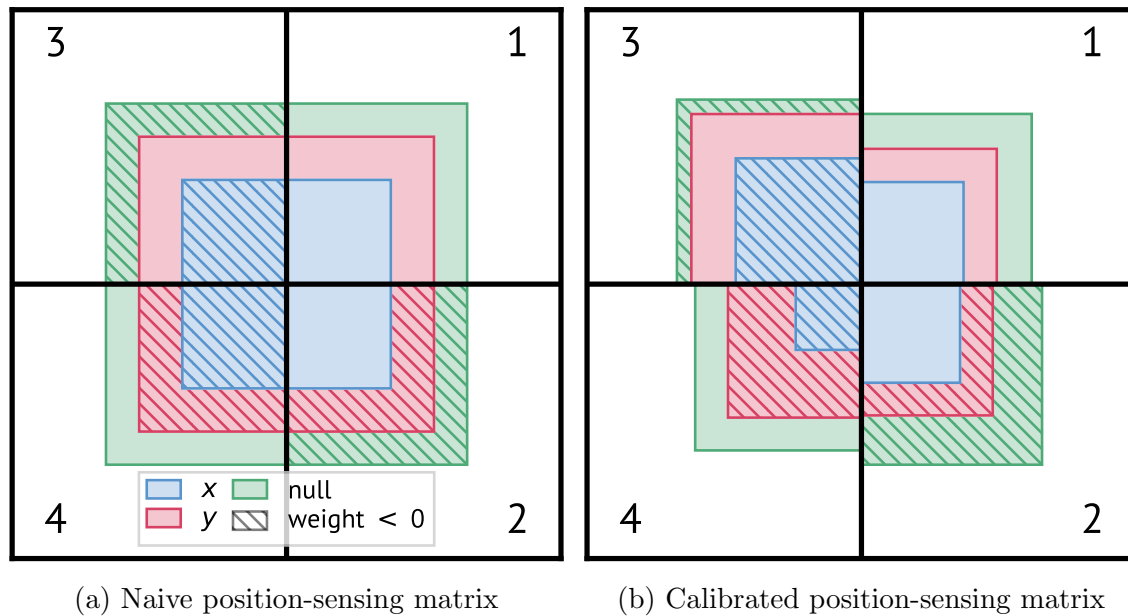


Figure 8.6: A visualization of the position-sensing matrices before and after calibration. The visible area in each color indicates the weight assigned to that quadrant for the data stream indicated. In the example shown here, quadrants 3 and 2 have disproportionate sensitivity to motion in x and the null stream, respectively, when compared to the naive matrix.

background in y must include scattered light in addition to any motion.

Finally, after having thoroughly checked the quality of the calibration, the null stream can be used as further verification of the hypothesis that the backgrounds originate predominantly from scattered light. Fig. 8.9 shows the measured force spectra in x , y , and the null stream for a physics dataset, confirming that the null stream sees apparent forces of comparable magnitude to the backgrounds in x and y .

Towards Background Subtraction

While the null stream is insensitive to motion, the calibrated motional data streams remain sensitive to scattered light. Indeed, backgrounds are visible in the motional data streams even when the proper calibration procedure has been followed. With no knowledge of the behavior of the backgrounds, they can be treated as random




Figure 8.7: The microsphere resonances as detected in the x , y , and null data streams before and after calibration. Without the proper calibration, the null stream sees significant microsphere motion; after calibration, it is insensitive to motion in x and y .

vectors in the four-dimensional sensor space. The projection of these vectors onto the motional plane describes the coupling of backgrounds into the signal channels, while the projection to the orthogonal axis describes their appearance in the null stream.

In the most general case, the background witness channel is useful as evidence that at least some of the backgrounds originate from scattered light, and nothing more. This is because the measurement of a vector along one axis says nothing about its projection onto orthogonal axes. With this information alone, there is no way to use the null streams to predict the backgrounds in the motional data streams — a required step in background subtraction. If the backgrounds display any consistent behavior, however, it may be possible to extract some relationship between the backgrounds as measured in the null streams and as measured in x and y . At minimum, one might expect there to be a consistent phase relationship between the two data streams. Preliminary investigations have shown that the transfer function relating the coupling of scattered light between the null stream and x and y is not stable, but with future improvements to the setup more consistent behavior may emerge.

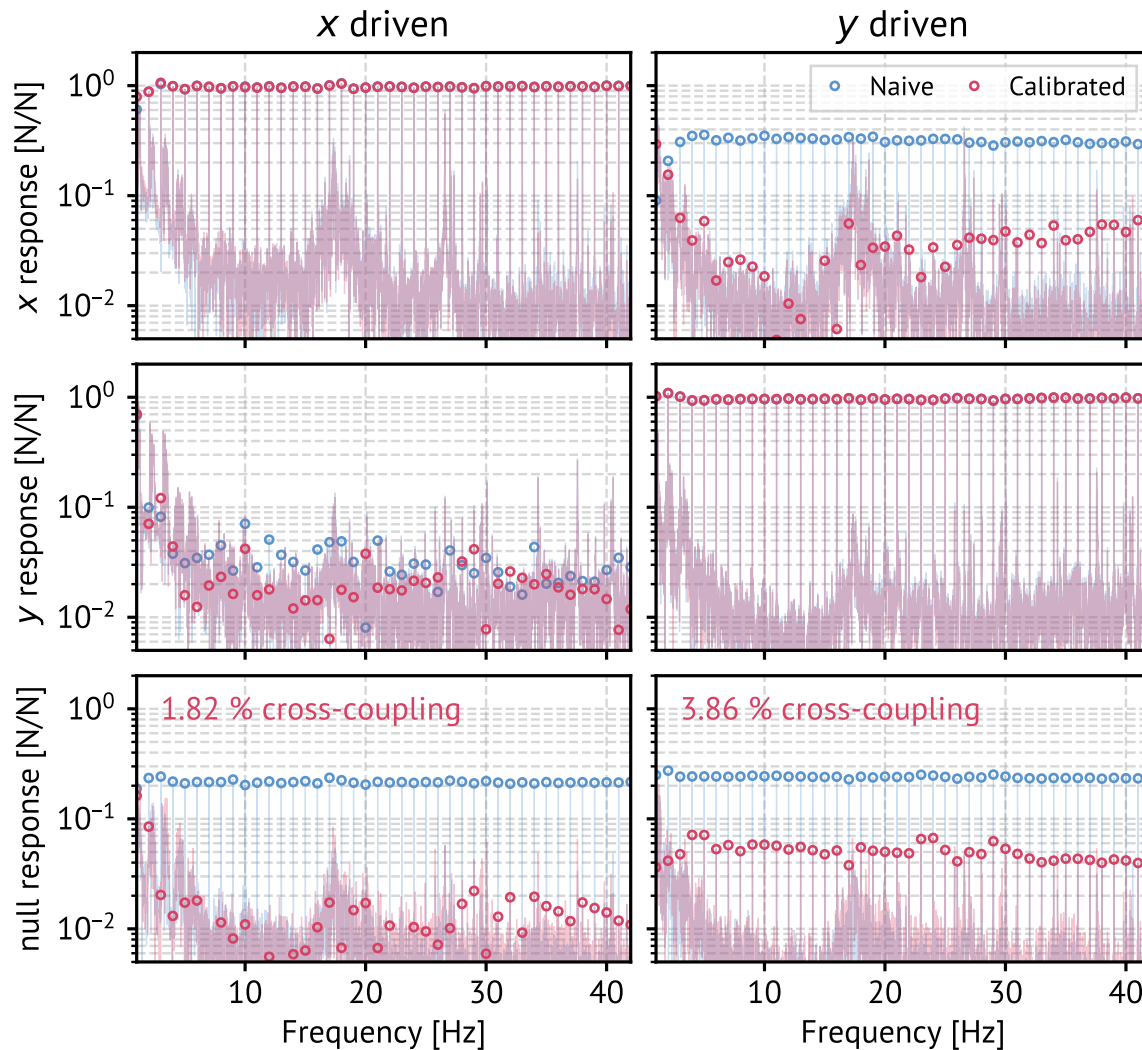


Figure 8.8: The responses to driven microsphere motion in the x , y , and null stream before and after the position-sensing calibration is applied. The columns correspond to frequency combs being driven in either x or y , while the rows show the responses. With calibration, the off-diagonal and null responses decrease significantly.

Scaling to Larger Sensors

The calibration technique demonstrated here has shown some utility in finding the sensor combinations that optimally capture microsphere motion for a four-pixel sensor. As the use of sensors with more pixels is explored, extending this method to more dimensions will yield even greater benefits.




Figure 8.9: Backgrounds observed in a physics dataset in the x , y , and null channels. The apparent force measured in the null stream is of roughly the same amplitude as those measured in x and y , indicating that scattered light backgrounds are responsible. However, as the spectral fingerprint of the backgrounds in the null stream is not predictive of the spectral fingerprint in the motional data streams, this spectrum cannot be used for background subtraction.

There are many potential strategies that could be used to identify the signatures of microsphere motion on a multi-pixel sensor, such as selective masking or blob detection. These can synthesize intuitive ideas about how microsphere motion should appear (e.g. it should show up in clusters of spatially-grouped pixels rather than be distributed randomly across the sensor) with empirical findings (e.g. there are four distinct blobs in most images). Unlike these, the approach described here is purely empirical and does not rely on any preconceptions of what motion in an image should look like. It is easily generalizable to an arbitrary number of pixels. For a sensor with

N_p pixels, the $x - y$ motion calibration matrix will be given by

$$(\mathbf{C}_{xy})_{2 \times N_p} = \mathbf{M}_{2 \times 2} \left(\mathbf{Q}_{2 \times N_p}^T \mathbf{Q}_{N_p \times 2} \right)^{-1} \mathbf{Q}_{2 \times N_p}^T \quad (8.12)$$

requiring the inversion of a 2×2 matrix as before. The rows of $(\mathbf{C}_{xy})_{2 \times N_p}$ contain weights for each of the N_p pixels in the sensor. Plotting these would enable a visualization of which regions of the sensor are most sensitive to microsphere motion. This can then inform other strategies that make use of grouping pixels, pre-thresholding, or other methods of identifying regions containing motional information on larger than single-pixel scales. With N_p pixels, there will be $N_p - 3$ null streams in which to witness backgrounds, all of which could also be used to identify light-sensitive but motion-insensitive regions of the image.

In addition to complementing other techniques, the effectiveness of this strategy should improve with increasing N_p . Treating the background as a random vector in the sensor space once again, we can determine its expected projection into the two-dimensional motion plane to determine the coupling of backgrounds into x and y . The projection scales like the square root of the number of pixels, so by replacing the QPD with a 100-pixel sensor, the amplitude of backgrounds in the x and y channels is expected to decrease by a factor of $\sqrt{100/4} = 5$. We can expect this level of improvement without doing any background subtraction or using any sophisticated image analysis algorithms.

Improvements to the Implementation

The position-sensing calibration matrix is currently computed from data collected immediately after a microsphere has been trapped, and then applied to all data collected with that microsphere. More regular calibrations should be performed to monitor potential drifts in the calibration matrix over time. In addition, the FPGA code which implements the feedback should be modified to use a user-specified calibration matrix rather than the hard-coded naive matrix, ensuring that the feedback is not introducing cross-coupling between x and y .

8.5 Coupling to the Electric Dipole Moment

Despite being discharged to neutrality, silica microspheres can have significant electric dipole moments (EDMs). In an electric field \vec{E} , the force on a dipole with an EDM of \vec{p} at position \vec{r} is given by

$$\vec{F}(\vec{r}) = (\vec{p} \cdot \nabla) \vec{E}(\vec{r}). \quad (8.13)$$

Changes in the electric field evidently cause changes in the force on the dipole. A trapped microsphere with a non-negligible EDM will therefore be sensitive to changes in the electric field at the trap focus. If this electric field depends on the position of the attractor, then the force will be modulated at the harmonics of the attractor drive, causing backgrounds. This section describes a technique for measuring and subtracting these backgrounds.

8.5.1 Electric Dipole Moment of Silica Microspheres

Silica microspheres can have both a permanent and an induced EDM. In a sufficiently strong electric field, these can be independently measured. The measurement technique, described in detail in Ref. [283], is summarized here briefly.

When a rotating electric field is applied to a microsphere, the microsphere will spin to keep its dipole moment aligned with the electric field vector. Rather than remaining perfectly aligned as the electric field spins, however, the direction of \vec{p} will undergo harmonic oscillations about the direction of \vec{E} in the plane in which \vec{E} is rotating. Motion in this degree of freedom is called libration. The libration frequency scales like $f \propto \sqrt{|\vec{p}||\vec{E}|/I}$ where I is the moment of inertia of the microsphere. \vec{p} is given by the sum of the permanent dipole moment, \vec{p}_0 , and the induced dipole moment, $\alpha_m \vec{E}$, where α_m is the polarizability of the microsphere. Scanning through a range of electric field strengths and measuring the libration frequency allows for the scaling with \vec{E} and the contributions from each component to be measured.

The rotation of the microsphere is measured with the spin PD which detects

s-polarized light exiting the trap. As the birefringent microsphere rotates, the s-polarized light is modulated at twice its rotation frequency. Libration can then be measured using the sidebands of the main rotation peak. Fig. 8.10 shows example spin PD spectra at a range of electric fields and the corresponding plot of the libration frequency as a function of the electric field.




Figure 8.10: Measurement of the dipole moment, both permanent and induced, of a microsphere. The left panel shows a set of spectra under different applied rotation fields, with the libration sidebands appearing farther from the main rotation peak with increasing field strength. On the right, the libration frequency is plotted as a function of the rotation field, with the fit to the summed contributions from the permanent and induced dipole moments shown in red.

8.5.2 Contact Potentials

During a measurement, the shield and attractor are both grounded through the same power supply. This equalizes the Fermi level between both devices. Despite this, there should still be a contact potential difference between the surfaces of the shield and attractor due to their different work functions. This would create electric field gradients at the trap focus which would then couple to the EDM of the microsphere. A primer on the physics underlying contact potentials is included in Appendix F. The

remainder of this chapter is devoted to a technique by which these backgrounds can be measured in a set of dedicated background witness channels and subtracted from the signal channels.

8.5.3 In Situ EDM Background Monitoring

We take the electric field to be caused by a contact potential difference V_c and an applied voltage V_{app} between the attractor and the shield. In the quasistatic approximation, this electric field scales linearly with the potential difference between the attractor and shield. It is therefore useful to compute the “reference” electric field assuming a 1 V potential difference, \vec{E}_{ref} , and define dimensionless scaling factors to incorporate the voltage dependence. The dimensionless voltages are denoted by ζ with appropriate subscripts. The electric field at position \vec{r} in the trap when the attractor is at position \vec{r}_a is then given by,

$$\vec{E}(\vec{r}, \vec{r}_a) = \vec{E}_{\text{ref}}(\vec{r}, \vec{r}_a) \cdot (\zeta_c + \zeta_{\text{app}}), \quad (8.14)$$

where we have now relegated the position dependence to the first term, while the second term contains the dimensionless contact potential difference, ζ_c , and the dimensional applied voltage, ζ_{app} , which set the absolute scale. Plugging this into Eq. (8.13) gives the force on the microsphere in terms of the applied potential difference:

$$\vec{F}(\vec{r}) = (\vec{p} \cdot \nabla) \vec{E}_{\text{ref}}(\vec{r}, \vec{r}_a) \cdot (\zeta_c + \zeta_{\text{app}}). \quad (8.15)$$

Each component of the force, indexed by i , can be expressed in terms of the individual components of the other vector quantities as

$$F_i = (\zeta_c + \zeta_{\text{app}}) \cdot [(p_x \partial_x + p_y \partial_y + p_z \partial_z) E_{\text{ref},i}(\vec{r}, \vec{r}_a)]. \quad (8.16)$$

For now, we assume that \vec{p} is the permanent dipole moment of the microsphere, so its components are constant in time³. The term in the square brackets can then

³An FEA study of the electric field produced by 1 V potential difference between the attractor and shield suggests that $|\vec{E}| \lesssim 10^3$ V/m at the trap focus. For a typical microsphere with $|\vec{p}| \sim 25 e \cdot \mu\text{m}$,

be labeled as $F_{\text{ref},i}(\vec{r}, \vec{r}_a)$, which again contains all the position dependence with no dependence on the absolute scale of the potentials.

We see that the force scales linearly with ζ_{app} and vanishes when $\zeta_{\text{app}} = -\zeta_c$, as expected. This means that backgrounds arising from a contact potential difference can be eliminated by tuning the applied voltage to a particular value. In principle, scanning different values of ζ_{app} until the force reaches a minimum would allow ζ_c to be identified. However, as DC measurements tend to be more difficult than measurements at higher frequencies, we modulate ζ_{app} continuously instead.

Producing the Background Witness Sidebands

We use a bias modulation frequency $f_{\text{bias}} \gg f_0$ and monitor the response induced by this modulation. Take the bias as a function of time to be given by a sinusoid,

$$\zeta_{\text{app}}(t) = \zeta_{\text{bias}} \sin(2\pi f_{\text{bias}} t), \quad (8.17)$$

where ζ_{bias} is the dimensionless amplitude of the modulation. The attractor stroke is also sinusoidal, with an amplitude of y_0 and a frequency of f_0 :

$$y(t) = y_0 \sin(2\pi f_0 t). \quad (8.18)$$

Keeping the position of the microsphere fixed, and assuming that the attractor moves only in y , $F_{\text{ref},i}$ can be expressed as a function of $y(t)$ as

$$F_{\text{ref},i} = F_{\text{ref},i}(y_0 \sin(2\pi f_0 t)). \quad (8.19)$$

This can be expanded as a Fourier series,

$$F_{\text{ref},i} = F_{0i} + \sum_{n=1}^{\infty} F_{ni} \cos(2\pi n f_0 t - \varphi_{ni}), \quad (8.20)$$

the torque is then $4 \times 10^{-20} \text{ N} \cdot \text{m}$, which over a half-period of the attractor stroke would cause a change in angular momentum of magnitude $2 \times 10^{-21} \text{ kg m}^2/\text{s}$. Given that microspheres are typically spun up to $\sim 100 \text{ kHz}$ ($|\vec{L}| = 10^{-18} \text{ kg m}^2/\text{s}$) prior to a measurement, the resulting change in the orientation would be $\sim 0.1^\circ$.

where F_{0i} is the average force in component i and F_{ni} and φ_{ni} are the amplitude and phase, respectively, of the modulations for each Fourier component indexed by n . Plugging this back into Eq. (8.14) gives

$$F_i = F_{0i} \zeta_c + \zeta_c \sum_{n=1}^{\infty} F_{ni} \cos(2\pi n f_0 t - \varphi_{ni}) \\ + F_{0i} \zeta_{\text{bias}} \sin(2\pi f_{\text{bias}} t) + \zeta_{\text{bias}} \sum_{n=1}^{\infty} F_{ni} \cos(2\pi n f_0 t - \varphi_{ni}) \sin(2\pi f_{\text{bias}} t). \quad (8.21)$$

The product in the second sum can be rewritten as,

$$\cos(2\pi n f_0 t - \varphi_{ni}) \sin(2\pi f_{\text{bias}} t) = \frac{1}{2} \left[\cos\left(2\pi (f_{\text{bias}} + n f_0) t - \varphi_{ni} - \frac{\pi}{2}\right) \right. \\ \left. + \cos\left(2\pi (f_{\text{bias}} - n f_0) t + \varphi_{ni} - \frac{\pi}{2}\right) \right], \quad (8.22)$$

and the force component then becomes,

$$F_i = F_{0i} \zeta_c \\ + F_{0i} \zeta_{\text{bias}} \cos\left(2\pi f_{\text{bias}} t - \frac{\pi}{2}\right) \\ + \zeta_c \sum_{n=1}^{\infty} F_{ni} \cos(2\pi n f_0 t - \varphi_{ni}) \\ + \frac{\zeta_{\text{bias}}}{2} \sum_{n=1}^{\infty} F_{ni} \cos\left(2\pi (f_{\text{bias}} + n f_0) t - \varphi_{ni} - \frac{\pi}{2}\right) \\ + \frac{\zeta_{\text{bias}}}{2} \sum_{n=1}^{\infty} F_{ni} \cos\left(2\pi (f_{\text{bias}} - n f_0) t + \varphi_{ni} - \frac{\pi}{2}\right). \quad (8.23)$$

The first term is a constant which represents the time-averaged force on the microsphere due to the contact potential. The second term gives the force on the microsphere resulting from the bias itself. The third term contains the forces resulting from the contact potential with frequencies at integer multiples of f_0 ; these are the backgrounds at the drive harmonics we are trying to witness. The fourth and fifth terms have frequencies at integer multiples above and below f_0 , respectively. Crucially, the last three terms — including both the backgrounds themselves and new

frequency-shifted terms — have amplitudes proportional to the harmonic-dependent coefficient F_{ni} , so the spectral fingerprint should be the same up to a scaling factor. These sideband frequencies $f_{\text{bias}} \pm nf_0$ constitute witness channels in which the backgrounds at the drive harmonics nf_0 can be measured. The upper and lower sidebands are phase shifted relative to the signal channels, with a relative phasing of $2\varphi_{ni}$ between opposite sidebands of the same n . Measurements of both the upper and lower sidebands will allow for F_{ni} and φ_{ni} to be measured for all n , ensuring complete reconstruction of the spectral fingerprint of the backgrounds.

Let $F_i^{\text{meas}}(f)$ represent component i of the complex force vector measured at frequency f . The magnitude, A_{ni}^{\pm} , and phase, ϕ_{ni}^{\pm} , of the sidebands are given by

$$A_{ni}^{\pm} = |F_i^{\text{meas}}(f_{\text{bias}} \pm nf_0)| \quad \text{and} \quad (8.24)$$

$$\phi_{ni}^{\pm} = \arg \{F_i^{\text{meas}}(f_{\text{bias}} \pm nf_0)\}. \quad (8.25)$$

where $+$ and $-$ correspond to the upper and lower sidebands, respectively. With these quantities, a background template B_{ni} for component i of harmonic n , defined as

$$B_{ni} = \frac{\zeta_c}{\zeta_{\text{bias}}} \left[A_n^- \exp \left\{ -i \left(\phi_n^- + \frac{\pi}{2} \right) \right\} + A_n^+ \exp \left\{ i \left(\phi_n^+ + \frac{\pi}{2} \right) \right\} \right], \quad (8.26)$$

can be used to predict the background as a function of ζ_c . Since the drive harmonics are assumed to consist of backgrounds plus a potential signal, ζ_c can then be determined through a joint fit of both the signal and background templates to the data.

