

Statistical Methods for a First MADMAX Axion Dark Matter Search and Beyond

Johannes H. B. Diehl



Technische Universität München
TUM School of Natural Sciences

Statistical Methods for a First MADMAX Axion Dark Matter Search and Beyond

Johannes H. B. Diehl

Vollständiger Abdruck der von der TUM School of Natural Sciences der Technischen Universität München zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften genehmigten Dissertation.

Vorsitz: Prof. Dr. Björn Garbrecht

Prüfende der Dissertation

1. Priv-Doz. Dr. Béla Majorovits
2. Prof. Dr. Elisa Resconi

Die Dissertation wurde am 29.08.2024 bei der Technischen Universität München eingereicht und durch die TUM School of Natural Sciences am 26.09.2024 angenommen.

Not from the stars do I my judgement pluck

– William Shakespeare, *Sonnet 14*

Abstract

The axion is a light, pseudoscalar boson holding the potential to resolve two major questions in fundamental physics: The unexpected absence of charge-parity violating strong interactions and the fundamental nature of dark matter. Most popular theoretical models of axions fall into the two major categories of DFSZ-type models, which incorporate further Higgs doublets, and KSVZ-type models, which add heavy quarks.

Recent years have witnessed a surge in experimental efforts to detect axions. Among them is the Magnetized Disk and Mirror Axion Experiment (MADMAX). MADMAX uses a dielectric haloscope design, featuring a stack of parallel, dielectric disks and a mirror to resonantly enhance conversion of axions into photons within a strong magnetic field. In thesis the frequentist analysis of a first axion dark matter search conducted with a MADMAX prototype, focusing on axion masses of $76.7 \mu\text{eV}$ and $79.4 \mu\text{eV}$, is presented. Finding no evidence for axions, the analysis sets new world-leading upper limits on the photon coupling of dark matter axions in these mass ranges.

Going beyond proof-of-principle measurements, understanding the full parameter space is crucial. While photon couplings for minimal extensions of the standard model and KSVZ-type models are well-established, in this thesis the first systematic catalogue of photon couplings for DFSZ-type axions is introduced. With this catalogue, joint photon coupling distributions for both major categories of axion models can be constructed, providing valuable inputs to inform sensitivity requirements for frequentist inference or prior probabilities in Bayesian frameworks.

As axion research progresses towards detecting dark matter axions, accurate parameter inference becomes essential. Current methods bias signal parameter estimates in case of an axion detection. A modification to the signal model is proposed in this thesis that facilitates bias-free inference.

Zusammenfassung

Axionen sind leichte, pseudoskalare Bosonen, die das Potenzial haben, zwei zentrale Fragen der fundamentalen Physik zu beantworten: das unerwartete Fehlen von CP-verletzenden starken Wechselwirkungen und die grundlegende Natur der Dunklen Materie. Die bekanntesten theoretischen Modelle für Axionen lassen sich in zwei Kategorien einteilen. Man spricht von Modellen des DFSZ-Typs, wenn dem Standardmodell der Teilchenphysik weitere Higgs-Dubletts hinzugefügt werden, und von Modellen des KSVZ-Typs, welche zusätzliche schwere Quarks enthalten.

In den letzten Jahren wurde eine Vielzahl experimenteller Bemühungen vorangetrieben, Axionen nachzuweisen, zu welchen auch das “Magnetized Disk and Mirror Axion Experiment” (MADMAX) gehört. MADMAX nutzt als Design ein dielektrisches Haloskop, bei welchem mehrere parallele dielektrische Platten und ein Spiegel verwendet werden, um die Umwandlung von Axionen in Photonen in einem starken Magnetfeld resonant zu verstärken. In dieser Arbeit wird die frequentistische Analyse einer ersten Suche nach Dunkler Materie in Form von Axionen mit einem MADMAX-Prototyp präsentiert, die sich auf Axionmassen von rund $76.7 \mu\text{eV}$ und $79.4 \mu\text{eV}$ konzentriert. Da keine Hinweise auf Axionen gefunden wurden, setzt die Analyse neue, weltweit führende obere Grenzen für die Photonen-Kopplung von Dunkle-Materie-Axionen in diesen Massenbereichen.

Um über erste Proof-of-Principle-Messungen hinauszugehen ist ein gutes Verständnis des vollständigen Parameterraums entscheidend. Während Kopplungskonstanten zu Photonen für minimale Erweiterungen des Standardmodells und Modelle des KSVZ-Typs bereits etabliert sind, wird in dieser Arbeit der erste systematische Katalog von Photonen-Kopplungen für Axionen des DFSZ-Typs vorgestellt. Mit diesem Katalog können gemeinsame Verteilungen der Photonen-Kopplung für beide Kategorien von Axionmodellen erstellt werden, die wertvolle Informationen liefern, um die Anforderungen an die Sensitivität frequentistischer Analysen oder a priori Wahrscheinlichkeiten im Kontext von Bayesscher Inferenz zu bestimmen.

Mit zunehmendem Fortschritt der Axionenforschung in Richtung des Nachweises von Dunkle-Materie-Axionen wird eine genaue Ermittlung von Signalparametern wichtiger. Im Fall der Detektion eines Axion Signals verzerren momentan verbreitete Methoden die Inferenz der Signalparameter. In dieser Arbeit wird eine Modifikation des Signalmodells vorgeschlagen, die eine verzerrungsfreie Inferenz ermöglicht.

Acknowledgements

Without the unwavering support of many people this thesis would not have been possible. First and foremost, I am immensely grateful to my supervisor, Béla Majorovits, for allowing me the space to grow at my own pace and explore my own research interests. His valuable guidance and the freedom he provided for independent research were instrumental to this work. Additionally, I want to express my sincerest gratitude to Frank Steffen for his mentorship throughout this project. His insightful advice and numerous helpful discussions were invaluable in navigating the complexities of my research. I also extend my thanks to Björn Garbrecht for his valuable input as an external expert.

Science is always a collaborative effort, and I had the pleasure of working with many bright and like-minded individuals. I want to offer my deepest appreciation to Oliver Schulz for sharing his limitless knowledge of computing and statistics, and to Emmanouil Koutsangelas for guiding me through the intricacies of axion theory. I am grateful to them, Jakob Knollmüller, and the members of the MADMAX collaboration, including Juan Maldonado, Stefan Knirck, Jacob Egge, Alberto Martini, Bernardo Ary dos Santos, Vijay Dabhi, Haotian Wang, Alireza Kazemipour, Dagmar Kreikemeyer, and Heesu Byun, for the projects we undertook and the enjoyable times we shared. Special thanks go to David Leppla-Weber for building the MADcode with me, and to the Max-Planck workshop apprentices for constructing an axion pool table for me. Furthermore, I appreciate the brilliant ideas from Akira Miyazaki and the helpful discussions with him, Gia Dvali, Allen Caldwell, Anton Ivanov, Sebastian Hoof, and Vaisakh Plakkot. This thesis was improved by the proofreading and valuable feedback from Frank Steffen, Georg Raffelt, Allen Caldwell, Anton Ivanov and Emmanouil Koutsangelas, for which I am grateful.

I am also grateful for the opportunity to participate in the International Max Planck Research School on Elementary Particle Physics. Interacting with many young researchers from other fields of physics at seminars as well as the legendary workshops at Ringberg castle was an invaluable experience.

Finally, I am thankful for all the support I received from friends and family over the years. I specifically want to thank Tim Gerbig and my other brethren as well as Marco Dembecki and the rest of Harlon's crew. Most importantly, I want to thank my parents. Thank you for your endless support and the unconditional love you raised me with. Thank you for always believing in me, even when I doubted myself.

This is for you!

Publications

The author of this thesis prepared the following publications over the course of the project as corresponding author:

- [1] *DFSZ-Type Axions and Where to Find Them.*
J. Diehl and E. Koutsangelas
Phys. Rev. D 107 (2023) 095020. arXiv: 2302.04667
- [2] *Bias-free Estimation of Signals on Top of Unknown Backgrounds.*
J. Diehl, J. Knollmüller, and O. Schulz.
NIM A 1063 (2024) 169259. arXiv: 2306.17667
- [3] *First search for axion dark matter with a MADMAX prototype*
MADMAX collaboration
B. Ary dos Santos Garcia, J. Diehl, D. Leppla-Weber, A. Martini et al.
in preparation

Furthermore, the author contributed to the following papers:

- [4] *Simulating MADMAX in 3D: Requirements for Dielectric Axion Haloscopes.*
MADMAX collaboration
S. Knirck, J. Schütte-Engel et al.
JCAP 10 (2021) 034. arXiv: 2104.06553
- [5] *First search for dark photon dark matter with a MADMAX prototype*
MADMAX collaboration
J. Egge, D. Leppla-Weber, S. Knirck et al.
arXiv: 2408.02368

Consequently, there is some overlap in content between this thesis and the papers mentioned above. In particular sec. 3.3 uses a plot from [4] to which the author contributed. The content of ch. 5 closely follows [1], which was prepared by the author in close collaboration with Emmanouil Koutsangelas and is therefore also featured in his PhD thesis [6]. In addition, ch. 6 closely follows [2]. Text and figures used in this thesis from the aforementioned publications were originally created by the author.

Contents

Introduction	1
1 QCD Axion Theory	5
1.1 Strong CP Problem and Peccei-Quinn Solution	6
1.2 Axion Mass and Photon Coupling	8
1.3 Axion Models	10
2 Axions as Dark Matter	15
2.1 Axions and Inflation	16
2.2 Production Mechanisms	17
2.3 The QCD Axion Window	19
2.4 The Axion Lineshape	21
3 Detecting Dark Matter Axions with Dielectric Haloscopes	23
3.1 Axion Electrodynamics	24
3.2 Power Boost Factor	26
3.3 The MADMAX Experiment	29
4 Foundational Statistical Methods	37
4.1 Methods for Frequentist Hypothesis Testing	38
4.2 Correcting Look-Elsewhere Effect	39
4.3 Axion Priors for Bayesian Inference	40
5 Non-minimal DFSZ-Type Axion Models	43
5.1 Calculating the PQ Charges in DFSZ-Type Models	44
5.2 What Makes the Model?	47
5.3 Anomaly Ratio Distributions	50
5.4 Implications for Axion Searches	63
5.5 Summary & Discussion	68
6 Bias-Free Signal Estimation	71
6.1 General Approach	72
6.2 Application to Simulated MADMAX Data	74
6.3 Summary & Discussion	84
7 The MADMAX Pipeline for Statistical Analysis	85
7.1 Runs Gaussianity Test	87

7.2	Digital Notch Filter	92
7.3	Constructing the Grand Spectrum	99
7.4	Limit Setting	113
7.5	Grand Spectrum Excesses	114
7.6	Summary & Discussion	116
8	First Results from a MADMAX Axion DM Search	119
8.1	Details of the Data Taking	120
8.2	Data Processing	123
8.3	Axion Limits	134
8.4	Summary & Discussion	136
	Conclusion & Outlook	139
	List of Figures	143
	List of Tables	145
	Bibliography	147

Introduction

Almost one hundred years ago Fritz Zwicky uncovered a startling discrepancy between the velocity dispersion of the constituents of the Coma galaxy cluster and the gravitational attraction inferred from their visible matter. Zwicky’s observation hinted at the existence of a mysterious, unseen component of the universe. This revelation laid the groundwork for our current understanding that some form of dark matter must exist and even dominate the matter content of the universe. The fundamental nature of this component is one of the largest open questions to both astro- as well as particle physics.

The axion emerges as a compelling candidate in this context, though it was initially proposed not to solve this cosmic enigma, but to address a profound issue in quantum chromodynamics (QCD). In QCD, interactions are predicted to violate charge-parity (CP) symmetry, yet no such violations have been observed with high precision. By promoting the parameter responsible for CP violation to a field, it can be dynamically set to zero, giving birth to the axion. The underlying non-thermal production mechanism can lead to a cosmologically relevant, stable and cold population of these axions. They are therefore a prime dark matter candidate offering a solution to two of modern physics’ biggest questions.

The initial axion proposal was quickly ruled out experimentally due to its sizeable couplings to the standard model. If, however, additional particles are present in the theory, couplings to the standard model can be reduced so much that the term “invisible axion models” was coined in the early 1980s. They are divided into two major categories depending on whether additional heavy quarks are added (KSVZ-type models) or additional Higgs particles (DFSZ-type).

To guide experiments to finding dark matter axions, precise knowledge of their theoretical predictions is needed. In fact, the models can produce a significant axionic contribution to dark matter for a wide but finite range of axion masses. In terms of their coupling to photons, for which the most extensive experimental program exists, every theoretical model makes a concrete prediction for a given axion mass. Considering only the minimal set of additional particles leads to only three distinct possible photon couplings for each mass. When this assumption is dropped, however, the parameter space opens up significantly making several orders of magnitude smaller or larger photon couplings viable.

Beyond minimal extensions an exhaustive set of predictions so far was only constructed for KSVZ-type models. In this thesis for the first time an analogous set for DFSZ-type models is presented, making it possible to compare these two major categories of invisible axion models and creating a joint model catalogue. The findings reveal that DFSZ-type models yield predictions closely resembling those of the KSVZ-

type case with axion-photon couplings that can exceed those of minimal models by several orders of magnitude. This comprehensive catalogue is particularly valuable for Bayesian statistical analysis, where it can be associated with probabilities, providing a necessary input for inference.

Today the invisible models seem far less invisible than forty years ago. So called “haloscope” experiments are using axion-photon couplings to gain sensitivity to relevant parts of the invisible axion parameter space. In the presence of a strong magnetic field, dark matter axions with a specific mass convert into radio-frequency photons of a specific wavelength λ . The photon emission can be enhanced through the use of cavities, dielectric structures as in this thesis, or similar resonators.

In haloscope experiments the power per frequency interval coming from the resonator is measured. This power is dominated by background resulting from the response of the system and the characteristics of the receiver chain. An axion signal would be visible as a small excess in several adjacent frequency bins. To detect such a small excess the background must first be estimated and removed. The resulting spectrum of power excesses is then expected to follow a Gaussian distribution and must not contain contributions of non-axionic origin. If no significant axion signal is found in this spectrum, a limit can be set.

Modelling the background *ab initio* is exceedingly difficult due to λ being comparable in size to the components of the system. Instead, typically filters are used to estimate the background, which are affected by the presence of a signal, altering the signal shape in the residuals after background removal. This leads to biased inference of signal parameters in case of an axion detection. In this thesis a way to account for the effect of the background removal in the signal model is demonstrated. With only minimal assumptions on the filter, which are usually fulfilled, the signal parameters can be inferred in a bias-free way. The method could readily be added to any haloscope analysis.

A major challenge for haloscopes is tuning, especially towards larger axion masses: Any given haloscope setup is only resonant for a small range of photon wavelengths, i.e. axion masses. Since invisible axion models give only very broad predictions for the range of possible axion masses, a tuning mechanism is vital. Tuning can be achieved by altering some length scale in the system to match different photon wavelengths. Reaching smaller λ requires smaller systems. This is detrimental to sensitivity, which scales linearly with the volume available for axion-photon conversion, i.e. for standard cavity haloscopes with λ^3 .

One relatively recent variant of the haloscope approach specifically designed to be able to reach larger axion masses is the Magnetized Disk and Mirror Axion Experiment (MADMAX). By using multiple parallel dielectric disks and a mirror with adjustable separations instead of a cavity, the sensitivity becomes largely independent of λ . A first axion dark matter search with a MADMAX prototype setup has been conducted in early 2024. By arranging its three disks in two different configurations world-leading limits have been achieved at axion masses in the ranges of (76.56 to 76.82) μeV and (79.31 to 79.53) μeV . These limits were achieved due to methods presented in this thesis, which contains all necessary details concerning the statistical analysis of the data.

The thesis is structured as follows: In the first four chapters necessary prerequisites are covered. Ch. 1 starts with the theoretical foundation of the axion in QCD. It also

specifically introduces previous work cataloguing axion-photon coupling predictions for KSVZ-type axions. Ch. 2 presents the axion as a dark matter candidate leading to a preferred range of axion masses, in which it could constitute the majority of the local dark matter. It also covers the expected shape of an axion dark matter signal. In ch. 3 the working principle of haloscopes is introduced with a specific focus on the MADMAX setup used for the measurements analysed later in the thesis. Ch. 4 covers some well-established but crucially important statistical methods. It illuminates how to interpret the results of the analysis conducted later and why a full catalogue of invisible axion models is vital for a statistical analysis.

Chs. 5 to 8 are dedicated to original work by the author. In ch. 5 axion-photon couplings are systematically calculated for a large number of DFSZ-type models. The values obtained for minimal extensions are still preferred for general DFSZ-type models. Associating a probability with each model, it is possible to obtain theory predictions for the probability distribution of photon couplings of DFSZ-type models as one-sided 68% or 95% lower bounds or two-sided bands. These are comparable to the KSVZ-type case. A subset of theoretically preferred DFSZ-type models is analysed separately and displays enhanced photon couplings compared to the full set of models.

Ch. 6 presents an approach to infer signal parameters in scenarios where no parametric model exists for the background, such as in haloscope experiments. Considering the effect of the background removal as part of the signal model allows for a bias-free inference, which is demonstrated in the concrete example of simulated MADMAX-like datasets.

Chs. 7 and 8 focus on the concrete analysis of the first search for axion dark matter with a MADMAX prototype setup. Here, ch. 7 abstractly introduces original methods to ensure and test Gaussianity of the spectra after background removal. It presents a pipeline to translate the measurements into a competitive limit on the axion-photon coupling.

In ch. 8 this pipeline is then applied to data taken by the MADMAX collaboration during a measurement campaign in 2024 using the MORPURGO magnet at CERN. Systematic uncertainties are taken into account. No hints for an axion signal are found, therefore a 95% confidence limit is set around $76.7 \mu\text{eV}$ and $79.4 \mu\text{eV}$. Assuming a local axion dark matter density of $\rho_a = 0.3 \text{ GeV}/\text{cm}^3$ it exceeds a previous limit at these masses by up to a factor of four.

1

QCD Axion Theory

This chapter is meant to highlight the theoretical ingredients necessary to understand the role of the axion in an astrophysical context in ch. 2 as well as laying the foundation for the phenomenology project [1] presented in ch. 5. Readers interested in a very detailed introduction to the axion are referred to [7] and [6], on which most of this chapter is based.

The chapter starts by explaining why the curious absence of CP violation in strong interactions is a problem and how the most prominent solution to this conundrum leads to the axion (sec. 1.1). The focus of this thesis lies on obtaining theory predictions for the coupling of axions to photons as well as experimental explorations in the axion mass vs. photon coupling parameter plane. In sec. 1.2 therefore theory independent expressions for those two parameters are obtained. This section also introduces the domain wall number and shows how to calculate anomaly coefficients which are crucial for obtaining axion-photon couplings. The two most prominent categories of QCD axion models, DFSZ- and KSVZ-type axions, are introduced in sec. 1.3. Here, emphasis is put on detailing how to obtain values for the anomaly coefficients in the standard DFSZ case, since these considerations will be expanded to theories with more than two Higgs doublets in ch. 5. Complementary investigations have already been conducted for the KSVZ case by other groups [8, 9]. In order to prepare for a later comparison their results are outlined here.

1.1 Strong CP Problem and Peccei-Quinn Solution

The Lagrangian of quantum chromodynamics (QCD) contains two terms that potentially violate charge-parity (CP) symmetry:

$$\mathcal{L}_m = - \sum_q \bar{q} m_q e^{i\theta_q \gamma^5} q \quad (1.1)$$

$$\mathcal{L}_\theta = \frac{\theta g_s^2}{32\pi^2} \text{Tr} G^{\mu\nu} \tilde{G}_{\mu\nu} . \quad (1.2)$$

Here, q are the quark fields that can be written in chiral representation as $q = q_L + q_R$ and $\bar{q} = q^\dagger \gamma^0$ is their conjugate. m_q is the mass of the quark q which can in general have a complex phase θ_q . g_s is the strong coupling constant and θ is called the vacuum angle. The matrix γ^5 distinguishes between the chiral components of q : $\gamma^5 q_{L/R} = \mp q_{L/R}$. $G_{\mu\nu}$ is the strong field strength tensor and $\tilde{G}^{\mu\nu} = \frac{1}{2} \epsilon^{\mu\nu\rho\sigma} G_{\rho\sigma}$ is its dual. The trace runs over the indices of the color group $SU(3)_c$, denoted in the following by Latin indices a, b , and c .

The two terms (1.1) and (1.2) are fundamentally connected: Observable parameters must be invariant under field redefinitions. For instance performing a global chiral transformation on the quark fields

$$q \rightarrow e^{i\gamma^5 \alpha} q \quad (1.3)$$

transforms θ_q in (1.1) as $\theta_q \rightarrow \theta_q + 2\alpha$, but also $\theta \rightarrow \theta - 2\alpha$ in (1.2) because the path integral measure is not invariant [10]. Therefore only the linear combination $\theta_{\text{eff}} = \theta + \theta_q$ can be an observable and it is possible to shift the full physical effect either to the fermion sector or to the gluon sector. In the latter case:

$$\theta_{\text{eff}} = \theta + \sum_q \arg m_q . \quad (1.4)$$

Before interpreting this result, it is important to take another closer look at (1.2). Terms in the Lagrangian such as (1.2) can often be considered unphysical. It is possible to write $\text{Tr} G^{\mu\nu} \tilde{G}_{\mu\nu}$ as a total derivative. With regard to the action, which integrates over the Lagrangian, this leads to a boundary term, which usually has no influence on physical results from perturbation theory. However, things are different for eq. (1.2): The QCD vacuum is not described by a vanishing vector potential, but allows for configurations that are gauge transformations of zero:

$$A_\mu = i g_s^{-1} U^{-1} \partial_\mu U , \quad (1.5)$$

with $U(x)$ being an element of the QCD gauge group $SU(3)_c$.

If all possible U can be smoothly deformed to the identity, then the contribution to the action vanishes. However, it can be shown that this is not the case. The $U(x)$ form topologically distinct classes $U_n(x)$ with $n \in \mathbb{Z}$ being called winding number or Pontryagin index.

Two vector potentials $A_\mu^{(n)}$ and $A_\mu^{(m)}$ constructed from U with different winding numbers $n \neq m$ correspond to two different vacuum states $|n\rangle$ and $|m\rangle$. Transitioning from $|n\rangle$ to $|m\rangle$ requires intermediate non-vanishing field strengths, i.e. an energy

barrier exists between $|n\rangle$ and $|m\rangle$. This implies that tunnelling can occur and that the true, physical vacuum is a superposition of $|n\rangle$:

$$|\theta\rangle = \sum_n e^{-in\theta} |n\rangle , \quad (1.6)$$

where θ is an angular parameter. Including potential transitions between states of differing winding number in the path integral leads to the appearance of a term of the form (1.2) proportional to this θ parameter.

The key take-away from this discussion is that the θ -term needs to be included in the Lagrangian due to the structure of the QCD vacuum, which was first realised by 't Hooft [11, 12]. The QCD Lagrangian therefore contains a term that is potentially CP violating.

Such a CP violating term would induce electric dipole moments in baryons, on which there are strong experimental bounds. In particular there is a tight constraint on the dipole moment of the neutron (90% confidence level)[13]:

$$d_n < 1.8 \times 10^{-26} e \text{ cm} , \quad (1.7)$$

where e denotes the electron electric charge. A similar limit can be obtained from mercury atoms [14]. Various independent theoretical calculations based on the aforementioned Lagrangian predict $d_n \sim 10^{-16} \theta_{\text{eff}} e \text{ cm}$ [15–21], which sets an upper limit on θ_{eff} of

$$|\theta_{\text{eff}}| \lesssim 10^{-10} . \quad (1.8)$$

This is a fine tuning problem because of the two contributions in eq. (1.4) [22, 23]. Attempts have been made to explain the smallness of θ_{eff} imposing unconventional dynamics [24, 25] or spontaneously breaking the CP symmetry [26–30]. The most prominent solution, however is the Peccei-Quinn mechanism:

Peccei and Quinn realised that by introducing a global chiral $U(1)$ symmetry (henceforth $U(1)_{\text{PQ}}$) the CP-violating angle θ_{eff} can be made dynamical [31, 32], and thus made to relax to the CP-conserving value $\theta_{\text{eff}} = 0$ [33]. The $U(1)_{\text{PQ}}$ symmetry is spontaneously broken with the axion being the corresponding pseudo-Nambu-Goldstone boson [34, 35].

“Pseudo” in this context indicates that the axion is not a truly massless Goldstone boson due to the $U(1)_{\text{PQ}}$ symmetry being anomalous. This means that the symmetry holds at the level of the classical action, but is broken for the quantum theory because the path integral measure is not invariant. Noether’s theorem states that every continuous global symmetry leads to a conserved current. An anomaly coefficient is the constant of proportionality indicating how much the current is not conserved. For J_μ^{PQ} , the current associated to the $U(1)_{\text{PQ}}$ symmetry, specifically

$$\partial^\mu J_\mu^{\text{PQ}} = \frac{g_s^2 N}{16\pi^2} \text{Tr} G^{\mu\nu} \tilde{G}_{\mu\nu} + \frac{e^2 E}{16\pi^2} F^{\mu\nu} \tilde{F}_{\mu\nu} , \quad (1.9)$$

where N and E are called color and electromagnetic (EM) anomaly coefficient, respectively. $F^{\mu\nu}$ is the EM field strength tensor and $\tilde{F}_{\mu\nu}$ its dual.

1.2 Axion Mass and Photon Coupling

This section is meant to give some intuition concerning the model independent axion potential and mass as well as the coupling to photons. The calculations up to eq. (1.17) assume energies below the Peccei Quinn and the electroweak breaking scales and only include the two lightest quarks for simplicity.

From a model building perspective using effective field theory, a complex scalar field $\Phi \propto \phi(x) \exp\left[-i\frac{a(x)}{v_a}\right]$ is introduced with v_a being the order parameter of the spontaneous symmetry breaking potential. The angular field $a(x)$ has an effective Lagrangian

$$\mathcal{L}_a = \frac{1}{2}(\partial_\mu a)^2 + \mathcal{L}_{\text{int}}(\partial_\mu a, q) + \frac{g_s^2 N}{16\pi^2} \frac{a}{v_a} \text{Tr} G^{\mu\nu} \tilde{G}_{\mu\nu} + \frac{e^2 E}{16\pi^2} \frac{a}{v_a} F^{\mu\nu} \tilde{F}_{\mu\nu} - \bar{q}_L M_q q_R + \text{h.c.} . \quad (1.10)$$

Here the second term encodes potential interactions with quark fields q and $M_q = \text{diag}(m_u, m_d)$ is the quark mass matrix. The θ term (1.2) is already absorbed in the third term in (1.10) due to a field redefinition $a \rightarrow a - \theta_{\text{eff}} v_a / (2N)$. One can write

$$f_a = \frac{v_a}{2N} , \quad (1.11)$$

which is called axion decay constant due to the role it plays in \mathcal{L}_{int} .

The term proportional to $G\tilde{G}$ can be removed from the Lagrangian using a field-dependent axial transformation of the quark fields:

$$q \rightarrow e^{i\gamma_5 \frac{a}{2f_a} Q_a} q . \quad (1.12)$$

Here Q_a is a generic matrix that acts on the quark fields. The only requirement on Q_a is that $\text{Tr} Q_a = 1$ in order for this transformation to cancel the $G\tilde{G}$ term. The transformation (1.12) affects the coupling to photons, since the transformation is anomalous under QED, as well as the quark mass matrix:

$$\mathcal{L}_a \supset \frac{1}{4} g_{a\gamma} a F^{\mu\nu} \tilde{F}_{\mu\nu} - \bar{q}_L e^{i\frac{a}{2f_a} Q_a} M_q e^{i\frac{a}{2f_a} Q_a} q_R + \text{h.c.} , \quad (1.13)$$

where

$$g_{a\gamma} = \frac{\alpha_{\text{em}}}{2\pi f_a} \left[\frac{E}{N} - 6\text{Tr}(Q_a Q^2) \right] \quad (1.14)$$

with fine structure constant $\alpha_{\text{em}} = e^2/(4\pi)$ and $Q = \text{diag}(2/3, -1/3)$ encoding the EM charges of the quarks. E/N is called the anomaly ratio.

Due to confinement at energies below the QCD scale pions among others become effective degrees of freedom of QCD instead of quarks and gluons. Using chiral perturbation theory it is therefore possible to write an effective chiral Lagrangian for the axion in terms of interactions with the pions. These calculations lead to an axion-pion potential. At the ground state of the pion and choosing $Q_a = \frac{1}{2}\text{diag}(1, 1)$ it reads

$$V(a) = -m_\pi^2 f_\pi^2 \sqrt{1 - \frac{4m_u m_d}{(m_u + m_d)^2} \sin^2\left(\frac{a}{2f_a}\right)} , \quad (1.15)$$

with m_π and f_π pion mass and decay constant, respectively. A Taylor expansion around $\frac{a}{f_a} = 0$ yields a quadratic term in a , the prefactor of which is associated as the axion mass

$$m_a = \frac{\sqrt{m_u m_d}}{m_u + m_d} \frac{m_\pi f_\pi}{f_a}. \quad (1.16)$$

For the axion-photon coupling the choice of $Q_a = M_q^{-1}/\text{Tr } M_q^{-1}$ is more convenient, since otherwise axion-pion mixing would have to be considered. With this choice and eq. (1.14) $g_{a\gamma}$ follows as

$$g_{a\gamma} = \frac{\alpha_{\text{em}}}{2\pi f_a} \left(\frac{E}{N} - \frac{2}{3} \frac{4m_d + m_u}{m_d + m_u} \right). \quad (1.17)$$

Going beyond this leading order estimation, the authors of [36] find a precise numerical value for the axion mass of

$$m_a = 5.70(7) \mu\text{eV} \left(\frac{10^{12} \text{ GeV}}{f_a} \right), \quad (1.18)$$

where the digit in brackets denotes the error on the last digit. For the axion-photon coupling they find

$$g_{a\gamma} = \frac{\alpha_{\text{em}}}{2\pi f_a} \left[\frac{E}{N} - 1.92(4) \right]. \quad (1.19)$$

These numerical values will be used in the remainder of this thesis. For convenience in later chapters it is useful to already define

$$\mathcal{C}_{a\gamma} \equiv E/N - 1.92(4). \quad (1.20)$$

As an angular variable the axion field a is defined over the domain $[0, 2\pi v_a[$, but the QCD induced potential (1.15) has a periodicity of $[0, 2\pi f_a[$ with $f_a = v_a/(2N)$. This means the potential has $2N$ local minima in the domain of a prompting the definition

$$N_{\text{DW}} \equiv 2N, \quad (1.21)$$

which is called domain wall number. This definition is needed in ch. 2 and revisited more thoroughly later in sec. 5.1.

The crucial takeaways from (1.16) and (1.17) are: Firstly, the axion mass is model dependent only via its inverse proportionality to the axion decay constant. Secondly, the coupling to photons has a model independent part derived from the quark mass matrix and a model dependent part derived from the anomaly coefficients E and N .

The anomaly coefficients can be calculated using triangle Feynman diagrams with one external line being the axion and the other two photons or two gluons, respectively (see [37] for a general review). Fermions ϕ contributing to the internal loop must be charged under the PQ symmetry and reside in a specific representation of the $SU(3)_c \times SU(2)_L \times U(1)_Y$ gauge group, which can be denoted via $(\mathcal{C}_\phi, \mathcal{I}_\phi, \mathcal{Y}_\phi)$. Here \mathcal{C}_ϕ , \mathcal{I}_ϕ and \mathcal{Y}_ϕ encode the color and weak isospin representation as well as the hypercharge, respectively. Each irreducible representation ϕ contributes to the anomaly coefficients as

$$N = \sum_{\phi} N_{\phi}, \quad E = \sum_{\phi} E_{\phi}, \quad (1.22)$$

particle type	weak hypercharge	
	left-handed	right-handed
up	+1/6	+2/3
down	+1/6	-1/3
lepton	-1/2	-1

Table 1.1: Weak hypercharges of SM fermions excluding neutrinos.

with

$$N_\phi = \chi_\phi \dim(\mathcal{I}_\phi) T(\mathcal{C}_\phi) , \quad (1.23)$$

$$E_\phi = \chi_\phi \dim(\mathcal{C}_\phi) \text{Tr}(q_\phi^2) . \quad (1.24)$$

Here $\dim(\mathcal{I}_\phi)$ and $\dim(\mathcal{C}_\phi)$ are the dimensions of the weak isospin and color representations, respectively. $T(\mathcal{C}_\phi)$ is called the color Dynkin index, which is defined via the $SU(3)_c$ group generators, importantly $T(3) = 1/2$. $q_\phi = \mathcal{I}_\phi^{(3)} + \mathcal{Y}_\phi$ is the EM charge generator. $\text{Tr}(q_\phi^2)$ effectively amounts to plugging the squared EM charge of the particle in question into the equation. χ_ϕ denotes the PQ charge. Contributions from right handed fermions obtain an overall minus sign.

1.3 Axion Models

In this section the most common benchmark models for UV completions of the axion effective Lagrangian are introduced and used to derive concrete values for the anomaly coefficients at low energies.

Eqs. (1.23) and (1.24) show that the anomaly coefficients and thus the axion-photon coupling depend on which particles are charged under the PQ symmetry. Broadly speaking either the standard model (SM) fermions can have non-zero PQ charges or new PQ-charged fermions can be added.

Starting with models of the former type, SM fermions can be divided into three types with respect to their weak hypercharge as summarised in tab. 1.1¹. This distinction is helpful, since the Lagrangian should retain a global $U(1)_Y \times U(1)_{\text{PQ}}$ symmetry, meaning that all hyper- and PQ charges in every term of the Lagrangian need to sum to zero.

One can write the Yukawa part of the Lagrangian for a generic model with a single Higgs doublet $H \sim (1, 2, -1/2)$ (like the SM Higgs boson) as

$$\mathcal{L}^Y \supset -y_U \bar{Q}_L u_R H - y_D \bar{Q}_L d_R \tilde{H} - y_E \bar{l}_L e_R \tilde{H} + \text{h.c.} , \quad (1.25)$$

where y_i , $i \in \{U, D, E\}$ are the Yukawa couplings, i_R is a right handed fermion field of up-, down- or lepton-type, and Q_L and l_L are the left handed quark and lepton doublets, respectively. The hyper- and PQ charges of $\tilde{H} = i\sigma_2 H^*$ have the opposite

¹Different conventions exist for how the weak hypercharge is calculated differing from the chosen one by a factor of 2 and/or -1 . The terms ‘‘hypercharge’’ and ‘‘isospin’’ in the following always refer to weak hypercharge and weak isospin, respectively.

sign as H . Denoting the PQ charge of H as χ_H and requiring PQ invariance for this Lagrangian gives the conditions

$$\chi_H = \chi_{u_L} - \chi_{u_R} \quad (1.26)$$

$$-\chi_H = \chi_{d_L} - \chi_{d_R} \quad (1.27)$$

$$-\chi_H = \chi_{e_L} - \chi_{e_R} . \quad (1.28)$$

Leptons are not color charged, i.e. they have zero color Dynkin index, and do not contribute to the color anomaly. The sum in eq. (1.22) goes over SM quarks only and therefore leads to $N = 0$. Coupling the same Higgs doublet to all colored fermions leads to a vanishing color anomaly and thus does not solve the strong CP problem. The historically first axion model due to Weinberg and Wilczek (PQWW) [34, 35] therefore requires the addition of a second Higgs doublet.

In the PQWW model the axion can be identified as a superposition of the angular degrees of freedom of the two Higgs doublets after electroweak symmetry breaking. This relates the PQ scale $\sim v_a$ to the electroweak scale ~ 246 GeV leading to relatively strong couplings via f_a , which have quickly been ruled out [38–41]. Making the two scales independent of each other is the goal of so called invisible axion models, which can be divided into two general classes:

Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) type models introduce an additional Higgs singlet for this purpose. In these models as in the original PQWW model the SM fermions carry PQ charges. Contrary, in Kim-Shifman-Vainshtein-Zakharov (KSVZ) type models only new heavy quarks are PQ charged. Since these added quarks can be singlets, the necessary additional Higgs can also be a singlet.

1.3.1 DFSZ Model

For models introducing more than one Higgs doublet, flavor changing neutral currents (FCNCs), that is transitions e.g. between up and charm quark, are a generic prediction. Experimentally they have been shown to be strongly suppressed [42, 43], therefore a theoretical mechanism must be found to avoid FCNCs at least at tree level. A simple way to achieve this is imposing the Glashow-Weinberg condition [44, 45], i.e. coupling only one specific Higgs to each fermion type. Later in ch. 5 this condition will be relaxed, however in this section only one Higgs doublet $H_u \sim (1, 2, -1/2)$ coupling to all up-type quarks and another one $H_d \sim (1, 2, +1/2)$ coupling to all down-type quarks are considered. The DFSZ model also introduces an additional SM-singlet complex scalar field $S \sim (1, 1, 0)$.

