

Theoretical aspects and data analysis for topological dark matter searches

By

Joseph A. Smiga

born in Rockville, MD, USA

Mainz // 1 October 2021

*Dissertation submitted
for the award of the title “Doctor of Natural Sciences”
to the Faculty of Physics, Mathematics, and Computer Science
of Johannes Gutenberg University Mainz*

Supervisor: Prof. Dr. Dmitry Budker

HIM
HELMHOLTZ
Helmholtz Institute Mainz



Abstract

Dark matter searches are a significant field of focus in modern physics. Despite the strong evidence for the existence of dark matter on galactic scales, its underlying composition remains a mystery. Theoretical efforts point to possible candidates, while experiments and observations place constraints on the theories. In this work, data from the Global Network of Optical Magnetometers for Exotic physics searches (GNOME) are analyzed for evidence of dark matter; namely axion/axion-like-particle domain walls which can couple to fermion spins. The underlying theory of this dark matter is connected to the observations from GNOME in order to constrain the physical parameter space. A thorough description of analysis methods is given along with quantitative meta-analysis of the experiment. Significant evidence of domain walls was not found which allows for a region of parameter space to be excluded.

To friends, family, and all that gives meaning

Acknowledgements

Studying physics has been a dream of mine for much of my life, and I am grateful for the opportunities that I have had to do so. Though the last few years of studying, researching, and working towards my doctoral degree have allowed me to build towards my goal, they have not been without the doubt, stress, and hesitations that many in similar positions endure. Unquestionably, the support of my friends, family, and colleagues have helped me make it through this process.

[Text removed for privacy rules]

Lastly, I would like to thank my parents and the rest of my family. They have always been supportive of my scientific ambitions from developing my interest in school to refining my knowledge in undergrad to moving to Germany for my PhD. They have always been there to encourage me, even if they may not have any idea what I am talking about — though this may not just be with science.

Conventions

- The metric signature is $(+, -, -, -)$.
- Units chosen such that $\hbar = c = 1$ unless otherwise stated.
- The continuous Fourier transform is $\{\mathcal{F}_C f\}(\omega) := \int_{-\infty}^{\infty} dt f(t)e^{-i\omega t}$ with inverse $\{\mathcal{F}_C^{-1} \tilde{f}\}(t) := \frac{1}{2\pi} \int_{-\infty}^{\infty} d\omega \tilde{f}(\omega)e^{+i\omega t}$.
- The discrete Fourier transform is $\{\mathcal{F}_D f\}[k] := \sum_{n=0}^{N-1} f[n]e^{-\frac{2\pi i}{N}nk}$ with inverse $\{\mathcal{F}_D^{-1} \tilde{f}\}[n] := \frac{1}{N} \sum_{k=0}^{N-1} \tilde{f}[k]e^{+\frac{2\pi i}{N}nk}$.
- Einstein summation notation is used; e.g., $x^\mu y_\mu = \sum_{\mu=0}^3 x^\mu y_\mu$.
- Vectors are written in boldface; e.g., \mathbf{x} .

Table 1: A list of mathematical conventions.

Symbol	Description
γ^μ	Dirac gamma matrices, $\mu \in \{0, 1, 2, 3\}$.
γ^5	Fifth gamma matrix, $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$.
ϕ^*	Complex conjugation of ϕ .
A^T	Transpose of matrix A .
ψ^\dagger	Adjoint of ψ , $\psi^\dagger = (\psi^*)^T$.
$\bar{\psi}$	The Dirac adjoint, $\bar{\psi} = \psi^\dagger\gamma^0$.
$\overleftrightarrow{\partial}_\mu$	$\psi \overleftrightarrow{\partial}_\mu \phi = \psi \partial_\mu \phi - (\partial_\mu \psi) \phi$.
$\langle x \rangle$	The average/expectation value of x .
$\mathcal{F}_C, \mathcal{F}_D$	Functionals for the continuous and discrete Fourier transform, respectively.
$[x]$	Ceiling function; round x up to the nearest integer.
$\lfloor x \rfloor$	Floor function; round x down to the nearest integer.
$B(k; n, p)$	Binomial mass function: probability of k out of n successes if each success has probability p .
$P(k; \mu)$	Poissonian mass function: probability of k events when expecting μ .
$\frac{\partial \mathbf{y}}{\partial \mathbf{x}}$	Jacobian matrix for (vector) function \mathbf{y} with respect to variables \mathbf{x} .
$\mathcal{O}(n)$	On the order of “ n ”; either a number, e.g. $\mathcal{O}(1)$, or scaling, e.g., $\mathcal{O}(n^2)$.

Table 2: List of constants with known/approximate values when applicable.

Symbol	Value	Description
c	299792458 m/s	Speed of light in a vacuum.
\hbar	6.582×10^{-16} eV s	Reduced Planck’s constant.
μ_B	5.788×10^{-5} eV/T	Bohr magneton.
f_π	93 MeV	Pion decay constant.
h		Hubble parameter in units of 100 km/s/Mpc.
R_\oplus	6378 km	Earth’s radius.
ρ_{DM}	4×10^5 GeV/m ³	Local dark matter energy density.
ρ_{DW}	$\lesssim \rho_{\text{DM}}$	Domain-wall energy density throughout the galaxy.
v_e	≈ 550 km/s	Galactic escape velocity.
v_c	$\approx 10^{-3}c$	Galactic orbital velocity.
\bar{v}	v_c	Average domain-wall velocity.
m_a		Axion/ALP mass.
N_{DW}		Domain-wall number/color anomaly.
\bar{f}_a		Axion/ALP decay constant.
f_a	\bar{f}_a/N_{DW}	
f_{int}		Coupling strength (in units of energy).
ξ	f_a/f_{int}	Ratio of energy constants. Other sources use $C = 2\xi$.
σ_{DW}	$\propto m_a f_a^2$	Domain-wall surface tension.
Δx	$\propto m_a^{-1}$	Domain-wall width.

Table 3: A list of abbreviations.

Name	Description
ALP	Axion-Like Particles
DFSZ	Dine-Fischler-Srednicki-Zhitnitsky
DM	Dark Matter
DW	Domain Wall
KSVZ	Kim-Shifman-Vainshtein-Zakharov
LSP	Lightest Supersymmetric Particle
MOND	Modified Newtonian Dynamics
NMOR	Nonlinear Magneto-Optical Rotation
PQWW	Peccei-Quinn-Weinberg-Wilczek
QCD	Quantum Chromodynamics
SERF	Spin-Exchange Relaxation-Free
SHM	Standard Halo Model
SUSY	Supersymmetry
WIMP	Weakly Interacting Massive Particle
<u>Experiments</u>	
ABRACADABRA	A Broadband/Resonant Approach to Cosmic Axion Detection with an Amplifying B-field Ring Apparatus
ADMX	Axion Dark Matter Experiment
ALPS	Any Light Particle Search
CASPEr	Cosmic Axion Spin Precession Experiment
CAST	CERN Axion Solar Telescope
GNOME	Global Network of Optical Magnetometers for Exotic physics searches
HAYSTAC	Haloscope at Yale Sensitive to Axion CDM
IAXO	International Axion Observatory
MADMAX	Magnetized Disc and Mirror Axion Experiment
<u>Statistics</u>	
CDF	Cumulative Distribution Function
EPDF	Event Probability Distribution Function
PDF	Probability Distribution Function
PMF	Probability Mass Function

Contents

Abstract	i
Acknowledgements	v
Conventions	vii
Contents	xi
1 Introduction	1
2 Background	3
2.1 Dark matter	3
2.2 Axions and ALPs	4
2.2.1 Strong- CP problem	4
2.2.2 Peccei-Quinn mechanism	5
2.2.3 PQWW, KSVZ, and DFSZ axions	6
2.2.4 Axions in string theory	9
2.2.5 Current constraints	9
2.3 WIMPs	10
2.4 Other dark matter candidates	12
2.5 Experimental searches	14
2.6 Macroscopic dark matter objects	16
2.6.1 Domain walls	16
2.6.2 Boson stars and Q -balls	17
2.7 Standard halo model	19
2.8 Magnetometer operation	23
3 Signal shape	25
3.1 Domain wall field	25
3.2 Spin coupling	27
3.2.1 Linear	28
3.2.2 Generalizations	29
3.3 Mass coupling	30

4	Analysis	33
4.1	Filtering data	33
4.1.1	Effects on signal	36
4.1.2	Effects on noise	36
4.1.3	Examples of signals and filters	38
4.2	Consistency check	41
4.2.1	Signal amplitude and sensitivity	42
4.2.2	Directional consistency	43
4.3	Geometric picture	46
4.4	Velocity lattice	47
4.4.1	Lattice requirements	48
4.4.2	Lattice optimization	51
4.5	Procedure	52
4.6	Choosing thresholds	55
5	Sensitivity	57
5.1	Network sensitivity	57
5.2	Connection to physical parameters	58
5.3	Sensitive region of parameter space	59
5.3.1	Loudest-event method	60
5.3.2	Extended method	61
5.3.3	Active time	61
5.4	Scaling	62
6	Experimental data	65
6.1	Data structure	65
6.2	Sample data	68
6.3	Run characteristics	70
6.4	Quality of the network	71
6.4.1	Ideal network	72
6.4.2	Quality of GNOME	75
7	Results	77
7.1	Selecting thresholds	78
7.2	Event rate	79
7.3	Exclusion region	79
8	Summary and outlook	85
	Bibliography	89

A	Useful derivations	101
A.1	Domain wall shape	101
A.2	χ^2 minimization	102
A.3	Sensitivity scaling	104
A.3.1	Coordinate system	104
A.3.2	Volume element	106
A.3.3	Directional sensitivity	108
B	Manipulating probability distributions	111
B.1	Transforming the distribution	112
B.2	Combining random values	115
C	Confidence and significance in Poissonian statistics	117
C.1	Poissonian distribution	118
C.2	Confidence intervals	118
C.3	Significance of a measurement	120

Chapter 1

Introduction

Dark matter remains a mystery to science. Despite strong astrophysical evidence of its gravitational influence on galactic scales — providing significantly higher mass to the galaxy than ordinary matter — its underlying nature remains unknown. Over the years, numerous theoretical models have been proposed in an attempt to understand its nature. Likewise, numerous experimental efforts have been made to test these models and place constraints on parameter space. Regardless, definitive evidence of dark matter beyond its gravitational influence continues to evade the best efforts of the physics community. The work described here focuses on how data from a particular dark matter experiment is used to better understand the nature of dark matter.

One new effort from the last decade to detect dark matter is the Global Network of Optical Magnetometers for Exotic physics searches (GNOME). First proposed in 2013 [1] with the first Science Run completed in 2017 [2], GNOME consists of shielded magnetometers around the world. This experiment aims to find evidence of exotic spin-coupling between atoms and dark matter that would appear as a pseudo-magnetic field in the network sensors. Particular interest for this experiment is in detecting domain-wall crossings from axion-like dark matter. Domain walls are two-dimensional macroscopic objects across which the underlying field changes significantly.

Substantial work has already led to results from GNOME. Much of the work described here expands on the analysis developed in Ref. [3] and the results described in Ref. [4]. Other analysis efforts of GNOME can be found in Refs. [5–7].

This dissertation is organized as follows: background information is summarized in Chapter 2; including an overview of dark matter theory, experimental efforts, and basic magnetometer operation. In Chapter 3, the expected shape of an expected signal is described. In Chapter 4, the analysis methods used to understand GNOME data are outlined. In Chapter 5, the sensitivity of the GNOME network is characterized along with a description of how to understand the types of signals that can be expected to be observed by the experiment. Some experimental data is presented in Chapter 6, along with meta-analysis on how well the experiment performed. Results of the experiment are summarized in Chapter 7, and concluding remarks on the findings, state, and future of GNOME are given in Chapter 8. Additionally, Appendix A includes various useful

mathematical derivations, Appendix B describes how probability distributions are and can be altered through various operations, and Appendix C describes various applications of Poissonian statistics.

Chapter 2

Background

Some foundation is needed to understand the basis and current state of dark matter searches. First, a general overview of dark matter is given in Sec. 2.1. Then, a more detailed overview of specific dark matter candidates is given with axions in Sec. 2.2, weakly interacting massive particles (WIMPs) in Sec. 2.3, and others in Sec. 2.4. Several experimental efforts to understand dark matter are described in Sec. 2.5. Macroscopic phenomena that may arise in dark matter fields — namely domain walls and boson stars — are described in Sec. 2.6 and a model of how these objects are distributed throughout the galaxy is described in Sec. 2.7. Finally, an overview of optical magnetometers, the device used by GNOME to search for dark matter, is given in Sec. 2.8.

2.1 Dark matter

One of the major outstanding problems in physics is in understanding the nature of dark matter. Various astronomical observations support the existence of matter that contributes to the mass of galaxies but is non-colliding and “dark” to direct observation [8]. The existence of this “dark matter” is supported by various astronomical measurements: galactic rotation curves of spiral galaxies [9–11], dynamics of galaxies in clusters [12], X-ray measurements of elliptical galaxies [13, 14] and galactic clusters [15], and gravitational lensing measurements [16–18] (notably from observations of the Bullet Cluster). These examples of dark matter observations all rely on combining observations of visible matter to measurements of mass and gravity. Though these measurements may allow one to understand how much dark matter there is in different parts of the universe, the constituent(s) of dark matter remain unknown.

Over the years, many candidates have been proposed as potential candidates for dark matter. Some of these candidates, such as the axion and WIMPs, arise naturally from theories that aimed to solve other problems in physics. However, it is entirely possible for dark matter to exist on its own without solving unrelated problems. Likewise, it is possible for multiple types of dark matter to combine to explain the missing matter. More details on proposed candidates are given in Sec. 2.2 (axions), Sec. 2.3 (WIMPs), and Sec. 2.4 (other candidates).

A major alternative to dark matter is Modified Newtonian Dynamics (MOND) [19, 20]. The premise of this theory is that Newton’s second law becomes non-linear in the limit of small acceleration, hence causing the rotation curves of galaxies to “flatten” at large distances where gravitational acceleration is very small. In particular,

$$\mathbf{F} = \begin{cases} m\mathbf{a} & a \gg a_0 \\ m\mathbf{a} \frac{a}{a_0} & a \ll a_0 \end{cases},$$

where a_0 is a constant with units of acceleration that defines the scale at which Newtonian dynamics fails. Though perhaps MOND explanations of astronomical observations are not as popular as dark matter explanations, some efforts on this front remain [21–24]. However, it seems that MOND would struggle to explain all features that could be explained by dark matter. For example, observations of gravitational lensing (e.g., the Bullet Cluster [16, 17]) seem to be at odds with MOND predictions. Some explanations for this discrepancy have been proposed [21, 22], but these are often incomplete or still require the inclusion of collisionless matter. Perhaps new MOND models will be refined capture these phenomena or could be considered along with dark matter models. For the current scope however, focus will be placed on dark matter models.

2.2 Axions and ALPs

As discussed in this section, the axion originally arose as a solution to the strong- CP problem in QCD. There are various mechanisms under which the axion can arise and it will generally be an extension of the Standard Model. However, one can consider a generalization of the axion as an axion-like particle (ALP) which has many of the same properties of the traditional axion but may not be a solution to the strong- CP problem. Axions that serve to solve the strong- CP problem are often called “QCD axions.” A useful review of axions is given in Ref. [25].

2.2.1 Strong- CP problem

The strong- CP problem concerns the fact that a CP -violating term is expected from QCD. In particular, the following theta-term should appear in the Standard Model Lagrangian [25],

$$\mathcal{L}_{\theta_{\text{QCD}}} := \frac{\theta_{\text{QCD}}}{32\pi^2} \text{Tr} G_{\mu\nu} \tilde{G}^{\mu\nu}, \quad (2.1)$$

where θ_{QCD} is a constant, $G_{\mu\nu}$ is the gluon field, and $\tilde{G}^{\mu\nu} := \epsilon^{\mu\nu\alpha\beta} G_{\alpha\beta}/2$ is its dual. The trace is taken over the adjoint representation of $SU(3)$. The term in Eq. (2.1) is a topological, or “total derivative,” term that will appear as a non-vanishing surface contribution to the action. However, experiments suggest that this term provides negligible contributions to the Standard Model, thus placing a very small upper-bound on θ_{QCD} .

Experimentally, a common way to measure the theta-term is through the neutron electric dipole moment (EDM). In particular, the neutron EDM is expected to be pro-

portional to θ_{QCD} [26–29]

$$|d_n| \approx \mathcal{O}(10^{-16} \text{ e cm}) \times \theta_{\text{QCD}}. \quad (2.2)$$

Measurements of the neutron EDM place the constraint, $|d_n| \lesssim \mathcal{O}(10^{-26} \text{ e cm})$ [30–32]. This leaves a constraint on the theta-parameter, $\theta_{\text{QCD}} \lesssim 10^{-10}$. Similar bounds can be made based on EDM measurements for ^{199}Hg [33].

Using “naturalness” arguments, one generally expects coupling constants to be $\mathcal{O}(1)$. Such a constraint points to potential new physics that suppress the effects of the theta-term $\mathcal{L}_{\theta_{\text{QCD}}}$; though it is still possible for $\theta_{\text{QCD}} \ll 1$ without additional phenomena, its unusually small value provides heuristic evidence of unknown physics. The “strong- CP problem” in QCD refers to the observation that the CP -violating theta-term, Eq. (2.1), appears negligible in experiments.

One should note, however, that the estimations for CP -violation rely on the case in which there are no massless quarks. One can show that the existence of a massless quark would, alone, preserve CP -invariance [34, 35]. However, barring the discovery of a new massless quark, this does not seem to be the case.

2.2.2 Peccei-Quinn mechanism

In 1977, R. D. Peccei and Helen R. Quinn proposed a solution to the strong- CP problem by introducing a dynamic pseudoscalar field that cancels the effects from θ_{QCD} [36, 37]. These ideas were expanded upon by F. Wilczek [38] and S. Weinberg [39]. To suppress the CP -violating effects, a complex (pseudo)scalar field ϕ is introduced. This field is equipped with the symmetry $U(1)_{\text{PQ}}$ associated with the complex phase $\phi \rightarrow e^{-2i\alpha}\phi$. This symmetry is associated with chiral rotations of fermions fields¹, $\psi \rightarrow e^{+i\alpha\gamma^5}\psi$ [37]. Further, the Lagrangian admits a potential term $V(|\phi|^2)$ which is minimized for $|\phi| = \frac{\bar{f}_a}{\sqrt{2}} \neq 0$. The magnitude $S := |\phi|$ is associated with a Higgs-mode while the phase is an axion-mode,

$$\phi := \frac{S}{\sqrt{2}} e^{ia/\bar{f}_a}, \quad (2.3)$$

where \bar{f}_a is the vacuum expectation value $\langle S \rangle$.

This new field will then couple to standard model fields such as quarks and leptons. This coupling will then induce an (effective) term

$$\mathcal{L}_{agg} := \frac{a}{32\pi^2 \left(\frac{\bar{f}_a}{N_{\text{DW}}}\right)} \text{Tr} G_{\mu\nu} \tilde{G}^{\mu\nu}, \quad (2.4)$$

which has the effect of making θ_{QCD} dynamic, $\theta_{\text{QCD}} \rightarrow \theta_{\text{QCD}}'(a)$. In particular, $\theta_{\text{QCD}}' = \theta_{\text{QCD}} + N_{\text{DW}}a/\bar{f}_a$ where N_{DW} , an integer, is the color anomaly or domain wall number².

¹A chiral rotation means that for spinor fields, the left-handed and right-handed components gain opposite phases. This is achieved through the addition of the γ^5 matrix in the exponent. The constant factor in the exponent defines the PQ charge of the field and could be zero for some fermions depending on the axion model.

²Elsewhere, such as Ref. [1], the Higgs-mode expectation value is defined as $S_0 := \langle S \rangle$ and $\bar{f}_a = S_0/N_{\text{DW}}$.

The value of N_{DW} depends on the particular model. Also, one could also re-define $a \rightarrow a - \bar{f}_a \theta_{\text{QCD}}$, or absorb this shift into S , so that $\theta_{\text{QCD}}' = N_{\text{DW}} a / \bar{f}_a$.

One can show that energy is minimized when $\theta_{\text{QCD}}' = 0$ [40–42]. In fact, the vacuum energy from the theta term is [40, 41],

$$E(\theta_{\text{QCD}}) \propto 1 - \cos(\theta_{\text{QCD}}') = 1 - \cos(N_{\text{DW}} a / \bar{f}_a). \quad (2.5)$$

Hence a theory with an axion will also have a vacuum that dynamically enforces strong CP -symmetry. Additionally, because a takes values in the range $[0, 2\pi\bar{f}_a)$ (i.e., to cover the complex phase in Eq. (2.3)), there are N_{DW} values for a that minimizes Eq. (2.5).

The QCD axion has an additional constraint relating the mass m_a and energy scale \bar{f}_a . This can be calculated using current algebra. Considering the contributions from the up and down quark³, the axion mass m_a and scale \bar{f}_a are related by [34, 43]

$$\begin{aligned} m_a &= \frac{f_\pi m_\pi}{2\bar{f}_a/N_{\text{DW}}} \left[\frac{4m_u m_d}{(m_u + m_d)^2} \right]^{1/2} [1 + \mathcal{O}(m_{u,d}/m_s)] \\ &\approx 5.4 \times 10^{-10} \text{ eV} \cdot \left(\frac{1.1 \times 10^{16} \text{ GeV}}{\bar{f}_a/N_{\text{DW}}} \right), \end{aligned} \quad (2.6)$$

where m_π is the pion mass, $f_\pi \approx 93 \text{ MeV}$ is the pion decay constant, m_u and m_d are the up- and down-quark masses respectively ($m_{u,d} \approx m_u \approx m_d$), and m_s is the strange-quark mass ($m_s \gg m_{u,d}$).

2.2.3 PQWW, KSVZ, and DFSZ axions

There are a few different models that introduce the axion to solve the strong CP problem. In each case, a complex scalar field is introduced with the $U(1)_{\text{PQ}}$ symmetry, the phase of which defines the axion field. Three such models are briefly described here: PQWW, KSVZ, and DFSZ.

PQWW axion

Perhaps the simplest axion model is the Peccei-Quinn-Weinberg-Wilczek (PQWW) axion [25, 36, 37]. In this model, the complex scalar field ϕ is introduced as a second Higgs doublet. In particular, one field couples to d -type quarks while the other couples to u -type quarks. There remains some choice as to which field couples to which quarks, as well as which field couples to leptons (or if there is an additional third Higgs-like field). The $U(1)_{\text{PQ}}$ chiral symmetry acts on this field via a complex phase. Likewise, parameterizing ϕ as in Eq. (2.3), the $U(1)_{\text{PQ}}$ symmetry acts on the axion field a as a constant shift. There is a potential acting on the new scalar field ϕ whose $U(1)_{\text{PQ}}$ symmetry is spontaneously broken, such as the conventional ‘‘sombbrero’’ potential,

$$V(\phi) \propto \left(|\phi|^2 - \frac{\bar{f}_a^2}{2} \right)^2. \quad (2.7)$$

³These quarks have the largest contribution to the axion mass as they are the lightest. The strange quark could also be included, though the contribution is very small.




Figure 2.1: (Left) Heavy quark triangle diagram connecting the axion a to the gluon field g that results in (Right) an effective term $\propto a \text{Tr} G_{\mu\nu} \tilde{G}^{\mu\nu}$ in the KSVZ model. A triangle diagram like the one on the left but with coupling to light is also of interest for canceling anomalies.

The vacuum expectation value is then $\langle \phi \rangle = \frac{\bar{f}_a}{\sqrt{2}}$

After electroweak symmetry breaking, the angular part of ϕ is left as a (pseudo) Goldstone boson, the axion a . The axion then couples to fermions via PQ charges with Lagrangian terms of the form (up to possibly a constant prefactor) [44],

$$\mathcal{L} \supset (a/\bar{f}_a) m_\psi \bar{\psi} i \gamma^5 \psi. \quad (2.8)$$

This coupling results in a triangle diagram (similar to that in Fig. 2.1) which generates the term given in Eq. (2.4) that acts to cancel the chiral anomaly.

KSVZ axions

The Kim-Shifman-Vainshtein-Zakharov (KSVZ) axion model includes the complex scalar ϕ (and symmetry-breaking potential, Eq. (2.7)) described in the PQWW model as well as a new heavy quark doublet Q_L and Q_R , both $SU(3)$ singlets and whose subscripts denote their respective PQ charge [25, 34, 45]. In this model, the complex scalar ϕ (and the axion a) only couples to the heavy quark Q . That is, the Lagrangian contains the Yukawa term giving mass $\sim \bar{f}_a$ to Q ,

$$\mathcal{L}_{\phi QQ} := f \bar{Q}_L \phi \bar{Q}_R + f^* \bar{Q}_R \phi^* \bar{Q}_L,$$

for constant f , but no such Yukawa terms for standard model particles. In particular, there are no tree-level couplings of the axion to standard model fields such as quarks, leptons, photons, and gluons. The heavy quark, however, does interact with both the axion and gluons. The axion-gluon interaction then arises as an effective term from the triangle diagram shown in Fig. 2.1. In the KSVZ model, one takes $N_{\text{DW}} = 1$ as there is only a single, unity-PQ-charge fermion Q [25].

The lack of direct interactions between the KSVZ axion and Standard Model particles leads to suppressed interactions [34, 45]. This means that experiments that would disprove the PQWW axion do not necessarily exclude the KSVZ axion. There could also be an effective electromagnetic coupling $a \rightarrow \gamma\gamma$ depending on the electromagnetic charge of Q (which is often assumed to be zero). Cavity experiments, for example, can probe this coupling for various heavy-quark charges and representations based on the $a \rightarrow \gamma\gamma$ coupling [46].

The KSVZ axion is not so readily excluded by experiment as the PQWW axion. Essentially, this is because the KSVZ model introduces a new field that loosens the connection between the axion and the Standard Model. The axion can then have both a lower mass and a weaker coupling to ordinary matter. A weaker coupling can explain why it would be difficult to observe with current experimental methods.

DFSZ axions

The Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) model was first proposed by A. R. Zhitnitsky [47] and later, independently by M. Dine, W. Fischler, and M. Srednicki [48]. The DFSZ model builds on the ideas of the KSVZ model, except that an additional Higgs field is introduced instead of a heavy quark. A complex scalar ϕ (and symmetry-breaking potential, Eq. (2.7)) whose phase describes the axion, as well as two Higgs doublets H_u and H_d . The H_u Higgs couples with the u -type quarks and H_d Higgs couples with the d -type quarks via the typical Yukawa terms. There is some choice as to which Higgs field couples to leptons; for example, DFSZ's paper [48] has H_d couple to the leptons.

The axion couples to the standard model via the Higgs sector with the scalar potential [25]

$$V_{\text{DFSZ}} = \lambda_H \phi^2 H_u H_d, \quad (2.9)$$

for some coupling constant λ_H . The H_u , H_d , and ϕ have charges X_u , X_d , and X_ϕ , respectively, under $U(1)_{\text{PQ}}$, so physics is invariant under the transform (for any constant α)

$$H_u \rightarrow e^{i\alpha X_u} H_u, \quad H_d \rightarrow e^{i\alpha X_d} H_d, \quad \phi \rightarrow e^{i\alpha X_\phi} \phi.$$

In order for Eq. (2.9) to be invariant under $U(1)_{\text{PQ}}$, one must demand $X_u + X_d + 2X_\phi = 0$. For example, $X_u = X_d = +1$ and $X_\phi = -1$. The DFSZ model works under the case in which the vacuum expectation value for the complex scalar field is much larger than that of the Higgs fields,

$$\langle \phi \rangle = \bar{f}_a / \sqrt{2} \gg \sqrt{\langle H_u \rangle + \langle H_d \rangle}. \quad (2.10)$$

Because quarks and leptons couple to the Higgs fields via Yukawa coupling, the quark and lepton doublets must also be charged $X_{u,d}$ under $U(1)_{\text{PQ}}$. The current for this charge is then (here, H_d couples to the leptons)

$$\begin{aligned} j_\mu^{\text{PQ}} = & X_\phi \phi^* \overleftrightarrow{\partial}_\mu \phi + X_u H_u^\dagger \overleftrightarrow{\partial}_\mu H_u + X_d H_d^\dagger \overleftrightarrow{\partial}_\mu H_d \\ & + X_u (\bar{u} \gamma_\mu \gamma^5 u + \bar{c} \gamma_\mu \gamma^5 c + \bar{t} \gamma_\mu \gamma^5 t) \\ & + X_d (\bar{d} \gamma_\mu \gamma^5 d + \bar{s} \gamma_\mu \gamma^5 s + \bar{b} \gamma_\mu \gamma^5 b) \\ & + X_d (\bar{e} \gamma_\mu \gamma^5 e + \bar{\mu} \gamma_\mu \gamma^5 \mu + \bar{\tau} \gamma_\mu \gamma^5 \tau). \end{aligned} \quad (2.11)$$

After electroweak symmetry breaking, this current leads to effective terms $\sim m_q (a/\bar{f}_a) \bar{q} i \gamma^5 q$ which, as before, lead to an anomalous term $\sim a G \tilde{G}$ via quark triangle diagrams like the one on the left of Fig. 2.1. Similarly, an anomalous term $\sim a F \tilde{F}$ appears due to triangle

diagrams with quarks and leptons coupling to photons. For the DFSZ model, $N_{\text{DW}} = 6$ for the six quarks.

In contrast to the KSVZ model in which axions couple to the Standard Model via an effective term from a loop diagram (and not to the Standard Model quarks), the DFSZ couples to the Standard Model quarks at tree-level. There would also be couplings between the axion and leptons similar to the axion-quark coupling because the leptons also gain a PQ charge [49]. Additionally, the DFSZ model is motivated by results that arise naturally in $SU(5)$ grand unification techniques [25, 46].

2.2.4 Axions in string theory

Axions — or, at the very least, ALPs — notably appear naturally in various realizations of string theory. In fact, many such particles may exist as a “string axiverse” [50]. Though a detailed exploration of these theories will not be given here, axions arise in string theory from the Kaluza-Klein zero-modes in the antisymmetric tensor when compactifying dimensions in the theory [43, 50, 51]. One finds that this mechanism can generate hundreds or even thousands of axion-like fields [50].

2.2.5 Current constraints

Numerous experiments and observations have been used to search for and constrain ALP dark matter. Astrophysical constraints, in particular, provide strong bounds on the characteristics of ALPs. Further discussion of dark matter experiments is provided in Sec. 2.5.

One of the most stringent bounds on ALP dark matter comes from supernova measurements. In particular, neutrino measurements from SN1987A, a supernova observed in 1987, provides strong bounds on ALPs and similar dark matter candidates [52–54]. These constraints follow from considering how the existence of ALPs would provide a cooling channel for hot objects as particles are produced and escape the object. For supernova with core temperatures around 30 MeV, particles with mass below about a few hundred MeV can have a significant cooling effect [54], though the extent of the cooling depends on how strongly the particle couples to the star matter. Neutrino measurements from SN1987A are consistent with cooling times from the Standard Model predictions [55] and can be used to place a conservative bound on the axion decay constant, [52, 54]

$$\bar{f}_a \gtrsim 10^8 \text{ GeV}. \quad (2.12)$$

If the decay constant were below this bound, the ALP luminosity would substantially affect the neutrino signal. It should be noted that the above constraint is somewhat model-dependent based on how axions couple to matter in neutron stars, and different models can also be considered [54, 56, 57]. Similar, though weaker, astrophysical bounds can be made based on the lifetime and luminosity of red giants, white dwarfs, etc.; see, e.g., the review in Ref. [53].

Various models have been proposed that could raise doubt on the constraint in Eq. (2.12). First, one may raise some questions about the exact mechanism under-

lying supernovae. For a rotating star undergoing a supernova, an accretion disk can form which would not have the same axion cooling as the star's denser core [58]. Neutrinos originating from the accretion disk could then arrive later than neutrinos originating from the core, alone. This would yield a deceptively long cooling time for the supernova as seen in neutrino measurements, thus allowing for a stronger axion coupling (or smaller decay constant) than Eq. (2.12). In addition, one may introduce new physics that suppresses the influence of ALPs in stars and supernovae. One such model would be to introduce an additional field that couples to ALPs and the Standard Model such that it causes the ALP to have a large effective mass in a dense environment [59]. The increased effective mass would then suppress ALP production, reduce the cooling effect, and weaken the astrophysical constraints.

Another bound on ALP dark matter arises from cosmology. Namely, the energy density of axions is limited by the critical density needed to close the universe. This establishes a bound on the axion decay constant [60–62]

$$\bar{f}_a \lesssim 10^{12} \text{ GeV}, \quad (2.13a)$$

$$\bar{f}_a \lesssim \frac{10^{12} \text{ GeV}}{\sqrt{m_a/10^{-5} \text{ eV}}}, \quad (2.13b)$$

where Eq. (2.13a) uses Eq. (2.6) to eliminate the mass-dependence. This bound relies on considering the dynamics of the expectation value $\langle a \rangle$ with the ansatz

$$\langle a \rangle = A(t) \cos(m_a t)$$

and determining how this evolves in an expanding universe. In particular, the bound (2.13) assumes an initial condition that $A(t_0)/\bar{f}_a = \mathcal{O}(1)$ when the temperature of the universe was $T_0 \approx 800 \text{ MeV}$. Additionally, axion influence on nucleons would affect the production of ${}^4\text{He}$ during Big Bang nucleosynthesis, which also provides cosmological bounds [63].

More modern cosmological measurements provide additional constraints. Measurements of the cosmic microwave background from the Planck telescope were used along with other measurements to determine an upper-bound on the axion mass and the density of cold dark matter [64, 65]

$$m_a < 0.529\text{--}2.09 \text{ eV}, \quad (2.14a)$$

$$\Omega_{\text{cdm}} h^2 \approx 0.12, \quad (2.14b)$$

where the mass bound depends on the data and model used in the calculation.

2.3 WIMPs

A common dark matter alternative to axions or ALPs is a class of particles known as Weakly Interacting Massive Particles (WIMPs). As the name suggests, these particles are massive — in contrast to ALPs which are typically ultralight — with weak interactions — as one would often expect from dark matter. Like axions, WIMPs arise naturally

from theories intended to solve an unrelated problem in physics; namely supersymmetry (SUSY) which can, in part, solve the hierarchy problem. Reviews of SUSY and WIMPs can be found in Ref.s [66–68] and are used to describe background information here.

The hierarchy problem concerns the vast discrepancy between the typical energy scales of different aspects of physics; namely between the electroweak ($\sim 10^2$ GeV) and Planck scales ($\sim 10^{19}$ GeV). Meanwhile, SUSY is a theory that extends the Standard Model by introducing a symmetry between bosons and fermions. As a result, each Standard Model fermion will have a bosonic superpartner and vice-versa. The Higgs particle mass, which reflects the electroweak energy scale, will be influenced by corrections from various Standard Model particles. However, with the introduction of SUSY, the corrections from particles and their super-partners will cancel, resulting in a massless Higgs boson. If SUSY is broken — which the lack of observed superpartners with the same mass seems to suggest — the Higgs boson will have some small mass tied to the symmetry-breaking scale. This consequentially implies that SUSY effects should be observable at energy scales of a few TeV [66, 69, 70]. It should be noted that many of the early foundational papers for SUSY [71–76] did not aim to solve the hierarchy problem and that SUSY is also often discussed in the context of string theory as its inclusion fixes certain inconsistencies with the theory.

WIMPs can naturally arise as the lightest supersymmetric particle (LSP) in a theory with SUSY. Notably, a symmetry known as “ R -parity” will preserve LSPs in a similar manner as how the lightest Standard Model particles cannot decay into lighter particles and are, thus, preserved. The R -parity assigns opposite charges to particles and their superpartners and conservation of this charge will prevent LSP from decaying into lighter, Standard Model particles. In typical SUSY models, LSPs are superpartners to some superposition of neutral, gauge bosons.

More motivation for WIMPs arises when considering cosmological factors. One can consider the presence of some heavy particle χ which annihilates or forms from light Standard Model particles, ℓ via $\bar{\chi}\chi \leftrightarrow \bar{\ell}\ell$. After the Universe cools below temperatures $T \sim m_\chi$, the annihilation interaction $\bar{\chi}\chi \rightarrow \bar{\ell}\ell$ will dominate, causing the density of χ to drop exponentially when in thermal equilibrium. The decay rate is $\Gamma_\chi = \langle\sigma_{Av}\rangle n_\chi$; where n_χ is the number density for χ and $\langle\sigma_{Av}\rangle$ is the thermally averaged product of the total cross-section and relative velocity. However, the Universe also expands at a rate defined by the Hubble parameter; conventionally denoted in units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ as h . When $\Gamma_\chi \lesssim h$, the annihilation “freezes-out” and a relic density of χ remains above what one would expect from thermal equilibrium. For the present critical density of the universe, $\rho_c \approx 10^{-1} h^2 \text{ GeV/cm}^3$, the density parameter for χ is given by [66]

$$\Omega_\chi h^2 = m_\chi n_\chi / \rho_c \approx 3 \times 10^{-27} \text{ cm}^3 \text{ s}^{-1} / \langle\sigma_{Av}\rangle . \quad (2.15)$$

Moreover, if the cross-section is dominated by weak interactions,

$$\langle\sigma_{Av}\rangle \sim \mathcal{O}(10^{-25} \text{ cm}^3 \text{ s}^{-1}) . \quad (2.16)$$

This implies that a WIMP associated with the electroweak force will inherently have a relic abundance that is very close to what you would expect for dark matter. The

coincidental agreement between cosmological relic density, the electroweak energy scale, and the observed dark matter density — phenomena that are not inherently linked — is known as the “WIMP miracle” and has motivated many searches for WIMPs.