Testing with Real Data

This procedure was tested with two different microspheres. Prior to the measurements, the microspheres were spun to > 100 kHz in order to confine the EDM to a single plane. In principle, this should mean that all components of the vector \vec{p} average to zero; in reality, however, drifts in the EDM can still result in some non-zero out-of-plane component. While this cannot be measured with the rotating electric field off, its magnitude is constrained to less than the measured value. Fig. 8.11 shows

the measured force spectrum for a dataset that used $\zeta_{\text{bias}} = 0.1$ and $f_{\text{bias}} = 46.5$ Hz. The force from the bias itself appears well above the noise in both x and y . From Eq. (8.23), the phase should be either 90° or -90° , depending on the sign of F_{0i} . The x component does have the expected phase; it is not clear why a different phase is observed in y . Eq. (8.23) also shows that the sidebands should be phase shifted by φ_{ni} above and below the phase at f_{bias} . For $n = 1$, the only sideband clearly visible above the noise, this relationship holds in both x and y .




Figure 8.11: The magnitude and phase of the measured forces at the sidebands of the bias modulation frequency, indicated by the vertical dashed green lines. The response is visible above the noise both at the bias modulation frequency and at the first sidebands. In the lower panel, the phase at the sidebands is shown, with the dashed lines indicating that the phases of the sidebands are shifted equally from the main modulation peak, as expected. The peaks halfway between the green dashed lines in the top panel are from scattered light backgrounds at the harmonics of the attractor drive.

The peaks between the bias modulation sidebands originate from scattered light

at the harmonics of the attractor drive. As the scattered light is evidently much greater in magnitude than the EDM backgrounds in the sidebands, subtraction of the EDM backgrounds should have little impact on the physics sensitivity.

Accounting for the Induced Dipole Moment

The calculation performed above requires that \vec{p} not depend on ζ_{bias} ; in other words, \vec{p} is the *permanent* rather than the *induced* dipole moment. If ζ_{bias} is sufficiently large as to induce a dipole moment comparable to the permanent dipole moment, the spectral fingerprint measured at the bias sidebands will simply reflect backgrounds created by the modulation itself. Fortunately, if this occurs, it will leave a distinct signature in the data. We repeat the calculation with the permanent dipole replaced by the induced dipole,

$$\vec{p} = \alpha_m \vec{E}_{\text{ref}} \cdot (\zeta_c + \zeta_{\text{app}}). \quad (8.27)$$

Picking up from Eq. (8.16), we have,

$$F_i = (\zeta_c + \zeta_{\text{app}})^2 \cdot [(\alpha_m E_{\text{ref},x} \partial_x + \alpha_m E_{\text{ref},y} \partial_y + \alpha_m E_{\text{ref},z} \partial_z) E_{\text{ref},i}(\vec{r}, \vec{r}_a)], \quad (8.28)$$

where now the force scales like $(\zeta_c + \zeta_{\text{app}})^2$. As before, we denote everything in the square brackets $F_{\text{ref},i}$ and perform a Fourier expansion. We do not need to repeat the full calculation to understand the result. With ζ_c constant and $\zeta_{\text{app}} = \zeta_{\text{bias}} \sin(2\pi f_{\text{bias}} t)$, expanding the squared term gives

$$(\zeta_c + \zeta_{\text{app}})^2 = \zeta_c^2 + 2\zeta_c \zeta_{\text{bias}} \sin(2\pi f_{\text{bias}} t) + \zeta_{\text{bias}}^2 \sin^2(2\pi f_{\text{bias}} t) \quad (8.29)$$

which, after simplifying and collecting terms, becomes

$$(\zeta_c + \zeta_{\text{app}})^2 = \left(\zeta_c^2 + \frac{\zeta_{\text{bias}}^2}{2} \right) + 2\zeta_c \zeta_{\text{bias}} \sin(2\pi f_{\text{bias}} t) - \frac{\zeta_{\text{bias}}^2}{2} \cos(2\pi (2f_{\text{bias}}) t). \quad (8.30)$$

As before, we have a DC term which, when multiplied by the Fourier series, gives backgrounds at the attractor drive harmonics. The amplitude is proportional to the first term which now includes ζ_{bias} , so we are creating backgrounds due to the bias

modulation. We have a term with a frequency of f_{bias} again as well, though this time it is proportional to the contact potential in addition to the bias modulation amplitude. This creates the sidebands of f_{bias} which can be used for background witnessing — though the backgrounds we witness this way are now partly of our own making. The final term gives something that wasn't present previously — sidebands of twice the modulation frequency and proportional to ζ_{bias}^2 , with no dependence on ζ_c . Peaks at these frequencies indicate that the bias modulation is causing an induced dipole moment of sufficient magnitude as to result in a measurable response in the microsphere.

Unfortunately, at the time the first data was collected to test this background witnessing scheme, the effect of the induced dipole was not considered. The bias modulation frequency was chosen to be halfway between two drive harmonics to ensure maximum visibility without any spectral leakage from the stray light backgrounds. This choice of frequency means that the sidebands of $2f_{\text{bias}}$ coincide with the drive harmonics and as a result, they measure significant levels of stray light above the noise. It is therefore not possible to determine whether the induced dipole had a significant effect. In a future measurement, a different choice of f_{bias} should be used.

Limits on the Background Modulation Technique

This technique of modulating an experimental parameter on which the background depends can be a useful tool for in situ measurement, but there is a critical limitation worth noting. It must not be possible for a true signal to couple into the background channels; otherwise, it could be measured there and then subtracted from the signal channels, resulting in a missed discovery. By extension, there must be no coupling between the signal and the experimental parameter being modulated of any kind, even indirect. For example, the following would *not* work: modulating the laser power at some frequency f_{laser} and reading the scattered light backgrounds, which one would expect to be proportional to the total amount of light, from the sidebands, $f_{\text{laser}} \pm n f_0$. While there is no *direct* interaction between the laser power and a Yukawa interaction, if there were a Yukawa interaction, it would cause the microsphere to move in response. This motion would then cause changes in the power measured in

each quadrant of the QPD at nf_0 . If the laser power were being modulated, then the changes in power as a result of true microsphere motion would also show up as sidebands of f_{laser} , coupling the Yukawa signal into the background channel indirectly.

8.6 Summary and Current Status

At the start of this chapter I introduced three strategies for background subtraction. The first strategy — modeling the backgrounds — is still beyond reach. Having now been thoroughly investigated, scattered light backgrounds do not seem amenable at all to modeling as they evolve in time with no discernible cause. The second technique, that of isolating and measuring the backgrounds, can be done for the mechanical vibration backgrounds since their presence is independent of whether the attractor is near enough to the microsphere to produce a signal. Witnessing the backgrounds in situ with a different set of data streams appears the most fruitful, at least for diagnostic purposes. A scattered light data stream which is insensitive to motion has been constructed, allowing for the amplitude of scattered light backgrounds to be measured on a file-by-file basis. Despite this, the inconsistent coupling of these backgrounds into the measured harmonics in x and y prohibits their direct subtraction. For electromagnetic backgrounds, in situ background subtraction appears more promising. A technique has been developed which allows for the spectral fingerprint of the backgrounds in both x and y to be measured and subtracted, with only the overall amplitude as a fit parameter. As these backgrounds are subdominant to scattered light, the utility of this technique in a physics run has yet to be demonstrated, though future reductions in scattered light backgrounds may enable such a demonstration.

Conclusions

The work described in this thesis spans two experiments that share the goal of advancing fundamental physics but take vastly different approaches to it. Each one targets a separate topic in which a gap in our current understanding offers the promise of a new discovery.

The field of neutrino physics has significant potential for groundbreaking discoveries, along with the tantalizing possibility that these discoveries could be used as stepping stones to future breakthroughs. Definitive evidence of $0\nu\beta\beta$ would be one such example. In addition to establishing that neutrinos are Majorana fermions, it could lead to an explanation of the matter-antimatter asymmetry via leptogenesis, it could lend support to Seesaw-based dark matter models, and it could point to physics originating from otherwise-inaccessible energy scales.

The nEXO experiment to search for $0\nu\beta\beta$ boasts a mature design based on a well-tested detector technology. LXe TPCs use the anticorrelated production of both charge and light signals to achieve excellent energy resolution. nEXO adds to this two additional levers for signal/background discrimination: machine-learning-based event topology reconstruction, which allows for multi-site γ backgrounds to be distinguished from the predominantly single-site $0\nu\beta\beta$ events, and standoff distance, which uses the fact that the natural self-shielding of LXe results in a nearly background-free inner volume. As the enriched xenon — the decay source — is removable, in the event of a possible discovery, it can be replaced with natural or depleted xenon to check that the signal disappears. nEXO is one of three proposed tonne-scale experiments, along with LEGEND and CUPID, that together constitute a robust, global, multi-isotope search program. While ever-present funding challenges have continually loomed large

over nEXO, I hope that the discussion of nEXO in this thesis adequately justifies its essential role in this program.

Meeting nEXO's energy resolution target of 0.8% at the $\beta\beta$ -decay Q -value requires that the charge and light responses be calibrated to sufficient precision. In Chapter 5, I showed how the electron capture decay of ^{127}Xe within the active volume could be used to calibrate the position-dependent light response, or lightmap. With colleagues at Stanford, I demonstrated the deployment of a ^{127}Xe calibration source in a test TPC for electron lifetime calibrations. With simulated data produced using the `nexo-offline` framework, I also showed how a machine learning algorithm could be used to reconstruct the lightmap in nEXO on a more frequent basis than in the baseline plan, which uses external γ calibration sources. I found that with a 1 Bq activity maintained in the TPC throughout a calibration campaign, less than six days would be required to meet nEXO's energy resolution target of 0.5% lightmap error in the inner two tonnes. This calibration could be conducted during data acquisition, as the low-energy ^{127}Xe decay does not interfere with the physics search, eliminating the need for dedicated lightmap calibration periods.

Xenon purification plays an important role in ensuring good charge signal quality in nEXO. Continuous purification is required to remove electronegative impurities which diffuse into the xenon from wetted surfaces. Standard purifiers, while excellent at their nominal purpose, often emanate radon into the flowing xenon gas at unacceptably high rates. My work described in Chapter 6 focused on developing a purifier made from high-purity zirconium to eliminate this source of backgrounds. The work involved multiple components: the design and construction of a prototype, the characterization of its purification performance, and radioassay measurements of the candidate material. The first two of these three were met with success. A prototype purifier was constructed which was found to have a purification efficiency of 83%; lower than the off-the-shelf alternative, but sufficiently high to allow for an increased flow rate to compensate for the difference. The radioassay measurements, on the other hand, showed that the high-purity zirconium was slightly more radioactive than the standard getter pellets. At the same time, there are indications that the radium resides on the surface of the pellets, potentially allowing for its removal

through a high-temperature bakeout. This possibility is under investigation, and the success or failure will determine the future of high-purity zirconium for use in nEXO's purifiers.

The second topic of this thesis concerns the possibility that our understanding of gravity is limited to the macroscopic world. Gravity has yet to be properly incorporated into a quantum mechanical framework, leaving it ripe for study at ever-shorter distances. Furthermore, many theoretical models postulate microscopic extra dimensions or new light bosons which could manifest as modifications to the gravitational ISL.

A search for new interactions of this kind, conducted using optically-levitated microspheres, is the focus of Chapter 7. The first search using this technique was performed in 2020, and since then, a number of upgrades have been made to the experimental setup. Most notably, the laser intensity noise was improved, and the attractor was made less reflective with the addition of a Platinum Black coating. Accompanying these hardware upgrades was my overhaul of the data analysis pipeline and statistical framework. With reduced scattered light backgrounds and intensity noise, measurements of both the x and y components of the force vector at multiple harmonics could be incorporated into the search. This enabled a major feature of the experiment — the ability to identify signals based on their unique spectral fingerprint — to be properly included in the statistical model, fully exploiting the inconsistencies between the reconstructed signal strengths in x and y and across all harmonics. The resulting constraints are $\sim 100 \times$ stronger than those previously reported.

Despite the clear improvement over the prior result, measurement backgrounds still limit the sensitivity of the experiment. Much of my work, discussed in Chapter 8, was centered on developing diagnostic tools that could be used to better understand these backgrounds. I developed a technique for calibrating the QPD response to the eigenmodes of the microsphere, allowing for cleaner force measurements, while at the same time providing a non-motional data stream in which the scattered light backgrounds can be monitored. This method gives a model-independent prescription for mapping out which pixels are most sensitive to motion compared to scattered light,

which will yield even greater benefits on a future sensor with many more pixels. I also demonstrated a method of witnessing backgrounds originating from electromagnetic interactions between the attractor and the dipole moment of a microsphere. This method, based on producing a copy of the background spectrum in the sidebands of a chosen modulation frequency, could allow for in situ background subtraction in the future, when scattered light backgrounds no longer dominate. While these techniques may provide some degree of background removal capability, even the most robust witnessing and subtraction methods cannot be expected to decrease backgrounds by multiple orders of magnitude. Subtraction of backgrounds in analysis is a poor substitute for their elimination in hardware. With these diagnostic tools in hand, reduction of the backgrounds should be the primary focus of future upgrades.

There are noticeable contrasts between the two experimental techniques described in this thesis. Experimental particle physics is a mature field in terms of both the hardware and software it employs. Its detector technologies have benefited from countless improvements over multiple generations of experiments. Alongside detector development, sophisticated tools for simulation and statistical analysis have been developed to maximize the returns from a given experiment. Meanwhile, optically-levitated sensors are only now passing the proof-of-principle stage and are beginning to catch up with the state of the art for tests of fundamental physics on the precision frontier. Many of their unique systematic effects are just beginning to be understood, and the most significant advancements are still to come.

The SM is known to be incomplete, though the next steps along the path to a more complete description of nature have yet to be found. Until a clear direction emerges, the daunting problem of discovering physics beyond the SM demands a wide range of strategies. The experiments described in this thesis represent two such strategies among many. This breadth of experimental techniques will be key to making the next breakthrough, whatever it may be.

Summary for Non-Scientists

The Standard Model of particle physics is the preeminent description of nature at the most fundamental level. It describes the world in terms of interactions among the known elementary particles, and it can be used to predict how these particles behave under specific conditions. These predictions can then be compared to experimental measurements to validate the theory. So far, the Standard Model has stood up to nearly every experimental test to which it has been subjected. Despite this, it is known to be an incomplete theory. It doesn't explain dark matter or dark energy, which are estimated to collectively account for 95% of the energy content of the Universe. It also fails to include gravity, which to this day is still best described by Albert Einstein's century-old theory of General Relativity. Finding experimental results that disagree with the Standard Model is therefore critical to understanding how it falls short and how it can be extended.

Majorana Neutrinos

Neutrinos are elementary particles produced in radioactive decays. They are second in abundance only to photons, they travel at near the speed of light, and they interact so rarely with the matter they pass through as to be virtually undetectable. They also represent a rare failure of the Standard Model. Neutrinos were first postulated in the 1930s to explain why energy appeared to be lost in some radioactive decays. They were later incorporated into the Standard Model as massless particles. It was only in the late 1990s that neutrinos were conclusively shown to have tiny but non-zero masses, and the Standard Model, by consequence, was shown to be incomplete.

Most other particles in the Standard Model have mass, so it may seem strange that neutrinos with mass pose such a challenge to the theory. Unlike other matter particles with mass, which are known as *Dirac* particles, neutrinos may be *Majorana*¹ particles instead. Loosely speaking, Dirac particles are distinct from their antimatter counterparts, while Majorana particles are their own antiparticles. Any particle which carries charge is forbidden from being a Majorana particle. This is because particles have opposite charges to their antiparticles, so for a particle to be its own antiparticle, it must have zero charge. As the only uncharged particles with mass, neutrinos are the sole candidates for Majorana status. A confirmation that neutrinos are Majorana particles would have profound implications for physics.

Neutrinos could be confirmed as Majorana particles through the observation of an exceedingly rare interaction called *neutrinoless double beta decay*. In regular double beta decay, two electrons (matter particles) and two antineutrinos (antimatter particles) are emitted from the nucleus of an atom. While this process is rare, it is known to occur in a number of isotopes. In the predicted neutrinoless version, the antineutrinos combine and vanish in a process known as *annihilation*, and only the electrons are emitted. Annihilation is only possible between a particle and its antiparticle, so a search for neutrinoless double beta decay is a direct test of whether a neutrino is its own antiparticle.

If observed, neutrinoless double beta decay would constitute the first evidence of the creation of matter (the electrons) without antimatter. This would have far-reaching implications for cosmology. According to the Standard Model, energy can be converted into equal quantities of matter and antimatter. Yet we know that the Universe is almost exclusively made of matter, with little antimatter to be found. This requires that some process in the early Universe created matter and antimatter asymmetrically. An observation of neutrinoless double beta decay would demonstrate that the laws of physics allow for this to occur.

¹Pronounced “my-uh-RAH-nuh”.

The nEXO Experiment

To search for neutrinoless double beta decay, an experiment should consist of a large amount of a candidate isotope and a way of instrumenting it to detect the decay. The nEXO experiment (the topic of Chapter 4) will use an isotope of xenon as the decay candidate. Five tonnes of xenon enriched in this isotope will be liquified inside a cylindrical copper vessel. When a decay occurs within the xenon, a flash of light and a cloud of charge will be produced at the decay site and detected by light and charge sensors covering the interior of the vessel, allowing the energy of the decay to be determined. A rendering of the experiment is shown in Fig. 4.2 on page 50. The experiment will be conducted at SNOLAB, a laboratory more than 2 km underground in a nickel mine near Sudbury, Ontario, Canada, to shield the sensitive detector from the showers of particles produced by cosmic rays in the atmosphere. My work for nEXO focused on two critical aspects of the detector design: calibration and xenon purification.

Development of a Calibration Scheme for nEXO

Calibration is required to ensure that the energy of decays is measured correctly. The particular calibration I focused on corrects the amount of light detected based on its location in the detector, since some regions of the detector are more visible to the light sensors than others.

I showed how a different radioactive isotope of xenon can be mixed throughout the vessel to produce decays with a known amount of light uniformly across the detector volume. The discrepancy between the amount produced and the amount measured could then be used to create a map of how efficiently light is collected throughout the detector. This work was done in two ways. In collaboration with colleagues in the lab at Stanford, I injected the radioactive xenon into a small-scale test detector as a proof of principle. I also ran simulations of the calibration process to validate its use for the full-scale nEXO detector. This work was published in Ref. [196] and is described in more detail in Chapter 5.

Design of a Xenon Purifier for nEXO

The energy of decays in the xenon volume is measured using the charge and light they produce. Xenon is transparent to the light produced, so the light reaches the sensors unimpeded. The charge, however, moves slowly through the xenon, guided by an electric field pulling it toward the sensor array. If impurities like oxygen and carbon dioxide are present in the xenon, they can capture a fraction of the charge produced in the decay before it is detected. To mitigate this, the xenon in nEXO will be continuously circulated through dedicated purifiers. The purifiers can be purchased off the shelf, though they are known to contain trace amounts of natural radioactivity that are problematic for sensitive detectors like nEXO.

My work, described in Chapter 6, involved designing an alternative purifier using high-purity zirconium metal, which we suspected of being lower in natural radioactivity than the standard purifiers. To test this, I procured multiple samples of high-purity zirconium and coordinated with colleagues within nEXO and elsewhere to subject the samples to a series of measurements. At the same time, I designed and built a prototype purifier in the lab at Stanford and brought it to SLAC National Accelerator Laboratory for testing. With the help of scientists at SLAC, I measured the purification capabilities of my custom purifier by deliberately loading xenon with impurities and using my purifier to remove them. As the results from the measurements of the zirconium samples started to come in, we were disappointed to learn that our suspicion was wrong, and the high-purity zirconium had roughly the same level of natural radioactivity as the standard purifiers. But we found encouraging evidence to suggest that the radioactive elements reside predominantly on the surface of the metal, so work is ongoing to determine if they can be removed. Depending on the success of this effort, purifiers using high-purity zirconium may be used by nEXO.

Micron-Scale Interactions

There are many reasons to search for new interactions at the micrometer scale. One is that gravity has never been measured over such short distances. To fully understand

gravity, its behavior at the micrometer scale must be confirmed, though due to its weakness a direct measurement is not feasible with current technology. Other reasons are more exotic in nature, with many theories predicting extra dimensions or new particles that could manifest as minuscule forces between objects a few micrometers apart.

A straightforward way to measure a force is to take a spring of known stiffness and measure how much it extends or compresses when the force is applied to it. This same principle underlies the measurement described here, but in this case the spring is replaced by a glass sphere levitated by a focused laser beam. Remarkably, this system behaves in much the same way as a mechanical spring. The laser applies forces on the sphere to keep it aloft and confined near the focus of the beam. When the sphere is pulled to one side, the light pulls it back, and the farther the sphere is pulled, the more the light resists. At the same time, the movement of the sphere causes measurable deflections of the laser light, allowing its motion to be reconstructed. The exact mechanism is complicated — it arises from a combination of reflections and refractions at the surface of the sphere, as illustrated in Fig. 7.1 on page 116 — but the result is that a sphere trapped in a laser beam can be used for precise measurements of very weak forces.

Search for New Interactions

Our search for new interactions uses a 10-micrometer-diameter sphere trapped at the focus of a laser as described above. A device with alternating strips of high and low density, called the attractor, acts as the source for new interactions. The complete setup is shown in Fig. 7.3 on page 122. By oscillating the attractor back and forth near the trapped sphere, any attractive or repulsive interaction with the density variation would cause the sphere to undergo a distinct sequence of motions in response. This would allow the response to be identified as the signature of a new interaction rather than something more mundane. My role in this project was to write the data analysis code and build the statistical model used to conduct the search. After conducting the search and looking at the data, we found that we had measured an apparent force,

but without the signature expected for a new interaction. We used this data to rule out new forces of a particular strength. Chapter 7 describes this work, and a paper summarizing it more concisely has been submitted for publication [244].

Identifying the Apparent Force

Since our measurements were not consistent with the signature of a new force, much of my work was on understanding their origin. This is described in Chapter 8. It seemed likely from the start that this apparent force was caused by stray laser light scattering from the attractor and landing on the sensor used to determine the sphere's position. I came up with a way to distinguish sensor readings originating from motion of the sphere from those caused by stray light. This method showed us that the apparent force was indeed originating predominantly from stray light. Another suspected source of the apparent force was from electromagnetic interactions between the sphere and the attractor. I developed a way to measure this effect and understand its contribution to our data. Efforts are currently underway to further mitigate the stray laser light interfering with our search for new interactions. If those are successful, the leftover electromagnetic interactions could be dealt with using my method for additional improvements in sensitivity.