There are three $U(1)$ symmetries, $U(1)_{H_u} \times U(1)_{H_d} \times U(1)_S$, associated with the phases of the three scalar fields in the kinetic and Yukawa sector of the Lagrangian, which need to be broken down into the $U(1)_Y \times U(1)_{\text{PQ}}$ symmetry mentioned earlier. This can be achieved by a term in the scalar potential of the form

$$V(H_u, H_d, S) \supset \lambda H_u H_d S^{\dagger 2} , \quad (1.29)$$

with arbitrary constant λ . Requiring PQ invariance of this term leads to a condition for the PQ charges:

$$\chi_u + \chi_d - 2\chi_S = 0 , \quad (1.30)$$

writing $\chi_{H_f} \equiv \chi_f$ with index $f \in \{u, d, S\}$ for sake of brevity. Via the potential it can also be ensured that the scalar fields develop vacuum expectation values (VEVs) v_f , expanding around which yields

$$H_u \supset \frac{v_u}{\sqrt{2}} e^{i\frac{a_u}{v_u}} \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad H_d \supset \frac{v_d}{\sqrt{2}} e^{i\frac{a_d}{v_d}} \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad S \supset \frac{v_S}{\sqrt{2}} e^{i\frac{a_S}{v_S}}, \quad (1.31)$$

where degrees of freedom not containing the axion have already been neglected.

Under PQ transformation the angular modes a_f transform as $a_f \rightarrow a_f + \kappa_f \chi_f v_f$ with κ_f being arbitrary constants, while the axion should transform as $a \rightarrow a + \kappa v_a$. The currents associated with the residual PQ and hypercharge symmetries can be written as

$$J_\mu^{\text{PQ}}|_a = \sum_f \chi_f v_f \partial_\mu a_f, \quad J_\mu^Y|_a = \sum_f Y_f v_f \partial_\mu a_f, \quad (1.32)$$

respectively, when only the relevant a_f fields are retained. It furthermore must be possible to write the PQ current in terms of the axion field as $J_\mu^{\text{PQ}}|_a = v_a \partial_\mu a$. This requirements together with (1.32) leads to finding the axion as a linear combination of the angular modes a_f [46]:

$$a = \frac{1}{v_a} \sum_f \chi_f v_f a_f, \quad (1.33)$$

$$v_a^2 = \sum_f \chi_f^2 v_f^2. \quad (1.34)$$

Choosing $v_S \gg v_{u,d}$ therefore makes the axion invisible.

To avoid kinetic mixing terms with the Z-boson one must impose an orthogonality condition on the currents (1.32):

$$\sum_f \chi_f Y_f v_f^2 = 0. \quad (1.35)$$

Thus, together with eq. (1.30) this determines the three PQ charges χ_f up to an overall normalisation that is usually fixed by choosing $\chi_S = 1$. One gets:

$$\chi_S = 1, \quad \chi_u = 2 \frac{v_d^2}{v_d^2 + v_u^2}, \quad \chi_d = 2 \frac{v_u^2}{v_d^2 + v_u^2}. \quad (1.36)$$

The discussion above just involved PQ charges of the scalar sector, the anomaly ratio E/N however sums over fermions contributing to the anomaly. So what about the PQ charges of the fermions? Yukawa terms such as in the Lagrangian (1.25) induce conditions like $\chi_u = \chi_{u_L} - \chi_{u_R}$. In general the anomaly coefficients depend only on the difference in PQ charges of right- and left-handed fermions.

It is necessary to specify whether to couple the leptons to H_d or to \tilde{H}_u . With this choice either $\chi_{e_L} - \chi_{e_R} \equiv \chi_e = \chi_d$ or $\chi_e = -\chi_u$, which defines the two benchmark models DFSZ₂-I and DFSZ₂-II, respectively². Since for SM fermions all variables

²In anticipation of the discussion for DFSZ-type models in ch. 5 the two Higgs doublets of the minimal models are explicitly referred to via the subscript. The standard nomenclature is DFSZ-I and DFSZ-II, respectively.

entering the calculation of eqs. (1.22) are known except for the PQ charges, one can write the anomaly ratio in general:

$$\frac{E}{N} = \frac{2}{3} + 2 \frac{\chi_u + \chi_e}{\chi_u + \chi_d}, \quad (1.37)$$

which for the benchmark models leads to

$$\text{DFSZ}_2\text{-I} : \frac{E}{N} = \frac{8}{3} \quad (1.38)$$

$$\text{DFSZ}_2\text{-II} : \frac{E}{N} = \frac{2}{3}. \quad (1.39)$$

1.3.2 KSVZ Model

The original KSVZ model adds a singlet fermion in the fundamental representation of color with zero hypercharge $Q \sim (3, 1, 0)$. Q gets its mass from a Yukawa coupling to a new Higgs singlet $S \sim (1, 1, 0)$. Analogously to the DFSZ case above, the S develops a VEV expanding around which immediately yields the axion as the angular degree of freedom:

$$S \supset \frac{v_S}{\sqrt{2}} e^{i \frac{a}{v_a}}. \quad (1.40)$$

SM fermions do not carry PQ charges and therefore do not contribute to the anomaly coefficients. E and N depend only on the properties of the added fermion. Specifically, the EM anomaly vanishes because in the original KSVZ model Q does not carry EM charge and thus

$$\text{KSVZ} : \frac{E}{N} = 0. \quad (1.41)$$

The three models (1.38), (1.39) and (1.41) are collectively called ‘‘minimal QCD axion’’.

1.3.3 KSVZ-Type Axion-Photon Coupling Distributions

It is possible to go beyond the minimal models presented above by introducing additional Higgs doublets to the theory for DFSZ-type models or by considering various representations of potentially multiple new heavy quarks for KSVZ-type models. DiLuzio et. al. [8, 47] as well as Plakkot and Hoof [9] constructed anomaly ratios for the latter case. Key results of their work are presented here, since this thesis completes the picture by doing the same for DFSZ-type models with multiple Higgs doublets in ch. 5 and also includes concrete comparisons between the resulting distributions.

A priori there is a vast space of possible theoretically allowed representations for every additional quark. However, in particular three in part interdependent phenomenological selection criteria can be found that limit the number of possible representations for each quark:

1. As sec. 2.3 will show, requiring axions to constitute a significant percentage of, but not more than, the DM energy density leads to preferred regions for the axion decay constant f_a . These regions differ substantially depending on which cosmological scenario is chosen. For simplicity the authors of [9] adopted $f_a < 5 \times 10^{11}$ GeV.

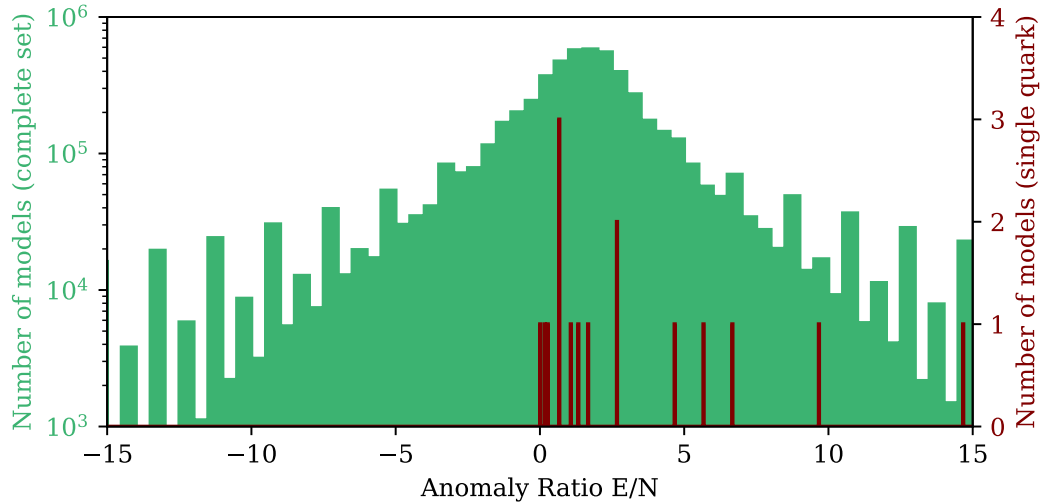


Figure 1.1: Anomaly ratio distributions for a single heavy quark in any representation (red) as well as the complete set of preferred models with up to 28 heavy quarks (green). Data taken from [8] and [9].

2. Constraints come from lifetimes of the heavy quarks beyond $\tau_Q \sim 10^{-2}$ s being ruled out by Big Bang Nucleosynthesis [48, 49] as well as gamma-ray observations [50].
3. Adding a particle $Q \sim (\mathcal{C}_Q, \mathcal{I}_Q, \mathcal{Y}_Q)$ with mass m_Q to a theory influences the running of the coupling constant at energies $> m_Q$. Large \mathcal{C}_Q , \mathcal{I}_Q and \mathcal{Y}_Q can lead to divergences in the coupling constants below the Planck scale already for a single heavy quark, leaving only 16 *preferred* representations when also taking the first two constraints into account. The running tends to grow stronger when adding more and more particles, leading to an upper bound on the possible number of additional heavy quarks of $N_Q = 28$.

The original KSVZ model as well as the 15 preferred single-quark representations from [8] are shown in fig. 1.1 as red lines. All anomaly ratios are unique, except for $8/3$ and $2/3$, which occur two and three times, respectively.

For multiple heavy quarks, allowing negative PQ charges, the authors of [9] find over 5×10^6 preferred models with only 820 different anomaly ratios. Their anomaly ratio distribution (fig. 1.1, green) is peaked closely around the value for which $g_{a\gamma}$ vanishes, making such axion models truly invisible. The peaked structure of the distribution at large absolute anomaly ratios will be more thoroughly investigated in sec. 5.3. The imposed selection criteria seem to also reject models leading to large axion-photon couplings. Defining

$$\widehat{E/N} \equiv \operatorname{argmax}_{E/N} (|E/N - 1.92|) \quad (1.42)$$

an analysis of all models with $N_Q \leq 9$ produces an extremal $\widehat{E/N} = -166/3$ for the preferred set of models whereas the full set has $\widehat{E/N} = -1312/3$.

2

Axions as Dark Matter

The universe is homogeneous and isotropic on large scales. Roughly 26% of its energy density is made up of elusive dark matter [51]. While it can be observed via its impact on the cosmic microwave background (CMB) and large structures, little is known about its impact and distribution on smaller, solar-system like scales. Dark matter could be both completely homogeneous as well as clumped in stellar-like objects as far as theoretical predictions are concerned. The universe expanded adiabatically from an early state of high density and high temperature, except during a short phase of exponential expansion called inflation and a subsequent reheating phase. The evolution of the universe is therefore often described as a function of its decreasing temperature.

This chapter shows that the theoretically well-motivated axion, introduced in ch. 1 to solve the strong CP problem, could also very well explain dark matter. Emphasis is put on aspects relevant to its detection using earth-based dark matter searches, specifically the relation of PQ scale to inflation and different potential production mechanisms. Ultimately, this chapter attempts to answer two questions: “Where should one look for an axion dark matter signal?” and “What does it look like?”.

Significant phenomenological differences arise whether the PQ symmetry is spontaneously broken before or after inflation (sec. 2.1). In Sec. 2.2 production mechanisms are discussed that could lead to axion densities and properties comparable to dark matter for both scenarios. The findings from the first two sections are used in sec. 2.3 to construct a mass range for the QCD axion which is preferred from a phenomenological point of view, effectively answering the first question. Sec. 2.4 then answers the second question by presenting a thorough calculation of the expected axion lineshape.

2.1 Axions and Inflation

Understanding the role of the axion in a cosmological context requires to consider the temperature evolution of the cosmos: As the universe cools down over time, the $U(1)_{\text{PQ}}$ symmetry gets spontaneously broken when $T \sim f_a$ and a specific value for the angular axion field a is chosen at random. Choosing a specific value for a can equivalently be expressed as choosing a specific initial angle θ_i . Only later when $T \sim \Lambda_{\text{QCD}}$ the PQ symmetry gets further broken down explicitly, the axion acquires its potential and one of the N_{DW} minima is selected. Regions in space, called Hubble volumes, that are not in causal contact at these times can randomly obtain different θ_i and select different minima.

Two energy scales play a pivotal role in axion phenomenology: Firstly, the Hubble rate during inflation H_I , and secondly, the maximal temperature reached in the universe after inflation T_{max} . The former can be constrained by CMB measurements [52], the latter is governed by the precise reheating scenario implemented in the inflationary model [53].

Pre-inflationary Scenario: If f_a is larger than both of these energy scales, the PQ symmetry is spontaneously broken during and never restored after inflation. Patches of the universe which are causally connected before inflation choose specific θ_i and get enlarged subsequently to leave only one specific value for θ_i in the observable universe. $\theta_i \in [-\pi, \pi[$ is therefore a free parameter in this scenario. Furthermore the axion is present as a massless field during, but not actively driving inflation, which leads to isocurvature perturbations. They leave distinct imprints in the CMB compared to adiabatic curvature perturbations [51], leading to strong constraints on the Hubble parameter during inflation in terms of the axion parameters [7]

$$H_I \lesssim 4 \times 10^6 \left(\frac{\Omega_a}{\Omega_{\text{DM}}} \right)^{-1} \left(\frac{f_a}{4 \times 10^{11} \text{ GeV}} \right) \theta_i . \quad (2.1)$$

Here Ω_a and Ω_{DM} are the density parameters for axions and dark matter, respectively, i.e. the normalised energy densities assuming a flat universe. Eq. 2.1 requires inflation at a relatively low energy scale for values of $\theta_i \propto \mathcal{O}(1)$.

Post-inflationary Scenario: If at least one of the scales H_I or T_{max} is larger than the PQ scale, the $U(1)_{\text{PQ}}$ gets spontaneously broken after inflation. Hubble volumes with different θ_i will come into causal contact at later times leading to a plethora of different initial misalignment angles in the observable universe. To obtain correct predictions for the axion energy density, one must average over all possible values taking into account anharmonic effects of the axion potential which leads to [36]

$$\theta_i \equiv \sqrt{\langle \theta_i^2 \rangle} \approx 2.15 . \quad (2.2)$$

At the boundaries of the patches topological defects arise (see fig. 2.1), which can decay into a non-relativistic axion population as we will see in the next section.

2.2 Production Mechanisms

2.2.1 Thermal Production

A thermal population of QCD axions could be produced via axion interactions with the SM. All axion models necessarily contain couplings to either SM particles or heavy quarks, and effective couplings to gluons and pions, which are all suppressed as $1/f_a$.

Axions are difficult to produce in this manner in sufficient abundance [54] and might still be relativistic at the present day. If this is the case they cannot serve as DM, but rather as dark radiation comparable to neutrinos with non-zero mass. For minimal QCD axion models decay constants leading to significant energy densities of thermal axions have been ruled out by stellar cooling constraints [55]. For these two reasons thermal axions are not further considered in this work.

2.2.2 Misalignment Mechanism

Misalignment production derives an energy density for the axion from the fact that the initial value for the theta angle θ_i usually deviates from the CP conserving vacuum state and has to approach $\theta = 0$ only after some time. It is therefore also referred to as vacuum realignment. Its contribution to Ω_{DM} is non-thermal which is why it can lead to a non-relativistic contribution to the cosmological energy density even for the lightest axion masses. In the following we will see that misalignment produced axions can constitute a significant percentage of the DM density.

Varying the action for an axion field a coupling minimally to gravity leads to the following equation of motion ([56] for a derivation):

$$\ddot{a} + 3H(T)\dot{a} + m_a(T)^2 a = 0, \quad (2.3)$$

with temperature dependent axion mass $m_a(T)$ and Hubble parameter $H(T)$. The approximations made in sec. 1.2 to obtain the axion mass (1.16) were only valid below the QCD scale $\Lambda_{\text{QCD}} \sim 200$ MeV. In general, effects leading to a potential and thus a mass for the axion only become relevant at this energy scale. For larger energy scales T , with $\Lambda_{\text{QCD}} \ll T < f_a$, the axion mass approaches $m_a = 0$. This means that m_a is not constant as the cosmological temperature decreases, which is made explicit in (2.3). Its exact scaling is being actively investigated [57–61].

To solve the differential equation (2.3), initial conditions

$$a_i = f_a \theta_i, \quad \dot{a}_i = 0, \quad (2.4)$$

need to be imposed. Here the latter condition is due to the absent axion potential at the high energies in the early universe. Qualitatively at early times, i.e. big temperatures, $m_a(T) \ll 3H(T)$ and the axion field and its energy density remain constant in time. During this epoch it effectively behaves as dark energy. When $m_a(T) \approx 3H(T)$, the axion field starts oscillating coherently around the vacuum $\theta = 0$ and its energy density decays over time due to the expansion of the universe, similarly to dark-matter. At the present day the energy density in axions can be approximated as

$$\rho_a \approx \frac{1}{2} m_a^2 a_0^2, \quad (2.5)$$

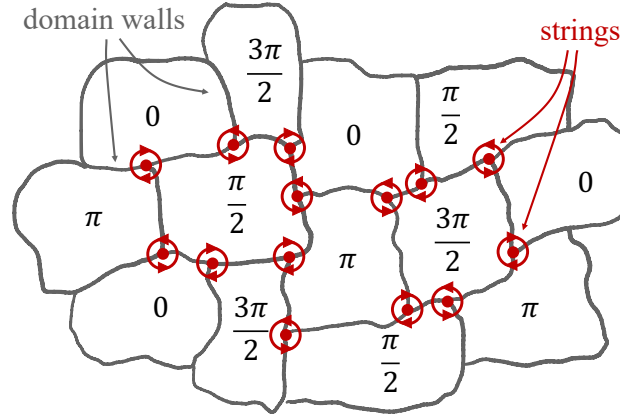


Figure 2.1: Schematic of topological defects in 2D after explicit breaking of the PQ symmetry for $N_{\text{DW}} > 1$. The size of the patches is governed by the size of Hubble volumes when $T \sim \Lambda_{\text{QCD}}$. Cosmic strings occur at the intersections of multiple patches, domain walls mark the boundaries between patches. Together they form a stable string-domain wall network for $N_{\text{DW}} > 1$.

where a_0 is the amplitude of the axion field. In terms of the density parameter Ω_a one finds [62]

$$\frac{\Omega_a^{\text{mis}}}{\Omega_{\text{DM}}} \approx \left(\frac{f_a}{4 \times 10^{11} \text{ GeV}} \right)^{7/6} \left(\frac{\theta_i}{1} \right)^2, \quad (2.6)$$

which is a good approximation below $f_a < 10^{15}$ GeV and ignores anharmonic corrections to the potential that arise as $\theta_i \gtrsim 1$. A thorough calculation is found in [63].

2.2.3 Decay of Topological Defects

As discussed in sec. 2.1, in the post-inflationary scenario topological defects can occur between regions of different θ_i . They can be classified in two categories:

Cosmic Strings. Consider a situation after spontaneous breaking of the PQ symmetry, when every Hubble volume has been assigned a specific initial misalignment angle and multiple patches have come into causal contact. In this situation paths may exist around which θ_i varies by $2\pi v_a$ (compare fig. 2.1). Enclosed in every such path is a point for which θ_i is not well-defined. In three spatial dimensions these points form long string-like objects or closed loops, which are called cosmic strings.

Domain Walls. When the PQ symmetry gets explicitly broken and the individual θ_i relax into distinct minima of the axion potential, a second type of topological defect forms. These so called domain walls are the reason why in sec. 1.2 the number of potential minima was called “domain wall number”. Domain walls arise, because it is not possible to smoothly transition from one minimum to another. Edges and intersections of domain walls are formed by cosmic strings. Even for $N_{\text{DW}} = 1$ domain walls exist due to the winding number discussed in sec. 1.1. However, in this case intersections of domain walls are not possible.

Individual, non-loop cosmic strings are stable. An interacting network of strings, however, will inevitably form closed loops that decay emitting non-relativistic axions. These radiated axions could form non-virialised axion streams, axion stars or miniclusters.

Domain walls in the $N_{\text{DW}} = 1$ scenario rapidly decay before they become cosmologically relevant. For $N_{\text{DW}} > 1$ however, they are stabilised by the intersections. Their energy density now would quickly reach unacceptably big values, which is known as the domain wall problem [64] (see [65] for a review).

Estimating the contribution to the energy density from the decay of topological defects requires numerical simulations capable of simultaneously resolving the Hubble scale as well as the width of the cosmic strings. Even with the advanced methods employed by various groups [66–74] the resulting densities vary a few orders of magnitude around Ω_a^{mis} .

2.3 The QCD Axion Window

The considerations above lead to a cosmologically preferred window, where the axion could constitute the entirety of DM without requiring fine-tuning of other parameters.

In the pre-inflationary scenario, the authors of [79] find a probability distribution for the axion mass with a 95% credible interval of $0.1 \mu\text{eV} \leq m_a \leq 160 \mu\text{eV}$ and a maximal value around $\sim 30 \mu\text{eV}$, which is shown as the green band on the pre-inflationary side of fig. 2.2. To obtain this result they assume $\Omega_a \sim \Omega_{\text{DM}}$, an initial misalignment angle randomly drawn from $\theta_i \in [-\pi, \pi]$, the local DM density, and E/N around the QCD window. They do not impose a preferred order of magnitude for f_a and also take existing limits into account.

In the post-inflationary scenario, due to the initial misalignment angle being specified by eq. (2.2), the axion decay constant for which misalignment production saturates the DM density can be calculated using (2.6) as $f_a^{\text{crit}} \approx 10^{11} \text{ GeV}$. Other computations in literature obtain similar results [57].

The uncertainty in this scenario comes from difficulties in estimating the axion abundance from decay of topological defects as described above. These decays most likely lead to scarce, compact objects with a very low rate for encounters with earth [80] making them inaccessible to earth-bound DM experiments like axion haloscopes. Therefore in the post-inflationary scenario axion masses with dominant misalignment production provide the best motivated parameter space with respect to these searches. In fig. 2.2 the range for which $\Omega_a^{\text{mis}} \geq 0.1 \Omega_{\text{DM}}$ is shown as a green band on the post-inflationary side.

In both scenarios the proposed MADMAX experiment [78, 81] covers a large portion of the well-motivated parameter space. This experimental effort will constitute a primary focus of this thesis and will be covered in detail in sec. 3.

More precise estimates for preferred and excluded regions than the ones shown in fig. 2.2 can be obtained by considering anharmonic effects of the axion potential [63] and a thorough investigation of axion properties during the QCD phase transition [82].

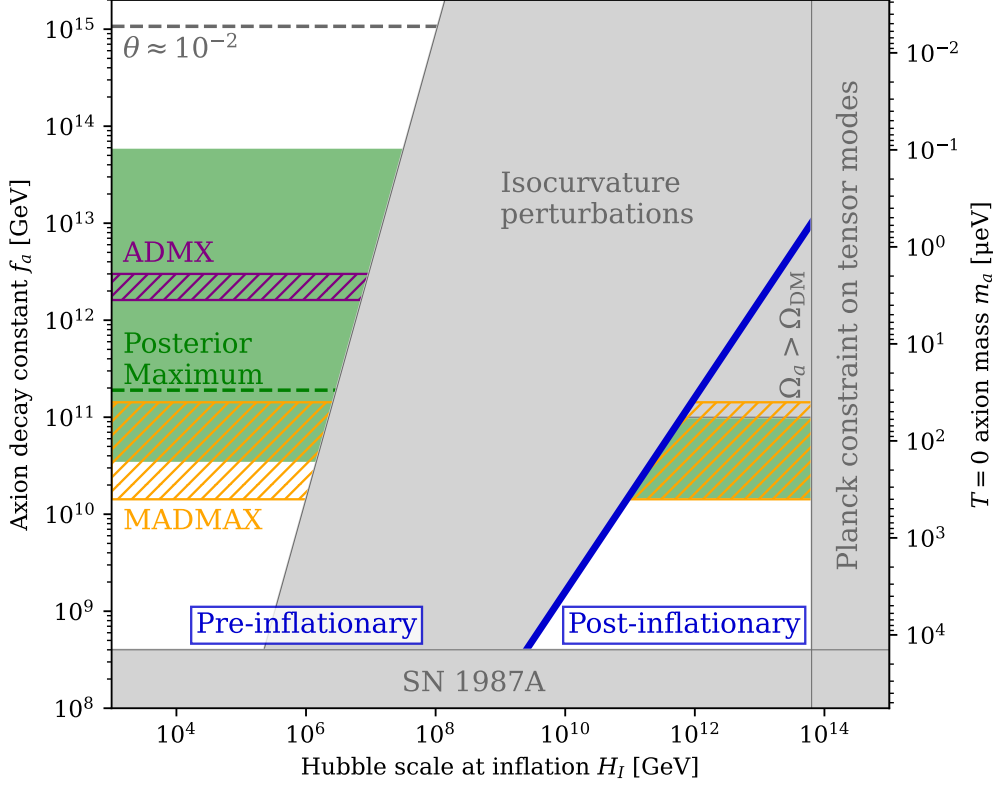


Figure 2.2: The QCD axion parameter space. Large H_I are constrained by the absence of tensor fluctuations in CMB measurements [52, 75], small f_a by supernova constraints [76]. Isocurvature perturbations constrain the pre-inflationary scenario via (2.1), whereas in the post-inflationary scenario for large f_a too much axion DM is produced (all grey regions). Fine-tuning to zero of θ_i by more than two orders of magnitude, indicated by the grey dashed line, makes $f_a > 10^{15}$ GeV undesirable. In purple the region is shown where the ADMX experiment reaches KSVZ sensitivity [77], in orange the proposed sensitivity range of the MADMAX experiment [78]. The blue line distinguishes pre- and post-inflationary scenario assuming inefficient thermalisation after inflation ($T_{\text{max}} < H_I$). Very efficient thermalisation ($T_{\text{max}} > H_I$) could open up the parameter space for the post-inflationary scenario down to $H_I \sim 10^4$ GeV for the projected MADMAX sensitivity range [53]. The green region indicates the preferred parameter space for pre- and post-inflationary scenarios as described in the text.

2.4 The Axion Lineshape

For direct DM searches such as axion haloscopes, where the full energy of a DM particle is converted into the energy of detectable photons, the photon frequency ω can be written as

$$\omega = \frac{m}{\sqrt{1-v^2}}, \quad (2.7)$$

where m and v are respectively the mass and the velocity relative to the lab of the DM particle. In particular, for a particle in the virialized DM halo v has two components: The movement of the DM particle in the restframe of the DM halo, which can be quantified by the DM velocity dispersion σ_v , and the velocity of the lab v_l with respect to the DM halo.

Assuming a standard halo model with spherical symmetry and a $1/r^2$ density profile, where r is the distance to the centre of the halo, leads to a Maxwell-Boltzmann distribution in DM velocities

$$f(\mathbf{v})d^3\mathbf{v} = \frac{1}{(2\pi)^{3/2}} \frac{1}{\sigma_v^3} e^{-\frac{|\mathbf{v}|^2}{2\sigma_v^2}} d^3\mathbf{v}. \quad (2.8)$$

In a lab frame boosted by v_l with respect to the restframe of the DM halo one can write in spherical coordinates without loss of generality

$$\begin{aligned} \tilde{v}_x &= v_x - v_l = v \sin(\theta) \cos(\phi) - v_l, \\ \tilde{v}_y &= v_y = v \sin(\theta) \sin(\phi), \\ \tilde{v}_z &= v_z = v \cos(\theta). \end{aligned} \quad (2.9)$$

Therefore

$$f(\tilde{\mathbf{v}})d^3\tilde{\mathbf{v}} \equiv f(\tilde{\mathbf{v}})d\tilde{v}_x d\tilde{v}_y d\tilde{v}_z = \sqrt{\frac{2}{\pi}} \frac{v}{\sigma_v v_l} \exp\left(-\frac{v^2 + v_l^2}{2\sigma_v^2}\right) \sinh\left(\frac{v v_l}{\sigma_v^2}\right) dv \equiv f_b(v)dv, \quad (2.10)$$

where the integrals $d\theta d\phi$ have already been evaluated. $f_b(v)$ is the probability distribution function (PDF) of a Maxwell-Boltzmann distribution with velocity dispersion σ_v boosted by v_l .

Realistic experiments have a non-vanishing resolution bandwidth (RBW) $\Delta\omega$, which means that signal power observed in a specific frequency bin receives contributions from all frequencies $[\omega - \Delta\omega/2, \omega + \Delta\omega/2]$. This effect becomes noticeable if $\Delta\omega$ is not much smaller than the FWHM of $f_b(v)$. Assuming a top hat shaped response, the signal strength in a given bin can be calculated by evaluating the cumulative distribution function (CDF)

$$F_b(v) \equiv \frac{\sigma_v}{\sqrt{2\pi}v_l} e^{-\frac{(v+v_l)^2}{2\sigma_v^2}} \left[1 - \exp\left(-\frac{2vv_l}{\sigma_v^2}\right) \right] + \frac{1}{2} \left[\operatorname{erf}\left(\frac{v-v_l}{\sqrt{2}\sigma_v}\right) + \operatorname{erf}\left(\frac{v+v_l}{\sqrt{2}\sigma_v}\right) \right] \quad (2.11)$$

at the edges of the frequency bin and subtracting the results from each other. The resulting signal shape can be seen in fig. 2.3 Note that the frequency dependence $v(\omega)$ enters via (2.7). For big RBWs the relative position of the frequency corresponding to $v = 0$ with respect to the edges of its bin can also play a relevant role (see [83] or subsec. 7.3.6).

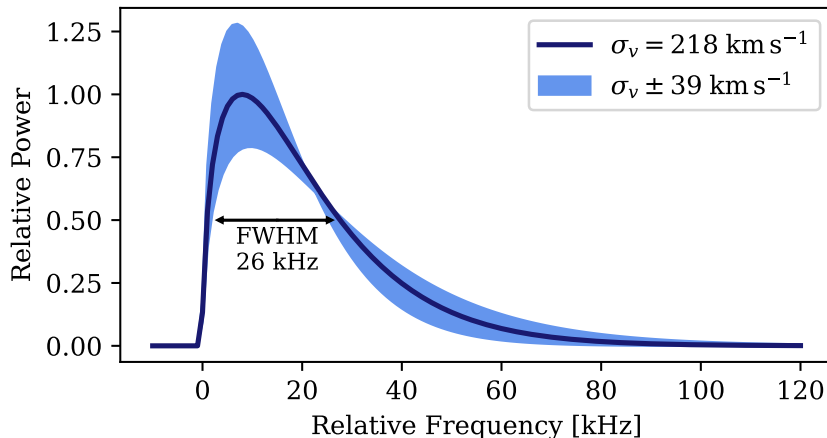


Figure 2.3: Axion signal shape computed from (2.11) for an axion at 19 GHz. For v_l and σ_v the values and uncertainties given in (2.12) are used.

To obtain a concrete signal model it is necessary to find realistic values and uncertainties for σ_v and v_l . The DM velocity dispersion σ_v is taken from [84] with a best fit value of $\sigma_v = 218 \pm 6 \text{ km s}^{-1}$. For v_l the solar velocity is used neglecting subdominant seasonal and daily modulations (see [85] for a thorough examination of these effects). From the vectorial velocity quoted in [86] the value $v_l = 242 \pm 2 \text{ km s}^{-1}$ is obtained.

The orbital velocity of earth around the sun varies around $29.8 \pm 0.5 \text{ km s}^{-1}$ [87], with rotational speeds being of the same order of magnitude as the seasonal variation of the orbital velocity.

For $v_l \sim \sigma_v$ the DM velocity distribution in the boosted lab frame still resembles a Maxwell-Boltzmann distribution reasonably well with uncertainties in v_l only leading to minor shifts in the distribution. In the context of a detection experiment it is therefore sufficient to translate the uncertainty on v_l into the σ_v parameter and assume v_l to be fixed. Propagating maximal errors then gives a conservative uncertainty estimate

$$v_l = 242 \text{ km s}^{-1}, \quad \sigma_v = 218 \pm 39 \text{ km s}^{-1}. \quad (2.12)$$

The considerations presented in this section so far apply equivalently to other virialised DM candidates. The uncertainty in σ_v will be leveraged in the Bayesian analysis of ch. 6, whereas the analysis in ch. 7 will assume best fit values.

The non-relativistic velocities of CDM lead to very light DM candidates such as the axion being well-described as classical fields. Taking the median of the boosted Maxwell-Boltzmann distribution (2.10) evaluated using the best fit values (2.12) gives $v_a \sim 403 \text{ km s}^{-1}$. For DM masses in the green range of fig. 2.2, $m_a \sim 0.1\text{-}400 \mu\text{eV}$, this leads to de Broglie wavelengths of

$$\lambda_{\text{dB}} = \frac{2\pi}{m_a v_a} \sim 2.3 \text{ m} - 9.2 \text{ km} \quad (2.13)$$

and with a local CDM abundance of $\rho_a \sim 300 \text{ MeV cm}^{-3}$ to a number density of $9 \times 10^{18} \lambda_{\text{dB}}^{-3}$ to $2 \times 10^{33} \lambda_{\text{dB}}^{-3}$, far in the classical regime.

3

Detecting Dark Matter Axions with Dielectric Haloscopes

Humanity has roughly 200 years of experience in building very sophisticated detectors for photons [88], but a successful detector for elusive dark matter particles has yet to be built. It is therefore unsurprising that many experimental efforts for detecting the axion focus on first converting it to photons, which can then be detected.

Sec. 2.4 already demonstrated the relationship between the mass of a dark matter particle and the frequency of the photon it converts into if any form of coupling exists. For the mass range relevant to axions $m_a \sim 0.1\text{-}400 \mu\text{eV}$, these photons are in the radiofrequency (RF) regime between 20 MHz and 100 GHz. RF engineering is therefore an integral part of this type of axion experiments.

Experiments trying to detect axions via its coupling to photons usually rely on a magnetic field to prompt axion-photon conversion. They can be divided into three broad categories:

Helioscopes like CAST [89] or IAXO [90] point a magnet at the sun in the hope that solar axions convert to photons inside of the magnetic field.

Instead of relying on the sun to produce axions, *light-shining-through-a-wall* (LSW) experiments produce axions themselves using a laser and a magnetic field. These axions can then travel through an obstacle that does not let any of the laser light pass and are subsequently back-converted into photons using a second magnet. While requiring less theory assumptions, LSW experiments such as ALPS-II [91] involve two axion-photon conversions and their sensitivity therefore scales with $g_{a\gamma}^4$.

The third type of experiment are *haloscopes* like ADMX [92] or HAYSTAC [93]. Contrary to the other two types of experiments, they rely on axions being locally present as dark matter. These axions are usually converted into photons in a resonator inside of a magnetic field, the resonance frequency of which has to be tuned in order to be sensitive to different axion masses.

Conventionally haloscopes use RF cavities as resonators, the size of which is determined roughly by the wavelength of the produced photons. The observed power from such a setup is proportional to the volume allowing for axion-photon conversion, which leads to a lower sensitivity for higher frequencies. Dielectric haloscopes are one attempt to alleviate this issue by combining multiple regions with length scales similar to the photon wavelength into one larger conversion volume.

This chapter describes the working principle behind dielectric haloscopes with special focus on the MADMAX haloscope and its already constructed prototype variants. It is structured as follows: First sec. 3.1 covers necessary basics of axion electrodynamics following the description in [94]. Sec. 3.2 then explains how dielectric interfaces can be used to enhance a potential axion signal, effectively serving as the aforementioned resonator. The MADMAX experiment, a haloscope using dielectric disks, is introduced in detail in sec. 3.3. Analysing data from a MADMAX prototype will be the focus in chs. 7 and 8. Its target design and currently existing prototypes are therefore described in detail, concluding with calibration strategies for the data analysed later.

3.1 Axion Electrodynamics

Including QED and axion specific terms from eq. (1.10) into the Lagrangian, one can derive as equations of motion the modified Maxwell equations [95, 96], which include the axion as an additional source with effective charge and current density

$$j_a^0 = -g_{a\gamma} \mathbf{B}_e \cdot \nabla a , \quad (3.1)$$

$$\mathbf{J}_a = g_{a\gamma} \mathbf{B}_e \partial_t a . \quad (3.2)$$

In the presence of a strong, static external magnetic field \mathbf{B}_e the leading terms of the macroscopic axion-Maxwell equations read [94]:

$$\epsilon \nabla \cdot \mathbf{E} = j^0 - j_a^0 , \quad (3.3)$$

$$\mu^{-1} \nabla \times \mathbf{B} - \epsilon \partial_t \mathbf{E} = \mathbf{J} + \mathbf{J}_a , \quad (3.4)$$

$$\nabla \cdot \mathbf{B} = 0 , \quad (3.5)$$

$$\nabla \times \mathbf{E} + \partial_t \mathbf{B} = 0 , \quad (3.6)$$

$$(\partial_t^2 - \nabla^2 + m_a^2) a = g_{a\gamma} \mathbf{E} \cdot \mathbf{B}_e , \quad (3.7)$$

where the magnetic field \mathbf{B} does not include the external field \mathbf{B}_e .

For detectors much smaller than the de Broglie wavelength of the axion λ_{dB} , one can assume a spatially constant axion field $a = a_0 \exp(-im_a t)$ with amplitude a_0 . Plugging (3.6) into the time derivative of (3.4) and assuming constant \mathbf{B}_e leads to the wave equation for the axion-induced electric field with solution

$$\mathbf{E}_a(t) = -\frac{g_{a\gamma} \mathbf{B}_e}{\epsilon} a(t) . \quad (3.8)$$

Note that the amplitude of this field in vacuum is given by $E_0 \equiv |g_{a\gamma} \mathbf{B}_e a_0|$. The field \mathbf{E}_a has discontinuities on interfaces inside the constant magnetic field \mathbf{B}_e where the dielectric constant ϵ is discontinuous (fig. 3.1¹). However, the parallel components of

¹For the remainder of this chapter the x direction is assumed to be perpendicular to the external magnetic field as defined in this figure.