Despite extensive experimental efforts, definitive evidence for WIMPs or SUSY has not been found. In particular, the first two experimental runs at the LHC have not found evidence for SUSY at beamline energies up to 14 TeV; placing significant constraints on theories of SUSY [70]. However, there remains some hope that some form of SUSY may be detected in a future 100 TeV accelerator [69, 77]. In addition to accelerator searches, experiments search for the scattering products from WIMPs interacting with nucleons; such as the XENON experiment discussed in Sec. 2.5.

2.4 Other dark matter candidates

There are numerous other dark matter candidates aside from axions and WIMPs. Here, a few examples are listed with a brief overview. Many of the dark matter candidates will also have modified versions not covered here. It is also entirely possible for dark matter to consist of a combination of different candidates. Even without discovery, continued experimental searches for dark matter and other exotic physics can constrain and rule-out some of dark matter candidates as well as inspire new models.

Neutrinos

Considering only Standard Model particles, one may initially suspect that dark matter could consist of neutrinos. These particles are electrically neutral and only weakly interact with other Standard Model particles, making them an excellent candidate. Though the Standard Model assumes that neutrinos are massless, observation of neutrino flavor oscillation [78–82] suggests that at least some neutrinos have mass.

However, further considerations find inconsistencies with neutrino dark matter and galactic-scale structure formation [83] and sky map data combined with cosmology models supports a relatively small contribution from dark matter neutrinos [84].

Dark photons

By introducing an additional $U(1)$ symmetry to the Standard Model, one obtains a new photon-like particle called the “dark photon” [85, 86]. The dark photon is similar to the photon except that it is associated with its own charge and can possess a mass. This is a bosonic candidate for dark matter as it can lack strong interactions. Additionally, dark photons would be produced during inflation [87–89].

Experimental studies and constraints on dark photons have been considered using the decay of strange hadrons [90], $g - 2$ measurements [90, 91], and cavity experiments [92]. New dark photon experiments continue, such as an exploratory search for dark photons with a global network of magnetometers using the Earth as a transducer [93, 94].

Relaxions

A potential ALP candidate for dark matter is the relaxion [95]. This particle was proposed as a solution to the hierarchy problem in which the ALP field couples to the Higgs field, allowing the Higgs mass dynamic scan over a range of values in the early universe. The relaxion field “rolls” down a bumpy potential. The bumps in the effective relaxion potential will grow when the Higgs mass is small. Combined with inflation which produces Hubble friction, allowing the relaxion field to stop rolling, this will solve the hierarchy problem as the Higgs mass will naturally settle at a small value as dynamics ease even if it began at some large cutoff scale. It should be noted that, unlike typical realizations of the QCD axion, the relaxion field is non-compact, though it retains an (approximate) discrete shift symmetry as the bumps in the potential are evenly spaced.

In addition to solving the hierarchy problem and possibly the strong- CP problem, it is also possible for relaxions to constitute dark matter [96,97]. Likewise, one could have gravitationally bound relaxion stars forming a halo around the Earth or Sun [98].

Quark nuggets

In 1984, Edward Witten described a mechanism by which dense quark matter could exist to the present day in the form of “quark nuggets” [99] which could explain dark matter.

Formation of quark nuggets arises naturally from a first-order phase transition in the vacuum. During the transition, regions of space change into the “low-temperature” state that gradually grow through nucleation⁴ within a “high-temperature” background state similar to how ice might form in cold water. Eventually, the low-temperature regions grow enough to dominate the universe, though high-temperature bubbles remain. If these bubbles cool further, not through evaporating baryons, but through, e.g., emitting neutrinos (that result from $q\bar{q}$ annihilation), then they could form a type of dense quark matter with a large (positive or negative) baryon number as macroscopic quark nuggets.

Unlike many other candidates for dark matter, quark nuggets would actually interact fairly strongly with other SM particles. However, a single quark nugget is dense enough that it would have a fairly small cross-section relative to its mass and massive enough to have a low number density; making interactions with ordinary matter rare enough to still be a dark matter candidate. Coincidentally, if the quark nuggets tend to consist of anti-quarks rather than quarks, this could explain the matter-antimatter asymmetry observed in the universe; i.e., the antimatter is “locked up” in quark nuggets [100–102].

Experimental constraints on quark nuggets have been considered using a variety of methods; such as neutrino flux from the sun [100], their acoustic signals when passing through the Earth [101, 103], through interactions with axions [102], as well as gravitational microlensing, gravitational waves, and cosmic rays [104].

⁴Witten also considers the case in which the transition occurs through spinodal decomposition instead of nucleation.

Primordial black holes

Unlike black holes that form at the end of the life-cycle of large stars, primordial black holes are hypothesized to form during the early universe [105,106]. The exact conditions under which these black holes form varies with different models, but the main idea is that early in the Universe, the typical energy density of space was large enough that small perturbations could cause enough energy to coalesce into a black hole.

It is possible for primordial black holes to contain at least some of the missing, dark matter in the universe. Similar to quark nuggets, black holes are very dense objects that could avoid direct detection due to their size and/or rarity, but still constitute a substantial amount of mass. A review of primordial black holes as dark matter can be found in Ref. [107] along with various astrophysical constraints.

Due to Hawking radiation, the black holes will gradually decrease in mass [108]. This condition places a lower-limit on the mass of primordial black holes to survive until the present day at $\gtrsim 10^{15}$ g. However, if these black holes formed at some time after Planck time $\sim 10^{-43}$ s, their mass would be at least 10^{-5} g [106,107]. The exact mass of primordial black holes that would survive to the present day is very model-dependent and may include a range of masses [107].

2.5 Experimental searches

Numerous experiments have been designed to better understand the nature of dark matter. Each experiment focuses on some type of interaction or affect that dark matter may have with visible matter. A few of such experiments are briefly described here. This by no means describes all dark matter experiments, and there are plenty of experiments whose data has/can be used to search for dark matter that were not initially designed to do so. However, this should provide an idea for various techniques being used.

The Cosmic Axion Spin Precession Experiment (CASPER) measures the effect of the ambient dark matter field on nuclear spin precession [109–113]. Using nuclear magnetic resonance techniques, CASEPER is a series of experiments able to precisely measure the affects of dark matter on various atoms and molecules through different coupling schemes. For example, a dark matter field can induce a coupling to matter similar to a magnetic field that oscillates at the dark matter’s Compton frequency. By matching this frequency with the Larmor frequency induced by a leading magnetic field, a resonance occurs that can be measured with magnetometers. CASPER has been used to search for ALPs [111–113] as well as dark photons [112].

The interaction between axions/ALPs and Standard Model particles can also extend to macroscopic scales. For example, using a magnetic torsion pendulum, one can measure the effects of galactic halo dark matter [114]. Measuring the collective effect of axions interacting with ordinary matter enables one to consider these macroscopic experiments; albeit still requiring precision techniques to measure the effects.

The ABRACADABRA experiment considers a magnetic field caused by axion-photon coupling [115,116]. In this experiment, a toroidal solenoid is used to produce a magnetic

field. A background axion field oscillating at its Compton frequency will result in an oscillating magnetic field circling the toroid which can be measured.

Another avenue for detecting dark matter is to convert it into visible matter, e.g., photons. The CAST [117] (and its proposed upgrade, IAXO [118]) and the proposed MADMAX [119] experiments work on this principle. In particular, a strong magnetic field is used to induce ambient light dark matter (e.g., axions/ALPs and dark photons) to transition into photons which can be detected. The CAST experiment focuses on dark matter generated in the Sun, while MADMAX focuses on dark matter present in the galactic halo.

Similarly, optical cavity experiments measure disturbances from axions on the cavity. Both the ADMX [120] and HAYSTAC [121] experiments search for galactic halo dark matter using this technique. The cavity is placed in a strong magnet, and the presence of an axion field induces an electromagnetic effect, disturbing the light in the cavity. Cavities have also been proposed as potential probes of apparent variations in fundamental constants due to dark matter interactions [7].

Light-shining-through-a-wall experiments consider an additional step: light is converted into dark matter before being converted back to light. The basic principle of a light-shining-through-a-wall experiment is that a light source is directed towards an opaque wall with a detector on the opposite side. If the detector observes a signal from the light source, that would imply that the light underwent some exotic interaction to some other particle (e.g., an axion) in order to traverse the wall. These experiments generally search for relatively light particles as heavier particles are difficult to create with a light source. For example, the ALPS experiment [122, 123] focuses on searches for axions/ALPs and dark photons. A strong magnetic field is applied on either side of the wall to induce transitions into and from dark matter.

As highly sensitive instruments, atomic clocks could also provide an avenue for dark matter detection. For example, the GPS.DM project uses the data from atomic clocks on board GPS satellites to search for macroscopic dark matter features (e.g., axion domain walls) in the galactic halo [124, 125]. As these features cross the Earth, one would expect “glitches” to propagate across the network. This experiment is able to capture useful information from archival data of existing devices.

The XENON project is a series of experiments using large, underground chamber of liquid xenon to detect dark matter [126]. As dark matter particles (particularly WIMPs) traverse the chamber, they can scatter off of xenon nuclei; resulting in photons and electrons. Photomultiplier tubes measure the photons, while an electric field across the chamber pulls electrons to the top of the chamber where they are measured by an array of detectors. By combining the delay between the photon and electron signal with the position of the detectors that measure the signal, the position of the scattering can be calculated. Events can also be distinguished by the relative strength of their signals. Notably, recent results from XENON1T — an iteration of the experiment with 2 tonnes (2.0×10^3 kg) of xenon in the detection volume — shows promising evidence of solar axions with mass $m_a \approx 2.3 \text{ keV}/c^2$ [127].

2.6 Macroscopic dark matter objects

Under certain circumstances, dark matter can form into stable macroscopic objects or features. Two notable examples will be discussed here: domain walls and boson stars. The domain wall is of particular interest for this work as GNOME attempts to observe these features. Though some effort has been made to observe boson stars with GNOME, as well [6].

2.6.1 Domain walls

Domains and domain walls are stable structures that form in a field with degenerate local vacuum states; that is, there are multiple states that a region of space can have that minimizes energy. A “domain” is a region of space with a given vacuum state, and “domain walls” separate different domains. A domain wall inherently contains some amount of energy because spatially transitioning from one vacuum state to another implies that a domain wall locally deviates from the vacuum state. Despite a system with a domain wall having more energy than a system without one, a domain wall can still be stable, especially in the absence of any biasing force that would cause one domain to be preferred (e.g., have lower energy) than the other, thus causing the preferred domain to expand and push away the domain wall. However, even if there are no preferred domains, a domain wall will have some tension that can lead to dynamics that minimize the surface area of domains.

Though domain walls can be stable, there is some tendency for a system to dynamically remove them. This then begs the question of whether these structures are truly stable on any meaningful timescale. The simplest answer comes from the observation that there do exist well-documented systems in which domain walls form; such as in ferromagnets.

Ferromagnets provide a common example of domains and domain walls, in part because the domains can be easily viewed with Kerr magnetometry. Ferromagnetic materials can be modeled by a lattice of small magnets from atomic spins. These materials have a rotational symmetry that is spontaneously broken in the presence of a biasing magnetic field; including the magnetic field generated by a region of the material in which the spins happen to align. In particular, spins tend to align with and amplify the local magnetic field. As a result, if one starts from random disorder, a ferromagnet will tend to order itself into domains as long as it is below the transition temperature at which its symmetry is broken. As different domains expand into disordered regions of the material, they may be oriented significantly differently than their neighboring domains and be unable to further expand. This results in a material speckled with regions with different orientations that are separated with a network of domain walls.

For this work, domain walls formed in an axion or ALP field is of particular interest. If $N_{\text{DW}} > 1$, ALPs will have multiple, unique vacuum states and can form domains separated by domain walls in space. It is possible that the energy contained in domain walls contributes appreciably to the dark matter energy density of the galaxy. In this picture, the galaxy could, for example, be contained in a foam-like network of domain

walls.

Even if $N_{\text{DW}} = 1$, domain walls can still form because of the topology of possible axion field values⁵. Recall that the axion arises as the phase of a complex field. A complex phase can be mapped to a circle S^1 ; typically with the axion taking on a real-valued field, $a(\mathbf{x}) \in \mathbb{R}$ but with the identification $a \cong a + 2\pi\bar{f}_a$. Because of this topology, a domain wall could separate the (identical) vacua a_0 and $a_0 \pm 2\pi\bar{f}_a$. While in the $N_{\text{DW}} > 1$ case, a finite domain wall forms a closed surface (e.g., a sphere with different vacuum states on the inside and outside), a finite domain wall in an $N_{\text{DW}} = 1$ would form an open surface with the same vacuum on either side. The edge of this surface is called the “axion string.” The axion field in a loop containing an axion string will continuously circle through the phase of the underlying complex scalar field.

Though ALP domain walls could be a promising potential avenue for dark matter searches, there has been some effort to develop models that avoid them altogether. The so-called “domain wall problem” refers to the inconsistency between the existence of axion domain walls and cosmological observation. In particular, the energy contained in axion domain walls would exceed the critical density for the universe [128,129]. There are a few ways to avoid the issue. First, one could imagine a scenario in which domain walls formed prior to inflation so that the edge of our domain lies outside of the observable universe [129]. Another solution is simply that the domain walls are especially unstable and will annihilate early on, possibly into dark matter [129–132]. This is a common motivation for models with $N_{\text{DW}} = 1$ (e.g., typical KSVZ axions) because the axion strings around the domain wall will generally exert some tension that shrinks the domain wall. It is also possible that there is simply a small bias at the cost of lightly breaking the PQ mechanism [128,129,131,133–135]. Finally, one can avoid the problem by simply considering ALPs that do not have the same mass restrictions as the axion to cause the domain wall problem in the first place.

For the most part, the specific mechanism by which some type of ALP domain wall could persist today will not be relevant to this work. The regime being considered here would be one in which the domain walls contain a relatively small energy density — on the order of dark matter, not exceeding critical density of the universe. Likewise, it is possible that a network of such domain walls would collapse under its own surface tension. It will be assumed that there is some mechanism to prevent collapse; possibly with an additional dark matter candidate. For example, pressure from the other dark matter resists domain walls akin to air preventing a balloon from collapsing.

A more quantitative discussion of domain walls is given in Chapter 3.1 and Appendix A.1.

2.6.2 Boson stars and Q -balls

Boson stars are generically a macroscopic, spherical, and stable field configuration of bosons. These configurations can be maintained through gravity, self-interaction, topol-

⁵These same principles generally apply to ALPs as well, but not in all models. For example, the relaxation field is described fully by a real number and not a value on a closed topology.

ogy, etc., and the exact shape and properties of these features can vary. One can consider boson stars made of, e.g., a gravitationally bound Bose-Einstein condensate of axions [136–138] or relaxions [98], though the stability of such features has also been questioned [139].

A feature related to the boson star is the Q -ball [140–142], which is a soliton that can arise in certain complex scalar fields. Further, the scalar field ϕ must possess a $U(1)$ symmetry $\phi \rightarrow e^{-i\theta} \phi$. The Lagrangian for this field has the form

$$\mathcal{L}_Q = \frac{1}{2} \partial^\mu \phi^* \partial_\mu \phi - U(\phi), \quad (2.17)$$

where the potential $U(\phi)$ has a global minimum at $\phi = 0$ and is invariant with respect to phase rotations of ϕ . The Noether's current and charge associated with the $U(1)$ symmetry are given by

$$j_Q^\mu = \frac{1}{2i} \phi^* \overleftrightarrow{\partial}^\mu \phi, \quad (2.18a)$$

$$Q = \int d^3\mathbf{x} j_Q^0. \quad (2.18b)$$

A Q -ball can be further understood by using a simple ansatz, $\phi(t, \mathbf{x}) = \tilde{\phi}(\mathbf{x}) e^{i\omega t}$ for the real, spatial function $\tilde{\phi}(\mathbf{x})$ and real constant ω . A stable solution can be obtained by minimizing the energy associated with Eq. (2.17) when inserting the ansatz,

$$E_Q = \int d^3\mathbf{x} \left[\frac{\omega^2}{2} |\tilde{\phi}|^2 + \frac{1}{2} |\nabla \tilde{\phi}|^2 + U(\tilde{\phi}) \right].$$

Consider a further restriction on the ansatz that the field is only non-zero in the region \mathcal{V} ,

$$\tilde{\phi}(\mathbf{x}) = \begin{cases} \phi_0 & \mathbf{x} \in \mathcal{V} \\ 0 & \mathbf{x} \notin \mathcal{V} \end{cases}, \quad (2.19)$$

To make the above ansatz continuous, one can allow the field to interpolate linearly with a slope of ϕ_0/ϵ within $\epsilon/2$ of the border $\partial\mathcal{V}$ for $\epsilon \rightarrow 0$. Plugging the ansatz (2.19) into Eq. (2.18), the charge is given by $Q = \omega V \phi_0^2$ for volume $V = \int_{\mathcal{V}} d^3\mathbf{x}$. The energy of the ansatz solution is given by

$$E_Q = \frac{Q^2}{2V\phi_0^2} + VU(\phi_0) + \frac{A\phi_0^2}{2\epsilon},$$

where A is the area of $\partial\mathcal{V}$ and the last term reflects the energy contributions from the boundary. The boundary term diverges as $\epsilon \rightarrow 0$. However, because the contribution is proportional to the area, it is clear that the \mathcal{V} that minimizes energy will also minimize surface area; i.e., \mathcal{V} is a sphere with some radius R . First, minimizing energy with respect to volume, $V = \sqrt{\frac{Q^2}{2\phi_0^2 U(\phi_0)}}$. Plugging this back into the energy, one obtains $E_Q/|Q| = \sqrt{2U(\phi_0)/\phi_0^2}$. Thus, if $U(\phi_0)/\phi_0^2$ has a minimum for $\phi_0 \neq 0$, one can find

a non-trivial solution to the simple ansatz presented here. This results in a simple relationship between the radius R of a Q -ball and physical parameters

$$R = \left(\frac{9Q^2}{32\pi^2\phi_0^2 U(\phi_0)} \right)^{1/6}. \quad (2.20)$$

The ansatz (2.19) described here is valid for a “thin-wall” Q -ball. However, the ansatz fails when $U(\phi)/\phi^2$ does not obtain a minimum for finite ϕ , one can still obtain a Q -ball by considering an ansatz with “smoother” walls. This scenario can occur, for example, if $U(\phi) \propto \phi^p$ for large ϕ and $p < 2$. This is a “thick-wall” Q -ball [142, 143]. Further studies of Q -ball shapes can be found in, e.g., Ref. [144].

2.7 Standard halo model

The standard halo model (SHM) provides a simple description of dark matter in the galaxy. The model assumes that dark matter exists as a homogeneous cloud encompassing the visible matter in the galaxy. Though simple, this model sufficiently explains the distribution of dark matter and its observed effects on galaxies.

In the SHM, dark matter is treated as a virialized (i.e., stable and gravitationally bound) cloud of point-like particles. For this work, we are particularly interested in surface-like features (i.e., domain walls), so the discussion will be extended to include such objects. In particular, a “foam” of domain walls whose bubbles contain different domains will be considered. A similar study of the SHM can be found, e.g., in Appendix A of Ref. [125]. Discussions of the SHM with point-like particles can be found in Ref. [145].

It is important to understand the typical speed and direction of dark matter objects passing the Earth. To do this, one can describe a probability distribution function (PDF) in “velocity space.” Further, an event probability distribution function (EPDF) derived from the PDF to distinguish the distribution of object velocities as they are observed from the velocities of all objects. General details on manipulating PDFs are explored in Appendix B.

To begin, consider the center-of-mass velocity \mathbf{v}_0 distribution of an ideal gas,

$$dP_{v_0} = (2\pi s^2)^{-3/2} e^{-\frac{\|\mathbf{v}_0\|^2}{2s^2}} d^3\mathbf{v}_0, \quad (2.21)$$

where s is the dispersion velocity and $d^3\mathbf{v}_0 = dv_{x0}dv_{y0}dv_{z0}$.

Inside this cloud, the Earth travels with the local velocity \mathbf{v}_L — which is approximately the galactic orbital velocity \mathbf{v}_c when the velocity of the Earth around the sun and us around the Earth is neglected. Further the dispersion velocity will satisfy⁶ $2s^2 \approx v_c^2$. In the Earth’s reference frame, the velocity is $\mathbf{v} = \mathbf{v}_0 - \mathbf{v}_L$ and the PDF, Eq. (2.21), becomes

$$dP_v = (\pi v_c^2)^{-3/2} e^{-\frac{\|\mathbf{v} + \mathbf{v}_L\|^2}{v_c^2}} d^3\mathbf{v}. \quad (2.22)$$

⁶A heuristic argument for this is that the galactic orbital speed v_c is the speed at which an orbit is stable. For a thermal gas to be stable, it should likewise be related to the same speed.

This, alone, is a good approximation for the velocity distribution for a gas of point-like particles.

One further correction can be made by demanding the speed be below the galactic escape velocity v_e . The effects of dark matter escaping the galaxy can be included into the PDF by multiplying Eq. (2.22) by a factor proportional to the Heaviside function $\Theta(v_e - \|\mathbf{v}_0\|) = \Theta(v_e - \|\mathbf{v} + \mathbf{v}_L\|)$. When doing so, one must also renormalize the equation to correct for the component of the probability lost at large speeds. However, demanding particles in the “gas” of the SHM remain below escape velocity implicitly assumes that no particles will enter the galaxy above escape velocity — only that they leave the galaxy. A full understanding of the contribution of high-speed objects “evaporating” out of the galaxy would also require an understanding of dark matter in intergalactic space.

For this work, it is particularly interesting to consider the case in which the objects are surface-like instead of point-like; particularly for observing domain walls. These objects are considered to be sufficiently large that they appear to be flat planes across the Earth. The large-scale structure of the surfaces — i.e., their shape, how they separate different domains, etc. — will not be considered here. Instead, we assume that the local orientation of these surfaces as they cross the Earth is random and isotropic. Further, one is only able to observe the component of the velocity perpendicular to the object’s surface v_\perp . However, there still should be some unobservable velocity component parallel to the surface. This is clear if the surfaces are of a large object, e.g., a spherical domain, which will have a three-dimensional velocity even if it cannot be observed locally⁷.

For surface-like objects, it becomes necessary to consider how the coordinate system is defined. For point-like particles, there are two vectors of interest: the velocity \mathbf{v} and the Earth’s velocity \mathbf{v}_L . It can then be useful to think of the velocity distribution in spherical coordinates with the Earth’s velocity as the pole. However, a surface will additionally have a perpendicular velocity \mathbf{v}_\perp that describes its orientation, and it is more natural to describe this vector in reference to the total velocity \mathbf{v} .

The coordinate system used here is given in Fig. 2.2 and described below. Before concretely defining all of the coordinate parameters, it helps to introduce a short-hand notation here that will simplify some of the equations. Namely, for a general polar angle θ , let the barred variable denote

$$\bar{\theta} := \cos \theta.$$

Denote the angle between the velocity \mathbf{v} and perpendicular velocity \mathbf{v}_\perp as $\eta \in [0, \pi/2]$ ⁸ so

$$v_\perp = v\bar{\eta},$$

where the non-bolded velocities denote the speed: $v := \|\mathbf{v}\|$ and $v_\perp := \|\mathbf{v}_\perp\|$. In the end, the direction of the (observable) perpendicular velocity is described with respect to the reference \mathbf{v}_L . The local velocity \mathbf{v}_L can then be understood in terms of the galactic

⁷Even if the objects are infinite planes, one can treat the parallel velocity as an internal degree of freedom, similar to spin for point-like objects.

⁸The upper bound is $\pi/2$ not π because the perpendicular velocity is a projection of the velocity, so it cannot point in the opposite direction of \mathbf{v} .




Figure 2.2: Pictorial view of the coordinate system for the (red) Earth's local velocity \mathbf{v}_L , the (blue) total velocity \mathbf{v} , and the (green) perpendicular velocity \mathbf{v}_\perp (which is perpendicular to the green plane). The dashed black arrow is perpendicular to \mathbf{v}_L and is chosen based on galactic reference.

reference frame (notably that the Earth travels towards the constellation, Deneb). Let $\xi \in [0, \pi]$ be the angle between \mathbf{v}_\perp and \mathbf{v}_L and $\phi_\perp \in [0, 2\pi)$ be the azimuthal angle for \mathbf{v}_\perp with respect to some direction perpendicular to \mathbf{v}_L . Finally, define an azimuthal angle $\phi_2 \in [0, 2\pi)$ for \mathbf{v} about \mathbf{v}_\perp and with respect to the projection of \mathbf{v}_L in the domain-wall plane (or any arbitrary reference in the plane, if \mathbf{v}_L is perpendicular to the plane). Altogether, \mathbf{v}_\perp is described in polar coordinates $(v_\perp, \phi_\perp, \xi)$ with the pole \mathbf{v}_L , \mathbf{v} is described with the polar coordinates (v, ϕ_2, η) with the pole \mathbf{v}_\perp , and \mathbf{v}_L is given with respect to galactic coordinates.

The orientation of the domain walls will also follow some probability distribution. It is assumed that this follows a flat distribution,

$$dP_{\text{ori}} = \frac{1}{2\pi} d\bar{\xi} d\phi_\perp. \quad (2.23)$$

This assumption may be imperfect and there may be some structure to a domain-wall network that leads to some preferred orientation. Such considerations would require a more detailed understanding of structure formation and considerations of various unknown characteristics of the underlying physics.

The final PDF is obtained by taking the product of Eq. (2.22) and Eq. (2.23) and integrating over η and ϕ_2 , representing the unobservable total velocity. First, Eq. (2.22)

must be expressed in terms of the coordinates $(v_\perp, \bar{\eta}, \phi_2)$.

$$\begin{aligned} dP_0 &= \frac{v_\perp^2}{\bar{\eta}^3 (\pi v_c^2)^{3/2}} e^{-\frac{\frac{v_\perp^2}{\bar{\eta}^2} + v_L^2 + 2v_\perp v_L (\sin(\xi) \cos(\phi_2) \sqrt{1/\bar{\eta}^2 - 1} + \cos(\xi))}{v_c^2}} dv_\perp d\bar{\eta} d\phi_2 \\ &= \frac{s v_\perp^2}{(\pi v_c^2)^{3/2}} e^{-\frac{s^2 v_\perp^2 + v_L^2 + 2v_\perp v_L (\sin(\xi) \cos(\phi_2) \sqrt{s^2 - 1} + \cos(\xi))}{v_c^2}} dv_\perp ds d\phi_2, \end{aligned}$$

for $s = 1/\bar{\eta} \in [1, \text{inf})$. Taking the product of the above equation with Eq. (2.23) and integrating over the unobservable parameters yields

$$\begin{aligned} dP &= dv_\perp d\bar{\xi} d\phi_\perp \frac{v_\perp^2}{2\pi (\pi v_c^2)^{3/2}} e^{-\frac{v_L^2 + 2v_\perp v_L \cos(\xi)}{v_c^2}} \\ &\quad \times \int_1^\infty ds \int_0^{2\pi} d\phi_2 s e^{-\frac{v_\perp^2 s^2 + 2v_\perp v_L \sin(\xi) \cos(\phi_2) \sqrt{s^2 - 1}}{v_c^2}}. \end{aligned} \quad (2.24)$$

Thusfar, the PDF has been described in terms of the velocity distribution of objects throughout the galaxy (or some generic volume). However, this is not necessarily the same as the distribution of velocities that are observed. In particular, domain walls that travel faster with respect to the network are more likely to be observed; thinking in terms of the volume ‘‘swept’’ by a domain wall over a given period of time, faster walls cover a larger volume and that volume is thus more likely to include the network. The EPDF is then given by the product of Eq. (2.24) and $v_\perp / \langle v_\perp \rangle$,

$$\begin{aligned} dP_{\text{evt}} &= dv_\perp d\bar{\xi} d\phi_\perp \frac{v_\perp^3}{2\pi (\pi v_c^2)^{3/2} \langle v_\perp \rangle} e^{-\frac{v_L^2 + 2v_\perp v_L \cos(\xi)}{v_c^2}} \\ &\quad \times \int_1^\infty ds \int_0^{2\pi} d\phi_2 s e^{-\frac{v_\perp^2 s^2 + 2v_\perp v_L \sin(\xi) \cos(\phi_2) \sqrt{s^2 - 1}}{v_c^2}}, \end{aligned} \quad (2.25)$$

where

$$\langle v_\perp \rangle = \int v_\perp dP.$$

This EPDF, as estimated with a Monte Carlo simulation, is shown in Fig. 2.3.

Once again, one can consider a speed cutoff due to domain walls traveling above escape velocity evaporating out of the galaxy. In the case of domain walls, there is some ambiguity as to whether the total velocity or perpendicular velocity is more relevant. If the domains can easily stretch out, the perpendicular velocity determines whether the wall will escape, because one could have part of the domain boundary escaping the galaxy, while another part remaining within the galaxy. However, if there is some resistance to stretching a domain when part of the boundary is moving particularly fast, it would effectively pull the slower part of the boundary. This latter scenario is the case in which the area of a domain wall contains some energy (and hence, stretching causes




Figure 2.3: The EPDF for domain walls from a Monte Carlo simulation. In this simulation, $v_c = v_L = 220$ km/s and $v_e = 550$ km/s. With this, about 10 % of domain walls are omitted due to exceeding escape velocity.

an increase in energy), as is discussed in Chapter 3.1. Thus, to remove velocities above the escape velocity from the EPDF amounts to multiplying by a factor proportional to

$$\Theta(v_e - \|\mathbf{v} + \mathbf{v}_L\|)$$

prior to the integration in Eq. (2.25).

2.8 Magnetometer operation

The previous sections focus on theoretical aspects behind dark matter searches as well as broad overviews of some experiments and measurements. A more focused overview of experimental devices is considered in this section; specifically for magnetometers. This work focuses on analyzing data from GNOME, which consists of optical magnetometers. To some extent, the specific operation of a magnetometer can be treated as a black box. However, knowing the underlying coupling to dark matter is crucial to understanding observations. A review of GNOME magnetometers used in the first Science Run is found in Ref. [2]. A more general review of optical magnetometry can be found in Ref. [146].

There are a few different types of magnetometers used in the GNOME network. Not only do the magnetometers probe different atoms/transitions, there are multiple magnetometry techniques implemented across the network. The types of sensors being used are amplitude modulated nonlinear magneto-optical rotation (NMOR), radio-frequency (rf)-driven, and spin exchange relaxation-free (SERF) magnetometers. Including various different types of magnetometers enables one to test different couplings at the cost

of adding additional model dependence. That is, because different magnetometers may couple differently to dark matter, the analysis of magnetometer data is sensitive to the choice in the coupling model.

Though there are differences, all current magnetometers in the GNOME network operate on common principles. They all probe a transition in alkali atoms held in a vapor cell. For the non-SERF magnetometers, these atoms are placed in a leading magnetic field, while SERF magnetometers operate at near-zero field. The geometry of the device determines the sensitive axis of the magnetometer. Using a pump- and probe-lasers, generally with a lock-in amplifier setup, a particular transition is measured that is sensitive to the magnetic field (i.e., $\Delta m \neq 0$ when using the sensitive axis as the quantization axis). A voltage representing the transition energy is the output of the magnetometer, and the data are calibrated to reflect the true magnetic field values.

The magnetometers in GNOME are set up with a few other features. Perhaps most notably, the magnetometers are placed within magnetic shielding. Because this experiment aims to observe non-magnetic phenomena that mimic the type of spin-coupling due to a magnetic field, it is important to suppress the spin-coupling linked to real magnetic fields. The sensors are also equipped with a GPS-synchronized data acquisition system [147] that ensures that the data streams from different stations can be synchronized. Further, various auxiliary sensors provide information on the reliability of each sensor. With this, one can better avoid questionable data in the analysis.

A notable difference in the operation of magnetometers is the use of SERF magnetometry [146, 148, 149]. A common problem in optical magnetometry is spin-exchange in which colliding atoms transfer angular momentum to each other; conserving total angular momentum. This reduces coherence and magnetometer sensitivity. Despite the total angular momentum being conserved, these collisions cause the atomic hyperfine states to change. In conventional optical magnetometers, this effect is reduced by operating with low atomic density at the cost of worse sensitivity. However, a way to improve overall sensitivity is to operate in a regime with high density, fast spin-exchange, and a weak magnetic field — the SERF regime. This leads to a non-vanishing average spin precession across the vapor cell that is slower than that of an individual atom. A pump beam can be used to orient the precession. Under appropriate conditions, a SERF magnetometer can have much better sensitivity than more traditional magnetometers due to the comparatively high density of alkaline atoms being probed.

Chapter 3

Signal shape

Understanding the types of signatures that dark matter can leave in an experiment is necessary for any dark matter search. Establishing the phenomenology of a theory enables one to understand the information gathered in an experiment. In this chapter, the phenomenology of ALP domain walls and axion coupling to matter is established. This allows one to understand a class of signal shapes that one may expect to see in an experiment.

First, the shape and characteristics of an ALP domain wall is described in Sec. 3.1. The shape of the domain wall is established in terms of the underlying degrees of freedom for the ALP. Then, using generic coupling terms, the signal induced by a domain wall is described in Sec. 3.2 for coupling to fermion spin and in Sec. 3.3 for Yukawa-like coupling to fermion mass. The spin-coupling is of particular interest to GNOME as it will result in a pseudo-magnetic field that can be observed by magnetometers.

3.1 Domain wall field

A common feature of ALPs is that they arise from considering a complex scalar field ϕ with a $\mathbb{Z}_{N_{\text{DW}}}$ symmetry, $\phi \rightarrow \exp(2\pi i k/N_{\text{DW}})\phi$ (for $k \in \mathbb{Z}$). The ALP field can have any potential $V(\phi)$ that respects the $\mathbb{Z}_{N_{\text{DW}}}$ symmetry. For example, the Lagrangian¹ may be

$$\mathcal{L} \supset |\partial_\mu \phi|^2 - \frac{\lambda}{\bar{f}_a^{2N_{\text{DW}}-4}} |2^{N_{\text{DW}}/2} \phi^{N_{\text{DW}}} + \bar{f}_a^{N_{\text{DW}}}|^2, \quad (3.1)$$

with slight variations by convention². The field can be parameterized by a real Higgs mode S and axion mode a as $\phi = \frac{S}{\sqrt{2}} \exp(ia/\bar{f}_a)$. Freezing the Higgs mode $S \rightarrow \bar{f}_a$ yields the axion Lagrangian,

$$\mathcal{L}_a = \frac{1}{2} (\partial_\mu a)^2 - 4\lambda \bar{f}_a^4 \cos^2 \left(\frac{a}{\bar{f}_a/N_{\text{DW}}} \right). \quad (3.2)$$

¹Note that the metric signature is $(+, -, -, -)$.

²Other references may use a minus sign in front of the $\bar{f}_a^{N_{\text{DW}}}$ term, which results in a similar potential, up to a phase. The end result of Eq. (3.1) is that the axion potential will have a maximum at zero, while the minus-convention will have a vacuum at zero.

Matching the second-order Taylor expansion around the vacuum $a = \frac{\pi \bar{f}_a}{N_{\text{DW}}}(1 + 2k)$ to the scalar mass term $\frac{m_a^2}{2}a^2$ yields

$$m_a = N_{\text{DW}} \bar{f}_a \sqrt{2\lambda}. \quad (3.3)$$

Written in terms of m_a and $f_a := \bar{f}_a/N_{\text{DW}}$, the axion Lagrangian is

$$\mathcal{L}_a = \frac{1}{2}(\partial_\mu a)^2 - 2m_a^2 f_a^2 \cos^2\left(\frac{a}{2f_a}\right). \quad (3.4)$$

The classical equation of motion, as determined by applying the principle of least action to Eq. (3.4) (assuming the boundary terms vanish), is

$$-\partial^2 a + m_a^2 f_a \sin\left(\frac{a}{f_a}\right) = 0. \quad (3.5)$$

Generically, domain walls form when a field can monotonically increase between vacuum states (possibly the same state). For simplicity, assume that the axion field is static in time and only varies in the x -direction, so $a = a(x)$ with a planar domain wall parallel to the yz -plane. The equation of motion, Eq. (3.5), becomes

$$\frac{a''}{f_a} + m_a^2 \sin\left(\frac{a}{f_a}\right) = 0,$$

because $-\partial^2 a = \frac{\partial^2 a}{\partial x^2} = a''$ for the chosen metric signature. Assume that $a(x \rightarrow \pm\infty)$ becomes constant. Then $a(x \rightarrow \pm\infty) = k_\pm \pi f_a$ from the equations of motion ($k_\pm \in \mathbb{Z}$). Note that $a = k\pi f_a$ are the extrema of the potential; minima for k odd and maxima for k even. With this in mind, define the boundary conditions $a(x \rightarrow \pm\infty) = \pm\pi f_a$ and $a(0) = 0$. That is, the axion wall intersects the origin and separates the vacua closest to $a = 0$.