Conclusions and Outlook

The two experiments described in this thesis are unified by the goal of discovering physics that cannot be explained under the Standard Model. The nEXO experiment will search for a predicted nuclear process that would have major implications for our understanding of nature. Taking an entirely different approach, levitated spheres can be used as sensors to search for the tiny forces that are predicted by many new theories. As we don't yet know where the next breakthrough will arise, this broad portfolio of experimental approaches is essential to further progress.

Appendix A

Bayesian Estimation of the $0\nu\beta\beta$ Discovery Probability

The calculation performed here is based on similar calculations from Refs. [284] and [285] but uses updated neutrino oscillation parameters. Experiments which measure neutrino oscillation or mass parameters often report a $\Delta\chi^2(\theta_i)$ curve, where θ_i is the parameter being measured. This can be used to compute the likelihood of obtaining the data, x_i , conditional on θ_i , as

$$\mathcal{L}_i(x_i|\theta_i) \propto \exp\left\{-\frac{\Delta\chi^2(\theta_i)}{2}\right\}. \quad (\text{A.1})$$

A global likelihood function can be constructed to represent the probability of obtaining all experimental data, X , conditional on the set of all parameters, Θ , as

$$\mathcal{L}(X|\Theta) = \prod_i \mathcal{L}_i(x_i|\theta_i). \quad (\text{A.2})$$

We are interested in using this likelihood to compute the posterior probability of Θ given X , which is related to the likelihood by Bayes' Theorem,

$$P(\Theta|X) \propto \mathcal{L}(X|\Theta) \pi(\Theta) \quad (\text{A.3})$$

where $\pi(\Theta)$ is the prior probability of Θ . This calculation can be done using a Markov Chain Monte Carlo (MCMC) approach. The calculation proceeds by choosing a basis set of parameters and assigning them priors. I have used the set

$$\Theta = \{\Sigma, \Delta m_{21}^2, \Delta m_{23}^2, \theta_{12}, \theta_{13}, \eta_2, \eta_3 - \delta_{\text{CP}}\}, \quad (\text{A.4})$$

with non-informative priors for all parameters: the mass terms have log-uniform priors while all angles have uniform priors on $[0, 2\pi]$. For computational convenience, the likelihoods and priors are implemented via their negative logarithms, with

$$-\log \pi(\phi) \propto \begin{cases} \text{constant}, & 0 \leq \phi \leq 2\pi \\ \infty, & \text{otherwise} \end{cases} \quad (\text{A.5})$$

for each $\phi \in \{\theta_{12}, \theta_{23}, \eta_2, \eta_3 - \delta_{\text{CP}}\}$, and

$$-\log \pi(\theta_i) \propto \begin{cases} \log m, & m \geq 0 \\ \infty, & \text{otherwise} \end{cases} \quad (\text{A.6})$$

for each $m \in \{\Sigma, \Delta m_{21}^2, \Delta m_{23}^2\}$. The results of this calculation are not sensitive to the parameter basis or the choice of priors used; see Ref. [284] for further discussion. The \mathcal{L}_i were calculated using published data for the oscillation parameters (from the NuFIT-5.2 release), m_e (from recent results from KATRIN) and $\langle m_{\beta\beta} \rangle$ (from KamLAND-Zen) [51, 286, 100]. The MCMC was implemented using the `emcee` Python package [287]. Sample chains consisting of 10^7 sampled Θ values for each of the NO and IO scenarios were produced.

To create the left panels on the lobster plot in Fig. 2.10, $\langle m_{\beta\beta} \rangle$ and m_{lightest} were computed for both the NO and IO scenarios using a number of the sampled Θ values. The calculated masses were then binned in a two-dimensional histogram with log-spaced bin edges. The same procedure was repeated with Σ in place of m_{lightest} to create the plots in the right panels.

Unlike the standard, uniformly-colored lobster plot, this version conveys how little probability density is contained in the downward-pointing tail under the NO scenario.

To see why this is so, one can fix all other parameters and plot $\langle m_{\beta\beta} \rangle$ as a function of the Majorana phases, η_2 and η_3 (Fig. A.1).




Figure A.1: Dependence of the effective Majorana mass on the two Majorana phases.

The required coincidence of phases becomes clear in this plot. $\langle m_{\beta\beta} \rangle$ is only driven to zero for precise combinations of η_2 and η_3 , while for the vast majority of the parameter space their values have almost no impact. Of course, there is nothing prohibiting some underlying physics that drives the phases toward these particular values, but in the absence of such a mechanism (or an unfortunate coincidence), $\langle m_{\beta\beta} \rangle$ should be expected to be at or above the meV scale.

Appendix B

The Stanford Liquid Xenon Lab

Much of the work described in Chapters 5 and 6 was done using the test facility in the lab at Stanford, which I helped to construct, maintain, and operate. This appendix describes the facility and some of its applications. It has been adapted from a manuscript submitted for publication of which I am a primary author [288].

B.1 Overview

The Stanford LXe TPC test facility, shown in Fig. B.1, includes two independent test stands, each with the infrastructure necessary to operate a variety of different LXe test chambers. The “Large System” is so named because it uses a cryostat with a 50 cm inner diameter, while the “Small System” uses a cryostat with a 40 cm inner diameter. The Small System and Large System can host test chambers with flange diameters of up to 30.5 cm (DN250CF) and 41.9 cm (DN350CF), respectively.

For each system, a flange mounted at the top of a Unistrut support frame acts as the feedthrough interface between the air side and cold space of the cryostat, sealing to a vacuum-insulated dewar using an o-ring. LXe test chambers are suspended from the underside of the top flange using low-thermal-conductivity rods, such as 1/4” SS316 all-threads or thin-wall SS316 tubing welded to fixture plates. Dewars are raised to mate with the top flange using a rope-and-pulley system. Each Unistrut frame includes a diamond-tread platform, providing access to the top flange, DAQ

rack, and xenon handling manifold. Both manifolds are connected to the same xenon storage bottle, with valves to direct xenon into and out of either system.

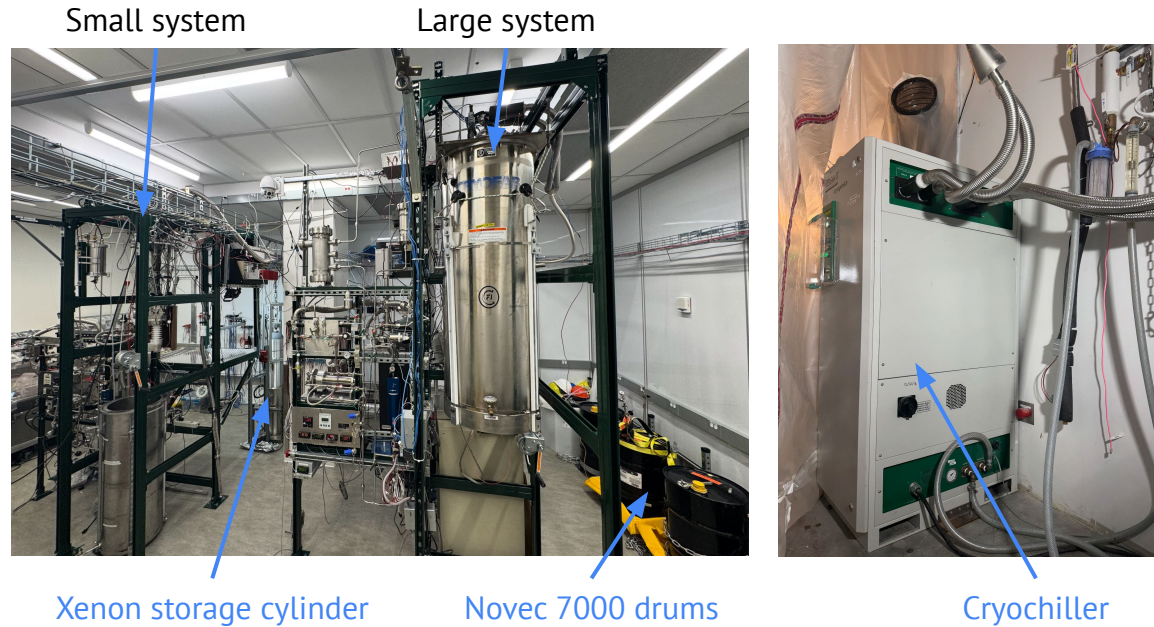


Figure B.1: Left: photo of the lab showing the two LXe test stands, the xenon storage cylinder, and the Novec 7000 supply. Right: the cryo chiller outside the lab.

B.2 Cryogenics

Unlike similar facilities, which typically use a cold finger conductively coupled to a test chamber in a vacuum space, the systems described here achieve cooling by immersing the test chambers into a temperature-controlled bath of cooling fluid. This alternative technique possesses a number of advantages, particularly with larger-scale detectors.

The large thermal mass of the cooling fluid, which far exceeds the thermal mass of LXe in the test chambers, results in excellent temperature stability; if cooling is lost due to a failure of the cryostat, the temperature of the LXe rises slowly and in sync with the temperature of the cooling fluid. For experiments using greater than 10 kg LXe masses, which can take multiple days to fill, this level of thermal stability is greatly appreciated by the operators that monitor for failures.

Furthermore, in the liquid-submerged cryostat, no mechanical or conductive thermal contact is required between the chamber and cooling element. Convective cooling between the liquid, cooling element, and submerged chamber removes the need for a copper LXe chamber or copper braids to achieve thermal uniformity. A chamber of any geometry can be cooled in this fashion, provided it fits inside the cryostat. This allows the LXe chambers to be constructed using off-the-shelf ultra-high-vacuum hardware such as ConFlat (CF), adding convenience and reducing cost.

B.3 Xenon Filling & Handling

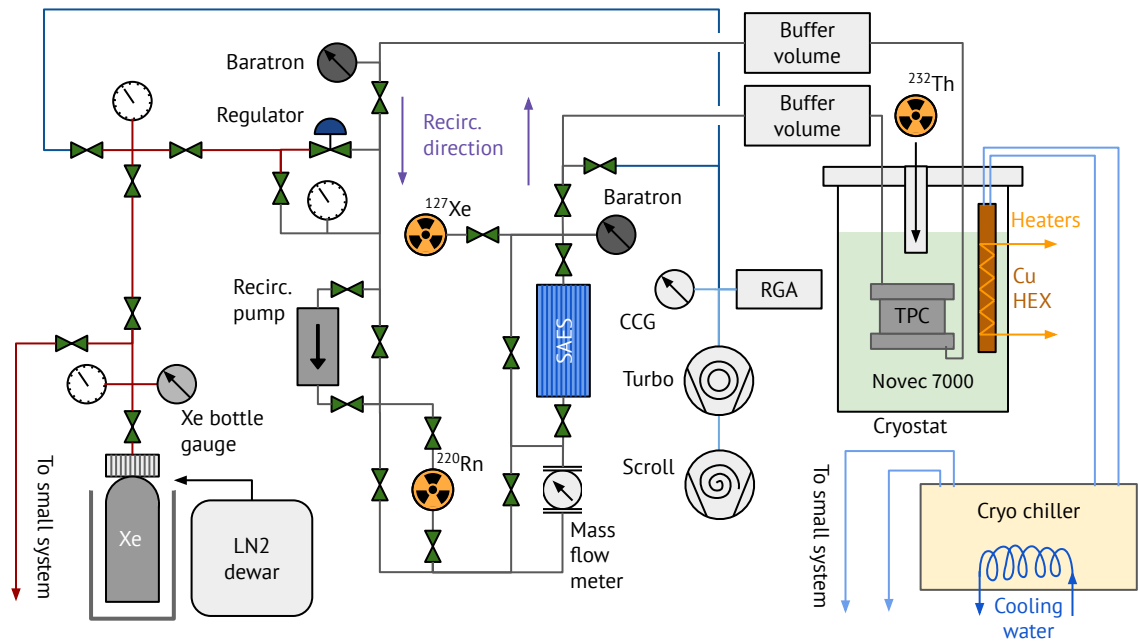


Figure B.2: Piping & instrumentation diagram for the Large System. Red lines represent high-pressure tubing while blue lines represent tubing maintained at vacuum. The Small System shares the xenon supply and the cryo chiller; otherwise the diagram for the Small System is the same.

A diagram illustrating one of the xenon systems in its entirety, including xenon

manifold, cryostat, and instrumentation, is shown in Fig. B.2. In addition to illustrating the context of the cryostat design and operation, this diagram will also be referred to in later subsections related to experimental capabilities like radioactive source deployment and purification.

Cooling power is provided to both systems by a dual-coil, Telemark TVP-2000 cryo chiller, which circulates a proprietary refrigerant in a closed-loop system. The fridge is supplied with room temperature water from the building for cooling its internal heat exchanger. Two sets of vacuum-jacketed refrigerant lines connect the fridge to feedthroughs on the cryostats, allowing for a closed-loop flow of refrigerant through a copper heat exchanger submerged within the cold space of each cryostat.

Each heat exchanger is suspended from the top flange, and consists of a copper tube with a 6 mm outer diameter brazed in a groove that meanders across a finned copper slab. The copper slabs are shaped as partial cylinders with a diameter close to the inner diameter of the cryostat, allowing for access to the central region where test chambers are suspended. This configuration is designed to promote convection of the cooling fluid. Cartridge heaters are embedded within the copper-slab heat exchangers to increase the rate of warming during xenon recovery. Thermocouples are fixed to the inlet tubing, outlet tubing, and at the top, middle, and bottom of the heat exchanger.

The cooling fluid used in these systems is the hydrofluoroether Novec 7000 produced by 3M [289]. Novec 7000 freezes at 151 K, enabling its use in the 161–175 K temperature range for LXe-based experiments. At 163 K, it has a kinematic viscosity of 8×10^{-6} m²/s, providing acceptable convective heat transfer in the LXe temperature range. The use of Novec 7000 as a cooling fluid for LXe TPCs was pioneered by EXO-200 [290] due to its intrinsic radiopurity, shielding power, viscosity at low temperatures, and thermal mass [291].

A typical experiment in the Stanford LXe facility involves filling a cryostat with 200–300 kg of Novec 7000; the substantial thermal mass ensures temperatures at the chamber remain stable, and provides a generous window in which the xenon in the chamber will remain at manageable pressures in the case of a loss of cooling power. Novec 7000 is added by simultaneously cooling the heat exchanger and keeping open a

tubing line between the cryostat (at rough vacuum) and a Novec 7000 storage drum. Filling is also facilitated by pressurizing the drum with nitrogen to overcome the hydrostatic column pressure of the fluid in the cryostat. Once the heat exchanger is partially submerged, refrigerator cooling rapidly reduces the Novec 7000 vapor pressure (9.5 psi at room temperature), reducing the pressure required to push liquid from the drum to the cryostat.

Cooling remains on continuously until the top thermocouple on the heat exchanger reaches the desired filling temperature of 163 K. Fig. B.3 shows an example cooling and warming cycle for an experiment in the Large System. Cooling typically takes about 10 hours, though ~ 5 more hours of normal on-off cooling are required for the temperature distribution throughout the Novec 7000 volume to reach its equilibrium state.

B.3.1 Performance

During operation, the temperature of a single thermocouple bolted to the heat exchanger is maintained by switching refrigerant flow on or off. A cooling duty cycle of 25% with a period of 5 minutes is typical. The Novec 7000 bath ensures that the chamber remains insensitive to temperature changes caused by the change in cooling state. xenon cell temperatures have been shown to remain stable to within 100 mK over many hours of operation (Fig. B.4 for example). There are no sensors measuring the LXe directly; thermocouples and resistance temperature detectors (RTDs) are instead placed in thermal contact with the body of the SS316 CF chamber.

Heat leaking into the cryostat is dominated by conduction through the cryostat top flange — a solid steel plate, 1.3 cm thick in the Small System and 1.6 cm thick in the Large System — coupled through gas convection of dry nitrogen between the Novec 7000 liquid surface and the top flange. The nitrogen is used to back-fill the cryostat volume to 1 atmosphere, mitigating potential water vapor leaks from the numerous KF connections and epoxied electrical feedthroughs. Permanent insulation of this region is impractical because of the density of cables and tubes extending through it.




Figure B.3: Cooling and warming temperatures and timescales for an experiment in the Large System cryostat. The temperatures during the multi-day experiment are redacted from the x-axis. A visible feature in a thermocouple at the bottom of the LXe cell associated with the time when LXe starts condensing into the chamber.

A rate of natural warming of 1 K per hour was measured by stopping cooling intentionally after establishing equilibrium at 168.5 K for multiple days (Fig. B.5). Based on the masses and heat capacities of objects within the cryostat, this implies a roughly 80 W heat leak. Because the LXe systems have been designed with a pressure rating of 1500 torr, operators have approximately 12 hours to remove the xenon in case of a loss of cooling power. While the system status is checked on a more frequent basis than every 12 hours, this flexibility has reduced the burden on operators who otherwise would have had to respond immediately to power outages. See Section B.4 for further detail about the monitoring and alarm system architecture.




Figure B.4: Stability of the xenon cell temperature measured by a thermocouple mounted on its top flange. Averaging is done over five on-off cooling cycles.

B.3.2 Filling & Recovery

A single aluminum-body storage cylinder containing ~ 30 kg of xenon, designed to be submerged in liquid nitrogen (LN2) for cryopumping recovery¹, is used to fill test chambers in both systems. This cylinder is permanently plumbed into a high-pressure manifold with valves that can be used to direct flow to one of the two xenon handling manifolds. Each of the handling manifolds uses a clean regulator² to supply xenon to the low-pressure region, while a bypass valve allows flow in the opposite direction during xenon recovery. The xenon storage bottle is suspended from a strain gauge, allowing for real-time monitoring of the mass of xenon that has been filled into an experimental chamber. In addition to the strain gauge, a mass flow

¹Luxfer SGS N265-SGS

²APTech two-stage, tied diaphragm regulators, models AP1720SM and AP1702SM.

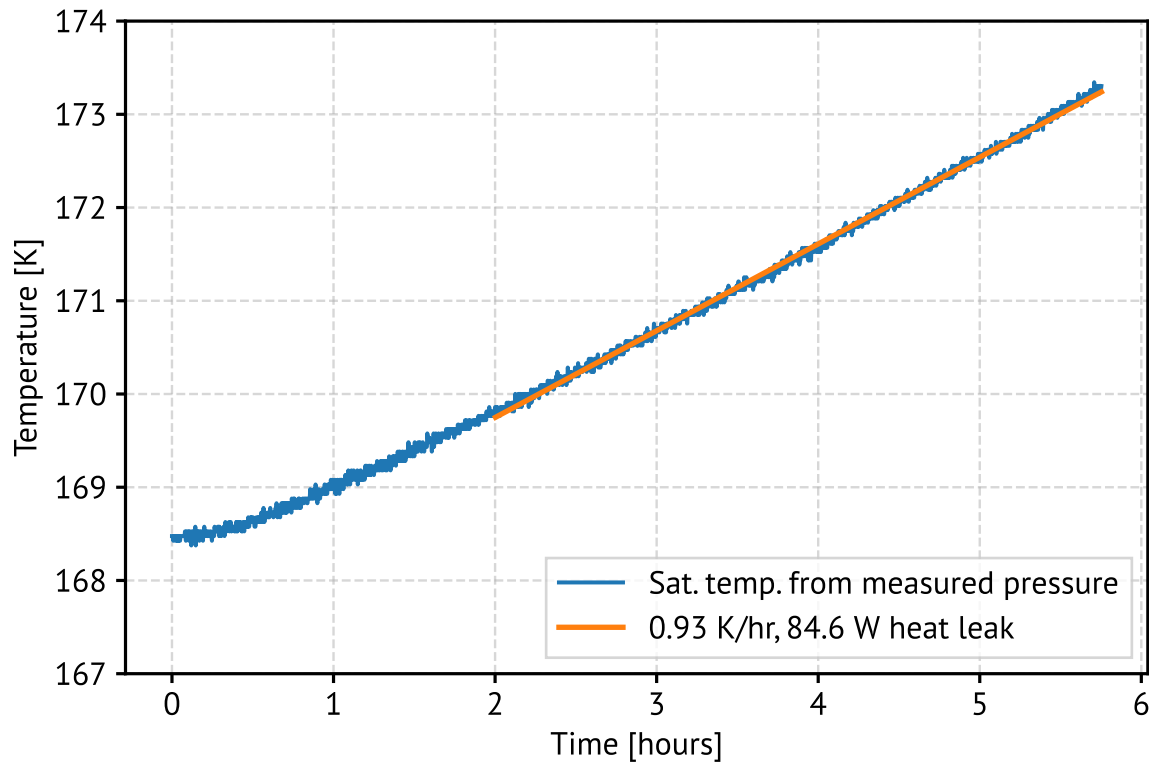


Figure B.5: Xenon temperature increase due to natural warming of a test cell while installed in the Large System. The linear fit is used to estimate an 84 W heat leak based on an estimate of the heat capacity of dominant masses in the cryostat. The Novec 7000 mass is 230 kg, the chamber is 38 kg, the copper heat exchanger is 16 kg, and the LXe mass is 9 kg, for a total of capacity of 326 kJ/K.

meter³ plumbed into the delivery manifold is used to measure xenon flow rates and reconstruct delivered mass. Below the suspended storage bottle is a dewar which can be raised and filled with LN₂, submerging the bottom portion of the aluminum storage bottle for use as a xenon cryopump.

Prior to filling, each system is pumped to vacuum with its own scroll pump and turbomolecular pump (turbopump). Each system is equipped with its own residual gas analyzer (RGA) and cold-cathode gauge (CCG) located near the turbopump. Under typical conditions, the vacuum achievable is $\sim 10^{-7}$ hPa at room temperature at the location of the gauge. However, the test chamber is separated from the pump and

³MKS Type 179A and 1479A mass flow meters.

CCG by ~ 4 meters of 13 mm diameter tubing, representing an effective conductance of about 0.03 L/s and a source of surface-water outgassing from 1200 cm² surface area. Lumped-element circuit models are maintained to represent each system's xenon delivery manifold, and are used to estimate pressure as a function of location and infer outgassing rates from control experiments. Depending on the outgassing state inside the chamber and in the manifold tubing, pressures inside the test chambers can fall between 10–100 \times larger than the value measured at the pump.