Figure 3.1: Sketch of axion-induced emission sourced at a dielectric interface in a magnetic field in the classical field picture. In the presence of an external magnetic field \mathbf{B}_e (blue arrows) the axion field a (dashed orange line) induces an electric field \mathbf{E}_a (green line) (eq. 3.8). \mathbf{E}_a depends on the dielectric constant ϵ of the medium, leading to a discontinuity at the interface between dielectric medium (grey area) and vacuum. This step is compensated by the emission of propagating EM waves on both sides of the interface with their E-field components shown in brown. Modified based on [97].

total electric and magnetic fields must still be continuous, which prompts the emission of propagating electromagnetic (EM) waves from the interfaces with amplitudes

$$|\mathbf{E}_1^\gamma| = \frac{E_0}{\sqrt{\epsilon_1}} \left| \frac{1}{\sqrt{\epsilon_2}} - \frac{1}{\sqrt{\epsilon_1}} \right|, \quad (3.9)$$

$$|\mathbf{E}_2^\gamma| = \frac{E_0}{\sqrt{\epsilon_2}} \left| \frac{1}{\sqrt{\epsilon_2}} - \frac{1}{\sqrt{\epsilon_1}} \right|, \quad (3.10)$$

$$|\mathbf{H}_{1,2}^\gamma| = E_0 \left| \frac{1}{\sqrt{\epsilon_2}} - \frac{1}{\sqrt{\epsilon_1}} \right|, \quad (3.11)$$

with $\mathbf{H} = \mathbf{B}/\mu_0$ assuming a negligible magnetic response. These propagating EM waves have energies corresponding to the axion mass and therefore, in the proposed MADMAX range of 40-400 μeV , wavelengths of $\lambda = 3.1\text{-}31$ mm.

The observable power density emitted by the dielectric interface into free space can be calculated from the cycle-averaged Poynting flux

$$\bar{S}_1^\gamma = \frac{1}{2} \overline{[\text{Re}(\mathbf{E}_1^\gamma) \times \text{Re}(\mathbf{H}_1^\gamma)]_x} = -\frac{E_0^2}{2\sqrt{\epsilon_1}} \left(\frac{1}{\sqrt{\epsilon_2}} - \frac{1}{\sqrt{\epsilon_1}} \right)^2, \quad (3.12)$$

where the minus sign indicates that it points in negative x direction. Using eqs. (2.5), (1.19) and (1.18) and assuming the entire local DM consisting of axions of $\rho_a \sim 300 \text{ MeV cm}^{-3}$, yields for the power emitted from a perfect mirror ($\epsilon_1 = 1$; $\epsilon_2 = \infty$) with surface area A :

$$P_\gamma = 2.28 \times 10^{-27} \text{ W} \left(\frac{A}{1 \text{ m}^2} \right) \left(\frac{\mathbf{B}_e}{10 \text{ T}} \right)^2 \mathcal{C}_{a\gamma}^2. \quad (3.13)$$

Experiments using this straightforward approach, like BREAD [98, 99] or DOSUE [100], are called “dish antenna haloscopes” [101]. They inherently are broadband searches, covering a wide range of potential axion masses at the same time. However, making the tiny power (3.13) detectable requires very long integration times or an enormous mirror surface. It is therefore desirable to boost the signal strength in some way.

3.2 Power Boost Factor

There is a way to increase the power of the emitted radiation compared to the aforementioned dish antenna approach. It leverages the coherence of the emission from multiple dielectric interfaces placed inside the static field \mathbf{B}_e at distances much smaller than the de Broglie wavelength of the axion. Such coherent emissions can be made to constructively interfere, significantly increasing the detected signal.

In the following, regions $r \in [0, m]$ with dielectric constant ϵ_r are labelled with increasing indices from left to right. Regions are assumed to extend infinitely in two spatial dimensions and to have interfaces perpendicular to the x direction (as shown in fig. 3.1), i.e. only 1D effects are considered. Furthermore assuming a spatially constant external magnetic field and only media with $\mu_r = 1$, the electric and magnetic fields inside each region become (see [94])

$$E_r^{\text{tot}}(x) = -\frac{E_0}{\epsilon_r} + R_r e^{+i\omega\sqrt{\epsilon_r}\Delta x} + L_r e^{-i\omega\sqrt{\epsilon_r}\Delta x}, \quad (3.14)$$

$$H_r^{\text{tot}}(x) = \sqrt{\epsilon_r} \left(R_r e^{+i\omega\sqrt{\epsilon_r}\Delta x} - L_r e^{-i\omega\sqrt{\epsilon_r}\Delta x} \right), \quad (3.15)$$

with Δx measuring the distance to the left boundary of region r . The first term in (3.14) is the constant axion-induced electric field. The remaining terms in E_r^{tot} as well as H_r^{tot} come from the left and right moving EM waves with amplitudes L_r and R_r , respectively. E and B field continuity at all m interfaces defines a linear system of equations for L_r and R_r . The total emitted power of a system closed off by a mirror on the left side, R_m , is then boosted by

$$\beta^2 \equiv \left| \frac{R_m}{E_0} \right|^2, \quad (3.16)$$

where β^2 is called (power) boost factor, which generally depends on the frequency ω and thus the wavelength of the emitted photons.

R_m can include constructively interfering contributions from all surfaces in the external magnetic field, if appropriate distances between the surfaces are chosen. Additionally, emissions can be reflected inside the system multiple times, resonantly enhancing the constructive interference. The boost factor is largest in systems with alternating regions of large and small dielectric constants that have an optical thickness corresponding to roughly half of the wavelength. An easy example for such a system are equidistant dielectric disks of appropriate thickness in vacuum, which is used as a baseline for the discussions in the following paragraphs.

For a typical experimental effort trying to probe different axion mass/ frequency ranges, having a large boost factor in a very small frequency range may not be optimal

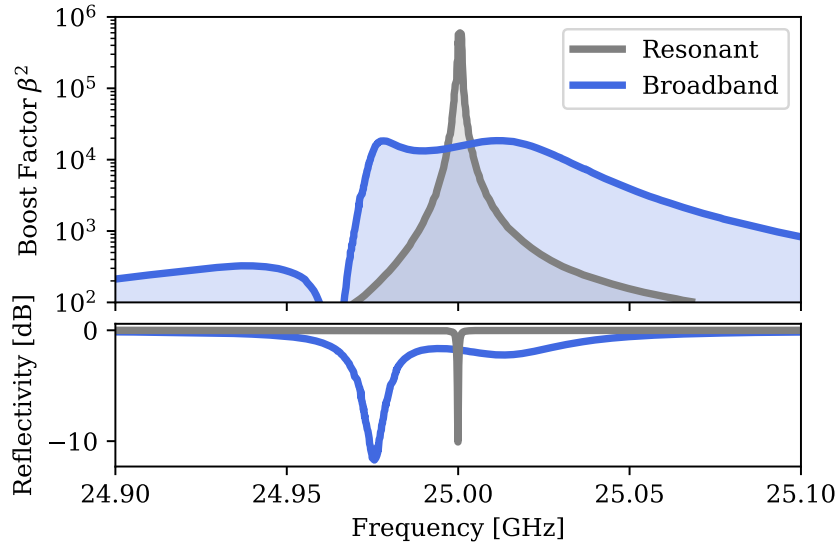


Figure 3.2: Simulated boost factor curves for two different configurations for a system with 20 dielectric disks ($\epsilon = 25$, 1 mm thickness). The configurations were optimised to yield the largest value possible for the minimum of β^2 in a 1 MHz (resonant) or a 50 MHz window (broadband), respectively. The bottom panel shows the simulated reflectivity for both configurations, clearly displaying dips at frequencies with large β^2 . Adapted from [94]

in terms of scanning speed. Realistic efforts contain some overhead when switching from one configuration to another, like calibration measurements or, especially at cryogenic temperatures, non-negligible durations for repositioning the dielectrics. Less resonant, more broadband configurations with top-hat like boost factor curves minimise this overhead. They could therefore be desirable to cover a larger frequency range in a smaller amount of time given a specific sensitivity target.

The authors of [94] showed, that it is indeed possible to trade big peak values for β^2 against an increased bandwidth by moving away from equidistant disk positions. For an example configuration with 20 dielectric disks this is demonstrated in fig. 3.2, upper panel. The effect is known as “area law” since the area under the square root of the boost factor curve integrated over the full frequency space remains approximately constant when losses are negligible.

A significant challenge for relying on a signal boost is that a direct measurement of the boost factor is not straightforward in the absence of an axion signal. However, eqs. (3.14), (3.15) and (3.16) suggest some correlation to the reflectivity \mathcal{R} response of the system. The interference conditions for propagating waves in the system are independent of whether they are sourced at the boundaries by non-negligible E_0 or injected via $L_m \neq 0$. Qualitatively, the simulations shown in fig. 3.2 indeed display a dip in reflectivity that coincides with frequencies of strong boost factors. In subsec. 3.3.4 reflectivity measurements will be used to quantitatively determine system parameters leading to a boost factor estimation.

Recently, a second strategy for obtaining boost factors has been developed, leveraging the electric fields E_R inside the system during a reflectivity measurement [102].

The fields E_R are measured using the bead pull method [103]: A small dielectric object, or bead, is inserted into the region of interest leading to a change in reflection coefficient, which is related to the field strength at its location. Probing E_R at multiple different x locations in the vacuum/ air gaps between regions with higher dielectric constant allows to reconstruct a reliable estimate for the boost factor in an open system with a single disk and mirror [104]. The method has already been applied for a first dark matter search [5].

In the 1D setup discussed above the axion emission from the mirror as well as its propagation through free space occur as plane waves. E and B field components are both perpendicular to the Poynting vector, therefore this solution is also called transverse electro-magnetic (TEM) mode.

In real, finite systems the existence of some boundary condition potentially changes the allowed modes. Consider for simplicity a circular mirror with radius r_0 , again potentially in combination with several dielectric disks. Such a system will in the following be referred to as booster. For the boundary conditions at $r = r_0$ the two extremes are free space, i.e. no reflection/ perfect absorption at the circular boundary, or a perfect circular waveguide, i.e. perfect reflection at the boundary. When absorbers are placed around the booster the free space condition is more applicable. Note, however, that even in this case inside the disks the boundary conditions are similar to that of a circular waveguide. Such systems need to be simulated taking the alternating boundary conditions into account. Alternatively, the booster can be placed inside of a metal cylinder leading to waveguide boundary conditions in the whole booster.

Waveguide modes differ substantially from the free space TEM mode. Firstly, the E and B fields of propagating EM fields cannot anymore both be perpendicular at the same time. Secondly, due to the fixed boundary conditions only a discrete set of modes can propagate inside the waveguide. One distinguishes between modes with a transverse E or a transverse B field and they are labelled as TE_{ij} or TM_{ij} , respectively with i indicating the number of radial and j indicating the number of azimuthal nodes of their field distribution. Higher order modes with larger i and j typically have smaller propagation constants as well as higher cutoff frequencies. The cutoff frequency determines up to which frequency a certain mode is still able to propagate inside a waveguide with a specific radius.

The booster has a relatively big radius, therefore multiple higher order modes are able to propagate in the system. For a full 3D simulation the propagation properties of these higher order modes need to be taken account. As a first estimate, however, a form factor η can be introduced that accounts for the mismatch between axion induced emission and the mode that couples to the receiver system by calculating the overlap integral between the two. An example is given in fig. 3.7.

This section explained in some generality how dielectric haloscopes are capable of boosting potential axion emission relative to dish antenna experiments. The considerations apply e.g. to the DALI [105, 106] and ORPHEUS [107, 108] experiments, or to the comparable approach by the LAMPOST experiment [109] probing axion dark matter at larger masses. For specificity the example of dielectric disks as the boosting components was employed, the working principle utilised by the MADMAX experiment, which is discussed in detail in the next section.




Figure 3.3: Baseline design of the MADMAX experiment. The green cylinder represents the booster consisting of a mirror on the left end and multiple dielectric disks. It is placed inside a cryostat (gray) to reduce noise levels. The booster part is enveloped by a dipole magnet, the coils of which are shown in red. Power emitted by the booster is guided by a focusing mirror (orange) into a horn antenna (yellow) and read out by a receiver chain (not shown). Not to scale. Taken from [78].

3.3 The MADMAX Experiment

As already discussed in sec. 3.2, one way of boosting a possible axion signal is by placing a mirror and multiple semi-transparent dielectric disks inside a magnetic field. The disks need to be positioned such that emissions from the various surfaces constructively interfere while being reflected multiple times within the system. This is the fundamental working principle of the MAgnetized Disk and Mirror Axion eXperiment (MADMAX).

The power emitted or reflected by the booster is focused into an antenna and coupled into a receiver system consisting of a low noise amplifier (LNA) as well as multiple stages of signal processing. The role of this receiver system is to amplify the signal and to convert the raw time-series spectrum, so it can be digitalised, Fourier transformed and stored as a real-valued frequency spectrum.

The goal of the MADMAX project is to search for QCD axions making up a majority of DM with DFSZ sensitivity ($E/N = 8/3$) in a mass range between $40 \mu\text{eV}$ and $400 \mu\text{eV}$ (i.e. 10 GHz - 100 GHz). An abstract schematic of the setup is shown in fig. 3.3.

3.3.1 MADMAX design target

Experimental parameters necessary for reaching this target sensitivity can be determined by setting the signal strength (3.13) into relation with the noise level from thermal and electronic backgrounds using the radiometer equation. The noise level depends on the effective system temperature T_{sys} , the resolution of the receiver system $\Delta\omega$, i.e. which frequencies are lumped together into one bin, and the integration time τ ,

i.e. the time in which the receiving system is actively measuring. For a DFSZ axion with $E/N = 8/3$, a local density of $\rho_a = 300 \text{ MeV cm}^{-3}$ and using the FWHM of a 19 GHz axion signal from fig 2.3, $\Delta\omega = 26 \text{ kHz}$, one obtains

$$\frac{S}{N} \approx 5 \left(\frac{A}{1 \text{ m}^2} \right) \left(\frac{B_e}{9.1 \text{ T}} \right)^2 \left(\frac{\beta^2}{5 \times 10^4} \right) \left(\frac{\eta}{0.69} \right) \left(\frac{2 \text{ K}}{T_{\text{sys}}} \right) \left(\frac{\tau}{4 \text{ days}} \right)^{1/2}. \quad (3.17)$$

The form factor η takes into account the mismatch between the axion induced emission and receiver system as described in sec. 3.2.

The reachable effective system temperature T_{sys} is limited by heat load from the motors responsible for moving the disks as well as noise contributions from the receiver chain and the booster itself. Results from first data-takings of a MADMAX prototype setup in a cryostat at $< 8 \text{ K}$ are to be published soon and construction of cryostat infrastructure at DESY is ongoing.

Concerning the magnetic field B_e , extensive design studies on the feasibility of a 9.1 T dipole magnet with an 1.35 m open bore at the centre have yielded positive results [110].

The size of the magnet bore may limit the maximal possible size A of the disks. The envisioned 1.35 m opening needs house cryostat as well as a contraption to mount and move the disks [111]. Requirements on manufacturing precision also play a role for determining the possible surface area of the disks depending on the material used. The simulations in this and the previous chapter assume $\epsilon = 25$ which is comparable with Lanthanum aluminate (LaAlO_3), the prototypes described below use Sapphire ($\epsilon \sim 9$). To reach the full 1 m surface area it may be necessary to construct the disks out of multiple individual tiles, simulations and manufacturing studies for which are ongoing [4, 81].

Manufacturing requirements can be broadly divided into three categories, which differ on the length scale of the inaccuracies: Tilts, planarity and surface roughness. Tilts describe placement of the disks which is not precisely parallel to the magnetic field. Planarity describes variations in the thickness of the disk on scales of 1/10 of the disk diameter. Surface roughness analogously describes variations of 1/4 of the photon wavelength. Studies on how the 3D power boost factor is affected by random variations at these three scales for a 20 and a 80 disk setup have been conducted [4]. The results can be seen in fig. 3.4.

More recently these investigations have been followed up by simulations for slightly curved disks. They show that taking disk inaccuracies into account when determining disk positions for optimal boost factors may alleviate the requirements imposed by fig. 3.4.

Reaching the MADMAX design target seems challenging but possible with significant steps toward it already having been taken. At target sensitivity it is possible to cover the target parameter space in a span of several years. To reach this goal, however, various intermediate milestones need to be passed, the most recent of which was taking physics measurements with multiple prototype setups. These setups are the focus of the next subsection.

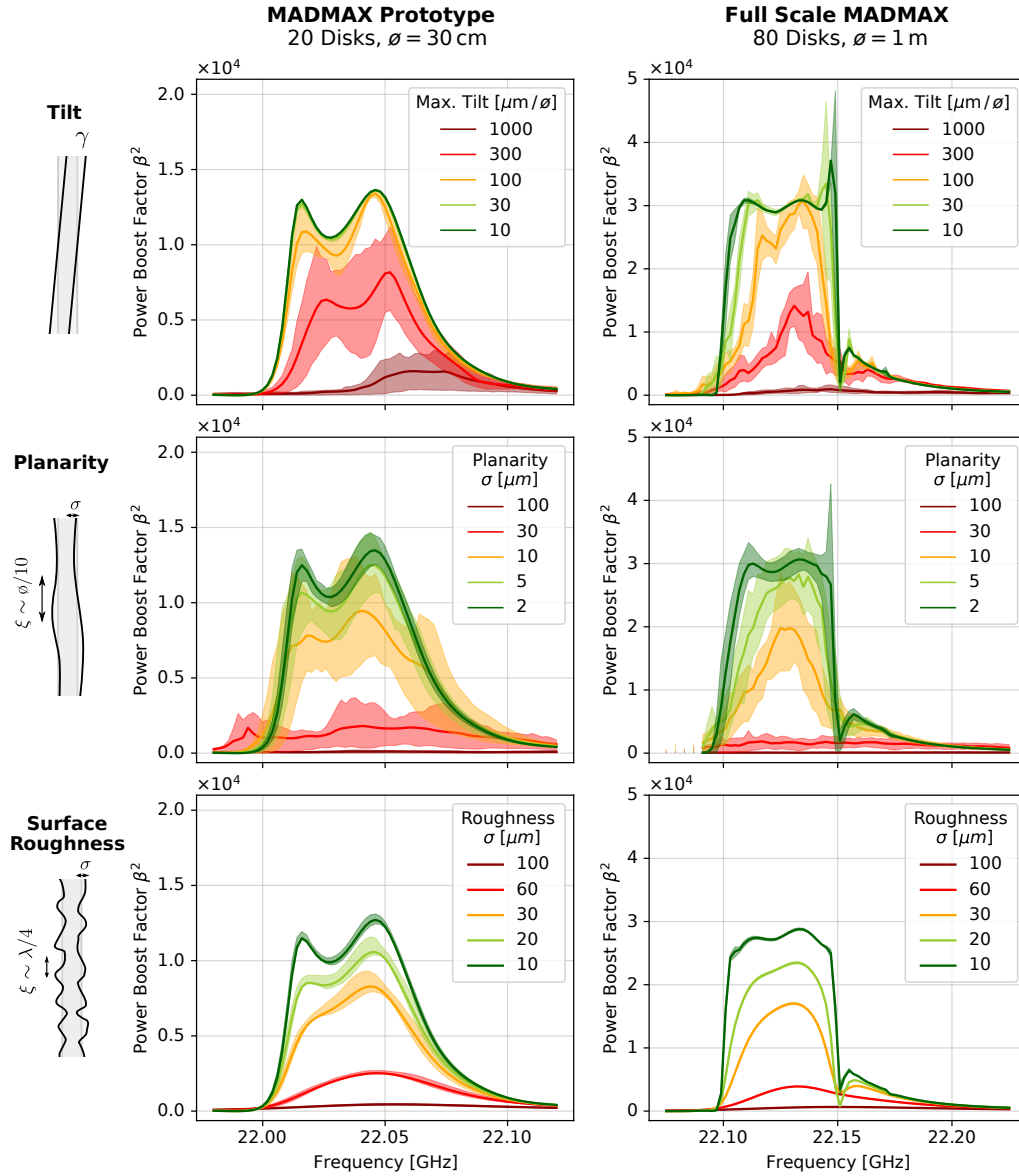


Figure 3.4: Influence of geometrical inaccuracies on the obtainable power boost factor including the form factor η for configurations with 20 LaAlO_3 disks with a diameter of 30 cm and 80 LaAlO_3 disks with 1 m diameter. The top row shows the effect of randomly tilting disks in any direction perpendicular to the surface of the mirror up to the value given in the captions. The middle and bottom rows show the effect of thickness variations of two different length scales. Shaded areas refer to 16 and 84 percentiles of the resulting boost factors with the solid lines indicating the ensemble mean. From [4].




Figure 3.5: Exploded view of the CB 200 booster. From left to right: Taper, dielectric lens (black), cylinder housing the disks and mirror, three sapphire disks (white) and spacer rings (grey) and aluminium mirror (golden). The system couples to a receiver system (not shown) on the left side. On the right side the contraption for putting pressure on the mirror was omitted. Adapted from [3].

3.3.2 Current MADMAX Prototypes

Between December 2023 and March 2024 several independent physics measurements have been conducted at CERN and at the University of Hamburg. The measurements used three distinct setups: CB 100, CB 200 and OB 300. The number refers to the diameter of the dielectric disks in millimetres. The naming scheme distinguishes between open booster (OB) setups, in which the dielectric disks are mounted on metallic suspensions surrounded by free space, and closed booster (CB) setups, in which the disks are placed inside a metal cylinder with their distances being determined by metallic spacer rings. OB systems are more similar to the target design for MADMAX and allow for easy measurement of internal electric fields using the bead pull method. CB systems on the other hand are easier to simulate due to the fixed boundary conditions provided by the metal cylinder.

OB 300 was used in December 2023 and January 2024 for measurements in Hamburg. It is an open system consisting of three sapphire disks with 30 cm diameter. Unlike the closed systems detailed below it makes use of a focusing mirror and a custom-made receiver system with components from Hamburg University and the Fermi National Accelerator Laboratory. No magnet was available to host a setup of this size at the time of measurement, therefore no axion search was possible. However, the measurement produced world-leading limits on dark photons, another dark matter candidate [5].

CB 100 is the first closed booster system constructed by the MADMAX collaboration. It was used in previous years for multiple test runs with and without magnetic field at CERN and at the Max-Planck Institute for Physics in Munich. During the measurement campaign at CERN in February and March 2024 it was used to obtain first cryogenic measurements inside of a magnetic field. To this end, as well as for the CB 200 measurements discussed in this thesis, the MORPURGO magnet was used with a magnetic field strength of up to 1.6 T. CB 100 consists of an aluminium mirror and three sapphire disks with a diameter of 10 cm and a thickness of 1 mm. The disks

are separated by gold-plated aluminium spacer rings and mounted inside a aluminium cylinder. The side opposite to the mirror is closed off by a gold-plated aluminium taper, the purpose of which is to couple any emissions of the booster into the receiver system.

CB 200 is a slightly larger, enhanced version of the CB 100 setup. Fig. 3.5 shows a schematic of the system. The most notable difference to CB 100 is the focusing lens (black in fig. 3.5) between booster and taper, which became necessary due to the increased size of the system to avoid unwanted modes at frequencies comparable to the boost factor peak. Additionally, CB 200 has some tuning capability: Different disk configurations can be chosen by using different sets of spacer rings, making it possible to choose different frequencies of maximal power boost factor in a ~ 1 GHz range. The distance between the disks is varied around 12-13 mm with exception of the first disk at ~ 8 mm distance from the mirror. For smaller shifts in frequency of order 10 MHz the high sensitivity of the booster system to variations of the air gap between mirror and first dielectric disk can be leveraged. The mirror is pressed onto the stack of disks by an adjustable screw, the pressure on which can be increased/ decreased to decrease/ increase the size of the first air gap on the μm level which shifts the boost factor distribution by several MHz to higher/ lower frequencies for measurements at ~ 20 GHz.

Ch. 7 features a first analysis of data from the CB 200 setup. Below some prerequisites are covered, which are necessary for the statistical analysis and treatment of systematics presented later.

3.3.3 Calibration of the CB 200 Setup

For physics and calibration measurements CB 200 was coupled to a Rode&Schwarz FSW50 spectrum analyser (SA) via multiple stages of amplification using a low noise amplifier (LNA), filtering, and mixing to convert the RF spectra down to lower frequencies. The data was then read out by a dedicated data acquisition system with negligible dead time. This receiver system was configured and tested by collaborators in Bonn and Aachen and is therefore called ‘‘Aachen/Bonn receiver’’ in the following. Spectra were taken at 8.94 Hz resolution with this setup. Individual spectra contain 15 min of time-series data. To measure its reflectivity response, CB 200 was directly connected to a vector network analyser (VNA).

The power recorded with an SA cannot easily be converted back into the power actually emitted by the booster. It needs to be calibrated, which can be done using the Y-factor calibration method [112]. Instead of the booster, a noise diode with known power output is attached to the SA via an attenuator. Then, the power output of the SA can be determined with the noise diode turned off and with the noise diode turned on. Because the power emitted by the noise diode for both, on and off, states is known for every frequency, the power measurement taken with the booster can be reconstructed. These power spectra are usually given in terms of T_{sys} . This temperature can be understood as the temperature a reference black body would need in order to lead to the observed power in the frequency bin. The procedure assumes a linear scaling between the output of the SA and the power emitted by the noise diode.

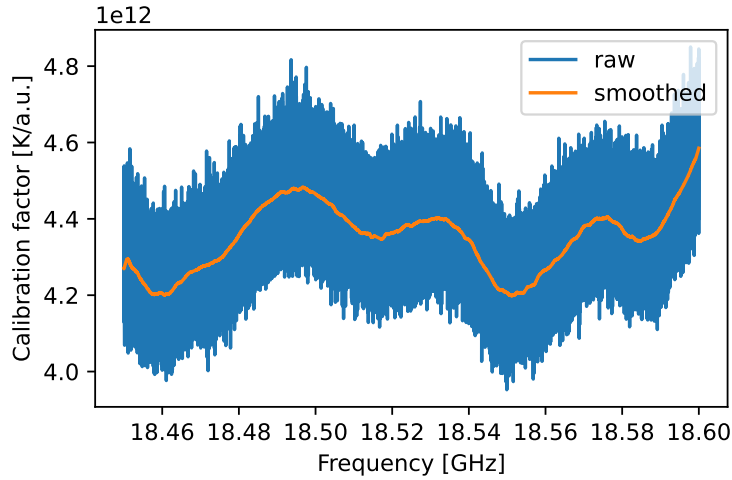


Figure 3.6: Example power calibration factor from the calibration measurement of Feb, 21st. The raw calibration factor obtained without smoothing is shown in blue, after applying an SG filter bin-to-bin fluctuations are significantly reduced (orange).

The output of this calibration measurement is a power calibration factor that is used to convert the arbitrary units of the raw SA output to temperature units. As an example the factor from the calibration performed on Feb, 21st is shown in fig. 3.6. The underlying calibration measurement has a total integration time much smaller than that of a typical physics measurement. For bin-to-bin fluctuations to not affect the sensitivity, the calibration factor is smoothed using a third-order Savitzky-Golay (SG) filter (see also subsecs. 6.2.1 and 7.3.5) with a width of 7.2 MHz.

The primary uncertainties playing a role in the power calibration procedure are coming from the precision of the noise diode and the attenuator. Both can be obtained from data sheets provided by the manufacturers. The attenuator can also be cross-checked via transmission measurements using the VNA. Additionally, measurement uncertainties play a role. They can be estimated as the standard deviation of the noise measurements around a mean given by the SG-filtered version of the dataset.

Lastly, the gain of the LNA is subject to environmental variation. This is reflected in uncalibrated spectra being shifted to slightly larger or smaller values compared to other spectra for the same experimental setup. In the following this effect is called gain drift. It is quantified by investigating the variation in the mean value taken over all frequency bins in one spectrum. The mean for individual spectra varies on the percent level, making this effect subdominant.

3.3.4 Boost Factor Estimation for the CB 200 Setup

For a closed system like CB 200 obtaining boost factors directly using a bead-pull measurement as described in sec. 3.2 is very challenging. Therefore another strategy was used to obtain boost factors, which is outlined below.

The two relevant components for the procedure are the booster itself as well as the LNA, which is the first component of the receiver chain. The idea is to first obtain and verify a 1D model for the booster, using VNA reflectivity measurements, and a

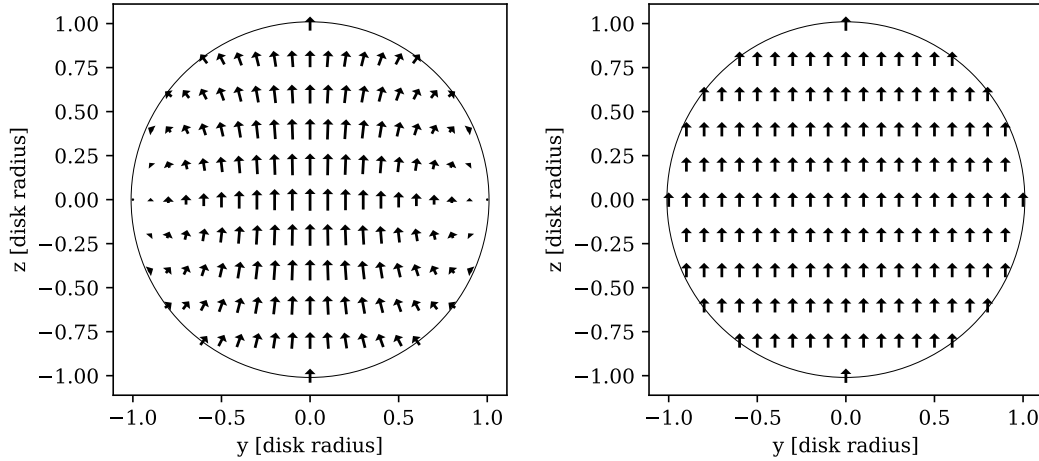


Figure 3.7: Comparison between TE_{11} mode of a circular waveguide (left) and axion induced electric field (right). The direction of the arrows indicates the direction of \mathbf{E} , their length the absolute value $|\mathbf{E}|$.

LNA noise model separately before combining them in order to evaluate the boost factor of the overall booster-LNA system.

A 1D model for the booster can yield a reliable value for the boost factor if at frequencies with high boost factors the booster contains only the mode used for the model. As mentioned earlier, in principle due to the large diameter of CB 200 multiple different modes can propagate in the system. A dip in the reflectivity spectrum like in fig. 3.2 indicates some resonance for one of the modes propagating in the booster. Multiple modes being allowed in the booster therefore implies the existence of many more dips in the reflectivity spectrum related to other resonances in the setup.

The receiver system is designed for the TE_{11} mode which has a large overlap with the axion induced current as shown in fig. 3.7. The latter is strictly (anti-)parallel to the external magnetic field with constant $|E|$ over the whole emitting surface. The TE_{11} mode however has to be perpendicular to the perfectly reflecting boundary. One has to account for this difference using the form factor η introduced earlier, which for the modes shown in fig. 3.7 amounts to $\eta \approx 0.84$.

To reliably single out the correct reflectivity dip from other structures in the reflectivity spectrum, one can leverage the fact that the boost peak is highly sensitive to the mirror location and less sensitive to the location of the booster relative to the receiver system. Adjusting the mirror position moves the frequency position of the reflectivity dip corresponding to the boost factor peak but leaves other dips relatively constant. The opposite effect is observed when moving the booster with respect to the taper connecting to the receiver system.

To further ensure that the feature observed in the power spectrum does indeed correspond to the wanted booster mode, again the bead-pull method can be of use: By performing a scan parallel to the disks between last disk and mirror, one can investigate if the field distribution at the frequency of the reflectivity dip matches the expectation for $|E|^2$ for the TE_{11} mode. For the configurations of CB 200 used at MORPURGO this was verified in advance using scans along $z = 0$ (as defined in fig. 3.7).




Figure 3.8: 1D power boost factors obtained via the fitting procedure described in the text. Different colours for individual boost factor curves are used only to enhance visibility.

A fit on the reflectivity spectrum using a 1D booster model is not capable of accounting for 3D effects and therefore cannot accurately reproduce the full reflectivity spectrum. However, this is not necessary since having singled out the relevant frequency range for the boost factor peak, it is sufficient to accurately reproduce the corresponding reflectivity dip. The model is constructed using PathWave Advanced Design System (ADS) by Keysight allowing mirror and disk properties to vary freely within their previously measured uncertainties.

The LNA model is also implemented using ADS. It contains several effective parameters, which are fitted using separate LNA calibration measurements. The uncertainty in these parameters is dominated by the power calibration corresponding to these measurements.

The distance between booster and LNA can be determined by using a fit on a booster noise measurement. A single parameter fit only on the distance is employed, all other relevant model parameters are drawn from their respective probability distributions as determined in previous steps. The fitting procedure is repeated multiple times resulting in boost factor curves such as the ones displayed in fig. 3.8.

A Normal distribution is fitted to these boost factor curves for every frequency bin individually. The resulting mean and standard deviation are then used in the further analysis process to assess the uncertainty corresponding to the boost factor.

4

Foundational Statistical Methods

There exists an age-old feud among statisticians on how to best interpret statistical findings. The gist of it can be summarised using the example of a coin toss:

Classically, statisticians assume a specific hypothesis and investigate how likely a specific measurement is under this hypothesis. They may use sentences like: “Assuming this is a fair coin, the five heads we got in a row are pretty unlikely.” This line of thinking implies knowledge of what would happen, given the hypothesis is true, if the measurement was repeated infinitely many times. Just as in the example, this is actually the case more often than not. This approach is known as frequentist statistics, since one can tell how frequently to expect a specific measurement.

A central criticism offered by other statisticians is that one usually does not care how likely a specific measurement was given some hypothesis which may or may not be true. What one should care about is what the measurement tells us about the hypothesis being true or not. One should say: “Given the five heads we got in a row, this being a fair coin is pretty unlikely.” A statement like this always depends on one’s beliefs about the hypothesis prior to the measurement. After all, one may be more inclined to trust the fairness of the coin if taking it from one’s own purse than if getting it from a shady magician. This dependence on prior belief is formalised in Bayes theorem, making this approach known as Bayesian statistics.

In this thesis both frequentist as well as Bayesian methods are used, the former in chs. 7 and 8, the latter in ch. 6. Given how crucial a good understanding of the statistical methods employed there is for the interpretation of the results, the relevant basics are introduced here. Starting with frequentist hypothesis testing in sec. 4.1, the concept of test statistics is introduced and type I and type II errors are covered. In sec. 4.2 the look-elsewhere effect is touched upon, explaining necessary corrections when conducting multiple investigations into equivalent hypotheses. Sec. 4.3 then




Figure 4.1: Example scenario of a frequentist hypothesis test. Under the hypothesis H_0 (H_1) the observed value for the test statistic T_{meas} is drawn from a Gaussian with standard deviation $\sigma = 1$ and mean $T_0 = 0$ ($T_1 = 5.65$), shown in green (orange). To determine which distribution it was drawn from a threshold T_{crit} (black, dashed line) can be employed. For $T_{\text{meas}} < T_{\text{crit}}$ one assumes H_0 to be true, H_1 is assumed to be true for $T_{\text{meas}} > T_{\text{crit}}$. The false positive and false negative rates can be obtained by integrating over the hatched green and orange areas, respectively.

switches to the Bayesian interpretation for parameter inference, explicitly illuminating the importance of physically well-motivated prior beliefs.

Textbooks covering frequentist and Bayesian statistical methods separately are abundant, for a detailed comparison of the two see e.g. [113].

4.1 Methods for Frequentist Hypothesis Testing

In hypothesis testing two hypotheses are compared with each other, which are usually called null hypothesis H_0 and alternative hypothesis H_1 . In the context of an axion haloscope one can for example choose H_0 as “no axion exists or its coupling to photons $g_{a\gamma}$ equals zero” and H_1 e.g. as “the axion exists at the observed frequency, makes up all of the local dark matter and couples to photons as $g_{a\gamma} = 10^{-14} \text{ GeV}^{-1}$ ”. Now an observable parameter T is needed, the distribution of which is known in the limit of infinitely many measurements, under the assumption of H_0 and H_1 . In the coin toss example above this was the number of heads in five coin tosses, for MADMAX the quantity is related to the measured power at a specific frequency.

The parameter T in this context is called a test statistics. In principle any metric can be chosen for which a probability distribution can be computed under H_0 and H_1 yielding different results for the two hypotheses. For most axion haloscopes including MADMAX the distributions at each measured frequency bin are Gaussians with equal standard deviations σ but different means T_0 and T_1 (see fig. 4.1). The experimental design determines the relationship between observable T and physically meaningful quantity $g_{a\gamma}$. This relationship must be tractable and a good design leads to large $|T_0 - T_1|/\sigma$. The standard deviation σ typically arises due to statistical fluctuations and therefore decreases with measurement time.