As derived in Appendix A.1, the classical domain-wall solution is given by

$$a(x) = 2f_a \arcsin(\tanh(m_a x)). \quad (3.6)$$

This is shown in Fig. 3.1.

It is also useful to consider the gradient of the field. As is discussed in Sec. 3.2, one can consider couplings between the axion field gradient and fermionic spins. The gradient is

$$a'(x) = \frac{2m_a f_a}{\cosh(m_a x)} \approx \frac{2m_a f_a}{1 + \left(\frac{x}{\frac{1}{2}\frac{2\sqrt{2}}{m_a}}\right)^2} + \mathcal{O}(m_a^4 x^4), \quad (3.7)$$

that is, it can be approximated by a Lorentzian. The full width at half maximum is

$$\Delta x = \frac{2 \cosh^{-1}(2)}{m_a} \approx \frac{2\sqrt{2}}{m_a}. \quad (3.8)$$




Figure 3.1: The classical axion field from a domain wall.

Finally, consider the energy density in the domain wall. The Hamiltonian density is

$$\begin{aligned}
 \mathcal{H}_a &= \frac{1}{2} (a')^2 + 2m_a^2 f_a^2 \cos^2\left(\frac{a}{2f_a}\right) \\
 &= \frac{2m_a^2 f_a^2}{\cosh^2(m_a x)} + 2m_a^2 f_a^2 \cos^2(\arcsin(\tanh(m_a x))) \\
 &= \frac{4m_a^2 f_a^2}{\cosh^2(m_a x)}.
 \end{aligned}$$

Integrating over the x yields the surface tension,

$$\sigma_{\text{DW}} = 8m_a f_a^2. \quad (3.9)$$

3.2 Spin coupling

The GNOME project is particularly interested in spin-coupling. The magnetometers used in GNOME rely on couplings involving atomic spins. However, the sensors are shielded from magnetic fields, so observed couplings are likely the result of an exotic field penetrating the shielding and coupling to atomic spins.

Generically, the spin coupling terms are of the form,

$$\mathcal{L}_{a\text{-spin}} = f_\mu(a) \bar{\psi} \gamma^\mu \gamma^5 \psi, \quad (3.10)$$

where $f_\mu(a)$ is an indexed function of the axion field and $\bar{\psi} \gamma^\mu \gamma^5 \psi$ is the axial-vector current for the fermion field ψ . One could further complicate this term by multiplying additional terms involving the fermion or other fields, but these couplings are generally suppressed.

3.2.1 Linear

Perhaps the simplest coupling between the axion field is the linear coupling of spins to the gradient $\partial_\mu a$. This is accomplished via the following coupling to the complex scalar field ϕ ,

$$\begin{aligned} \mathcal{L}_{\text{int}} &= \frac{i\phi \overleftrightarrow{\partial}_\mu \phi^*}{f_a f_{\text{int}}} \bar{\psi} \gamma^\mu \gamma^5 \psi \\ &\quad \xrightarrow{S \rightarrow \bar{f}_a} \frac{\partial_\mu a}{f_{\text{int}}} \bar{\psi} \gamma^\mu \gamma^5 \psi, \end{aligned} \quad (3.11)$$

where f_{int} is the coupling scale.

The axial-vector current is related to the spin \mathbf{S} , so that the interaction Hamiltonian becomes

$$H_{\text{int}} = \frac{1}{f_{\text{int}}} \nabla a \cdot \frac{\mathbf{S}}{\|\mathbf{S}\|}, \quad (3.12)$$

i.e., for spin- $1/2$ particles, $1/\|\mathbf{S}\| = 2$.

Magnetometers used for GNOME measure the change between two atomic energy levels. Denote the electron spin \mathbf{S}_e , nuclear spin \mathbf{I} , orbital angular momentum \mathbf{L} , and total angular momentum $\mathbf{F} = \mathbf{S}_e + \mathbf{I} + \mathbf{L}$. The magnetometers measure the magnetic field along some sensitive axis (e.g., the direction of a biasing field). Denote the projection of the spin along this axis as m_k for k representing the relevant spin (e.g., m_F for \mathbf{F}).

Plugging in the maximum gradient from Eq. (3.7) into Eq. (3.12) yields the change in energy between atomic energy levels

$$\Delta E = \sum_{i \in \{e, p, n, \dots\}} \frac{2\eta \sigma_{(i)} \Delta m_F}{\|\mathbf{S}_{(i)}\|} \frac{f_a}{f_{\text{int}}^{(i)}} m_a, \quad (3.13)$$

where i labels the species of fermion, $\sigma_{(i)} = \frac{\langle \mathbf{S}_{(i)} \cdot \mathbf{F}_{(i)} \rangle}{F_{(i)}^2}$ is the projected spin coupling, $\eta = \cos \theta$ for the angle between the axion gradient and sensitive axis θ , and $f_{\text{int}}^{(i)}$ is the interaction coupling to particle i . In general, we will add up $\frac{\sigma_{(i)}}{\|\mathbf{S}_{(i)}\| f_{\text{int}}^{(i)}}$ into an effective ratio $\frac{2\sigma_j}{f_{\text{int}}}$, where j labels the magnetometer.

The energy shift due to the Zeeman effect is

$$\Delta E_B = g_F \mu_B \Delta m_F B, \quad (3.14)$$

for g -factor, g_F , and Bohr magneton, μ_B . The GNOME magnetometers are shielded from magnetic field, so the observed “magnetic field” observed is fictitious. For this reason, B_j refers to the pseudo-magnetic field in the j^{th} magnetometer. Equating Eq. (3.13) and Eq. (3.14) yields

$$\frac{g_{F,j} B_j}{\sigma_j \eta_j} = \frac{4}{\mu_B} m_a \xi =: \mathcal{B}_p, \quad (3.15)$$

where $\xi := \frac{f_a}{f_{\text{int}}}$ and the j subscript refers to the j^{th} magnetometer. Here, \mathcal{B}_p is the normalized pseudo-magnetic field; i.e., normalized to be independent of the magnetometer

type and relative orientation. One can understand ξ as the (unitless) coupling strength relative to f_a . Values of $\sigma_j/g_{F,j}$ based on proton-coupling for GNOME magnetometers is given in Table 6.1.

3.2.2 Generalizations

Though much of the work described here will focus on the linear pseudoscalar ALP coupling with spins described above, one can further generalize the coupling. Consider spin coupling terms similar to Eq. (3.10) where $f_\mu(a) = \partial_\mu g(a)$.

For ALP fields, there is an unbreakable symmetry $a \rightarrow a + 2\pi\bar{f}_a$ imposed by the topology of the field; recall that an ALP field is originally defined as the phase of a complex field. This periodic condition must also occur in the function $f_\mu(a)$; note that this condition is only true up to a constant in $g(a)$, because a constant vanishes upon differentiation. After $U(1)_{\text{PQ}}$ symmetry-breaking, the ALP field retains the discrete symmetry $a \rightarrow a + 2\pi f_a = a + 2\pi\bar{f}_a/N_{\text{DW}}$. Strictly speaking, there is no reason to demand that $f_\mu(a)$ has this periodic condition. An interaction term without this $\mathbb{Z}_{N_{\text{DW}}}$ symmetry will explicitly break it, which may cause the presence of fermions to drive the motion of domain walls as one domain could become more energetically preferable.

Considering the discrete symmetries needed in the coupling term, a Fourier expansion of the function $g(a)$ is a natural choice to enforce this symmetry. The minimal symmetry requirement demands that $\partial_\mu g(a) = \partial_\mu g(a + 2\pi\bar{f}_a)$. Requiring that the $\mathbb{Z}_{N_{\text{DW}}}$ symmetry remain unbroken can be achieved by omitting terms in the Fourier series. Defining the Fourier coefficients $\{\alpha_i\}$ and $\{\beta_i\}$, the Fourier expansion is given by

$$g(a) = \alpha_0 a + \sum_{n=1}^{\infty} \left[\frac{\alpha_n \bar{f}_a}{n} \sin\left(\frac{n}{\bar{f}_a} a\right) - \frac{\beta_n \bar{f}_a}{n} \cos\left(\frac{n}{\bar{f}_a} a\right) \right]. \quad (3.16)$$

The $\frac{\bar{f}_a}{n}$ factors acting on the Fourier coefficients are chosen to simplify the equations in a later step. The $\mathbb{Z}_{N_{\text{DW}}}$ symmetry can be preserved if the only non-zero terms in the Fourier series are multiples of N_{DW} .

The observed signal pattern for a domain wall can be understood by the energy shift ΔE induced by the domain wall field, Eq. (3.6). Inserting the field into Eq. (3.10) and expanding, one obtains

$$\Delta E = 2m_a f_a (\bar{\psi} \gamma^3 \gamma^5 \psi) \operatorname{sech}(m_a z) \left\{ \alpha_0 + \sum_{n=1}^{\infty} [\alpha_n \cos(2n \arcsin(\tanh(m_a z))) + \beta_n \sin(2n \arcsin(\tanh(m_a z)))] \right\}. \quad (3.17)$$

This defines a series of signal shapes defined by each Fourier component. For example, the α_0 term defines the shape of the signal due to linear coupling. Example signals are given in Fig. 3.2.




Figure 3.2: Sample signal shapes for different components described in Eq. (3.16) and Eq. (3.17). The dashed gray lines are sech functions (i.e., the α_0 term).

In addition to the class of signal shapes described above, one can consider other signal shapes for non-ALP scalar fields. If the field is not periodic, one can instead use a Taylor series to expand the function $g(a)$. Further, one can consider terms with higher-orders in the fermion field (or other Standard Model particles).

If only very small perturbations in the ALP field are considered, one can still use a Taylor expansion to approximate the coupling. However, the field across a domain wall inherently changes significantly. Second-order and higher terms in a (finite) Taylor expansion would significantly break the $\mathbb{Z}_{N_{\text{DW}}}$ symmetry over the range of two vacua and possibly possibly the unbreakable $a \rightarrow a + 2\pi\bar{f}_a$ symmetry unless $N_{\text{DW}} \gg 1$ (so the field across the domain wall does not vary significantly relative to \bar{f}_a). Other experiments could still attempt to measure small oscillations in the ALP field within a domain. For these experiments, it could be of interest to include a quadratic coupling as an approximation for higher-order terms. For example, there is a proposal to use GNOME to search for Q -balls via their linear and quadratic spin coupling [6].

3.3 Mass coupling

In addition to the spin-coupling described in the previous section, one could also consider a coupling between the ALP field and a mass-like Standard Model term. The form of this coupling is

$$\mathcal{L}_{a-\psi\psi} = g(a)\bar{\psi}\psi, \quad (3.18)$$

where $g(a)$ is some function of the ALP-field.

This work does not consider the mass-like coupling described by Eq. (3.18), though this is a common type of coupling to consider. The shape of an observed signal from a domain wall, as before, is described by simply plugging Eq. (3.6) into Eq. (3.18). For example, a linear coupling would induce a signal similar to Fig. 3.1 with the position axis replaced by time. Notably, one would observe an effective variation in mass for fermions as the ALP field varies over time. In a similar coupling to photons with $\bar{\psi}\psi \rightarrow F_{\mu\nu}^2$, the variation in the ALP field would appear as a variation in the fine-structure constant.

Various experimental efforts and proposals have aimed to search for mass-like coupling of the axion to fermions. The apparent variation in mass or fine-structure constant would be observable in atomic clocks [7, 150], cavities [7], spectroscopy [151, 152], and gravimeters [153, 154]. These experiments among others can provide bounds for mass-like coupling.

Because the space occupied by Earth possesses a high-density of mass relative to the galaxy, there may be significant back action of the Earth on the ALP field due to a mass-coupling term. This idea was explored in detail in Ref. [155] for the case where $g(a)$ is a quadratic polynomial. The accumulation of mass can result in a significant effective potential term for the ALP field. This could result in one domain being favored in the presence of matter. In a more extreme case, focusing on a pair of local minima in the ALP potential as a (local) double-well, the presence of the effective potential from background matter could combine the double-well into a single-well between the double-well. In this case, there would be a “half” domain wall surrounding the Earth as the effective potential transitions from the single-well case to the double-well case away from dense regions of matter. In contrast, because fermion spins do not substantially align across the Earth, the back action from the spin-coupling term will not play such a large role.

Chapter 4

Analysis

Data from the GNOME network are studied for signs of ALP domain walls crossing the network. For this, methods were developed to search through the data to detect for signatures of such objects. This process involves applying filters to the data as well as scanning the data for domain-wall signals. In scanning the data, various thresholds must also be selected to ensure that the results are statistically rigorous. These analysis methods are also described in Ref. [3].

The effects of applying filters to the data are described in Sec. 4.1. Methods for testing the consistency between the measurements and a domain wall signal are described in Sec. 4.2, and a geometric picture of this process is presented in Sec. 4.3. Part of the analysis involves testing for domain walls with different velocities. This requires one to develop a set of velocities to test that adequately covers possible domain-wall velocities. Requirements, design, and optimization of this set is described in Sec. 4.4. An overview of the analysis procedure as it is applied to real data is described in Sec. 4.5. Finally, methods for determining the various thresholds used in the analysis are described in Sec. 4.6.

4.1 Filtering data

Filtering is often used to reduce noise in order to better observe a signal. This noise can be random/stochastic (e.g., Gaussian noise in which each data point has some random offset from the true value) or systematic (e.g., a 50 or 60 Hz oscillation due to powerline interference with electronics). Generically, a filter can be anything that removes some information with the aim to leave only useful information. An example of filtering is given in Fig. 4.1 with simulated data to illustrate its benefits.

GNOME aims to search for signals in data streams over time, and the relevant filters are linear impulse filters. These filters can be described by an impulse response $h(t)$ (or $h[i]$ for discrete data) with which the data are convolved or, equivalently a transfer function $H(\omega)$ (or $H[k]$ for discrete data) which describes how the data are attenuated




Figure 4.1: (a) Simulated data with a Lorentzian signal, Gaussian-distributed noise, and a 10 Hz background noise. (b) A notch frequency filter is applied to remove the frequency component around 10 Hz followed by (c) averaging/down-sampling by 0.2 s.

in frequency space. In particular, filtering the continuous (discrete) data $y(t)$ ($y[i]$),

$$y_f(t) = \{y * h\}(t) := \int ds y(s)h(t-s), \quad (4.1)$$

$$y_f[t] = \{y * h\}[t] := \sum_s y[s]h[t-s], \quad (4.2)$$

where the integral (sum) is over the relevant domain and the discrete indices are cyclic. Without loss of generality, the domain for continuous data will be \mathbb{R} and for discrete data will be $0, 1, \dots, N-1$ so cyclicity implies $t \sim t + N$. It should also be noted that the convolution operator $*$ is commutative,

$$y * h = h * y,$$

and associative

$$h_2 * (h_1 * y) = (h_2 * h_1) * y.$$

An example of an impulse filter is a rolling average. In this case, h would be a square pulse whose length is the averaging time T and whose height is the normalization $1/T$. Plugging h into Eq. (4.1) or (4.2), the filtered signal is described by the integral/sum of the signal around the respective point, up to a normalization constant; i.e., an average.

The transfer function arises when considering the convolution theorem which connects the functional pointwise product \cdot of functions with the convolution operator $*$ via Fourier transforms. To be explicit, the continuous and discrete Fourier transform \tilde{f} of f will be defined here, respectively, as

$$\{\mathcal{F}_C f\}(\omega) := \int_{-\infty}^{\infty} dt f(t)e^{-i\omega t} \quad \text{and} \quad \{\mathcal{F}_C^{-1} \tilde{f}\}(t) := \frac{1}{2\pi} \int_{-\infty}^{\infty} d\omega \tilde{f}(\omega)e^{+i\omega t}, \quad (4.3)$$

$$\{\mathcal{F}_D f\}[k] := \sum_{n=0}^{N-1} f[n]e^{-\frac{2\pi i}{N}nk} \quad \text{and} \quad \{\mathcal{F}_D^{-1} \tilde{f}\}[n] := \frac{1}{N} \sum_{k=0}^N \tilde{f}[k]e^{+\frac{2\pi i}{N}nk}. \quad (4.4)$$

For the continuous case, the Fourier transform yields a function in terms of the angular frequency, while for the discrete case, if the sampling rate is r_s , the frequency is given in units of $r_s/2N$. Connecting the two, $\omega \sim \frac{2\pi r_s}{N}k$. The convolution theorem for functions f and h is then written¹

$$\mathcal{F}(f * h) = \mathcal{F}f \cdot \mathcal{F}h, \quad (4.5)$$

where \mathcal{F} can be either discrete or continuous.

With the convolution theorem in mind, one can better understand the practical use of defining a quantity related to a Fourier transform of h . Because impulse filters affect signals via convolution and moving to Fourier space replaces convolutions with a simple pointwise product, this transformation provides a useful computational tool. The transfer function H is given by²

$$H(i\omega) := \{\mathcal{F}_C h\}(\omega) \quad \text{and} \quad H\left[e^{\frac{2\pi i}{N}k}\right] := \{\mathcal{F}_D h\}[\omega]. \quad (4.6)$$

From Eq. (4.5), one can interpret the transfer function as defining the attenuation factors in frequency space. One can also define a frequency-response function

$$\tilde{h} := \mathcal{F}h, \quad (4.7)$$

so $\tilde{h}(\omega) = H(i\omega)$ and $\tilde{h}[k] = H\left[e^{\frac{2\pi i}{N}k}\right]$. Setting H or \tilde{h} to zero or one at certain frequencies will, respectively, exclude or include those frequencies in the filtered data. In general, the transfer function can be complex, though requiring the input and output data from a filter be real is achieved by constraining

$$H^*(+i\omega) = H(-i\omega) \quad \text{or} \quad H^*\left[e^{+\frac{2\pi i}{N}k}\right] = H\left[e^{-\frac{2\pi i}{N}k}\right], \quad (4.8)$$

where H^* is the complex conjugate of H and the discrete indices are cyclic (i.e., $k \sim k + N$). Heuristically, a complex phase on the transfer function will cause a time-shift in the filtered data; though perhaps different shifts for different frequency components. As an example, consider the impulse function $h_{t_0}(t) = \delta(t - t_0)$ that translates the data by t_0 . The corresponding transfer function is $H_{t_0}(i\omega) = e^{-i\omega t_0}$, which is completely defined by its complex phase.

A useful feature of the filters described here is their linearity. Notably, the convolution operator is bilinear and the Fourier transform is linear,

$$\begin{aligned} f * (ag + h) &= a(f * g) + (f * h), \\ \mathcal{F}(ag + h) &= a(\mathcal{F}g) + (\mathcal{F}h), \\ \mathcal{F}^{-1}(ag + h) &= a(\mathcal{F}^{-1}g) + (\mathcal{F}^{-1}h), \end{aligned}$$

¹Different conventions for the Fourier transform may include a factor $\sim 2\pi$.

²Normally, the transfer function is written with the Laplace transform (continuous) and Z -transform (discrete) in mind. In particular, $\{\mathcal{L}h\}(s) = \int dt h(t)e^{-st} =: H(s)$ and $\{\mathcal{Z}h\}[z] = \sum_t h[t]z^{-t} =: H[z]$. The Fourier transform is a special-case of these transforms.

where f , g , and h are functions, and a is a constant (acting on a function as the typical pointwise product). *Bilinearity* in the convolution follows trivially from the fact that it is a commutative operator. The fact that these operators are linear means that a measurement can be split into signal y_s and noise y_n components, $y = y_s + y_n$, and the filter effects on each component can be understood separately.

One may apply multiple filters to data. Consider applying the two filters h_1 and h_2 to y ,

$$y_f = h_2 * h_1 * y,$$

noting that $*$ is associative, so the order of these operations does not matter. For the transfer functions $H_1 \sim \mathcal{F}h_1$ and $H_2 \sim \mathcal{F}h_2$, the transfer function for the combined filter is simply the product

$$H = H_1 \cdot H_2. \quad (4.9)$$

This can be seen by applying the convolution theorem, Eq. (4.5), twice.

Before describing exactly how filters will effect measurements, it should be noted that “filtering” as described here can occur at various stages in the experiment. Often, one considers digitally applying filters to recorded data or to a live data stream (though the latter can result in unwanted time-shifting). However, the experiment itself may have some analog filters either by design or experimental limitations. Bandwidth measurements are often used to characterize this “filter” as they describe how an experiment responds to signals at different frequencies.

4.1.1 Effects on signal

The effects on signals will be considered in an approximation in which data collection is continuous. From a practical standpoint, this removes complications of considering the position of the signal within a sampling period. This allows one to ignore subtle effects that arise from discretizing the data. For example, when considering a peak-shaped signal, one would need to consider whether the peak is centered on or off of a discrete point. When returning to the more realistic, discrete data, one can sample points from the continuous function periodically.

With this approximation in mind, the effects on signals becomes straight-forward. The filtered signal shape can be obtained either with the impulse response $h(t)$ and Eq. (4.1) or the frequency response $\tilde{h}(\omega)$ with the aid of the convolution theorem, Eq. (4.5). The linearity of the filter means that one can understand the shape of a signal with different amplitudes; i.e., if, for constant A , $y_0(t) = Ay_1(t)$ then $y_{f0}(t) = Ay_{f1}(t)$. Specific examples of filter effects on signals are given in Sec. 4.1.3.

4.1.2 Effects on noise

When calculating the effects of filtering on noise, discrete data are considered. Moreover, only stochastic noise will be considered; essentially, it will be assumed that systematic noise is entirely removed by the filters and calibration. Additionally, it is assumed that

the noise is described by a normal distribution. What is needed, then, is to calculate the propagation-of-errors (i.e., noise) across the filter.

Because filtering is a linear operation in this case, propagation-of-errors is a fairly simple operation. For full generality, let the noise be fully described by the covariance matrix Σ , though this is often diagonal as data points at different times are assumed to be uncorrelated. Specifically, for data $y[t]$, $\Sigma_{ij} = \Sigma[i, j]$ represents the covariance between $y[i]$ and $y[j]$ while Σ_{ii} is the variance in point $y[i]$. The covariance is often estimated by calculating the standard error of $y[t]$,

$$\sigma^2 = \frac{1}{N} \sum_{t=0}^{N-1} (y[t] - \bar{y})^2 \quad \text{for} \quad \bar{y} = \frac{1}{N} \sum_{t=0}^{N-1} y[t],$$

so $\Sigma_{ij} = \delta_{ij}\sigma^2$ in this estimation³.

Propagation of errors follows the typical methods. The act of filtering (Eq. (4.2)) is equivalent to a linear operation with the (real and circulant) Jacobian matrix $J[t, s] = h[t - s]$,

$$\mathbf{y}_f = J\mathbf{y},$$

treating the data stream $y[t]$ as an N -vector \mathbf{y} . If the transfer function $H(i\omega)$ is real, then $J[t, s]$ will be symmetric. This follows from the idea that a real transfer function does not shift the data, because a symmetric J (or where $h[t] = h[-t]$) means that each filtered point has equal contribution from a time before and the same time after that point. The covariance in the filtered data is $J\Sigma J^T$ or

$$\begin{aligned} \Sigma_f[t, s] &= \sum_{m,n=0}^{N-1} J[t, m]\Sigma[m, n]J[s, n] \\ &= \sum_{m,n=0}^{N-1} h[t - m]h[s - n]\Sigma[m, n]. \end{aligned} \quad (4.10)$$

If the noise is constant and uncorrelated, $\Sigma_{ij} = \delta_{ij}\sigma^2$, this relationship simplifies to

$$\begin{aligned} \Sigma_f[t, s] &= \frac{\sigma^2}{N} \sum_{k=0}^{N-1} \left| \tilde{h}[k] \right|^2 e^{\frac{2\pi i}{N}k(t-s)} \\ &= \sigma^2 \left\{ \mathcal{F}_D^{-1} \left| \tilde{h} \right|^2 \right\} [t - s]. \end{aligned} \quad (4.11)$$

The above equation is obtained by first moving to Fourier space $h = \mathcal{F}^{-1}\tilde{h}$. The filtered data's covariance matrix Σ_f is symmetric, real, and circulant for constant, uncorrelated noise. If the filter roughly leaves $y_f[t] \sim y[t]$ and the initial data are uncorrelated, the filtered data will be roughly uncorrelated.

³A better approximation for Σ can be made by having M uncorrelated, repeated measurements of $y[t]$, $\{y_i[t]\}$, and calculating the covariance as $\Sigma[t, s] = \frac{1}{M} \sum_i (y_i[t] - \bar{y}[t]) (y_i[s] - \bar{y}[s])$ for $\bar{y}[t] = \frac{1}{M} \sum_i y_i[t]$, though this is often either difficult or impossible.

One limitation that remains for the discrete data is that it relies on the size of the data, N . Ideally, N would be sufficiently large such that considering a larger data-set would not substantially alter the results. In particular, if the filter has significant features at low-frequency, one could miss these effects when scaling down to filter a short data stream; e.g., a high-pass filter to remove background drifts on scales larger than the data being filtered. If the noise that is fairly stable and uncorrelated and $h[t]$ is described by a sharp peak, then Eq. (4.11) suggests that the filtered covariance differs by the unfiltered covariance by a constant factor. Choosing an N sufficiently large that all the features of the desired filter are obtained (and the appropriate $h[t]$), will yield a good approximation of this factor. This is particularly useful when considering a sensor that may have many on- and off-times but fairly stable noise during its on-time. Having a single factor would allow one to predict how a filter attenuates (or increases) noise regardless of the size of the data segment.

One can also consider the effect of a filter on noise in the continuum limit. In this case, the covariance matrix can be thought of as a function $\Sigma(t, s)$. Of course, one should be careful about what it means for continuous data to have random Gaussian noise, though these ideas are explored in areas like stochastic calculus. Normally, a Jacobian matrix describes the derivatives of a vector function $J = \partial \mathbf{y}_f / \partial \mathbf{y}$ whose matrix multiplication is performed via summation, while in the continuum limit, this matrix describes the functional derivative $J = \delta \mathbf{y}_f / \delta \mathbf{y}$ whose matrix multiplication is performed via integration. Uncorrelated data would have a covariance matrix of the form $\Sigma(t, s) = \delta(t - s)\sigma^2(t)$. It follows that Eq. (4.10) and Eq. (4.11) become, respectively

$$\begin{aligned} \Sigma_f(t, s) &= \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy J(t, x) \Sigma(x, y) J(s, y) \\ &= \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy h(t - x) h(s - y) \Sigma(x, y), \end{aligned} \quad (4.12)$$

and (if $\Sigma(t, s) = \sigma^2 \delta(t - s)$)

$$\begin{aligned} \Sigma_f(t, s) &= \frac{\sigma^2}{2\pi} \int_{-\infty}^{\infty} d\omega \left| \tilde{h}(\omega) \right|^2 e^{i\omega(t-s)} \\ &= \sigma^2 \left\{ \mathcal{F}_C^{-1} \left| \tilde{h} \right|^2 \right\} (t - s). \end{aligned} \quad (4.13)$$

Interestingly, the continuum limit no longer includes N . Thinking in terms of frequencies, the continuum would include frequencies larger than the sampling rate and oscillations longer than the total sampling time. It follows that so long as the filter does not have many high-frequency features that would be lost on the discrete data, the factor in Eq. (4.13) provides a good description of how the filter attenuates the noise.

4.1.3 Examples of signals and filters

The methods for determining the effects of filtering on the data are described in the previous two sections. It is useful to explore some specific examples of signals, filters, and how filters affect the data.

It is of particular interest here to explore peak-shaped signals, though one can easily explore different signal shapes that may arise from some theoretical predictions. To give some specific examples of signal shapes here, one can respectively define a Lorentzian, Gaussian, sech, and square pulse as

$$y_L(t; \Gamma, A, t_0) := \frac{A \left(\frac{1}{2}\Gamma\right)^2}{\left(\frac{1}{2}\Gamma\right)^2 + (t - t_0)^2}, \quad (4.14a)$$

$$y_G(t; \sigma, A, t_0) := \frac{A}{\sqrt{2\pi\sigma^2}} e^{-\frac{(t-t_0)^2}{2\sigma^2}}, \quad (4.14b)$$

$$y_{\text{sh}}(t; \Gamma, A, t_0) := A \operatorname{sech}\left(\frac{t - t_0}{\Gamma}\right), \quad (4.14c)$$

$$y_{\text{sq}}(t; \Gamma, A, t_0) := \begin{cases} A & |t - t_0| \leq \Gamma/2 \\ 0 & |t - t_0| > \Gamma/2 \end{cases}, \quad (4.14d)$$

where the variables to the right of the semicolon characterize the shape/width, amplitude, and position of a signal. Writing $\mathcal{F}y = \tilde{y}$, the Fourier transforms of these signal shapes are

$$\tilde{y}_L(\omega; \Gamma, A, t_0) = A\Gamma \frac{\pi}{2} e^{-\Gamma|\omega|/2 - i\omega t_0}, \quad (4.15a)$$

$$\tilde{y}_G(\omega; \sigma, A, t_0) = A e^{-\sigma^2\omega^2/2 - i\omega t_0}, \quad (4.15b)$$

$$\tilde{y}_{\text{sh}}(\omega; \Gamma, A, t_0) = A\Gamma\pi \operatorname{sech}\left(\frac{\pi}{2}\omega\Gamma\right) e^{-i\omega t_0}, \quad (4.15c)$$

$$\tilde{y}_{\text{sq}}(\omega; \Gamma, A, t_0) = A\Gamma \operatorname{sinc}\left(\frac{\omega\Gamma}{2}\right) e^{-i\omega t_0}. \quad (4.15d)$$

In addition to signals, it is useful to consider various filters explicitly. As with the signals, there are many different types of filters that can be applied, so several examples are given here.

Consider a rolling average for discrete data⁴ with N points and continuous data,

$$h_{\text{avg}}[t; T, t_0] := \begin{cases} \frac{1}{T} & 0 \leq t - t_0 \leq T - 1 \\ 0 & \text{else} \end{cases}, \quad (4.16a)$$

$$\tilde{h}_{\text{avg}}[t; T, t_0] = \frac{\operatorname{sinc}\left(2\pi\frac{k}{2N/T}\right)}{\operatorname{sinc}\left(2\pi\frac{k}{2N}\right)} e^{-\frac{2\pi i}{N}\left(t_0 + \frac{T-1}{2}\right)k}, \quad (4.16b)$$

$$h_{\text{avg}}(t; T, t_0) := \begin{cases} \frac{1}{T} & |t - t_0| \leq T/2 \\ 0 & |t - t_0| > T/2 \end{cases}, \quad (4.16c)$$

$$\tilde{h}_{\text{avg}}(\omega; T, t_0) = \operatorname{sinc}\left(\frac{\omega T}{2}\right) e^{-i\omega t_0}, \quad (4.16d)$$

⁴Expressions and inequalities with discrete indices are “cyclic” so that they refer to the equivalent value in $[0, N - 1]$. For example, if $n \in [0, N - 1]$, then the inequality $n + N < A$ should be understood as $n < A$.

where discrete variables are all integers in $[0, N - 1]$ and continuous variables are real. Note that the center of the rolling average given here is different for the discrete versus continuous case.

Likewise, a simple band-pass filter that accepts (positive and negative) frequency between two values

$$h_{\text{bp}}[t; k_0, k_1] = \begin{cases} \frac{\sin\left(2\pi\frac{k_1-1/2}{N}t\right) - \sin\left(2\pi\frac{k_0-1/2}{N}t\right)}{N \sin\left(2\pi\frac{t}{2N}\right)} & t \neq 0 \\ 2(k_1 - k_0)/N & t = 0 \\ -\frac{\delta_{k_0 0}}{N} - \begin{cases} (-1)^t/N & N \text{ even, } k_1 = \frac{N}{2} + 1 \\ 0 & \text{else} \end{cases} & \end{cases}, \quad (4.17a)$$

$$\tilde{h}_{\text{bp}}[k; k_0, k_1] := \begin{cases} 1 & k_0 \leq \pm k < k_1 \\ 0 & \text{else} \end{cases}, \quad (4.17b)$$

$$h_{\text{bp}}(\omega; \omega_0, \omega_1) = \frac{1}{\pi} [\omega_1 \text{sinc}(\omega_1 t) - \omega_0 \text{sinc}(\omega_0 t)], \quad (4.17c)$$

$$\tilde{h}_{\text{bp}}(\omega; \omega_0, \omega_1) := \begin{cases} 1 & \omega_0 \leq |\omega| < \omega_1 \\ 0 & \text{else} \end{cases}, \quad (4.17d)$$

recalling $\omega \sim \frac{2\pi r_s}{N} k$. As before, discrete variables are integers and continuous variables are real. Here, there is the additional constraint that $0 \leq k_0 \leq k_1 < \frac{N}{2} + 1$ and $0 \leq \omega_0 < \omega_1$. Further, the band-pass filter describes a low-pass filter if $k_0 = 0$ or $\omega_0 = 0$ and a high-pass filter if⁵ $k_1 = \lceil \frac{N+1}{2} \rceil$ or $\omega_1 \rightarrow \infty$ (and $\frac{\sin(\omega_1 t)}{\pi t} \rightarrow \delta(t)$). Further, a low- or high-pass filter can be used to describe all such band-pass filters. For example, denoting $y_{\text{hp}}(t; \omega_0)$ as the data $y(t)$ with an ω_0 high-pass filter applied to it, an ω_1 low-pass filter is given by $y_{\text{lp}}(t; \omega_1) = y(t) - y_{\text{hp}}(t; \omega_1)$ and an (ω_0, ω_1) band-pass filter is given by $y_{\text{bp}}(t; \omega_0, \omega_1) = y_{\text{hp}}(t; \omega_0) - y_{\text{hp}}(t; \omega_1)$.

There are many different potential signals beyond those given in Eq. (4.14) and filters beyond those given in Eq. (4.16) and Eq. (4.17), as well as combinations of filters. For example, applying a T rolling average filter, Eq. (4.16), centered at $t_0 = 0$ to the example

⁵Note the use of the ceiling function here, $\lceil x \rceil$, which rounds up to the nearest integer.

signals, Eq. (4.14), also centered at $t_0 = 0$, respectively yields

$$y_{L,\text{avg}}(t; \Gamma, A, 0) = \frac{A\Gamma}{2} \left[\arctan\left(\frac{t+T/2}{\Gamma/2}\right) - \arctan\left(\frac{t-T/2}{\Gamma/2}\right) \right], \quad (4.18a)$$

$$y_{G,\text{avg}}(t; \sigma, A, 0) = \frac{A}{2} \left[\operatorname{erf}\left(\frac{t+T/2}{\sqrt{2}\sigma}\right) - \operatorname{erf}\left(\frac{t-T/2}{\sqrt{2}\sigma}\right) \right], \quad (4.18b)$$

$$y_{\text{sh},\text{avg}}(t; \Gamma, A, 0) = A\Gamma \left[\arctan\left(\sinh\frac{t+T/2}{\Gamma}\right) - \arctan\left(\sinh\frac{t-T/2}{\Gamma}\right) \right], \quad (4.18c)$$

$$y_{\text{sq},\text{avg}}(t; \Gamma, A, 0) = \frac{A}{T} \times \begin{cases} \frac{T+\Gamma}{2} + t & -\frac{\Gamma+T}{2} \leq t \leq -|\frac{\Gamma-T}{2}| \\ \frac{T+\Gamma}{2} - t & |\frac{\Gamma-T}{2}| \leq t \leq \frac{\Gamma+T}{2} \\ \min\{T, \Gamma\} & |t| \leq |\frac{\Gamma-T}{2}| \\ 0 & |t| > \frac{\Gamma+T}{2} \end{cases}. \quad (4.18d)$$

Likewise, for a high-pass filter accepting frequencies in the range (ω_0, ∞) , Eq. (4.17),

$$y_{L,\text{hp}}(t; \Gamma, A, 0) = A \frac{(\frac{1}{2}\Gamma)^2}{t^2 + (\frac{1}{2}\Gamma)^2} \times e^{-\frac{1}{2}\Gamma\omega_0} \left(\cos(\omega_0 t) - \frac{t \sin(\omega_0 t)}{\frac{1}{2}\Gamma} \right), \quad (4.19a)$$

$$y_{G,\text{hp}}(t; \sigma, A, 0) = \frac{A}{\sqrt{2\pi}\sigma^2} e^{-\frac{t^2}{2\sigma^2}} \operatorname{Re} \left[\operatorname{erf}\left(\frac{\sigma^2\omega_1 + it}{\sqrt{2}\sigma}\right) \right], \quad (4.19b)$$

$$y_{\text{sh},\text{hp}}(t; \Gamma, A, 0) = A\Gamma \int_{\omega_0}^{\infty} \frac{\cos(\omega t) d\omega}{\cosh\left(\frac{\pi}{2}\Gamma\omega\right)}, \quad (4.19c)$$

$$y_{\text{sq},\text{hp}}(t; \Gamma, A, 0) = \frac{A\Gamma}{\pi} \int_{\omega_0}^{\infty} \operatorname{sinc}\left(\frac{\omega\Gamma}{2}\right) \cos(\omega t) d\omega. \quad (4.19d)$$

As shown before, the filters will also have an effect on the noise. For simplicity, the noise will be assumed to be constant and uncorrelated $\Sigma[m, n] = \sigma^2\delta_{mn}$. First, for the rolling-average, Eq. (4.16), applying Eq. (4.10) yields the filtered covariance matrix

$$\Sigma_{\text{avg}}[m, n] = \frac{\sigma^2}{T^2} \times \begin{cases} T - |m - n| & |m - n| \leq T \\ 0 & |m - n| > T \end{cases}. \quad (4.20)$$

For this, the correlation between two points decreases linearly to zero the further separated the two points are. Observe that $\Sigma_{\text{avg}}[m, m] = \sigma^2/T$ and down-sampling every T points removes all correlation as one would expect. For the band-pass filter, Eq. (4.17), applying Eq. (4.11) and observing that $|\tilde{h}_{\text{bp}}|^2 = \tilde{h}_{\text{bp}}$, yields

$$\Sigma_{\text{bp}}[m, n] = \sigma^2 h_{\text{bp}}[m - n]. \quad (4.21)$$

4.2 Consistency check

In order to determine if an event signal is detected across the GNOME network, some system is needed to determine whether the data across the network are consistent with

what one would expect. Ideally, this system would minimize the rate of false-positive events due to noise while reliably detecting true-positive events.