During filling, the regulator is opened until pressure gauges⁴ at the cell inlet and outlet read the desired cell pressure. A heated non-evaporable getter (NEG) purifier⁵ in the flow path is used to purify the xenon during filling. The filling rate is limited by the rate of condensation at the liquid surface and stainless steel walls, which in turn depends on the surface area in the detector. For this reason, the larger test chambers can be filled at higher flow rates. A typical filling time is ~ 24 hours for ~ 30 kg.

Recovery of the liquefied xenon is limited by the temperature of the LXe surface inside the test chamber. To decrease the time it takes to recover LXe, limited by the natural warming rate of 1 K per hour, heaters embedded in the copper heat exchanger apply up to 1 kW of power to the Novec 7000. A constant supply of LN2 is maintained to keep the xenon storage cylinder cold during the recovery process.

Because the recovery rate exceeds the range of the mass flow meter, and the cylinder mass cannot be measured accurately due to the changing buoyant force from LN2, the pressure in the chamber is used as an indication of the progress of recovery. Early in the recovery process when the xenon is coldest, the flow is throttled through a single valve to ensure that the LXe surface does not drop below freezing due to evaporative cooling. A finned aluminum heat sink at the storage bottle inlet passively prevents xenon freezing as gas enters and expands in the bottle during fast recovery. Multiple other tube and valve heaters consisting of ceramic heaters embedded in aluminum blocks can be placed around the manifold to prevent freezing at regions most susceptible to cooling by the fast-moving xenon gas.

As the Novec 7000 surrounding the chamber warms, the valve can be fully opened,

⁴MKS Baratron Capacitance Manometer

⁵SAES Monotorr PS3-MT3-R-1.

at which point the end of the recovery process is indicated with the chamber pressure dropping to below the solidification pressure, solid-xenon vapor pressure, and eventually to ≤ 1 mbar. A typical recovery rate is 10 kg per 3 hours, but can reach as high as 10 kg per hour in cases where the Novec 7000 is pre-warmed prior to recovery.

B.3.3 Recirculation

In both systems, a xenon recirculation pump is used to force flow of xenon gas through a heated NEG purifier to remove electronegative impurities which can capture drifting charges. Both pumps use hermetic volumes in which the only contact with xenon is with stainless steel or teflon. Actuation of a piston draws xenon into the pump inlet and out the pump outlet through one-way valves which set the flow direction.

In the Small System, the pump consists of a large bellows which is driven externally by a pneumatic piston. The Large System uses the pump described in Ref. [241, 242], with a magnetically-coupled piston driven via a stepper motor coupled to a linear translation stage. The controller for the stepper motor can run customized programs for different flow conditions, depending on the pressure, impedance, and flow rate. Flow rate readings in both systems are continually recorded, allowing for a measurement of the time to recirculate the entire mass of xenon contained within the chamber, which is a key parameter in determining equilibrium with impurities. An example of instantaneous mass flow rate over a 5 minute period of time using the Large System xenon pump is shown in Fig. B.6.

Recirculation is aided by a resistive heater at the outlet of each chamber ensuring the evaporation rate is sufficient to maintain the desired flow rate. These heaters are coupled to a copper block surrounding the outlet tubing, all in a vacuum volume to thermally isolate this region from the Novec 7000.

Recirculation is also used for the deployment of internal calibration sources, as described in Section B.5.5. In this method, forced flow of xenon over a radioactive source sweeps radon atoms into the test chamber. Additional ports along the recirculation manifold are flexible enough to allow testing of alternative purifying technologies for cross comparison.




Figure B.6: Mass flow rate measured during recirculation through a heated NEG purifier. Each linear stroke of the pump creates a positive pressure differential, but at the turning points the flow decreases, creating the periodic triangular shape shown above. The pump is turned off for a short period of time to allow the xenon gas at the inlet to condense at high pressures, resulting in a duty cycle that is typically greater than 80% (the off period is shown in shaded red). The average flow rate is calculated by integrating two periods of this cycle (shaded gray).

B.4 Slow Controls & Monitoring

B.4.1 LabVIEW Code

LabVIEW programs provide the main slow controls interface for both systems. Each system has its own National Instruments chassis through which all instrumentation for that system is controlled and read out. This includes:

- Temperatures measured at 5–10 locations covering the test chamber, heat exchanger, and cryostat using T-type thermocouples and Pt100 RTDs

- Pressures measured at the xenon cell inlet, xenon cell outlet, in the Novec 7000 volume, and by a two-stage vacuum gauge in the pumping manifold
- Strain gauge mass measurements for the xenon storage bottle and Novec 7000 storage drum
- Mass flow measurements through the recirculation manifold

The LabVIEW programs also read and write data to and from the cryo chiller via RS232. Commands to turn on and off the cryo chiller, start the precooling sequence, and turn on or off cooling for a specified system can be sent using dedicated buttons on the LabVIEW front panel. The LabVIEW programs continually read the suction and discharge pressures in the cryo chiller along with the temperature at five locations. While the cryo chiller does have an internal interlock to shut down in the event of loss of cooling water flow, an additional flow switch is added externally to allow LabVIEW to redundantly handle and monitor this failure condition.

All data described above is written to a text file at a loop frequency that is typically set to be between 1 and 10 seconds. A Python-based data parsing class is used both for fast analysis of slow controls data by users, and by the monitoring system which regularly checks these parameters against a set of thresholds.

The temperature setpoint, electrode voltages, xenon pump, and data acquisition system are all controllable from the lab computers which can be accessed remotely. The only regular tasks that require the physical presence of an operator involve the operation of manual valves, typically used only during xenon filling and recovery. This ensures that between filling and recovery, the system can be operated entirely remotely.

B.4.2 Monitoring Code

The test facilities described here are designed to allow for operation by as few as 2 people. To enable that level of short-handed operation, a Python-based monitoring system was developed that automates the process of checking critical parameters and alerting operators.

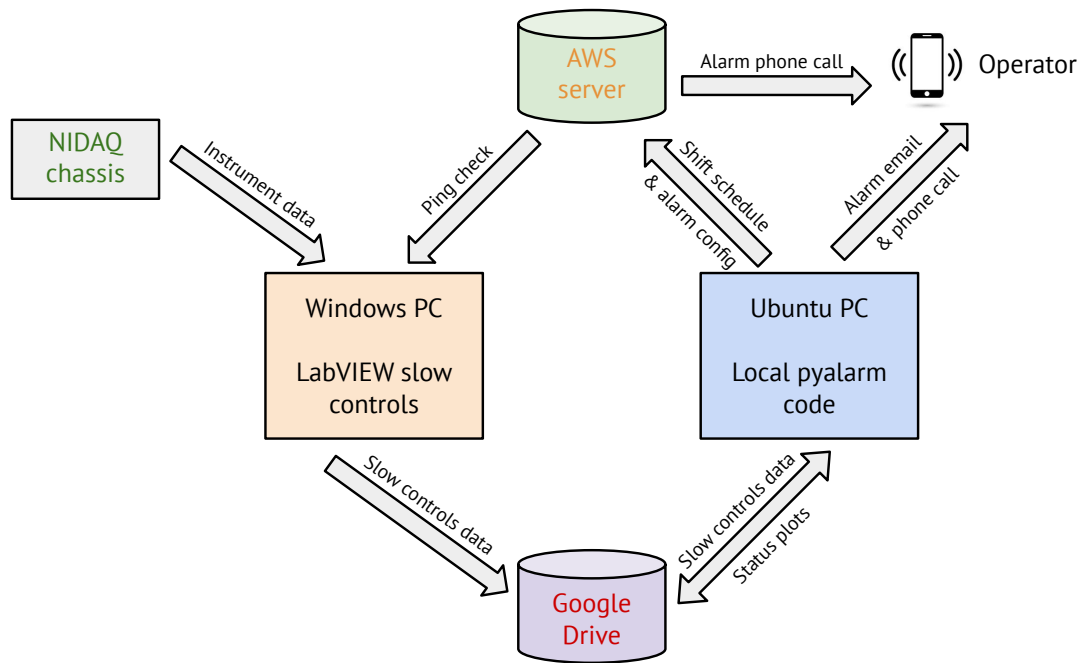


Figure B.7: Functional schematic of the distributed monitoring system.

Fig. B.7 shows a schematic of the monitoring system. The primary monitoring script runs on a desktop computer in the lab, while a second script runs on an off-site server managed by Amazon Web Services. Two configuration files may be updated as the monitoring program operates:

- **Shift schedule:** a YAML file that contains a list of operators and a run coordinator, along with their email addresses, phone numbers, and start/end times for their shift. The monitoring code has been designed to ensure that there is always a designated operator that can be reached by phone.
- **Alarm config file:** a YAML file that allowable ranges for thermocouples, pressures, and any other parameter logged by the LabView system. If a parameter departs from its allowed range, the operator will be alerted by email and with regular phone calls until the issue is resolved.

The main loop of the monitoring code includes the following steps. The shift

schedule and alarm configuration are loaded and synchronized with the remote monitoring server. The operator currently on shift is identified, and if they were not on shift previously, they are alerted with a phone call. Next, the most recent LabVIEW data file is downloaded from a Google Drive folder, the file is parsed, and the parameters are checked against the thresholds in the alarm config file. At the same time, a number of diagnostic plots are produced and uploaded to a Google Drive folder for quick interpretation by the operator. The monitoring loop will also ensure that the slow controls data is up to date and that the LabVIEW code is still writing current parameters as expected. If no errors occur, the script will sleep for the specified time until the next iteration (typically 5 minutes).

Both the shift schedule and the alarm configuration can be modified at any time without interrupting the monitoring code. The entire code is wrapped in a loop with error-catching cases that call the operator upon encountering any unexpected problems to ensure that the monitoring never crashes without notifying.

If an alarm is triggered by the monitoring code, almost all potential problems can be resolved remotely. All computers are remotely accessible, and a series of cameras placed around the lab allow for remote viewing of most hardware and instrumentation. Only errors originating from the cryo chiller, which require a manual reset, and power failures necessitate the physical presence of the operator.

In the event of a power outage in the lab, the remote monitoring server will detect that the local monitoring computer is no longer pingable and alert the operator. Some essential instruments in the lab are powered through an uninterruptible power supply (UPS) to ensure the xenon can be safely recovered without building power.

B.5 Experiments Currently Supported

The test platforms described above are designed to be general-purpose systems capable of supporting a variety of R&D efforts relevant to LXe detectors. A few of the xenon test chambers that are frequently operated within the Stanford xenon facility are described in this section, as well as the supporting electronics and sensor technologies that they exercise.

B.5.1 Test TPCs

Three single-phase test TPCs (Fig. B.8) are frequently used in the Stanford xenon facility. These detectors were designed primarily to test prototype charge and light sensors for the nEXO experiment [194].

All three detectors have an active volume enclosed by a cathode grid at one end, a charge-sensitive anode plane at the other, and copper field-shaping rings spaced out along the length of the cylindrical drift region. The drift region is enclosed in a stainless steel CF spool piece with the charge-sensitive anode mounted to the top flange and a light sensor array mounted to the bottom flange beneath the cathode. The cathode grids are constructed from stainless steel etched hexagonal meshes with $\sim 95\%$ optical transparency. Four Vespel rods screwed to the bottom flange of the test chamber extend vertically upwards around the drift region. Regularly-spaced notches in these rods hold the cathode and field shaping rings. The TPCs and sensor arrays are fully modular; any test TPC can use any charge or light sensor array with the appropriate CF adapters.

Short TPC

The “Short TPC” (STPC) is contained within a CF spool piece between two DN250CF flanges at either end. The drift region is 2.1 cm long, resulting in a drift time of $8\ \mu\text{s}$ at the nominal drift field of $380\ \text{V/cm}$ [186]. The STPC holds $\sim 10\ \text{kg}$ xenon.

Long TPC

Like the STPC, the “Long TPC” (LTPC) has a 19.3 cm diameter drift region between two DN250CF flanges. The LTPC has a significantly longer drift region of 13.8 cm, corresponding to a $\sim 70\ \mu\text{s}$ drift time at $380\ \text{V/cm}$. The drift region is surrounded by five field-shaping rings which keep the electric field uniform to within 80% in the central 55% fiducial volume. It holds $\sim 34\ \text{kg}$ xenon.

Wide TPC

The “Wide TPC” (WTPC) has a 16.1 cm drift length. It is contained within a DN350CF vessel and has a drift region with a diameter of 29.6 cm. With this larger diameter, the WTPC can hold ~ 64 kg xenon. With 7 rings surrounding the drift region, an electric field uniformity of 95% in the central 55% fiducial volume is achieved. This TPC was designed specifically to test arrays of large area sensors which require the largest commercially available CF diameter.




Figure B.8: CAD renderings of the Short TPC (left), the Long TPC (middle) and the Wide TPC (right), at a 1:1 scale with one another.

B.5.2 Data Acquisition

Both test stands have a dedicated data acquisition (DAQ) rack, distinct from the slow controls, which provides flexible triggering and digitization schemes. The most commonly used waveform digitizers are four Struck SIS3316 VME digitizers. Each unit consists of 16 channels with configurable 14-bit ADC dynamic range and sampling

rate from 25 to 250 MS/s, allowing for simultaneous readout of the multi-channel charge and light sensors described in Section B.5.3.

B.5.3 Light & Charge Sensing

Scintillation light is collected by a square array of 24 FBK VUV-HD1 SiPMs grouped together in pairs, providing 12 light-collection channels. The SiPM array measures $6.5\text{ cm} \times 6.5\text{ cm}$. It is mounted to a double-sided DN100CF flange, which can be bolted to the underside of the bottom flange of any test chamber. The bottom flanges of all test chambers have a 10.6 cm bore to accommodate the SiPM tile.

The small area of the SiPM tile compared to the test TPCs results in a low light collection efficiency. For this reason, the light signals are primarily used for triggering but not energy reconstruction. Alternatively, the flange carrying the SiPM array can be replaced with a UV quartz viewport and a UV photomultiplier tube.

A sub-D feedthrough routes SiPM signals from the LXe volume into a dry-nitrogen-filled volume containing preamplifier electronics [292]. This volume is sealed using CF to isolate the electronics from the Novec 7000. The preamplifier outputs are transmitted through feedthroughs to the Novec 7000 space, where RG58 cables transmit to the top flange and air space.

The prototype ionization-charge detection module, or “charge tile”, described in Ref. [188] has been used with both the LTPC and the STPC. In nEXO, charge tiles similar to this prototype will be arrayed to cover the entire anode plane, maximizing charge collection efficiency and single-site/multi-site discrimination [293].

One strength of a mid-scale LXe facility is in its ability to test arrays of prototype sensors, exploring potential unknown impacts on signal quality driven by inter-sensor interactions. The WTPC is specifically designed to explore a 2-by-2 array of $10\text{ cm} \times 10\text{ cm}$ charge-tile modules, requiring a large chamber diameter.

B.5.4 LXe Compatibility Test Chambers

Xenon Purity Monitor

The xenon Purity Monitor (PM) (Fig. B.9, left) consists of a 12 cm drift region, a gold-coated photocathode, an anode, and two Frisch grids. A charge signal is produced by the photoelectric effect using UV light injected through a 600 μm -core polyimide coated fiber in the xenon space. Light is sourced by a 60 W xenon flash lamp, liberating on the order of 100–1000 fC of charge per flash from the photocathode. A Frisch grid 1 cm from both the photocathode and anode isolate them from induction current while charges drift in the central region, and 14 field rings maintain a constant electric field of up to 400 V/cm in the central region.

The electron lifetime is determined by the ratio of charge collected by the anode to charge liberated by the photocathode. Materials and test samples installed into two sample holders outside of the field cage allow for measurements of electronegative out-diffusion of elements to be included in large-scale xenon detectors. The design of the PM is based on a similar device described in Ref. [294].

High-Voltage Test Chamber

The High-Voltage Test Chamber (HVTC) (Fig. B.9, right) is a 10 kg-scale LXe experiment that has observed a variety of high-voltage phenomena (HVP) using multiple pairs polished electrodes with 16 cm^2 field-stressed area oriented in a plane-to-plane geometry with the ability to explore fields up to 60 kV/cm. The emphasis of the experiment is to explore the impact on HVP mitigation by depositing thin films of metals and insulators onto the surfaces of electrodes. A comparison of the performance of electrodes with surfaces of bare stainless steel, platinum, and aluminum-magnesium-fluoride has been explored. A publication detailing this experiment is in progress.

Electrodes specifically shaped to create a large, uniform electric field in the central region are held at a fixed position, separated by 3.3 mm [295]. The stressed effective area of the surfaces of the electrodes is about 16 cm^2 , and within that area the field

is uniform to 3%⁶. The stressed volume, a cylinder of radius equal to the stressed area contour, is 5.7 cm³ with a field uniformity of 3%.

The cathode, attached to a 100 kV-rated high-voltage feedthrough, is biased at a ramp rate of 2 V/s. The anode is connected to a passive integration circuit with a 400 μ s time constant, and observes charge depositions across the gap at an equivalent noise level of 1 pC. Two 2.54 cm diameter photomultiplier tubes⁷ with a 90 degree azimuthal separation view the electrode gap through re-entrant, VUV-transparent vacuum windows. The photomultipliers record scintillation coincident with charge depositions, as well as cosmic rays.

B.5.5 Calibrations

Both systems allow for multiple types of calibration sources to be used, including internal α sources and internal and external γ sources. An overview of the deployment system for each source is outlined below.

Internal Sources

To calibrate the detector response for a particular test chamber, internal sources which generate events uniformly throughout the LXe volume are preferred. Dissolved ²²⁰Rn sources have been used to calibrate large-scale LXe TPCs [296, 297, 206] by releasing α particles in the LXe volume. ¹²⁷Xe has been shown to be a useful calibration source for low-energy events in LXe TPCs [298, 196]. The test stands described here have infrastructure to facilitate the deployment of both of these sources.

The α -decays for calibration originate from an electrodeposited and diffusion-bonded oxide disc of ²²⁸Th from Eckert & Ziegler⁸. The disc is held in a CF enclosed volume equipped with particle filters on either end. The source volume is installed into the xenon recirculation path with a bypass valve and lockable radiation source valves.

⁶Standard deviation over the surface of the electrode

⁷Hamamatsu R8520-406

⁸AF-228-PM at 18 kBq original activity.




Figure B.9: CAD renderings of the purity monitor (left) and high-voltage test chamber (right), at a 1:1 scale to one another.

The decay chain of ^{228}Th includes a decay to ^{220}Rn , which emanates from the surface of the disc into the flowing xenon gas. The ^{220}Rn atoms are carried into the xenon chambers before they decay with a 55 second half life. Following a sufficiently long recirculation period, an equilibrium concentration of the 10.64 hour half life ^{212}Pb is reached in the chamber. Short-lived daughters of ^{212}Pb result in an α -decay of ^{212}Po which may be used for the electron lifetime and energy calibration.

The internal ^{127}Xe calibration source is produced through neutron activation on ^{126}Xe by irradiating a sample cylinder of natural xenon, as described in Ref. [196]. A 5.6 mL intermediate volume between the high-pressure sample cylinder and the

recirculation manifold is used to inject quantities of the activated xenon into the recirculation path. The ^{127}Xe mixes uniformly throughout the TPC, allowing uniform detector illumination with 408 keV electron recoil events which can be used for electron lifetime and energy calibrations. Due to the longer half life of 36.3 days, this calibration source is useful for repeated calibrations without having to frequently re-introduce the active isotope into the TPC.

External Sources

External γ sources are also useful as they allow for rapid deployment and removal. The Large System includes a removable source tube submerged within the Novec 7000 to allow external γ sources to be inserted in close proximity to the test chamber. The source tube is isolated from the Novec 7000 and the air using an UltraTorr fitting inserted into a welded, closed tube. The source, which is attached to the end of a plastic rod, can be positioned anywhere along the 80 cm length of the source tube. A sealed ^{232}Th source is typically used in this system, with calibrations performed on events from the 2614 keV ^{208}Tl γ transition.

Appendix C

Microsphere Properties and Measurement Configurations

The results discussed in Chapters 7 and 8 were obtained from datasets collected with four microspheres. Some datasets were collected without a trapped microsphere. Measured properties of the microspheres used are shown in Table C.1, while the datasets and measurement conditions are summarized in Table C.2.

Table C.1: Properties of the four microspheres with which the data used in this work was collected. Date refers to the date the microsphere was trapped. It is also used as the name of the directory in which all datasets collected with it are saved (see Table C.2 for specific locations). Diameters are the nominal values, verified using SEM imaging of ensembles of microspheres. Masses are measured in the trap; unequal inferred densities between microspheres are attributed to differences in the synthesis process. $|\vec{p}_0|$ is the magnitude of the permanent electric dipole moment and α_m is the polarizability of the microsphere.

Index	Date	Diameter [μm]	Mass [pg]	$ \vec{p}_0 $ [$e \cdot \mu\text{m}$]	α_m [$e \cdot \mu\text{m}/(\text{kV}/\text{m})$]
1	20230303	7.56	225 ± 1	—	—
2	20230614	9.98	543 ± 5	19 ± 9	19.5 ± 0.4
3	20240909	7.56	312 ± 12	19 ± 3	12.4 ± 0.2
4	20241009	9.98	1059 ± 37	38 ± 10	32.6 ± 0.5

There are four types of dataset included here. Gravity datasets are those collected with the attractor oscillating in proximity to the microsphere. Transfer function

calibration (TF cal.) datasets are those in which a frequency comb with 1 Hz spacing from 1 Hz to 700 Hz is applied separately along the x , y , and z axes. Position-sensing calibration (Pos. cal.) datasets are those in which the microsphere undergoes free resonant motion with the vacuum chamber at anywhere from 1 mbar to 2 mbar. In distance ladder (DL) datasets, data is collected with the attractor at a range of separations from the microsphere.

Many of the Gravity datasets included in Table C.2 were not used for any of the figures or the physics result presented in this thesis. Datasets in the first ten rows were selected as candidate datasets to be included in the global analysis, but not all of them produced constraints that were competitive. The constraints shown in Fig. 7.8 were set with Gravity dataset 1 for $\alpha < 0$ and 3 for $\alpha > 0$, both from sphere 4. Datasets from the eleventh row of Table C.2 onward were collected for calibration or diagnostic purposes only and were not used to produce the physics result. More information on the conditions under which each of these datasets were collected can be found on the elog.