Equipped with the two probability distributions for the test statistics under H_0 and H_1 , one can introduce a threshold value T_{crit} . For a measurement T_{meas} below this value H_0 is accepted, above it H_0 is rejected and H_1 is accepted instead. It is now possible to analytically calculate how often one makes the wrong choice by integrating over the relevant parts of the probability distributions for H_0 or H_1 . Rejecting H_0 even though it is true is called false positive or type I error (hatched green in fig. 4.1). Accepting H_0 even though it is false is called false negative or type II error (hatched orange). Shifting the rejection threshold trades type I error for type II error or vice versa.

In physics, H_0 is typically considered to be the scenario involving no new physics. Rejecting H_0 therefore requires extraordinary evidence typically considered to be between 3σ and 5σ globally, corresponding to a false positive rate between 2.7×10^{-3} and 5.7×10^{-7} . Requirements for wrongly rejecting H_1 are typically much less strict, usually a false negative rate between 5% and 10% is deemed acceptable.

If H_0 is not rejected this threshold based approach leads to an exclusion limit which is independent of statistical fluctuations in the data. In an axion context for every individual frequency bin T_{crit} is used to construct the limit instead of the individual T_{meas} , leading to a “smooth” exclusion limit. This approach was employed e.g. by the HAYSTAC experiment [83].

The ADMX experiment employs a simpler, yet equally valid approach [114], as will be done by the MADMAX experiment [3]. Instead of comparing two hypotheses it is possible to claim incompatibility of the data with a single hypothesis. After carefully investigating data with large T_{meas} in order to not accidentally overlook potential signals, it is statistically sound to interpret any T_{meas} as the data being incompatible with a hypothesis H_1 which is expected to lead to $T > T_{\text{meas}}$ in e.g. 95% of cases.

There is one caveat when using this approach: Assume the no-axion hypothesis H_0 to be true, which means that T_{meas} is a sample of the corresponding distribution around T_0 . Trivially, 5% of the T_{meas} will then be smaller than 95% of the values of T expected under H_0 , wrongly excluding H_0 for these measurements. For many independent frequency points this implies that 5% of them would show an exclusion limit down to $g_{a\gamma} = 0 \text{ GeV}^{-1}$. The interpretation of such a limit is not intuitive, in spite of it being accurate under the stated assumptions. The MADMAX limits shown in ch. 8 are therefore conservatively restricted to not exceed the median expected limit T_0 by more than 1σ .

4.2 Correcting Look-Elsewhere Effect

Above it was stated that claiming detection of new physics requires a very small false positive rate down to $f_p = 5.7 \times 10^{-7}$. If only one data point is observed, that is sampled from a Normal distribution, this f_p corresponds to a 5σ excess. For a set of n independent data points, however, the statistical probability that any one of them exceeds a specific σ level increases as (e.g. [115])

$$f_p^{\text{global}} = 1 - \left(1 - f_p^{\text{local}}\right)^n, \quad (4.1)$$

where the superscripts “global” and “local” distinguish between the probability when considering an excess in any one of the data points or a specific data point only.

When data points are correlated, an effective number of independent points n_{eff} needs to be found. A standard procedure for an ordered set of equidistant data points X_i is to divide the number of correlated data points n by the the sum of the auto-correlation function for all possible offsets [116]:

$$n_{\text{eff}} = \frac{n}{1 + 2 \sum_o \text{cor}(X_i, X_{i+o})}, \quad (4.2)$$

with offset o . This procedure can be applied to the actually measured set of data points or, as is done in sec. 7.5, be leveraged in a Monte-Carlo simulation for simulated datasets. The latter approach is beneficial to increase the statistical significance for sets containing only a small number of data points or to avoid unwanted experimental artefacts.

4.3 Axion Priors for Bayesian Inference

In the frequentist framework above the probability distributions showed how often to expect a specific measurement for a parameter T when repeating an experiment assuming some hypothesis to be true. A confidence interval or limit obtained with this method should be interpreted in terms of the compatibility of the data with some set of model parameters. It does not mean that there is a 95% probability that the true model parameter is within the interval or limit.

The latter interpretation is only valid for Bayesian credible intervals or limits. They are obtained from the posterior probability, or simply posterior, $P(\theta|x)$ of a set of model parameters θ given a measurement x . To calculate the posterior one has to apply Bayes' theorem

$$P(\theta|x) \propto P(x|\theta)P(\theta), \quad (4.3)$$

where the probability of a measurement given a set of model parameters $P(x|\theta)$ is called the likelihood and $P(\theta)$ is the probability distribution of the model parameters prior to the measurement, commonly simply called prior.

The choice of the prior distribution can substantially affect credible intervals and must therefore be conducted with care. It can be informed by theoretical predictions and/or previous measurements. In the absence of any prior information for one of the model parameters an uninformative prior must be chosen. It is subject to open debate what the most uninformative prior is for different scenarios.

In an axion haloscope context there are three independent parameters to consider: the axion mass m_a , determining the frequency at which an axion signal can be observed, the dark matter velocity dispersion σ_{DM} determining the FWHM of the axion peak, and the axion-photon coupling $g_{a\gamma}$, determining the height of the peak¹ (see also secs. 2.4 and 6.2).

The axion mass is only slightly constrained from theoretical considerations. For global fits the considerations of sec. 2.3 can inform a potential prior choice. In practice only a very small piece of the axion mass parameter space can be constrained

¹Technically also the local DM density and the fraction of DM made out of axions affect the peak height in the same way as $g_{a\gamma}$. A Bayesian analysis should therefore construct a prior for the combination of these three degenerate contributions. In practice, however, axion haloscopes typically assume axions to constitute all of DM and a fixed local density.

by most haloscopes. It can therefore be helpful to restrict the prior range to the observable values. While this does not lead to an accurate global credible limit, locally, i.e. assuming an axion to exist in the measurable mass range, the obtained limits are still accurate.

The DM velocity dispersion is well-determined from measurements leading to tight priors. For details see sec. 2.4.

The axion-photon coupling for a given m_a is determined by the anomaly ratio E/N according to eq. (1.17). Determining a prior distribution for $g_{a\gamma}$ can therefore be translated into finding all possible theoretical predictions for E/N . For the minimal QCD axion this “distribution” would only consist of three discrete values. In subsec. 1.3.3 previous efforts to compute anomaly ratios for KSVZ-type models were mentioned. Doing the same also for DFSZ-type models is necessary in order to be able to compute a general non-minimal QCD axion prior distribution. Such a prior is crucial for applying a Bayesian analysis to any haloscope dataset. This is the main motivation behind ch. 5.

Beyond the QCD axion, particles with similar properties can arise from many theories beyond the SM, among them string theories [117]. These axion-like particles (ALPs) do not solve the strong CP problem but include an effective coupling to photons. Specific models can lead to various predictions for the relation between ALP mass and photon coupling (see e.g. [118] for some examples). Often, however, ALPs are considered as model-independent, which means that the photon coupling is assumed to be independent of the axion mass and could be arbitrarily small/ big within existing experimental bounds. When the magnitude of the photon coupling is unknown, arguably a prior that treats all decades equally (i.e. $P(10^{-19} \leq g_{a\gamma} [\text{GeV}^{-1}] \leq 10^{-18}) = P(10^{-18} \leq g_{a\gamma} [\text{GeV}^{-1}] \leq 10^{-17}) = \dots$) is the most natural choice. Such a prior has an infinite volume as $g_{a\gamma} \rightarrow 0$. If the data does not contain conclusive evidence for an axion, the likelihood does not vanish in this limit. The resulting posterior will therefore not be normalisable [119] making it impossible to obtain Bayesian exclusion limits for model-independent ALPs using the standard methodology.

The constraining capability of the data can instead be investigated by calculating the ratio between prior and posterior probability. Parameter space in which the posterior is suppressed compared to the prior can be deemed disfavoured by the data, making it possible to construct limits using the data, even though they are not the standard Bayesian credible limits. Methods using this approach exist [119] and similar methods have already been used in the context of axion research [120].

5

Non-minimal DFSZ-Type Axion Models

In subsec. 1.3.1 minimal DFSZ axion models have been investigated, which introduce two Higgs doublets as well as a Higgs singlet to solve the strong CP problem with an invisible axion. The minimal KSVZ model has been described, introducing next to the Higgs singlet an additional heavy quark in a specific representation, and extensive efforts to generalise this class of models by allowing multiple heavy quarks in various representations have been recapitulated (subsec. 1.3.3 or [9]). This leaves the glaring questions: How can the minimal DFSZ models be generalised? Which axion-photon couplings are generated in such models and how do these findings compare to the case of non-minimal KSVZ-type models?

The paper [1], co-written by the author and followed closely by the present chapter, systematically constructs non-minimal DFSZ-type models. It exploits the fact that the only model dependence of axion-photon couplings is given by the ratio of EM and QCD anomaly coefficients (see eq. (1.19)), which in turn only depend on representations of the contributing fields. For DFSZ-type models where only SM fermions contribute to the anomalies, calculating axion-photon couplings reduces to finding possible PQ charges for the SM fermions by solving small linear systems of equations (LSEs). This allows the construction of a vast catalogue of non-minimal DFSZ-type models, considering models with up to $n_D = 9$ Higgs doublets corresponding to the nine right-handed fermions of the SM. For convenience, a new notation is introduced, explicitly stating the number of Higgs doublets in specific models as DFSZ_{n_D} .

Key findings gleaned from this catalogue include axion-to photon couplings of minimal DFSZ models remaining favoured in spite of numerous other photon couplings being allowed. Unlike in the minimal case, photon couplings for non-minimal DFSZ-type models are found to be similar to the KSVZ-type case. Models with domain wall numbers of unity, which are phenomenologically favoured, may also be easier

to detect, displaying significantly enhanced photon couplings. These statements are quantified by constructing bands and lower limits in the parameter space of axion mass vs. photon coupling to be reached by experiments that aspire to be sensitive to a certain fraction of non-minimal models. Furthermore, the systematic method presented here facilitates to specifically search for models with large photon couplings, making it possible to surpass previous maximal values [8].

The considerations above are interesting in their own right from a model builder's perspective, but become crucial when applying Bayesian statistics on QCD-sensitive axion experiments. As seen in sec. 4.3, this statistical framework requires theoretical prior knowledge that is updated by experimental data. Only considering minimal invisible axion models in this framework is equivalent to setting the probability of additional Higgs doublets or heavy quarks equal to exactly zero. Undoubtedly, this is a valid standpoint some physicists will take, invoking occam's razor as an argument. However, one may argue that assigning zero probability to anything is a very strong statement, especially considering one additional Higgs doublet or heavy quark has already been added to the SM in order to construct the invisible axion model. As any discussion on prior belief, this debate does not have a definitive solution, which is why it is important to provide raw data and generating code for the findings presented here at the end of subsec. 5.4.2.

To start with, in sec. 5.1 it is shown that the calculation of anomaly ratios is similar to the minimal DFSZ case. Sec. 5.2 then covers the question what in detail constitutes a valid model, introducing the concept of multiplicity and discussing possible selection criteria. In sec. 5.3 the procedure to compute all possible anomaly ratios is applied explicitly to the case of three Higgs doublets before moving on to larger numbers of doublets. For more than seven Higgs doublets the full model catalogue cannot be computed, therefore this case is treated separately. Assumptions are stated which are necessary to interpret the distributions from a probabilistic standpoint. The results are then compared with the literature values for KSVZ-type models. Axion bands and limits are constructed in sec. 5.4 together with a discussion on other implications for axion searches. The chapter concludes with sec. 5.5.

5.1 Calculating the PQ Charges in DFSZ-Type Models

In eq. (1.37) the anomaly ratio for the DFSZ axion model containing two Higgs doublets and one Higgs singlet was expressed in terms of PQ charges of fermions. When the Glashow-Weinberg condition is not imposed, the equation generalises to

$$\frac{E}{N} = \frac{2}{3} + 2 \frac{\sum_{u_i} \chi_{u_i} + \sum_{e_i} \chi_{e_i}}{\sum_{u_i} \chi_{u_i} + \sum_{d_i} \chi_{d_i}}, \quad (5.1)$$

depending on PQ charges $\chi_j \equiv \chi_{jL} - \chi_{jR}$, with $j \in \{u_i, d_i, e_i\}$ and i running over the three fermion generations. For SM fermions excluding neutrinos, $n_D \leq 9$ unique χ_j can be allowed, where n_D corresponds to the number of Higgs doublets introduced in a specific DFSZ-type theory.

These PQ charges need to be determined. The coupling in the Yukawa sector together with the requirement of PQ invariance translates the PQ charges of the Higgs doublets to the fermions. If one Higgs doublet couples to multiple fermions this

immediately implies that their PQ charges are equal. The PQ charges of the n_D Higgs doublets are fixed by the following requirements, similarly to what was discussed in sec. 1.3:

Orthogonality between PQ current and weak hypercharge current

The relations imposed by the orthogonality condition, eqs. (1.32)-(1.35), hold equivalently for more than two Higgs doublets if f runs over all Higgs fields in the theory.

Yukawa sectors of the form of eq. (1.25) directly fix the weak hypercharges of the doublets to

$$-Y_{H_{u_i}} = Y_{H_{d_i}} = Y_{H_{e_i}} = \frac{1}{2} \quad (5.2)$$

and the singlet has $Y_S = 0$. One can furthermore set the vacuum expectation values (VEVs) to $v_f = 1$ without loss of generality because the anomalies do not depend on them.

Orthogonality between PQ and hypercharge current yields one condition on the PQ charges of the Higgs fields, however for n_D doublets $n_D + 1$ conditions are required to fix the PQ charges of all Higgs and thereby all fermion fields. The next paragraph introduces $n_D - 1$ additional conditions, which is sufficient for calculating the anomaly ratio. With n_D conditions all PQ charges of the Higgs doublets can be expressed in terms of χ_S . Since both E and N are linear in the PQ charge, χ_S cancels out. However other quantities such as the DW number do depend on χ_S and below it will be used to get constraints on the singlet charge.

PQ invariance of explicit symmetry breaking potential

The standard kinetic scalar Lagrangian displays a $U(1)^{n_D+1}$ symmetry, which needs to be explicitly broken down to $U(1)_{\text{PQ}} \times U(1)_Y$, similar to the case of the minimal DFSZ model. Otherwise the PQ current is not well-defined and Goldstone bosons with decay constants of order of the electroweak scale can arise. It is therefore necessary to introduce a potential V_{eb} that includes at least $n_D - 1$ terms which explicitly break this $U(1)^{n_D+1}$ symmetry.

These terms can be divided into two categories: Terms with two Higgs doublets and two times the singlet which are henceforth called bilinear terms and denoted symbolically as $HHSS$, and quadrilinear terms $HHHH$ consisting of four Higgs doublets. Higher orders in the scalar are not considered, since these terms are not renormalizable. At least one term of the form $HHSS$ is necessary to render the axion invisible, the form of the remaining terms is in principle free.

For the minimal DFSZ model only one term, $\propto H_u H_d S^{\dagger 2}$ leads to a unique, non-trivial condition on the PQ charges. For $n_D > 2$ in theory there can be almost infinitely many terms in the potential, but note that the requirement of PQ invariance translates each term into a condition on the PQ charges of the Higgs particles. Treating them as unknowns there are $n_D + 1$ free parameters. With the conditions above and a (for now) arbitrary normalisation for χ_S , $n_D - 1$ linearly independent conditions from the explicit symmetry breaking potential lead to a well-defined, solvable system. Less conditions lead to undesired massless states since the $U(1)^{n_D+1}$ symmetry is not broken down enough, more conditions usually lead to inconsistencies. However, on top of the minimal potential containing exactly $n_D - 1$ terms which

are all linearly independent, additional terms leading to linearly dependent conditions can be added. This will be investigated further in the next section.

Well-definiteness of DW number

A consistency condition on χ_S follows from the DW number N_{DW} being integer-valued in theories where the axion is a linear combination of fields.

In the low-energy regime the QCD anomaly induces a potential to the axion (compare eq. (1.15) and the discussion beneath), which is periodic in $[0, 2\pi f_a[$, with $f_a = v_a/2N$, while the axion is an angular field defined over $[0, 2\pi v_a[$. The periodic nature of the potential leads to discrete vacua, and the number of vacua within a single 2π loop is known as the DW number N_{DW} , which can be read off to be $N_{\text{DW}} = 2N$. In the language of symmetries, this potential explicitly breaks the original $U(1)_{\text{PQ}}$ down to the discrete group $\mathbb{Z}_{N_{\text{DW}}}$ under which the axion transforms as $a \rightarrow a + 2\pi n f_a$ with $n \in \mathbb{Z}$. The DW number is encoded in this transformation and represents the value of n that results in a single loop with a circumference of $2\pi v_a$.

However, there is a caveat in theories where the axion is a linear combination of angular modes a_f [121]. Each angular mode possesses its own residual cyclic symmetry due to explicit breaking, given by $a_f \rightarrow a_f + 2\pi n_f v_f$ where $n_f \in \mathbb{Z}$. Applying these symmetry transformations on both sides of the definition for the axion field (eq. (1.33)) one can extract the DW number as

$$N_{\text{DW}} = 2N \frac{\sum n_f \chi_f v_f^2}{\sum \chi_f^2 v_f^2}. \quad (5.3)$$

In order for the DW number to be integer, the fraction in this expression has to be integer (which, as it turns out, can be chosen to be one without loss of generality). Setting $n_f = \chi_f$ is the simplest way to achieve this. However, a less restrictive alternative can be obtained by using the orthogonality condition to remove one of the v_f terms in both the numerator and the denominator of the fraction, and then comparing terms with the same v_f^2 . Performing this in DFSZ₂ for simplicity, with $f = u, d, S$, and requiring the fraction to be unity, one finds

$$n_S = \chi_S, \quad (5.4)$$

$$n_u + n_d = \chi_u + \chi_d = 2\chi_S. \quad (5.5)$$

In the second equality the PQ invariance from the unique V_{eb} term was used, i.e. $H_u H_d S^{\dagger 2}$. The consideration above demonstrates that the residual cyclic symmetries of the underlying angular modes result in the condition $\chi_S \in \mathbb{Z}$ in the minimal DFSZ model.

When repeating this procedure for larger numbers of doublets, one finds that eq. (5.4) is always present along with additional relations similar to eq. (5.5). These relations imply that for the DW number to be integer, both χ_S and certain combinations of PQ charges must also be integer. In particular, $2N = \sum_i \chi_{u_i} + \chi_{d_i}$ and $\sum_i \chi_{u_i} + \chi_{e_i}$ have to satisfy this criterion. The key difference for non-minimal models is that fulfilling all conditions becomes more restrictive compared to the minimal case. For instance, it could require the minimal value of χ_S to be integer and greater than one.

To summarise, in theories where the axion is a linear combination of fields, the DW number can still be expressed as $N_{\text{DW}} = 2N$, provided that the fraction in eq. (5.3) equals one. This additional condition introduces the requirement for χ_S to be an integer, although not necessarily one. For the purposes of the discussion below it is convenient to set χ_S to its minimum possible value. Consequently, the DW number is defined as

$$N_{\text{DW}} = \min \text{ positive integer } \{2N\} , \quad (5.6)$$

with this definition being used for the remainder of this work.

5.2 What Makes the Model?

This section presents an overview over all relevant criteria determining what constitutes a valid theory.

5.2.1 Higgs Sector

Firstly, only $n_D \leq 9$ are considered, together with a single singlet to make the axion invisible. The anomaly of the PQ current is dependent solely on the difference between the PQ charges of left- and right-handed fermions. One can therefore without loss of generality set the PQ charges of the left-handed fermions to zero, leaving nine PQ-charged, right-handed fermions. If two Higgs doublets couple to the same fermion, they must have equal PQ charge to satisfy PQ invariance of the Yukawa sector. Therefore maximal freedom is reached for nine Higgs doublets from the standpoint of PQ charges.

Neutrinos could alter this picture. Dirac mass terms for neutrinos could be generated from some of the Higgs doublets, altering the PQ charges of the other fermions and thus indirectly affect the anomaly ratio. Since it is currently unknown how neutrino masses are generated, this analysis neglects these intricacies and sets PQ charges of neutrinos to zero. Subsec. 5.3.4 will show that the conclusions of this chapter are most likely unaffected by this omission.

5.2.2 Multiplicity

One of the standard principles in constructing models is to incorporate all feasible terms that adhere to the given symmetries. If, for some reason, certain terms are not included at tree-level, these terms will be generated at higher orders without protection from an underlying symmetry. In terms of V_{eb} , this implies that potentials generating the same PQ charges should not be considered distinct as they can simply be added. This concept can be explained in terms of conditions and LSEs. In the construction detailed in sec. 5.1 exactly $n_D - 1$ terms were required in the explicit breaking potential. Fewer terms would result in under-determined systems and for too many independent terms the resulting system would be over-determined, with $\chi_f = 0$ for all f , and would thus fail to solve the strong CP problem. Linearly dependent terms giving rise to redundant conditions, however, can always be added to the potential without leading to over-determined systems. Potentials with linearly dependent terms can be broken down to smaller, still well-defined sub-systems, which lead to different underlying LSEs. However, all of these LSEs have the same solution,

i.e. lead to the same PQ charges, and should therefore not be counted as different theories.

This reasoning also has implications for the Yukawa sector. Although the construction described here involves coupling a single doublet to each right-handed fermion, in principle, multiple doublets can be coupled to the same right-handed fermion. Therefore, the Yukawa sector for each possible set of PQ charges is completed post hoc. For instance, one can complete the Yukawa sector as follows: If one solution of the LSE yields $\chi_{d1} = \chi_{e1}$, then the corresponding completed Yukawa sector is given by

$$\begin{aligned} y_{1j}^d H^{d1} \bar{Q}_L^1 d_R^j &\longrightarrow (y_{1j}^d H^{d1} + \tilde{y}_{1j}^d H^{e1}) \bar{Q}_L^1 d_R^j, \\ y_{1j}^e H^{e1} \bar{E}_L^1 e_R^j &\longrightarrow (\tilde{y}_{1j}^e H^{d1} + y_{1j}^e H^{e1}) \bar{E}_L^1 e_R^j. \end{aligned} \quad (5.7)$$

This ensures that all Yukawa terms compatible with a given solution are included (such as cross-couplings where for instance up-type doublets couple to down-type fermions). Furthermore, since each set of PQ charges becomes unique after adding the potentials, the completed Yukawa sector with all compatible couplings is determined uniquely. Therefore, for counting the multiplicity, a model is defined as a Yukawa sector that was completed as demonstrated above and it includes a V_{eb} for which all potentials with a specific set of PQ charges were added. Crucially, the multiplicity as defined above does not count theories as distinct, if they lead to the same PQ charges. The final step in the construction procedure is to calculate the anomaly ratio for each model and count how often each anomaly ratio arises.

5.2.3 Selection Criteria for DFSZ-Type Models

Having specified what exactly constitutes a model, the next question becomes whether there are any (phenomenological) selection criteria that can be used in order to extract a subset of preferred axion models. An equivalent discussion for KSVZ-type models can be found in [8] or [9] and is recapitulated in subsec. 1.3.3.

$N \neq 0$

Sets of PQ charges leading to $N = 0$ do not solve the strong CP problem and are therefore discarded.

A Priori Model Selection

One could explore models with multiple singlets or more than $n_D = 9$ doublets, which do not couple to the SM fermions. Models like this leave eq. (5.1) unchanged, but allow to obtain very large PQ charges [8, 122] by introducing additional solutions for the PQ charges. Adhering to the more narrow definition of DFSZ-type models given in the previous section narrows down the range of models under consideration and could therefore be considered as a kind of selection criterion.

Running of EW Coupling

The presence of additional Higgs doublets can potentially affect the running of the electroweak gauge coupling. A strict upper limit on the number of doublets, determined by the requirement of preserving asymptotic freedom and avoiding Landau poles

below the Planck scale, however only arises at $n_D \sim 50$ [8], which is much higher than the number of doublets considered in this study.

Domain Wall Problem

The periodic low energy potential of the axion (1.15) explicitly breaks the original PQ symmetry, reducing it to the discrete group $\mathbb{Z}_{N_{\text{DW}}}$, which is subsequently spontaneously broken by one of the vacua. As a consequence DWs form as described in subsec. 2.2.3, which can dominate the energy density of the universe, if $N_{\text{DW}} > 1$. Therefore, one possible criterion for selecting axion models is to impose $N_{\text{DW}} = 1$.

However, if the PQ symmetry is broken during or before inflation, the DWs do not form within our Hubble sphere, effectively avoiding the issue. Also in scenarios where the PQ symmetry breaking occurs after inflation, solutions have been proposed to circumvent the DW problem [123, 124]. Therefore, $N_{\text{DW}} = 1$ is here not considered to have a sufficient level of generality to represent a necessary selection criterion for the main analysis. Nevertheless, the DW numbers for DFSZ₃ to DFSZ₇ are calculated, demonstrating the influence of this selection criterion.

Flavour-Changing-Neutral-Currents (FCNC)

FCNCs are a generic feature of models with multiple Higgs doublets. They are subject to strong experimental constraints [125], which could potentially impose significant limitations on the number of viable DFSZ-type models. In the following various known methods to avoid these FCNCs are discussed (see [126] for a review):

FCNCs can arise only when the Yukawa couplings break generation invariance, i.e. when right-handed fermions of a given electric charge couple to different Higgs doublets. The Weinberg-Glashow-Paschos condition [44, 45] demands that this is not the case. Imposing this condition is known as natural flavour conservation and effectively sets several Yukawa couplings to zero, resulting in $n_D - 3$ decoupled Higgs doublets for $n_D > 3$. Hence, for DFSZ-type models as defined above, natural flavor conservation can only be achieved when $n_D \leq 3$.

To avoid FCNCs at tree level it is sufficient to impose an alignment condition [127–129]. Requiring the Yukawa matrices of each right-handed fermion to be proportional to each other makes them simultaneously diagonalisable in the fermion mass eigenbasis.

Going beyond flavour alignment, specific forms of Yukawa matrices in flavour space can be found, which allow viable SM mass and mixing phenomenology while sufficiently suppressing tree-level FCNCs [130]. The implementation of these mass matrix ansätze as well as of natural flavour conservation often relies on imposing (discrete) symmetries, which also protect the flavour structure from quantum corrections. However, imposing additional symmetries on the scalar potential spoils the so-called decoupling property of general multi-Higgs doublet models [131]. In order to avoid FCNCs using these solutions, it would be necessary to systematically determine which models possess discrete symmetries that simultaneously avoid FCNCs and allow for a decoupling limit. Due to the large number of models and the absence of a catalogue of possible symmetries for $n_D > 3$ [126], conducting such an analysis is beyond the scope of this work.

5.3 Anomaly Ratio Distributions

5.3.1 A Cooking Recipe

The previous sections have presented a review on which conditions fix the PQ charges in a DFSZ framework. They outlined how different V_{eb} can lead to the same set of PQ charges and many different sets of charges to the same anomaly ratio and what exactly should be counted as a single, unique theory. The different steps of the procedure can be summarised in a concise fashion:

1. Specify the Yukawa sector for a fixed n_D by coupling one doublet to each right-handed fermion. This exhausts the maximal freedom regarding the anomaly ratio.
2. Write down all possible V_{eb} with $n_D - 1$ terms.
3. Solve all associated LSEs to find all possible sets of PQ charges. Underdetermined systems are discarded.
4. Add the potentials of all equal PQ charges to get the most general potential associated with a particular solution. This defines one model for the sake of counting the multiplicity.
5. For each model complete the Yukawa sector by adding all Yukawa terms compatible with the PQ- and hypercharges.
6. For each model calculate the anomaly ratio and count its multiplicity.

Step 5 is only relevant from a model building perspective and not for the acquired solutions. It is therefore skipped in practice and whenever a Yukawa sector is mentioned in the remainder of this work, it refers to the one specified in the first step. The above constitutes a recipe to calculate all possibilities for the PQ charges of the fermions given a specific Yukawa sector. How this is applied in practice is shown below for a simple example with three Higgs doublets.

5.3.2 Example: Three Higgs Doublets

For DFSZ₃ with the Weinberg-Glashow-Paschos condition imposed, i.e. with one Higgs doublet per type of fermion, it is possible to explicitly demonstrate the approach outlined in the previous section. For this model there are three possible bilinears, namely $(H_u H_d)$, $(H_u H_e)$ and $(H_d H_e^\dagger)$, together with their complex conjugates. Each bilinear can either be coupled to the singlet, which results in 6 different terms of the form $HHSS$, or to another bilinear, which results in 36 different quadrilaterals of the form $HHHH$. For the latter case, removing terms that are related by Hermitian conjugation and terms that result in no condition reduces the number to 9 (see tab. 5.1). For $n_D = 3$, the breaking potential consists of either one $HHSS$ and one $HHHH$ term or two $HHSS$ terms. For the former, there are *a priori* 54 possibilities, and for the latter 15, totalling to 69 possibilities for V_{eb} (see tab. 5.2).

The resulting 3×3 LSEs consist of the orthogonality relation, $\chi_u v_u^2 - \chi_d v_d^2 - \chi_e v_e^2 = 0$, and the two conditions coming from the potential. Solving the LSEs yields the

and write the subscript of the doublets that couple to a certain fermion to the corresponding position. If one doublet couples to multiple fermions, use the first subscript in the order presented above. For more comprehensive notation use fermion type (up-, down-, or lepton-type, short u , d , or e) and generation (1 to 3). For DFSZ₉, this row vector would be $[u1, u2, u3, d1, d2, d3, e1, e2, e3]$ while for the original DFSZ₂-I model it would be $[u1, u1, u1, d1, d1, d1, d1, d1, d1]$.

5.3.3 Choices for a Statistical Interpretation

Above a procedure is outlined for obtaining many different solutions for the Higgs charges. Applying this procedure one can count and histogram the number of models leading to a specific anomaly ratio. However, ultimately it is desirable to translate this catalogue of models with specific E/N values to a *probability distribution* of anomaly ratios. To this end, relative probabilities of the solutions are required, which are subject to some theoretical prior beliefs. These beliefs guide multiple decisions regarding:

- **The concept of multiplicity**

What exactly counts as a single model? Here the choice is made to count each unique solution for the PQ charges, effectively adding all compatible potential terms, as described in the last two paragraphs of sec. 5.2.

- **The relative probability of different Yukawa sectors given a specific n_D .**

A reasonable choice is to demand all solutions with a given n_D to be equally probable. Similarly, it is also reasonable to consider different initial, uncompleted Yukawa sectors as equally probable. Unfortunately, it is impossible for both conditions to hold simultaneously since different Yukawa sectors can lead to different numbers of possible solutions. In this analysis, the principle of treating solutions as equally probable is adopted, given equal multiplicity and the same n_D . Consequently, no “beauty” arguments are applied, e.g. in favour of equal coupling patterns for different fermion types.

- **The relative probabilities of different n_D .**

In the total anomaly ratio distribution, equal probability is assigned to all values of n_D in the range $2 \leq n_D \leq 9$. This implies that any individual solution for, e.g., DFSZ₃ (which has 16 solutions) is considered to be much more probable than any individual solution for DFSZ₅ (which has 9.7×10^4 solutions). Alternative considerations could be made, such as penalising models with higher n_D , giving larger probability to models that satisfy symmetry arguments (such as DFSZ₃ with one Higgs doublet per fermion type or DFSZ₉ with one Higgs per right-handed fermion), or considering all charge solutions equally probable. In the latter scenario, the resulting histogram would most likely be completely dominated by DFSZ₉ due to its significantly larger number of unique solutions.

The aforementioned arguments all imply a probabilistic approach to model selection, wherein nature “chooses” one of the possible realisations at random. While this notion




Figure 5.1: Anomaly ratio distributions for DFSZ-type models with 4 Higgs doublets. Two Higgs couple to the fermions specified in the panels with the other two Higgs covering the remaining two fermion types invariant with respect to fermion generation. For example, the Yukawa sectors $[u1, u1, u3, d1, d1, d1, e1, e1, e1]$, $[u1, u2, u1, d1, d1, d1, e1, e1, e1]$, and $[u1, u2, u2, d1, d1, d1, e1, e1, e1]$ are all equivalent and have anomaly ratio distributions as shown in the left panel.

itself may be open to critique, in the absence of any compelling underlying physical argument that singles out a particular model, it could be deemed satisfactory.

Even under the assumption of probabilistic model selection, any of these choices is to some extent a matter of taste. For this reason it is important to provide the raw catalogues and generating code as a supplement, so the reader is not dependent on the choice made here (see subsec. 5.4.2).

5.3.4 Results for Four to Seven Higgs Doublets

The exploration of higher numbers of Higgs doublets involves the investigation of various Yukawa sectors. PQ charges and anomaly ratios are computed numerically using the programming language “Julia” [132]. The “StaticArrays” package [133] allows computation of an extremely large number of LSEs very fast without heap memory allocation. Nevertheless, calculating all possible solutions proved prohibitive beyond $n_D = 7$. Therefore the analysis of this case is postponed to the next section, with the focus here being on DFSZ₄ to DFSZ₇.

Fig. 5.1 presents an overview over the anomaly ratio distributions for DFSZ₄ models grouped by the different Yukawa sectors. Each of the histograms shows all models of the specified coupling with the explicit symmetry breaking potential V_{eb} consisting of $k \geq 1$ $HHSS$ - and $3 - k$ $HHHH$ -terms. The result does not depend on fermion generation since the construction of the Higgs charges as well as eq. (5.1) treat all generations equally. Yukawa sectors with special coupling to a lepton have histograms symmetric around $5/3$, while the histograms for up- and down-type special couplings are mirrored around $5/3$.

The reason for this is a symmetry in the construction as well as in eq. (5.1): For all n_D , since all possible Yukawa sectors are considered as mentioned above, every solution has a corresponding one with

$$\begin{aligned} \chi_{\tilde{u}_i} &\rightarrow -\chi_{d_i} \\ \chi_{\tilde{d}_i} &\rightarrow -\chi_{u_i} \end{aligned} \quad (5.10)$$

This is due to up-type and down-type quarks being treated equally in the construction except for the sign of their hypercharges. In the example above, all solutions for the

Yukawa sector $[u1, u1, u3, d1, d1, d1, e1, e1, e1]$ have a corresponding solution in the Yukawa sector $[u1, u1, u1, d1, d1, d3, e1, e1, e1]$ under the above mentioned transformation. Solutions that relate via eq. (5.10) can easily be seen to have anomaly ratios relating by

$$\frac{\widetilde{E}}{N} \rightarrow \frac{10}{3} - \frac{E}{N}, \quad (5.11)$$

which is a mirror symmetry around $\frac{5}{3}$.

If all nine histograms of fig. 5.1 are added up, i.e. if no Yukawa sector is treated preferentially, one obtains the distribution shown in fig. 5.3 (second row, left). Due to the symmetries of the nine contributing Yukawa sectors, the distribution is symmetric around $5/3$ as well. The largest number of models coincides with the two possible values for the DFSZ₂ model: $2/3$ and $8/3$. Both of these statements are true for $n_D \in [4, 7]$, as fig. 5.3 shows (second row, third row left).

With increasing n_D , the number of unique anomaly ratios increases and more extreme E/N values can be found. Anomaly ratios $E/N = 5/3 + k$ with $k \in \mathbb{Z}$ are highly favoured for $n_D \geq 5$ compared to other E/N values, especially for odd k . This very characteristic, peaked spectrum evolves with n_D : E/N values with high probability tend to have their probabilities shrink with increasing n_D , whereas low probability E/N values behave in the opposite manner. In fig. 5.3, one can most easily see this evolution at big anomaly ratios $E/N \gtrsim 10$.

5.3.5 Anomaly Ratios From a Mathematical Perspective

This trend can be understood from a purely mathematical perspective: For DFSZ-type models the anomaly ration (5.1) is a function with nine variables, the values of each of which can be thought of as being drawn from a specific distribution. In fig. 5.2 the effect of using different distributions for the variables is shown on the resulting E/N distribution. A continuous, flat charge distribution of arbitrary width produces a smooth, fat-tailed E/N distribution. If the median of the charges is 0, the median of the distribution is at $5/3$ (fig. 5.2, top three rows). Allowing only positive values for the charges shifts the distribution to higher values, with a median of $8/3$ and makes $E/N < 0$ impossible (fig. 5.2 bottom row). The fewer distinct input values for the charges are used, the more peaked the anomaly ratio structure becomes, i.e. anomaly ratios with high relative probability see their likelihood increased and vice versa. This also leads to fewer possible unique E/N -values.

The continuum limit with its vanishing skewness and positive kurtosis can be approximated in analytic form via a Pearson type VII distribution [134],

$$p\left(\frac{E}{N}\right) = \frac{1}{\alpha B(m - \frac{1}{2}, \frac{1}{2})} \left[1 + \left(\frac{\frac{E}{N} - \lambda}{\alpha}\right)^2 \right]^{-m}, \quad (5.12)$$

with parameters $\lambda = 5/3$, $\alpha = 7/4$, and $m = 1$, and Beta function B with $B(1/2, 1/2) = \pi$.