Regardless of its nature, a signal in a GNOME magnetometer will have some amplitude that will depend on the vector field or gradient direction along the sensitive axis of the magnetometer. For example, coupling between atomic spins and axion-field gradient, Eq. (3.12), has an amplitude proportional to the axion field gradient and the dot product between the sensitive axis and gradient direction. A localized feature in the axion field, e.g. a domain wall, that crosses the Earth will have some direction, speed, and characteristic magnitude profile. For planar objects and objects much larger than the Earth/network, the network will observe the same magnitude profile — up to a predictable attenuation factor determined by direction — with some signal delay between sensors.

Information about signal amplitudes in various sensors and their relative delays can be used to understand the magnitude and velocity of the originating signal. This information can be used to identify potential signals, determine their (most-likely) characteristics, and determine whether they are consistent with an expected signal. Identifying a class of signals (e.g., signal magnitudes and widths) that can be identified with the analysis algorithm will describe the sensitivity of the experiment. The sensitive “region” in terms of observed parameters (e.g., magnetic field magnitude and signal width) can then be translated into a region of theoretical parameters (e.g., coupling strength and axion mass).

In this section, we discuss two measurements that can be used to determine consistency between what is measured and what one would expect: signal magnitude and direction. For magnitude, the signal amplitudes in each sensor are compared to see if they are consistent with what one expects from a single object passing the network. For direction, the velocity of the (planar) object is compared to the direction of its pseudo-magnetic-field gradient. Due to the approximate planar symmetry of these objects at the scale of the network size, only the component of the velocity/gradient perpendicular to the plane is measurable. In either case, some quantitative measure of consistency is determined. A threshold on this measurement will determine whether or not a signal was detected. The appropriate thresholds are determined through false-positive and false-negative studies of the data and are described later in Sec. 4.6 and Chapter 7.1.

4.2.1 Signal amplitude and sensitivity

The sensitivity of GNOME is partially dependent on the exact manner by which the data are analyzed. In any case, GNOME consists of multiple magnetometers with different noise characteristics and with directional sensitivity to fields. Additionally, the noise characteristics can change and magnetometers can turn on or off over time. It is then useful to define a flexible definition of network sensitivity.

For simplicity, we consider the situation in which neither the network (or Earth) nor the object crossing the network are accelerating. A major implication of this assumption is that the Earth is not rotating; or at least that the object crosses the Earth on a timescale much less than one day. To put this in perspective, if an object traveling took

two hours to cross the Earth, it would need to travel at less than 1% of the galactic orbital velocity, $v_c = 220$ km/s — the typical speed of dark matter objects according to the SHM.

Under the above assumption, the strength of the signal can be described by a single vector \mathbf{m} . The signal in each magnetometer is described by some function of \mathbf{m} , $s_i = d_i(\mathbf{m})$, though this signal will generally appear at different times for different magnetometers. Typically, $d_i(\mathbf{m})$ is linear so that $s_i = \mathbf{d}_i \cdot \mathbf{m}$, where \mathbf{d}_i is in the sensitive direction and $\|\mathbf{d}_i\|$ reflects the coupling strength to the i^{th} magnetometer⁶. One can combine the signals from all magnetometers \mathbf{s} into a single linear equation with the matrix D whose rows are \mathbf{d}_i ,

$$D\mathbf{m} = \mathbf{s}. \quad (4.22)$$

In practice, the signals have some statistical noise so Eq. (4.22) is not an exact equation. The noise in the magnetometers can be described by the covariance matrix $\Sigma_{\mathbf{s}}$. Assuming the noise in the different magnetometers are uncorrelated, $\Sigma_{\mathbf{s}}$ is a diagonal matrix whose i^{th} element in the diagonal is the variance in the noise of the i^{th} magnetometer. Minimizing

$$\chi^2 := (D\mathbf{m} - \mathbf{s})^T \Sigma_{\mathbf{s}}^{-1} (D\mathbf{m} - \mathbf{s}) \quad (4.23)$$

(see Appendix A.2) yields

$$\mathbf{m} = \Sigma_{\mathbf{m}} D^T \Sigma_{\mathbf{s}}^{-1} \mathbf{s} \quad \text{for} \quad \Sigma_{\mathbf{m}} = (D^T \Sigma_{\mathbf{s}}^{-1} D)^{-1}. \quad (4.24)$$

Further, the χ^2 from Eq. (4.23) yields useful statistical information for testing the consistency of the signal with what one would expect. In particular, for N active sensors, the right tail of the χ^2 -distribution with $(N - 3)$ degrees-of-freedom yields a p -value. A p -value close to zero (large χ^2) implies that the signal is inconsistent, while a p -value close to one (small χ^2) implies good consistency.

The magnitude-to-uncertainty ratio, ζ is the strength of the signal with respect to its uncertainty. This can also be thought of as the signal-to-noise ratio for the collective signal across the network. In other words,

$$\zeta = \sqrt{\mathbf{m}^T \Sigma_{\mathbf{m}}^{-1} \mathbf{m}}. \quad (4.25)$$

This quantity will be useful to determine the significance of a signal.

4.2.2 Directional consistency

In the previous section, consistency is discussed with respect to the amplitude of a signal as observed in magnetometers in the network. In addition to signal strength, there is

⁶For example, different magnetometers may couple to the axion field in different but predictable ways. The magnitude of \mathbf{d}_i can be used to normalize this difference so that if a signal passes any two magnetometers with the same direction relative to their respective sensitive axis, then the two magnetometers will effectively measure the same amplitude after this correction.

also timing information that can be used to test the consistency of a signal with the data; namely as it relates to the velocity of the observed object. For planar objects, the direction of the field gradient and velocity are the same. In practice, there are two main ways of determining the velocity of an object which are described below: either calculating a best-fit velocity from the data or through “scanning” the data with different velocities.

If calculated, the timing of at least five signals in individual magnetometers is needed: one to use as reference, three for each spatial dimension, and at least one more for statistics. This will generally require more than five sensors as magnetometers whose sensitive axes are perpendicular to the gradient will not observe a signal. This method also requires some guarantee that the signals in the individual sensors are from the same event. Once the timings are known, the velocity \mathbf{v} can be determined up to uncertainty by solving a linear equation. The direction of \mathbf{v} can be then be compared to \mathbf{m} for statistical consistency.

Alternatively, one can scan over different velocities. In this case, the velocity of the object is assumed, and the analysis is repeated for many different velocity assumptions; in other words, we “scan” over different object speeds and directions. The direction of \mathbf{m} is then compared to the respective scanned direction. Scanning over more velocities is needed to better distinguish quickly changing signals (e.g., due to short duration and large amplitude). Because the velocity is not being measured, this method will generally only require four sensor; three for each spatial dimension and at least one more for statistics.

These two methods come with their own advantages and disadvantages. One major disadvantage to calculating the velocity is that it requires a signal be identified in individual magnetometers. This further limits searches to large, localized signals and may miss smaller signals. Additionally, because the magnetometers are insensitive to vector fields perpendicular to their sensitive axis, more active sensors would be needed to test directional consistency by calculating the velocity compared to scanning velocities. However, scanning velocities comes with the disadvantage that it requires analyzing the entire data set for each scanned velocity. Additionally, some signals may be lost if the object’s velocity is missed because not enough velocities are scanned (determining how many velocities to scan is discussed in Sec. 4.4). At present, the analysis method takes use of scanning velocities for the consistency check. The simplicity from not needing to find signals in individual sensors is a major advantage. However the large number of calculations needed to scan over velocity space remains a limitation in the analysis.

Some method must also be devised to determine if two vectors, \mathbf{m} and \mathbf{v} , have the same direction. For this case, the two vectors will agree if they are either parallel or anti-parallel. Calculating the absolute angle between two vectors, \mathbf{m} and \mathbf{v} , amounts to calculating

$$\theta = \arccos \left(\frac{\mathbf{m} \cdot \mathbf{v}}{\|\mathbf{m}\| \|\mathbf{v}\|} \right).$$

One would like to then determine a statistical measure of how much θ deviates from 0 or π . The main issue with simply applying the typical propagation-of-errors to this equa-




Figure 4.2: Pictorial representation of comparing the directions between two vectors \mathbf{m} (blue) and \mathbf{v} (red). The ellipsoids at the end of the vectors represent their uncertainty, and a plane is drawn containing the two vectors and normal to \mathbf{n} . The angle between \mathbf{m} and \mathbf{v} is given by θ . This angle can be positive or negative depending on whether it runs counterclockwise with respect to the plane's normal \mathbf{n} (here, it is positive).

tion is that propagation-of-errors relies on the assumption that the formula is roughly linear around where it is being evaluated. However, $\cos \theta$ around where $\cos \theta = \pm 1$ is precisely where the function has no linear component; a small perturbation in the vectors results in a vanishingly small change in θ . This problem is avoided when working in two dimensions, because there is a canonical orientation; one can meaningfully define a clockwise and counter-clockwise angle. In this case, the calculation of the angle remains linear, notably a linear perturbation in the vectors may result in a sign change. One solution to this is to define an orientation via an axis of rotation whose perpendicular plane contains \mathbf{m} and \mathbf{v} with the orientation defined via a right-hand rule. A pictorial representation of this is given in Fig. 4.2. The main problem with this method, aside from the added complexity, is that it relies on defining a plane of rotation that may also have uncertainty in its definition.

A simpler alternative would be to apply some condition on the absolute angle between the two directions regardless of the uncertainty in the directions. The main concern then is in how one might designate a meaningful threshold for agreement. This becomes less of an issue when scanning over the velocities because the typical angular spacing in

the scanned velocities provides a meaningful reference for how much the two angles can deviate. If \mathbf{m} deviates from \mathbf{v} by an angular spacing much larger than the angular spacing in the scanned velocities, then one would expect, if the signal were real, that looking at an adjacent scanned velocity would yield better agreement. The directional consistency threshold will be given in terms of the angular spacing such that increasing the threshold will reduce how strict the directional consistency check is.

4.3 Geometric picture

To gain a novel understanding of the amplitude consistency check, it helps to imagine the χ^2 -minimization problem in geometric terms. This is described first in general terms and then applied to the relevant problem.

The main geometric picture to consider is a metric space⁷ on the set of measurements, X , which includes all measurements involved (e.g., magnetic field measurements in all GNOME sensors). Let \mathbf{x} be a given measurement with uncertainty given by the covariance matrix Σ_x . The distance between \mathbf{x} and \mathbf{y} is defined by

$$d(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{y} - \mathbf{x})^T \Sigma_x^{-1} (\mathbf{y} - \mathbf{x})}. \quad (4.26)$$

A “theory” is parameterized by $\mathbf{p} \in P$ (e.g., the velocity of an object) and is described by the function $\mathbf{y} : P \rightarrow X$, where $\mathbf{y}(\mathbf{p})$ describes the expected measurements for the parameter, \mathbf{p} . The set $\mathbf{y}(P) \subset X$ describes a hypersurface of possible measurements assuming the theory is true. From this picture, it is clear that the dimension of the measurement space X must be larger than the dimension of the parameter space P in order to be able to use a measurement to determine the appropriate parameter \mathbf{p} so that $\mathbf{y}(P)$ “fits” in the space X .

The χ^2 -minimization problem is to minimize (as in Eq. (4.23))

$$\chi^2(\mathbf{p}) := (\mathbf{y}(\mathbf{p}) - \mathbf{x})^T \Sigma_x^{-1} (\mathbf{y}(\mathbf{p}) - \mathbf{x}) \quad (4.27)$$

as a function of $\mathbf{p} \in P$. Observe, then, that the χ^2 acts as the distance function $\chi^2(\mathbf{p}) = d^2(\mathbf{x}, \mathbf{y}(\mathbf{p}))$, so minimizing χ^2 is equivalent to finding the point(s) in $\mathbf{y}(P)$ that is closest to \mathbf{x} . Likewise, the χ^2 value for \mathbf{p} is simply the square of the distance between $\mathbf{y}(\mathbf{p})$ and the measurement \mathbf{x} .

It may help to define a “normalized” measurement space \tilde{X} where $\mathbf{x} \rightarrow \tilde{\mathbf{x}} := \Sigma_x^{-1/2} \mathbf{x}$ and $\mathbf{y} \rightarrow \tilde{\mathbf{y}} := \Sigma_x^{-1/2} \mathbf{y}$. Here, $\Sigma_x^{1/2}$ is the “square-root” of the covariance matrix⁸, Σ_x .

⁷Here, “metric space” refers to the formal mathematical object consisting of a set and a real, non-negative function $d(x, y)$ that takes two elements of the set and returns the distance between those elements. This function is commutative $d(x, y) = d(y, x)$, zero if and only if the two inputs are the same, and follows the triangle inequality $d(x, y) + d(y, z) \geq d(x, z)$.

⁸This will be well defined in this case, because Σ_x is real, symmetric, and (strictly) positive definite. This means that one can apply the spectral decomposition theorem to decompose $\Sigma_x = QDQ^T$ where D is a diagonal matrix whose elements are the eigenvalues (all positive) and Q is an orthogonal matrix whose columns are the respective eigenvectors. Then, $\Sigma_x^{1/2} = QD^{1/2}Q^T$ where the square-root acting on the diagonal matrix will act on the elements in its diagonal.

Under this construct, the distance formula becomes $d(\mathbf{x}, \mathbf{y}) = \tilde{d}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = \|\tilde{\mathbf{y}} - \tilde{\mathbf{x}}\|$, the typical L^2 norm. In effect, what has been done is that the signal was normalized by the noise; if the noise in each sensor is uncorrelated, Σ_x is a diagonal matrix whose elements are the variance, so $\Sigma_x^{-1/2}$ is a diagonal matrix whose elements are the inverse standard deviation.

Now consider a linear theory where $\mathbf{y}(\mathbf{m}) = D\mathbf{m}$ for some known matrix D , as in the the amplitude consistency check. Here, $\mathbf{y}(P)$ is a flat hyperplane that intersects the origin. In the linear case, as previously shown, χ^2 -minimization is explicitly solvable. What is particularly interesting is that the magnitude-to-uncertainty ratio from Eq. (4.25) becomes

$$\begin{aligned} \zeta &= \sqrt{\mathbf{m}^T \Sigma_m^{-1} \mathbf{m}} \\ &= \sqrt{(D\mathbf{m})^T \Sigma_s^{-1} (D\mathbf{m})} \\ &= \sqrt{\mathbf{y}^T \Sigma_s^{-1} \mathbf{y}} \\ &= \|\tilde{\mathbf{y}}\|. \end{aligned}$$

Now, both the magnitude-to-uncertainty ratio ζ and χ^2 can be interpreted as distance measurements. Specifically, decomposing the measurement \mathbf{x} into a component perpendicular and parallel to the $\mathbf{y}(P)$ hyperplane, then $\sqrt{\chi^2}$ is the perpendicular distance while ζ is the parallel distance.

4.4 Velocity lattice

The analysis described here relies on considering signals caused by objects coming from different directions. In order to ensure objects are observed regardless of their velocity (within reason), the analysis algorithm scans over a region in velocity space; i.e., the three-dimensional space of possible velocities. In order to accomplish this, a scan over discrete velocity in a lattice⁹ is defined. This lattice must be sufficiently dense as to not miss signals while not so dense as to be computationally impractical.

Each velocity in the lattice will serve two primary functions: correcting for delays between signals observed by different sensors and directional consistency. To understand how to adjust for delays, let magnetometer i be at¹⁰ \mathbf{x}_i with the signal $s_i(t)$, and let some domain wall have perpendicular velocity \mathbf{v}_\perp . The timing delay with respect to the

⁹Here, “lattice” is being used informally to describe some array with some conditions on typical spacing between points, but may not have a simple relationship between the location of points (e.g., via a discrete translational symmetry).

¹⁰The exact coordinate system is not particularly relevant. For example, a galactic coordinate system is useful for orienting with respect to the SHM. However, it is often more practical to define an Earth-centered coordinate system wherein the origin is placed at the center of Earth, the z -axis is along the Earth’s axis, the x -axis points from the Earth’s center to (0° N, 0° E), and the y -axis follows from the right-hand rule.

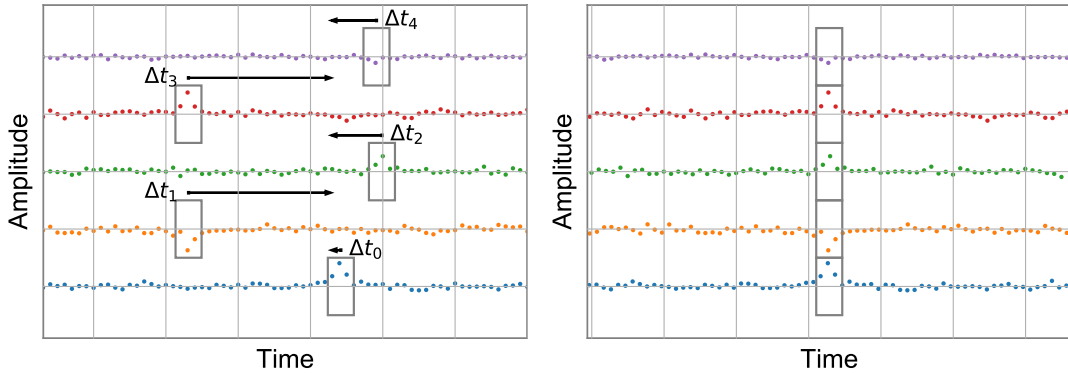


Figure 4.3: Visualization of time-shifting signals on simulated data. Each color represents a different sensor and vertical offsets are used for clarity. (left) Signals are inserted in the position marked by the boxes. (right) Signals are shifted over by their respective Δt .

origin would then be

$$t_i = \frac{\mathbf{v}_\perp \cdot \mathbf{x}_i}{\|\mathbf{v}_\perp\|^2}. \quad (4.28)$$

Shifting signals

$$s_i(t) \rightarrow s_i(t - t_i) \quad (4.29)$$

will remove delays between signals (see Fig. 4.3). The time t is then the time with respect to a hypothetical sensor at the origin. It should be noted that additional corrections may be needed to account for the Earth's rotation¹¹ (or other acceleration and relative motion of sensors in the network). In addition to adjusting for delays, the velocity is also used as a directional consistency check. This is discussed in Sec. 4.2.2.

Each point on the velocity lattice will effectively cover some region of velocity space. Covering all speeds and directions would require an infinitely large lattice; small speeds require asymptotically more lattice points due to large signal delays. One solution is to consider an expected distribution of object velocities and attempt to cover some percentage of that distribution at the cost of a higher false-negative rate. This distribution is the SHM (described in Sec. 2.7).

4.4.1 Lattice requirements

The density of points needed for a given volume of velocity space can be determined by considering by how much the relative timing of a signal varies based on the direction and speed of the signal. Consider two sensors located at \mathbf{x}_1 and \mathbf{x}_2 , so $\Delta \mathbf{x} = \mathbf{x}_2 - \mathbf{x}_1$,

¹¹For particularly slow domain walls, $2R_\oplus/v_\perp \gtrsim 1$ day, merely shifting the time of a signal will fail because a single domain wall can be observed multiple times.

and a domain wall with (perpendicular) velocity \mathbf{v}_\perp . The relative timing is then

$$\Delta t = \frac{\Delta \mathbf{x} \cdot \mathbf{v}_\perp}{\mathbf{v}_\perp \cdot \mathbf{v}_\perp} = \frac{\Delta x}{v_\perp} \cos(\theta), \quad (4.30)$$

where θ is the angle between \mathbf{v}_\perp and $\Delta \mathbf{x}$. Now, consider a small perturbation dv_\perp or $d\theta$,

$$d\Delta t = -\frac{\Delta x}{v_\perp^2} \cos(\theta) dv_\perp - \frac{\Delta x}{v_\perp} \sin(\theta) d\theta.$$

Now, the density for the lattice can be determined by demanding that the relative timing change by less than τ so $|d\Delta t| < \tau$. Further, let all the sensors be contained in a sphere of radius R (so $\Delta x \leq 2R$) and make no assumption on the angle θ (so $|\sin \theta| \leq 1$ and $|\cos \theta| \leq 1$). This leads to the following conditions on the spacing between the angles and speeds of points in the velocity lattice,

$$\Delta v_\perp < \frac{\tau v_\perp^2}{2R}, \quad (4.31a)$$

$$\Delta \theta < \frac{\tau v_\perp}{2R}. \quad (4.31b)$$

A lattice that satisfies both of these conditions will be sufficiently dense to guarantee that error in the timing will not exceed τ . Some natural choices for this is to have $R \gtrsim R_\oplus \approx 6371$ km because GNOME is located on the Earth and τ approximately the sampling period of the data (or less).

Defining a lattice that satisfies Eq. (4.31) can be accomplished using a variety of algorithms. Though there are many well-established, three-dimensional lattices, they do not typically follow a spherical pattern that would be necessary here. Instead, a Fibonacci lattice is used to cover spheres at different radii with sufficient angular density. The idea of this lattice is that each sequential point on the surface of the sphere differs by some angle that is the product of the irrational number $2\pi\varphi$ for golden ratio $\varphi := \frac{1+\sqrt{5}}{2}$. The polar angle is then chosen to evenly cover the sphere. A sphere with N_θ points covering a polar range up to Θ will then consist of the points

$$(\theta, \phi) \in \left\{ \left(\arccos \left(1 - \frac{1 - \cos \Theta}{N} (j + j_0) \right), 2\pi\varphi j \right) \right\}_{j=0}^{N_\theta-1}, \quad (4.32)$$

where $j_0 \in [0, 1)$ is a phase choice. If $j_0 = 0$, the pole is included, and $j = 0.5$ is often used to give a more “random” lattice.

In order to match the density requirements, Eq. (4.31), there must be a sufficiently large number of points in each Fibonacci lattice and enough lattices are placed with different radii (i.e., speeds). To understand the needed density of points on a lattice, consider the solid angle covered by a circular region whose radius spans a $\Delta\theta/2$ arc on the sphere. This solid angle is $4\pi \sin^2(\Delta\theta/4)$ steradians. Observe that this can be thought of as the solid angle covered by a single point of the lattice with the appropriate spacing, because two such points with adjacent circular regions will be $\Delta\theta$ apart. Considering




Figure 4.4: Example of velocity lattice using a series of Fibonacci spheres. Color designates different speeds and grey spheres of constant speed are used as visual guides.

that there are $4\pi \sin^2(\Theta/2)$ steradians covered by the polar angle Θ , the number of points should be

$$N_\theta \gtrsim \frac{\sin^2(\Theta/2)}{\sin^2(\Delta\theta/4)}. \quad (4.33)$$

The lower bound here assumes that the lattice points are perfectly placed to exactly cover the desired region of the sphere. However, this is not generally true, particularly for $N_\theta > 2$, which would leave holes and redundancies if the lower bound was used exactly.

In addition to ensuring some angular density, enough speeds are needed to satisfy Eq. (4.31a). Let $\{v_n\}$ be the set of scanned speeds for $n = 0, 1, \dots, N_v$. These satisfy the equation

$$\int_{v_n}^{v_m} \frac{dv_\perp}{\Delta v_\perp} = m - n,$$

noting that Δv_\perp is a function of v_\perp . With this constraint in mind and setting v_{N_v-1} with a desired v_0 (i.e., to cover a range that includes $[v_0, v_{N_v-1}]$, but may also include

lower speeds), the scanned speeds are

$$\begin{aligned} v_n &= \left[\frac{1}{v_{N_v-1}} + (N_v - 1 - n) \frac{\tau}{2R} \right]^{-1} \\ \text{for } N_v &= \left\lceil \left(\frac{1}{v_0} - \frac{1}{v_{N_v-1}} \right) \frac{2R}{\tau} + 1 \right\rceil. \end{aligned} \quad (4.34)$$

A visualization of a velocity lattice generated as described here is shown in Fig. 4.4.

Because larger lattice size corresponds to additional resource requirements, it helps to gain a better understanding how lattice size scales with the input parameters: Θ , v_0 , v_{N_v-1} , τ , and R . Without any approximations or constraints, the total number of points will be

$$N = \left\lceil \left(\frac{1}{v_0} - \frac{1}{v_{N_v-1}} \right) \frac{2R}{\tau} \right\rceil \left\lceil \frac{\sin^2(\Theta/2)}{\sin^2\left(\frac{1}{4}\left(n + \frac{2R}{\tau v_{N_v-1}}\right)^{-1}\right)} \right\rceil. \quad (4.35)$$

This equation is somewhat cumbersome in terms of understanding how lattice scales with the various parameters. Consider the limit $v_{N_v-1} \rightarrow \infty$ (and omitting the lattice point at infinity) as well as the case wherein $\frac{2R}{\tau} \gg v_0$. Under these conditions,

$$N \approx \frac{16}{3} \left(\frac{2R}{v_0\tau} \right)^3 \sin^2(\Theta/2). \quad (4.36)$$

4.4.2 Lattice optimization

Above, a method was described to generate a velocity lattice for covering some volume of velocity space. However, the exact volume to cover is a matter of choosing a region containing some percent of the EPDF, Eq. (2.25). There are many ways of accomplishing this, so the region will be chosen to minimize the number of lattice points.

The shape of the region in velocity space to use is entirely arbitrary. However, keeping in mind that there is an axial symmetry in the distribution about the Earth's galactic velocity, \mathbf{v}_L , it is natural to require the region to likewise have this symmetry. Meanwhile, the velocity lattice requires more points at lower speeds; requiring smaller spaces between both the angular and speed spacing. Thus, the family of regions to be considered here will consist of the volume bounded within some polar angle Θ of $-\mathbf{v}_L$ and further from the origin than some speed v_\perp .

First, one can constrain the shape of the region covered by the velocity lattice by observing the Cumulative Distribution Function (CDF); see Fig. 4.5. Demanding that the CDF achieve (at minimum) some percentage will constrain the polar angle and minimal speed of the region. The larger this percentage, the smaller chance of missing a signal (assuming that the signals follow the SHM) but at the cost of requiring more lattice points.

After constraining the shapes of the region covered by the velocity lattice to guarantee some coverage of the SHM EDPF, what remains is to minimize the resource requirements;




Figure 4.5: The CDF in terms of speed and polar angle Θ relative to $-\mathbf{v}_L$. That is, the probability that an observed event is oriented within the given polar angle and whose speed is greater than the given speed. The thick black line is drawn at 95%. For this CDF, $v_c = v_L = 220$ km/s and $v_e = 550$ km/s.

namely the number of lattice points needed to cover the region. After the previous constraints, the family of regions is one-dimensional and can be parameterized by polar angle coverage. Applying Eq. (4.36) to the minimum speed and polar angle coverage of this family of regions, one obtains the number of lattice points needed for the different regions (see Fig. 4.6). One observes that the case in which all directions are covered is optimal.

4.5 Procedure

Throughout this chapter, the concepts needed to analyze the data from GNOME were established. In this section, these ideas will be summarized.

The first step is naturally data collection. Sensors in the GNOME network are operated by various groups around the world. The collaboration will often designate specific times, referred to as “Science Runs,” to operate as many sensors in the network as possible. Other times are used for upgrades, maintenance, as well as off-run operation. The data have been collected for the following science runs:

- Science Run 1 (6 June–5 July 2017) consisted of six sensors.
- Science Run 2 (29 November–22 December 2017) consisted of nine sensors.
- Science Run 3 (1 June 2018–10 May 2019) consisted of ten sensors.

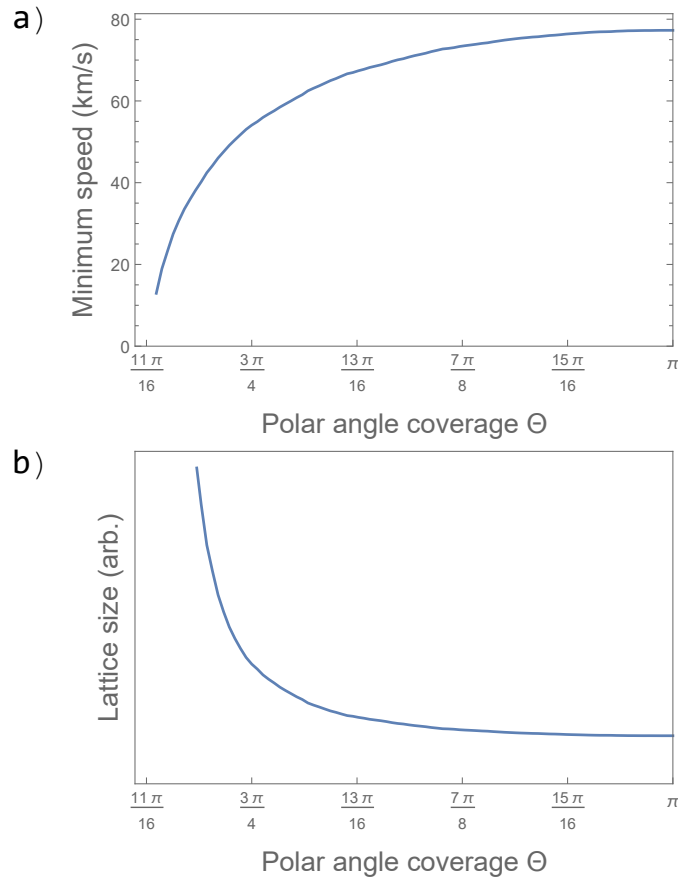


Figure 4.6: Plots relevant for optimizing the lattice size. The SHM uses the same conditions as in Fig. 4.5 with 95 % probability. (a) The minimum speed vs polar angle coverage when demanding 95 % coverage of the SHM. For full polar angle coverage, the minimal speed would be 77 km/s. (b) The lattice size as a function of polar angle coverage. The size is arbitrary as it varies approximately by a proportionality factor based on τ .

- Science Run 4 (30 January–30 April 2020) consisted of nine sensors.
- Science Run 5 (23 August–31 October 2021¹²) consists of eleven sensors.

The data are recorded in centralized servers to allow for analysis of every sensor’s data. The raw magnetometer data is stored along with sanity channel data that states whether the data at that time is reliable based on ancillary information. Additionally, metadata are recorded such as calibration information, sensor location and orientation, and timing.

After data are collected, they need to be read by the analysis algorithm. First the data are read into the program and calibrated into magnetic field units. Then, data with bad sanity are removed. Filters are applied to remove systematic noise. This includes powerline frequencies at 50 Hz or 60 Hz depending on the local electric grid and a high-pass filter to remove long-term drifts and constant offset in the magnetic field measurement. This step of reading and processing the data can be performed through a buffer. This allows one to read segments of data into memory as opposed to all data.

A velocity lattice is generated using requirements and techniques described in Sec. 4.4. Each velocity corresponds to a different time-shifting, Eq. (4.29). If there is directional bias in the lattice based around galactic coordinates (e.g., to prefer velocities opposite to that of Earth’s motion in the galaxy), then the distribution would change as the Earth rotates. In this case, the amount by which time is shifted also varies with time, and there may be a slightly more complicated process for synchronizing signals. However, this issue is moot if all directions are covered. From this stage, one can make efficient use of multi-processing with a process for each velocity in the lattice. Analyses with different time-shifts are essentially independent.

Once the data are time-shifted so that potential signals are aligned, the collective data across the network are tested for consistency with a domain wall (or other signal pattern of interest). This amounts to following the procedures described in Sec. 4.2 at every time to determine a magnitude vector that best agrees with the observed signals as well as a p -value quantifying this agreement. Consistency amounts to demanding that the magnitude vector is in roughly the same direction (within some threshold) as the respective velocity vector used to shift the signals and that the p -value is below some threshold. How this threshold is determined is described in Sec. 4.6.

There may be observed, consistent signals at consecutive times or with different velocities at the same time. These are likely the result of a single event and should be grouped together as a single event to prevent double-counting. Once grouped, one can take, e.g., the time and velocity with the largest p -value as representative because it shows good agreement with expectation. One may further require that the event have some minimal magnitude-to-uncertainty ratio in order to avoid counting the many “null events” in which a signal is found to be consistent with zero. The characteristics of the events are recorded: time, velocity, magnitude vector, p -value, and magnitude-to-uncertainty ratio.

Once recorded, the potential events can be used to argue for discovery. However, as one would expect the background to produce some number of false-positive events

¹²Ongoing.

regardless of whether or not a domain wall was discovered, some care is needed. Either a very large signal or many smaller signals would point to such a discovery. In either case, a discovery of ALP domain walls or lack thereof would allow a better understanding of the possible characteristics of ALPs or similar particles. Further discussion of translating results to parameter space is given in Chapter 5.

4.6 Choosing thresholds

There are a few thresholds discussed as part of the analysis algorithm, especially for testing consistency between observation and the expected signal from an ALP domain wall; namely p -value and directional consistency. The directional consistency threshold is defined as the proportion of the angular separation between the magnitude and velocity vectors and the angular separation of velocities in the lattice (at the respective speed). Additionally, understanding the distribution of magnitude-to-uncertainty ratios is needed to determine the existence of ALP domain walls and determine constraints.

Determining appropriate thresholds to use in the analysis relies on understanding the characteristics of the data being analyzed. This includes understanding what type of signals will reliably be observed in the data as well as how likely it is to find false-positive signals. These heavily rely on understanding the characteristics of the background data. In many experiments, background data can be taken as part of an ancillary experiment by taking data with foreground signals either removed or blocked. However, this is not possible for GNOME. Instead, background data are generated by combining data from random times in all the sensors in the network (including times in which sensors were inactive) via “time-shuffling.” These data are then treated as if they were taken synchronously. As long as the data are shuffled by times much larger than a domain-wall-crossing time, any foreground signals in an individual magnetometer should be rejected as they do not appear in other magnetometers.

A false-negative analysis is used to determine how reliable the analysis algorithm can identify signals. This analysis consists of inserting signals into time-shuffled background data. The false-negative rate is the percent of inserted signals that are missed, and the percent of inserted signals that are observed is the detection probability. The rate varies with the thresholds and characteristics of the inserted signals. So long as the inserted signal is large enough to be distinguishable from noise over the time used in the shuffled data, the choice of thresholds should not play a major role; determining what types of signals should be observable is explored in Chapter 5. One can demand a minimal detection probability to restrict the p -value and directional consistency thresholds. Further restrictions can be made to minimize the false-positive rate.

The false-positive analysis is used to characterize the background data. This analysis consists of running the analysis algorithm on time-shuffled background data without inserted signals and determining the rate of potential events; particularly with respect to their magnitude-to-uncertainty ratios. This information is used to determine whether the signals measured by GNOME can be explained by the background alone.

Chapter 5

Sensitivity

The sensitivity region describes the class of signals that could be observed by an experiment. This can be described in terms of experimental parameters, such as the directly observed signal, or theoretical parameters, such as coupling strength. Generally, one wants to understand the region of physical parameter space to which an experiment is sensitive. There could be multiple reasons that a signal is not observed; perhaps the experiment is simply not sensitive enough to observe the signal, or the signal is too rare to be observed in the finite data available. These conditions restrict the sensitivity region up to some level of confidence. If nothing is observed from the experiment, the sensitive region can be understood as the excluded region of physical parameter space.

The sensitivity is formally defined in Sec. 5.1. Physical parameters related to the ALP theory described in Chapter 3 are connected to the class of observable domain-wall signals in Sec. 5.2. Methods for generating an exclusion/sensitivity region of parameter space for an experiment are described in Sec. 5.3. Finally, the manner in which the sensitivity of the network scales with its size and other parameters is discussed in Sec. 5.4.