Table C.2: The datasets used in this work. The first column refers to the index in Table C.1. Directory gives the location of the data on the dataserer, with the footnote indicating the full path, and `date` referring to the string in column 2 of Table C.1. Position is that of the microsphere, or for datasets collected without a trapped microsphere, the trap focus, relative to the edge of the central gold finger on the attractor. Plane is that in which the electric dipole moment of the microsphere is rotated (at $\sim 10^6$ rad/s) prior to a measurement. V_{bias} is the amplitude of the bias on the attractor relative to the grounded shield, f_{bias} is the frequency of the bias modulation, and t_{int} is the total integration time. The final column indicates the figures in this work produced using data from the given dataset.

Sphere	Type	Directory	Position [μm]	Plane	V_{bias} [mV]	f_{bias} [Hz]	t_{int} [10^3 s]	Figures
2	Gravity	5 ^a	(14.7, 0.6, 0.2)	<i>yz</i>	0	0	100	—
3	Gravity	2 ^a	(9.6, 1.0, -1.3)	<i>yz</i>	0	0	45	—
3	Gravity	5 ^a	(9.6, 1.1, -2.9)	<i>yz</i>	0	0	54	—
3	Gravity	6 ^a	(10.3, 1.1, -3.6)	<i>yz</i>	0	0	45	—
3	Gravity	7 ^a	(10.2, 1.1, -3.1)	<i>yz</i>	-50	0	49.5	—
3	Gravity	8 ^a	(10.3, 1.1, -3.1)	<i>xz</i>	-50	0	49.5	—
3	Gravity	10 ^a	(9.2, 1.1, -3.1)	<i>xz</i>	150	40.5	49.5	—
4	Gravity	1 ^a	(11.0, 1.1, -1.6)	<i>yz</i>	0	0	45	7.8
4	Gravity	3 ^a	(11.0, 1.1, -1.6)	<i>yz</i>	200	46.5	45	7.8, 8.11
4	Gravity	5 ^a	(13.0, 1.1, -1.6)	<i>yz</i>	0	0	45	7.5, 7.6, 8.1, 8.9
4	Gravity	7 ^a	(~ 4000 , 1.1, -1.5)	<i>yz</i>	0	0	45	8.2
3	Pos. cal.	1 ^b	—	—	—	—	0.25	8.5, 8.6b, 8.7
3	TF cal.	1 ^c	—	—	—	—	0.18	8.8
1	DL	2-24 ^d	(14.0 - 27.8, 0, 0)	<i>yz</i>	0	0	3.5 each	8.4
None	DL	1, 3, 5, 7 ^e	(9 - 27, 0, 23)	—	0	0	10 each	8.3
None	DL	1-4 ^f	(11 - 34, 0, 24)	—	0	0	3.5 each	8.3
None	DL	5-8 ^g	(11 - 34, 0, 23)	—	0	0	3.5 each	8.3

^a/data/new_trap/{date}/Bead0/Gravity/{dataset}/

^b/data/new_trap/{date}/Bead0/QPDDiag/{dataset}/

^c/data/new_trap/{date}/Bead0/TransFunc/trapFocus/{dataset}/

^d/data/new_trap/20230330/Shaking/YZ/0mV/ and /data/new_trap/20230330/Shaking/YZ/distanceLadder/{dataset}

^e/data/new_trap/20230218/Backgrounds/ScatteredLight/Shaking/{dataset}

^f/data/new_trap/20230323/scatteredLight/100um/Shaking/{dataset}

^g/data/new_trap/20230327/scatteredLight/50um/Shaking/{dataset}

Appendix D

The OptLevAnalysis Package

The `OptLevAnalysis` package includes Python modules used for analysis of data from the levitated microspheres experiment. The code is hosted on GitHub [299] and the basic functionality is described in this appendix. Only the most important steps and analysis options are included in this overview. Thorough descriptions of the functionality of the code can be found in the relevant docstrings. A high-level summary of the basic analysis pipeline and the roles of selected modules is shown in Fig. D.1.

D.1 The `data_processing` Module

The `data_processing` module contains the two main data analysis classes. The `FileData` class contains methods for loading, filtering, applying calibrations, and reducing individual 10-second data files. The `AggregateData` class is used for analysis of complete datasets.

D.1.1 The `FileData` Class

A `FileData` object is initialized by passing the path to a binary file in HDF5 format. These files are typically stored in subdirectories of `/data/new_trap/` on the dataservert. The `load_data` method is then used to extract the data and populate

the attributes of the object. This data includes the following fields with an entry for each timestep:

- **quad_data**: interleaved data from the QPD, z PD, and FPGA clock. The first five entries are the heterodyne amplitudes for the four QPD quadrants and the z PD, the next five are their heterodyne phases, and the final two make up the timestamp, broken up into two separate 32-bit objects.
- **XYPD**: the x and y positions from the DC QPD or, in older datasets, a position-sensing photodiode.
- **pos_data**: the measured x , y , and z positions reconstructed from QPD data on an FPGA and used for feedback, along with the corresponding feedback signals.
- **cant_data**: the x , y , and z position coordinates of the Aerotech cantilever stage on which the attractor is mounted.
- **laser_power**: the power measured by the input PD used to infer the power entering the vacuum chamber.
- **p_trans**: the power measured by the transmitted PD used to infer the power transmitted through the vacuum chamber.
- **spin_data**: (only recorded for every 10th file) the power measured by the spin-monitoring photodiode. This is sampled at ten times the usual sampling rate in order to capture the microsphere's spin frequency.
- **accel**: the accelerations measured by the three seismic accelerometers fixed to the granite table.
- **microphone**: noise measured by the microphones placed around the lab. Currently, only a single microphone is used, but the code can handle any number of microphone channels.

Other data, including the sampling rate, camera status, and microsphere height are also recorded in each file. After the data is extracted, it is processed as follows:

1. The x , y , and z positions and the timestamps are reconstructed from the `quad_data`, using either the naive or calibrated position-sensing matrices depending on the `diagonalize_qpd` input argument.
2. The data is downsampled by a factor of 20. This step may be skipped depending on the `downsample` input argument.
3. The attractor stage position is calibrated to position units.
4. The DFT of the microsphere position data is computed.
5. The microsphere position data is filtered using Wiener filters constructed to remove coherent noise sensed by the accelerometers or the microphones. This step may be skipped depending on the `wiener` input argument¹.
6. The force response transfer functions are used to calibrate the microsphere response measured in the QPD, XYPD, and zPD from raw ADC counts to force units. This step may be skipped depending on the `no_tf` input argument.
7. A mask is constructed to keep only the DFTs at the harmonics of the attractor drive.
8. If the `signal_model` input argument was provided, signal templates are computed for each of the signal parameters (typically λ values).
9. Depending on the `lightweight` input argument, raw data may be dropped to reduce the size of the object.

The `FileData` class has been designed to be compatible with multiple types of datasets. For processing of many dataset types, the downsampling, filtering, transfer function calibration, and signal template computation steps should be skipped with the appropriate input arguments.

¹The filtering step may be done prior to the computation of the DFTs if the `time_domain` input argument is provided and time-domain filters have been computed.

D.1.2 The AggregateData Class

An `AggregateData` object is initialized with a path or paths to the directory or set of directories containing the HDF5 data files, and optionally:

- File prefixes used to identify the files to be loaded.
- A description used to identify the dataset.
- The number of files to be loaded.
- The index of the first file to be loaded.
- The config files to be used for the data to be loaded.

If the data is to be used for the Yukawa interaction search, the `load_yukawa_model` method should be called to initialize the signal model. This will then allow for signal templates to be computed for each `FileData` object.

The `load_file_data` method is then used to create `FileData` objects for each of the files specified by the arguments passed to the initialization method. This method takes a number of input arguments, most of which are then passed to the individual `FileData` objects to specify options for downsampling, calibration, filtering, etc. It also takes an argument `num_cores` which sets the number of CPU cores used to parallelize the file loading process.

The config file for each dataset contains data on the experimental configuration that is used to compute the signal templates. This includes the microsphere mass, diameter, its position relative to the attractor, and the calibrated position sensing matrix. A config file is created by default for each dataset when the LabVIEW program to begin data collection is launched. The config file is written in the YAML format.

Once `FileData` objects have been constructed for each of the files to be included, the data is extracted and added to the `agg_dict` attribute. This dictionary is the main container for data which can be used by the plotting and statistical analysis modules described below. Depending on the `lightweight` input argument to the

`load_file_data` method, the individual `FileData` objects may either be kept as attributes of the `AggregateData` object, or dropped.

The `AggregateData` class also contains the `diagonalize_qpd` method, which is used to compute the calibrated position-sensing matrix, as described in Section 8.4.1. This method can be applied to an `AggregateData` object containing uncalibrated datasets collected with the microsphere undergoing free resonant motion.

Once data has been loaded into an `AggregateData` object, the object can be saved to HDF5 format using the `save_to_hdf5` method. To load a previously-saved `AggregateData` object, an empty object can be created by initializing an `AggregateData` object with no input arguments, before calling the `load_from_hdf5` method with the path to the HDF5 file specified.

D.2 Other Modules

D.2.1 The plotting Module

Once an `AggregateData` object has been constructed for a dataset or collection of datasets, the `agg_dict` can be passed to functions provided by the `plotting` module to produce a number of diagnostic plots. These include:

- `spectra`: spectra or spectral densities for the measured forces in the QPD, z PD, and XYPD.
- `env_noise`: spectra or spectral densities for the environmental sensors (accelerometers and microphones).
- `spectrogram`: spectrograms showing the time evolution of the spectra measured in any of the QPD, z PD, XYPD, or environmental sensors.
- `time_evolution`: time evolution of the forces in x , y , and z at a number of harmonics as measured by the QPD, z PD, and XYPD.
- `polar_plots`: polar plots showing the complex force responses at a number of harmonics for a specified force component and sensor.

- `mles_vs_time`: time evolution of the MLEs for α at a number of harmonics for the specified force component and sensor.
- `alpha_limit`: the constraints on α as a function of λ for the given dataset or collection of datasets.
- `limit_vs_integration`: the relationship between the integration time and the sensitivity to α in order to determine the point at which the experiment becomes background limited.

The `plotting` module also includes functions to visualize the force response transfer functions, Wiener filters, and position-sensing matrices.

D.2.2 The stats Module

Statistical analysis is performed on `AggregateData` objects by passing the `agg_dict` to functions from the `stats` module. Before doing any statistical analysis, the `fit_alpha_all_files` function must be used to compute the log-likelihood functions for each harmonic, force component, and integration. The calculation is done as described in Appendix E. The result is a `numpy` array of the log-likelihood coefficients a , b , and c parameters for each measurement in a single harmonic, force component, and integration. Since the signal templates depend on the particular λ value used, there are entries for each λ included in the signal model.

Many other functions in the `stats` module take the array of log-likelihood coefficients as an input argument and perform operations to combine the individual log-likelihoods or compute limits with them. These include:

- `combine_likeliheids_over_dim`: takes an array of log-likelihood coefficients and sums the log-likelihood parabolas along a given dimension to produce a combined log-likelihood.
- `get_limit_from_likeliheids`: computes a limit on α at a single λ from one or more log-likelihood functions. The test statistic is computed independently for each set of log-likelihood coefficients provided. If a single set of log-likelihood

coefficients is given (meaning the log-likelihoods have been summed over all dimensions), the calculation is done analytically as described in Appendix E. If multiple are provided, the calculation is done with the sum of test statistics (this is the “best harmonic” approach described in Section 7.4.4 and used in Ref. [266]). This requires computing the limit numerically, as the asymptotic behavior of multiple summed test statistics is no longer described by Wilks’ theorem.

- `get_alpha_vs_lambda`: computes a limit curve for a range of λ by invoking the `get_limit_from_likelihoods` function for each λ .
- `group_likelihoods_by_parameter`: tests all combinations of a specified number of harmonics and chooses those that either minimize the χ^2 or maximize signal power (two harmonic selection approaches discussed in Section 7.4.4).
- `add_systematic_uncertainty`: adds systematic uncertainty to the global log-likelihood function as described in Section 7.4.3 by broadening the parabola.

D.2.3 The signals Module

Signal models are handled by the `signals` module. This module defines the `SignalModel` class, which loads the force grids that have been pre-computed using the finite-element analysis method described in Section 7.3. This class can also be used to compute background templates, such as those originating from electromagnetic interactions with the attractor, for example. The `SignalModel` object is passed to the `FileData` object when each file is loaded, allowing a unique template to be computed based on the attractor stroke in that particular file.

D.2.4 The synthetic_data Module

To test a statistical model, it is useful to inject known signals into noise-only data and ensure that the statistical model reconstructs the signal at the correct level. The `synthetic_data` module contains the `SynthFile` and `SynthAggregate` classes,

subclasses of `FileData` and `AggregateData`, respectively. These subclasses can be used in the same way as their parent classes to load and process data; however, they also accept input arguments specifying the α and λ to be used for injected synthetic signals. The signals are injected into the raw QPD or PSPD data in the time domain. When signals are injected into noise-only datasets, a synthetic attractor stroke has to be added as well so that the template for the given file is generated correctly as if the attractor were actually moving.

D.2.5 External Code

Fitting of the force response transfer functions and the Wiener filters are both done separately using compiled code. For the transfer functions, the code implements a modified version of the method described in Ref. [300] to fit the complex domain data to poles and zeros, from which some physical interpretation can be drawn. The Wiener filters are computed using code which implements the formulae in Refs. [268, 301]. The transfer functions and filters produced using this code are stored in subdirectories of `/data/new_trap_processed/` on the dataserer. The `FileData` code automatically looks for the transfer functions and filters in standard locations when processing a file.

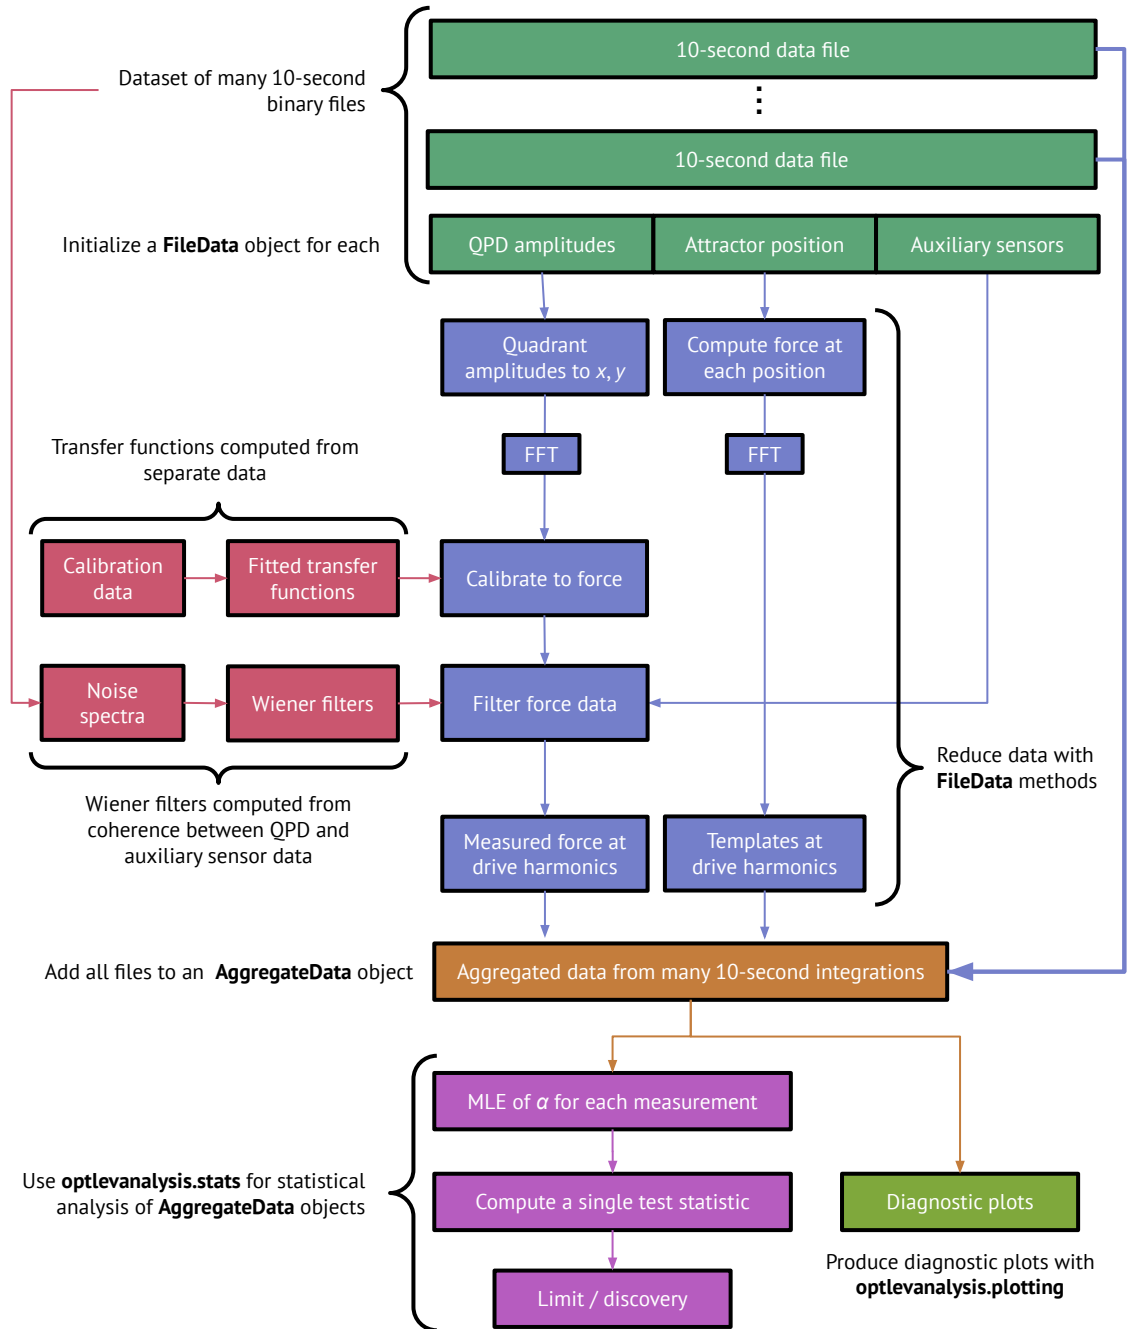


Figure D.1: Flow chart showing the analysis pipeline for the micron-scale interaction search.

Appendix E

Calculation of the Exclusion Limits on Yukawa Interactions

This appendix describes explicitly how the statistical model introduced in Section 7.4 can be used to place constraints on α as a function of λ given a set of measured complex amplitudes, F_{ijk} , and their corresponding statistical uncertainties, σ_{ijk} .

E.1 Computing the Log-Likelihood Functions

To minimize clutter in the following calculation, I do not explicitly write out the dependence on α and λ except where it is necessary to avoid ambiguity. We start from the definition of an individual likelihood function in Eq. (7.6) and take the logarithm to obtain

$$\log \mathcal{L}_{ijk} = -\log(2\pi\sigma_{ijk}^2) - \frac{[\Re(F_{ijk} - \alpha \tau_{ijk})]^2}{2\sigma_{ijk}^2} - \frac{[\Im(F_{ijk} - \alpha \tau_{ijk})]^2}{2\sigma_{ijk}^2}. \quad (\text{E.1})$$

By expanding the squared terms and factoring out α and α^2 where necessary, we

can rewrite the equation for this parabolic function in standard form as

$$\begin{aligned} \log \mathcal{L}_{ijk} = & - \left(\frac{\Re(\tau_{ijk})^2 + \Im(\tau_{ijk})^2}{2\sigma_{ijk}^2} \right) \alpha^2 \\ & - \left(\frac{\Re(F_{ijk})\Re(\tau_{ijk}) + \Im(F_{ijk})\Im(\tau_{ijk})}{\sigma_{ijk}^2} \right) \alpha \\ & - \left(\frac{\Re(F_{ijk})^2 + \Im(F_{ijk})^2}{2\sigma_{ijk}^2} \right) \\ & - \log(2\pi\sigma_{ijk}^2). \end{aligned} \quad (\text{E.2})$$

We then make the following definitions:

$$a_{ijk} \equiv \frac{\Re(\tau_{ijk})^2 + \Im(\tau_{ijk})^2}{2\sigma_{ijk}^2}, \quad (\text{E.3})$$

$$b_{ijk} \equiv \frac{\Re(F_{ijk})\Re(\tau_{ijk}) + \Im(F_{ijk})\Im(\tau_{ijk})}{\sigma_{ijk}^2}, \text{ and} \quad (\text{E.4})$$

$$c_{ijk} \equiv \frac{\Re(F_{ijk})^2 + \Im(F_{ijk})^2}{2\sigma_{ijk}^2}, \quad (\text{E.5})$$

where a_{ijk} , b_{ijk} , and c_{ijk} are directly calculable from the data. At this stage, one can use these coefficients of the parabola to define $\hat{\alpha}_{ijk}$, the MLE for the particular measurement indexed by i , j , and k , as

$$\hat{\alpha}_{ijk} = -\frac{b_{ijk}}{2a_{ijk}}. \quad (\text{E.6})$$

To calculate $-2 \log \Lambda$, we note that the triple product in Eq. (7.7) has become a triple sum,

$$-2 \log \Lambda = -2 \sum_i \sum_j \sum_k [\log \mathcal{L}_{ijk}(\alpha|\lambda) - \log \mathcal{L}_{ijk}(\hat{\alpha}|\lambda)], \quad (\text{E.7})$$

which we can then rewrite in terms of the newly-defined parameters as

$$-2 \log \Lambda = -2 \sum_i \sum_j \sum_k (a_{ijk} \alpha^2 + b_{ijk} \alpha + c_{ijk} - a_{ijk} \hat{\alpha}^2 - b_{ijk} \hat{\alpha} - c_{ijk}) \quad (\text{E.8})$$

$$= 2 (\alpha^2 - \hat{\alpha}^2) \sum_i \sum_j \sum_k a_{ijk} + 2 (\alpha - \hat{\alpha}) \sum_i \sum_j \sum_k b_{ijk}. \quad (\text{E.9})$$

We can group terms and make the following definitions,

$$a \equiv 2 \sum_i \sum_j \sum_k a_{ijk} \quad (\text{E.10})$$

$$b \equiv 2 \sum_i \sum_j \sum_k b_{ijk} \quad (\text{E.11})$$

which results in a much cleaner equation for $-2 \log \Lambda$,

$$-2 \log \Lambda = a (\alpha^2 - \hat{\alpha}^2) + b (\alpha - \hat{\alpha}). \quad (\text{E.12})$$

By definition, $\hat{\alpha}$ is the value of α which minimizes this quantity. Using this fact, one can relate a , b , and $\hat{\alpha}$ as

$$b = -2a\hat{\alpha}, \quad (\text{E.13})$$

which, when plugged back in to Eq. (E.12), gives

$$-2 \log \Lambda = a (\alpha - \hat{\alpha})^2. \quad (\text{E.14})$$

By comparing this equation to the log-likelihood for a Gaussian-distributed quantity, we can identify a as relating to the statistical uncertainty, and make the definition

$$\sigma_{\text{stat}} = \frac{1}{\sqrt{a}}. \quad (\text{E.15})$$

This can then be added in quadrature with σ_{sys} , calculated as described in Section 7.4.3, to give the overall uncertainty, $\sigma_{\hat{\alpha}}$. To recap, we started with a series of complex force amplitude measurements, F_{ijk} , and the corresponding force templates, τ_{ijk} , and showed how $-2 \log \Lambda$ can be calculated directly from that data.