Following these insights from a mathematical perspective it is expected that the histograms for larger n_D should be smoother, considering that there are more unique solutions (tab. 5.3). Note, however, that this effect neglects the influence of choosing different probabilities for different solutions. Non-uniform probabilities reduce the

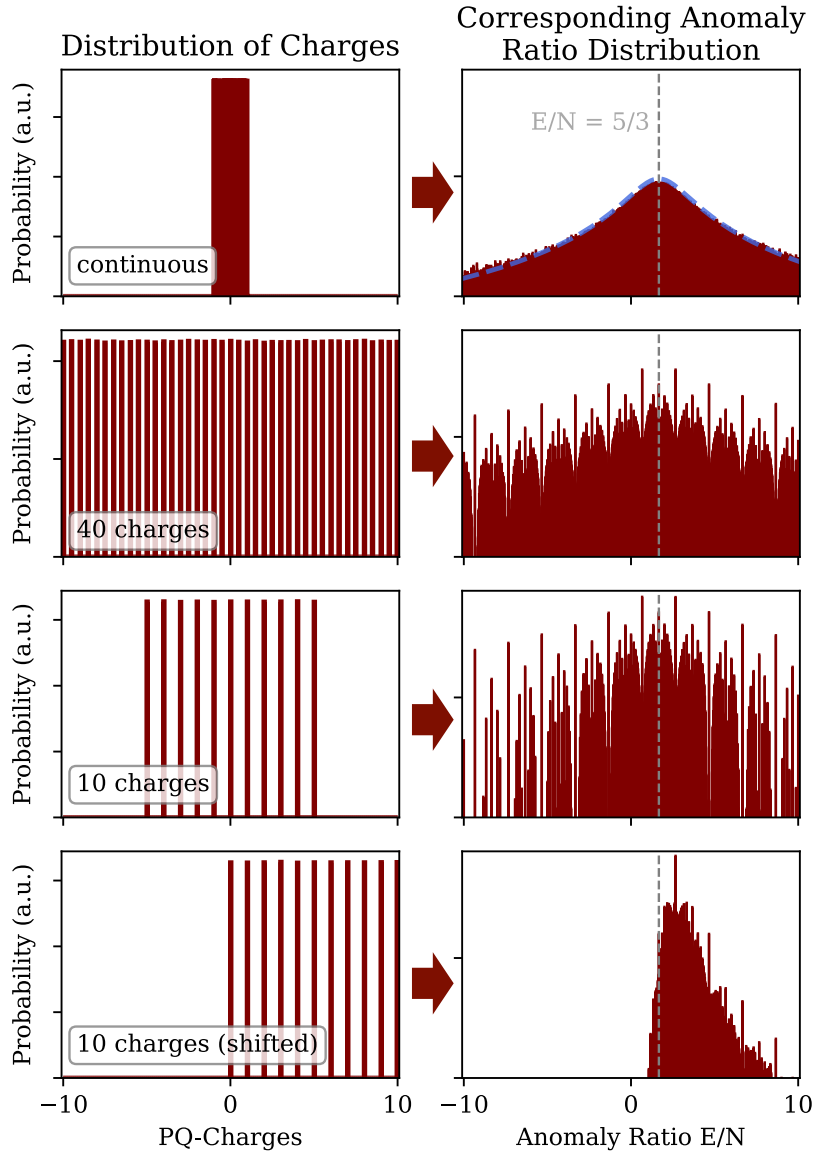


Figure 5.2: Influence of drawing PQ charges from different distributions on the resulting anomaly ratio distribution, using eq. (5.1). More unique charges lead to a smoother anomaly ratio distribution, irrespective of their distribution. Charge distributions centred around 0 produce anomaly ratio distributions centred around $5/3$. The dashed blue line in the top right panel denotes the fit presented in eq. (5.12).




Figure 5.3: Anomaly ratio distributions for different numbers of Higgs doublets. All histograms are symmetric around $5/3$. $n_D \geq 5$ display a characteristic peaked structure, which becomes smoother with increasing n_D . DFSZ₈ and DFSZ₉ could not be fully calculated, the two semi-transparent colours denote the two estimates as discussed in the text.

effective number of different solutions.¹ Using the approach of adding all possible potential terms for one solution of charges leads to more comparable probabilities for the charges than if one had separately considered all potentials with the minimal amount of terms to fix the PQ charges (minimal potentials). Therefore the effect of non-uniform charge probabilities is clearly subdominant for DFSZ₅ to DFSZ₇. This most likely still is the case even for DFSZ₈ and DFSZ₉.

5.3.6 Extrapolation to More than Seven Higgs Doublets

While the anomaly ratios in principle can be calculated for any number of doublets, calculating all possible anomaly ratios for larger n_D involves solving an extremely large number of LSEs. In order to see how many, one can estimate the number of all

¹Just think of the extreme case of say, a charge distribution with 100 unique solutions, in which 10 solutions are $1000\times$ more probable than the other 90. The resulting E/N distribution will behave more as if it came only from 10 unique charges than as if it had 100.

Table 5.3: Important statistics of DFSZ-type models broken down by number of Higgs doublets n_D . Included is information on the model with maximal photon coupling $\widehat{E/N}$ from eq. (1.42) and the percentage of models that have minimal photon coupling (photophobic, $|E/N - 1.92| < 0.04$). ‘x’ denotes values that could not be estimated.

n_D	$\#V_{\text{eb}}$	Unique solutions	Unique E/N s	$\widehat{E/N}$	% Photophobic	% $N_{\text{DW}} = 1$
2	2	2	2	2/3	0	0
3	54	16	5	-4/3	0	0
4	52614	996	83	-52/3	1.4	6.00
5	6.65×10^7	9.7×10^4	432	-112/3	1.52	6.64
6 ¹	$\lesssim 4 \times 10^9$	$> 2.19 \times 10^6$	1680	-238/3	1.37	5.83
7 ¹	$\lesssim 7 \times 10^{12}$	x	6256	-466/3	1.39	5.19
8 ²	$\lesssim 2 \times 10^{16}$	x	> 11617	$< -628/3$	x	x
9 ²	$\lesssim 1 \times 10^{20}$	x	$\gg 14122$	$< -1216/3$	x	x

¹For $n_D \geq 6$, $\#V_{\text{eb}}$ and “unique solutions” are estimates. Number of minimal potentials calculated via eq. (5.14), many of which will be unphysical and not produce valid solutions for PQ charges. “unique solutions” gives the number of solution found in sample, for which data exists.

²For $n_D \geq 8$, not all possible models were calculated, therefore no exact values are available for the number of unique E/N , the percentage of photophobic models or models with $N_{\text{DW}} = 1$. $\widehat{E/N}$ was estimated as shown in subsec. 5.3.7.

possible terms for step 2 in the cooking recipe in subsec. 5.3.1 with an arbitrary n_D . Since the number of possible bilinears is $n_B = \binom{n_D}{2}$ plus their Hermitian conjugate, there are $2n_B$ terms of the form $HHSS$. Regarding the quadrilinears this results in $(2n_B)^2$ possible terms. Written as a matrix, this yields

$$\begin{array}{cc} & \begin{array}{cc} HH & (HH)^\dagger \end{array} \\ \begin{array}{c} HH \\ (HH)^\dagger \end{array} & \begin{pmatrix} A & B \\ C & D \end{pmatrix}, \end{array} \quad (5.13)$$

where A denotes the submatrix formed by all terms of the form $HHHH$, B by $HH(HH)^\dagger$, and so on. However, as in the DFSZ₃ example, there are several equal terms in this matrix that should not be counted. First of all, the whole matrix is symmetric. Secondly, since Hermitian conjugated terms are equal, D is completely redundant with respect to A . Lastly, B is anti-symmetric, so that the number reduces to n_B^2 quadrilinears.

From the set of all terms, one needs to pick $n_D - 1$ terms where at least one must be of the form $HHSS$. Hence one can pick between 1 and $n_D - 1$ terms of the form $HHSS$, then fill up with $HHHH$ terms, and repeat this for all possible amounts of $HHSS$ terms (ignoring equivalences in the case of multiple $HHSS$ terms). The total number of possible V_{eb} can then be estimated by

$$N_{\text{tot}}(n_D) \sim \sum_{j=0}^{n_D-2} \binom{2n_B}{1+j} \binom{n_B^2}{n_D-2-j}, \quad (5.14)$$

which at the same time is the number of LSEs that needs to be solved.

Simplifications used for the DFSZ₃ example apply here as well, namely setting all VEVs to one and not fixing χ_S , but regardless of these simplifications the computation time rises exponentially with n_D . While $N_{\text{tot}}(n_D = 3) = 69$ is easily manageable, for e.g. $n_D = 8$ the number of possibilities becomes $N_{\text{tot}}(n_D = 8) \approx 2 \cdot 10^{16}$. Thus, computing requirements for solving all LSEs beyond DFSZ₇ are prohibitive.

An easy solution to the computationally prohibitive number of LSEs would be to sample the (minimal) potentials. However, due to step 4 in the approach, this is not possible without introducing a bias: A multitude of minimal potentials can lead to the same solution, thus all of them belong to the same model. This model therefore has a very long potential and is likely to be found by any sampling algorithm. On the opposite side, there are also models containing just one or two of the minimal potentials. Sampling in the space of minimal potentials therefore leads to biased sampling in the space of models.

An alternative estimation for the DFSZ₈ and DFSZ₉ distributions can come from the following considerations. If a large enough number of theories is considered, fig. 5.2 (top, right) and fig. 5.2 (third row, right) can be viewed as extremal cases for the anomaly ratio distribution. “Extremal” in this context should not be understood in terms of an upper or lower limit on individual E/N bins; after all (normalised) probability measures are being considered. Rather, fig. 5.2 (top row, right) is very smooth, whereas fig. 5.2 (third row, right) is very peaked. To quantify this criterion one can look at the cumulative sum of anomaly ratios below a specific value. Similarly to the two sample Kolmogorov-Smirnov-test, “smoothness” of an anomaly ratio distribution $f(E/N)$ can be defined as

$$\max_x \left| \sum_{E/N < x} f(E/N) - \sum_{E/N < x} c(E/N) \right|, \quad (5.15)$$

where $c(E/N)$ represents the continuous distribution as shown in fig. 5.2 (top, right). Eq. (5.15) defines smoothness via the maximum of the difference for all anomaly ratios in the cumulative sum of the distribution compared to the case of continuous charges. In subsec. 5.4.1 the close connection of this metric to the relevant observable will be demonstrated. The metric runs from one to zero (by construction for the continuous distribution), with values closer to zero indicating a smoother distribution. For DFSZ₃, the value is 17%, for DFSZ₄ already 5.7%, and down to 1.4% for DFSZ₇.

The DFSZ₈ and DFSZ₉ anomaly ratio distributions can be roughly estimated by constraining their smoothness. For DFSZ₃ to DFSZ₇ a higher number of doublets leads to a smoother anomaly ratio distribution. From investigations of the biased sampling for $n_D = 6$ and $n_D = 7$, where the true distributions were available, it was observed that sampling leads to less smooth distributions. This means that the distribution for $n_D = 8$ or $n_D = 9$ is expected to be smoother than their respective sampled distribution and the $n_D = 7$ distribution. In fig. 5.3, the $n_D = 7$ distribution is used as one estimate, denoted as “limit 1”.

The other estimate, overestimating the smoothness, can come from the observation that the difference in smoothness of the distributions is smaller between DFSZ₆ and DFSZ₇ than between DFSZ₅ and DFSZ₆. Extrapolating the histograms beyond $n_D = 7$ using the difference of the distributions of DFSZ₆ and DFSZ₇ should therefore yield anomaly ratio distributions which are smoother than expected. In fig. 5.3, the

difference is subtracted once to reach the estimate for $n_D = 8$ and twice for $n_D = 9$, denoted as “limit 2”. For the metric (5.15) the values 0.73% and 0.71% are found for DFSZ₈ and DFSZ₉, respectively.

The above approach should not be viewed as representing hard limits for the anomaly ratio distributions for eight or nine Higgs doublets, but rather a rough estimate. The difference in probability in fig. 5.3 looks substantial only due to the logarithmic axis. Both estimates are much closer to the continuous case of fig. 5.2 (top, right) than to the peaked one of fig. 5.2 (third row, right) in the sense that only very little of their probability mass lies at unique E/N values and most in a continuum.

5.3.7 Constructing extreme $|g_{a\gamma}|$

Another problem with sampling potentials as described in the previous paragraphs is that it is very unlikely to find the anomaly ratio corresponding to the maximum axion-photon coupling $\widehat{E/N} = \operatorname{argmax}_{E/N} (|E/N - 1.92|)$ (see also eq. (1.42)).

This anomaly ratio, however, is very useful for constraining the parameter region of DFSZ-type models. For this reason a procedure is laid out here to construct an estimate for it. Note that due to the symmetry around $E/N = 5/3$, in absence of selection criteria, $\widehat{E/N}$ is not given by the largest possible anomaly ratio but the smallest.

The procedure is based on observations of the LSEs that led to $\widehat{E/N}$ for the smaller numbers of doublets. Any of the LSEs leading to $\widehat{E/N}$ of DFSZ₄ can be extended to an LSE leading to $\widehat{E/N}$ for DFSZ₅. The same behaviour can be seen from DFSZ₅ to DFSZ₆ and in a slightly different form from DFSZ₃ to DFSZ₄. No rigorous mathematical reason has been found for this observation, so applying it to larger n_D is more of an educated guess. However, it turns out to give extreme anomaly ratios, so it is used to systematically estimate $\widehat{E/N}$.

The procedure goes as follows. First, take all LSEs that lead to $\widehat{E/N}$ for a number of doublets where all solutions are known, say $n_D = 6$. Secondly, add one additional Higgs doublets by specifying the Yukawa sector for the new doublet. Thirdly, adjust the orthogonality relation appearing in all LSEs depending on what type of doublet is added. Then, add one additional relation to the LSEs, solve them and calculate the anomaly ratio. After that repeat this for every possible relation and every possible Yukawa sector. Finally, extract the LSEs with the smallest anomaly ratio.

This results in highly negative anomaly ratios. However, for DFSZ₉ by taking the resulting LSEs and systematically exchanging one (or more if the runtime is acceptable) of the relations, new LSEs can be found that give even smaller anomaly ratios. In DFSZ₉ for instance, the smallest anomaly ratio constructed in this way is $E/N = -1216/3$, and it is generated by the terms

$$\begin{aligned} & (H_{d_2}^\dagger H_{e_1})(H_{d_2}^\dagger H_{d_1}), (H_{u_1} H_{d_1})(H_{u_1} H_{d_2}), \\ & (H_{u_3}^\dagger H_{u_1})(H_{u_3}^\dagger H_{u_2}), (H_{e_1}^\dagger H_{d_1})(H_{e_1}^\dagger H_{e_2}), \\ & (H_{e_2}^\dagger H_{d_1})(H_{e_2}^\dagger H_{e_3}), (H_{u_2} H_{d_3})(H_{u_1}^\dagger H_{u_2}), \\ & (H_{d_3} H_{u_1})(H_{d_1}^\dagger H_{d_3}), (H_{d_1} H_{u_1})S^\dagger S^\dagger. \end{aligned}$$

5.3.8 Comparison with KSVZ-Type Models

Having constructed and reviewed an extensive anomaly ratio catalogue of DFSZ-type models, it is now possible to draw a comparison to the analogue catalogue of KSVZ-type models constructed by the authors of [9] (see also subsec. 1.3.3). Even though the theoretical framework for KSVZ-type models differs substantially from DFSZ-type models, the calculation of the anomaly ratio is quite comparable. Thus the catalogues are naively expected to display a certain amount of similarity, the investigation of which is the focus of this section.

In their paper on the KSVZ-type model catalogue, the authors add all anomaly ratios of phenomenologically allowed KSVZ-type models, irrespective of the number of quarks, allowing to add or subtract quark representations. This means that a single model with $N_Q = 9$ quarks, of which there are $> 1 \times 10^5$, is deemed equally probable as a single model with $N_Q = 1$, of which there are only 15. The distribution is therefore dominated by $7 \lesssim N_Q \lesssim 21$. If a similar approach was used for DFSZ-type models, extrapolating the evolution of unique solutions with increasing n_D , the resulting distribution would be indistinguishable from the DFSZ₉ case. In subsec. 5.3.3 it was argued to instead use an approach in which all separate values for n_D are equally probable. For consistency one must demand the same for N_Q in the KSVZ-type case. Since raw data was provided by [9], it was possible to weight their KSVZ-type data in a way that gives equal probability to all values of N_Q .²

Using this weighting, their data can be compared with the DFSZ-type results on a fair basis and the result is shown in fig. 5.4. Nevertheless differences remain: The authors of [9] were able to apply strict selection criteria, significantly reducing the number of viable models. Here, no similarly stringent selection criteria were found, so the DFSZ-type catalogue reflects the full set of models rather than a preferred set. A comparison between the two types of models or a combined axion band should therefore not be seen as final, but only as incorporating all selection criteria known so far. Also note that in the DFSZ-type case, a model with higher n_D is always less likely than a model with lower n_D , which can be seen as an appropriate penalty for introducing more degrees of freedom to the model. In the weighting scheme used here for KSVZ-type data from [9] this is not the case, since for e.g. $N_Q = 28$ they find only 510 preferred models, much less than for $N_Q = 9$, making a single model with $N_Q = 28$ more likely than a single model with $N_Q = 9$ when equal weighting is applied for all N_Q .

One can clearly see the effect of the equal weights for all n_D in fig. 5.4 in the region around $E/N = 5/3$: The five E/N values of the DFSZ₂ and DFSZ₃ models show highly elevated probability due to their large relative probabilities (compare fig. 5.3). The effect of the two estimates for DFSZ₈ and DFSZ₉ only becomes substantial at low absolute probabilities and above $|E/N| \gtrsim 20$. The KSVZ-type results are also found to form a peaked structure similar to the DFSZ-type case, which only becomes visible in a very finely binned histogram. In fact, for E/N values excluding DFSZ₂ and DFSZ₃, the DFSZ-type histograms are less peaked than the KSVZ-type ones, with decreased probability at moderately large $|E/N|$ and significantly increased probability for $|E/N| \gtrsim 40$. This trend does not translate to the largest possible axion-

²So now the 15 models with $N_Q = 1$ combined are equally likely as all $> 1 \times 10^5$ models with $N_Q = 9$ combined.

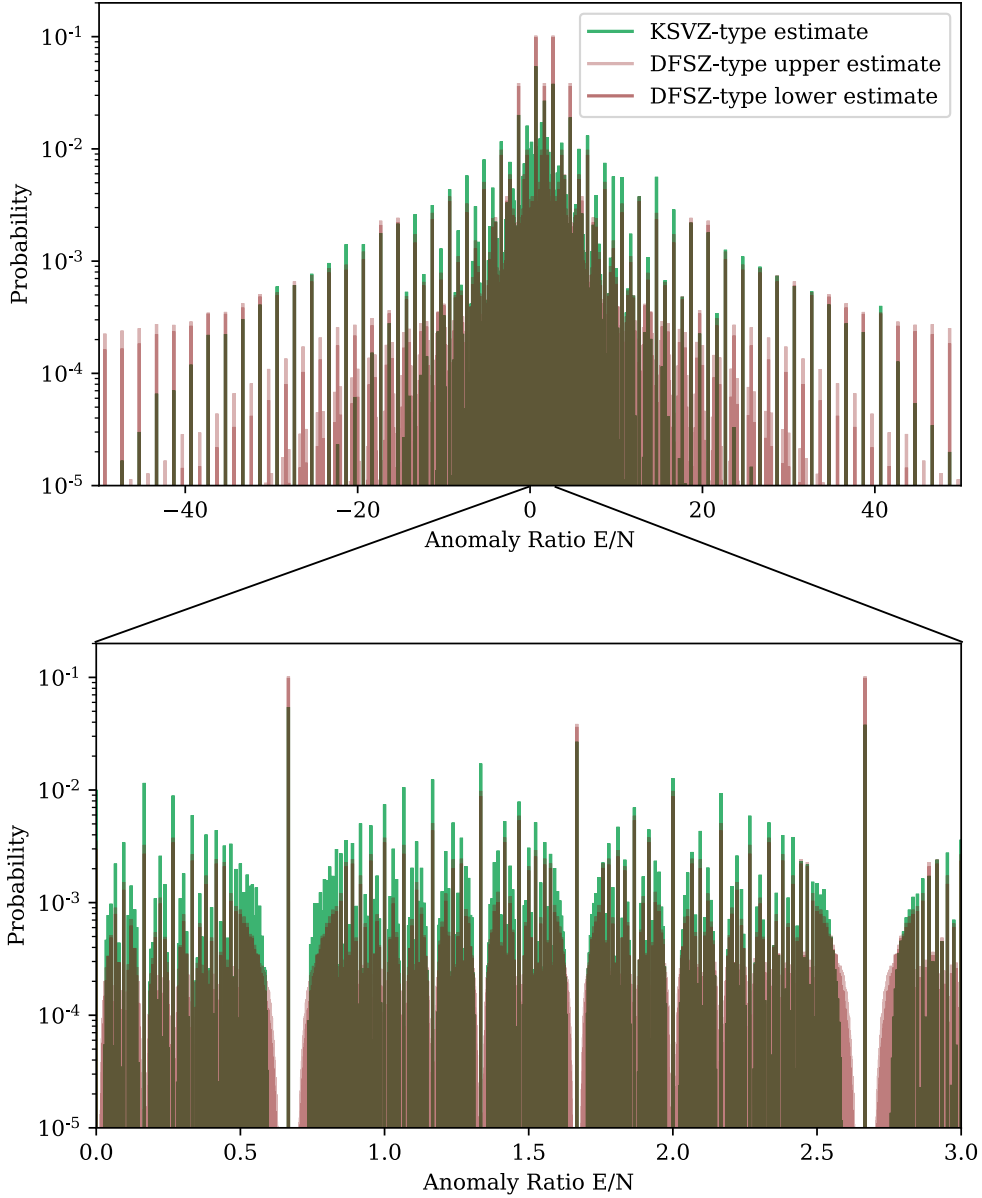


Figure 5.4: Comparison between anomaly ratio distributions for KSVZ-type and DFSZ-type axion models. The KSVZ-type estimate of [9] includes all phenomenologically allowed models, adding and subtracting quark representations and assumes every model to be equally likely. The DFSZ-type results presented here include calculations for DFSZ_2 to DFSZ_7 and estimates for DFSZ_8 and DFSZ_9 giving equal probability to each n_D . For the DFSZ-type distribution different shades denote maximum and minimum for each bin under the two limits for DFSZ_8 and DFSZ_9 described above.

photon coupling, however. A maximal $|g_{a\gamma}|$ is found at $\widehat{E/N} = -1216/3$, which is comparable to the KSVZ case for $N_Q \leq 9$ before any phenomenological constraints ($\widehat{E/N} = -1312/3$).

Models with smallest axion to photon couplings, in the following photophobic models, are defined the same way as in [9] by their anomaly ratio E/N being compatible with vanishing $g_{a\gamma}$ within 1σ theoretical uncertainty (see eq. (1.19)). Tab. 5.3 shows that there is no clear trend toward a higher or lower percentage of photophobic models with increasing n_D . As discussed in subsec. 5.3.4 the anomaly ratio distribution becomes smoother with increasing n_D : Peaks become less pronounced and anomaly ratios with low probability become more likely. The absence of a clear trend hints at the photophobic region being right in the middle between those two extremes. Overall the percentage of photophobic models found for DFSZ-type models with $n_D \leq 7$ is similar to the KSVZ-type case.

In both, KSVZ- and DFSZ-type of model probability distributions, the probability close to the highest peaks is strongly suppressed (fig. 5.4, bottom). The effect is less severe for DFSZ-type models than for KSVZ-type ones, because as noted before the former are less peaked if the effect of DFSZ₂ and DFSZ₃ is removed.

Upon closer inspection the distribution of KSVZ-type models is not symmetric around $5/3$, unlike the DFSZ-type one. Median and mean anomaly ratios are $E/N|_{\text{mean}} = 1.43$ and $E/N|_{\text{median}} = 1.30$, respectively, whereas for DFSZ-type models both are exactly $E/N|_{\text{mean}} = E/N|_{\text{median}} = 5/3$. These values remain unchanged, even if only considering the subset of $N_{\text{DW}} = 1$ models. The deviation from $5/3$ in the KSVZ-type models of [9] may arise due to the phenomenological selection criteria they impose.

5.4 Implications for Axion Searches

5.4.1 $\mathcal{C}_{a\gamma}$ Cumulative Distribution Function

In the previous sections probability mass functions have been derived for the anomaly ratio from theoretical assumptions for different DFSZ-type theories. To be able to understand the implications for axion searches, these E/N distributions must be mapped into $g_{a\gamma}$ space via $g_{a\gamma} = \frac{\alpha_{\text{em}}}{2\pi f_a} \mathcal{C}_{a\gamma}$ with $\mathcal{C}_{a\gamma}$ as defined in (1.20). To recapitulate,

$$\mathcal{C}_{a\gamma} = \frac{E}{N} - 1.92(4) \equiv \frac{E}{N} - \mathcal{C}_{a\gamma\gamma}^{(0)}, \quad (5.16)$$

where the correction term independent of E/N has been dubbed $\mathcal{C}_{a\gamma\gamma}^{(0)}$. In order to be independent of the axion mass, results are plotted with respect to the unitless quantity $|\mathcal{C}_{a\gamma}|$.

Traditionally, two-sided axion bands centred around the region of maximal probability are given in this case [8, 9, 135–137]. However, usually an experiment is sensitive to all axion-photon couplings above a certain threshold $|\mathcal{C}_{a\gamma}|^{\text{min}}$. It could therefore also be deemed relevant to experimental efforts to post a one-sided limit that has to be reached in order to be sensitive to, e.g. 68% of all DFSZ-type models given a specific axion mass. For this purpose a cumulative distribution function (CDF) is used, plotted against $|\mathcal{C}_{a\gamma}|$, which can be understood as the combined theoretical prior probability of models with $|\mathcal{C}_{a\gamma}|(\text{model}) > |\mathcal{C}_{a\gamma}|^{\text{min}}$.




Figure 5.5: Cumulative probability of models with $|\mathcal{C}_{a\gamma}|$ larger than the indicated values. The plot includes DFSZ-type models of arbitrary domain wall number N_{DW} with DFSZ₃ to DFSZ₇ as well as $N_{\text{DW}} = 1$ models for DFSZ₄ to DFSZ₇ (for smaller n_D no $N_{\text{DW}} = 1$ models exist). The CDFs become smoother with increasing n_D , with DFSZ₆ and DFSZ₇ already being almost indistinguishable. $N_{\text{DW}} = 1$ models have systematically larger $|\mathcal{C}_{a\gamma}|$, shifted by almost one order of magnitude. The dashed line indicates the analytic fit on the continuum limit from eq. (5.17).

Since the anomaly ratio is treated as a random variable coming from a distribution that is to be determined, one has to treat $\mathcal{C}_{a\gamma}^{(0)}$ in the same way. Its uncertainty can be modelled as a normal distribution with mean 1.92 and standard deviation 0.04. In effect, this smooths out steps in the CDF from high probability E/N values, especially for anomaly ratios close to the mean value of $\mathcal{C}_{a\gamma}^{(0)}$.

5.4.2 Experimental Constraints

Under the assumptions outlined above the anomaly ratios of the DFSZ₂ and DFSZ₃ models still are the most notable features in the probability distribution, even for all possible DFSZ-type models. However, since only one value of the anomaly ratio is realised in nature, reaching sensitivity to these models may be either not necessary or not sufficient.

Fig. 5.5 shows the resulting theoretical prior probability of DFSZ-type axion models with $|\mathcal{C}_{a\gamma}|$ larger than a specific value. The results are broken down by possible values of n_D . Considering first the “all N_{DW} ”-case, in which the domain wall number does not present a meaningful selection criterion, DFSZ₃ models have zero probability above $\log|\mathcal{C}_{a\gamma}| \gtrsim 0.5$. Should an axion be found above this value and be determined to be of DFSZ-type, this would imply the existence of $n_D > 3$ Higgs doublets. The impact of the prominent peaks of maximal probability between $E/N = -4/3$ or $E/N = 14/3$ on the cumulative probability is only minor for theories with $n_D \geq 5$.




Figure 5.6: Cumulative probability of models with $|\mathcal{C}_{a\gamma}|$ larger than the indicated values for the complete set of DFSZ-type and KSVZ-type models as well as for models with $N_{\text{DW}} = 1$ specifically (thin lines). One sided 95% and 68% limits for both cases are given with coloured vertical dotted lines. The arrows at the top indicate the location of DFSZ_{2-I} and DFSZ_{2-II}.

Since the CDFs for DFSZ₆ and DFSZ₇ are already almost indistinguishable, estimates for higher n_D are not included in the plot. In fact, the relative difference on $|\mathcal{C}_{a\gamma}|$ exclusion limits between the two ways of estimating the smoothness of the DFSZ₈ and DFSZ₉ distributions is below the percent level. For the purpose of $|\mathcal{C}_{a\gamma}|$ exclusion limits the two estimates are therefore virtually equivalent. In the following limit 2, the extrapolation estimate, is used.

It is possible to obtain a reasonable analytic estimate for the cumulative probability distribution by going back to the analytic anomaly ratio fit from eq. (5.12). For $|\mathcal{C}_{a\gamma}|$, it translates to

$$p(|\mathcal{C}_{a\gamma}|) = 1 - \frac{\tan^{-1} \left[\frac{4}{7} (|\mathcal{C}_{a\gamma}| - \frac{19}{75}) \right] + \tan^{-1} \left[\frac{4}{7} (|\mathcal{C}_{a\gamma}| + \frac{19}{75}) \right]}{\pi}, \quad (5.17)$$

which is plotted as a dotted line in fig. 5.5.

Now contrast the full set of DFSZ₄ to DFSZ₇ models with the respective subsets with unity domain wall number. The latter models could be considered preferred in the post-inflationary scenario due to cosmological energy density arguments (see subsec. 2.2.3). $N_{\text{DW}} = 1$ models display $|\mathcal{C}_{a\gamma}|$ values almost an order of magnitude higher on average than the full set and are therefore much easier to detect. Similarly to the “all N_{DW} ”-case, higher n_D values tend to lead to smoother distributions. It therefore seems reasonable to introduce estimates for $n_D \geq 8$ similarly as before, again leading to a difference with respect to the $|\mathcal{C}_{a\gamma}|$ limits between the two estimates below the percent level. Again limit 2, the extrapolation estimate, is used.

In fig. 5.6 a comparison of the CDFs for DFSZ- and KSVZ-type models is shown. In general, both types display very similar cumulative probabilities for all $|\mathcal{C}_{a\gamma}|$. Only




Figure 5.7: Probability density in $\log|C_{a\gamma}|$ -space of the combined DFSZ-type and KSVZ-type “all N_{DW} ”-case. Central 68% and 95% regions used for fig. 5.8 (bottom) are indicated in different shades of yellow. Note that the underlying distribution is discrete and any illustration will in part depend on the chosen binning.

for DFSZ-type models with $N_{\text{DW}} = 1$ a significant percentage of models is above $\log|C_{a\gamma}| \gtrsim 1.5$. The lines of $E/N = 2/3$ and $E/N = 8/3$ are clearly visible for DFSZ-type models, but also for KSVZ-type models. The relative difference between the 68% limits of KSVZ- and DFSZ-type axions is only $\sim 3\%$ and $\sim 19\%$ for the 95% limits with the DFSZ-type limit being higher in the latter case. Taking into account possible effects from diverging theory assumptions, this relative difference can be seen as negligible.

While the investigation of different theoretical assumptions is beyond the scope of this thesis, note that other assumptions on the full set of models only modify the relative importance of the prominent DFSZ₂ and DFSZ₃ peaks. Consider for example a different definition of multiplicity based on minimal potentials. This dramatically increases their probability mass but does not shift the overall cumulative probability to higher or lower $|C_{a\gamma}|$ values. In this sense any variation of theoretical assumptions (excluding model selection criteria) should lie between the cumulative probabilities of the $n_D = 2$ and continuum cases.

Translating the aforementioned limits to $g_{a\gamma}$ and plotting them over a range of axion masses yields fig. 5.8 (top). An experimental exclusion limit touching the 68% line excludes 68% of the probability mass over the model space under the assumptions outlined above given a specific mass range. An experiment targeting sensitivity down to the 95% line will be sensitive to 95% of the probability for all models in the targeted mass range. These and the 99% limit for DFSZ-type as well as KSVZ-type models are included as well as the combined case of $N_{\text{DW}} = 1$. For the full cumulative probability from which these three limits are taken, see fig. 5.6. Note that all experimental limits shown in fig. 5.8 are frequentist in nature and should therefore only be seen as a rough

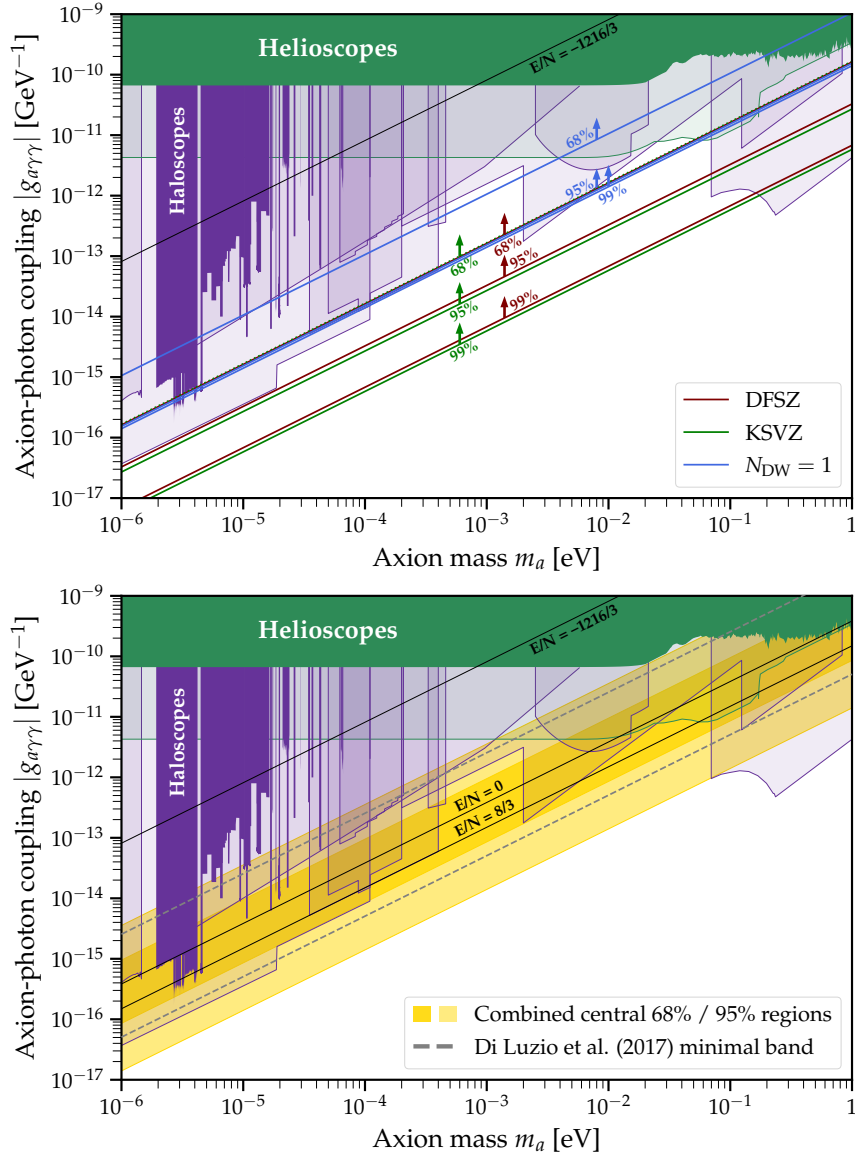


Figure 5.8: **Top:** 68%, 95% and 99% limits for the complete preferred KSVZ-type case [9], the complete DFSZ-type case (using extrapolation for DFSZ₈ and DFSZ₉) as well as the combined $N_{\text{DW}} = 1$ case. The largest DFSZ-type coupling found is shown in black ($E/N = -1216/3$). Helioscope limits and forecasts [89, 90, 138] are shown in green as well as limits and forecasts from various haloscope experiments [81, 92, 93, 98, 109, 139–177] in purple. **Bottom:** Central 68% and 95% regions for the case combining all preferred KSVZ-type and all DFSZ-type models together with a previous band from [8] for comparison (dashed). For reference lines corresponding to $E/N = 0$ and $E/N = 8/3$ are shown in black. (Plotted with tools by O’Hare [178].)

Table 5.4: $|\mathcal{C}_{a\gamma}|$ lower prior limits for selected combinations of models. All limits shown are one sided, so a central 68% band can be constructed with values given for 16% and 84% and similar for 95%. **KSVZ-type** denotes re-weighted results from [9], **DFSZ-type** results presented here. Both are combined with equal probability for the case **Combined**. The combination only considering models with DW number of unity is shown as $\mathbf{N}_{\text{DW}} = 1$.

$ \mathcal{C}_{a\gamma} $			68% band		95% band	
One-sided limit	68%	95%	16%	84%	2.5%	97.5%
KSVZ-type	0.833	0.135	4.684	0.427	15.274	0.068
DFSZ-type	0.809	0.164	4.529	0.482	19.272	0.08
Combined	0.819	0.148	4.593	0.451	17.285	0.074
$N_{\text{DW}} = 1$	5.294	0.769	22.773	2.733	36.729	0.731

comparison with respect to the theoretical prior probabilities presented here which are more in line with a Bayesian interpretation of statistics. In black the maximal $\widehat{E/N}$ value found for DFSZ₉ is included. In addition to being excluded by experiments for a large fraction of the m_a range, this model may likely be subject to phenomenological constraints (see sec. 5.2.3).