5.1 Network sensitivity

In order to establish the types of signals that can be detected, first the sensitivity of the network must be assessed. The exact sensitivity depends on the analysis methods, so first an abstract approach will be used. Consider a function $\mathcal{A}(\mathbf{m}; \Sigma_s, D)$ which returns the strength of the collective response given the magnetometer characteristics D (namely the sensitive axes) and magnetometer noise Σ_s for a signal \mathbf{m} . One could potentially consider other information (e.g., the time of the signals) but this should suffice for this work. A signal is observable if $\mathcal{A}(\mathbf{m}; \Sigma_s, D) > \zeta_0$, where ζ_0 is some threshold.

The sensitivity of the network can be described by the minimal signal needed to induce an observable signal. Typically, \mathcal{A} is absolutely scalable¹, so one can re-write the

¹Absolutely scalable means that one can “pull-out” a positive scalar from the function. For example, let f be an absolutely scalable function, then $f(\mathbf{v}) = \|\mathbf{v}\| f(\hat{\mathbf{v}})$. As long as $\mathcal{A}(\mathbf{m})$ is strictly monotonically increasing in $\|\mathbf{m}\|$, one can always rescale it to be absolutely scalable. The only conditions where this breaks down is when larger signals do not imply a signal is more observable.

condition for a signal to be observable as

$$\|\mathbf{m}\| > \beta_{\zeta_0}(\hat{\mathbf{m}}) := \frac{\zeta_0}{\mathcal{A}(\hat{\mathbf{m}}; \Sigma_{\mathbf{s}}, D)}, \quad (5.1)$$

where $\beta_{\zeta_0}(\hat{\mathbf{m}})$ is the minimal magnitude needed to observe a signal in the $\hat{\mathbf{m}}$ direction with the threshold ζ_0 . Because $\beta_{\zeta_0} = \zeta_0 \beta_1$, one can define $\beta := \beta_1$ for simplicity.

The sensitivity described in Eq. (5.1) is still dependent on the direction of the signal, $\hat{\mathbf{m}}$. It is often useful to consider a direction-independent sensitivity. There are a couple manners by which this can be achieved. For example, the sensitivity could be averaged over some distribution of directions $\rho(\hat{\mathbf{m}})d^2\Omega$,

$$\bar{\beta} = \int d^2\Omega \rho(\hat{\mathbf{m}}) \beta(\hat{\mathbf{m}}), \quad (5.2)$$

or the largest (and “worst”) sensitivity,

$$\bar{\beta} = \max_{\hat{\mathbf{m}}} \beta(\hat{\mathbf{m}}). \quad (5.3)$$

The latter case will be of particular interest as it is also an upper-bound for the former. That is, any bounds on parameter space made with Eq. (5.3) will contain the same bound made with Eq. (5.2).

The generalized method of describing sensitivity can be applied to the analysis methods used for the GNOME analysis. For this, let \mathcal{A} be the magnitude-to-uncertainty ratio from Eq. (4.25), so $\bar{\beta}$ becomes

$$\bar{\beta} = \max_{\hat{\mathbf{m}}} \left(\hat{\mathbf{m}}^T (D^T \Sigma_{\mathbf{s}}^{-1} D) \hat{\mathbf{m}} \right)^{-1/2}. \quad (5.4)$$

This can be maximized by solving for the eigenvalues for $(D^T \Sigma_{\mathbf{s}}^{-1} D)$. If $\hat{\mathbf{m}}$ is an eigenvector with eigenvalue λ , then $\beta(\hat{\mathbf{m}}) = \lambda^{-1/2}$. By expressing $\hat{\mathbf{m}}$ in the eigenbasis, it is clear that the lowest and highest eigenvalues will determine the extrema for $\beta(\hat{\mathbf{m}})$. In other words, the above equation has a maximum of $\bar{\beta} = \lambda_{\min}^{-1/2}$ where λ_{\min} is the smallest eigenvalue of $(D^T \Sigma_{\mathbf{s}}^{-1} D)$. Likewise, the corresponding eigenvector would be the least-sensitive direction.

5.2 Connection to physical parameters

The network sensitivity will translate into some set of physical parameter space of signals that can be observed. In this section, focus will be placed on whether or not an object will be observed given that it crossed the network and not on whether an object is simply too rare to be observed over the experiment time.

Particular focus will be placed on axion domain walls (Sec. 3.1) with linear-spin coupling (Sec. 3.2.1). At the center of the domain wall, where the signal is strongest, the

signal strength is characterized by Eq. (3.15), defining the normalized pseudo-magnetic field strength,

$$\frac{g_{F,j}B_j}{\sigma_j\eta_j} = \frac{4}{\mu_B}m_a\xi =: \mathcal{B}_p. \quad (5.5)$$

The magnitude vector \mathbf{m} is defined precisely so that

$$\mathcal{B}_p = \|\mathbf{m}\|$$

by including a factor of $\frac{g_{F,j}}{\sigma_j}$ in each row of D .

The network sensitivity is related to physical parameters by noting the connection between \mathbf{m} given by Eq. (5.5) and the condition for a signal to be observable, Eq. (5.1). Namely that if a signal is observable with magnitude $\zeta_0\bar{\beta}$, then one should be able to observe a domain wall in an axion field with physical parameters (m_a, ξ) if the following is satisfied,

$$\frac{1}{\zeta_0 A_{\text{filt}}(m_a)} \frac{4m_a\xi}{\mu_B} = \frac{\mathcal{B}'_p}{A_{\text{filt}}(m_a)} > \bar{\beta}, \quad (5.6)$$

where A_{filt} is an attenuation factor due to filtering that depends on signal mass (or width) and $\mathcal{B}'_p := \mathcal{B}_p/\zeta_0$ is the strength of a the signal needed to induce a magnitude-to-uncertainty ratio of one (prior to filter attenuation).

Though linear coupling between fermionic spins and the gradient of an axion domain wall was considered here, other objects and interaction terms could be included with similar methods. All that one needs is a method of translating physical parameters to the strength of an induced signal. The analysis method described here focuses on the largest magnitude that a signal induces, independent of the signal shape. This means that the ability to observe a signal does not depend on the signal shape; though filters may significantly attenuate signals.

5.3 Sensitive region of parameter space

Beyond determining if a given object could be observed, one must consider how likely the object would pass through the network. This is based largely on the density of the objects. The typical rate of events is given by

$$r = \frac{\bar{v}\rho_{\text{DW}}}{\sigma_{\text{DW}}} = \frac{\bar{v}\rho_{\text{DW}}}{8m_a f_a^2}, \quad (5.7)$$

where $\sigma_{\text{DW}} = 8m_a f_a^2$ is the surface tension (Eq. (3.9)).

Understanding the likelihood of observing an event relies largely on Poissonian statistics. These ideas are explored in-depth in Appendix C. The Poissonian probability mass function (PMF), i.e., the probability of observing k events given that μ are expected, is given by

$$P(k; \mu) = \frac{\mu^k}{k!} e^{-\mu}. \quad (5.8)$$

There are two approaches for determining the significance of an event considered here. In the first approach, one considers only the probability of observing at least one event. The signal from this event must be large enough that it could not be explained by background alone. For this, the minimum magnitude-to-uncertainty ratio threshold ζ_0 must be fairly large. However, if no such events are found, one can generate bounds using the largest-observed magnitude-to-uncertainty ratio because there were no larger event. This method is the “loudest-event” method. The second approach involves comparing the rate of events from the background measurements to the observed rate at all magnitude-to-uncertainty ratios. This method will be referred to as the “extended” method as the loudest-event method can be understood as a particular limit of the extended method. The basis for the extended method is described in Appendix C.3.

5.3.1 Loudest-event method

In the first approach, one must consider the probability of observing no signals. In addition, even if one would expect an event to be observable, it is possible for it to be missed due to stochastic noise or other limitations in the analysis. Let $\epsilon \in [0, 1]$ be the detection efficiency and $\mu = r\tilde{T}$ be the expected number of events for event rate r (including those missed by the analysis) and active time \tilde{T} , then the probability that no events are observed is

$$\sum_{k=0}^{\infty} (1-\epsilon)^k \frac{(r\tilde{T})^k}{k!} e^{-r\tilde{T}} = e^{-\epsilon r\tilde{T}}. \quad (5.9)$$

Note that the factor of ϵ can be absorbed into r .

One can then define a lower bound on observable event rates $R_{\tilde{C}}$ at confidence level \tilde{C} by demanding that there be $1 - \tilde{C}$ chance of observing no events, or

$$r \geq R_{\tilde{C}} = \frac{-\log(1 - \tilde{C})}{\epsilon\tilde{T}}. \quad (5.10)$$

If no events are found, the rate $R_{\tilde{C}}$ can be understood as an upper bound on the event rate. Combining Eq. (5.7) and Eq. (5.10) yields the bound on observable physical parameter space (recovering factors of \hbar and c)

$$f_{\text{int}} \leq \frac{\hbar}{\xi} \sqrt{\frac{-\bar{v}\rho_{\text{DW}}\epsilon}{8m_a \log(1 - \tilde{C})}} \tilde{T}, \quad (5.11)$$

recalling $f_{\text{int}} := f_a/\xi$. Also, the time \tilde{T} is dependent on the time during which the network is sensitive to the given parameter space. This may seem to be circular, but the left-hand side of Eq. (5.11) depends on the interaction strength f_{int} while the right-hand side depends on the mass m_a and ratio ξ ; noting from Eq. (5.5) that the strength of a signal depends on (m_a, ξ) and filter attenuation depends only on m_a . Further discussion of the active time is given in Sec. 5.3.3.

5.3.2 Extended method

Beyond the loudest-event method, the extended method tries to find any excess event rate above the background; including under conditions where there is a large false-positive rate. For this, one establishes a false-positive rate (with respect to total experiment time, T) as a function of the magnitude-to-uncertainty ratio threshold ζ_0 . Likewise, using a false-positive background data analysis, one can establish the expected rate of positive signals in the absence of domain walls. Using methods described in Appendix. C.3, one can establish an upper-bound on the rate $R_{\tilde{C}}$ at confidence \tilde{C} to exclude the rates $\epsilon\bar{r} := \frac{\tilde{T}}{T}\epsilon r > R_{\tilde{C}}$. With Eq. (5.7), this translates to the bound (recovering factors of \hbar and c)

$$f_{\text{int}} \leq \frac{\hbar}{\xi} \sqrt{\frac{\bar{v}\rho_{\text{DW}}\epsilon}{8m_a T R_{\tilde{C}}}} \tilde{T}. \quad (5.12)$$

The loudest-event method yields roughly the same results as the extended method bound when only large magnitude-to-uncertainty ratio thresholds are considered. Prior to analyzing the data, one can determine the smallest magnitude-to-uncertainty ratio needed for a single observed event over the experiment time to be significant. If the loudest event observed is slightly below this threshold, and the extended method is only considers the excess event rate at the largest magnitude-to-uncertainty ratio, then these two methods are identical. However, because the extended method also includes smaller magnitude-to-uncertainty ratios that are especially of interest when testing for low-mass ALPs.

5.3.3 Active time

The time during which the network is sensitive to domain walls with a particular set of physical parameters is important because it is entirely possible for domain walls to be missed because their signal is too weak, regardless of how many cross the network. The active time can be parameterized using two variables: the signal duration Δt (or, equivalently the mass, because $m_a \propto (\bar{v}\Delta t)^{-1}$) and the pseudo-magnetic field needed to induce a magnitude-to-uncertainty ratio of one (prior to filter attenuation) $\mathcal{B}'_p := \mathcal{B}_p/\zeta_0$. Based on the definitions in Eq. (3.8) and Eq. (5.5), these can be understood in terms of the physical parameters as follows (recovering factors of \hbar and c),

$$\Delta t = \frac{2 \cosh^{-1}(2)}{\bar{v}m_a c/\hbar} \quad \text{and} \quad \mathcal{B}'_p = \frac{4m_a c^2 \xi}{\mu_B \zeta_0}. \quad (5.13)$$

The function $\tilde{T}(\Delta t, \mathcal{B}'_p)$ can be determined by considering the set of signals observable at any given time — which is determined by the network sensitivity and filter attenuation factor A_{filt} — then integrating over the experiment time T .

5.4 Scaling

The manner in which the size of the network affects the sensitivity. The scaling of the sensitivity is formally derived in Appendix A.3. If σ is the typical noise (i.e., standard deviation) of an individual sensor, there are n magnetometers in the network, and the magnetometers' sensitive axes are evenly distributed, then the network sensitivity is

$$\beta \approx \frac{\sigma}{\sqrt{n/3}}. \quad (5.14)$$

This follows the typical \sqrt{n} scaling. In addition to an improvement in overall sensitivity, additional sensors will improve the directional coverage of the network and improve reliability of the network. If there are many sensors, as long as they are not injecting bad data into the network, an individual sensor can go offline without dramatically affecting the network and the effects from small errors in an individual sensor can be suppressed by data from other sensors.

The manner in which the sensitive region of physical parameter space changes with the characteristics of the network is practically important when trying to observe unexplored parameter space. The physical parameter space can be parametrized by mass m_a and any two of $\{f_a, f_{\text{int}}, \xi := f_a/f_{\text{int}}\}$. Often, one reduces the parameter space by one dimension by fixing the ratio ξ . The class of signals observed from an experimental perspective are compared to physical parameters in Eq. (5.13).

The mass is related to signal duration with a simple relationship $m_a \propto \Delta t^{-1}$, so, holding the rate and signal magnitude constant, sensitivity to longer signals is equivalent to sensitivity to lower mass (and vice versa). This can be accomplished by choosing filters that do not remove lower-frequency components. When filtering, consider two extremes; no/minimal filtering and matched filter. Without filtering or only filtering enough to remove systematic noise, sensitive parameter space will be largely independent of mass (limited only by the few unavoidable filters and limits in the sensors such as sample rate and bandwidth). In the matched filter case, a convolution is performed between the data and the shape of the desired signal — i.e., a peak of some known width. This is an optimal filter for amplifying the magnitude-to-uncertainty ratio for a particular signal width (or mass) at the cost of worsening sensitivity to all other signal shapes.

The magnitude of a signal does not depend on the energy scales, f_a or f_{int} , but on the mass m_a and ratio ξ as $\mathcal{B}_p \propto m_a \xi$. Thus, a more massive ALP would result in a larger signal. A larger signal would also be observable for more time over the course of the experiment, but would also imply the existence of fewer domain walls. From Eq. (5.11) and Eq. (5.12), the upper bound on observable coupling/decay constant scales $\propto m_a^{-1/2}$, though as long as a given mass is observable, there is also a $\propto \tilde{T}^{+1/2}$ scaling in these energy scales. For the same active time \tilde{T} and ratio ξ , increasing the number of sensors n will then improve sensitivity to a smaller mass as $n^{1/2}$ and to the energy scales, f_a or f_{int} , as $n^{1/4}$ at these lower masses. Likewise, for constant mass, the sensitivity to a smaller ratio ξ improves as $n^{1/2}$ meaning that, though the axion-decay-constant bounds will not improve, the coupling scale f_{int} will improve by $n^{1/2}$. The exact scaling of the effective time with network sensitivity will depend also on the choice of filters.

Another useful scaling to understand is related to how many velocities need to be scanned in the analysis algorithm. This lattice was described in Sec. 4.4 to be dependent on some time τ describing the maximal error in time-shifting. In practice, this is typically based on the sample rate of the data (which is also typically down-sampled by the averaging time). From Eq. (4.36), the number of points in the lattice scales as $N \propto \tau^{-3}$. Thus, by doubling the averaging time (i.e., halving the sampling rate), one also requires an eighth as many lattice points, which directly translates into computational requirements. The longer averaging time will also decrease the mass to which the network is sensitive.

Computational complexity scale linearly with the size of the network. Many operations are performed on the individual magnetometers such as filtering and noise estimations. The data from the different magnetometers are used together in the consistency check, including matrix operations such as inversions; which typically have quadratic scaling with size. However, the only n -dimensional matrix that is inverted is the covariance matrix Σ_s , which is diagonal for uncorrelated noise, so inverting this matrix scales linearly. The other inversion is of the $d = 3$ -dimensional matrix, $D^T \Sigma_s^{-1} D$, whose computational complexity is constant. The remaining operations include various matrix multiplications, which have either linear or constant scaling with n . Similarly, memory requirements will generally scale linearly as one needs to read-in each additional sensor; though there may be implementations of this analysis wherein quadratic memory scaling is exchanged for faster computation of some matrix operations. Similar arguments can be made to find that computational complexity also scales linearly with time.

Chapter 6

Experimental data

Before discussing the implementation of the analysis described in the previous chapters, it is useful to understand how the data are stored as well as establish some idea of what the data look like.

The sensors in GNOME are located around the world, each with their own design. The types of magnetometers are summarized in Table 6.1 while the location and orientations of the magnetometers are listed in Table 6.2. Some magnetometers are upgraded, replaced, and adjusted, so these characteristics may change.

This chapter is organized as follows: the data structure used to store GNOME data is described in Sec. 6.1. Examples of data from the GNOME network is given in Sec. 6.2, while data describing an entire Science Run are presented in Sec. 6.3. The quality of a network is described in Sec. 6.4 along with a discussion of optimizing the network.

6.1 Data structure

In order to facilitate analysis, data from different sensors in the GNOME network is centrally stored. Each sensors will periodically synchronize collected data with a server located in Mainz, Germany. Servers in other locations can be used to create backups of the data and provide a lower latencies at distant locations.

The magnetometer data are stored with some basic standards, though these may be subject to change depending on needs determined by the GNOME collaboration. The magnetic field data and sanity channel data (designating the reliability of the data) are collected and stored at 512 Hz and 1 Hz sampling rates, respectively. These data are stored along with relevant metadata including date and time of data collection, sampling rate, units, and calibration information. Each file contains 1 minute of data and the files are stored in folders based on the station, year, month, and day of the form

```
.../station/yyyy/mm/dd/.
```

Each one-minute file contains roughly 100 kB of data.

Data from GNOME are stored using the HDF5 data structure. This enables an efficient and flexible means of storing relevant information. Notably, the magnetic field

Table 6.1: The different types of magnetometers used in GNOME and the corresponding factor for proton spin coupling. The factor is the ratio of the fractional spin polarization σ_p and the Landé g -factor. The value reported is used here (using Ref. [156]) while the errors include the range of values from other calculations [157–161]. Further information can be found in Ref. [162]. “Runs” refers to the science runs during which the magnetometer was active. Magnetometers that were changed (i.e., Hayward, Krakow, and Mainz) have multiple lines.

Station	Sensor Type	Probed transition	σ_p/g	Runs
Beijing	NMOR	^{133}Cs D2 F=4	$-0.39_{-0.00}^{+0.19}$	2,3,5
Belgrade	rf-driven	^{133}Cs D1 F=4	$-0.39_{-0.00}^{+0.19}$	5
Berkeley 1	NMOR	^{133}Cs D2 F=4	$-0.39_{-0.00}^{+0.19}$	1,3,4
Berkeley 2	NMOR	^{133}Cs D2 F=4	$-0.39_{-0.00}^{+0.19}$	1–4
Canberra	SERF	^{87}Rb D1	$0.70_{-0.15}^{+0.00}$	5
Daejeon	NMOR	^{133}Cs D2 F=4	$-0.39_{-0.00}^{+0.19}$	2–5
Fribourg	rf-driven	^{133}Cs D1 F=4	$-0.39_{-0.00}^{+0.19}$	1–3
Hayward	NMOR	^{87}Rb D1	$-0.36_{-0.00}^{+0.05}$	1–3
	SERF	^{87}Rb D1	$0.70_{-0.15}^{+0.00}$	4,5
Hefei	SERF	Rb/K D1	$-0.38_{-0.00}^{+0.05}$	2,3
Krakow	NMOR	^{87}Rb D1 F=2	$0.50_{-0.11}^{+0.00}$	1,2,4
	SERF	^{87}Rb D1	$0.70_{-0.15}^{+0.00}$	5
Lewisburg	SERF	^{87}Rb D2	$0.70_{-0.15}^{+0.00}$	2–5
Los Angeles	rf-driven	^{85}Rb D2 F=2	$0.50_{-0.07}^{+0.00}$	5
Mainz	NMOR	^{85}Rb D2 F=3	$0.50_{-0.11}^{+0.00}$	1–4
	SERF	^{87}Rb D1	$0.70_{-0.15}^{+0.00}$	5
Moxa	rf-driven	^{133}Cs D1 F=4	$-0.39_{-0.00}^{+0.19}$	4,5
Oberlin	SERF	K/Rb D1	$-0.38_{-0.00}^{+0.05}$	3–5
Stuttgart	rf-driven	^{85}Rb D2 F=3	$0.50_{-0.11}^{+0.00}$	

Table 6.2: The location and orientation of the GNOME magnetometers. The orientation is given in terms of the “altitude” angle with respect to the horizon and “azimuthal” angle with respect to North (clockwise). Some magnetometers were rotated at various times, so the current orientation may be different.

Station	Longitude	Location		Orientation	
		Latitude	Elevation	Alt	Az
Beijing	116.1868° E	40.2457° N	107.46 m	0°	251°
Belgrade	20.3909° E	44.8556° N	153.70 m	70°	295°
Berkeley 1	122.2572° W	37.8722° N	88.54 m	0°	28°
Berkeley 2	122.2570° W	37.8723° N	99.31 m	90°	
Canberra	149.1185° E	35.2745° S	593.42 m	90°	
Daejeon	127.3987° E	36.3909° N	71.14 m	90°	
Fribourg	7.1581° E	46.7930° N	244.84 m	0°	190°
Hayward	122.0539° W	37.6564° N	155.20 m	90°	
Hefei	117.2526° E	31.8429° N	110.43 m	0°	90°
Krakow	19.9048° E	50.0289° N	263.55 m	0°	45°
Lewisburg	76.8825° W	40.9557° N	126.94 m	90°	
Los Angeles	118.4407° W	34.0705° N	149.68 m	0°	270°
Mainz	8.2354° E	49.9915° N	193.02 m	-90°	
Oberlin	82.2204° W	41.2950° N	216.63 m	0°	300°
Stuttgart	9.1036° E	48.7443° N	525.37 m	48.7°	9.1°

and sanity data are stored as “datasets” and the metadata are stored as “attributes” in each file. The data were read into the analysis algorithm using a Python library developed for GNOME.

6.2 Sample data

To give a better idea on what GNOME data “looks” like, it helps to present some raw data. In this section, examples of data from the network will be described. Particular focus will be placed on Science Run 2 from December 2017 as this was the focus of much of the data analysis thus far.

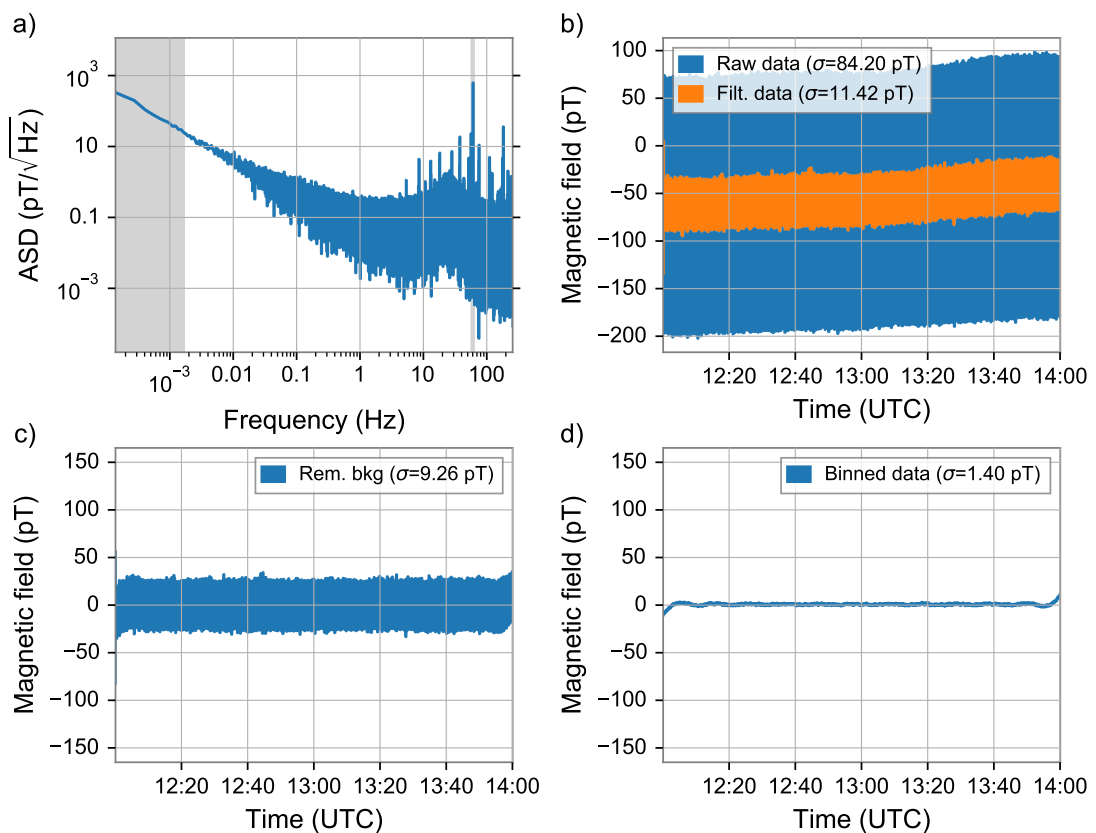


Figure 6.1: Example of filtering and averaging data. Two hours of data are shown from the Berkeley 2 magnetometer starting on 14 December 2017. (a) Amplitude spectral density of data with filtered frequencies grayed-out — the width of the line at 60 Hz does not reflect the exact size of the filter. (b) Raw data and data with a 60 Hz notch filter. (c) Background removed with a 1.67 mHz high-pass filter. (d) Data binned in 1 s bins.

First, the data from an individual magnetometer will be considered. Data from the

Berkeley 2 magnetometer is shown in Fig. 6.1. It should be noted that the data here were chosen to emphasize how filters can be used to reduce noise and does not necessarily reflect data typical in other sensors nor the same sensor at different times. For example, this magnetometer had a very strong powerline-frequency noise at this time. In any case, one sees a substantial reduction in the standard deviation after each subsequent filter is applied; dropping from 84 pT initially to 1.4 pT after filtering. However, because the data were sampled at 512 Hz and the averaging was over 1 s, one would expect the noise to be reduced by a factor of $\sqrt{512} \approx 22.6$ from this step. However, the reduction is only by a factor of 6.6. This suggests that there is some non-Gaussianity characteristics of the noise remaining, despite attempts to remove these with filters. Note also that the filters can lead to some artifacts on the edges of the data, but, in practice, data on the edge of a filtered subset is not used in the search for events.




Figure 6.2: Data collected by all the GNOME sensors over Science Run 2 in 2017. Data are averaged across 30 minute bins to emphasize long-term drifts. Only available, reliable data are shown.

It is also important to consider what GNOME data look like on long-term time scales. This can be useful in understanding how the background drifts over time. Figure 6.2 shows several weeks of data with 30 minute averaging for different GNOME sensors.

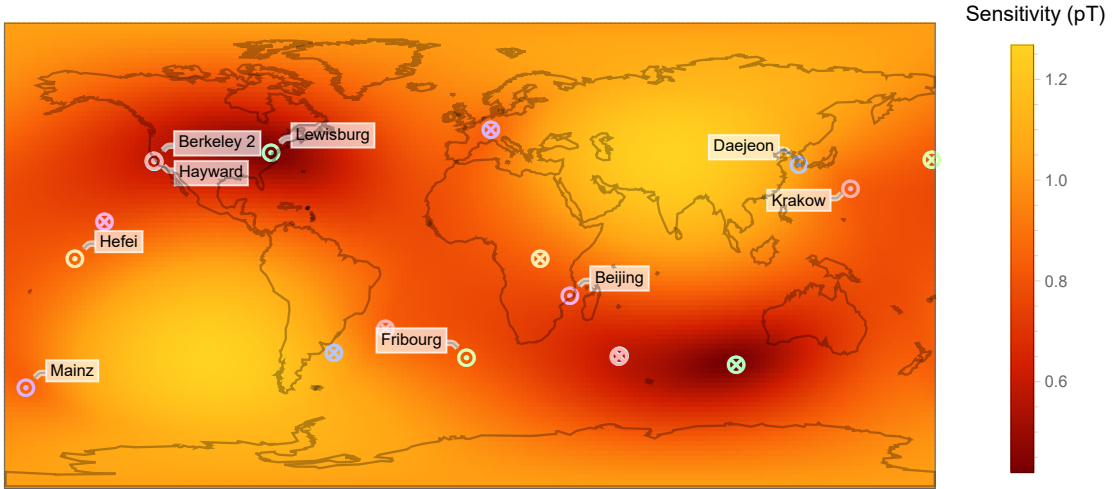


Figure 6.3: Average directional sensitivity over Science Run 2. Position on the map corresponds to the direction from the center of the Earth to that point on the map. Sensitivity is for proton coupling, and data are averaged into 20 s bins with a 1.67 MHz high-pass filter and notch filters at powerline frequency. Average over time is calculated as $\langle \beta^{-1} \rangle^{-1}$ with $\beta^{-1} = 0$ when fewer than four sensors are active. The forward \odot and backward \otimes directions of the various sensors are marked with labels on the forward directions; the two points corresponding to a given sensor have the same color.

One can see, for example, that some of the sensors have daily fluctuations. Though, it should be noted that analysis is generally designed on timescales much shorter than a day to avoid the effects of Earth’s rotation. In any case, this can point to other means of improving sensor stability. The drifts can be removed with high-pass filter at the cost of reduced sensitivity to longer signals.

6.3 Run characteristics

Over the course of a science run, the sensitivity of the network will inevitably vary as magnetometers join the network or become inactive. If the network has good directional coverage and is fairly stable, one would expect the sensitivity to have little variation in different directions. The average directional sensitivity over time for Science Run 2 is shown in Fig. 6.3. This type of map can motivate future network improvements by pointing towards directions that could be better-covered; by adding sensors, rotating existing sensors, or improving existing sensors currently pointing in the given direction.

Just as sensitivity varies with direction, the network sensitivity also varies over time. This is due both changes in an individual sensor and sensors becoming active and inactive; including due to data being flagged as unreliable by auxiliary measurements. The sensitivity over time, as well as the number of active magnetometers, is shown in Fig. 6.4.




Figure 6.4: (a) The number of active magnetometers and (b) sensitivity over time using proton coupling. The solid lines are one-day rolling average. The data are filtered with a 20 s averaging, 1.67 MHz high-pass filter, and notch filters at powerline frequency.

Though not always, one generally sees that when more magnetometers are active, the network becomes sensitive to weaker signals.

6.4 Quality of the network

The network used in the GNOME experiment may be improved without needing additional sensors or improvements in the existing sensors. This can be accomplished simply by reorienting the sensitive axes of the magnetometers. The “quality” of the network will be referred to here as how close the arrangement of the network is to the optimal arrangement; not the absolute sensitivity of the network, so equally improving all sensors would not change the network quality. A quantitative measure of the network quality is presented here.

The orientation of sensors in an optimized network depends on how the network sensitivity is defined (see Chapter 5.1). Two definitions of network sensitivity are described in this work: as an average over a known distribution of directions (Eq. (5.2)) and as the worst-case-scenario (Eq. (5.3)). The latter and more conservative definition for network sensitivity is used. The worst sensitivity $\bar{\beta}_0$ is determined by the smallest eigenvalue of $(D^T \Sigma_s^{-1} D)$ while the best sensitivity $\bar{\beta}_1$ is given by the largest eigenvalue — specifically, $\bar{\beta} = \lambda^{-1/2}$ where λ is an eigenvalue. The respective eigenvectors denote the worst and best directions for a signal.

If the network quality is optimized, the sensitivity $\beta(\hat{\mathbf{m}})$ would not depend on the direction $\hat{\mathbf{m}}$. This is because the existence of some preferred direction implies that one could improved sensitivity in a less-preferred direction by rotating the sensitive axis of a magnetometer away from the preferred direction and towards the less-preferred

direction. In practice, it is not possible for GNOME to maintain optimal quality under this definition, because the noise of the magnetometers varies over time along with what sensors are active.

To better assess the quality of the network, it helps to define some quantitative “quality factor.” This factor would ideally reflect only how optimally the sensitive axes are oriented regardless of the network sensitivity. One possibility is to define the quality factor as the quotient of the best and worst sensitivity,

$$q_0 := \bar{\beta}_1 / \bar{\beta}_0. \quad (6.1)$$

This factor will lie in the range $(0, 1]$ with $q_0 = 1$ being ideal.

There are a few characteristics of an ideal network that are useful to consider. First, for an ideal network, $(D^T \Sigma_s^{-1} D) = \bar{\beta}^{-2} \mathbb{1}$ is proportional to the identity matrix. Consider two uncorrelated ideal networks: A and B . If the two networks are combined,

$$\begin{bmatrix} D_A^T & D_B^T \end{bmatrix} \begin{bmatrix} \Sigma_A & 0 \\ 0 & \Sigma_B \end{bmatrix}^{-1} \begin{bmatrix} D_A \\ D_B \end{bmatrix} = (\bar{\beta}_A^{-2} + \bar{\beta}_B^{-2}) \mathbb{1}.$$

Thus, the combined network is also ideal with sensitivity $(\bar{\beta}_A^{-2} + \bar{\beta}_B^{-2})^{-1/2}$. Observe that the relative orientation of the two sub-networks in the combined network is irrelevant. Further, this reflects the lack of uniqueness in an ideal network. The set of ideal networks is invariant under any global rotations of independent ideal sub-networks as well as parity reflection of any independent sets of sensors (i.e., reversing the sensitive axis of an independent sensor).

6.4.1 Ideal network

Some heuristic approximation can be made for the sensitivity of an ideal network. Consider a network of n uncorrelated magnetometers with coupling $\{\kappa_j\}$ ¹, standard deviation (i.e., noise) $\{\sigma_j\}$, and sensitive axis in the direction $\{\hat{\mathbf{d}}_j\}$. Thus, the j^{th} sensor will observe a signal with amplitude $s_j = \kappa_j \hat{\mathbf{d}}_j \cdot \mathbf{m}$. If the angle θ_j between the sensitive axis $\hat{\mathbf{d}}_j$ and the field direction \mathbf{m} is random for each sensor, so $(\hat{\mathbf{d}}_j \cdot \hat{\mathbf{m}})^2 \approx \langle \cos^2 \theta \rangle = 1/d$ for $d = 3$ spatial dimensions (this calculation is given towards the end of Appendix A.3.3), then the optimal sensitivity is

$$\bar{\beta}_{\text{opt}} \approx \sqrt{d / \sum_{j=0}^{n-1} (\kappa_j^2 / \sigma_j^2)}. \quad (6.2)$$

Note that this may not be the exact value for an optimized network of uncorrelated sensors due to the assumptions made in the calculation. In particular, it is assumed that

¹Here, $\kappa_j = \frac{g_{F,j}}{\sigma_j}$ where σ_j is the projected spin coupling and $g_{F,j}$ is the g -factor for the j^{th} magnetometer. Note that elsewhere in this section σ_j is used to refer to the standard deviation in the j^{th} magnetometer data.

the direction between the sensitive axis $\hat{\mathbf{d}}_j$ and the field direction \mathbf{m} is independently random for each sensor, though correlations may exist. Deviation from Eq. (6.2) and the actual optimal sensitivity is likely to be small in a network consisting of many, similar sensors. In this case, the optimal network would have sensitive axes evenly distributed in all directions, making the set of dot products $\hat{\mathbf{d}}_j \cdot \hat{\mathbf{m}}$ describe a similar distribution as if the angle between the signal vector and sensitive axis was random for each sensor.

It is possible to determine an ideal network configuration under certain simplified conditions. In particular, for independent, identical sensors, there are a few configurations that optimize the network. The simplest of which would be a set of three sensors with orthogonal sensitive axes. More optimal networks can be found by taking inspiration from Platonic solids. In particular, the vertices of Platonic solids are evenly distributed in all directions relative to the center of the polyhedron. One can construct a network configuration by orienting the sensitive axis of a sensor from the center of the Platonic solid to each of its vertices. One can verify explicitly that each of these arrangements is an optimal network with sensitivity $\beta = \sigma\sqrt{3}/n$ for n vertices/sensors in the network. Each of these networks is given in Table 6.3.

The optimal networks defined by the Platonic solids can be further reduced into optimal sub-networks. For example, the octahedral network consists of two sub-networks of three orthogonal sensors. Due to the parity symmetry in the octahedron, cube, icosahedron, and dodecahedron, one can separate each such network into two optimized networks whose sensitive axes are anti-parallel. Together, this defines optimal configurations of three, four, six, and ten sensors; denoted by N_3 , N_4 , N_6 , and N_{10} . One can also show that none of these networks can be described as the combination of sub-networks with the other configurations (e.g., N_6 cannot be reduced into two N_3 networks). That is, these networks are irreducible.

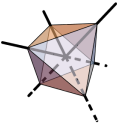
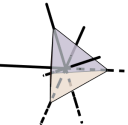
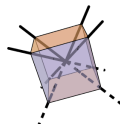
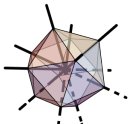
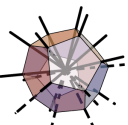
The arrangements of sensors described above can be combined into any optimal network with three, four, six, or more identical, uncorrelated magnetometers. Further, there is some liberty in how these networks are constructed. For example, a network with $n = 6$ sensors could be arranged as N_6 or two N_3 sub-networks.