E.2 Limits from the Test Statistic

For any value of α , the observed value of q_α , which we denote as $q_{\alpha,\text{obs}}$, can be calculated from the data using Eq. (7.11). We can then compare $q_{\alpha,\text{obs}}$ to the distribution of q_α under the null hypothesis to quantify the incompatibility between the data and the null hypothesis at each value of α . To be quantitative, we can define the p -value, p_α , as the probability under the null hypothesis of obtaining a value of q_α of greater than $q_{\alpha,\text{obs}}$ for a particular value of α :

$$p_\alpha = \int_{q_{\alpha,\text{obs}}}^{\infty} f(q_\alpha|\alpha) dq_\alpha = 1 - F(q_\alpha|\alpha) \quad (\text{E.16})$$

where $f(q_\alpha|\alpha)$ is the distribution of q_α under the null hypothesis and F is its cumulative distribution function (CDF). Wilks' theorem tells us that $f(q_\alpha|\alpha)$ is a half- χ^2 distribution with one degree of freedom,

$$f(q_\alpha|\alpha) = \frac{1}{2}\delta(q_\alpha) + \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{q_\alpha}} e^{-q_\alpha/2}, \quad (\text{E.17})$$

which has a CDF given by

$$F(q_\alpha|\alpha) = \Phi(\sqrt{q_\alpha}), \quad (\text{E.18})$$

where Φ is the Gaussian CDF. Therefore, we get that

$$p_\alpha = 1 - \Phi(\sqrt{q_\alpha}). \quad (\text{E.19})$$

Solving for q_α gives the value of the test statistic corresponding to an exclusion at a confidence level of $1 - p_\alpha$:

$$q_\alpha = (\Phi^{-1}(1 - p_\alpha))^2. \quad (\text{E.20})$$

Plugging in the definition of q_α in terms of $\hat{\alpha}$ and $\sigma_{\hat{\alpha}}$, we find

$$\frac{(\alpha - \hat{\alpha})^2}{\sigma_{\hat{\alpha}}^2} = (\Phi^{-1}(1 - p_\alpha))^2, \quad (\text{E.21})$$

which tells us that the critical value of α for a 95% CL limit, α_{95} , is

$$\alpha_{95} = \hat{\alpha} + \text{sgn}(\hat{\alpha}) \cdot \sigma_{\hat{\alpha}} \Phi^{-1}(1 - p_{\alpha}), \quad (\text{E.22})$$

where $\text{sgn}(\hat{\alpha})$ ensures that the limit is only set on the sign of α corresponding to the sign preferred by the data. Our chosen exclusion confidence level is 95%, so $p_{\alpha} = 0.05$, and the equation used to set the constraints becomes

$$\alpha_{95} = \hat{\alpha} + \text{sgn}(\hat{\alpha}) \cdot 1.64 \sigma_{\hat{\alpha}}. \quad (\text{E.23})$$

Repeating this calculation of α_{95} using the selected harmonics for each value of λ gives the 95%CL limit shown in Fig. 7.8.

Appendix F

The Physics of Contact Potentials

A potential difference can arise between the surfaces of two metals in electrical contact due entirely to their surface properties, creating an electric field in their vicinity [302]. This appendix provides a brief overview of why this occurs.

In what follows, zero potential is taken to be at infinity. The electrostatic potential at two points can then be defined: φ is the potential in the bulk phase of a metal, called the inner potential, while ψ is the potential just outside the surface of the metal, called the outer potential. The outer potential describes the work required to bring a unit charge to just outside the metal's surface, before any image charge effects become significant. At the surface, the distribution of electrons extends partially into the vacuum, creating a dipole layer separating the vacuum from the bulk. We call the resulting potential difference the surface potential,

$$\chi \equiv \varphi - \psi. \tag{F.1}$$

When two metals are isolated from one another, there can be an arbitrary potential difference between them. On being brought into contact, electrons will flow between them to equalize the Fermi level everywhere. The Fermi level, E_F , is given by

$$E_F = \mu + e\varphi \tag{F.2}$$

where μ is the chemical potential for an electron, φ is the electrostatic potential, and

e is the negative elementary charge. Equating the Fermi levels in two phases gives the electrostatic potential difference arising from the difference in chemical potential of the two species:

$$\varphi_2 - \varphi_1 = \frac{\mu_1 - \mu_2}{e}. \quad (\text{F.3})$$

This potential difference between the bulk phases of two metals is known as the Galvani potential. It results in a charge double layer forming at the metal interface, yet the compensating difference in chemical potential means that no energy is required or gained for an electron crossing from one phase into the other. As it does not result in current flow, it also means that the Galvani potential is unmeasurable with a voltmeter.

By contrast, energy is required for an electron to be brought from the Fermi level across the dipole layer at the metal-vacuum interface. The work required is called the work function, denoted by Φ , and calculated as,

$$\Phi = e\psi - E_F. \quad (\text{F.4})$$

Now that we have defined the work required to cross the barrier between each of the metals and vacuum, and established that no work is required to cross the metal-metal interface, we can determine the potential difference that exists between the outer surfaces of both metals, which we denote V_c . The net work to bring an electron from material 1, to vacuum, to material 2, and back to material 1 should be equal to zero:

$$0 = \Phi_1 + eV_c - \Phi_2. \quad (\text{F.5})$$

V_c is therefore set by the difference in work function between the two metals,

$$V_c = \frac{\Phi_2 - \Phi_1}{e}. \quad (\text{F.6})$$

V_c is known as the Volta potential, or the “contact potential.” The name is misleading; the contact potential has nothing to do with the contact or connection between the metals, provided there exists some conductive path between them. It is an effect arising from the surface properties of the two metals. The work function

difference gives rise to a potential difference and a corresponding electric field in the vacuum between the metals. Fig. F.1 summarizes the relationships between φ , ψ , χ , μ , E_F , and Φ , and shows how when isolated metals are brought into contact, a contact potential can be produced.




Figure F.1: Schematic showing how the contact potential difference arises. When two metals are isolated, their Fermi levels will in general not be the same. When brought into contact, electrons flow to keep the Fermi levels the same, and as a result, the difference in work function creates a potential difference between the metal surfaces.

If a bias V_{bias} is applied between the two metals, the Fermi level will differ by eV_{bias} . The potential difference between the two exterior surfaces then becomes,

$$V_{21} = \frac{\Phi_2 - \Phi_1}{e} + V_{\text{bias}}. \quad (\text{F.7})$$

Applying a bias of $\Phi_1 - \Phi_2$ therefore cancels the potential difference due to the contact potential and eliminates the resulting electric field between the surfaces, as shown in Fig. F.2.

Note that like the Galvani potential, the contact potential cannot be measured using a voltmeter. A voltmeter measures currents driven by a difference in the Fermi level at the two terminals. Since the Fermi level is the same throughout the system, no

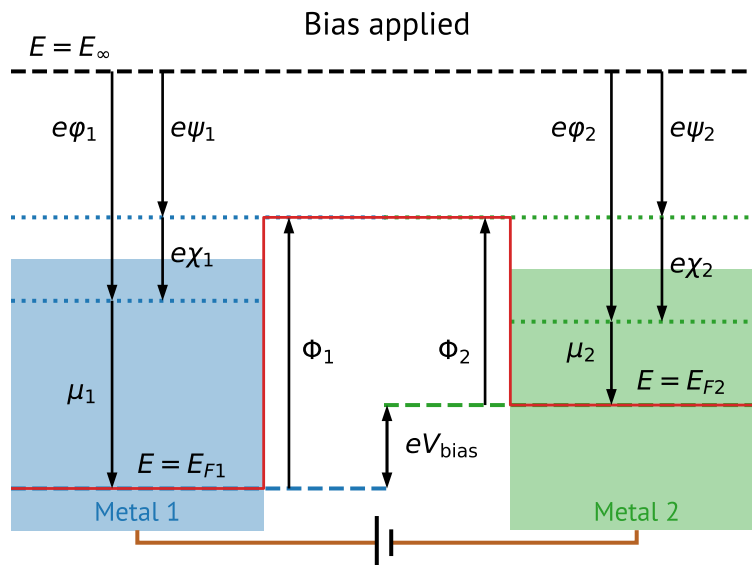


Figure F.2: When a voltage V_{bias} is applied, the Fermi levels of the two metals differ by eV_{bias} . Tuning the bias voltage to $-V_c$ can eliminate the potential difference between the surfaces arising from the contact potential.

currents will arise when voltmeter probes are attached. The contact points between the metals and the probes will simply cause new charge double layers to form at the interfaces to compensate the differing chemical potentials, while the Fermi level will remain unchanged.

Unlike the Galvani potential, the contact potential creates a real electric field between two surfaces; this can be determined from measurements using Kelvin probe force microscopy (KPFM) [303, 304]. The principle behind this measurement is that variations in distance between two surfaces that are electrically connected changes their capacitance and causes the flow of charge. KPFM can be used to measure the work function of a metal on small length scales by modulating the bias between a probe and the metal. This changing force on the probe causes a vibration at the modulation frequency proportional to the potential difference between the surfaces. Varying the DC bias until the vibration is minimized allows the contact potential between the surface and the probe to be determined. Repeating this procedure for two different materials allows the work function difference between them to be determined.

References

- [1] C. Burgess and G. Moore, *The Standard Model: a primer* (Cambridge University Press, Cambridge, 2006).
- [2] M. Srednicki, *Quantum field theory* (Cambridge University Press, Cambridge, 2007).
- [3] M. Gell-Mann, “The interpretation of the new particles as displaced charge multiplets”, *Il Nuovo Cimento* (1955-1965) **4**, 848 (1956).
- [4] P. A. M. Dirac, “The quantum theory of the electron”, *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* **117**, 610 (1928).
- [5] E. Majorana, “Teoria simmetrica dell’elettrone e del positrone”, *Il Nuovo Cimento* (1924-1942) **14**, 171 (1937).
- [6] L. M. Brown, “The idea of the neutrino”, *Physics Today* **31**, 23 (1978).
- [7] F. Reines and C. L. Cowan, “Detection of the free neutrino”, *Physical Review* **92**, 830 (1953).
- [8] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, “Detection of the free neutrino: a confirmation”, *Science* **124**, 103 (1956).
- [9] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, “Experimental test of parity conservation in beta decay”, *Physical Review* **105**, 1413 (1957).

- [10] T. D. Lee and C. N. Yang, “Question of parity conservation in weak interactions”, *Physical Review* **104**, 254 (1956).
- [11] M. Goldhaber, L. Grodzins, and A. W. Sunyar, “Helicity of neutrinos”, *Physical Review* **109**, 1015 (1958).
- [12] J. N. Bahcall, “Solar neutrinos. I. Theoretical”, *Physical Review Letters* **12**, 300 (1964).
- [13] R. Davis, “Solar neutrinos. II. Experimental”, *Physical Review Letters* **12**, 303 (1964).
- [14] R. Davis, D. S. Harmer, and K. C. Hoffman, “Search for neutrinos from the Sun”, *Physical Review Letters* **20**, 1205 (1968).
- [15] J. N. Bahcall and R. L. Sears, “Solar neutrinos”, *Annual Review of Astronomy and Astrophysics* **10**, 25 (1972).
- [16] V. Trimble and F. Reines, “The solar neutrino problem—a progress(?) report”, *Reviews of Modern Physics* **45**, 1 (1973).
- [17] J. C. Evans, R. Davis, and J. N. Bahcall, “Brookhaven solar neutrino detector and collapsing stars”, *Nature* **251**, 486 (1974).
- [18] B. Pontecorvo, “Mesonium and antimesonium”, *Soviet Physics JETP* **6**, 429 (1957).
- [19] Q. R. Ahmad *et al.* (SNO Collaboration), “Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory”, *Physical Review Letters* **89**, 011301 (2002).
- [20] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), “Evidence for oscillation of atmospheric neutrinos”, *Physical Review Letters* **81**, 1562 (1998).
- [21] Z. Maki, M. Nakagawa, and S. Sakata, “Remarks on the unified model of elementary particles”, *Progress of Theoretical Physics* **28**, 870 (1962).

- [22] S. M. Bilenky and B. Pontecorvo, “Lepton mixing and neutrino oscillations”, *Physics Reports* **41**, 225 (1978).
- [23] L. Wolfenstein, “Neutrino oscillations in matter”, *Physical Review D* **17**, 2369 (1978).
- [24] L. Wolfenstein, “Effect of matter on neutrino oscillations”, in *Proceedings of Neutrino '78* (Purdue University, 1978) pp. C3–C6.
- [25] L. Wolfenstein, “Neutrino oscillations and stellar collapse”, *Physical Review D* **20**, 2634 (1979).
- [26] S. P. Mikheyev and A. Yu. Smirnov, “Resonance enhancement of oscillations in matter and solar neutrino spectroscopy”, *Soviet Journal of Nuclear Physics* **42**, 913 (1985).
- [27] S. P. Mikheyev and A. Yu. Smirnov, “Resonant amplification of ν oscillations in matter and solar-neutrino spectroscopy”, *Il Nuovo Cimento C* **9**, 17 (1986).
- [28] S. P. Mikheev and A. Yu. Smirnov, “Neutrino oscillations in a variable density medium and neutrino bursts due to the gravitational collapse of stars”, *Soviet Physics JETP* **64**, 4 (1986).
- [29] A. Yu. Smirnov and S. P. Mikheev, “Neutrino oscillations in matter with varying density”, in *6th Moriond Workshop: Massive Neutrinos in Particle Physics and Astrophysics* (1986) pp. 355–372.
- [30] B. T. Cleveland, T. Daily, J. Raymond Davis, J. R. Distel, K. Lande, C. K. Lee, P. S. Wildenhain, and J. Ullman, “Measurement of the solar electron neutrino flux with the Homestake chlorine detector”, *The Astrophysical Journal* **496**, 505 (1998).
- [31] F. Kaether, W. Hampel, G. Heusser, J. Kiko, and T. Kirsten, “Reanalysis of the Gallex solar neutrino flux and source experiments”, *Physics Letters B* **685**, 47 (2010).

- [32] M. Altmann *et al.*, “Complete results for five years of GNO solar neutrino observations”, *Physics Letters B* **616**, 174 (2005).
- [33] J. N. Abdurashitov *et al.* (SAGE Collaboration), “Measurement of the solar neutrino capture rate with gallium metal. III. Results for the 2002–2007 data-taking period”, *Physical Review C* **80**, 015807 (2009).
- [34] H. H. Chen, “Direct approach to resolve the solar-neutrino problem”, *Physical Review Letters* **55**, 1534 (1985).
- [35] B. Aharmim *et al.* (SNO Collaboration), “Combined analysis of all three phases of solar neutrino data from the Sudbury Neutrino Observatory”, *Physical Review C* **88**, 025501 (2013).
- [36] K. Abe *et al.* (Super-Kamiokande Collaboration), “Solar neutrino measurements using the full data period of Super-Kamiokande-IV”, *Physical Review D* **109**, 092001 (2024).
- [37] G. Bellini *et al.* (Borexino Collaboration), “Measurement of the solar ^8B neutrino rate with a liquid scintillator target and 3 MeV energy threshold in the Borexino detector”, *Physical Review D* **82**, 033006 (2010).
- [38] M. Agostini *et al.* (Borexino Collaboration), “Comprehensive measurement of pp-chain solar neutrinos”, *Nature* **562**, 505 (2018).
- [39] T. Wester *et al.* (Super-Kamiokande Collaboration), “Atmospheric neutrino oscillation analysis with neutron tagging and an expanded fiducial volume in Super-Kamiokande I–V”, *Physical Review D* **109**, 072014 (2024).
- [40] M. G. Aartsen *et al.* (IceCube Collaboration), “Measurement of atmospheric tau neutrino appearance with IceCube DeepCore”, *Physical Review D* **99**, 032007 (2019).
- [41] R. Abbasi *et al.* (IceCube Collaboration), “Measurement of atmospheric neutrino oscillation parameters using convolutional neural networks with 9.3 years of data in IceCube DeepCore”, *Physical Review Letters* **134**, 091801 (2025).

- [42] Y. Abe *et al.* (Double Chooz Collaboration), “Indication of reactor $\bar{\nu}_e$ disappearance in the Double Chooz experiment”, *Physical Review Letters* **108**, 131801 (2012).
- [43] J. K. Ahn *et al.* (RENO Collaboration), “Observation of reactor electron antineutrinos disappearance in the RENO experiment”, *Physical Review Letters* **108**, 191802 (2012).
- [44] F. P. An *et al.*, “Observation of electron-antineutrino disappearance at Daya Bay”, *Physical Review Letters* **108**, 171803 (2012).
- [45] A. Gando *et al.* (KamLAND Collaboration), “Reactor on-off antineutrino measurement with KamLAND”, *Physical Review D* **88**, 033001 (2013).
- [46] P. Adamson *et al.* (MINOS Collaboration), “Measurement of neutrino and antineutrino oscillations using beam and atmospheric data in MINOS”, *Physical Review Letters* **110**, 251801 (2013).
- [47] P. Adamson *et al.* (MINOS Collaboration), “Electron neutrino and antineutrino appearance in the full MINOS data sample”, *Physical Review Letters* **110**, 171801 (2013).
- [48] M. A. Acero *et al.* (NOvA Collaboration), “Improved measurement of neutrino oscillation parameters by the NOvA experiment”, *Physical Review D* **106**, 032004 (2022).
- [49] K. Abe *et al.*, “Measurements of neutrino oscillation parameters from the T2K experiment using 3.6×10^{21} protons on target”, *The European Physical Journal C* **83**, 782 (2023).
- [50] “NuFIT 6.0” (2024), <http://www.nu-fit.org/>, (Accessed: April 2025).
- [51] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. P. Pinheiro, and T. Schwetz, “NuFit-6.0: updated global analysis of three-flavor neutrino oscillations”, *Journal of High Energy Physics* **2024**, 216 (2024).

- [52] M. Aker *et al.* (KATRIN Collaboration), “Direct neutrino-mass measurement based on 259 days of KATRIN data”, *Science* **388**, 180 (2025).
- [53] M. Aker *et al.*, “The design, construction, and commissioning of the KATRIN experiment”, *Journal of Instrumentation* **16**, T08015.
- [54] A. Ashtari Esfahani *et al.* (Project 8 Collaboration), “Cyclotron radiation emission spectroscopy of electrons from tritium β decay and $^{83\text{m}}\text{Kr}$ internal conversion”, *Physical Review C* **109**, 035503 (2024).
- [55] A. Ashtari Esfahani *et al.*, “Determining the neutrino mass with cyclotron radiation emission spectroscopy—Project 8”, *Journal of Physics G: Nuclear and Particle Physics* **44**, 054004 (2017).
- [56] Z. Sakr, “A short review on the latest neutrinos mass and number constraints from cosmological observables”, *Universe* **8**, 284 (2022).
- [57] N. Aghanim *et al.* (Planck Collaboration), “Planck 2018 results - V. CMB power spectra and likelihoods”, *Astronomy & Astrophysics* **641**, A5 (2020).
- [58] N. Aghanim *et al.*, “Planck 2018 results - VI. cosmological parameters”, *Astronomy & Astrophysics* **641**, A6 (2020).
- [59] F. J. Qu *et al.*, “The Atacama Cosmology Telescope: a measurement of the DR6 CMB lensing power spectrum and its implications for structure growth”, *The Astrophysical Journal* **962**, 112 (2024).
- [60] M. S. Madhavacheril *et al.*, “The Atacama Cosmology Telescope: DR6 gravitational lensing map and cosmological parameters”, *The Astrophysical Journal* **962**, 113 (2024).
- [61] A. Adame *et al.*, “DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations”, *Journal of Cosmology and Astroparticle Physics* **2025**, 021.

- [62] J.-Q. Jiang, W. Giarè, S. Gariazzo, M. G. Dainotti, E. Di Valentino, O. Mena, D. Pedrotti, S. Santos da Costa, and S. Vagnozzi, “Neutrino cosmology after DESI: tightest mass upper limits, preference for the normal ordering, and tension with terrestrial observations”, *Journal of Cosmology and Astroparticle Physics* **2025**, 153.
- [63] L. Herold and M. Kamionkowski, “Revisiting the impact of neutrino mass hierarchies on neutrino mass constraints in light of recent DESI data”, *Physical Review D* **111**, 083518 (2025).
- [64] S. Weinberg, “Baryon- and lepton-nonconserving processes”, *Physical Review Letters* **43**, 1566 (1979).
- [65] M. Magg and Ch. Wetterich, “Neutrino mass problem and gauge hierarchy”, *Physics Letters B* **94**, 61 (1980).
- [66] P. Minkowski, “ $\mu \rightarrow e\gamma$ at a rate of one out of 10^9 muon decays?”, *Physics Letters B* **67**, 421 (1977).
- [67] M. Gell-Mann, P. Ramond, and R. Slansky, “Complex spinors and unified theories”, in *Supergravity: Proceedings of the Supergravity Workshop at Stony Brook, 27-29 September 1979*, edited by Peter Van Nieuwenhuizen, Daniel Z. Freedman (North-Holland Publishing Company, 1979) pp. 315–321, arXiv:1306.4669.
- [68] T. Yanagida, “Horizontal symmetry and masses of neutrinos”, *Progress of Theoretical Physics* **64**, 1103 (1980).
- [69] M. Goeppert-Mayer, “Double beta-disintegration”, *Physical Review* **48**, 512 (1935).
- [70] A. Barabash, “Precise half-life values for two-neutrino double- β decay: 2020 review”, *Universe* **6**, 159 (2020).
- [71] W. Huang, M. Wang, F. Kondev, G. Audi, and S. Naimi, “The AME 2020 atomic mass evaluation (I). Evaluation of input data, and adjustment procedures”, *Chinese Physics C* **45**, 030002 (2021).