The results presented in this chapter make it possible for the first time to give values of one-sided limits or axion bands for the combined KSVZ-type and DFSZ-type case, assuming a DFSZ-type axion to be equally likely as a KSVZ-type one. The associated PDF for the combined “all N_{DW} ”-case is shown in fig. 5.7. In $\log|\mathcal{C}_{a\gamma}|$ -space with the chosen binning, the distributions look roughly Gaussian with the exception of several notable peaks, at $E/N = 5/3, 8/3, 2/3, 14/3$ and $-4/3$ (from left to right). Note, however, that the true underlying distribution is comprised out of a multitude of delta peaks and thus fundamentally discrete. Central 68% and 95% bands from this distribution are used in fig. 5.8 (bottom) together with a previous estimate for the same band from [8]. Previous work was either limited to very few extensions of DFSZ-type [8] or to the KSVZ-type case [9]. Even now many caveats have to be kept in mind, like the imprecise prediction for DFSZ₈ and DFSZ₉ models or the lack of selection criteria in the DFSZ-type case. Nevertheless, for the reader access to data of typical limits and bands for a variety of scenarios may be useful. An overview can be found in tab. 5.4 and more detailed information, such as the full model catalogues as well as axion limits/ bands together with usable Bayesian theory priors, is hosted on the website “zenodo” <https://doi.org/10.5281/zenodo.7656939>. The generating code can be found at <https://github.com/jhbdiel/DFSZforest>.

5.5 Summary & Discussion

In this chapter axion-photon couplings for a large number of DFSZ-type models were systematically calculated. This resulted in limits that have to be reached in order to be sensitive to a certain fraction of the probability mass of these models.

First, (phenomenological) selection criteria were discussed, such as the absence of FCNCs and the DW problem, to extract preferred DFSZ-type models. In contrast

to KSVZ-type axions, where all selection criteria follow from cosmological bounds on additional fermions, for DFSZ-type axions no criteria display a sufficient level of generality, merely desirable features were found.

Next, a recipe was put forth for calculating all anomaly ratios and hence all axion-photon couplings. This recipe is based on the fact that the PQ charges are not free but fixed by linear consistency and phenomenology relations. For the sake of calculating the anomaly ratio, this reduces the procedure of DFSZ-type model building to solving LSEs. Thus, systematically going through all Yukawa sectors and solving all possible LSEs for each, all possible anomaly ratios could be derived for up to seven Higgs doublets.

In addition, by counting how many models lead to a certain anomaly ratio and establishing relative probabilities of these models, probabilities could be assigned to each anomaly ratio. For this counting of models, a model was considered to correspond to the Lagrangian that arises by combining different potentials that give rise to the same set of PQ charges and by adding the Yukawa couplings compatible with the resulting set of PQ charges. This takes into account the general mantra that all terms allowed by symmetry should be included in the Lagrangian and overcounting is avoided.

The resulting anomaly ratio distributions have their medians at $E/N = 5/3$, their maximum values at $E/N = 2/3$ and $E/N = 8/3$, and a characteristic shape that is similar to the one of KSVZ-type models. These observations can be explained by abstractly modelling the resulting sets of PQ charges as discrete charge distributions with uniform probability and symmetry around zero.

While the aforementioned recipe works in principle for an arbitrary number of Higgs doublets, the necessary computational time becomes too large for eight or more doublets. Simple sampling of potential terms leads to a significant bias, so that estimates for $n_D > 7$ had to be constructed based on the expected smoothness of the distributions. Moreover, by using an incremental construction procedure, it was possible to find a maximal anomaly ratio that is more than a factor of two higher than in previous estimates [8].

Regarding the axion experimental program, the anomaly ratio distributions confirm the experimental importance of the values dictated by the minimal DFSZ models, namely $E/N = 2/3$ and $E/N = 8/3$, since they are also favoured for every number of Higgs doublets (except $n_D = 3$ with the Weinberg-Glashow-Paschos condition imposed). However, they also show that plenty of viable parameter space lies above and below these lines. Overall, this means that a non-observation at these favoured values is not enough to declare the axion excluded, while an observation above these values would be a hint for more than one additional Higgs doublet from the DFSZ-type point of view. The statistical interpretation also reveals that both, KSVZ- and DFSZ-type models, set very similar sensitivity requirements on experiments.

For $n_D \geq 4$ a subset of models with DW number $N_{\text{DW}} = 1$ exists, making DFSZ-type models theoretically more viable in post-inflationary scenarios. This subset even displays a significantly enhanced axion-photon coupling compared to the minimal scenarios for both invisible axion classes, hence making these models on average easier to probe.

The analysis presented here can be extended in multiple directions. For instance, it would be interesting to perform a similar analysis for models with a right-handed neut-

rino or KSVZ/DFSZ hybrid models, namely models with additional Higgs singlets, Higgs doublets, and heavy quarks. From purely mathematical arguments regarding the anomaly ratio, even though eq. (5.1) would change, similar shapes of the resulting distributions and axion mass vs. axion-photon coupling exclusion lines are expected. This expectation does, however, not make an explicit analysis dispensable. Furthermore, it would be interesting to investigate other axion couplings, such as the axion-electron coupling. Additionally, it would be desirable to find a better estimate for the anomaly distribution of eight or more doublets, or even an unbiased way to calculate it.

This project provides a first well-motivated theory prior combining both major categories of invisible axions beyond the minimal models. While so far most axion experiments have subscribed to a frequentist interpretation of statistics, the catalogue presented here is a valuable resource for analysis efforts moving towards a Bayesian framework like [79, 120]. This prior is also used for the Bayesian analysis presented in ch. 6, where a novel modification to the standard background removal technique is validated using generated mock-data resembling a full-scale MADMAX setup. The MADMAX data-taking of 2024 which is analysed in ch. 8 by design does not reach sensitivity of photon couplings close to typical values of this prior. Therefore applying the prior in that context is not beneficial.

Lastly, this analysis is useful for axion searches irrespective of the statistical interpretation. With all possible E/N values for up to seven doublets and a full catalogue for up to five doublets, in the case of an axion detection one can proceed to do hypothesis testing with the compatible models. Since all E/N values for preferred KSVZ-type models are also known, this could be used for the purpose of model comparison between these two model classes. Hence, with or without a statistical perspective, this project presents another step forward in the understanding and mapping of the landscape of axion models.

6

Bias-Free Signal Estimation

Axion haloscope experiments face a common but challenging problem: A small-amplitude signal S of unknown location is to be found in a dominating background B . Because the relevant wavelengths are comparable to the size of the system, modelling the background from first principles poses a virtually insurmountable challenge.

To be able to employ a forward-modelling approach, i.e. to construct a model that is able to simultaneously consider signal and background components, parametric models for signal as well as background need to be available.

It could still be possible to construct such a parametric background model using an empirical function without relying on first principles. However, such a function would have to be able to accommodate a range of possible background shapes to a high degree of precision as determined by the expected signal amplitude while not introducing a vast number of nuisance parameters.

In practice it is more feasible to approximate the background using a parameter-free background filter f_{bg} and subtract this estimate from the observation. The residuals then ideally only retain some statistical noise N as well as the signal component, which can be found using standard statistical methods.

However, such a background estimator is usually affected by the presence of a signal, which in turn means that the signal shape in the residuals is altered. Signal parameters inferred from such residuals are therefore biased unless measures are taken to correct for this effect.

In an axion context, simulated experiments have been used to correct for the bias in signal amplitude, but not for the other parameters related to signal shape and position [83, 114]. This chapter improves on these previous efforts by presenting a method that allows the use of parameter-free background filters, while still being able to obtain inherently unbiased signal estimates.

Such a method is vitally important when a signal is found in order to not infer biased physical parameters. However, the approach could also provide benefits in

the absence of measurable signals: If strong prior knowledge exists on one or more of the parameters in a Bayesian context, a bias in the signal model could lead to decreased sensitivity due to inferred, biased signal parameters corresponding to regions in parameter space with lower prior probability.

First, sec. 6.1 presents the general principle of the approach in an abstract manner, before it is applied in sec. 6.2 to a concrete physics use-case in the context of the MADMAX haloscope using Savitzky-Golay filters as background estimators. Sec. 6.3 concludes the chapter. The results presented here have already been published by the author in [2], which this chapter follows closely.

6.1 General Approach

Formally, the approach for bias-free signal estimation can be applied whenever the following ingredients are present:

1. An experimental observation X , which can be written as a linear combination of three distinct components signal S , background B and noise fluctuations N (compare fig. 6.1):

$$X = S + B + N . \quad (6.1)$$

S and B are formally only distinguished by the fact that the former should be retained and the latter removed. N however has the special role of accounting for the full statistical uncertainty. In particular, for the probability distribution of X , given a known signal and background, it is required that

$$p(X|S, B) = p_{S+B}(N) , \quad (6.2)$$

with known $p_{S+B}(N)$, the probability distribution of the noise which is parameterised by the sum of S and B . A simple example for this distribution would be a Normal distribution with mean 0 and standard deviation depending on $S + B$. Also, a vanishing expectation value $E(N) = 0$ is required, so that

$$E(X) = S + B . \quad (6.3)$$

2. An unbiased and effective background estimator f_{bg} . Many mathematical constructs could be used as f_{bg} , for example a moving average or any analytic function fulfilling all requirements. Ideally, it is constructed to yield $f_{\text{bg}}(X) = B$ or equivalently $X - f_{\text{bg}}(X) = S + N$, however this is not necessary. It only needs to be unbiased in the sense that it reconstructs the background faithfully and to be effective in suppressing the noise:

$$B - f_{\text{bg}}(B) \approx 0 , \quad (6.4)$$

$$f_{\text{bg}}(N) \approx 0 . \quad (6.5)$$

The approximations above need to hold relative to possible signal amplitudes and length scales. In particular, for the application below it must be impossible for $B - f_{\text{bg}}(B)$ to mimic a signal

$$|B - f_{\text{bg}}(B)| \ll |\text{ampl}(S)| \quad (6.6)$$

and $X - f_{\text{bg}}(X)$ must have approximately the same noise as X

$$|f_{\text{bg}}(N)| \ll \sqrt{\text{var}(N)}. \quad (6.7)$$

3. The background estimator must also be linear, in particular

$$f_{\text{bg}}(S + B + N) = f_{\text{bg}}(S) + f_{\text{bg}}(B + N). \quad (6.8)$$

4. Furthermore, it must not reconstruct non-vanishing signals perfectly

$$S \neq 0 \implies f_{\text{bg}}(S) \neq S. \quad (6.9)$$

If this were not the case, the background removal could potentially completely eliminate a signal. Crucially, this requirement is much less strict than the naive assumption of the background estimator leaving the signal approximately invariant, i.e. $S - f_{\text{bg}}(S) \approx S$. The method works even for substantial changes, even though practically this could lead to a loss in sensitivity.

5. A signal that can be parameterised in tractable form using a set of parameters ϕ . This is made explicit in the following by writing the signal as S_ϕ . Less formally: Possible signal shapes need to be known.

Having a tractable parameterisation of the signal is a prerequisite for attempting parameter inference in virtually all cases. In a typical inference scenario where signal and background are separable, the first condition is also usually satisfied. The central requirements are therefore conditions two to four, which limit the choice of potential filters in a given scenario.

Note, that the domains of S and B can vary from the domain of X . For instance, if p_{S+B} follows a Poisson distribution, then S and B have \mathbb{R}^n as their domain, with $n \in \mathbb{N}$, while the observation X lies in \mathbb{N}_0^n . However, additions or subtractions of S , B or X as well as the background estimator need to be well-defined on the respective domains. In practice, this criterion is usually met.

Under the above conditions it is possible to construct a forward model of the experiment without a parameterised background model. One can make the background estimator f_{bg} a part of the experiment, replacing the original observation X with a virtual observation

$$X' \equiv X - f_{\text{bg}}(X), \quad (6.10)$$

estimating the background

$$B_{\text{est}} \equiv f_{\text{bg}}(X), \quad (6.11)$$

and defining a modified signal model

$$S'_\phi \equiv S_\phi - f_{\text{bg}}(S_\phi). \quad (6.12)$$

Using the requirements for the filter eqs. (6.4), (6.5) and (6.8), one can rewrite eq. 6.1 as

$$\begin{aligned} N &= X - S - B \\ &= X' + f_{\text{bg}}(X) - S - B \\ &= X' + f_{\text{bg}}(B + N) + f_{\text{bg}}(S) - S - B \\ &\approx X' + B + f_{\text{bg}}(S) - S - B \\ &= X' - S'_\phi. \end{aligned} \quad (6.13)$$

Since $S_\phi + B \approx S'_\phi + B_{\text{est}}$, this successfully approximately transformed $p(X|S, B)$ from eq. (6.2) into an unbiased statistical model, which is parameterised by signal parameters, but independent of the unknown background: $p(X|S, B) \approx p_{S'_\phi + B_{\text{est}}}(X' - S'_\phi)$. Under the assumptions above and given an actual observation X one can use this result to approximate the likelihood function of the signal parameters

$$\mathcal{L}_X(\phi) \approx p_{S'_\phi + B_{\text{est}}}(X' - S'_\phi) \quad (6.14)$$

and infer signal parameters ϕ from an observation X using standard statistical tools.

Note that for any application it is imperative to thoroughly verify if conditions two and four hold for the selected background estimator f_{bg} . For condition two, eqs. (6.6) and (6.7) or similar estimations may be used to check that the filter is unbiased and effective. Condition four could pose problems in connection with eq. (6.6). It may be impossible to approximately satisfy both requirements if signal and background fluctuate on similar scales, i.e. occupy similar regions in Fourier space. Even without an explicit parameterisation of the background, its general properties like amplitude and scale of fluctuations usually are available. The difficulty of a given application primarily depends on this similarity between background and signal as well as the signal-to-noise ratio.

For arbitrarily small signals it is difficult to a priori rule out $B - f_{\text{bg}}(B)$ matching allowed signal shapes in some region of the dataset, potentially leading to erroneously detecting a signal. This is one of the reasons why detectable signal strengths are assumed in the application below. Goodness-of-fit checks however should allow to reveal such false detections.

The application in the following section demonstrates the capability of the approach using Bayesian parameter inference. The method however is independent of a specific statistical framework and generally valid under the conditions described above.

6.2 Application to Simulated MADMAX Data

6.2.1 Data Generation and Analysis

A MADMAX setting is used to test the bias-free signal estimation approach on a real-world example. One thousand MADMAX-like mock-datasets are generated, a few of which are shown as grey lines in fig. 6.1. These datasets are then analysed with bias-free signal models that take the effect of the background subtraction into account and with standard signal models that do not. For the analysis a Bayesian approach is chosen, however the bias-free signal estimation procedure can be applied to other inference methods, e.g., a maximum likelihood estimate of the signal parameters.

One dataset consist of 25001 data points with 2 kHz spacing at frequencies slightly above $f_{\text{ref}} = 11$ GHz. It has three components (see fig. 6.1):

- **Signal.** An axion signal with shape determined by eq. (2.11). Fixed ρ_a and v_{lab} are assumed and m_a , σ_v and $\mathcal{C}_{a\gamma}$ are drawn from the prior. For $\mathcal{C}_{a\gamma}$, a priori models with $\mathcal{C}_{a\gamma} < 0.86$ are excluded to ensure that the signals have photon couplings large enough to be detectable at the noise level given below. The $\mathcal{C}_{a\gamma}$ cutoff was chosen such that at least 99% of simulated signals could be reliably recovered by the analysis. This cutoff corresponds to an SNR of 2.38

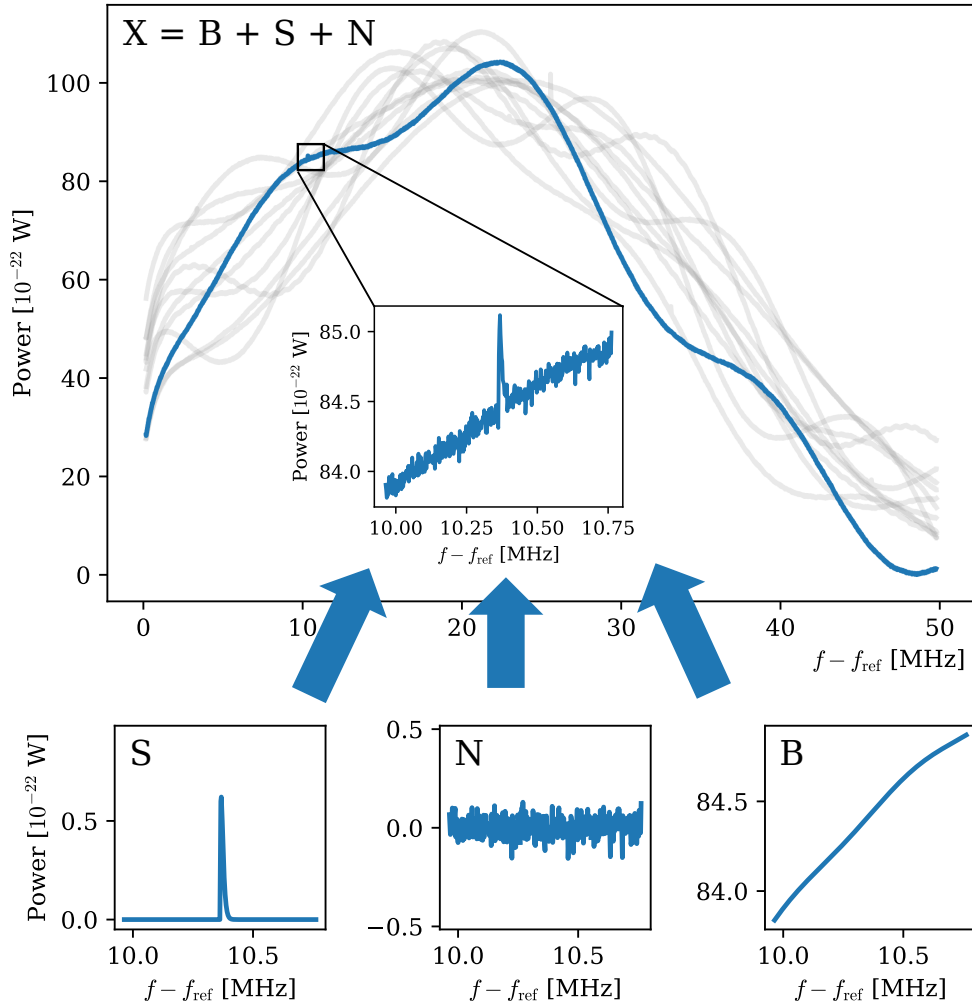


Figure 6.1: Components of a simulated example dataset. The upper panel shows multiple different full datasets X consisting of three components background B , simulated axion signal S and uncorrelated noise N as defined above, in order to visualise the variation in the simulated backgrounds. A zoom into the region of the axion signal for one of them is given. The three lower panels contain the three individual components of a dataset.

for the signal amplitude with respect to the standard deviation of the noise. So the scenario is considered where an axion signal has been discovered but its quantitative properties have not been inferred yet.

- **Noise.** Uncorrelated Gaussian background noise with standard deviation $\sigma = 5 \times 10^{-24}$ W, corresponding to a realistic integration time of two weeks assuming a noise temperature of roughly 9 K.
- **Background.** Correlated, non-thermal background. Its shape is modelled after spectra shown in [179], as well as unpublished MADMAX testruns in 2022, but does not involve fits to actual datasets. The following formula is used:

$$B(f)[\text{W}] = 10^{-20} \left[\operatorname{erf} \left(\frac{f - f_0}{5 \text{ MHz}} \right) \left(\frac{f_0}{f} \right)^3 + \exp \left(- \left(\frac{f - 25 \text{ MHz} (1 + \frac{r_1}{15})}{20 \text{ MHz} (1 + \frac{r_2}{10})} \right)^2 \right) \right] \\ + 4 \times 10^{-22} (1 + r_3) \sin \left(\frac{f + r_4 f_0}{2.5 \text{ MHz}} \right) + 5 \times 10^{-24} \left[(1 + r_5) \sin \left(\frac{f + r_6 f_0}{0.25 \text{ MHz}} \right) \right. \\ \left. + (1 + r_7) \sin \left(\frac{f + r_8 f_0}{0.1 \text{ MHz}} \right) \right], \quad (6.15)$$

where $f_0 = 4.218$ MHz and f are frequencies relative to f_{ref} and all r_i are independent random variables drawn from a Gaussian $\mathcal{N}(\mu = 0, \sigma = 1)$. The first line of eq. (6.15) and the first sinusoidal component determine the large-scale shape of the background, but are easy to distinguish from an axion signal with a FWHM of $\gtrsim 11$ kHz. Therefore two more sine-like components are introduced (last of the three terms in (6.15)) with random phase, amplitudes of order of the uncorrelated noise and fixed periods of 100 kHz and 250 kHz. This choice serves as a proxy for the expected, more intricate, structure of fluctuations in physically realistic scenarios. It represents a bad-case scenario where signal extraction nevertheless remains feasible. The background is notably different from spectra obtained in the data-taking of 2024 (see ch. 8) as the shape of the background crucially depends on the signal receiving chain. As long as the scales at which the background fluctuates are comparable to realistic values, differences in the overall shape are of little consequence.

Tab. 6.1 summarises the relevant parameters for the analysis.

The signals in the example application are small in the sense that the median of their signal to background ratios (SBR) is at 0.016. Fig. 6.2 shows the SBRs and signal to noise ratios (SNRs) for the mock datasets, represented by signal amplitude divided by background at the signal frequency and signal amplitude divided by noise standard deviation, respectively.

The data consisting of these three components is filtered with a fourth-order Savitzky-Golay (SG) filter [180] of a width of 221 data points, corresponding to 442 kHz. SG filters are moving window polynomial fits on an odd number of data points. They have two free parameters: The polynomial order o and the number of data points in a window w . For each window the polynomial value at the centre data point j is used as the filtered value $f_{\text{bg}}(X_j)$. SG filters leverage the fact that data points are equidistant. Under this condition the central value of any polynomial can be computed analytically from the data values making the method very efficient for

Table 6.1: Overview over the fixed parameters in the analysis and the priors used for the three non-fixed parameters. $f_{\min/\max}$ are the extremal absolute frequencies at which the measurement is sensitive.

Parameter	Value/ Prior
m_a	Uniform(f_{\min}, f_{\max})
$\mathcal{C}_{a\gamma}$	From [1, 9], excluding $\mathcal{C}_{a\gamma} < 0.86$
σ_v	Normal(218, 39) [km s ⁻¹]
ρ_a	0.3 GeV cm ⁻³
β^2	5×10^4
v_{lab}	242 km s ⁻¹
B_e	10 T
A	1 m ²

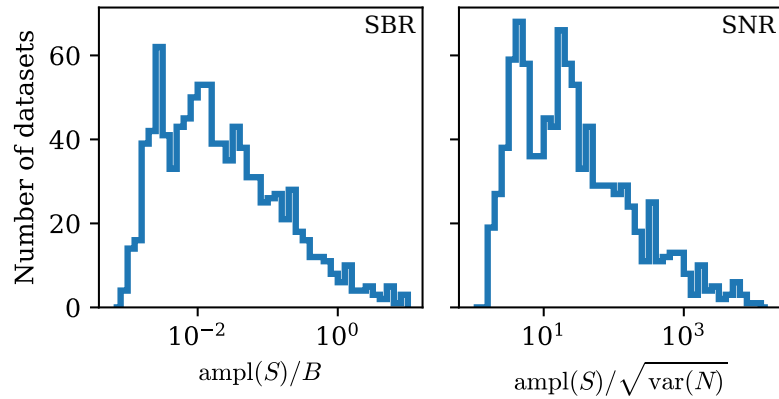


Figure 6.2: Distribution of the signal amplitude across the mock datasets, relative to the background and the noise level. **Left:** Signal to background ratio, given by the amplitude of the signal divided by the background at that frequency. **Right:** Signal to noise ratio, given by the amplitude of the signal divided by the standard deviation of the noise component. The bimodal shape of the SNR histogram is caused by the $\mathcal{C}_{a\gamma}$ prior the signals are sampled from.

window sizes up to several thousand data points. By construction it is difficult to obtain good estimates for the first and last $w/2$ data points in a dataset. For this reason these data points are removed from the dataset after the SG filter has been applied.¹

Subtracting the filtered from the raw data removes the correlated background component almost completely, but also slightly distorts the signal shape, as shown below. The parameters of the SG fit were chosen to yield a good reduction of the background while leaving signal and noise almost unchanged. The exact choice of SG parameters does not affect the validity of the approach. On the contrary: While the parameter bias induced by previous approaches crucially depends on filter parameters, the bias-free inference is independent of them, as long as the SG filter fulfils eqs. (6.4) and (6.5).

As the signal model for the Gaussian likelihood eq. (2.11) is used, multiplied by $\rho_a (g_{a\gamma}/m_a)^2 B_e^2 A \beta^2 q_e / \hbar$ to obtain the correct signal amplitude. Here, m_a is the mass of the axion, ρ_a its local energy density and $g_{a\gamma}$ its coupling to photons. B_e is the external magnetic field, A the surface area of the dielectric disks in the setup and β^2 the power boost factor (compare sec. 3.3). q_e and \hbar are electron charge and reduced Planck constant.

As standard deviation the amplitude of the uncorrelated noise is used, which is unknown in realistic scenarios and therefore has to be inferred from the data. If one would simply take the standard deviation of the background reduced data, the presence of a signal would lead to a slight bias for the noise towards higher values. To prevent the signal from biasing the noise-level estimation, first the SG filter is used to remove a background estimate from the data. Then each spectrum is partitioned into three frequency regions of equal size (roughly 8333 bins). Because the width of each of these regions is much larger than that of the localised filtered signal, it can only be present in up to two of the regions simultaneously. The region with the smallest standard deviation is selected for the noise-level estimation. This removes the bias caused by the presence of the signal. In the test case the inferred noise level has a negligible difference compared to the ground truth. This procedure works only if the background filter is unbiased and the noise level of the residuals is constant, which was verified for the present test case. In some applications relative residuals may need to be used, especially for frequency dependent N or if background oscillations depend on its absolute magnitude.

For Monte Carlo analysis the Reactive Nested Sampling algorithm [181–183] is employed via the Bayesian Analysis Toolkit in Julia (BAT.jl) [184].

6.2.2 Requirement Validation

Sec. 6.1 states several requirements for the bias-free method to be applicable. This subsection will show that they are all fulfilled in the example use case.

The first two requirements are fulfilled by construction of the mock datasets. The red curve in fig. 6.3 (bottom left panel) corresponds to $S - f_{\text{bg}}(S)$ and clearly shows, that it is not equal to zero everywhere. The signal S can also be parameterised as

¹SG filters will become important again in sec. 7.3, where the influence of the filter is further investigated and also transfer functions for different filter parameters are given.

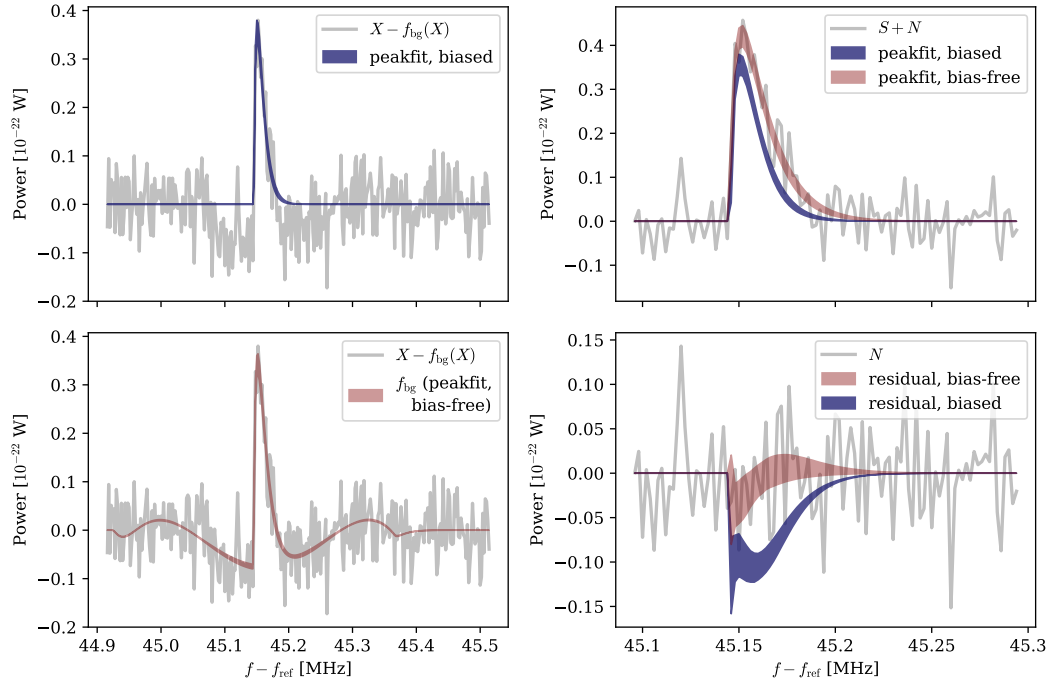


Figure 6.3: Effect of using a bias-free signal model for the fit on the data. 68% central posterior predictive intervals for the bias-free and biased fit are plotted in red and blue, respectively. **Top Left:** Data after background reduction via SG filter and biased fit without SG filter applied. **Bottom Left:** Same data and bias-free fit with SG filter applied. **Top Right:** Signal and noise components of the data as well as biased and bias-free peak-fit without applying an SG filter. **Bottom Right:** Deviation of posterior predictive of bias-free and biased peak-fit from the true signal (“residual”) compared to noise component of the data. The peak-fits in all panels have been performed on background reduced data as shown in the left panels. The data components shown in the right panels are for comparison only. Note the smaller frequency range plotted in the right panels.

described in the previous subsection. It remains to be shown that eqs. (6.6) and (6.7) approximately hold, which hinges on the choice of SG filter parameters.

As stated in sec. 6.1 already, the SG filter is considered to be a good background estimator if $|B - f_{bg}(B)| \ll |\text{ampl}(S)|$ and to be efficient at removing the noise component if $|f_{bg}(N)| \ll \sqrt{\text{var}(N)}$. Fig. 6.4 shows the distribution for both of these conditions using an SG filter with a width of 221 data points and polynomial order 4 on the mock datasets.

Fig. 6.4, left panel, demonstrates the chosen SG filter to be an excellent background estimator with practically all residuals being below 5×10^{-4} of the signal amplitude. The SG filter applied to noise retains the Gaussian behaviour of the noise component, but reduces its standard deviation by almost one order of magnitude to $0.13\sqrt{\text{var}(N)}$, as can be seen in fig. 6.4, right panel. The chosen SG parameters therefore represent a filter that fulfils all requirements listed in sec. 6.1.

In a more realistic simulated setting one could allow for proportionality between




Figure 6.4: Validation of the chosen SG filter parameters in terms of background estimation and noise removal. **Left:** Deviation between filtered and unfiltered true backgrounds for all data points normalised to the signal amplitude for each mock dataset (see eq. (6.6)). **Right:** Comparison between the filtered noise for all data points and the standard deviation of the unfiltered noise component for each mock dataset (see eq. (6.7)).

noise and background components. If their relationship is linear, instead of only subtracting the background estimate from the data, additionally dividing by the background estimate, in the absence of an axion, leads to a constant standard deviation over the whole frequency range for the residual data. In such a scenario, if the background is improperly removed or due to a signal, the noise level may vary slightly. The simplification applied here, assuming constant noise power independent of the background, does not lead to relevant errors. It was shown that the deviation from the truth of the background estimator is much smaller than the signal amplitude (fig. 6.2, left), which in turn is much smaller than the background. Therefore, small noise variations due to improper background estimation and the small underestimation of the noise at the signal location are not expected to be noticeable using any inference method.

6.2.3 Results

Fig. 6.3 shows the result for one of the datasets, demonstrating the effect of the bias-free vs biased method qualitatively. When the standard signal model is used to fit the signal peak (fig. 6.3, top left), the effect of the background removal via SG filter is not taken into consideration. Due to the presence of the signal, the background around the signal is overestimated, leading to a systematic decrease in signal amplitude and adjacent data points below the baseline (grey). A peak-fit with its 68 percentile central posterior interval (blue) is perfectly capable of fitting this modified signal, but cannot fit the surrounding data points - and therefore does not retain the true signal parameters.

The bottom-left plot in fig. 6.3 shows the same, but with a bias-free fit on the signal peak that takes the effect of SG filtering into account. A good fit is obtained over the whole frequency range, which displays the characteristic effect of an SG filter on the signal. Unbiased signal parameters can be obtained based on this fit, as fig. 6.3, top right shows. The non-filtered and background free signal peak, i.e. the signal




Figure 6.5: Reduced chi-square distributions for all mock datasets, obtained from fits using the standard biased method (blue) as well as the bias-free method (red). The reduced chi-square was calculated for the posterior median using the inferred noise level as described above and only in a range of 350 data points around the known signal frequency, because only in this region pronounced differences between the biased and the bias free method are to be expected.

component S , is fitted well by the posterior predictive central interval of the bias-free peak-fit on which no SG filter has been applied. The standard, biased peak-fit however underestimates signal height and width. This underestimation is made more visible in fig. 6.3, bottom right, where the deviation relative to the true signal is shown in comparison with the noise level, both for the biased and the bias-free peak-fit. While the 68 percentile of the bias-free fit includes the true signal over almost the full frequency range, the standard, biased fit displays significant deviation.

As a first quality-of-fit check, a histogram of reduced chi-square values for all mock datasets is shown in fig. 6.5. The distribution of chi-square values for the bias-free method is compatible with unity (1.005 ± 0.08). For the biased method, however, a distribution of chi-square values systematically above one is observed. While most datasets display reasonably good $\chi_{\text{red}}^2 \in [0.8, 1.5]$, some reduced chi-square values exceed even the upper bound of fig. 6.5, indicating extremely bad fits.

To further assess the quality of the fits, Bayesian coverage testing is performed and the Bayesian posterior is shown to be compatible with the true parameter values, but only when using the bias-free approach. From a frequentist perspective, repeating an unbiased analysis multiple times should lead to the true signal parameters being uniformly distributed over all marginal posterior quantiles. As simulated data are used access to the ground truth of the signal parameters is available, and as 10^3 equivalent mock-datasets have been generated, it is possible to verify the Bayesian results in this manner. The outcome is shown in fig. 6.6. The standard, biased peak-fit that does not take the effect of the SG filter into account systematically underestimates $\mathcal{C}_{a\gamma}$ and σ_v . It also shifts the axion mass to slightly larger values. The bias-free fit, in comparison, shows no significant deviation from the expected uniform distribution.

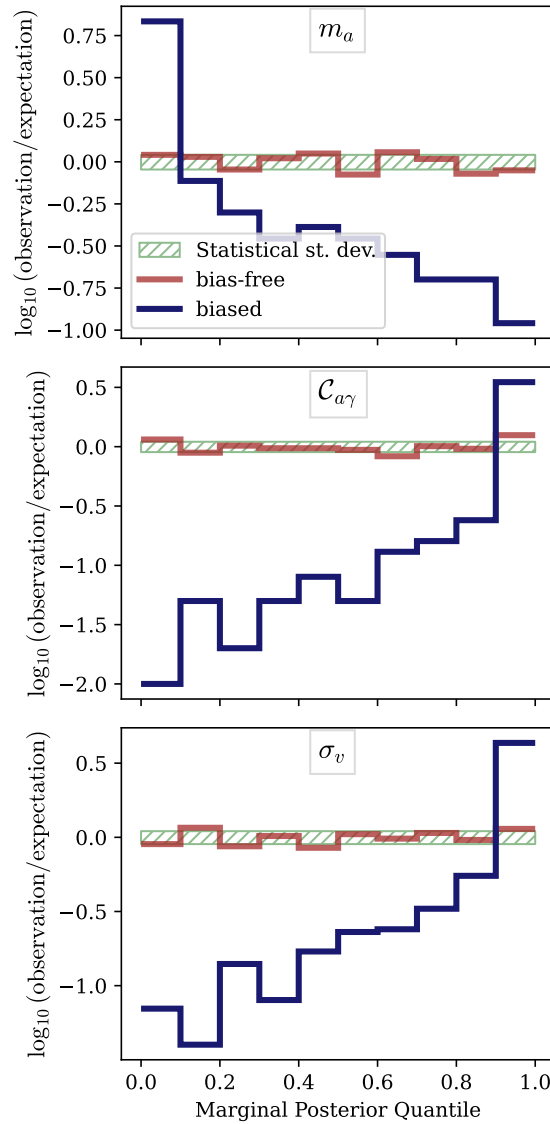


Figure 6.6: Distribution of the posterior quantile bin in which the true signal parameter resides, for the three parameters. The red curves show the bias-free parameter inference, the blue curves show the biased case. For 10^3 mock-datasets and 10 quantile bins 100 ± 10 cases are expected for each bin, which is shown as the green dashed region. The width of the green band corresponds to the expected statistical fluctuation. The distribution of the parameters in the bias-free approach shows no significant deviations from the expectation. For the standard, biased approach, a significant and systematic shift away from the true signal parameters can be seen. In this case almost all true signal parameters reside in the first or last quantile of the posterior distribution.

Table 6.2: Coverage properties of the Bayesian credible intervals: The table shows the percentage of cases in which true signal parameters reside in the 68% (resp. 95% and 99.7%) posterior credible interval, for each of the three free fit parameters.