Though not based on a Platonic solid, one can explicitly construct an (irreducible) optimal network with five identical, independent sensors. In fact, two such arrangements are given by

$$D_{5a} = \begin{bmatrix} 0 & \sqrt{\frac{5}{6}} & \sqrt{\frac{1}{6}} \\ 0 & -\sqrt{\frac{5}{6}} & \sqrt{\frac{1}{6}} \\ \sqrt{\frac{5}{6}} & 0 & \sqrt{\frac{1}{6}} \\ -\sqrt{\frac{5}{6}} & 0 & \sqrt{\frac{1}{6}} \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad D_{5b} = \begin{bmatrix} 1 & 0 & 0 \\ -\sqrt{\frac{1}{3}} & -\sqrt{\frac{2}{3}} & 0 \\ -\sqrt{\frac{1}{3}} & \sqrt{\frac{2}{3}} & 0 \\ 0 & \sqrt{\frac{1}{6}} & \sqrt{\frac{5}{6}} \\ 0 & -\sqrt{\frac{1}{6}} & \sqrt{\frac{5}{6}} \end{bmatrix}.$$

These networks are unique with respect to parity reversal of individual sensitive axes or global rotations; this is evident because the latter has orthogonal sensitive axes while the former does not. These networks are also irreducible because an optimal network

Table 6.3: Examples of ideal networks based on the Platonic solids. The orientations of sensitive axes in an ideal network are given as lines (dashed in one direction, solid in the other). A sum of vertices is used to describe multiple ideal subnetworks. For example, the cube has eight vertices, and the corresponding ideal network consists of two ideal networks with four sensors each; hence “4+4” is listed. In each case here, $X + X$ vertices describes two ideal networks with X sensitive axes in opposite directions.

Name	Vertices	Shape
Octahedron	3 + 3	
Tetrahedron	4	
Cube	4 + 4	
Icosahedron	6 + 6	
Dodecahedron	10 + 10	

needs at least three sensors and one cannot separate five sensors into two or more ideal subnetworks with at least three sensors. From these examples, it is evident that optimal networks are not necessarily unique, even if they are irreducible.

Though, as shown above, there exists a network of at least three independent, identical sensors exists with $q_0 = 1$, it is not necessarily the case if the sensors are not identical. For example, consider $n \geq 2$ independent, identical sensors with noise σ_0 and another, better sensor with noise $\sigma_1 < \sigma_0 \sqrt{2/n}$. Choosing some direction to orient the better sensor and orienting the remaining sensitive axes to evenly cover orthogonal directions, the most-sensitive direction will be along the sensitive axis of the better sensor. Thus, it is not possible to achieve $q_0 = 1$ in this hypothetical network, because there will always be a preferred direction. However, it is still possible to optimize the network sensitivity.

The process of optimizing a given, arbitrary network can be accomplished in various ways. For example, a “greedy” algorithm may operate by adding sensors such that their

sensitive axis points in the least-sensitive direction for the rest of the network. Such an algorithm could be iterated on a complete network by, in each step, removing some sensor and adding it back in the least-sensitive direction. Likewise, if $\beta(\hat{\mathbf{x}})$ is the sensitivity function for a network, one can slightly rotate one/several sensitive axes as to locally “ascend” $\beta(\hat{\mathbf{x}})$; i.e., rotate towards less-sensitive directions. The exact amount to rotate the sensors in each step can be chosen from any one of numerous optimization methods; e.g., annealing or gradient descent.

6.4.2 Quality of GNOME

Unlike the ideal networks discussed thusfar, the noise in the GNOME magnetometers varies over time, and stations will become active or inactive over time. Even if the network was optimized at some time, it will become suboptimal as the characteristics of GNOME change. The average network conditions can be used to understand the best and worst direction; see Fig. 6.3. If one wants to add a new sensor to the network, orienting its sensitive axis in the least-sensitive direction according to this average would be ideal.




Figure 6.5: The “quality” of the GNOME network over time. Solid lines are the one-day rolling average. Proton couplings are used, and data are filtered with 20s averaging, 1.67 MHz high-pass filter, and notch filters to remove powerline frequency. (a) The number of active sensors. (b) The quality q_0 of the network. (c) The factor by which the sensitivity could be improved relative to the theoretical sensitivity, Eq. (6.2).

The quality of GNOME over the various Science Runs is given in Fig. 6.5 along with the factor by which the sensitivity could be improved relative to the approximate optimum, Eq. (6.2). In contrast to the sensitivity, the quality factor (assuming at least three sensors are active) does not have a clear dependence on the number of active sensors; the quality decreases significantly towards the end of Science Run 3 because there was often fewer than three active sensors. Additionally, one can see that Science Run 2 often had a nearly-optimal network.

The GNOME network can be algorithmically optimized to better understand the network's potential. For example, one can consider a network consisting of the nine magnetometers from Science Run 2 whose noise is their respective average noise during the run after adjusting for proton coupling. Without altering the orientations of any magnetometer, the network has a worst-sensitivity of $\bar{\beta} = 1.0$ pT and a best-sensitivity of 0.43 pT ($q_0 = 0.4$). For an optimized network, $q_0 = 0.7$ and $\beta = 0.70$ pT. The sensitivity of this network is improved by a factor of 1.5 when optimizing the network. However, it should be noted that this optimization assumes all magnetometers are active which is not often the case in usual operation.

Chapter 7

Results

Through analyzing the data from GNOME, regions of ALP parameter space have been probed for signals arising from domain walls. In the absence of evidence, one can exclude the probed region of parameter space at a given confidence level. In particular, Science Run 2 — having taken place during December 2017 — is studied. Though Science Runs 3 and 4 contain more data, Science Run 2 had excellent and fairly consistent sensitivity.

The data were analyzed using methods described in Chapter 4. First, a false-negative analysis was conducted to determine appropriate thresholds to guarantee some detection efficiency. Then, a false-positive analysis was combined with the observed event rate to determine whether domain walls were detected and to place bounds on ALP parameter space.

For the analysis discussed here, data from Science Run 2 were studied with a particular set of filters. Zero-phase notch filters (quality factor $Q = 60$) were applied to the data at the local powerline frequency; e.g., sensors in North America had 60 Hz noise removed while sensors in Europe had 50 Hz noise removed. Likewise a zero-phase 1.67 mHz first-order high-pass Butterworth filter is applied to data from every sensor to remove drifts and noise on timescales over 10 min. In both cases, “zero-phase” is accomplished by applying the filter both forward and backward. The data were also averaged into 20 s bins, which both improves sensitivity to signals longer than the averaging time and reduces the computational requirements.

In Sec. 7.1, the methods described in Chapter 4.6 are applied to GNOME data. The results of the full GNOME analysis, namely the event rate in terms of magnitude-to-uncertainty, are presented in Sec. 7.2. The event rate did not significantly exceed what one would expect from background alone, so an excluded region of parameter space is presented in Sec. 7.3. The excluded region is presented using the “extended” method and compared to the region generated with the “loudest-event” method. Results using the loudest-event method are also presented in Ref. [4].

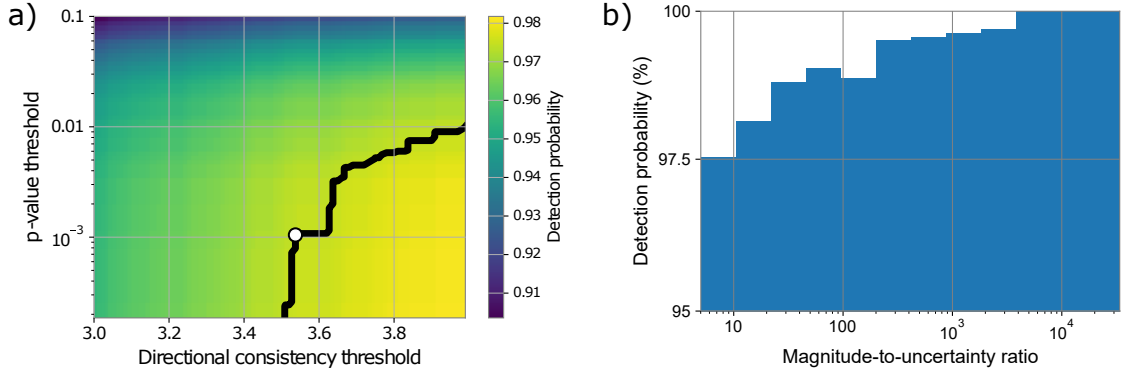


Figure 7.1: Probability of detecting an expected signal inserted into background data. (a) Considers different p -value and directional consistency thresholds. The black line corresponds to 97.5% detection efficiency and the white dot reflects which threshold was used in the final study. (b) Relates the detection probability to magnitude-to-uncertainty ratio. This figure is adapted from one found in Ref. [4].

7.1 Selecting thresholds

To begin, thresholds must be determined to optimize the efficiency by which events are detected while minimizing the background of false-positive events. For this analysis, a goal is set to detect and identify $\epsilon = 97.5\%$ of signals from domain-wall-crossing events. This is accomplished by inserting signals into background noise and testing the algorithm's ability to detect them as part of a false-negative analysis.

The characteristics of the signals inserted into the data are chosen to cover the range of signals that we hope to detect. The background data was generated from measurements by time-shuffling the data. Thus, the domain-wall-crossing signals were inserted at random times and directions with magnitudes \mathcal{B}_p between 0.1 pT and 10^3 pT and widths Δt between 0.01 s and 10^3 s with a flat PDF on a logarithmic scale.

The results of the false-negative analysis can be seen in Fig. 7.1. In Fig. 7.1a, the probability of detection ϵ is shown as a function of the p -value threshold and directional consistency threshold (i.e., the deviation between the measured signal gradient and scanned velocity in units of angular spacing). The p -value threshold is an upper-bound while the directional consistency threshold is a lower-bound for rejection. A black line highlights the combinations of these thresholds that yields $\epsilon = 97.5\%$. The white dot on this line is the choice of thresholds used in the final analysis. This choice was used to reduce the background event rate, as calculated in the false-positive analysis.

The detection probability can vary based on the magnitude of the signal, as well (Fig. 7.1b). In order to exclude parameter space with either the loudest-event or extended method, one must guarantee some minimal detection efficiency for all events at least as large as some magnitude-to-uncertainty ratio. In the loudest-event method, this is the detection efficiency for the loudest event, while in the extended method, this is the efficiency for any given magnitude-to-uncertainty ratio. In this analysis, one finds that

the detection efficiency is at least 97.5% for magnitude-to-uncertainty ratios above 5.

7.2 Event rate

Comparing the event rate observed in the experiment to that which one expects from background is crucial for constraining physical parameter space. An excess of events over the background event rate would suggest the existence of ALP domain walls — or at least the existence of some class of large structures with exotic spin coupling. The absence of a significant excess could likewise be used to exclude a region of physical parameter space for a given theory.

To determine what constitutes a “significant excess” in the event rate, a false-positive analysis is performed on background data. As discussed in Chapter 4, background data are generated by sampling data from each sensor in the network at random times through “time-shuffling.” One can generate virtually unlimited background data using this method. In the end, the amount of background simulated is determined by computational resources and the diminishing benefit of simulating additional background. Ideally, there would be significantly more simulated background data than the amount of data used in the experiment. For results discussed here, the data were taken from Science Run 2 (a 23-day run) and represent 10.7 years of time-shuffled data.

The event rate for GNOME Science Run 2 is shown as a function of minimum magnitude-to-uncertainty ratio in Fig. 7.2 for the time-shuffled background data and the measurement data. At no point does the 90%-confidence lower-bound on the excess event rate exceed zero — yet alone at the 5σ confidence typically needed for discovery. In fact, the most significant excess occurred for the loudest event with only 42%-confidence that the lower-bound was positive. Thus, we can conclude that evidence for ALP domain walls was not found in this analysis of GNOME’s Science Run 2 data. The upper-bound on the excess event rate can be used with the extended method to generate a bound on physical parameter space.

7.3 Exclusion region

Given that the analysis of GNOME data did not find evidence of ALP domain walls, bounds on physical parameter space can be established. This will be based on the sensitivity of the network to different signals over the course of the experiment. This will define a region of parameter space that should yield positive results in the experiment. In the absence of positive results, sets of parameters in this region can be excluded.

First, one can consider the types of signals that one would expect from the perspective of the experiment. The GNOME sensors do not directly measure ALP parameters like mass and coupling constants; instead they measure the pseudo-magnetic field induced by exotic spin coupling of an ALP feature. For example, if one expects to observe Lorentzian signals in the network, one can consider the widths and magnitudes that can be observed. The experimental sensitivity over time was already discussed (see Chapter 6.3 and Fig. 6.4), but these did not also consider that filters will affect signals

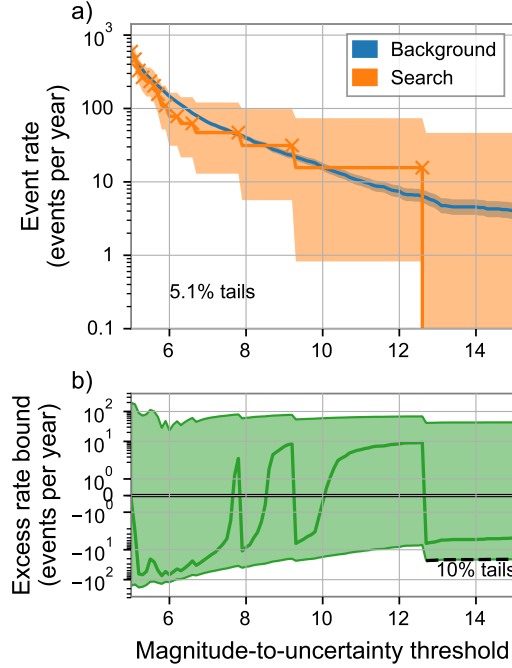


Figure 7.2: Event rates from the analysis of Science Run 2. (a) The background and search data event rate with $1 - \sqrt{90\%} \approx 5\%$ tails; i.e., the upper-bound and lower-bound (when defined) are at $\approx 95\%$, and the confidence that the rate is between the bounds is $\approx 90\%$. Markers on the search-data event rate emphasize individual events. (b) The excess event rate with 10% tails. The dashed black line indicates that no events were found in the measurement data, so the search data event rate is zero and the lower-tail is $\approx 5\%$.

differently depending on their width. Considering the filters used on the various magnetometer data sets with the sensitivity and integrating over time yields Fig. 7.3. A “U”-shaped region of sensitivity appears on this plot. The edge of this “U” is blurred by changes in sensitivity over time. Roughly speaking, the left side can be understood as resulting from thin signals being averaged-out, and the right side is the result of long signals being removed by the high-pass filter. If no filters were applied, the plot would appear like a horizontal line separating a sensitive and insensitive region.

At this point, all the information needed to generate an exclusion plot; either using the loudest-event method, Eq. (5.11), or the extended method, Eq. (5.12). In either case, Fig. 7.3 is the function for \bar{T} in terms of variables directly observable through the experiment; these are related to theoretical parameters via Eq. (5.13). Also, the false-negative analysis was used to guarantee the detection efficiency $\epsilon = 97.5\%$ or better.

The expected speed is taken to be the Earth’s galactic rotation speed, $\bar{v} = 10^{-3}c$, and the energy density of domain walls is taken to be the dark matter density, $\rho_{\text{DW}} = \rho_{\text{DM}} \approx 4 \times 10^5 \text{ GeV/m}^3$. It is worth remarking that equating the domain-wall and dark-matter

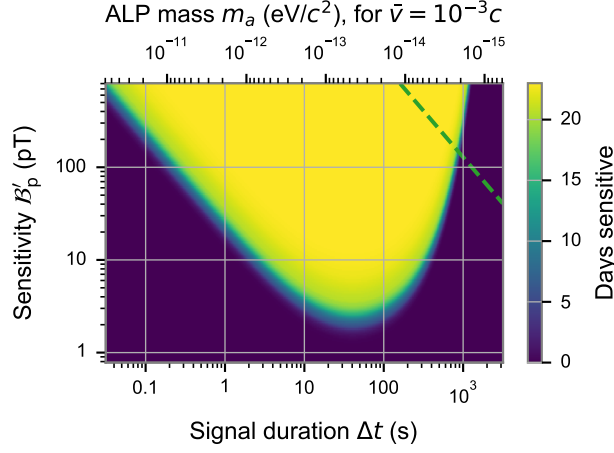


Figure 7.3: Time during which GNOME is sensitive to different signal widths and magnitudes. Here, one assumes that the signals are Lorentzian with FWHM of Δt . The sensitivity \mathcal{B}'_p is defined as the magnitude (i.e., height of the Lorentzian) needed to induce a magnitude-to-uncertainty ratio of one. The dashed green line represents the part of the physical parameter space relevant when constraining $f_{\text{int}} = f_a$.

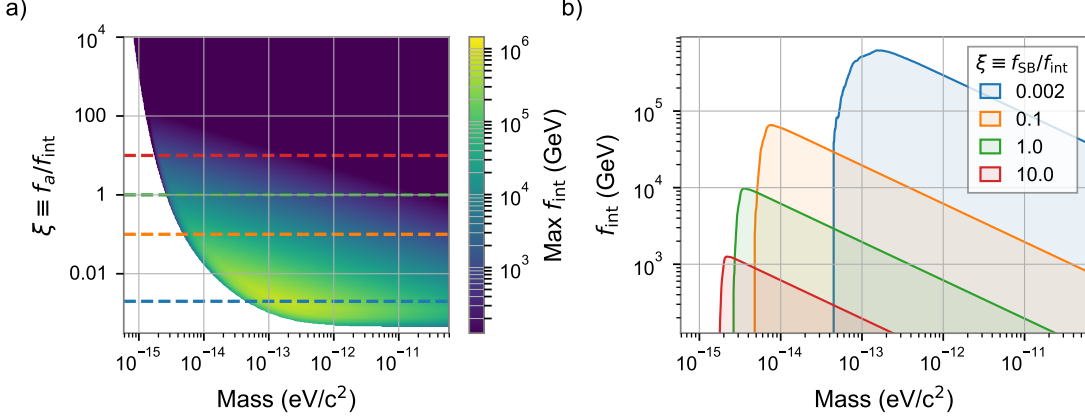


Figure 7.4: Region of physical parameter space excluded by GNOME Science Run 2 using the extended method with a 90% confidence level. (a) The upper-bound on excluded interaction scale f_{int} as a function of ALP mass and the ratio ξ . The colored regions in (b) are marked by the respective dashed lines. (b) The excluded interaction scale f_{int} as a function of ALP mass for different ratios ξ . The ratio $\xi = 0.002$ corresponds to the cross-section with the highest bound on coupling; excluding $f_{\text{int}} < 6 \times 10^5 \text{ GeV}$ at $m_a = 1.5 \times 10^{-13} \text{ eV}/c^2$.

energy densities is not a trivial one. It is entirely possible that domain walls constitute only a fraction of the energy density attributed to dark matter — or even that domain walls do not exist at present time. The dark-matter energy density provides a bound on ρ_{DW} as the amount of energy not yet accounted for in the galaxy. If the actual energy density is much smaller, the frequency of crossing events would also be smaller, which would weaken the constraints presented here.

For the loudest-event method, the magnitude-to-uncertainty ratio ζ_0 can be read off of Fig. 7.2 as the ratio of the largest event, $\zeta_0 = 12.6$. For the extended method, Fig. 7.2b includes the bound $R_{\tilde{c}}$ as a function of the lower-bound on the ratio ζ_0 . Each value of ζ_0 can be used to generate an exclusion plot — the union of the excluded regions is the final result. One only needs to consider the bound on the excess event rate when the rate of measurement events decreases.

The region of parameter space excluded by this analysis of Science Run 2 using the extended method is shown in Fig. 7.4. The largest exclusion for the coupling occurs for $m_a = 1.5 \times 10^{-13} \text{ eV}/c^2$ and $\xi = 0.002$ where $f_{\text{int}} < 6 \times 10^5 \text{ GeV}$ is excluded. Likewise setting $\xi = 1$ (or $f_{\text{int}} = f_a$), the coupling $f_{\text{int}} < 10^4 \text{ GeV}$ is excluded for $m_a = 3.6 \times 10^{-15} \text{ eV}/c^2$.

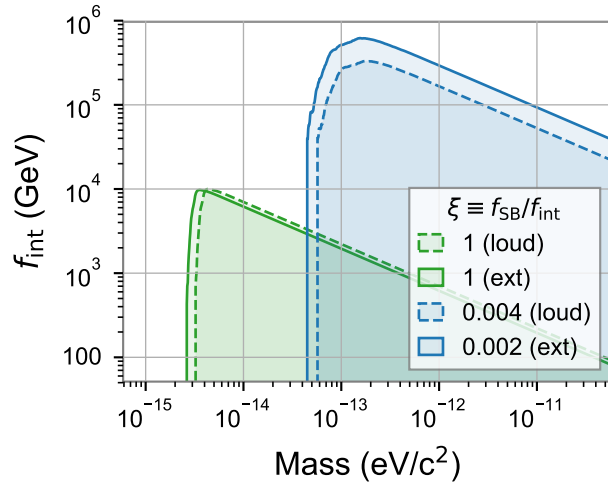


Figure 7.5: Comparison between the excluded region using the loudest-event (dashed edge) and extended (solid edge) methods setting $\xi = 1$ (green) and to the value that maximizes the largest upper bound for f_{int} (blue).

The excluded region from the extended method is compared to the loudest-event method in Fig. 7.5. Recall that the magnitude of the domain-wall signal is proportional to $m_a \xi$ and the extended method aims to include weaker signals. Thus, as expected, the extended method covers smaller masses than the loudest-event method. For the same ratio $\xi = 1$, the loudest-event method has a slightly better bound on f_{int} at larger masses. This likely the result of being “lucky” with how small the largest event happened to be,

and one would expect this relative improvement to diminish with more data. When considering the value of ξ that maximizes the respective bounds on f_{int} , the region excluded by the loudest-event method is entirely contained in the region excluded by the extended method. As with mass, the extended method excludes regions of parameter space that have smaller ratios ξ — this is why the extended method is optimized when this ratio is half the ratio that optimizes the loudest-event method bound. The smaller ratio ξ also corresponds to a theory with a weaker coupling between ALPs and fermions relative to the axion decay constant.

Chapter 8

Summary and outlook

This work applies a theoretical framework of dark matter to understand data gathered by experiment. Particular focus is placed on domain walls in an ALP field that couples to fermions via their spin, but the analysis is set up in such a way to allow for similar methods to be applied to other searches. This can be understood generically as following the following steps of experimental design:

0. Decide on what type of dark matter to detect and what type of experiment to design. Some of such models and experiments are described in Chapter 2.
1. Understand what type of signal is expected: the coupling type, pattern, shape, etc. For this work, this is described in Chapter 3.
2. Devise a method for testing the consistency between experimental data and expected observation — including consideration for things like filtering and expected background events. This step enables one to make a statement about whether something was detected. These methods are described in Chapter 4.
3. Relate what is directly observable in the experiment to the parameter space for the theoretical unknowns. Here, this involved relating what signal widths/magnitudes can be reliably observed/expected to the ALP mass m_a and the energy scales f_a and f_{int} . These relations are described in Chapter 5.

These steps can then be applied to data. A review of the GNOME data is given in Chapter 6 with the results of analyzing the data given in Chapter 7. The observed data was found to be consistent with what one would expect from background data alone; meaning that strong evidence of ALP domain walls could not be confirmed.

Given that the GNOME experiment could not show the existence of ALP domain walls during Science Run 2, a region of physical parameter space can be excluded at some confidence level. The excluded region at 90% confidence is presented in Fig. 7.4. The set of parameters within the excluded region would have resulted in the detection of ALP domain walls, while the parameters not excluded could have been missed by the experiment. The strictest bound on the coupling scale from this analysis is $f_{\text{int}} >$

6×10^6 GeV for $\xi = 0.002$ and $m_a = 10^{-13}$ eV/ c^2 . For $\xi = 1$, the coupling is bounded by $f_{\text{int}} > 10^4$ GeV when $m_a = 4 \times 10^{-15}$ eV/ c^2 .

These constraints fall short of the conservative astrophysical bound on the axion from SN1987A [52, 54] of 10^8 GeV. However, this bound is highly model-dependent, assuming a coupling based on the KSVZ and DFSZ axion models. The work presented here is comparatively model-independent, with the coupling strength allowed to vary freely for any axion or ALP model. The value of the ratio is typically $\xi \approx \mathcal{O}(0.1)$ [56, 57, 163] based on calculations from the underlying model¹.

Further, the results presented here reflect a truly novel approach to a dark matter search. This is the first full result from the GNOME network. There are some examples of networks being used to search for ALP domain walls (e.g., GPS.DM [124] and gravimeter [153]) using mass-like coupling. However, GNOME is used here to search for spin-coupling.

Continued operation of GNOME will provide new and improved bounds on dark matter. The bounds presented here can be improved by taking more data, improving the sensors (both the sensitivity and reliability), or simply through new analysis. For example, one can apply a different set of filters that suppress/emphasize different types of signals in the same data.

Because evidence of ALP domain walls were not found, most of the focus was on describing excluded bounds. However, if a future study were to find evidence of an excess event rate in the search data as compared to the expected background rate, one could use this information to generate a bound on physical parameters. First, similar to how an upper-bound on the event rate was used to generate the excluded region (see Eq. 5.12), a band of possible ranges can be generated using the same methods described in Appendix C. This range translates to a band of parameter space. The magnitude-to-uncertainty ratio of the excess can be used to restrict $\mathcal{B}_p \propto m_a \xi$, and the width of the observed signals restricts the mass. Taken together, one can then constrain the parameters $\{m_a, f_{\text{int}}, \xi\}$ to some volume. However, one may also find that the shape of the observed signals is not consistent with that of a domain wall; meaning that the excess event rate could be the result of other objects with exotic spin coupling. Such a discovery would nevertheless be exciting and point to new physics.

The analysis described in this work was intended to search for ALP domain walls, but the same analysis would also yield results for other objects. In particular, if an object with an exotic spin coupling crosses the network and is large enough to appear planar on the scale of the Earth, one may observe an excess in event rates. As an example, there has been some interest in using GNOME to detect boson stars [6, 98], especially in the “thin-wall” regime. If there are boson stars much larger than the Earth, its edge would appear similar to a plane to the network. Likewise, the analysis described here would apply to different signal shapes caused by the same object but with different spin-coupling.

The GNOME collaboration continues to grow and adopt new magnetometer tech-

¹Other references (e.g., Refs. [52–54, 56, 57, 163]) will use the dimensionless constant C_i to describe the coupling strength to particle i that can be related to the constant used here with $C/2 = \xi$.

nologies to improve sensitivity. A set of planned upgrades referred to as “Advanced GNOME” aims to improve sensitivity by adding SERF comagnetometers to the network [149, 164, 165]. This upgrade could also include careful consideration of how the sensors are oriented to optimize sensitivity; as per Chapter 6.4.

Dark matter searches have played a major role in modern theoretical and experimental physics. These efforts have led to new models of physics as well as expansions on existing models to attempt to explain this unknown phenomenon. New experiments have been designed and old experiments have been retooled in the effort to probe this frontier. Despite its influence on galactic scales, dark matter continues to evade our best efforts to understand the basic operations of the universe. Hopefully, as experimental and theoretical methods advance, the true nature of this mysterious matter will be revealed.

Bibliography

- [1] Pospelov, M., Pustelny, S., Ledbetter, M. P. *et al.* Detecting Domain Walls of Axionlike Models Using Terrestrial Experiments. *Phys. Rev. Lett.* **110**, 021803 (2013).
- [2] Afach, S., Budker, D., DeCamp, G. *et al.* Characterization of the global network of optical magnetometers to search for exotic physics (GNOME). *Physics of the Dark Universe* **22**, 162–180 (2018).
- [3] Masia-Roig, H., Smiga, J. A., Budker, D. *et al.* Analysis method for detecting topological defect dark matter with a global magnetometer network. *Physics of the Dark Universe* **28**, 100494 (2020).
- [4] Afach, S., Buchler, B. C., Budker, D. *et al.* Search for topological defect dark matter using the global network of optical magnetometers for exotic physics searches (GNOME). *arXiv:2102.13379 [astro-ph, physics:hep-ex, physics:physics]* (2021).
- [5] Pustelny, S., Jackson Kimball, D. F., Pankow, C. *et al.* The Global Network of Optical Magnetometers for Exotic physics (GNOME): A novel scheme to search for physics beyond the Standard Model. *Annalen der Physik* **525**, 659–670 (2013).
- [6] Jackson Kimball, D. F., Budker, D., Eby, J. *et al.* Searching for axion stars and Q-balls with a terrestrial magnetometer network. *Phys. Rev. D* **97**, 043002 (2018).
- [7] Dailey, C., Bradley, C., Jackson Kimball, D. F. *et al.* Quantum sensor networks as exotic field telescopes for multi-messenger astronomy. *Nat Astron* **5**, 150–158 (2021).
- [8] Gorenstein, P. & Tucker, W. Astronomical Signatures of Dark Matter. *Advances in High Energy Physics* **2014**, 1–10 (2014).
- [9] Rubin, V. C., Burstein, D., Ford, W. K., Jr. & Thonnard, N. Rotation velocities of 16 SA galaxies and a comparison of Sa, Sb, and SC rotation properties. *ApJ* **289**, 81 (1985).
- [10] Corbelli, E. & Salucci, P. The extended rotation curve and the dark matter halo of M33. *Monthly Notices of the Royal Astronomical Society* **311**, 441–447 (2000).

-
- [11] Sofue, Y. & Rubin, V. Rotation Curves of Spiral Galaxies. *Annu. Rev. Astron. Astrophys.* **39**, 137–174 (2001).
- [12] Newman, A. B., Treu, T., Ellis, R. S. & Sand, D. J. The density profiles of massive, relaxed galaxy clusters. II. Separating luminous and dark matter in cluster cores. *ApJ* **765**, 25 (2013).
- [13] Fabricant, D. & Gorenstein, P. Further evidence for M87’s massive, dark halo. *ApJ* **267**, 535 (1983).
- [14] Mathews, W. G. & Brighenti, F. Hot Gas in and Around Elliptical Galaxies. *Annu. Rev. Astron. Astrophys.* **41**, 191–239 (2003).
- [15] Schmidt, R. W. & Allen, S. W. The dark matter haloes of massive, relaxed galaxy clusters observed with Chandra. *Monthly Notices of the Royal Astronomical Society* **379**, 209–221 (2007).
- [16] Markevitch, M., Gonzalez, A. H., David, L. *et al.* A Textbook Example of a Bow Shock in the Merging Galaxy Cluster 1E 0657-56. *ApJ* **567**, L27–L31 (2002).
- [17] Clowe, D., Gonzalez, A. & Markevitch, M. Weak-Lensing Mass Reconstruction of the Interacting Cluster 1E 0657-558: Direct Evidence for the Existence of Dark Matter. *ApJ* **604**, 596–603 (2004).
- [18] Tam, S.-I., Jauzac, M., Massey, R. *et al.* The distribution of dark matter and gas spanning 6 Mpc around the post-merger galaxy cluster MS 0451-03. *Monthly Notices of the Royal Astronomical Society* **496**, 4032–4050 (2020).
- [19] Milgrom, M. A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. *ApJ* **270**, 365 (1983).
- [20] Famaey, B. & McGaugh, S. S. Modified Newtonian Dynamics (MOND): Observational Phenomenology and Relativistic Extensions. *Living Rev. Relativ.* **15**, 10 (2012).
- [21] Angus, G. W., Famaey, B. & Zhao, H. S. Can MOND take a bullet? analytical comparisons of three versions of MOND beyond spherical symmetry: Comparing three versions of MOND. *Monthly Notices of the Royal Astronomical Society* **371**, 138–146 (2006).
- [22] Angus, G. W., Shan, H. Y., Zhao, H. S. & Famaey, B. On the Proof of Dark Matter, the Law of Gravity, and the Mass of Neutrinos. *ApJ* **654**, L13–L16 (2007).
- [23] Milgrom, M. Testing MOND over a Wide Acceleration Range in X-Ray Ellipticals. *Phys. Rev. Lett.* **109**, 131101 (2012).
- [24] Chae, K.-H., Lelli, F., Desmond, H. *et al.* Testing the Strong Equivalence Principle: Detection of the External Field Effect in Rotationally Supported Galaxies. *ApJ* **904**, 51 (2020).

-
- [25] Marsh, D. J. E. Axion Cosmology. *Physics Reports* **643**, 1–79 (2016).
- [26] Crewther, R., Di Vecchia, P., Veneziano, G. & Witten, E. Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics. *Physics Letters B* **88**, 123–127 (1979).
- [27] Morgan, M. A. & Miller, G. A. The neutron electric dipole moment in the cloudy bag model. *Physics Letters B* **179**, 379–384 (1986).
- [28] Aoki, S. & Hatsuda, T. Strong CP violation and the neutron electric dipole moment reexamined. *Phys. Rev. D* **45**, 2427–2436 (1992).
- [29] Pospelov, M. & Ritz, A. Theta-Induced Electric Dipole Moment of the Neutron via QCD Sum Rules. *Phys. Rev. Lett.* **83**, 2526–2529 (1999).
- [30] Harris, P. G., Baker, C. A., Green, K. *et al.* New Experimental Limit on the Electric Dipole Moment of the Neutron. *Phys. Rev. Lett.* **82**, 904–907 (1999).
- [31] Baker, C. A., Doyle, D. D., Geltenbort, P. *et al.* Improved Experimental Limit on the Electric Dipole Moment of the Neutron. *Phys. Rev. Lett.* **97**, 131801 (2006).
- [32] Pendlebury, J. M., Afach, S., Ayres, N. J. *et al.* Revised experimental upper limit on the electric dipole moment of the neutron. *Phys. Rev. D* **92**, 092003 (2015).
- [33] Romalis, M. V., Griffith, W. C., Jacobs, J. P. & Fortson, E. N. New Limit on the Permanent Electric Dipole Moment of ^{199}Hg . *Phys. Rev. Lett.* **86**, 2505–2508 (2001).
- [34] Shifman, M., Vainshtein, A. & Zakharov, V. Can confinement ensure natural CP invariance of strong interactions? *Nuclear Physics B* **166**, 493–506 (1980).
- [35] Kim, J. E. Light pseudoscalars, particle physics and cosmology. *Physics Reports* **150**, 1–177 (1987).
- [36] Peccei, R. D. & Quinn, H. R. CP Conservation in the Presence of Pseudoparticles. *Physical Review Letters* **38**, 1440–1443 (1977).
- [37] Peccei, R. D. & Quinn, H. R. Constraints imposed by CP conservation in the presence of pseudoparticles. *Phys. Rev. D* **16**, 1791–1797 (1977).
- [38] Wilczek, F. Problem of Strong P and T Invariance in the Presence of Instantons. *Physical Review Letters* **40**, 279–282 (1978).
- [39] Weinberg, S. A New Light Boson? *Physical Review Letters* **40**, 223–226 (1978).
- [40] Callan, C. G., Dashen, R. & Gross, D. J. Toward a theory of the strong interactions. *Phys. Rev. D* **17**, 2717–2763 (1978).
- [41] Coleman, S. R. The uses of instantons. *Subnucl. Ser.* **15**, 805 (1979).

- [42] Vafa, C. & Witten, E. Parity Conservation in Quantum Chromodynamics. *Phys. Rev. Lett.* **53**, 535–536 (1984).
- [43] Svrcek, P. & Witten, E. Axions in string theory. *J. High Energy Phys.* **2006**, 051–051 (2006).
- [44] Srednicki, M. Axion couplings to matter. *Nuclear Physics B* **260**, 689–700 (1985).
- [45] Kim, J. E. Weak-Interaction Singlet and Strong CP Invariance. *Phys. Rev. Lett.* **43**, 103–107 (1979).
- [46] Kim, J. E. Constraints on very light axions from cavity experiments. *Phys. Rev. D* **58**, 055006 (1998).
- [47] Zhitnitsky, A. R. On possible suppression of the axion hadron interactions. (in russian). *Sov. J. Nucl. Phys.* **31**, 260 (1980).
- [48] Dine, M., Fischler, W. & Srednicki, M. A simple solution to the strong CP problem with a harmless axion. *Physics Letters B* **104**, 199–202 (1981).
- [49] Ferreira, R. Z., Notari, A. & Rompineve, F. Dine-Fischler-Srednicki-Zhitnitsky axion in the CMB. *Phys. Rev. D* **103**, 063524 (2021).
- [50] Arvanitaki, A., Dimopoulos, S., Dubovsky, S. *et al.* String axiverse. *Phys. Rev. D* **81**, 123530 (2010).
- [51] Witten, E. Some properties of $O(32)$ superstrings. *Physics Letters B* **149**, 351–356 (1984).
- [52] Raffelt, G. & Seckel, D. Bounds on exotic-particle interactions from SN1987A. *Phys. Rev. Lett.* **60**, 1793–1796 (1988).
- [53] Raffelt, G. G. Astrophysics probes of particle physics. *Physics Reports* **333–334**, 593–618 (2000).
- [54] Chang, J. H., Essig, R. & McDermott, S. D. Supernova 1987A constraints on sub-GeV dark sectors, millicharged particles, the QCD axion, and an axion-like particle. *Journal of High Energy Physics* **2018**, 51 (2018).
- [55] Burrows, A. & Lattimer, J. M. Neutrinos from SN 1987A. *ApJ* **318**, L63 (1987).
- [56] Janka, H.-T., Keil, W., Raffelt, G. & Seckel, D. Nucleon Spin Fluctuations and the Supernova Emission of Neutrinos and Axions. *Phys. Rev. Lett.* **76**, 2621–2624 (1996).
- [57] Keil, W., Janka, H.-T., Schramm, D. N. *et al.* Fresh look at axions and SN 1987A. *Phys. Rev. D* **56**, 2419–2432 (1997).
- [58] Bar, N., Blum, K. & D’Amico, G. Is there a supernova bound on axions? *Phys. Rev. D* **101**, 123025 (2020).