- [72] M. Wang, W. Huang, F. Kondev, G. Audi, and S. Naimi, “The AME 2020 atomic mass evaluation (II). Tables, graphs and references”, *Chinese Physics C* **45**, 030003 (2021).
- [73] W. H. Furry, “On transition probabilities in double beta-disintegration”, *Physical Review* **56**, 1184 (1939).
- [74] R. N. Mohapatra and J. D. Vergados, “New contribution to neutrinoless double beta decay in gauge models”, *Physical Review Letters* **47**, 1713 (1981).
- [75] W. Rodejohann, “Neutrino-less double beta decay and particle physics”, *International Journal of Modern Physics E* **20**, 1833 (2011).
- [76] M. Duerr, M. Lindner, and A. Merle, “On the quantitative impact of the Schechter-Valle theorem”, *Journal of High Energy Physics* **2011**, 91 (2011).
- [77] J. Schechter and J. W. F. Valle, “Neutrinoless double- β decay in $SU(2)\times U(1)$ theories”, *Physical Review D* **25**, 2951 (1982).
- [78] J. T. Suhonen, “Value of the axial-vector coupling strength in β and $\beta\beta$ decays: a review”, *Frontiers in Physics* **5**, Article 55 (2017).
- [79] J. Kotila and F. Iachello, “Phase-space factors for double- β decay”, *Physical Review C* **85**, 034316 (2012).
- [80] S. Stoica and M. Mirea, “New calculations for phase space factors involved in double- β decay”, *Physical Review C* **88**, 037303 (2013).
- [81] S. Stoica and M. Mirea, “Phase space factors for double-beta decays”, *Frontiers in Physics* **7**, Article 12 (2019).
- [82] J. Engel and J. Menéndez, “Status and future of nuclear matrix elements for neutrinoless double-beta decay: a review”, *Reports on Progress in Physics* **80**, 046301 (2017).
- [83] V. I. Tretyak and Y. G. Zdesenko, “Tables of double beta decay data—an update”, *Atomic Data and Nuclear Data Tables* **80**, 83 (2002).

- [84] H. Primakoff and S. P. Rosen, “Double beta decay”, *Reports on Progress in Physics* **22**, 121 (1959).
- [85] M. Berglund and M. E. Wieser, “Isotopic compositions of the elements 2009 (IUPAC technical report)”, *Pure and Applied Chemistry* **83**, 397 (2011).
- [86] M.-M. Bé *et al.*, *Table of radionuclides*, Monographie BIPM-5, Vol. 3 (Bureau International des Poids et Mesures, Sèvres, France, 2006).
- [87] M.-M. Bé *et al.*, *Table of radionuclides*, Monographie BIPM-5, Vol. 5 (Bureau International des Poids et Mesures, Sèvres, France, 2010).
- [88] M. J. Dolinski, A. W. P. Poon, and W. Rodejohann, “Neutrinoless double-beta decay: status and prospects”, *Annual Review of Nuclear and Particle Science* **69**, 219 (2019).
- [89] M. Agostini, G. Benato, J. A. Detwiler, J. Menéndez, and F. Vissani, “Toward the discovery of matter creation with neutrinoless $\beta\beta$ decay”, *Reviews of Modern Physics* **95**, 025002 (2023).
- [90] J. J. Gómez-Cadenas, J. Martín-Albo, J. Menéndez, M. Mezzetto, F. Monrabal, and M. Sorel, “The search for neutrinoless double-beta decay”, *La Rivista del Nuovo Cimento* **46**, 619 (2023).
- [91] K.-H. Ackermann *et al.*, “The GERDA experiment for the search of $0\nu\beta\beta$ decay in ^{76}Ge ”, *The European Physical Journal C* **73**, 2330 (2013).
- [92] C. Aalseth *et al.* (Majorana Collaboration), “Search for neutrinoless double- β decay in ^{76}Ge with the MAJORANA DEMONSTRATOR”, *Physical Review Letters* **120**, 132502 (2018).
- [93] R. Brugnera (LEGEND Collaboration), “Neutrinoless double-beta decay search with the LEGEND experiment”, *Journal of Instrumentation* **20**, C03050.
- [94] N. Abgrall *et al.* (LEGEND Collaboration), “LEGEND-1000 preconceptual design report” (2021), arXiv:2107.11462 [physics].

- [95] D. R. Artusa *et al.*, “Searching for neutrinoless double-beta decay of ^{130}Te with CUORE”, *Advances in High Energy Physics* **2015**, 879871 (2015).
- [96] K. Alfonso *et al.*, “CUPID: the next-generation neutrinoless double beta decay experiment”, *Journal of Low Temperature Physics* **211**, 375 (2023).
- [97] V. Alenkov *et al.*, “First results from the AMoRE-Pilot neutrinoless double beta decay experiment”, *The European Physical Journal C* **79**, 791 (2019).
- [98] C. Kraus and S. J. M. Peeters, “The rich neutrino programme of the SNO+ experiment”, *Progress in Particle and Nuclear Physics Neutrinos in Cosmology*, in *Astro, Particle and Nuclear Physics*, **64**, 273 (2010).
- [99] J. Hartnell (SNO+ Collaboration), “Neutrinoless double beta decay with SNO+”, *Journal of Physics: Conference Series* **375**, 042015 (2012).
- [100] S. Abe *et al.* (KamLAND-Zen Collaboration), “Search for the majorana nature of neutrinos in the inverted mass ordering region with KamLAND-Zen”, *Physical Review Letters* **130**, 051801 (2023).
- [101] R. Nakamura, H. Sambonsugi, K. Shiraishi, and Y. Wada, “Research and development toward KamLAND2-Zen”, *Journal of Physics: Conference Series* **1468**, 012256 (2020).
- [102] S. Ajimura *et al.* (CANDLES Collaboration), “Low background measurement in CANDLES-III for studying the neutrinoless double beta decay of ^{48}Ca ”, *Physical Review D* **103**, 092008 (2021).
- [103] R. Arnold *et al.* (NEMO-3 Collaboration), “Results of the search for neutrinoless double- β decay in ^{100}Mo with the NEMO-3 experiment”, *Physical Review D* **92**, 072011 (2015).
- [104] V. Álvarez *et al.* (NEXT Collaboration), “The NEXT-100 experiment for neutrinoless double beta decay searches (conceptual design report)” (2011), arXiv:1106.3630 [physics].

- [105] P. Novella *et al.* (NEXT Collaboration), “Demonstration of neutrinoless double beta decay searches in gaseous xenon with NEXT”, *Journal of High Energy Physics* **2023**, 190 (2023).
- [106] G. Anton *et al.* (EXO-200 Collaboration), “Search for neutrinoless double- β decay with the complete EXO-200 dataset”, *Physical Review Letters* **123**, 161802 (2019).
- [107] N. Ackerman *et al.* (EXO Collaboration), “Observation of two-neutrino double-beta decay in ^{136}Xe with the EXO-200 detector”, *Physical Review Letters* **107**, 212501 (2011).
- [108] M. A. Luty, “Baryogenesis via leptogenesis”, *Physical Review D* **45**, 455 (1992).
- [109] W. Buchmüller, R. D. Peccei, and T. Yanagida, “Leptogenesis as the origin of matter”, *Annual Review of Nuclear and Particle Science* **55**, 311 (2005).
- [110] A. Kusenko, F. Takahashi, and T. T. Yanagida, “Dark matter from split seesaw”, *Physics Letters B* **693**, 144 (2010).
- [111] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “Tests of general relativity with the binary black hole signals from the LIGO-Virgo catalog GWTC-1”, *Physical Review D* **100**, 104036 (2019).
- [112] R. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog”, *Physical Review D* **103**, 122002 (2021).
- [113] L. Susskind, “Dynamics of spontaneous symmetry breaking in the Weinberg-Salam theory”, *Physical Review D* **20**, 2619 (1979).
- [114] N. Craig, “Naturalness: past, present, and future”, *The European Physical Journal C* **83**, 825 (2023).
- [115] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “The hierarchy problem and new dimensions at a millimeter”, *Physics Letters B* **429**, 263 (1998).

- [116] E. G. Floratos and G. K. Leontaris, “Low scale unification, Newton’s law and extra dimensions”, *Physics Letters B* **465**, 95 (1999).
- [117] A. Kehagias and K. Sfetsos, “Deviations from the $1/r^2$ Newton law due to extra dimensions”, *Physics Letters B* **472**, 39 (2000).
- [118] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “Phenomenology, astrophysics, and cosmology of theories with submillimeter dimensions and TeV scale quantum gravity”, *Physical Review D* **59**, 086004 (1999).
- [119] E. A. Mirabelli, M. Perelstein, and M. E. Peskin, “Collider signatures of new large space dimensions”, *Physical Review Letters* **82**, 2236 (1999).
- [120] G. F. Giudice, R. Rattazzi, and J. D. Wells, “Quantum gravity and extra dimensions at high-energy colliders”, *Nuclear Physics B* **544**, 3 (1999).
- [121] J. L. Hewett, “Indirect collider signals for extra dimensions”, *Physical Review Letters* **82**, 4765 (1999).
- [122] E. G. Adelberger, B. R. Heckel, and A. E. Nelson, “Tests of the gravitational inverse-square law”, *Annual Review of Nuclear and Particle Science* **53**, 77 (2003).
- [123] J. Murata and S. Tanaka, “A review of short-range gravity experiments in the LHC era”, *Classical and Quantum Gravity* **32**, 033001 (2015).
- [124] S. Chatrchyan *et al.* (CMS Collaboration), “Search for new physics with a monojet and missing transverse energy in pp collisions at $\sqrt{s} = 7$ TeV”, *Physical Review Letters* **107**, 201804 (2011).
- [125] G. Aad *et al.* (ATLAS Collaboration), “Search for new phenomena with the monojet and missing transverse momentum signature using the ATLAS detector in $\sqrt{s} = 7$ TeV proton–proton collisions”, *Physics Letters B* **705**, 294 (2011).
- [126] S. Chatrchyan *et al.* (CMS Collaboration), “Search for dark matter and large extra dimensions in monojet events in pp collisions at $\sqrt{s} = 7$ TeV”, *Journal of High Energy Physics* **2012**, 94 (2012).

- [127] G. Aad *et al.* (ATLAS Collaboration), “Search for dark matter candidates and large extra dimensions in events with a photon and missing transverse momentum in pp collision data at $\sqrt{s} = 7$ TeV with the ATLAS detector”, *Physical Review Letters* **110**, 011802 (2013).
- [128] V. Khachatryan *et al.* (CMS Collaboration), “Search for dark matter, extra dimensions, and unparticles in monojet events in proton–proton collisions at $\sqrt{s} = 8$ TeV”, *The European Physical Journal C* **75**, 235 (2015).
- [129] A. Sirunyan *et al.* (CMS Collaboration), “Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13$ TeV”, *Physical Review D* **97**, 092005 (2018).
- [130] G. Aad *et al.* (ATLAS Collaboration), “Search for new phenomena in events with an energetic jet and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”, *Physical Review D* **103**, 112006 (2021).
- [131] A. Tumasyan *et al.* (CMS Collaboration), “Search for new particles in events with energetic jets and large missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Journal of High Energy Physics* **2021**, 153 (2021).
- [132] G. Aad *et al.* (ATLAS Collaboration), “Search for quantum black hole production in lepton + jet final states using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”, *Physical Review D* **109**, 032010 (2024).
- [133] A. Hayrapetyan *et al.* (CMS Collaboration), “Search for new physics in high-mass diphoton events from proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Journal of High Energy Physics* **2024**, 215 (2024).
- [134] L. Randall and R. Sundrum, “Large mass hierarchy from a small extra dimension”, *Physical Review Letters* **83**, 3370 (1999).
- [135] W. D. Goldberger and M. B. Wise, “Modulus stabilization with bulk fields”, *Physical Review Letters* **83**, 4922 (1999).

- [136] C. Csáki, M. L. Graesser, and G. D. Kribs, “Radion dynamics and electroweak physics”, *Physical Review D* **63**, 065002 (2001).
- [137] C. Csáki, J. Hubisz, and S. J. Lee, “Radion phenomenology in realistic warped space models”, *Physical Review D* **76**, 125015 (2007).
- [138] L. Randall and R. Sundrum, “An alternative to compactification”, *Physical Review Letters* **83**, 4690 (1999).
- [139] A. Arvanitaki, S. Dimopoulos, V. Gorbenko, J. Huang, and K. Van Tilburg, “A small weak scale from a small cosmological constant”, *Journal of High Energy Physics* **2017**, 71 (2017).
- [140] I. Antoniadis, N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “New dimensions at a millimeter to a fermi and superstrings at a TeV”, *Physics Letters B* **436**, 257 (1998).
- [141] C. Vafa, “The string landscape and the swampland” (2005), arXiv:hep-th/0509212.
- [142] E. Palti, “The swampland: introduction and review”, *Fortschritte der Physik* **67**, 1900037 (2019).
- [143] M. Montero, C. Vafa, and I. Valenzuela, “The dark dimension and the swampland”, *Journal of High Energy Physics* **2023**, 22 (2023).
- [144] F. Marchesano, G. Shiu, and T. Weigand, “The Standard Model from string theory: what have we learned?”, *Annual Review of Nuclear and Particle Science* **74**, 113 (2024).
- [145] H. Leeb and J. Schmiedmayer, “Constraint on hypothetical light interacting bosons from low-energy neutron experiments”, *Physical Review Letters* **68**, 1472 (1992).
- [146] Yu. N. Pokotilovski, “Constraints on new interactions from neutron scattering experiments”, *Physics of Atomic Nuclei* **69**, 924 (2006).

- [147] V. V. Nesvizhevsky, G. Pignol, and K. V. Protasov, “Neutron scattering and extra-short-range interactions”, *Physical Review D* **77**, 034020 (2008).
- [148] Y. Kamiya, K. Itagaki, M. Tani, G. N. Kim, and S. Komamiya, “Constraints on new gravitylike forces in the nanometer range”, *Physical Review Letters* **114**, 161101 (2015).
- [149] V. V. Voronin, I. A. Kuznetsov, and D. D. Shapiro, “Search for novel short-range forces between elementary particles in neutron scattering”, *JETP Letters* **107**, 1 (2018).
- [150] C. C. Haddock *et al.*, “Search for deviations from the inverse square law of gravity at nm range using a pulsed neutron beam”, *Physical Review D* **97**, 062002 (2018).
- [151] G. L. Klimchitskaya, “Recent breakthrough and outlook in constraining the non-Newtonian gravity and axion-like particles from Casimir physics”, *The European Physical Journal C* **77**, 315 (2017).
- [152] B. Heacock *et al.*, “Pendellösung interferometry probes the neutron charge radius, lattice dynamics, and fifth forces”, *Science* **373**, 1239 (2021).
- [153] E. Fischbach and C. L. Talmadge, *The search for non-Newtonian gravity* (Springer, New York, NY, 1999).
- [154] S. M. Merkowitz, “Tests of gravity using lunar laser ranging”, *Living Reviews in Relativity* **13**, 7 (2010).
- [155] R. S. Decca, D. López, H. B. Chan, E. Fischbach, D. E. Krause, and C. R. Jamell, “Constraining new forces in the casimir regime using the isoelectronic technique”, *Physical Review Letters* **94**, 240401 (2005).
- [156] A. A. Geraci, S. J. Smullin, D. M. Weld, J. Chiaverini, and A. Kapitulnik, “Improved constraints on non-Newtonian forces at 10 microns”, *Physical Review D* **78**, 022002 (2008).

- [157] Y.-J. Chen, W. K. Tham, D. E. Krause, D. López, E. Fischbach, and R. S. Decca, “Stronger limits on hypothetical Yukawa interactions in the 30–8000 nm range”, *Physical Review Letters* **116**, 221102 (2016).
- [158] R. Spero, J. K. Hoskins, R. Newman, J. Pellam, and J. Schultz, “Test of the gravitational inverse-square law at laboratory distances”, *Physical Review Letters* **44**, 1645 (1980).
- [159] J. K. Hoskins, R. D. Newman, R. Spero, and J. Schultz, “Experimental tests of the gravitational inverse-square law for mass separations from 2 to 105 cm”, *Physical Review D* **32**, 3084 (1985).
- [160] D. J. Kapner, T. S. Cook, E. G. Adelberger, J. H. Gundlach, B. R. Heckel, C. D. Hoyle, and H. E. Swanson, “Tests of the gravitational inverse-square law below the dark-energy length scale”, *Physical Review Letters* **98**, 021101 (2007).
- [161] A. O. Sushkov, W. J. Kim, D. A. R. Dalvit, and S. K. Lamoreaux, “New experimental limits on non-Newtonian forces in the micrometer range”, *Physical Review Letters* **107**, 171101 (2011).
- [162] S.-Q. Yang, B.-F. Zhan, Q.-L. Wang, C.-G. Shao, L.-C. Tu, W.-H. Tan, and J. Luo, “Test of the gravitational inverse square law at millimeter ranges”, *Physical Review Letters* **108**, 081101 (2012).
- [163] W.-H. Tan *et al.*, “Improvement for testing the gravitational inverse-square law at the submillimeter range”, *Physical Review Letters* **124**, 051301 (2020).
- [164] J. Lee, E. Adelberger, T. Cook, S. Fleischer, and B. Heckel, “New test of the gravitational $1/r^2$ law at separations down to $52 \mu\text{m}$ ”, *Physical Review Letters* **124**, 101101 (2020).
- [165] M. V. Moody and H. J. Paik, “Gauss’s law test of gravity at short range”, *Physical Review Letters* **70**, 1195 (1993).

- [166] E. G. Adelberger, J. H. Gundlach, B. R. Heckel, S. Hoedl, and S. Schlamminger, “Torsion balance experiments: a low-energy frontier of particle physics”, *Progress in Particle and Nuclear Physics* **62**, 102 (2009).
- [167] J. Chiaverini, S. J. Smullin, A. A. Geraci, D. M. Weld, and A. Kapitulnik, “New experimental constraints on non-Newtonian forces below 100 μm ”, *Physical Review Letters* **90**, 151101 (2003).
- [168] H. Cavendish, “XXI. Experiments to determine the density of the earth”, *Philosophical Transactions of the Royal Society of London* **88**, 469 (1798).
- [169] L. Magrini, P. Rosenzweig, C. Bach, A. Deutschmann-Olek, S. G. Hofer, S. Hong, N. Kiesel, A. Kugi, and M. Aspelmeyer, “Real-time optimal quantum control of mechanical motion at room temperature”, *Nature* **595**, 373 (2021).
- [170] U. Delić, M. Reisenbauer, K. Dare, D. Grass, V. Vuletić, N. Kiesel, and M. Aspelmeyer, “Cooling of a levitated nanoparticle to the motional quantum ground state”, *Science* **367**, 892 (2020).
- [171] D. C. Moore and A. A. Geraci, “Searching for new physics using optically levitated sensors”, *Quantum Science and Technology* **6**, 014008 (2021).
- [172] C. Gonzalez-Ballester, M. Aspelmeyer, L. Novotny, R. Quidant, and O. Romero-Isart, “Levitodynamics: levitation and control of microscopic objects in vacuum”, *Science* **374**, eabg3027 (2021).
- [173] S. Bose, A. Mazumdar, G. W. Morley, H. Ulbricht, M. Toroš, M. Paternostro, A. A. Geraci, P. F. Barker, M. S. Kim, and G. Milburn, “Spin entanglement witness for quantum gravity”, *Physical Review Letters* **119**, 240401 (2017).
- [174] E. Conti *et al.*, “Correlated fluctuations between luminescence and ionization in liquid xenon”, *Physical Review B* **68**, 054201 (2003).
- [175] E. Aprile, C. E. Dahl, L. de Viveiros, R. J. Gaitskell, K. L. Giboni, J. Kwong, P. Majewski, K. Ni, T. Shutt, and M. Yamashita, “Simultaneous measurement

- of ionization and scintillation from nuclear recoils in liquid xenon for a dark matter experiment”, *Physical Review Letters* **97**, 081302 (2006).
- [176] E. Aprile, “Liquid xenon detectors for particle physics and astrophysics”, *Reviews of Modern Physics* **82**, 2053 (2010).
- [177] E. Aprile *et al.*, “First dark matter search with nuclear recoils from the XENONnT experiment”, *Physical Review Letters* **131**, 041003 (2023).
- [178] Z. Bo *et al.*, “Dark matter search results from 1.54 tonne · year exposure of PandaX-4T”, *Physical Review Letters* **134**, 011805 (2025).
- [179] J. Aalbers *et al.*, “First dark matter search results from the LUX-ZEPLIN (LZ) experiment”, *Physical Review Letters* **131**, 041002 (2023).
- [180] J. Jortner, L. Meyer, S. A. Rice, and E. G. Wilson, “Localized excitations in condensed Ne, Ar, Kr, and Xe”, *The Journal of Chemical Physics* **42**, 4250 (1965).
- [181] J. Thomas and D. A. Imel, “Recombination of electron-ion pairs in liquid argon and liquid xenon”, *Physical Review A* **36**, 614 (1987).
- [182] G. Anton *et al.* (EXO-200), “Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment”, *Physical Review C* **101**, 065501 (2020).
- [183] L. Baudis, P. Sanchez-Lucas, and K. Thieme, “A measurement of the mean electronic excitation energy of liquid xenon”, *The European Physical Journal C* **81**, 1060 (2021).
- [184] S. Al Kharusi *et al.*, “nEXO pre-conceptual design report” (2018), arXiv:1805.11142 [physics].
- [185] B. Aharmim *et al.* (SNO Collaboration), “Measurement of the cosmic ray and neutrino-induced muon flux at the Sudbury Neutrino Observatory”, *Physical Review D* **80**, 012001 (2009).