	Coverage biased			Coverage bias-free		
	68%	95%	99.7%	68%	95%	99.7%
m_a	0.241	0.405	0.589	0.690	0.945	0.994
$\mathcal{C}_{a\gamma}$	0.055	0.131	0.426	0.643	0.925	0.993
σ_v	0.141	0.25	0.387	0.669	0.946	0.992




Figure 6.7: Posterior and prior probability density for signal parameter $\mathcal{C}_{a\gamma}$ for a mock dataset with no simulated axion present. The prior extends far beyond the region shown for $\mathcal{C}_{a\gamma} > 2$. As expected, the posterior is compatible with a no-signal hypothesis, sets an upper limit on $\mathcal{C}_{a\gamma}$, and shows no gain of information with respect to the prior at low values.

A more standard test to quantify coverage is given in tab. 6.2. With the standard biased method, only a small fraction of true signal parameters lies within the 68%, 95% or 99.7% credible intervals of the posterior, whereas the coverage for the bias-free method is close to the expected values.

Fig. 6.7 shows the prior and posterior distributions for a simulated dataset with no axion present. The probability density of the posterior is clearly severely suppressed for $\mathcal{C}_{a\gamma} > 1.5$, which hints at the limit setting potential of a MADMAX-like experiment with parameters as chosen above. Below $\mathcal{C}_{a\gamma} < 1.4$ posterior and prior follow similar distributions, no information has been gained. The $\mathcal{C}_{a\gamma}$ posterior is therefore compatible with a no-axion scenario, as expected.

Positive statistical fluctuations in several adjacent bins can potentially mimic faint signals in scenarios with no detectable axion present. To avoid erroneous detections in these scenarios it is crucial to not underestimate the noise level implemented in the likelihood and to design a statistically sound detection criterion. The latter is a highly relevant question on its own right, but requires extensive validation and is

beyond the scope of this project. For this reason $\mathcal{C}_{a\gamma} > 0.86$ is assumed a priori in the mock dataset, corresponding to detectable signal strengths for the present analysis (compare tab. 6.1).

6.3 Summary & Discussion

This chapter presents a method for fitting small-amplitude signals on top of an unparameterised background without incurring a bias on the signal parameters. If an unbiased, effective and linear background estimator is used to remove the background, its effect on the signal can be incorporated into the signal model. The problem then reduces to finding a single signal component in the presence of statistical noise.

One possible application for the method are axion haloscope measurements. In order to demonstrate the effectiveness of the bias-free approach, it was applied to 10^3 simulated MADMAX-like datasets using an SG filter as background estimator and compared with an analogous simulation using the standard, biased method. For inference a Bayesian approach using nested sampling was chosen. For the bias-free approach indeed all inferred signal parameters deviated from the ground truths only within the expected statistical fluctuations. The standard, biased approach, however, showed severe, systematic deviations. Should an axion signal be detected, applying the bias-free method instead of the standard approach therefore becomes necessary.

Beyond parameter inference in the case of an axion detection, it is reasonable to assume that the inclusion of the presented approach could improve the sensitivity of detection procedure, as the modified signal model will more closely resemble the respective features in the data. This exploration is left for future work.

7

The MADMAX Pipeline for Statistical Analysis

The previous chapter described in detail a method which in its scope is oriented towards future applications. It can be leveraged in any present axion haloscope search, but becomes crucially important only when using Bayesian methodology or in case of an axion detection.

Recently, however, a very present application became available in form of the first axion dark matter measurements with a MADMAX prototype setup. In order to analyse it, a set of methods had to be adapted or developed: A frequentist approach was chosen for the analysis, an existing analysis pipeline [83] had to be modified to be used on the data and additional checks and pre-processing steps were developed in order to obtain a competitive and reliable sensitivity. These methods are introduced in the present chapter, before being applied in ch. 8.

As in the previous chapter, it is again assumed that the analysed data resemble a set of frequency spectra containing background, noise and signal components, which will be defined similarly as before. In an axion haloscope context after the background component is removed, e.g. using a non-parametric filter as before, all bins in the residual spectra should be uncorrelated samples from a Gaussian distribution with zero mean. Bins containing a potential signal component should follow a distribution with positive mean, but make up at best a tiny fraction of the total number of bins in the spectrum. They therefore do not significantly alter the expected Gaussian distribution for the overall residual. The bins of the spectra have to be uncorrelated because correlations between adjacent bins with correlation lengths comparable to the FWHM of an axion signal have the potential to significantly impair the sensitivity of an experiment. Deviations due to other effects such as residual radio-frequency (RF) structures may lead to bins which are samples from distributions with positive mean potentially mimicking an axion signal and are therefore also to be avoided.

Usually to test Gaussianity a histogram is constructed from a set of data points, which is then compared with a Normal distribution. This test, however, is not sensitive to correlations between individual data points. An additional method based upon the work in [185] is therefore presented here that is sensitive to Non-Gaussianity and better-suited to identify correlations between data points. The method is called *Runs Gaussianity* test in the following. To demonstrate its validity and sensitivity, in the present chapter the Runs Gaussianity test is only applied to abstract sets of data points not intended to carry physical meaning. It will be applied in the context of the MADMAX haloscope in ch. 8, in which case the data points are individual bins of one frequency spectrum.

Should the Runs Gaussianity or similar tests indicate deviations from the expected behaviour for the residual spectra, these deviations need to be investigated and if possible avoided. Some background filters may be better suited than others to remove specific contributions to the background. Savitzky-Golay filters which are typically used for this task in axion haloscopes can serve as high-pass filters, suppressing all components of a spectrum that fluctuate on large scales. They are, however, not ideal for dealing with sinusoidal contributions to the background that fluctuate on a specific scale, i.e. components with a period of a specific number of data points. This is true especially if their scale is close to the FWHM of an axion signal. When such contributions are present, a notch filter is required, capable of only removing fluctuations of a specific scale. For the analysis of the first MADMAX axion measurement such a filter proved necessary. A *digital notch filter* (DNF) is therefore developed here, comparing different possible implementations and investigating its effect on potential axion signals.

Having processed a set of spectra to the point where only an uncorrelated noise and potentially a signal component remain, spectra covering the same frequency regions have to be combined in a way that enhances the signal-to-noise ratio (SNR) as much as possible. A possible procedure is described in [83], most parts of which are followed in the present work. With a combined frequency spectrum the search for an axion signal becomes possible, for which either frequentist or Bayesian methods could be applied. One simple option from the frequentist tool-box is to cross-correlate the spectrum with the expected axion lineshape and to investigate the statistical significance of the largest excess. If it is not deemed to be statistically significant, a limit can be set. If it is, a set of strategies needs to be in place to rule out any non-axionic origin.

This chapter is structured as follows: First, the Runs Gaussianity test and digital notch filter are introduced in secs. 7.1 and 7.2, respectively. Sec. 7.3 then covers combining individual spectra and cross-correlation with the axion lineshape for simulated datasets. Here important systematic effects as well as deviations from the standard procedure in [83] are pointed out. In sec. 7.4 a limit is constructed for the simulated datasets before sec. 7.5 adds a discussion on procedures for dealing with significant excesses. The chapter concludes in sec. 7.6.

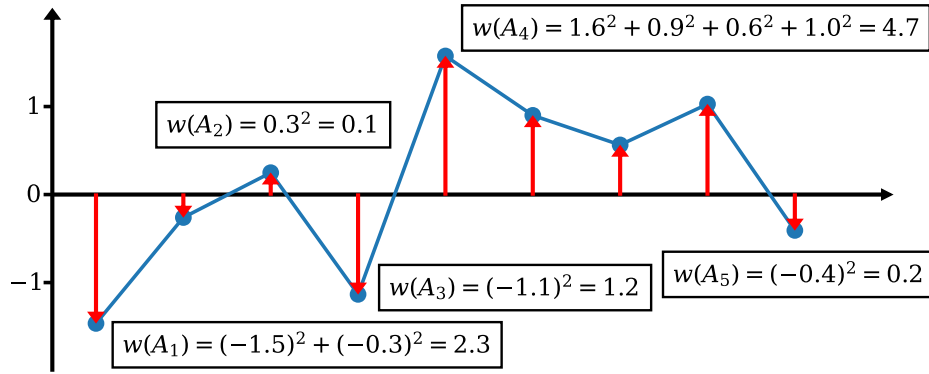


Figure 7.1: Exemplary calculation of runs weights for 9 data points drawn from independent standard normal distributions ($\mu = 0$, $\sigma = 1$). Adjacent points above and below the mean are summed up to obtain the weights defined in (7.1).

7.1 Runs Gaussianity Test

7.1.1 Description and Validation

In [185] the authors define a novel test statistics T for a set of independent data points X_i , which are normally distributed with known means μ_i and variances σ_i . The basic idea is to split the set of N data points into individual “runs” A_j defined as sets of consecutive data points above or below their respective mean values. In fig. 7.1 for example $A_4 = \{X_5, X_6, X_7, X_8\}$. In the context of the MADMAX haloscope the X_i are frequency bins in a spectrum after the background has been removed and their mean value is zero. For each run a weight is calculated as

$$w(A_j) = \sum_i \frac{(X_i - \mu_i)^2}{\sigma_i^2}, \quad (7.1)$$

where the sum runs over all data points contained in the run A_j . A graphical representation for a simple example is given in fig. 7.1. The authors of [185] continue to define a test statistics T as $T \equiv \max_j w(A_j)$ for only runs above the mean and spend significant effort calculating p -values for it. Notably, small p values corresponding to big T do neither require nor give information on the shape of the structure that caused the big value for w beyond it being some sort of localised excess.

This test statistic has already been successfully applied to find simulated axion signals in the context of early MADMAX-like datasets [179]. It is certainly possible to find potential axion signals, the shape of which is well known (see sec. 2.4) using this method. Its shape-agnostic nature, however, makes it especially well suited for applications in which the user does not precisely know what they are looking for.

The following application leverages this fact: Looking at the full distribution of weights for all runs can serve as a powerful test for any kind of variation from Gaussianity as will be shown below. The statement could be extended to other distributions, but would require the generation of dedicated reference datasets following said distributions.




Figure 7.2: Probability density distributions for w given reference datasets of various lengths N .

The abstract idea for the method presented here is as follows:

1. Generate a reference distribution of $w(A_j)$ for a simulated dataset with N data points, where all data points independently follow normal distributions with $\mu = 0$ and $\sigma = 1$. This means the simulated dataset contains only Gaussian white noise. Contrary to [185], this analysis considers weights for runs above the mean as well as runs below the mean.
2. Calculate the distribution of $w(A_j)$ for an actual measurement that has been processed so only a Gaussian white noise component should be present in the data.
3. Compare the two distributions visually and using the quantitative metric G introduced below.

For the application on a haloscope like MADMAX a dataset could be a single frequency spectrum after background removal, that consists of N frequency bins.

The test statistic T of [185] is applied in [179], where the authors divide their dataset into many chunks with $N = 96$ data points each in order to be able to analytically calculate p -values for each chunk. The present analysis does not impose such scrutiny and is thus not limited to small parts of the dataset. Nevertheless, reference datasets may feature a different number of data points as the measurement the user wants to compare it to. This means an investigation is in order to determine the influence of N on the dataset. Fig. 7.2 shows reference datasets for a broad range of N . To generate the distributions shown there the weights for runs are calculated in datasets containing only Gaussian white noise with length N . This process is repeated until for each N roughly 10^7 weights are obtained. This number of weights was chosen to get reasonable statistics down to probability densities of 10^{-6} corresponding to $w \approx 30$. No significant differences are visible between $N = 2 \times 10^5$ (blue) and

$N = 10^3$ (orange). Only for $N = 100$ large values for w show decreased probability, which could be explained by edge effects.

To quantify how well a distribution resembles the reference, one can introduce a metric G as follows. Compute the weights for all runs A_j^{meas} in a given measured dataset “meas”. Analogously to the histograms in fig. 7.2 define a set of K ranges r_k between 0 and the maximal weight in the measured dataset

$$r_k = \left[\frac{k-1}{K} \max_j w(A_j^{\text{meas}}), \frac{k}{K} \max_j w(A_j^{\text{meas}}) \right], \quad (7.2)$$

with $k \in [1, K]$. Define ν_k^{meas} as the number of runs weights with values in a range r_k for the dataset. Find ν_k^{ref} , the number of runs weights $w(A_j^{\text{ref}})$ in the same range for a reference dataset “ref” and scale it to the same number of runs as the measured dataset, i.e. $\sum_k \nu_k^{\text{ref}} = \sum_k \nu_k^{\text{meas}}$. One can then construct a metric for the difference between the two datasets as

$$G \equiv \frac{1}{K} \sum_k \frac{(\nu_k^{\text{meas}} - \nu_k^{\text{ref}})^2}{\nu_k^{\text{ref}}}. \quad (7.3)$$

Should $\nu_k^{\text{ref}} = 0$ for one of the ranges r_k , e.g. because the measured dataset includes one very big weight, an extrapolation of the exponential trend visible in fig. 7.2 is used.

To investigate which values of G are compatible with a dataset being distributed according to Gaussian white noise, the distribution of G is numerically approximated using simulated datasets with a known Gaussian distribution.

For fig. 7.3 10^4 abstract mock datasets are simulated, each with 4×10^4 data points independently sampled from a normal distribution with $\mu = 0$ and $\sigma = 1$. G is then calculated for each dataset varying the number of data points N or the number of ranges K used in the calculation of G .

The resulting distributions show peaks around $G = 1$ with slight variations in mode and full width at half maximum (FWHM) of the peak depending on the investigated parameters. Selected properties of the G distributions are given in tab. 7.1. All distributions display significant tails for large G , therefore even for $G \gg 1$ it cannot be entirely ruled out that the observed value is due to a statistical fluctuation and does not come from a deviation from Gaussianity. When G is constructed using only $K = 20$ ranges, the resulting histogram is shifted compared to using larger K . In the following therefore $K = 50$ ranges are used for the calculation of G . With this K a conservative value of $0.7 < G < 3.5$ is demanded to confidently state that no significant deviations from Gaussian white noise behaviour are observable.

When Poisson statistics is assumed, ν_k^{ref} in the denominator of the sum in (7.3) can be seen as a variance and G should resemble a reduced χ^2 . Investigations into the exact role of G in standard statistics terms as well as defining a concrete framework for associating a p -value with a specific value of G are left for future work. Doing so may require slightly modifying the calculation of G to yield more predictable values in its big G tail.

For an implementation of the method presented above for the programming language Julia see <https://gitlab.desy.de/johannes.diehl/runsgaussianity>.




Figure 7.3: Dependence of the metric G as defined in (7.3) on parameters N , the number of data points in the dataset, and K , the number of ranges used. The panels show distributions of the metric G for 10^4 simulated Gaussian white noise test datasets. **Left:** Dependence of G on N , each using $K = 50$ ranges. **Right:** Dependence of G on K , each for 10^3 data points. The black dashed line marks unity, which would be median and mode of the distribution if ν_k^{ref} would accurately follow a Poisson distribution. The grey dashed lines mark the boundaries of the region deemed compatible with the assumption of a Gaussian white noise distribution. Datasets with $G > 5$ are not shown to improve readability of the plot. They make up roughly 2-3% of all datasets.

Table 7.1: Selected properties of the G distributions shown in fig. 7.3.

N	K	min G	max G	G 5th-percentile	G 95th-percentile
10^3	20	0.14	49	0.51	1.77
10^3	50	0.51	517	0.78	3.34
10^3	100	0.56	301	0.84	3.14
2.5×10^4	50	0.46	120	0.75	3.49
4×10^4	50	0.51	532	0.82	3.04

7.1.2 Example Applications

In the following the Runs Gaussianity test described above is applied to multiple simulated datasets in order to demonstrate its capability in determining deviations from Gaussian white noise behaviour.

To cover a variety of scenarios three datasets are simulated, each with 2.5×10^4 data points, and compared with a reference dataset containing the same amount of data points. They are intentionally kept as abstract as possible, but in a haloscope context the datasets would correspond to frequency spectra and the index of the data point to the index of a frequency bin, i.e. a specific frequency. $K = 50$ ranges are used for the calculation of G . The three datasets each include a component of Gaussian white noise and additionally one of the following features:

- A Gaussian peak centred around the 1300th data point with a standard deviation of 10 data points and an amplitude of $1.5 \times$ the noise standard deviation (see inset in top row of fig. 7.4).

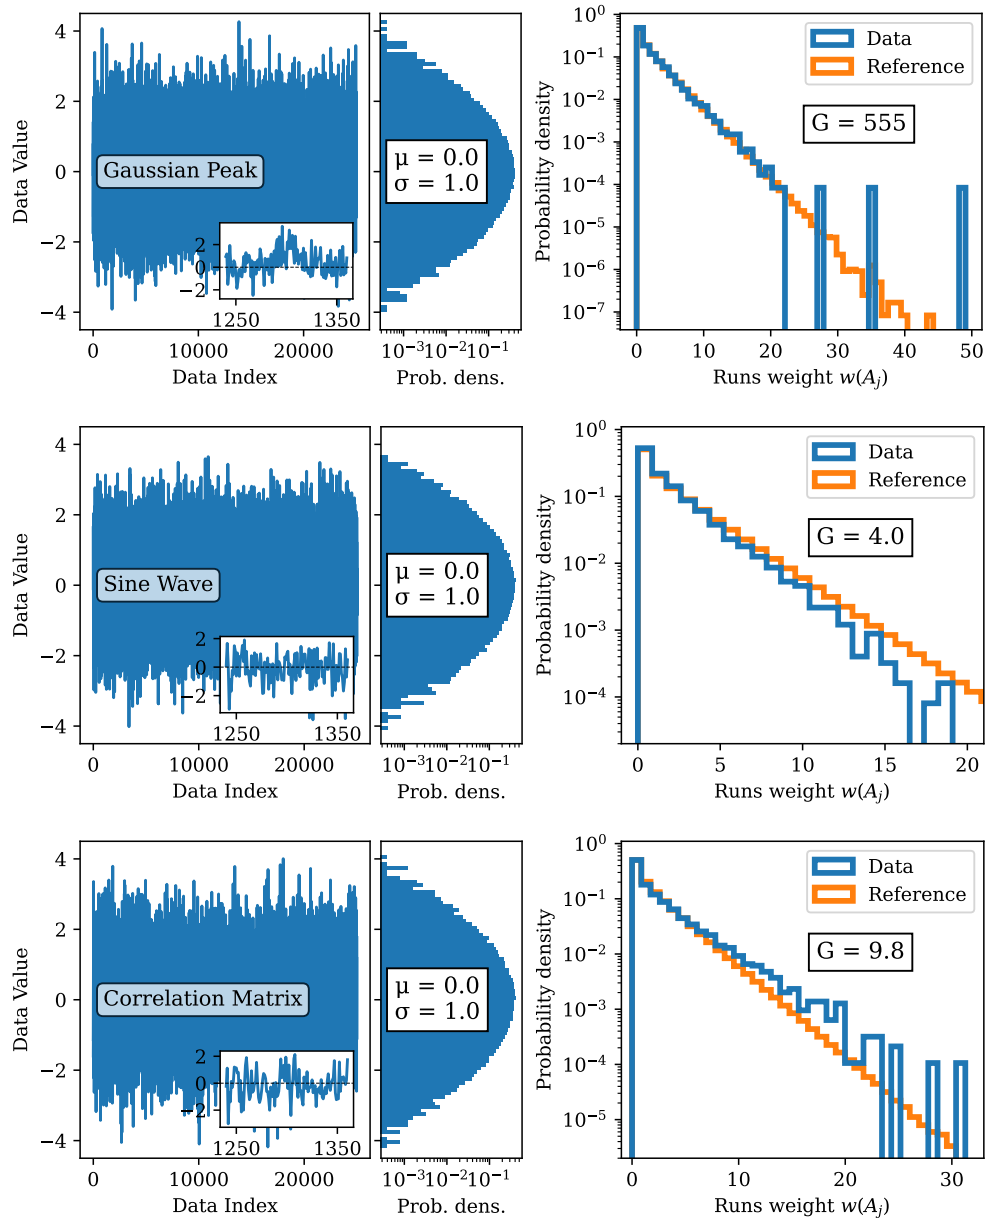


Figure 7.4: Three example use cases demonstrating the power of the runs Gaussianity test. From top to bottom: Peak search, inserted sine wave, correlated noise. For each application the full dataset is plotted with an inset showing a small relevant region (left). Centre panels contain histograms of the dataset often used to investigate Gaussianity. μ and σ give best fit values for fitting a Gaussian to the histogram. Panels on the right show results for the runs Gaussianity test together with G defined in (7.3).

- A sine wave with wavelength of 2 data points and amplitude $0.5\times$ the noise standard deviation. Note that a wavelength this small effectively anti-correlates adjacent data points.
- A correlation matrix $C_{ij} = 10^{-|i-j|}$ is multiplied to the white noise. This leads to adjacent data points having some correlation (Pearson coefficient: 0.19) which drops off exponentially for more distant data points.

The results can be seen in fig. 7.4¹. While in all three cases a simple histogram of the datasets perfectly agrees with a standard Gaussian distribution (middle panels), the Runs Gaussianity method is able to detect the deviation from pure white noise behaviour.

A peak in the dataset manifests in one or two bins with significantly increased runs weights (top right panel in fig. 7.4). The sine wave injected in the second example makes it likely that adjacent data-points have opposite signs. Therefore the probability density is suppressed for large runs weights (middle right panel). The correlation induced in the third example has the opposite effect. Adjacent bins are more likely to have the same sign, therefore the average run contains slightly more data points leading to increased probability for obtaining large runs weights (bottom right panel).

In all three examples deviations from Gaussianity are very challenging to spot by visual inspection or using a histogram of the dataset. The runs Gaussianity method, however, is able to reliably spot the deviation. In all three cases is the metric G larger than the values deemed compatible with the Gaussian hypothesis, $G \in [0.7, 3.5]$, hinting at a Non-Gaussian component in the dataset. The test was repeated with multiple mock datasets to ensure the results shown in fig. 7.4 are representative.

7.2 Digital Notch Filter

The Aachen-Bonn receiver system as used for the 2024 data taking at MORPURGO introduced a small non-thermal background contribution manifesting as a sine-like pattern in the measured frequency spectra with scales similar to the axion linewidth.

In a Fourier transformed version of the measured frequency spectra, i.e. in time domain, the contribution shows up as a series of localised peaks which can be removed while only minimally affecting the rest of the spectrum as well as, crucially, a potential axion signal. This method of removing peaks in the Fourier spectrum corresponding to unwanted background components is called digital notch filter in the following. The filter is digital in contrast to the physical frequency filters used in the receiver chain. It is a notch filter, because it suppresses only a small part of the Fourier spectrum, leaving the remaining spectrum untouched.

7.2.1 Description and Example

The digital notch filter is applied to a Fourier transform of a dataset. It is therefore helpful to introduce Fourier transforms using a mock frequency spectrum resembling an axion haloscope measurement before presenting the method.

¹In the left panels showing the full datasets it may seem like the standard deviation is larger than one. This effect is only due to the line width used in the plot being non-negligible, making it impossible to resolve all 2.5×10^4 data points in the plot.

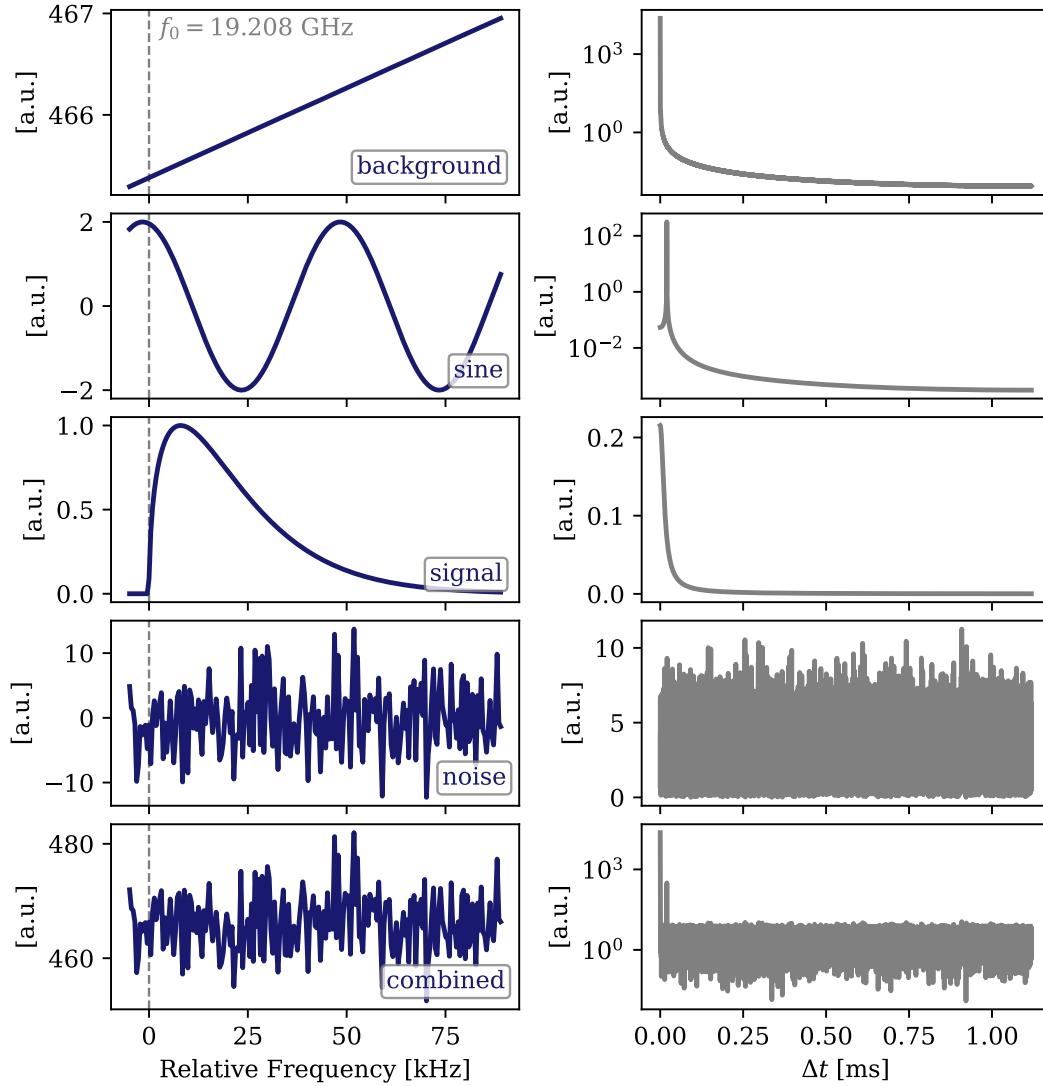


Figure 7.5: Fourier transform of a mock frequency spectrum. **Left:** A piece of the mock frequency spectrum around the simulated axion signal (at $f_0 \approx 19.208$ GHz) with frequencies relative to the signal. The top four panels show individual components as described in the text and the bottom plot shows the spectrum with all components combined. **Right:** Absolute values of the Fourier transforms of the individual components and of the combined mock spectrum are shown.

A part of such a spectrum dissected into its components can be seen in fig. 7.5 (left column). The full spectrum has 10^5 data points with a bin width of 447 Hz covering a 45 MHz range between 19.186 GHz and 19.231 GHz. It consists of

- a background obtained from part of a single timestamp of a rescaled real measurement, to which a smoothing Savitzky-Golay filter² was applied (width 4999 data-points, polynomial order 3)
- a sinusoidal component with a “wavelength” of 50 kHz and amplitude of two times the amplitude of the axion signal
- an axion signal at 19.208 GHz
- a Gaussian white noise component with $\mu = 0$ and a standard deviation of 1% of the background level

The right column of fig 7.5 shows the absolute values of the Fourier transforms of the individual components as well as the full dataset (bottom). Fourier transforms generally produce complex-valued information, from the real and imaginary part of which the original dataset can be recovered. The frequency spectra obtained in MADMAX measurements are themselves Fourier transformed versions of a discrete data stream in time. A Fourier transform of these spectra would therefore yield the initial data stream if the full complex frequency spectrum would be retained. Since only a real valued spectrum is available, the complex phase information is lost and the Fourier transform of the frequency spectra cannot be interpreted as corresponding to real, physical time.

Nevertheless, as they are Fourier transforms of spectra that carry frequency units, the Fourier spectra on the right of fig. 7.5 carry units of time. The scaling of the x axis is determined by the frequency spectrum: The bin width of the mock frequency spectrum under consideration of 447 Hz leads to a range in the Fourier spectrum of $\frac{1}{2} \times \frac{1}{447 \text{ Hz}} = 1.12 \text{ ms}$. The full range of the frequency spectrum of 44.7 MHz leads to a resolution of the Fourier spectrum of 22.4 ns.

Fig. 7.5 only shows the absolute values of the Fourier spectra, because the individual complex components do not carry physical meaning. From top to bottom, the Fourier transform of the background component alone shows a dominant peak at small times with a strong drop-off towards bigger times. This is consistent with the construction allowing only the background to contain fluctuations on broad scales in the frequency spectrum. Secondly, the sine wave component as expected shows up as a localised peak in the Fourier spectrum at $\Delta t = 20 \mu\text{s}$. This value was chosen to roughly correspond to the Fourier transform of the axion signal (third row) having dropped to $1/e$ of its amplitude. The Fourier transform of the axion signal is similar to a Gaussian for small Δt which can be explained by the axion peak also being similar to a Gaussian. The white noise component (forth row) has the real and imaginary parts of its Fourier spectrum independently following Gaussian white noise distributions as well. The absolute value shown in the figure therefore resembles the absolute value of Gaussian white noise. The two peaks from the first two components are still visible over the noise floor in the Fourier transform of the full mock spectrum (last row).

²More on this class of filters in secs. 6.2 and 7.3.




Figure 7.6: Fourier peak of the sinusoidal component before and after the digital notch filter was applied. The figure shows part of the full Fourier spectrum (see also fig. 7.5, bottom right) before (grey line) and after replacing the peak with various methods as described in the text.

The purpose of the digital notch filter is to remove the sinusoidal component in the mock frequency spectrum while ideally leaving all other components untouched. To do so the following steps are applied:

1. Obtain some information about the spectral features that are to be removed. This information can but does not need to come from the spectrum under investigation. For each feature i , information on the range of the Fourier spectrum at which the feature appears is needed. The example presented here appears at $l_i = 20 \mu\text{s}$ and has a width of $w_i \approx 90 \text{ ns}$.
2. Apply a Fourier transformation to the dataset under investigation. This results in a complex-valued Fourier spectrum.
3. Remove bins in the regions $l_i \pm w_i$ from the real and imaginary parts of the Fourier spectrum. Replace them with values drawn randomly from real and imaginary components of the Fourier spectrum in regions with negligible deviations from a noise only spectrum. This replacement method is called “random” in the following.³
4. Apply a back transformation to obtain the original spectrum without the removed spectral features.

The method as presented above has some pre-requisites. The dataset under investigation may not contain gaps or other sorts of missing data. The method only works if w_i is much smaller than the distances between the features and if $\sum_i w_i$ is much smaller than the full range of the Fourier spectrum.

Other methods for replacing the spectral feature could include setting both real and imaginary part of the Fourier spectrum equal to zero in the regions $l_i \pm w_i$ or separately

³Note that standard numerical implementations of Fourier transformations result in two mirrored spectra such as the one shown in fig. 7.5, right. Therefore in practice to remove one spectral feature this procedure has to be applied to the Fourier spectrum twice.

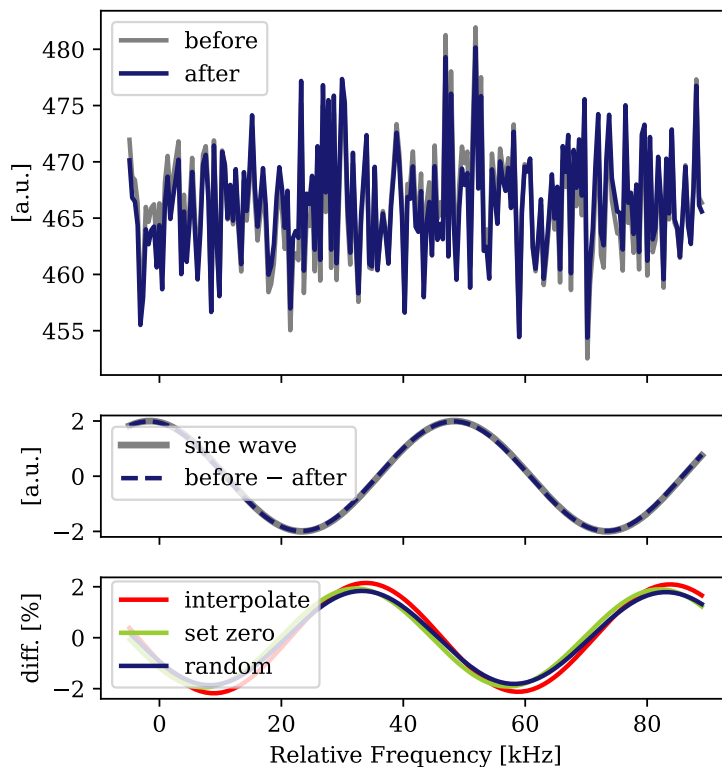


Figure 7.7: Effect of an error from the digital notch filter. **Top:** A slice of the mock frequency spectrum before (grey) and after (blue) the digital notch filter using random replacement has been applied. **Middle:** Absolute difference between the two spectra shown in the top panel (blue, dashed) compared with the sinusoidal component (grey) that should be removed by the digital notch filter. **Bottom:** Difference between two curves in the middle panel, i.e. between component subtracted by digital notch filtering and the sinusoidal component, divided by the amplitude of the sine wave. All three replacement methods show similar results. The amplitude of the error made by the digital notch filter is very small compared with the other scales involved.

linearly interpolating real and imaginary parts between the first data points outside of the region. These methods are called “set zero” and “interpolate” in the following. The effect of the digital notch filter on the Fourier spectrum of the example with all three peak replacement methods can be gleaned from fig. 7.6. With the values chosen above, 9 points in the Fourier spectrum were replaced with random values from other parts of the spectrum. Replacing fewer data points may not completely remove the feature, while replacing more could affect other components of the data set that are to be left untouched.

Fig. 7.7 shows the effect of the filter on a part of the mock frequency spectrum. In the top panel, one can clearly see parts of the spectrum being shifted to higher/ lower values in an alternating way. This difference between spectrum before and after the digital notch filter has been applied follows a sinusoidal pattern closely resembling the sine wave that is supposed to be removed by the filter (middle panel). The filter does not perfectly capture amplitude or phase of the sine wave component as the bottom




Figure 7.8: Effect of the digital notch filter on the shape of a potential axion signal. The top panel shows unaltered axion signal (grey) as well as axion signal after filtering using the random replacement method (dashed blue). The bottom panel shows the signal shape modification from the digital notch filter, i.e. the difference between “before” and “after” in the top panel, for all three replacement methods mentioned above.

panel demonstrates by showing the difference between filter effect and sine wave. It is, however, very close with the amplitude of the error made by the filter in this example being at $\sim 2\%$ of the amplitude of the removed component or $\sim 0.1\%$ compared to the absolute values of the dataset.

The digital notch filter has so far been shown to be very efficient at removing a specific sinusoidal component. For experimental efforts, however, it is also crucial to not affect potential signals. Applying the filter to a signal-only dataset only introduces deviations at the permille level compared to the signal amplitude as fig. 7.8 shows. This result holds for all three peak replacement methods, interpolate, set-zero and random.

7.2.2 Validation

In order to be able to use the digital notch filter in the context of the MADMAX haloscope, one needs to make sure that the findings obtained from the example in the previous subsection generalise. The error introduced by subtracting the sine wave (fig. 7.7, bottom) as well as the signal shape alteration induced by the filter (fig. 7.8, bottom) have been investigated. It remains to be shown how both quantities evolve with respect to amplitude and wavelength of the subtracted sinusoidal component. The result can be seen in fig. 7.9 as a function of amplitude of the sine wave (left) and its wavelength (right).

Firstly, for changes in signal shape (top panels) the maximal deviation between axion signals of unity amplitudes before and after applying the digital notch filter using the random replacement method is plotted. The maximal deviation is used,




Figure 7.9: Influence of amplitude (left) and location (right) of the Fourier peak to be removed on signal shape (top) and error for the removal of the sine wave (bottom). For the top panels the maximal deviation between true signal shape and signal shape after filtering is shown, while the bottom ones use the median of the absolute deviation between inserted sine wave and component removed by the filter.

since the deviation is largest at frequencies with high signal strength where the impact of deviations is greatest.

Varying the amplitude of the sine wave has no influence since signal alteration depends only on parameters of the digital notch filter (top left). An increase in the amplitude of the sinusoidal component could affect the axion signal shape only by making it necessary to choose a larger value for w_i . When the wavelength of the sinusoidal component is increased, shifting the Fourier peak to larger times, the influence on a potential axion signal decreases exponentially (top right).

Secondly, the error in the subtracted sine wave is estimated by taking the median deviation between the inserted sine wave and the difference of the spectra before and after filtering, divided by the amplitude of the sine wave (bottom panels). Again, random replacement was used. The median is used since frequencies at the low- and high-frequency boundaries of the datasets tended to display larger deviations. These parts of the dataset are not too relevant in the analysis of MADMAX data (ch. 8), since they display low boost factor values and may be cut in later stages of the data processing.