-
- [59] DeRocco, W., Graham, P. W. & Rajendran, S. Exploring the robustness of stellar cooling constraints on light particles. *Phys. Rev. D* **102**, 075015 (2020).
- [60] Preskill, J., Wise, M. B. & Wilczek, F. Cosmology of the invisible axion. *Physics Letters B* **120**, 127–132 (1983).
- [61] Abbott, L. & Sikivie, P. A cosmological bound on the invisible axion. *Physics Letters B* **120**, 133–136 (1983).
- [62] Dine, M. & Fischler, W. The not-so-harmless axion. *Physics Letters B* **120**, 137–141 (1983).
- [63] Blum, K., D’Agnolo, R. T., Lisanti, M. & Safdi, B. R. Constraining axion dark matter with Big Bang Nucleosynthesis. *Physics Letters B* **737**, 30–33 (2014).
- [64] Giusarma, E., Di Valentino, E., Lattanzi, M. *et al.* Relic neutrinos, thermal axions, and cosmology in early 2014. *Phys. Rev. D* **90**, 043507 (2014).
- [65] Di Valentino, E., Giusarma, E., Lattanzi, M. *et al.* Cosmological axion and neutrino mass constraints from Planck 2015 temperature and polarization data. *Physics Letters B* **752**, 182–185 (2016).
- [66] Jungman, G., Kamionkowski, M. & Griest, K. Supersymmetric dark matter. *Physics Reports* **267**, 195–373 (1996).
- [67] Martin, S. P. A Supersymmetry Primer. *arXiv:hep-ph/9709356* (2016).
- [68] Haber, H. E. & Haskins, L. S. Supersymmetric Theory and Models. *arXiv:1712.05926 [hep-ph]* (2020).
- [69] Berlin, A., Lin, T., Low, M. & Wang, L.-T. Neutralinos in vector boson fusion at high energy colliders. *Phys. Rev. D* **91**, 115002 (2015).
- [70] Zyla, P. *et al.* Review of Particle Physics. *PTEP* **2020**, 083C01 (2020).
- [71] Yu. A. Gol’fand & E. P. Likhtman. Extension of the algebra of Poincare group generators and violation of P invariance. *JETP Lett.* **13**, 323 (1971).
- [72] Ramond, P. Dual Theory for Free Fermions. *Phys. Rev. D* **3**, 2415–2418 (1971).
- [73] Neveu, A. & Schwarz, J. Factorizable dual model of pions. *Nuclear Physics B* **31**, 86–112 (1971).
- [74] Gervais, J.-L. & Sakita, B. Field theory interpretation of supergauges in dual models. *Nuclear Physics B* **34**, 632–639 (1971).
- [75] Volkov, D. & Akulov, V. Is the neutrino a goldstone particle? *Physics Letters B* **46**, 109–110 (1973).

- [76] Wess, J. & Zumino, B. Supergauge transformations in four dimensions. *Nuclear Physics B* **70**, 39–50 (1974).
- [77] Chapter 3: Beyond the Standard Model Phenomena. In Mangano, M., Golling, T., Hance, M. *et al.* (eds.) *Physics at the FCC-Hh, a 100 TeV Pp Collider*, vol. 3 of *CERN Yellow Reports Monographs*, 441–634 (CERN, Geneva, 2017).
- [78] Pontecorvo, B. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *JETP* **26**, 984–988 (1968).
- [79] Gribov, V. & Pontecorvo, B. Neutrino astronomy and lepton charge. *Physics Letters B* **28**, 493–496 (1969).
- [80] Mann, A. K. & Primakoff, H. Neutrino oscillations and the number of neutrino types. *Phys. Rev. D* **15**, 655–665 (1977).
- [81] Fukuda, Y., Hayakawa, T., Ichihara, E. *et al.* Evidence for Oscillation of Atmospheric Neutrinos. *Physical Review Letters* **81**, 1562–1567 (1998).
- [82] Gonzalez-Garcia, M. & Maltoni, M. Phenomenology with massive neutrinos. *Physics Reports* **460**, 1–129 (2008).
- [83] Bond, J. R., Efstathiou, G. & Silk, J. Massive Neutrinos and the Large-Scale Structure of the Universe. *Phys. Rev. Lett.* **45**, 1980–1984 (1980).
- [84] Jarosik, N., Bennett, C. L., Dunkley, J. *et al.* Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Sky maps, systematic errors, and basic results. *ApJS* **192**, 14 (2011).
- [85] Holdom, B. Two U(1)’s and ϵ charge shifts. *Physics Letters B* **166**, 196–198 (1986).
- [86] Nelson, A. E. & Scholtz, J. Dark light, dark matter, and the misalignment mechanism. *Phys. Rev. D* **84**, 103501 (2011).
- [87] Graham, P. W., Mardon, J. & Rajendran, S. Vector dark matter from inflationary fluctuations. *Phys. Rev. D* **93**, 103520 (2016).
- [88] Ahmed, A., Grzadkowski, B. & Socha, A. Gravitational production of vector dark matter. *J. High Energ. Phys.* **2020**, 59 (2020).
- [89] Kolb, E. W. & Long, A. J. Completely dark photons from gravitational particle production during the inflationary era. *J. High Energ. Phys.* **2021**, 283 (2021).
- [90] Pospelov, M. Secluded U(1) below the weak scale. *Phys. Rev. D* **80**, 095002 (2009).
- [91] Endo, M., Hamaguchi, K. & Mishima, G. Constraints on hidden photon models from electron $g - 2$ and hydrogen spectroscopy. *Phys. Rev. D* **86**, 095029 (2012).
- [92] Dixit, A. V., Chakram, S., He, K. *et al.* Searching for Dark Matter with a Superconducting Qubit. *Phys. Rev. Lett.* **126**, 141302 (2021).

- [93] Fedderke, M. A., Graham, P. W., Kimball, D. F. J. & Kalia, S. The Earth as a transducer for dark-photon dark-matter detection. *arXiv:2106.00022 [hep-ph]* (2021).
- [94] Fedderke, M. A., Graham, P. W., Kimball, D. F. J. & Kalia, S. Search for dark-photon dark matter in the SuperMAG geomagnetic field dataset. *arXiv:2108.08852 [astro-ph, physics:hep-ph]* (2021).
- [95] Graham, P. W., Kaplan, D. E. & Rajendran, S. Cosmological Relaxation of the Electroweak Scale. *Physical Review Letters* **115** (2015).
- [96] Fonseca, N. & Morgante, E. Relaxion dark matter. *Phys. Rev. D* **100**, 055010 (2019).
- [97] Banerjee, A., Kim, H. & Perez, G. Coherent relaxion dark matter. *Phys. Rev. D* **100**, 115026 (2019).
- [98] Banerjee, A., Budker, D., Eby, J. *et al.* Relaxion stars and their detection via atomic physics. *Commun Phys* **3**, 1 (2020).
- [99] Witten, E. Cosmic separation of phases. *Phys. Rev. D* **30**, 272–285 (1984).
- [100] Gorham, P. W. & Rotter, B. J. Stringent neutrino flux constraints on antiquark nugget dark matter. *Phys. Rev. D* **95**, 103002 (2017).
- [101] Budker, D., Flambaum, V. V. & Zhitnitsky, A. Infrasonic, acoustic and seismic waves produced by the Axion Quark Nuggets. *arXiv:2003.07363 [hep-ph]* (2020).
- [102] Budker, D., Flambaum, V. V., Liang, X. & Zhitnitsky, A. Axion quark nuggets and how a global network can discover them. *Phys. Rev. D* **101**, 043012 (2020).
- [103] VanDevender, J. P., VanDevender, A. P., Sloan, T. *et al.* Detection of magnetized quark-nuggets, a candidate for dark matter. *Sci Rep* **7**, 8758 (2017).
- [104] Bai, Y., Long, A. J. & Lu, S. Dark quark nuggets. *Phys. Rev. D* **99**, 055047 (2019).
- [105] Zel’dovich, Y. B. & Novikov, I. B. The Hypothesis of Cores Retarded during Expansion and the Hot Cosmological Model. *Soviet Astronomy* **10**, 602 (1967).
- [106] Carr, B. J. & Hawking, S. W. Black Holes in the Early Universe. *Monthly Notices of the Royal Astronomical Society* **168**, 399–415 (1974).
- [107] Carr, B., Kühnel, F. & Sandstad, M. Primordial black holes as dark matter. *Phys. Rev. D* **94**, 083504 (2016).
- [108] Hawking, S. W. Black hole explosions? *Nature* **248**, 30–31 (1974).
- [109] Budker, D., Graham, P. W., Ledbetter, M. *et al.* Proposal for a Cosmic Axion Spin Precession Experiment (CASPER). *Physical Review X* **4** (2014).

-
- [110] Garcon, A., Aybas, D., Blanchard, J. W. *et al.* The cosmic axion spin precession experiment (CASPER): A dark-matter search with nuclear magnetic resonance. *Quantum Sci. Technol.* **3**, 014008 (2018).
- [111] Wu, T., Blanchard, J. W., Centers, G. P. *et al.* Search for Axionlike Dark Matter with a Liquid-State Nuclear Spin Comagnetometer. *Phys. Rev. Lett.* **122**, 191302 (2019).
- [112] Garcon, A., Blanchard, J. W., Centers, G. P. *et al.* Constraints on bosonic dark matter from ultralow-field nuclear magnetic resonance. *Sci. Adv.* **5**, eaax4539 (2019).
- [113] Aybas, D., Adam, J., Blumenthal, E. *et al.* Search for Axionlike Dark Matter Using Solid-State Nuclear Magnetic Resonance. *Phys. Rev. Lett.* **126**, 141802 (2021).
- [114] Terrano, W. A., Adelberger, E. G., Hagedorn, C. A. & Heckel, B. R. Constraints on Axionlike Dark Matter with Masses Down to 10^{-23} eV / c^2 . *Phys. Rev. Lett.* **122**, 231301 (2019).
- [115] Kahn, Y., Safdi, B. R. & Thaler, J. Broadband and Resonant Approaches to Axion Dark Matter Detection. *Phys. Rev. Lett.* **117**, 141801 (2016).
- [116] Ouellet, J. L., Salemi, C. P., Foster, J. W. *et al.* First Results from ABRACADABRA-10 cm: A Search for Sub- μ eV Axion Dark Matter. *Phys. Rev. Lett.* **122**, 121802 (2019).
- [117] CAST Collaboration. New CAST limit on the axion–photon interaction. *Nature Phys* **13**, 584–590 (2017).
- [118] Armengaud, E., Avignone, F. T., Betz, M. *et al.* Conceptual design of the International Axion Observatory (IAXO). *J. Inst.* **9**, T05002–T05002 (2014).
- [119] MADMAX Collaboration, Brun, P., Caldwell, A. *et al.* A new experimental approach to probe QCD axion dark matter in the mass range above 40 μ eV. *Eur. Phys. J. C* **79**, 186 (2019).
- [120] Du, N., Force, N., Khatiwada, R. *et al.* Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment. *Phys. Rev. Lett.* **120**, 151301 (2018).
- [121] Backes, K. M., Palken, D. A., Kenany, S. A. *et al.* A quantum enhanced search for dark matter axions. *Nature* **590**, 238–242 (2021).
- [122] Ehret, K., Frede, M., Ghazaryan, S. *et al.* New ALPS results on hidden-sector lightweights. *Physics Letters B* **689**, 149–155 (2010).
- [123] Bähre, R., Döbrich, B., Dreyling-Eschweiler, J. *et al.* Any light particle search II — Technical Design Report. *J. Inst.* **8**, T09001–T09001 (2013).

-
- [124] Roberts, B. M., Blewitt, G., Dailey, C. *et al.* Search for domain wall dark matter with atomic clocks on board global positioning system satellites. *Nature Communications* **8** (2017).
- [125] Roberts, B. M., Blewitt, G., Dailey, C. & Derevianko, A. Search for transient ultralight dark matter signatures with networks of precision measurement devices using a Bayesian statistics method. *Physical Review D* **97** (2018).
- [126] Aprile, E., Arisaka, K., Arneodo, F. *et al.* The XENON100 dark matter experiment. *Astroparticle Physics* **35**, 573–590 (2012).
- [127] Aprile, E., Aalbers, J., Agostini, F. *et al.* Excess electronic recoil events in XENON1T. *Phys. Rev. D* **102**, 072004 (2020).
- [128] Sikivie, P. Axions, Domain Walls, and the Early Universe. *Phys. Rev. Lett.* **48**, 1156–1159 (1982).
- [129] Sikivie, P. Axion Cosmology. *arXiv:astro-ph/0610440* (2008).
- [130] Hiramatsu, T., Kawasaki, M., Saikawa, K. & Sekiguchi, T. Production of dark matter axions from collapse of string-wall systems. *Phys. Rev. D* **85**, 105020 (2012).
- [131] Hiramatsu, T., Kawasaki, M., Saikawa, K. & Sekiguchi, T. Axion cosmology with long-lived domain walls. *J. Cosmol. Astropart. Phys.* **2013**, 001–001 (2013).
- [132] Barr, S. M. & Kim, J. E. New Confining Force Solution of the QCD Axion Domain-Wall Problem. *Phys. Rev. Lett.* **113**, 241301 (2014).
- [133] Larsson, S. E., Sarkar, S. & White, P. L. Evading the cosmological domain wall problem. *Phys. Rev. D* **55**, 5129–5135 (1997).
- [134] Caputo, A. & Reig, M. Cosmic implications of a low-scale solution to the axion domain wall problem. *Phys. Rev. D* **100**, 063530 (2019).
- [135] Ibe, M., Kobayashi, S., Suzuki, M. & Yanagida, T. T. Dynamical solution to the axion domain wall problem. *Phys. Rev. D* **101**, 035029 (2020).
- [136] Braaten, E., Mohapatra, A. & Zhang, H. Dense Axion Stars. *Phys. Rev. Lett.* **117**, 121801 (2016).
- [137] Eby, J., Leembruggen, M., Street, L. *et al.* On Approximation Methods in the Study of Boson Stars. *Physical Review D* **98** (2018).
- [138] Braaten, E. & Zhang, H. *Colloquium* : The physics of axion stars. *Reviews of Modern Physics* **91** (2019).
- [139] Visinelli, L., Baum, S., Redondo, J. *et al.* Dilute and dense axion stars. *Physics Letters B* **777**, 64–72 (2018).

- [140] Rosen, G. Particlelike Solutions to Nonlinear Complex Scalar Field Theories with Positive-Definite Energy Densities. *Journal of Mathematical Physics* **9**, 996–998 (1968).
- [141] Coleman, S. Q-balls. *Nuclear Physics B* **262**, 263–283 (1985).
- [142] Kusenko, A. & Steinhardt, P. J. Q -Ball Candidates for Self-Interacting Dark Matter. *Physical Review Letters* **87** (2001).
- [143] Dvali, G., Kusenko, A. & Shaposhnikov, M. New physics in a nutshell, or Q-ball as a power plant. *Physics Letters B* **417**, 99–106 (1998).
- [144] Heeck, J., Rajaraman, A., Riley, R. & Verhaaren, C. B. Understanding Q -balls beyond the thin-wall limit. *Phys. Rev. D* **103**, 045008 (2021).
- [145] Freese, K., Lisanti, M. & Savage, C. *Colloquium* : Annual modulation of dark matter. *Rev. Mod. Phys.* **85**, 1561–1581 (2013).
- [146] Budker, D. & Romalis, M. Optical magnetometry. *Nature Phys* **3**, 227–234 (2007).
- [147] Włodarczyk, P., Pustelny, S., Budker, D. & Lipiński, M. Multi-channel data acquisition system with absolute time synchronization. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **763**, 150–154 (2014).
- [148] Ledbetter, M. P., Savukov, I. M., Acosta, V. M. *et al.* Spin-exchange-relaxation-free magnetometry with Cs vapor. *Phys. Rev. A* **77**, 033408 (2008).
- [149] Padniuk, M., Kopciuch, M., Cipolletti, R. *et al.* Self-compensating co-magnetometer vs. spin-exchange relaxation-free magnetometer: Sensitivity to non-magnetic spin couplings. *arXiv:2107.05501 [physics.ins-det]* (2021).
- [150] Roberts, B. M., Delva, P., Al-Masoudi, A. *et al.* Search for transient variations of the fine structure constant and dark matter using fiber-linked optical atomic clocks. *New J. Phys.* **22**, 093010 (2020).
- [151] Leefer, N., Weber, C. T. M., Cingöz, A. *et al.* New Limits on Variation of the Fine-Structure Constant Using Atomic Dysprosium. *Phys. Rev. Lett.* **111**, 060801 (2013).
- [152] Antypas, D., Tretiak, O., Zhang, K. *et al.* Probing fast oscillating scalar dark matter with atoms and molecules. *Quantum Sci. Technol.* **6**, 034001 (2021).
- [153] McNally, R. L. & Zelevinsky, T. Constraining domain wall dark matter with a network of superconducting gravimeters and LIGO. *Eur. Phys. J. D* **74**, 61 (2020).
- [154] Figueroa, N. L., Budker, D. & Rasel, E. M. Dark matter searches using accelerometer-based networks. *Quantum Sci. Technol.* **6**, 034004 (2021).

- [155] Stadnik, Y. V. New bounds on macroscopic scalar-field topological defects from nontransient signatures due to environmental dependence and spatial variations of the fundamental constants. *Phys. Rev. D* **102**, 115016 (2020).
- [156] Schmidt, T. Über die magnetischen Momente der Atomkerne. *Z. Physik* **106**, 358–361 (1937).
- [157] Engel, J. & Vogel, P. Spin-dependent cross sections of weakly interacting massive particles on nuclei. *Phys. Rev. D* **40**, 3132–3135 (1989).
- [158] Iachello, F., Krauss, L. M. & Maino, G. Spin-dependent scattering of weakly interacting massive particles in heavy nuclei. *Physics Letters B* **254**, 220–224 (1991).
- [159] Engel, J., Ressel, M. T., Towner, I. S. & Ormand, W. E. Response of mica to weakly interacting massive particles. *Phys. Rev. C* **52**, 2216–2221 (1995).
- [160] Flambaum, V. V. & Tedesco, A. F. Dependence of nuclear magnetic moments on quark masses and limits on temporal variation of fundamental constants from atomic clock experiments. *Phys. Rev. C* **73**, 055501 (2006).
- [161] Toivanen, P., Kortelainen, M., Suhonen, J. & Toivanen, J. Large-scale shell-model calculations of elastic and inelastic scattering rates of lightest supersymmetric particles (LSP) on ^{127}I , ^{129}Xe , ^{131}Xe , and ^{133}Cs nuclei. *Phys. Rev. C* **79**, 044302 (2009).
- [162] Kimball, D. F. J. Nuclear spin content and constraints on exotic spin-dependent couplings. *New J. Phys.* **17**, 073008 (2015).
- [163] di Cortona, G. G., Hardy, E., Vega, J. P. & Villadoro, G. The QCD axion, precisely. *J. High Energ. Phys.* **2016**, 34 (2016).
- [164] Kornack, T. W. & Romalis, M. V. Dynamics of Two Overlapping Spin Ensembles Interacting by Spin Exchange. *Phys. Rev. Lett.* **89**, 253002 (2002).
- [165] Kornack, T. W., Ghosh, R. K. & Romalis, M. V. Nuclear Spin Gyroscope Based on an Atomic Comagnetometer. *Phys. Rev. Lett.* **95**, 230801 (2005).
- [166] Roston, G. & Obaid, F. Exact analytical formula for Voigt spectral line profile. *Journal of Quantitative Spectroscopy and Radiative Transfer* **94**, 255–263 (2005).

Appendix A

Useful derivations

A.1 Domain wall shape

Consider the second-order differential equation,

$$a''(x) + \mu^2 \sin(a(x)) = 0 . \quad (\text{A.1})$$

Replacing $a \rightarrow a/f_a$ and $\mu \rightarrow m_a$ reproduces Eq. (3.5) from Sec. 3.1. Consider the boundary conditions,

$$\begin{aligned} a(0) &= 0 , \\ a(x \rightarrow \pm\infty) &= \pm\pi , \\ a'(x \rightarrow \pm\infty) &= 0 . \end{aligned}$$

This corresponds to a solution in which a varies from $-\pi$ to $+\pi$, centered at $x = 0$. Note that the constraint on $a(x \rightarrow \pm\infty)$ implies the constraint on $a'(x \rightarrow \pm\infty)$. Finally, demand that a be monotonically increasing.

Eq. (A.1) is solved with these boundary conditions by

$$a(x) = 2 \arcsin(\tanh(\mu x)) . \quad (\text{A.2})$$

Proof. First, one can reduce this second-order equation to a first-order one by invoking the $x \rightarrow -\infty$ constraints:

$$\begin{aligned} 0 &= a'' + \mu^2 \sin(a) \\ 0 &= 2a'a'' + 2a'\mu^2 \sin(a) \\ &= \frac{d}{dx} \left[(a')^2 - 2\mu^2 \cos(a) \right] \\ \int_{-\infty}^x d\bar{x} \cdot 0 &= \int_{-\infty}^x d\bar{x} \frac{d}{d\bar{x}} \left[(a')^2 - 2\mu^2 \cos(a) \right] \\ 0 &= (a')^2 - 2\mu^2 [\cos(a) + 1] . \end{aligned}$$

Demanding that a is monotonically increasing, so $a' \geq 0$, this can be simplified to

$$\frac{a'(x)}{2} = \mu \cos\left(\frac{a(x)}{2}\right).$$

Solving this equation can be done with the aid of a change-of-variables. In particular,

$$\frac{a}{2} = \arcsin c \quad \text{so} \quad \frac{a'}{2} = \frac{c'}{\sqrt{1-c^2}}.$$

Now, the $x = 0$ constraint is used so $c(0) = \sin(a(0)/2) = 0$. Then,

$$\begin{aligned} \frac{c'}{\sqrt{1-c^2}} &= \mu \sqrt{1-c^2} \\ \int_0^x d\bar{x} \frac{c'(\bar{x})}{1-c^2(\bar{x})} &= \int_0^x d\bar{x} \mu \\ \int_0^x d\bar{x} \frac{d}{d\bar{x}} \tanh^{-1}(c(\bar{x})) &= \mu x \\ \tanh^{-1}(c(x)) &= \mu x \\ \sin\left(\frac{a(x)}{2}\right) &= \tanh(\mu x). \end{aligned}$$

Finally, one obtains the solution,

$$a(x) = 2 \arcsin(\tanh(\mu x)).$$

□

A.2 χ^2 minimization

Consider a minimization problem wherein one has the measurements \mathbf{y} and a set of unknown parameters \mathbf{x} . The unknown parameters will model the measurements with the function $\boldsymbol{\mu}(\mathbf{x})$ that predicts \mathbf{y} ,

$$\boldsymbol{\mu}(\mathbf{x}) \approx \mathbf{y}, \tag{A.3}$$

where deviation from equality can occur due to measurement errors. These errors can be described by the covariance matrix $\Sigma_{\mathbf{y}}$. The goal then is to minimize

$$\chi^2(\mathbf{x}) := (\mathbf{y} - \boldsymbol{\mu}(\mathbf{x}))^T \Sigma_{\mathbf{y}}^{-1} (\mathbf{y} - \boldsymbol{\mu}(\mathbf{x})) \tag{A.4}$$

to find the most-likely \mathbf{x} to describe the observed \mathbf{y} . The most-likely parameters are given by

$$\mathbf{x}_0 = \Sigma_{\mathbf{x}} J^T \Sigma_{\mathbf{y}}^{-1} (\mathbf{y} - \mathbf{b}) \quad \text{for} \quad \Sigma_{\mathbf{x}} = (J^T \Sigma_{\mathbf{y}}^{-1} J)^{-1}, \tag{A.5}$$

where $J_{ij} = \frac{\partial \mu_i(\mathbf{x})}{\partial x_j}$ is the Jacobian matrix and $\mathbf{b} := \boldsymbol{\mu} - J\mathbf{x}$. Observe that \mathbf{b} is exactly the affine component of $\boldsymbol{\mu}$ if $\boldsymbol{\mu}$ is of the form $\boldsymbol{\mu}(\mathbf{x}) = A\mathbf{x} + \mathbf{b}$. Likewise, $J = A$ in this case (independent of \mathbf{x}).

Proof. Denote $\frac{\partial}{\partial x_k} = \partial_k$ so minimizing Eq. (A.4) requires that

$$\begin{aligned}
0 &= \partial_k \chi^2(\mathbf{x}) \\
&= \sum_{i,j} \partial_k (\mu_i(\mathbf{x}) - y_i) (\Sigma_{\mathbf{y}}^{-1})_{ij} (\mu_j(\mathbf{x}) - y_j) \\
&= \sum_{i,j} (\partial_k \mu_i(\mathbf{x})) (\Sigma_{\mathbf{y}}^{-1})_{ij} (\mu_j(\mathbf{x}) - y_j) \\
&\quad \sum_{i,j} (\mu_i(\mathbf{x}) - y_i) (\Sigma_{\mathbf{y}}^{-1})_{ij} (\partial_k \mu_j(\mathbf{x})) \\
&= 2 \sum_{i,j} \frac{\partial \mu_i(\mathbf{x})}{\partial x_k} (\Sigma_{\mathbf{y}}^{-1})_{ij} (\mu_j(\mathbf{x}) - y_j),
\end{aligned}$$

where the last line relabels $i \leftrightarrow j$ in the second sum and uses the symmetry of the covariance matrix $(\Sigma_{\mathbf{x}}^{-1})_{ij} = (\Sigma_{\mathbf{x}}^{-1})_{ji}$. Noting the Jacobian matrix $J_{ik}(\mathbf{x}) = \frac{\partial \mu_i(\mathbf{x})}{\partial x_k}$, the above equation can be simplified in matrix form,

$$J^T(\mathbf{m}) \Sigma_{\mathbf{x}}^{-1} \boldsymbol{\mu}(\mathbf{x}) = J^T(\mathbf{m}) \Sigma_{\mathbf{x}}^{-1} \mathbf{y}. \quad (\text{A.6})$$

Let \mathbf{x}_0 minimize χ^2 with $\boldsymbol{\mu}_0 = \boldsymbol{\mu}(\mathbf{x}_0)$ and $J(\mathbf{x}_0) = J_0$. Expanding $\boldsymbol{\mu}(\mathbf{x})$ around \mathbf{x}_0 ,

$$\boldsymbol{\mu}(\mathbf{x}) = \boldsymbol{\mu}_0 + J_0(\mathbf{x} - \mathbf{x}_0) + \mathcal{O}((\mathbf{x} - \mathbf{x}_0)^2).$$

Assuming non-linear terms are negligible, Eq. (A.6) becomes

$$\begin{aligned}
J_0^T \Sigma_{\mathbf{x}}^{-1} (\boldsymbol{\mu}_0 + J_0(\mathbf{x} - \mathbf{x}_0)) &\approx J_0^T \Sigma_{\mathbf{x}}^{-1} \mathbf{y} \\
J_0^T \Sigma_{\mathbf{x}}^{-1} \mathbf{y} + J_0^T \Sigma_{\mathbf{x}}^{-1} (J_0 \mathbf{x}_0 - \boldsymbol{\mu}_0) &\approx J_0^T \Sigma_{\mathbf{x}}^{-1} J_0 \mathbf{x} \\
(J_0^T \Sigma_{\mathbf{x}}^{-1} J_0)^{-1} J_0^T \Sigma_{\mathbf{x}}^{-1} (\mathbf{y} - \mathbf{b}) &\approx \mathbf{x},
\end{aligned} \quad (\text{A.7})$$

for $\mathbf{b} := (\boldsymbol{\mu}_0 - J_0 \mathbf{x}_0)$.

Now, in general, for $\mathbf{u}(\mathbf{v}) = A\mathbf{v} + \mathbf{b}$, the error in \mathbf{v} , $\Sigma_{\mathbf{v}}$, is propagated into error in \mathbf{u} , $\Sigma_{\mathbf{u}}$, as

$$\begin{aligned}
\Sigma_{\mathbf{u}} &= \left(\frac{\partial \mathbf{u}}{\partial \mathbf{v}} \right) \Sigma_{\mathbf{v}} \left(\frac{\partial \mathbf{u}}{\partial \mathbf{v}} \right)^T \\
&= A \Sigma_{\mathbf{v}} A^T.
\end{aligned} \quad (\text{A.8})$$

Plugging Eq. (A.7) into Eq. (A.8) yields

$$\begin{aligned}
\Sigma_{\mathbf{x}} &= (J_0^T \Sigma_{\mathbf{x}}^{-1} J_0)^{-1} J_0^T \Sigma_{\mathbf{x}}^{-1} \cdot \Sigma_{\mathbf{x}} \cdot \left((J_0^T \Sigma_{\mathbf{x}}^{-1} J_0)^{-1} J_0^T \Sigma_{\mathbf{x}}^{-1} \right)^T \\
&= (J_0^T \Sigma_{\mathbf{x}}^{-1} J_0)^{-1} J_0^T \Sigma_{\mathbf{x}}^{-1} J_0 (J_0^T \Sigma_{\mathbf{x}}^{-1} J_0)^{-1} \\
&= (J_0^T \Sigma_{\mathbf{x}}^{-1} J_0)^{-1},
\end{aligned}$$

where the second equality uses the fact that $\Sigma_{\mathbf{x}}^{-1}$ and $(J_0^T \Sigma_{\mathbf{x}}^{-1} J_0)^{-1}$ are symmetric. Using this result to simplify Eq. (A.7) yields the desired

$$\mathbf{x} = \Sigma_{\mathbf{x}} J_0^T \Sigma_{\mathbf{x}}^{-1} (\mathbf{y} - \mathbf{b}) \quad \text{for} \quad \Sigma_{\mathbf{x}} = (J_0^T \Sigma_{\mathbf{x}}^{-1} J_0)^{-1}.$$

□

A.3 Sensitivity scaling

It is of interest to the growth of GNOME to understand how the addition of new sensors affects the sensitivity of the network as a whole. For this derivation, the coordinates will be generalized in d dimensions and the network will be assumed to have random, evenly distributed sensitive directions. The uncertainty scaling will be expressed with respect to the noise σ in an individual magnetometer. The collective sensitivity β then scales as

$$\beta \approx \sigma \sqrt{\frac{d}{n}}, \quad (\text{A.9})$$

for n sensors.

The derivation will be split into multiple parts to establish some background. In Sec. A.3.1, spherical coordinates will be generalized in d dimensions, and in Sec. A.3.2, the volume element will be described in these coordinates. The proof of Eq. (A.9) is given in Sec. A.3.3.

A.3.1 Coordinate system

It helps to generalize spherical coordinates for d dimensions. Denote $\mathbf{x} \in \mathbb{R}^d$ with components x_0, \dots, x_{d-1} . Define the $d-1$ angles $\{\varphi_1, \dots, \varphi_{d-1}\}$ and radius r such that

$$x_0 = r \prod_{j=1}^{d-1} \sin(\varphi_j) \quad \text{and} \quad x_i = r \cos(\varphi_i) \prod_{j=i+1}^{d-1} \sin(\varphi_j). \quad (\text{A.10})$$

It remains still to show that this is in fact a bijection; in particular, for $r \geq 0$, $\varphi_1 \in [0, 2\pi)$, and $\varphi_i \in [0, \pi]$ (where $i > 1$). It will sometimes be useful to compare the coordinates in different dimensions, for this, the dimension will be places in a superscript, e.g., $\mathbf{x}^{(d)}$.

Proof. Observe that the defined coordinates can be restated inductively by dimension for $\mathbf{x}^{(d)} \in \mathbb{R}^d$. Define $r_i^{(k)} = \sqrt{\sum_{j=0}^{i-1} (x_j^{(k)})^2}$ as the radial distance projected into i -dimensions and $r^{(k)} = r_k^{(k)}$. For $d \geq 2$, the coordinates are defined inductively as

$$x_i^{(k)} = \begin{cases} x_i^{(k-1)} \sin(\varphi_{k-1}) & i \neq k-1 \\ r^{(k)} \cos(\varphi_{k-1}) & i = k-1 \end{cases}.$$

As an initial condition for the inductive definition, demand that $x_0^{(1)} = r^{(d)}$. Then

$$\begin{aligned} \left(r^{(k+1)}\right)^2 &= \left(x_k^{(k+1)}\right)^2 + \sum_{i=0}^{k-1} \left(x_i^{(k+1)}\right)^2 \\ &= \left(r^{(k+1)}\right)^2 \cos^2(\varphi_k) + \sin^2(\varphi_k) \sum_{i=0}^{k-1} \left(x_i^{(k)}\right)^2 \\ &= \left(r^{(k+1)}\right)^2 \cos^2(\varphi_k) + \left(r^{(k)}\right)^2 \sin^2(\varphi_k). \end{aligned}$$

The above equation holds for any φ_k , so $r^{(k+1)} = r^{(k)} = r^{(d)} = r$; that is, the radius parameter is constant across each inductive step. Moreover $r = r^{(d)}$ parameterizes radius (clearly $r \geq 0$); $r^{(d)}$ is the radius by definition, but this shows that r is also the radius.

To prove that the map to the new coordinate system is bijective (or, equivalently, invertible), it remains to show that the angles $\{\varphi_i\}$ can be described as a function of the Cartesian coordinates \mathbf{x} . Observe that the angles defined in the inductive coordinate definition are independent of the dimension. Consider first $d = 2$ wherein $x_0^{(2)} = r \sin(\varphi_1)$ and $x_1^{(2)} = r \cos(\varphi_1)$, this can be inverted using a function similar to $\arctan(x_0^{(2)}/x_1^{(2)})$ so that $\varphi_1 \in [0, 2\pi)$ as opposed to $(-\pi, \pi)$. In particular, consider the conventional xy -Cartesian coordinate system for $x = x_1^{(2)}$ and $y = x_0^{(2)}$ and let φ be the angle of (x, y) measured counter-clockwise from the x -axis. Denote the function used to obtain the angle as $\arctan(x, y)$. Observe that $x_1^{(d)}$ and $x_0^{(d)}$ will generate the same angle, because they are the same as their lower-dimension counterparts up to a shared factor of $\prod_{i=2}^{d-1} \sin \varphi_i$. The remaining angles $\varphi_2, \dots, \varphi_{d-1}$ are similarly defined by considering $x = x_{k-1}^{(k)}$ and $y = x_0^{(k)}$ for $k > 2$; in particular, $x = r \cos(\varphi_{k-1})$, $y = r \sin(\varphi_{k-1})$, and $\varphi_{k-1} = \arctan(x, y)$. However $y \geq 0$, so only $\varphi_{k-1} \in [0, \pi]$ are possible. Again, this definition works for the higher-dimensional counter-parts because the two coordinates scale by the same amount, $\prod_{i=k-1}^{d-1} \sin(\varphi_i)$.

In summary, the inverse of the spherical- to Cartesian-coordinates is given by

$$\begin{aligned} r &= \sqrt{\sum_{j=0}^{d-1} x_j^2} \geq 0, \\ \varphi_1 &= \arctan(x_1, x_0) \in [0, 2\pi), \\ \varphi_i &= \arctan\left(x_i, \sqrt{\sum_{j=0}^{i-1} x_j^2}\right) \in [0, \pi] \quad \text{for } i > 1. \end{aligned} \tag{A.11}$$

The only possible issue with the ranges is that if \mathbf{x} is zero, then $r = 0$ and the angles can take any value. However, this will have no effect on integration because the origin is a set of measure zero, and the restriction $r = 1$ is of interest here. \square

Before continuing, it may help to relate this generalized coordinate system to the conventional polar and spherical coordinate system. For polar coordinates ($d = 2$),

$x_0 = r \sin(\phi_0)$ and $x_1 = r \cos(\phi_0)$, which is the typical convention for $(x, y) = (x_1, x_0)$. Likewise, for spherical coordinates, $x_0 = r \sin(\phi_0) \sin(\phi_1)$, $x_1 = r \cos(\phi_0) \sin(\phi_1)$, and $x_2 = r \cos(\phi_1)$; so $(x, y, z) = (x_1, x_0, x_2)$, ϕ_0 is the azimuthal angle, and ϕ_1 is the polar angle. The main change in convention is that, in both cases, the order of the x and y coordinates is switched.