- [186] J. B. Albert *et al.*, “Measurement of the drift velocity and transverse diffusion of electrons in liquid xenon with the EXO-200 detector”, *Physical Review C* **95**, 025502 (2017).
- [187] O. Njoya *et al.*, “Measurements of electron transport in liquid and gas xenon using a laser-driven photocathode”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **972**, 163965 (2020).
- [188] M. Jewell *et al.*, “Characterization of an ionization readout tile for nEXO”, *Journal of Instrumentation* **13**, P01006.
- [189] G. Gallina *et al.*, “Performance of novel VUV-sensitive silicon photo-multipliers for nEXO”, *The European Physical Journal C* **82**, 1125 (2022).
- [190] J. H. Hubbell and S. M. Seltzer, “Tables of X-ray mass attenuation coefficients and mass energy-absorption coefficients 1 keV to 20 MeV for elements $Z = 1$ to 92 and 48 additional substances of dosimetric interest” (1995), <http://physics.nist.gov/PhysRefData/XrayMassCoef/cover.html>, (Accessed: June 2025).
- [191] M. J. Berger and S. M. Seltzer, “Stopping powers and ranges of electrons and positrons (2nd Ed.):”, National Institute of Standards and Technology, Gaithersburg, MD (1983), <https://doi.org/10.6028/NBS.IR.82-2550A>, (Accessed: June 2025).
- [192] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition”, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016) pp. 770–778.
- [193] R. H. M. Tsang *et al.*, “An integrated online radioassay data storage and analytics tool for nEXO”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **1055**, 168477 (2023).

- [194] G. Adhikari *et al.* (nEXO Collaboration), “nEXO: neutrinoless double beta decay search beyond 10^{28} year half-life sensitivity”, *Journal of Physics G: Nuclear and Particle Physics* **49**, 015104 (2021).
- [195] M. Szydagis *et al.*, “Noble Element Simulation Technique v2.0” (2018), <https://doi.org/10.5281/zenodo.1314669>.
- [196] B. Lenardo *et al.* (nEXO Collaboration), “Development of a ^{127}Xe calibration source for nEXO”, *Journal of Instrumentation* **17**, P07028.
- [197] R. Saldanha, *Electron collection efficiency and charge resolution v1.1*, Tech. Rep. (nEXO Collaboration Wiki, 2021).
- [198] V. N. Solovov, V. Chepel, M. I. Lopes, A. Hitachi, R. Ferreira Marques, and A. J. P. L. Policarpo, “Measurement of the refractive index and attenuation length of liquid xenon for its scintillation light”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **516**, 462 (2004).
- [199] A. Baldini *et al.*, “Absorption of scintillation light in a 100 l liquid xenon γ -ray detector and expected detector performance”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **545**, 753 (2005).
- [200] A. Manalaysay *et al.*, “Spatially uniform calibration of a liquid xenon detector at low energies using $^{83\text{m}}\text{Kr}$ ”, *Review of Scientific Instruments* **81**, 073303 (2010).
- [201] L. W. Kastens, S. B. Cahn, A. Manzur, and D. N. McKinsey, “Calibration of a liquid xenon detector with $^{83}\text{Kr}^{\text{m}}$ ”, *Physical Review C* **80**, 045809 (2009).
- [202] D. S. Akerib *et al.* (LUX Collaboration), “ $^{83\text{m}}\text{Kr}$ calibration of the 2013 LUX dark matter search”, *Physical Review D* **96**, 112009 (2017).
- [203] D. S. Akerib *et al.* (LUX Collaboration), “Tritium calibration of the LUX dark matter experiment”, *Physical Review D* **93**, 072009 (2016).

- [204] E. M. Boulton *et al.*, “Calibration of a two-phase xenon time projection chamber with a ^{37}Ar source”, *Journal of Instrumentation* **12**, P08004.
- [205] K. Ni, R. Hasty, T. M. Wongjirad, L. Kastens, A. Manzur, and D. N. McKinsey, “Preparation of neutron-activated xenon for liquid xenon detector calibration”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **582**, 569 (2007).
- [206] R. Lang, A. Brown, E. Brown, M. Cervantes, S. Macmullin, D. Masson, J. Schreiner, and H. Simgen, “A ^{220}Rn source for the calibration of low-background experiments”, *Journal of Instrumentation* **11**, P04004.
- [207] E. Aprile *et al.* (XENON Collaboration), “Results from a calibration of XENON100 using a source of dissolved radon-220”, *Physical Review D* **95**, 072008 (2017).
- [208] M.-M. Bé *et al.*, *Table of radionuclides*, Monographie BIPM-5, Vol. 8 (Bureau International des Poids et Mesures, Sèvres, France, 2017).
- [209] Laboratoire National Henri Becquerel, “Nucléide - Lara”, <http://www.lnhb.fr/Laraweb/index.php>, (Accessed: July 2025).
- [210] M. B. Chadwick *et al.*, “ENDF/B-VII.0: next generation evaluated nuclear data library for nuclear science and technology”, *Nuclear Data Sheets Evaluated Nuclear Data File ENDF/B-VII.0*, **107**, 2931 (2006).
- [211] A. J. Koning, D. Rochman, J. C. Sublet, N. Dzysiuk, M. Fleming, and S. van der Marck, “TENDL: complete nuclear data library for innovative nuclear science and technology”, *Nuclear Data Sheets Special Issue on Nuclear Reaction Data*, **155**, 1 (2019).
- [212] L. Fabris, N. W. Madden, and H. Yaver, “A fast, compact solution for low noise charge preamplifiers”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **424**, 545 (1999).

- [213] S. Seibert and A. Latorre, “Fast optical monte carlo simulation with surface-based geometries using chroma” (2011), <https://pdfs.semanticscholar.org/33ad/1bae64007a43a840288a888eba7bc3e3a37a.pdf>.
- [214] G. Bakale, U. Sowada, and W. F. Schmidt, “Effect of an electric field on electron attachment to sulfur hexafluoride, nitrous oxide, and molecular oxygen in liquid argon and xenon”, *The Journal of Physical Chemistry* **80**, 2556 (1976).
- [215] G. Plante, E. Aprile, J. Howlett, and Y. Zhang, “Liquid-phase purification for multi-tonne xenon detectors”, *The European Physical Journal C* **82**, 860 (2022).
- [216] C. Benvenuti, “Molecular surface pumping: the getter pumps”, in *CAS - CERN Accelerator School : Vacuum Technology* (CERN, 1999) pp. 43–50.
- [217] C. Benvenuti, “Getter pumping”, in *CAS - CERN Accelerator School : Vacuum in Accelerators* (CERN, 2007) pp. 313–320.
- [218] B. Ferrario, “Chemical pumping in vacuum technology”, *Vacuum 1st International Workshop on Interaction of Gases with Solids Affecting Vacuum Processes*, **47**, 363 (1996).
- [219] J. P. Pemsler, “Diffusion of oxygen in zirconium and its relation to oxidation and corrosion”, *Journal of The Electrochemical Society* **105**, 315 (1958).
- [220] C. Morant, J. M. Sanz, L. Galán, L. Soriano, and F. Rueda, “An XPS study of the interaction of oxygen with zirconium”, *Surface Science* **218**, 331 (1989).
- [221] J. S. Foord, P. J. Goddard, and R. M. Lambert, “Adsorption and absorption of diatomic gases by zirconium: studies of the dissociation and diffusion of CO, NO, N₂, O₂ and D₂”, *Surface Science* **94**, 339 (1980).
- [222] N. Stojilovic, E. T. Bender, and R. D. Ramsier, “Surface chemistry of zirconium”, *Progress in Surface Science* **78**, 101 (2005).
- [223] C. Boffito, B. Ferrario, P. della Porta, and L. Rosai, “A nonevaporable low temperature activatable getter material”, *Journal of Vacuum Science and Technology* **18**, 1117 (1981).

- [224] Y. P. Lee, S. Yokouchi, Y. Morimoto, H. Sakamoto, T. Nishidono, N. Hinago, and S. H. Be, “The performance characteristics of a St 707 non-evaporable getter strip for SPring-8”, *Shinku* **33**, 154 (1990).
- [225] C. Benvenuti and P. Chiggiato, “Pumping characteristics of the St707 nonevaporable getter (Zr 70 V 24.6-Fe 5.4 wt %)”, *Journal of Vacuum Science & Technology A* **14**, 3278 (1996).
- [226] S. Avdiaj, B. Šetina Batič, J. Šetina, and B. Erjavec, “Oxygen diffusion in the non-evaporable getter St 707 during heat treatment”, *Materiali in Tehnologije* **45**, 33 (2011).
- [227] A. Dobi, D. S. Leonard, C. Hall, L. J. Kaufman, T. Langford, S. Slutsky, and Y. R. Yen, “Study of a zirconium getter for purification of xenon gas”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **620**, 594 (2010).
- [228] D. S. Akerib *et al.* (LZ Collaboration), “The LUX-ZEPLIN (LZ) radioactivity and cleanliness control programs”, *The European Physical Journal C* **80**, 1044 (2020).
- [229] E. Aprile *et al.* (XENON Collaboration), “ ^{222}Rn emanation measurements for the XENON1T experiment”, *The European Physical Journal C* **81**, 337 (2021).
- [230] C. Hall, “Radon burden and Xe purification performance of the LZ hot zirconium getter”, in *Low Radioactivity Techniques (LRT2022) Poster Session* (Sanford Underground Research Facility (SURF), Rapid City, SD, USA, 2022).
- [231] E. Aprile *et al.* (XENON Collaboration), “Material radiopurity control in the XENONnT experiment”, *The European Physical Journal C* **82**, 599 (2022).
- [232] F. Jörg, *From ^{222}Rn measurements in XENONnT and HeXe to radon mitigation in future liquid xenon experiments*, Ph.D. thesis, Heidelberg University, Heidelberg, Germany (2022).

- [233] D. O. Northwood, “The development and applications of zirconium alloys”, *Materials & Design* **6**, 58 (1985).
- [234] L. Xu, Y. Xiao, A. van Sandwijk, Q. Xu, and Y. Yang, “Production of nuclear grade zirconium: a review”, *Journal of Nuclear Materials* **466**, 21 (2015).
- [235] L. Xu, Y. Xiao, A. van Sandwijk, Q. Xu, and Y. Yang, “Separation of zirconium and hafnium: a review”, in *Energy Materials 2014* (Springer International Publishing, Cham, 2016) pp. 451–457.
- [236] O. B. Malyshev, K. J. Middleman, J. S. Colligon, and R. Valizadeh, “Activation and measurement of nonevaporable getter films”, *Journal of Vacuum Science & Technology A* **27**, 321 (2009).
- [237] B. Einfeld and K. Schnitzlein, “The influence of confining walls on the pressure drop in packed beds”, *Chemical Engineering Science* **56**, 4321 (2001).
- [238] E. M. Moghaddam, E. A. Foumeny, A. I. Stankiewicz, and J. T. Padding, “Fixed bed reactors of non-spherical pellets: importance of heterogeneities and inadequacy of azimuthal averaging”, *Chemical Engineering Science: X* **1**, 100006 (2019).
- [239] S. Ergun, “Fluid flow through packed columns”, *Chemical Engineering Progress* **48**, 89 (1952).
- [240] K. Wamba, *Aspects of the R&D for the Enriched Xenon Observatory for double beta decay*, Ph.D. thesis, Stanford University, Stanford, CA (2007).
- [241] F. LePort *et al.*, “A magnetically driven piston pump for ultra-clean applications”, *Review of Scientific Instruments* **82**, 105114 (2011).
- [242] E. Brown *et al.*, “Magnetically-coupled piston pump for high-purity gas applications”, *The European Physical Journal C* **78**, 604 (2018).
- [243] R. A. Pierotti, “The solubility of gases in liquids”, *The Journal of Physical Chemistry* **67**, 1840 (1963).

- [244] G. Venugopalan *et al.*, “Search for new interactions at the micron scale with a vector force sensor” (2024), arXiv:2412.13167 [hep-ex].
- [245] A. Ashkin, “Acceleration and trapping of particles by radiation pressure”, *Physical Review Letters* **24**, 156 (1970).
- [246] A. Ashkin, J. M. Dziedzic, and T. Yamane, “Optical trapping and manipulation of single cells using infrared laser beams”, *Nature* **330**, 769 (1987).
- [247] A. Ashkin and J. M. Dziedzic, “Optical trapping and manipulation of viruses and bacteria”, *Science* **235**, 1517 (1987).
- [248] A. Ashkin, “Trapping of atoms by resonance radiation pressure”, *Physical Review Letters* **40**, 729 (1978).
- [249] S. Chu, L. Hollberg, J. E. Bjorkholm, A. Cable, and A. Ashkin, “Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure”, *Physical Review Letters* **55**, 48 (1985).
- [250] K. C. Neuman and S. M. Block, “Optical trapping”, *Review of Scientific Instruments* **75**, 2787 (2004).
- [251] A. Ashkin, “Optical trapping and manipulation of neutral particles using lasers”, *Proceedings of the National Academy of Sciences* **94**, 4853 (1997).
- [252] A. Ashkin, “Forces of a single-beam gradient laser trap on a dielectric sphere in the ray optics regime”, *Biophysical Journal* **61**, 569 (1992).
- [253] P. A. M. Neto and H. M. Nussenzveig, “Theory of optical tweezers”, *Europhysics Letters* **50**, 702 (2000).
- [254] A. Mazolli, P. A. Maia Neto, and H. M. Nussenzveig, “Theory of trapping forces in optical tweezers”, *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **459**, 3021 (2003).

- [255] A. Ashkin, J. M. Dziedzic, J. E. Bjorkholm, and S. Chu, “Observation of a single-beam gradient force optical trap for dielectric particles”, *Optics Letters* **11**, 288 (1986).
- [256] A. Ashkin and J. M. Dziedzic, “Optical levitation by radiation pressure”, *Applied Physics Letters* **19**, 283 (1971).
- [257] A. Ashkin and J. M. Dziedzic, “Stability of optical levitation by radiation pressure”, *Applied Physics Letters* **24**, 586 (1974).
- [258] A. Ashkin and J. M. Dziedzic, “Feedback stabilization of optically levitated particles”, *Applied Physics Letters* **30**, 202 (1977).
- [259] A. Cavalleri, G. Ciani, R. Dolesi, M. Hueller, D. Nicolodi, D. Tombolato, S. Vitale, P. J. Wass, and W. J. Weber, “Gas damping force noise on a macroscopic test body in an infinite gas reservoir”, *Physics Letters A* **374**, 3365 (2010).
- [260] microparticles GmbH, “microParticles Onlineshop”, <https://www.microparticles-shop.de/index.php?language=en>, (Accessed: December 2024).
- [261] A. Kawasaki, A. Fieguth, N. Priel, C. P. Blakemore, D. Martin, and G. Gratta, “High sensitivity, levitated microsphere apparatus for short-distance force measurements”, *Review of Scientific Instruments* **91**, 083201 (2020).
- [262] E. G. Acheson, “AquaDAG patent US844989A” (1907), <https://patents.google.com/patent/US844989A>.
- [263] D. C. Moore, A. D. Rider, and G. Gratta, “Search for millicharged particles using optically levitated microspheres”, *Physical Review Letters* **113**, 251801 (2014).
- [264] C. P. Blakemore, A. D. Rider, S. Roy, Q. Wang, A. Kawasaki, and G. Gratta, “Three-dimensional force-field microscopy with optically levitated microspheres”, *Physical Review A* **99**, 023816 (2019).

- [265] Q. Wang, A. D. Rider, D. C. Moore, C. P. Blakemore, L. Cao, and G. Gratta, “A density staggered cantilever for micron length gravity probing”, in *2017 IEEE 67th Electronic Components and Technology Conference (ECTC)* (2017) pp. 1773–1778.
- [266] C. P. Blakemore, A. Fieguth, A. Kawasaki, N. Priel, D. Martin, A. D. Rider, Q. Wang, and G. Gratta, “Search for non-Newtonian interactions at micrometer scale with a levitated test mass”, *Physical Review D* **104**, L061101 (2021).
- [267] G. Venugopalan and G. Gratta, “Platinum black for stray-light mitigation on high-aspect-ratio micromechanical cantilever”, *Review of Scientific Instruments* **96**, 075004 (2025).
- [268] B. Allen, W. Hua, and A. Ottewill, “Automatic cross-talk removal from multi-channel data” (1999), arXiv:gr-qc/9909083.
- [269] G. Vajente, “Data mining and machine learning improve gravitational-wave detector sensitivity”, *Physical Review D* **105**, 102005 (2022).
- [270] C. P. Blakemore, A. D. Rider, S. Roy, A. Fieguth, A. Kawasaki, N. Priel, and G. Gratta, “Precision mass and density measurement of individual optically levitated microspheres”, *Physical Review Applied* **12**, 024037 (2019).
- [271] A. D. Rider, *Measurements with optically levitated microspheres*, Ph.D. thesis, Stanford University, Stanford, CA (2019).
- [272] K. Kohn, *Characterizing optical and electromagnetic signal-like backgrounds in measurements of gravity with optically levitated microspheres*, Bachelor’s thesis, Stanford University, Stanford, CA (2025).
- [273] J. Neyman, E. S. Pearson, and K. Pearson, “IX. On the problem of the most efficient tests of statistical hypotheses”, *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* **231**, 289 (1933).

- [274] S. S. Wilks, “The large-sample distribution of the likelihood ratio for testing composite hypotheses”, *The Annals of Mathematical Statistics* **9**, 60 (1938).
- [275] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics”, *The European Physical Journal C* **71**, 1554 (2011).
- [276] E. Hough, *Novel techniques to measure micron-scale gravity*, Bachelor’s thesis, Stanford University, Stanford, CA (2022).
- [277] F. Monteiro, G. Afek, D. Carney, G. Krnjaic, J. Wang, and D. C. Moore, “Search for composite dark matter with optically levitated sensors”, *Physical Review Letters* **125**, 181102 (2020).
- [278] E. Kilian *et al.*, “Dark matter searches with levitated sensors”, *AVS Quantum Science* **6**, 030503 (2024).
- [279] D. Carney, K. G. Leach, and D. C. Moore, “Searches for massive neutrinos with mechanical quantum sensors”, *PRX Quantum* **4**, 010315 (2023).
- [280] N. Priel, A. Fieguth, C. P. Blakemore, E. Hough, A. Kawasaki, D. Martin, G. Venugopalan, and G. Gratta, “Dipole moment background measurement and suppression for levitated charge sensors”, *Science Advances* **8**, eabo2361 (2022).
- [281] J. Millen, T. S. Monteiro, R. Pettit, and A. N. Vamivakas, “Optomechanics with levitated particles”, *Reports on Progress in Physics* **83**, 026401 (2020).
- [282] Acktar, “Magic Black™ Coating”, <https://acktar.com/product/magic-black/>, (Accessed: March 2024).
- [283] C. P. Blakemore, D. Martin, A. Fieguth, N. Priel, G. Venugopalan, A. Kawasaki, and G. Gratta, “Librational feedback cooling”, *Physical Review A* **106**, 023503 (2022).

- [284] M. Agostini, G. Benato, and J. A. Detwiler, “Discovery probability of next-generation neutrinoless double- β decay experiments”, *Physical Review D* **96**, 053001 (2017).
- [285] I. Esteban, M. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou, “The fate of hints: updated global analysis of three-flavor neutrino oscillations”, *Journal of High Energy Physics* **2020**, 178 (2020).
- [286] M. Aker *et al.* (KATRIN Collaboration), “Direct neutrino-mass measurement with sub-electronvolt sensitivity”, *Nature Physics* **18**, 160 (2022).
- [287] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, “emcee: the MCMC hammer”, *Publications of the Astronomical Society of the Pacific* **125**, 306 (2013).
- [288] E. Angelico, J. Dalmasson, R. DeVoe, G. Gratta, C. A. Hardy, B. Lenardo, L. Si, M. Vidal, and S. Wu, “A flexible test facility for liquid xenon detector development” (2025), arXiv:2508.05853 [physics].
- [289] 3M, “3M™ Novec™ 7000 Engineered Fluid”, https://www.3m.com/3M/en_US/p/d/b5005006004/, (Accessed: August 2025).
- [290] N. Ackerman *et al.*, “The EXO-200 detector, part II: auxiliary systems”, *Journal of Instrumentation* **17**, P02015.
- [291] D. S. Leonard *et al.*, “Systematic study of trace radioactive impurities in candidate construction materials for EXO-200”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **591**, 490 (2008).
- [292] L. Fabris, *Novel readout design criteria for SiPM-based radiation detectors*, Ph.D. thesis, Università degli studi di Bergamo, Bergamo, Italy (2016).
- [293] Z. Li *et al.*, “Simulation of charge readout with segmented tiles in nEXO”, *Journal of Instrumentation* **14**, P09020.

- [294] G. Carugno, B. Dainese, F. Pietropaolo, and F. Ptohos, “Electron lifetime detector for liquid argon”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **292**, 580 (1990).
- [295] N. G. Trinh, “Electrode design for testing in uniform field gaps”, *IEEE Transactions on Power Apparatus and Systems* **PAS-99**, 1235 (1980).
- [296] W. Ma *et al.*, “Internal calibration of the PandaX-II detector with radon gaseous sources”, *Journal of Instrumentation* **15**, P12038.
- [297] F. Jörg, S. Li, J. Schreiner, H. Simgen, and R. F. Lang, “Characterization of a ^{220}Rn source for low-energy electronic recoil calibration of the XENONnT detector”, *Journal of Instrumentation* **18**, P11009.
- [298] D. Akerib *et al.*, “Ultralow energy calibration of lux detector using ^{127}Xe electron capture”, *Physical Review D* **96**, 112011 (2017).
- [299] StanfordBeads, “OptLevAnalysis: analysis code for the Stanford optical levitation lab” (2025), <https://github.com/stanfordbeads/OptLevAnalysis>.
- [300] B. Gustavsen, “Relaxed vector fitting algorithm for rational approximation of frequency domain responses”, in *2006 IEEE Workshop on Signal Propagation on Interconnects* (2006) pp. 97–100.
- [301] J. C. Driggers *et al.* (The LIGO Scientific Collaboration Instrument Science Authors), “Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO”, *Physical Review D* **99**, 042001 (2019).
- [302] L. B. Loeb, “The contact potential difference or Volta potential”, in *Static Electrification*, edited by L. B. Loeb (Springer Berlin Heidelberg, Berlin, Heidelberg, 1958) pp. 32–58.
- [303] M. Nonnenmacher, M. P. O’Boyle, and H. K. Wickramasinghe, “Kelvin probe force microscopy”, *Applied Physics Letters* **58**, 2921 (1991).

- [304] W. Melitz, J. Shen, A. C. Kummel, and S. Lee, “Kelvin probe force microscopy and its application”, *Surface Science Reports* **66**, 1 (2011).