A.3.2 Volume element

For this work, the set of unit vectors in \mathbb{R}^d is of interest. The set of unit vectors forms a $(d-1)$ -sphere, S^{d-1} , and can be formed by setting $r = 1$ and allowing all angles to vary. In order to integrate over unit vectors, the induced volume element in the generalized spherical coordinates is of interest. As will be proven here, the volume element in the coordinate system defined in Sec. A.3.1 is

$$\begin{aligned} dx_0 \cdots dx_{d-1} &= r^{d-1} dr d\Omega^{d-1} \\ \text{for } d\Omega^{d-1} &= \prod_{j=1}^{d-1} (\sin^{j-1}(\varphi_j) d\varphi_j). \end{aligned} \quad (\text{A.12})$$

Proof. The typical change-of-coordinate rules is

$$dx_0 \cdots dx_{d-1} = \left| \frac{\partial(x_0, \dots, x_{d-1})}{\partial(r, \varphi_1, \dots, \varphi_{d-1})} \right| dr d\varphi_1 \cdots d\varphi_{d-1}.$$

With this in mind, consider the Jacobian matrix inductively (by dimension),

$$\begin{aligned} J^{(d)} &= \frac{\partial(x_0^{(d)}, \dots, x_{d-1}^{(d)})}{\partial(r, \varphi_1, \dots, \varphi_{d-1})} \\ &= \left[\begin{array}{c|ccc} J^{(d-1)} \sin(\varphi_{d-1}) & x_0^{(d-1)} \cos(\varphi_{d-1}) & & \\ & \vdots & & \\ & x_{d-2}^{(d-1)} \cos(\varphi_{d-1}) & & \\ \hline \cos(\varphi_{d-1}) & 0 & \cdots & 0 \end{array} \right] \begin{array}{c} \\ \\ \\ \\ \hline -r \sin(\varphi_{d-1}) \end{array}. \end{aligned} \quad (\text{A.13})$$

Now, the determinant of the Jacobian matrix is

$$\begin{aligned} \det J^{(d)} &= -r \det \left(J^{(d-1)} \right) \sin^d(\varphi_{d-1}) \\ &\quad + (-1)^{d-1} \cos(\varphi_{d-1}) \\ &\quad \times \det \left[\begin{array}{c|ccc} \frac{\partial(x_1^{(d-1)}, \dots, x_{d-2}^{(d-1)})}{\partial(\varphi_1, \dots, \varphi_{d-2})} \sin(\varphi_{d-1}) & x_0^{(d-1)} \cos(\varphi_{d-2}) & & \\ & \vdots & & \\ & x_{d-2}^{(d-1)} \cos(\varphi_{d-2}) & & \end{array} \right]. \end{aligned}$$

Focus on the remaining determinant of the $(d-1) \times (d-1)$ matrix,

$$\begin{aligned}
\det(\dots) &= (-1)^{d-2} \det \left[\begin{array}{c|c} r \frac{x_0^{(d-1)}}{r} \cos(\varphi_{d-2}) & \frac{\partial(x_1^{(d-1)}, \dots, x_{d-2}^{(d-1)})}{\partial(\varphi_1, \dots, \varphi_{d-2})} \sin(\varphi_{d-1}) \\ \vdots & \\ r \frac{x_{d-2}^{(d-1)}}{r} \cos(\varphi_{d-2}) & \end{array} \right] \\
&= (-1)^{d-2} \det \left[\begin{array}{cc} r \cos(\varphi_{d-1}) & 0 \\ 0 & \sin(\varphi_{d-1}) \\ & \ddots \\ & \sin(\varphi_{d-1}) \end{array} \right] \\
&\quad \times \det \left[\begin{array}{c} \frac{x_0^{(d-1)}}{r} \\ \vdots \\ \frac{x_{d-2}^{(d-1)}}{r} \end{array} \quad \frac{\partial(x_1^{(d-1)}, \dots, x_{d-2}^{(d-1)})}{\partial(\varphi_1, \dots, \varphi_{d-2})} \right] \\
&= (-1)^{d-2} r \cos(\varphi_{d-1}) \sin^{d-2}(\varphi_{d-1}) \det J^{(d-1)},
\end{aligned}$$

where the second line used the fact that the determinant of a product is the product of determinants, $\det AB = \det A \cdot \det B$. Together,

$$\begin{aligned}
\det J^{(d)} &= -r \det \left(J^{(d-1)} \right) \sin^d(\varphi_{d-1}) - r \cos^2(\varphi_{d-1}) \sin^{d-2}(\varphi_{d-1}) \det \left(J^{(d-1)} \right) \\
&= -r \sin^{d-2}(\varphi_{d-1}) \det \left(J^{(d-1)} \right).
\end{aligned}$$

The initial condition for this inductive expression is $\det J^{(2)} = r dr d\varphi_0$. Thus,

$$\begin{aligned}
dx_0 \cdots dx_{d-1} &= r^{d-1} dr \prod_{j=1}^{d-1} (\sin^{j-1}(\varphi_j) d\varphi_j) \\
&= r^{d-1} dr d\Omega^{d-1},
\end{aligned} \tag{A.14}$$

where

$$\begin{aligned}
d\Omega^{d-1} &= \prod_{j=1}^{d-1} (\sin^{j-1}(\varphi_j) d\varphi_j) \\
&= \sin^{d-2}(\varphi_{d-1}) d\varphi_{d-1} d\Omega^{d-2}
\end{aligned} \tag{A.15}$$

is the surface element on S^{d-1} . \square

Finally, it is useful to define the surface area of S^d as

$$A_d := \int_{S^d} d\Omega^d. \tag{A.16}$$

An exact equation for this integral can be determined (e.g., by considering the integral over a $(d+1)$ -dimensional Gaussian) but will not be needed.

A.3.3 Directional sensitivity

This section will use the coordinate system described in the previous sections to prove Eq. (A.9).

Proof. The sensitivity of GNOME is defined in Chapter 5. Signals in the GNOME network are affected by the directional sensitivity of the magnetometers in the network. A signal is described by a $d = 3$ -dimensional vector \mathbf{m} and the sensitive axis of the i^{th} magnetometer is \mathbf{d}_i . The magnitude $\|\mathbf{m}\|$ describes the strength of the wall, while the magnitude of \mathbf{d} can be chosen to account for the particular coupling of the sensor to the signal. Let D be a $n \times d$ matrix, where n is the number of sensors in the network, whose rows are the $\{\mathbf{d}_i\}$ vectors. The set of signals $\{s_i\}$ from the sensors is described by

$$D\mathbf{m} = \mathbf{s}. \quad (\text{A.17})$$

Denote the uncertainty in the n sensors in the network with the covariance matrix $\Sigma = \text{diag}\{\sigma_i^2\}$. The sensitivity of the network is defined in Eq. (5.1) and Eq. (5.4) as the strength of a signal needed to guarantee a signal-to-noise ratio of 1. The sensitivity is then

$$\beta = 1/\sqrt{\hat{\mathbf{m}}^T D^T \Sigma^{-1} D \hat{\mathbf{m}}}. \quad (\text{A.18})$$

To better understand how the sensitivity scales with the size of the array, it helps to make some simplifying assumptions. First, assume that all the sensors have the same noise $\Sigma = \sigma^2 \mathbb{1}$ and $\|\mathbf{d}_i\| = \kappa$. Further, consider the vectors \mathbf{d}_i as evenly distributed along S^{d-1} . The expected squared sensitivity is then

$$\begin{aligned} \beta^2 &\approx \sigma^2 / \langle \|D\hat{\mathbf{m}}\|^2 \rangle \\ &= \sigma^2 / \left\langle \sum_{j=1}^n \kappa^2 \cos^2 \theta_j \right\rangle \\ &= \frac{\sigma^2}{n\kappa^2 \langle \cos^2 \theta \rangle}, \end{aligned}$$

where θ_i is the angle between \mathbf{m} and \mathbf{d}_i .

To calculate the expectation value of $\cos \theta$, one can orient x_{d-1} in the direction $\hat{\mathbf{m}}$ so that $\varphi_{d-1} = \theta$. The expected value is obtained by integrating against the PDF $\frac{d\Omega^{d-1}}{A_{d-1}}$ over the sphere,

$$\begin{aligned} \langle \cos^2(\varphi_{d-1}) \rangle &= \frac{1}{A_{d-1}} \int_{S^{d-1}} d\Omega^{d-1} \cos^2(\varphi_{d-1}) \\ &= \frac{1}{A_{d-1}} \int_{S^{d-1}} d\Omega^{d-2} d\varphi_{d-1} \sin^{d-2}(\varphi_{d-1}) \cos^2(\varphi_{d-1}). \end{aligned}$$

Focus now on the integral over φ_{d-1} (dropping the subscript),

$$\begin{aligned} \int_0^\pi d\varphi \sin^{d-2}(\varphi) \cos^2(\varphi) &= \int_0^\pi d\varphi \left[\frac{d}{d\varphi} \left(\frac{\sin^{d-1}(\varphi)}{d-1} \cos(\varphi) \right) + \frac{\sin^d(\varphi)}{d-1} \right] \\ &= \int_0^\pi d\varphi \frac{\sin^{d-2}(\varphi)}{d-1} (1 - \cos^2(\varphi)) \\ \frac{d}{d-1} \int_0^\pi d\varphi \sin^{d-2} \varphi \cos^2 \varphi &= \frac{1}{d-1} \int_0^\pi d\varphi \sin^{d-2} \varphi \\ \int_0^\pi d\varphi \sin^{d-2} \varphi \cos^2 \varphi &= \frac{1}{d} \int_0^\pi d\varphi \sin^{d-2} \varphi. \end{aligned}$$

Plugging this back into the equation,

$$\begin{aligned} \langle \cos^2(\varphi_{d-1}) \rangle &= \frac{1}{dA_{d-1}} \int_{S^{d-1}} d\Omega^{d-2} d\varphi_{d-1} \sin^{d-2}(\varphi_{d-1}) \\ &= \frac{1}{dA_{d-1}} \int_{S^{d-1}} d\Omega^{d-1} \\ &= 1/d. \end{aligned}$$

Finally, the sensitivity is

$$\beta \approx \frac{\sigma}{\kappa} \sqrt{\frac{d}{n}}. \quad (\text{A.19})$$

One can absorb κ into the noise σ or set $\kappa = 1$ to get Eq. (A.9).

This is the typical \sqrt{n} improvement that one often finds. The dependence on the dimension d can be understood as one needing a sensor pointing in every dimension to achieve the same sensitivity as a single sensor, except in all directions. \square

Appendix B

Manipulating probability distributions

A probability distribution function (PDF) can be thought of as a differential n -form¹ $P(\mathbf{x})d^n\mathbf{x}$ over an n -dimensional oriented manifold \mathcal{M} such that $\int_{\mathcal{M}} P(\mathbf{x})d^n\mathbf{x} = 1$ and whose integral over any n -dimensional submanifold is non-negative (with appropriate orientation). Moreover, for an n -dimensional submanifold $\mathcal{U} \subset \mathcal{M}$, the integral $\int_{\mathcal{U}} P(\mathbf{x})d^n\mathbf{x}$ is interpreted as the probability that one obtains a result in \mathcal{U} when randomly sampling from \mathcal{M} .

It is useful to understand how to manipulate and combine probability distributions. This has physical applications such as understanding the shape of spectral lines [166] and the standard halo model [125]. Additionally, there are practical applications such as when one has a random number from a simple PDF and would like a random number sampled from a different PDF.

One simple type of transformation is to consider the surjective map $\psi : \mathcal{M} \rightarrow \mathcal{N}$ for n -dimensional manifolds \mathcal{M} and \mathcal{N} where \mathcal{M} has the PDF, $P_0(\mathbf{x})d^n\mathbf{x}$. One can understand the PDF $P_1(\mathbf{y})d^n\mathbf{y}$ of $\mathbf{y} = \psi(\mathbf{x})$ by understanding that the probability that \mathbf{x} is in the submanifold $\mathcal{U} \subset \mathcal{M}$ is the same as in the image² $\psi(\mathcal{U}) \subset \mathcal{N}$,

$$\int_{\mathcal{U}} P_0(\mathbf{x})d^n\mathbf{x} = \int_{\psi(\mathcal{U})} P_1(\mathbf{y})d^n\mathbf{y}. \quad (\text{B.1})$$

In this appendix, methods for transforming probability distributions will be explored. First, in Sec. B.1, a method for transforming a known distribution to a desired distribution will be explored. Then, in Sec. B.2, the effects of transforming one or more random variables is considered.

Before continuing, it is important to note that focus will be placed on PDFs, which are defined for a manifold \mathcal{M} . However, it is often important to consider discrete,

¹ $d^n\mathbf{x} \equiv dx^1 \wedge dx^2 \wedge \dots \wedge dx^n$ with appropriate parameterization.

²If ψ is not injective, then one may need to double-count the repeated points in $\psi(\mathcal{U})$. For example, if $\mathcal{M} = \mathbb{R}$, $\mathcal{N} = \mathbb{R}^+$, and $\psi(x) = |x|$, Eq. (B.1), setting $\mathcal{U} = [a, b]$ for $a < 0$ and $b > 0$, would read as $\int_a^b P_0(x)d^n\mathbf{x} = \left(\int_0^{-a} + \int_0^b\right) P_1(\mathbf{y})d^n\mathbf{y}$.

Probability Mass Functions (PMFs) on a set M . PMFs are described by a function $Q : M \rightarrow \mathbb{R}$ such that $Q(k)$ is the probability of selecting k . Some examples of PMFs include the Binomial distribution and Poisson distribution. Comparing transformations of PMFs to PDFs often amounts to simply switching sums and integrals. One could also consider an injective function $\phi : M \rightarrow \mathcal{M}$ from a set M to a manifold \mathcal{M} . A PMF Q on M would then become a PDF $P(\mathbf{x})d^m\mathbf{x}$ on \mathcal{M} where

$$P(\mathbf{x}) = \sum_{k \in M} Q(k) \delta(\mathbf{x} - \phi(k)) d^m\mathbf{x},$$

and $\delta(\mathbf{x})$ is the m -dimensional Dirac delta function. Further, transformations between PMFs, the function $f : M \rightarrow N$ becomes the function $\psi : \mathcal{M} \rightarrow \mathcal{N}$ between manifolds, such that the following diagram commutes

$$\begin{array}{ccc} M & \xrightarrow{f} & N \\ \downarrow \phi_M & & \downarrow \phi_N \\ \mathcal{M} & \xrightarrow{\psi} & \mathcal{N} \end{array}$$

where $\phi_M : M \rightarrow \mathcal{M}$ and $\phi_N : N \rightarrow \mathcal{N}$ are injective.

B.1 Transforming the distribution

In this section, we will consider the case in which an initial distribution is known and there is a different desired distribution. This can be particularly useful when building Monte Carlo simulations in which tools are available to produce a random number from a known distribution (often a flat distribution on $[0, 1)$) but a different distribution is desired.

In general, the problem of transforming between known distributions $P_0(\mathbf{x})d^m\mathbf{x}$ on \mathcal{M} and $P_1(\mathbf{y})d^n\mathbf{y}$ on \mathcal{N} amounts to finding the appropriate function $\psi(\mathbf{x}) = \mathbf{y}$ such that Eq. (B.1) is satisfied. This will be described for the $n = 1$ case because the 1-dimensional topologies are easily characterized and the regions \mathcal{U} can be easily generalized.

For $n = 1$ -dimension, transforming between different probability distributions can be accomplished in generality³. To see this, it is useful to consider $\mathcal{M} = [0, 1)$ with a flat distribution dx (i.e., $P_0(x) = 1$). Now, consider some other distribution $P_1(y)dy$ on $\mathcal{N} = (a, b)$. Let $\psi : \mathcal{M} \rightarrow \mathcal{N}$ be a monotonic function with $\lim_{x \rightarrow 0^+} \psi(x) = a$. Define the cumulative distribution function (CDF) as

$$C_1(y) = \int_a^y P_1(\bar{y}) d\bar{y}. \quad (\text{B.2})$$

³Note that, because S^1 and \mathbb{R} differ topologically by the inclusion/removal of a set of measure zero, one can show that a distribution on any real interval can be extended to any 1-dimensional topology. Thus, only a real interval will be considered here.

Consider the integral over the family of intervals $\mathcal{U}_x = [0, x) \subset \mathcal{M}$ according to Eq. (B.1),

$$\int_0^x P_0(\bar{x})d\bar{x} = \int_a^{\psi(x)} P_1(\bar{y})d\bar{y},$$

$$x = C_1(\psi(x)).$$

From this, one obtains the PDF transformations,

$$y = C_1^{-1}(x) \quad \text{and} \quad x = C_1(y), \quad (\text{B.3})$$

to transform from dx to $P_1(y)dy$ and vice-versa, respectively.

Strictly speaking, it remains to be seen that Eq. (B.1) is satisfied for all submanifolds of $[0, 1)$ using the transformation from Eq. (B.3). It suffices to show that Eq. (B.1) is satisfied for any interval $[x_0, x)$ since all submanifolds will be the union of such intervals. This is accomplished by considering $[x_0, x) = \mathcal{U}_x \cap (\mathcal{M} - \mathcal{U}_{x_0})$:

$$\begin{aligned} \int_{x_0}^x P_0(\bar{x})d\bar{x} &= \int_0^x P_0(\bar{x})d\bar{x} - \int_0^{x_0} P_0(\bar{x})d\bar{x} \\ &= \int_a^{\psi(x)} P_1(\bar{y})d\bar{y} - \int_a^{\psi(x_0)} P_1(\bar{y})d\bar{y} \\ &= \int_{\psi(x_0)}^{\psi(x)} P_1(\bar{y})d\bar{y}. \end{aligned}$$

Thus, the proposed transformation works for all subsets \mathcal{U} .

This procedure can be extended to transform any two 1-dimensional PDFs by first transforming to $[0, 1)$ then any other PDF. For example, consider the transformation from any $P_0(x)dx$ on \mathcal{M} to $P_1(y)dy$ on \mathcal{N} can be accomplished by

$$y = C_1^{-1}(C_0(x)), \quad (\text{B.4})$$

where $C_0(x)$ is the CDF for $P_0(x)dx$ and $C_1(y)$ is the CDF for $P_1(y)dy$.

To illustrate an $n = 2$ -dimensional case, a specific example will be considered. Consider the initial distribution is flat on the square $\mathcal{M} = [0, 1) \times [0, 1)$, $P_0(\alpha, \beta)d\alpha d\beta = d\alpha d\beta$. We will consider a distribution on a triangular region \mathcal{N} bounded by $y = \pm x$ and $x = 1$ given by

$$P_1(x, y)dx dy = \left(1 + \frac{y}{x}\right) dx dy.$$

The goal, then, is to find a surjective map $\psi(\alpha, \beta) = (x(\alpha, \beta), y(\alpha, \beta))$. There are many ways to map a square region to a triangular region. To begin, we can make some choices to constrain the map. First, we will consider a map that maps together the boundaries $\partial\mathcal{M} \mapsto \partial\mathcal{N}$ and that is bijective on the interiors of these regions. Further, one can demand that $x = x(\alpha)$ with $x(0) = 0$ and $x(1) = 1$. Likewise, let $y(\alpha, \beta)$ be monotonically increasing in β so $y(\alpha, 0) = -x(\alpha)$ and $y(\alpha, 1) = x(\alpha)$. Then, $x(\alpha)$ can




Figure B.1: Visualization of $n = 2$ -dimensional transformation of (a) a flat distribution $P_0(\alpha, \beta)d\alpha d\beta$ to a (b) non-flat distribution $P_1(x, y)dx dy$ described in the text. The shading indicates the value of $P_1(x, y)$.

be defined by integrating over β ,

$$\begin{aligned} \int_0^\alpha d\bar{\alpha} \int_0^1 d\bar{\beta} &= \int_0^{x(\alpha)} d\bar{x} \int_{-\bar{x}}^{\bar{x}} d\bar{y} \left(1 + \frac{\bar{y}}{\bar{x}}\right) \\ \alpha &= \int_0^{x(\alpha)} d\bar{x} 2\bar{x} \\ &= x^2(\alpha), \end{aligned}$$

so $x(\alpha) = \sqrt{\alpha}$. For y , one can add a further constraint so that $y(\alpha, \beta) = f(\beta)\sqrt{\alpha} = f(\beta)x(\alpha)$, so that lines of constant β in \mathcal{M} form lines through the origin in \mathcal{N} . Then,

$$\begin{aligned} \int_0^\alpha d\bar{\alpha} \int_0^\beta d\bar{\beta} &= \int_0^{\sqrt{\alpha}} d\bar{x} \int_{-\bar{x}}^{f(\beta)\bar{x}} d\bar{y} \left(1 + \frac{\bar{y}}{\bar{x}}\right) \\ \alpha\beta &= \int_0^{\sqrt{\alpha}} d\bar{x} \frac{1}{2} (1 + f(\beta))^2 \\ \beta &= \left(\frac{1 + f(\beta)}{2}\right)^2 \\ f(\beta) &= 2\sqrt{\beta} - 1. \end{aligned}$$

Taken together,

$$\begin{aligned}\psi(\alpha, \beta) &= \left(\sqrt{\alpha}, (2\sqrt{\beta} - 1)\sqrt{\alpha} \right), \\ \psi^{-1}(x, y) &= \left(x^2, \frac{1}{4} \left(1 + \frac{y}{x} \right)^2 \right).\end{aligned}$$

A visualization of this transformation is given in Fig. B.1.

B.2 Combining random values

In the previous section, the problem was to generate a desired distribution $P_1(\mathbf{y})d^m\mathbf{y}$ from a given distribution $P_0(\mathbf{x})d^n\mathbf{x}$ by finding the appropriate transformation $\psi(\mathbf{x}) = \mathbf{y}$. In this section, we consider the case in which a transformation is known and the final probability distribution is constructed from this. For this, it is possible to generalize the problem slightly more such that the function $\psi : \mathcal{M} \rightarrow \mathcal{N}$ is a surjective map from an m -dimensional space \mathcal{M} to an n -dimensional space \mathcal{N} for $n \leq m$. This generalization allows one to consider the case of combining independent random variables $x_0 \in \mathcal{M}_0$ from $P_{00}(x_0)dx_0$ and $x_1 \in \mathcal{M}_1$ from $P_{01}(x_1)dx_1$ by imagining a them as a single variable $(x_0, x_1) \in \mathcal{M}_0 \times \mathcal{M}_1$ with the distribution $P_0(\mathbf{x})d^2\mathbf{x} = P_{00}(x_0)P_{01}(x_1)dx_0dx_1$. These variables are then combined via a function $\psi : \mathcal{M}_0 \times \mathcal{M}_1 \rightarrow \mathcal{N}$ into a single variable y with distribution $P_1(y)dy$.

To begin, the case where the dimensions of the two spaces are the same $n = m$. Generating a random number with a known distribution and then applying some transformation to the output will yield a new distribution. Given the PDF $P_0(\mathbf{x})d^n\mathbf{x}$ and an invertible transformation $\psi(\mathbf{x}) = \mathbf{y}$, the distribution $P_1(\mathbf{y})d^n\mathbf{y}$ would simply follow the change-of-coordinates,

$$P_1(\mathbf{y}) = P_0(\psi(\mathbf{y})) \left| \frac{\partial\psi^{-1}(\mathbf{y})}{\partial\mathbf{y}} \right|, \quad (\text{B.5})$$

where $\left| \frac{\partial\psi^{-1}(\mathbf{y})}{\partial\mathbf{y}} \right|$ is the Jacobian for ψ^{-1} . If ψ is piecewise-invertible, one defines a partition $\{\mathcal{U}_i\}$ of \mathcal{M} such that the restriction $\psi_i = \psi|_{\mathcal{U}_i} : \mathcal{U}_i \rightarrow \mathcal{V}_i \subset \mathcal{N}$ is invertible. Then Eq. (B.5) includes a sum,

$$P_1(\mathbf{y}) = \sum_{i, \mathbf{y} \in \mathcal{V}_i} P_0(\psi_i(\mathbf{y})) \left| \frac{\partial\psi_i^{-1}(\mathbf{y})}{\partial\mathbf{y}} \right|. \quad (\text{B.6})$$

Further generalizations are also possible but will not be explored in detail here. For instance, if ψ is constant for some region, the distribution $P_1(\mathbf{y})d^n\mathbf{y}$ will include some values with discrete probabilities. This can be accomplished with including a Dirac delta function in $P_1(\mathbf{y})$.

Now, consider the case where the dimension of the target space is less than the initial space $n < m$. It follows that ψ cannot be an injective function so one must integrate over a $(m-n)$ -dimensional space given by $\psi^{-1}(\mathbf{y})$. The idea then is to reparameterize \mathbf{x} such

into the pair (\mathbf{y}, \mathbf{t}) where \mathbf{t} parameterizes the submanifold $\psi^{-1}(\mathbf{y})$. That is, $\psi(\phi(\mathbf{y}, \mathbf{t})) = \mathbf{y}$ for the bijective function $\phi(\mathbf{y}, \mathbf{t}) \mapsto \mathbf{x}$. Given this function, the resulting PDF is

$$P_1(\mathbf{y})d^n\mathbf{y} = \left[\int_{\psi^{-1}(\mathbf{y})} \psi(\phi(\mathbf{y}, \mathbf{t})) \left| \frac{\partial\phi(\mathbf{y}, \mathbf{t})}{\partial(\mathbf{y}, \mathbf{t})} \right| d^{m-n}\mathbf{t} \right] d^n\mathbf{y}. \quad (\text{B.7})$$

Comparing Eq. (B.5) to Eq. (B.7), one sees that ϕ acts like an inverse of ψ after extending \mathcal{N} to the m -dimensional manifold $\mathcal{N} \times \psi^{-1}(\mathbf{y})$.

Consider the example of adding up two random numbers x from $P_x(x)dx$ and y from $P_y(y)dy$. The result will be $z = \psi(x, y) = x + y$ from an unknown distribution $P_z(z)dz$. For this case, $\phi(z, t) = (t, z - t)$ will satisfy $z = \psi(\phi(z, t))$. The Jacobian determinant is $\left| \frac{\partial\phi(\mathbf{y}, \mathbf{t})}{\partial(\mathbf{y}, \mathbf{t})} \right| = 1$, so the final distribution is described by

$$P_z(z) = \int P_x(t)P_y(z - t)dt = (P_x * P_y)(z),$$

which is simply a convolution of P_x and P_y .

Appendix C

Confidence and significance in Poissonian statistics

There are numerous cases in which one counts the number of instances of something and would like to either make a statement on how common that thing is or if that thing is occurring more/less than expected. For example, a Geiger counter can be used to measure the rate/intensity of ionizing radiation. One can then use this information to determine the rate at which a radioactive substance decays or compare the expected decay rate to the observed one. The mathematics used to meaningfully understand these counts and rates is Poissonian statistics.

For the sake of this appendix, the “thing” being counted will be called an event, without distinguishing whether an “event” is due to background or the thing one is trying to measure, because a false-positive event is, by its nature, indistinguishable from a true-positive event. Poissonian statistics assumes that

- the events are countable and one can properly distinguish events (no double-counting or missed events).
- the presence of one event does not affect the presence/absence of another. If the measurements occur over time, for example, one would not expect events to be clustered together.
- there is some predictable, inherit expected number of events in a given study. This condition could fail, for example, if the event rate varied over time in an unpredictable fashion.

This appendix will consider the statistics underlying such situations. A particular focus will be placed on understanding experiments in which the underlying rate is unknown and not predicted by theory. This will have direct application to GNOME wherein we search for domain-wall-crossing events and would like to understand the rate of events, especially when compared to an expected rate of background events.

C.1 Poissonian distribution

Fundamental to Poissonian statistics is understanding a probability distribution of possible experimental outputs. Such a distribution will be parametrized by a single parameter: the expected event count μ , which can be any positive real number.

Before describing the Poissonian probability distribution, consider the binomial distribution. This distribution describes the outcome of an experiment in which there are n chances of an event occurring, where there is a p probability of the event occurring in each chance. Though combinatoric arguments, the binomial distribution is given by the probability of k events occurring,

$$B(k; n, p) = \frac{n!}{k!(n-k)!} p^k (1-p)^{n-k}, \quad (\text{C.1})$$

this is the binomial Probability Mass Function (PMF). The mean of this distribution is $\mu = np$.

The Poissonian distribution is the limit of the binomial distribution as $n \rightarrow \infty$ while keeping the mean constant. For this, we replace $p = \mu/n$,

$$\begin{aligned} \lim_{n \rightarrow \infty} B(k; n, \mu/n) &= \lim_{n \rightarrow \infty} \frac{n!}{k!(n-k)!} (\mu/n)^k (1 - \mu/n)^{n-k} \\ &= \lim_{n \rightarrow \infty} \frac{\mu^k}{k!} \frac{n!}{(n-k)! n^k} (1 - \mu/n)^n \\ &= \frac{\mu^k}{k!} e^{-\mu}. \end{aligned}$$

Thus, the probability of measuring k events given that μ events are expected is given by the Poissonian PMF,

$$P(k; \mu) = \frac{\mu^k e^{-\mu}}{k!}. \quad (\text{C.2})$$

Because a common application of Poissonian statistics is in considering the rate r of events over some time T , one often finds Eq. (C.2) written with $\mu = rT$.

If μ is known, one can determine the significance of measuring n events by considering the probability of measuring more- and/or less-than n events. Here, the case in which the expected number of events is not known is considered. First, by establishing a confidence band for μ given a measurement. Then, a scheme for determining if two measurements are consistent will be described.

C.2 Confidence intervals

The main idea of building a confidence interval is that one would like to know what the expected number of events is, given that some number of events are measured. Because the PMF (Eq. (C.2)) is non-vanishing over the natural numbers regardless of μ , it is technically possible for μ to be anything. However, one would like to make a statement on a “reasonable” range of expected event counts within some confidence interval.




Figure C.1: Confidence intervals for $n = 5$ observed events and bounds given by $X = Y = \frac{1-C}{2}$.

It should also be noted that, for large μ , the probability of observing any given number of events becomes very small. Thus, it makes sense to build the bounds of the confidence interval using the probability that one measures at least as many or as few events as measured, as opposed to the likelihood of measuring exactly as many events as observed.

With these ideas in mind, the confidence interval is defined as the range of expected event counts such that the observed count falls near the center of the distribution. Given that n events were observed, define two thresholds X and Y for the bounds of the confidence interval (μ_L, μ_H) such that

$$X = \sum_{k=0}^n P(k; \mu_H) \quad \text{and} \quad Y = \sum_{k=n}^{\infty} P(k; \mu_L). \quad (\text{C.3})$$

Here, values for X and Y are generally chosen and the above equation defines μ_H and μ_L . For the first constraint, the upper bound μ_H is determined by finding the largest expected count such that the probability of observing n or fewer events is at least X . The second constraint analogously determines a lower bound¹. Together, the confidence level for this interval is given by

$$C = 1 - (X + Y). \quad (\text{C.4})$$

Keeping the confidence level constant still leaves some choice in how far the observed count can appear in the tails of the distribution. One could, for example, define the bounds such that $X = Y = \frac{1-C}{2}$ to get an interval around the observed count (e.g.,

¹There is some question as to whether one should include n in the tail (as in this case) or not. For example, the upper bound is defined such that the observed count would lie squarely in the tail (hence the interval being open to exclude the edge cases). As it is defined here, the interval may not vanish for $C \rightarrow 0$ (see, e.g., Fig. C.1).

Fig. C.1). Setting $(X, Y) = (0, 1 - C)$, the upper-bound disappears, leaving only a lower-bound, and similar for $(X, Y) = (1 - C, 0)$ defining an upper-bound. Also, if no events are observed, there is no lower-bound solution because $\sum_{k=0}^{\infty} P(k; \mu) = 1$, so it only makes sense to define the upper-bound at confidence level $C = 1 - X$.

C.3 Significance of a measurement

Consider running an experiment counting the number of events with the aim of determining if there is an excess of events with respect to some background (or other baseline). In many cases, there may not be a clear, underlying expected rate for the background; with the background defined by a combination of experimental design and uncontrollable factors. The background could be calculated by an ancillary experiment wherein the signals are removed or suppressed. In this case, the comparison will be between two values that contain some uncertainty: the “background measurement” n_0 and the “search measurement” n_1 (being the combination of background and foreground). The naïve approach is to simply approximate the two values as part of a Gaussian distribution with a \sqrt{n} standard deviation. This approach works particularly well for large counts as the Poissonian distribution looks Gaussian. However, a more rigorous method will be explored here that works in all cases.

One is often interested in comparing the rate of events and may use different durations for the background and search measurement. For example, one may not be able to take as much search-measurement data due to resource limitations. Thus, this section will allow this additional flexibility. With this in mind, the following quantities are defined,

$$\mu_0 = r_0 T_0, \quad (\text{C.5})$$

$$\mu_1 = r_1 T_1 = r_0 T_1 + r \tilde{T}_1, \quad (\text{C.6})$$

where μ_0 and μ_1 are the expected number of background- and search-measurement events, r_0 and r_1 are the underlying background- and search-measurement rates, T_0 and T_1 are the times used for the background- and search-measurement experiments, r is the underlying excess rate in the foreground, and \tilde{T}_1 is the time during which the search-measurement experiment is sensitive to foreground events. In general, the “excess rate” could be negative, and proving the existence of a significant excess will depend on excluding the negative case. Often, one simply considers conditions under which $\tilde{T}_1 = T_1$ so $r = r_1 - r_0$, however a bit more generality is included here to consider experiments in which there are times that certain events are missed (e.g., due to temporary reduction in sensitivity). With this in mind, define

$$\bar{r} = \frac{\tilde{T}_1}{T_1} r$$

to be the observed rate of events over the search-measurement experiment.

In an abstract view, the picture described here is of a theory with two parameters: the event rates r_0 and r_1 (likewise, the excess rate r can replace r_1). The theory states that

one would expect an observation to sample from a Poissonian PMF with the respective event rate. In an experiment, the times T_0 and T_1 are independent variables, while the measured number of events n_0 and n_1 are the dependent variables. Based on observation, one can define a “confidence region” in (r_0, r) parameter space at some confidence level. A visualization of the parameter space is given in Fig. C.2. The question, then, is to ask what combinations of r_0 and r can explain the observed rate within some confidence level. Since the background is not physically relevant, particular focus will be placed on defining bounds for the excess rate r given any background r_0 .

Using the methods described in Sec. C.2, the measurements n_0 and n_1 will define confidence intervals on μ_0 and μ_1 with confidence levels C_0 and C_1 , respectively. Likewise, dividing by the corresponding time yields confidence intervals for the rates r_0 and r_1 ,

$$\begin{aligned} r_{0L} < r_0 < r_{0H}, \\ r_{1L} < r_1 < r_{1H}. \end{aligned} \tag{C.7}$$

The confidence that both of these inequalities hold is simply

$$C = C_0 \cdot C_1. \tag{C.8}$$

Further, taking the difference in the two rate bounds (Eq. (C.7)) yields a bound on the observed excess rate,

$$-(r_{0H} - r_{1L}) < \bar{r} < r_{1H} - r_{0L}, \tag{C.9}$$

noting that $r_{1L} - r_{0H}$ will only be positive if there is a significant excess rate. This interval corresponds to a confidence of at least $C = C_0 C_1$. The exact confidence would be a bit higher to include cases in which the inequalities in Eq. (C.7) do not hold, but Eq. (C.9) holds.

Recall that the bound on \bar{r} can be translated in terms of r by multiplying each element in the inequality by T_1/\tilde{T}_1 , where T_1 is the total time used to measure μ_1 and \tilde{T}_1 is the time during which the search measurement is sensitive to signals. This is necessary to meaningfully interpret the rate.

Just as there is some choice in intervals for possible rates, there is some flexibility in how the confidence interval for \bar{r} is given. For example, an upper-bound on r_0 (so $r_{0L} \rightarrow 0$) and a lower-bound on r_1 (so $r_{1H} \rightarrow \infty$) leaves only the lower-bound

$$\bar{r} > -(r_{0H} - r_{1L}). \tag{C.10}$$

This extreme is particularly useful if one wants to claim that it is unlikely that there is no excess rate, because the presence of an excess would exclude the possibility that the search measurement can be explained by background, alone ($\bar{r} = 0$). The confidence level at which $r_{0H} = r_{1L}$ will describe the significance of a discovery. Alternatively, considering a lower-bound on r_0 (so $r_{0H} \rightarrow \infty$) and an upper-bound on r_1 (so $r_{1L} \rightarrow 0$) leaves only the upper-bound,

$$\bar{r} < r_{1H} - r_{0L}. \tag{C.11}$$




Figure C.2: A diagram of the theoretical rate parameter space with confidence intervals. (a) Parameter space in terms of (r_0, r_1) with confidence intervals in blue and red and \bar{r} in grey. (b) Parameter space in terms of (r_0, \bar{r}) with r_1 in grey. The horizontal black dashed lines represent the bounds on the excess rate \bar{r} .

This bound is particularly useful when there was not a significant excess rate. In this case, the bound can be used to exclude scenarios in which a large excess rate is expected. Keeping the confidence level constant and considering other (connected) intervals on the measurements will result in a confidence interval/band on r contained by the two above extreme cases.