A larger sine amplitude at constant noise level makes it easier to remove leading to the exponential drop off observed in fig. 7.9, bottom left panel. Shifting the Fourier peak to larger times also makes it slightly broader. Keeping w constant, this leads

to slightly worse inference of the sinusoidal component, which explains the upwards trend in the bottom right panel.

Given the replacement method is inherently statistical, the mean of 50 repetitions of each procedure is taken to ensure the results are representative.

Additional tests to further improve confidence in the digital notch filtering method could include varying backgrounds, noise levels or signal strengths and positions. Also removing multiple spectral features could be investigated. Given the smallness of the errors involved, however, such efforts are left for future work.

7.3 Constructing the Grand Spectrum

This and the subsequent section describe in detail how to obtain a limit for a set of raw measurement data. The procedure is similar to the analysis pipeline established by the HAYSTAC experiment, presented in detail in [83]. Multiple steps however could be simplified or had to be slightly altered, which is why it is described here in its entirety.

As explained in sec. 4.1, a frequentist analysis on an abstract level requires a test statistic with a known distribution given a specific hypothesis about nature. Furthermore, a way to translate parameters of this distribution into physically meaningful quantities, $g_{a\gamma}$ in the present case, is needed. The following explains how to obtain both from available raw measurements.

It is instructive to briefly go over all data processing steps before returning to each one in more detail. Frequency spectra are taken in chunks integrating over 900 s of measurement time. They are called *raw spectra* and can be thought of as consisting of two or three components (compare Fig. 6.1): A correlated noise background, uncorrelated Gaussian white noise and potentially an axion signal. For data processing the background component needs to be removed to obtain *processed spectra*, displaying Gaussian white noise behaviour with zero mean in the absence of an axion signal and a standard deviation depending on the integration time. The processed spectra are combined using maximum likelihood (ML) weights, which depend on the relative strength of a reference axion signal in each bin. The resulting *combined spectrum* is then cross-correlated with the axion lineshape to obtain a better signal to noise ratio (SNR) in the *grand spectrum*.

Simultaneously, several sets of calibration data (see sec. 3.3) need to be used to determine the potential strength of a reference axion signal in every bin of the spectra in order to fulfil the second requirement of a frequentist analysis.

This section uses mock spectra simulated to resemble actual spectra to understand the reasoning behind all processing steps, similarly to ch. 6. Ch. 8 will then introduce in detail the full CB200 dataset from measurements in 2024 and apply the pipeline to obtain physical limits.

Tab. 7.2 gives an overview over the nomenclature used for the various spectra introduced in this section.

7.3.1 Raw to processed spectra

Two example spectra used in this section are shown in fig. 7.10 (top). Every spectrum represents the receiver output after 15 min of integration time. The spectra have units

Table 7.2: Variables introduced for different spectra used in the analysis. Index i always indicates a specific timestamp and j a frequency bin, therefore combined and grand spectra only include the latter index.

Variable	spectrum
u_{ij}	uncalibrated spectrum
b_{ij}	estimated background
p_{ij}	processed spectrum
r_{ij}	rescaled spectrum
c_j	combined spectrum
g_j	grand spectrum

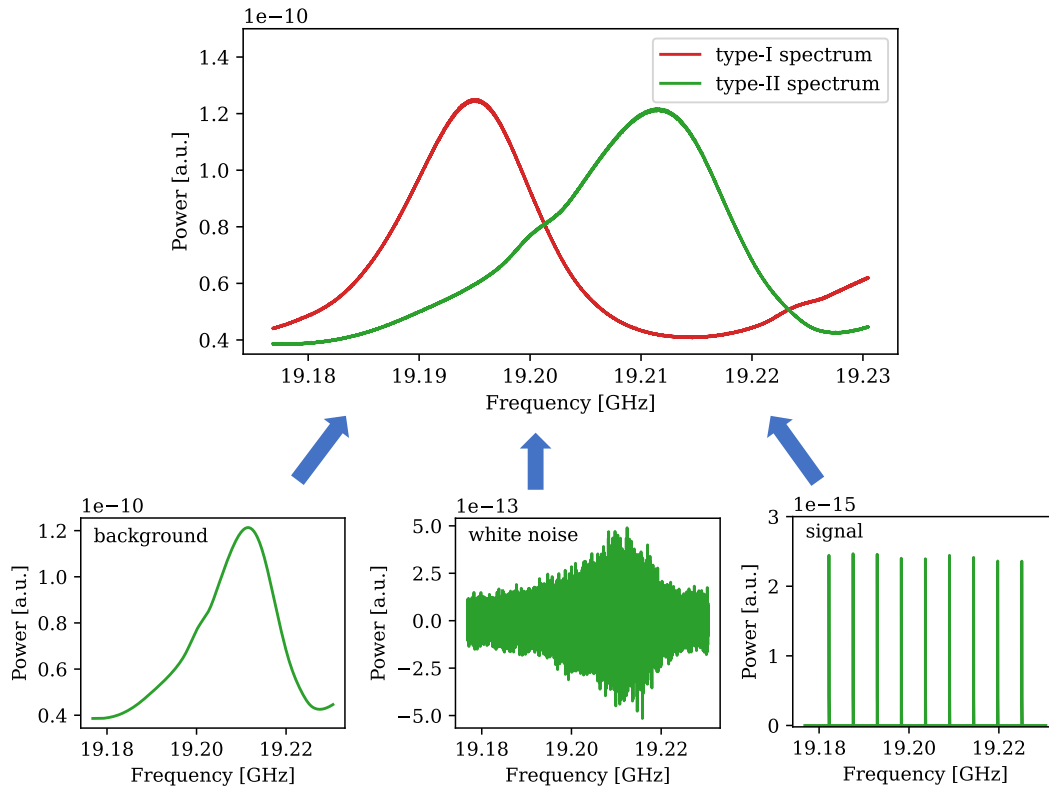


Figure 7.10: Datasets simulated for this section consisting of three components: Background, white noise and equidistant signals (bottom, from left to right). Backgrounds are smoothed versions of actual 15 min spectra from two different booster states. The spectra have 6×10^4 data points each, with bin widths of $\Delta\omega = 894$ Hz being equal to the resolution bandwidth.

of power per frequency bin, but since they are uncalibrated, the numerical value does not contain physical information. The goal of this subsection is to convert these raw spectra to Gaussian random samples⁴.

This section aims at developing a good understanding of the pipeline before applying it to real data. Therefore the spectra are constructed from known components as shown in fig. 7.10 (bottom). Parts of smoothed versions of raw data from two distinct datasets, i.e. data taken using two different booster states, of the CERN 2024 run with CB200 are used as background (bottom left). Spectra from the same dataset resemble each other closely, therefore they are denoted as type-I and type-II spectra in the following and only one of the 246 and 164 spectra available, respectively, is shown as an example in fig. 7.10. Statistical fluctuations on top of this background are modelled as Gaussian white noise with a standard deviation of 0.11% of the background level at each frequency. Additionally, nine signals are added to the datasets at equidistant frequencies. Their signal strengths are chosen to be equal after applying the calibration in both datasets individually. For the dataset using type-I baselines, $B_e = 1.0$ T and a constant $\beta = 2.5 \times 10^3$ are used. For the other dataset resembling type-II measurements, $B_e = 1.5$ T and a constant $\beta^2 = 0.9 \times 10^3$ are chosen. These choices purposefully lead to unrealistic exclusion limits, the spectral shape of which will be dominated by variations in system temperature. They are, however, useful to validate the procedure.

The spectral shape of the correlated noise background of the raw spectra depends on the exact RF characteristics of the receiver chain, which are prohibitively difficult to simulate. Since an analytic, parametric description of the background is not available, a non-parametric fit is employed. This makes simultaneous signal and background estimation impossible, introducing a bias on signal parameters (more details Ch. 6). Once the fit has been performed, the power excess p_{ij} in bin j of the processed spectrum p_i can be computed via

$$p_{ij} = \frac{u_{ij}}{b_{ij}} - 1 \quad (7.4)$$

with u_{ij} being the value in bin j of the uncalibrated raw spectrum and b_{ij} the corresponding value of the fit. This means in terms of the notation introduced in sec. 6.1 $b_{ij} = f_{\text{bg}}(X)$ and $u_{ij} = X$. Eq. (7.4) is used instead of $u_{ij} - b_{ij}$ because the standard deviation of the uncorrelated noise is proportional to the background level. It ensures the processed spectra to have a constant standard deviation over the whole frequency range.

A fit on the background should be fast and capable of fitting all structures with widths larger than the axion lineshape while leaving smaller structures untouched. One way of approaching this challenge is by making use of a filter. Usually filters are applied to smooth noisy data, suppressing small scale noise while retaining large scale structures. In the present case, however, the roles are reversed: The goal is to suppress large scale RF background structures while retaining a potential localised axion signal. Roles of passband and stopband are switched compared to usual filter applications by employing eq. (7.4) instead of proceeding to analyse the filter output b_{ij} .

From the variety of possible filters, again Savitzky-Golay (SG) filters were chosen for this analysis. An SG filter is a moving polynomial fit defined by its two parameters

⁴Avid readers may wonder why raw spectra are used and not calibrated ones that display sensible physical units. Subsec. 7.3.4 motivates this choice.

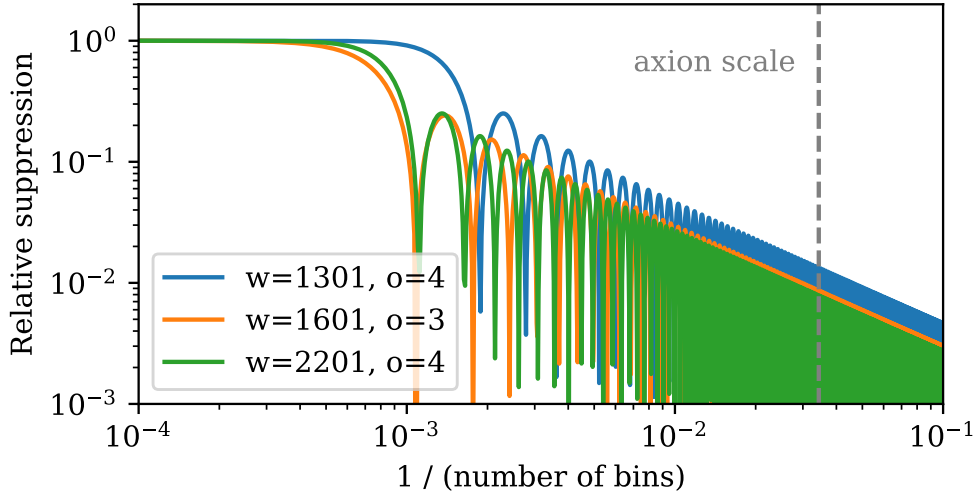


Figure 7.11: Transfer functions of different SG filters. For $\Delta\omega = 849$ Hz the axion FWHM corresponds to roughly 29 bins, indicated as the vertical dashed grey line.

polynomial order o and window size w . As in ch. 6, whenever it is applied the first and last $w/2$ data points are removed from the dataset due to worse fitting at the edges of the dataset.

The effect of SG filters with different parameters o and w are demonstrated in fig. 7.11. With the chosen filter parameters, structures with widths beyond ~ 1000 bins are in the passband of the filters, meaning they will be fit and not retained in the processed spectra. Scales around the axion FWHM of ~ 26 kHz are well in the stopband. Nevertheless, an axion signal is still slightly suppressed by the SG filter. In fact, this effect warrants its own section later (subsec. 7.3.5).

Ideally, in the absence of axions each bin p_{ij} after SG filtering is an independent random sample drawn from a Gaussian distribution with zero mean and a standard deviation

$$\sigma_p = 1/\sqrt{t_{\text{int}}\Delta\omega}, \quad (7.5)$$

that depends only on the integration time t_{int} and resolution bandwidth (RBW) $\Delta\omega$ of the measurement and is equal for every bin of the spectrum. The spectrum is unitless, but the y -axis quantifies the white noise level compared to the background level, so for the simulated spectra $\sigma_p = 0.11\%$ corresponds to $t_{\text{int}} = 924$ s for the chosen RBW $\Delta\omega = 894$ Hz. One example spectrum is shown in fig. 7.12 (blue) together with a magnified version of the signal component present in the spectrum (orange). The shape of the axion signals in the measurement is altered by the background filter as demonstrated in the right panel of fig. 7.12.

7.3.2 Processed Spectra to Combined Spectrum

Ultimately one spectrum should be constructed which at every frequency contains information from all 15 min measurements at that frequency⁵ It is tempting to naively just bin-wise sum over individual processed spectra, however this is clearly not ideal.

⁵In [83] this process is called *vertical stacking*.




Figure 7.12: Type-II example spectrum (blue) and isolated signal component of the spectrum (orange) after processing. The signal component is magnified by $\times 100$ for visibility. The right panel shows a zoom on the lowest frequency signal in the spectrum, displaying alteration in the signal shape due to the effect of the background filter.

An excess in a spectrum bin e.g. with a large external magnetic field B_e carries more information than an excess in a bin that is associated with lower B_e and should therefore get a larger emphasis. The authors of [83] find it optimal to define *maximum likelihood weights* for combining the spectra. Each bin of each spectrum is multiplied by a factor such that a reference axion with its total signal power in this bin would lead to $\mu = 1$ for that bin.

This implies that for such rescaled spectra individual bins have different standard deviations. In the previous subsection it was claimed that in the absence of axions every bin of a processed spectrum could be seen as a random sample of the same Gaussian distribution with zero mean and standard deviation σ_p . This statement is important, because it implies that after multiplying each frequency bin with a different factor one still knows the standard deviation of the distribution it was drawn from, even though only one sample from each of the now differing Gaussian distributions is available. Individual frequency bins get the index j and spectra from the multiple measurements that are to be combined are labelled with index i .

Quantitatively, the j th bin of a rescaled spectrum r_i is obtained from the processed spectrum p_i via

$$r_{ij} = p_{ij} \cdot \frac{b_{ij} P_{ij}^f k_B \Delta\omega}{s_{ij}}, \quad (7.6)$$

where k_B is the Boltzmann constant and P_{ij}^f the j th bin of the power calibration factor. s_{ij} is the strength of a hypothetical single-bin axion signal given by

$$s_{ij} = \frac{q_e}{\hbar} \rho_a \left(\frac{g_{a\gamma}}{m_a} \right)^2 B_{e,i}^2 A \beta_{ij}^2. \quad (7.7)$$

Here q_e is the electron charge, \hbar the reduced Planck constant, ρ_a the local axion energy density, $g_{a\gamma}$ is the axion-photon coupling and m_a the axion mass. Experimental parameters affecting s_{ij} are external magnetic field B_e , which could vary between measurements i , and surface area of the dielectric disks A as well as the boost factors β_{ij}^2 .

In practice, β_{ij}^2 and P_{ij}^f will not change significantly from one 15 min measurement to the next. Therefore it is not necessary to constantly recalibrate and recompute the boost factors. They however clearly vary for different booster states, which warrants the inclusion of the i index. Both are also the only quantities beyond the raw data that realistically display a frequency dependence.⁶

The rescaled spectra are defined so that in the absence of an axion every bin is a random sample from a Gaussian distribution with zero mean and standard deviation $P_{ij}^f k_B \Delta \omega \sigma_{p,i} / s_{ij}$. A single-bin axion would lead to the same standard deviation, but $\mu = 1$ when filter effects from the processing step before are neglected. An example rescaled spectrum can be found in fig. 7.13 (top).

As weights the inverse variance in every bin is used, i.e.

$$w_{ij} = \left(\frac{b_{ij} P_{ij}^f k_B \Delta \omega}{s_{ij}} \sigma_{p,i} \right)^{-2}. \quad (7.8)$$

To normalise the weights defined above one has to divide eq. 7.8 by $\sum_i w_{ij}$. The combined spectrum is now obtained using a weighted sum of the rescaled spectra and dividing every bin by the appropriate standard deviation $\sigma_{c,j}$:

$$c_j = \sigma_{c,j}^{-1} \sum_i r_{ij} \frac{w_{ij}}{\sum_i w_{ij}}, \quad (7.9)$$

where

$$\sigma_{c,j} = \left(\sum_i w_{ij} \right)^{-\frac{1}{2}}. \quad (7.10)$$

This implies that the combined spectrum c_j is normalised to a standard Gaussian distribution and should have zero mean and a standard deviation of one. An axion in a single-bin would not alter the standard deviation, but lead to an expected mean value of $\mu_j = \sigma_{c,j}^{-1}$ when filter effects are neglected. Eq. 7.10 only holds if p_{aj} and p_{bj} are uncorrelated for all j if $a \neq b$. The combined spectrum for the simulated datasets can be found in fig. 7.13 (middle). It is expected to approximately follow a Gaussian white noise distribution which is demonstrated explicitly in fig. 7.15 (top) using the Runs Gaussianity test.

7.3.3 Combined to Grand Spectrum

Unlike for the HAYSTAC analysis of [83], in the present case the bin width of 894 Hz is comparable to the magnitude of the axion FWHM (26 KHz) making their rebinning scheme at this point superfluous. This step of the HAYSTAC analysis was therefore skipped.

However, in the combined spectrum, as defined above, an axion would still lead to a deviation from the mean $\mu = 0$ in multiple adjacent bins. It would be desirable to use this knowledge and condense the full axion signal information into one bin, while still retaining the native frequency resolution. This can be done by cross-correlating the combined spectrum with the axion lineshape to obtain the grand spectrum g .

⁶ $g_{a\gamma}$ and m_a are both proportional to the frequency of the bin, therefore their ratio is independent.

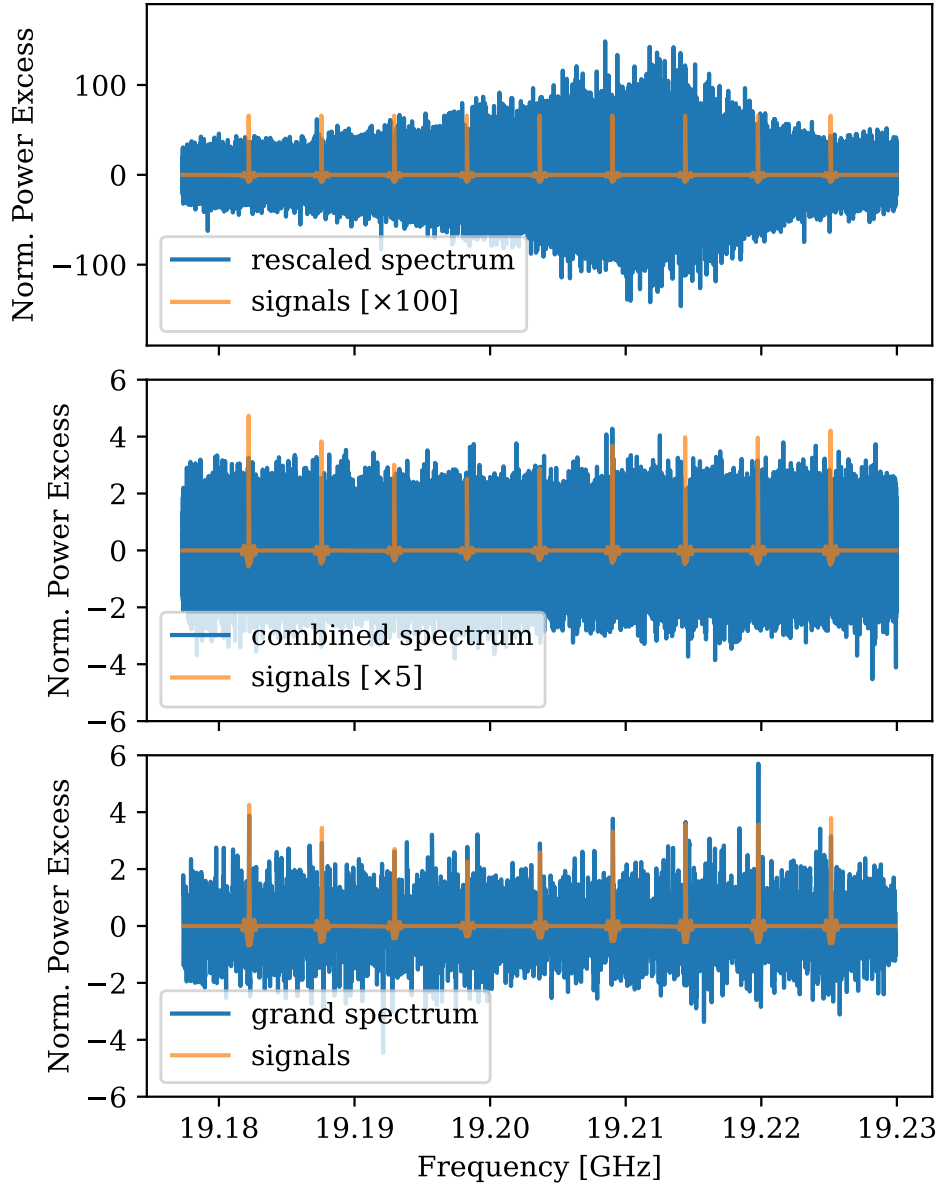


Figure 7.13: Spectra at intermediate steps towards limit setting compared with simulated axion signal components at each stage. **Top:** One example rescaled spectrum of the type-II simulated dataset. Rescaled spectra (blue) are normalised to the amplitude of a simulated axion signal (orange) with reference $\mathcal{C}_{a\gamma}^{\text{ref}} = 2 \times 10^3$. **Middle:** Combined spectrum (blue) including all timestamps of both the type-I and II datasets. The combined spectrum is normalised for all bins to have an equal standard deviation, therefore signals at different frequencies have different amplitudes (orange). **Bottom:** Grand spectrum (blue) obtained by applying an axion lineshape filter to the combined spectrum which leads to an increase of the signal amplitudes compared to the standard deviation of the noise. Note the changing scaling of the signal component in the different panels.

In practice this is implemented as in [83] by splitting the grand spectrum into two components

$$g_j = \frac{D_j}{R_j}, \quad (7.11)$$

which are defined as

$$D_j = \sum_m \frac{c_{j+m}}{\sigma_{c,j+m}} L_m, \quad (7.12)$$

$$R_j = \sqrt{\sum_m \left(\sigma_{c,j+m}^{-1} L_m \right)^2}. \quad (7.13)$$

Here L_m is the m th bin of the theoretical expectation for the axion lineshape defined in terms of the axion CDF $F_b(v)$ from (2.11). Note that the number of bins in g_j is reduced compared to c_j by the number of bins in the lineshape used to cross-correlate.

When the bias-free approach presented in ch. 6 is used, the lineshape L_m simply has to be updated to incorporate the effect from the concrete SG filter used. Effectively, one has to cross-correlate with a lineshape of the form of fig. 6.3 (bottom left) instead of fig. 6.3 (top left) as is usually done. This deviation from the standard pipeline followed in this section has not been found to increase sensitivity for earlier test runs and was therefore not employed here.

The grand spectrum for the simulated datasets can be found in fig. 7.13 (bottom). Since grand as well as combined spectrum normalise every bin to an equal noise level, the amplitude of reference signals in grand or combined spectrum indicate the SNR an axion signal with $C_{a\gamma}^{\text{ref}}$ would have at the respective frequencies. Comparing the middle and bottom panels of fig. 7.13 and considering the different scaling factors for the signals in both panels shows that taking the axion lineshape into account increases the amplitude of the reference signals compared to the combined spectrum by roughly a factor of five.

The reference signals also still display larger amplitudes at the edges of the dataset than at around 19.20 GHz. Fig. 7.14 explains this effect quantitatively. In red and green the expected signal amplitudes are shown for the two individual booster states. One can see that the expected amplitudes are inversely correlated with the total power in each frequency bin for both booster states, shown in fig. 7.10. Because the boost factor for both states was assumed to be constant in frequency, any variation in the SNR can only come from differences in the absolute noise level, which is proportional to the total power.

The SG filter not only affects the power excess from potential signals, but also the noise level, i.e. the standard deviation of the grand spectrum. In fact the grand spectrum shown above has a standard deviation $\sigma_g = 0.93 < 1$. When true white noise would get cross-correlated with the axion lineshape, the standard deviation would still equal unity. The reason why this is not the case here can be found in correlations that are introduced by the SG filter employed to obtain the processed spectra. The effect is explained in detail in [83] and investigated for the present datasets below in subsec. 7.3.5. The expected signal amplitudes in fig. 7.14 have already been rescaled to take the attenuation from the SG filter into account. They agree with the expectation at the signal frequency below the percent level (fig. 7.14, bottom panel). The systematic offset visible could be explained by uncertainty in the calculation of the attenuation from the SG filter.




Figure 7.14: Expected signal amplitude for the simulated dataset. Dashed lines indicate the individual expectations for the two booster states in the simulated dataset (red and green), as well as the expectation for the combined dataset (blue). The latter coincides perfectly with the amplitudes of the simulated signals in the grand spectrum (obtained by subtracting the grand spectra with and without signals from each other). The bottom panel shows the difference between signal amplitudes and expectation at the frequency of the signal with deviations below the percent level.

7.3.4 On Calibrating Spectra

The procedure above employed raw, uncalibrated spectra that are SG filtered to obtain processed spectra and then combined and rescaled taking the calibration into account. Arguably it would be more straightforward to calibrate the spectra first and work with physically meaningful units in the rest of the analysis pipeline. This subsection shows that both methods are equivalent in the absence of errors and why the former was employed. In the following only a single power calibration is assumed for all spectra i , i.e. the spectra are part of the same dataset. Then compared to the previous definition the index i can be dropped from the power calibration factor, so $P_{ij}^f \rightarrow P_j^f$.

Define hypothetical, true values for the power calibration factor \tilde{P}_j^f , which would be obtained after an infinite measurement time. Similarly to sec. 6.1 assume that this ideal power calibration factor can be perfectly fit using the SG filter f_{bg} , i.e.

$$f_{\text{bg}}\left(\tilde{P}_j^f\right) = \tilde{P}_j^f. \quad (7.14)$$

Due to the linearity of SG filters, it is then trivial to show that p_{ij} is the same no

matter if uncalibrated or calibrated spectra are used:

$$p_{ij} = \frac{u_{ij}^{\text{cal}}}{b_{ij}^{\text{cal}}} - 1 = \frac{u_{ij} \tilde{P}_j^f}{f_{\text{bg}}(u_{ij} \tilde{P}_j^f)} - 1 = \frac{u_{ij} \tilde{P}_j^f}{f_{\text{bg}}(u_{ij}) f_{\text{bg}}(\tilde{P}_j^f)} - 1 = \frac{u_{ij}}{b_{ij}} - 1, \quad (7.15)$$

where again $b_{ij} = f_{\text{bg}}(u_{ij})$ and the superscript “cal” denotes calibrated spectra in contrast to raw spectra without superscript. r_{ij} and w_{ij} furthermore depend only on the product $b_{ij} \tilde{P}_j^f$, which can similarly be shown to equal b_{ij}^{cal} . If eq. (7.14) holds, calibrating the raw spectra before any other processing step is therefore mathematically equivalent to using uncalibrated spectra to obtain p_{ij} and using $b_{ij} \tilde{P}_j^f$ in the calculation of c_{ij} .

In practice, however, only an imperfect estimate of \tilde{P}_j^f is available:

$$P_j^f = \tilde{P}_j^f + \epsilon_j, \quad (7.16)$$

$$f_{\text{bg}}(P_j^f) = \tilde{P}_j^f + \epsilon'_j, \quad (7.17)$$

where $|\epsilon_j| \ll \tilde{P}_j^f$ denotes a small error and $\epsilon'_j \equiv f_{\text{bg}}(\epsilon_j) \neq \epsilon_j$. P_j^f is obtained by applying an SG filter, here with width $w = 8001$ data points and polynomial order $o = 3$, the values of ϵ_j and ϵ'_j therefore depend on the parameters of the SG filter. With these deviations from the true power calibration factor \tilde{P}_j^f using calibrated spectra in the calculation of the processed spectra leads to an imperfect estimate:

$$\begin{aligned} \frac{u_{ij}^{\text{cal}}}{b_{ij}^{\text{cal}}} - 1 &= \frac{u_{ij} (\tilde{P}_j^f + \epsilon_j)}{f_{\text{bg}}(u_{ij}) (\tilde{P}_j^f + \epsilon'_j)} - 1 = \frac{u_{ij}}{b_{ij}} \left(\frac{1 + \epsilon_j / \tilde{P}_j^f}{1 + \epsilon'_j / \tilde{P}_j^f} \right) - 1 = \\ &= \frac{u_{ij}}{b_{ij}} - 1 + \frac{u_{ij}}{b_{ij}} \frac{\epsilon_j}{\tilde{P}_j^f} - \frac{u_{ij}}{b_{ij}} \frac{\epsilon'_j}{\tilde{P}_j^f} = p_{ij} + \mathcal{O} \left(\frac{\epsilon_j}{\tilde{P}_j^f}, \frac{\epsilon'_j}{\tilde{P}_j^f} \right), \end{aligned} \quad (7.18)$$

where the third equality uses a Taylor series of the fraction in brackets. Constructing the processed spectra using calibrated data leads to offsets of $\mathcal{O}(\epsilon_j / \tilde{P}_j^f, \epsilon'_j / \tilde{P}_j^f)$. Using uncalibrated spectra on the other hand leads only to multiplicative errors in r_{ij} and w_{ij} via

$$P_j^f = \tilde{P}_j^f + \epsilon_j = \tilde{P}_j^f \left[1 + \mathcal{O} \left(\frac{\epsilon_j}{\tilde{P}_j^f} \right) \right]. \quad (7.19)$$

In practice one power calibration will be used for many different spectra i . When summing over all spectra for which the same power calibration with P_j^f is used, eq. 7.18 becomes

$$\begin{aligned} c_j &\propto \sum_{i=1}^n \left(\frac{u_{ij}^{\text{cal}}}{b_{ij}^{\text{cal}}} - 1 \right) = \sum_{i=1}^n \left[p_{ij} + \mathcal{O} \left(\frac{\epsilon_j}{\tilde{P}_j^f}, \frac{\epsilon'_j}{\tilde{P}_j^f} \right) \right] = \\ &= n \times \mathcal{O} \left(\frac{\epsilon_j}{\tilde{P}_j^f}, \frac{\epsilon'_j}{\tilde{P}_j^f} \right) + \sum_{i=1}^n p_{ij}. \end{aligned} \quad (7.20)$$

Since p_{ij} should be uncorrelated between individual spectra i , this implies an increasing error relative to the standard deviation of $\sum_{i=1}^n p_{ij}$ as n increases, whereas the same does not happen for the factor giving the error incurred when calibrating after processing in (7.19). This effect can clearly be seen in fig. 7.15 for the example spectra also used in the previous subsections. The top row shows the method employed here leading to a combined spectrum displaying Gaussian white noise behaviour. In the bottom row, however, which shows spectra that were calibrated before processing, significant deviations from white noise can be seen. The Runs Gaussianity test (bottom right) introduced in sec. 7.1 shows a clear discrepancy between the combined spectrum (blue) and the white noise reference (orange). The deviation is so big that it can even be seen by eye when zooming into a small region of the spectrum (inset in bottom left plot).

In conclusion, using uncalibrated spectra to obtain the processed spectra is a viable method. The calibration then enters later when the spectra are combined and normalized power excesses are converted into a limit on $g_{a\gamma}$. Using uncalibrated spectra avoids deviations from $\mu_j = 0$ in the combined spectrum, which otherwise grow with increasing number of spectra for which the same power calibration is used, i.e. increasing measurement time. It can easily be shown that both methods produce equal processed spectra in the absence of errors on the power calibration factor. A typical Y-factor calibration as used for MADMAX in addition to the power calibration factor allows for an additional additive component called power calibration offset. The considerations above only hold as long as this component is negligible.

7.3.5 SG filter induced SNR attenuation

When no physically meaningful background model is available, fitting the data to obtain processed spectra as done above will lead to some attenuation on the signal-to-noise ratio (SNR). SG filtering induces some suppression on all scales. Fluctuations with big widths deep in the passband of the filter are efficiently suppressed when eq. (7.4) is applied. But also fluctuations with widths of only a few bins or even the bin-to-bin fluctuations of a white noise component incur some suppression as the transfer function of the possible filters is non-zero for all scales (see fig. 7.11).

To investigate the SNR attenuation induced by the SG filter again an approach similar to what has been presented in [83] is used. Mock datasets are constructed similarly as above in the current section, obtaining again realistic backgrounds from raw data of the type-I dataset smoothed with an SG filter of polynomial order $o = 3$ and a window size $w = 6001$ bins (i.e. ~ 5.4 MHz). To these backgrounds Gaussian white noise is added, again for each frequency point with a standard deviation of 0.11% of the background value at this frequency. Every spectrum additionally contains one signal, which is placed at a random but known location within the central 80th percentile of the dataset to avoid possible edge effects. Signal parameters are tuned to yield a significant, but not unrealistic, excess for the grand spectrum at the signal location.

The analysis pipeline is then applied to the spectrum multiple times: First the known background is subtracted from the spectrum representing the case of an ideal filter that does not induce any attenuation. Then SG filters with various window sizes are used to obtain processed spectra. Since only one spectrum is used in each case the

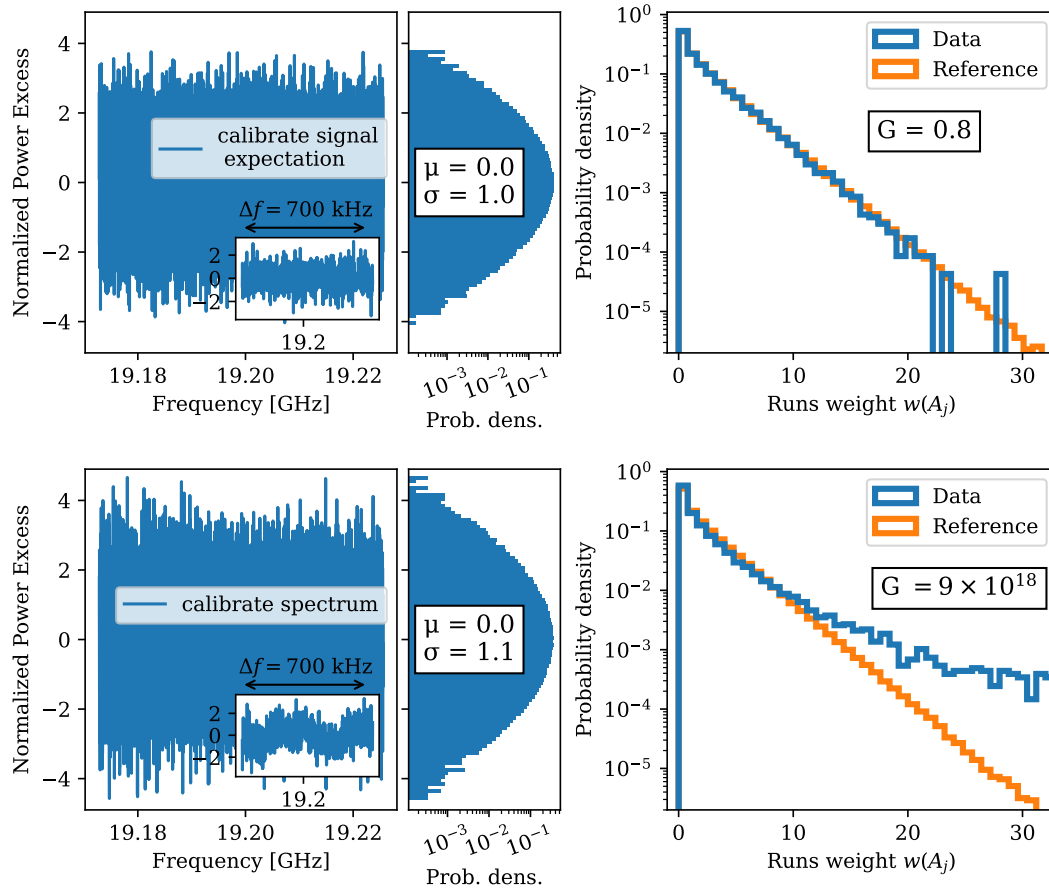


Figure 7.15: Gaussianity of the combined spectrum when accounting for the power calibration before or after constructing the processed spectra. **Top:** Combined spectrum using the method presented here, where uncalibrated spectra are processed and information on calibration enters via the signal expectation. **Bottom:** Combined spectrum when calibrating the spectra before processing. The resulting spectra are shown on the left with a zoom into an arbitrary region with a width of ~ 700 kHz in the inset. The central panels show histograms of the combined spectra giving means and standard deviations. The right panels show the Runs Gaussianity test described above including the parameter G , clearly displaying deviations from Gaussian white noise for the second method (bottom right).




Figure 7.16: Result of the SG attenuation simulation conducted on a type-I simulated dataset using an SG filter with polynomial order $o = 4$ and a window size of $w = 1201$ data points. Both mean and standard deviation of the resulting grand spectrum are reduced by applying the SG filter (orange) compared to the ideal case where the known background is subtracted (grey). This leads to a reduction in SNR governed by eq. (7.21).

combined spectra are equivalent to the processed spectra. After constructing grand spectra the grand spectrum value at the signal location is saved for all scenarios. This process is repeated 10^4 times. A histogram of the resulting excesses for the ideal scenario and a filter width of $w = 1101$ can be seen in fig. 7.16.

In the ideal, effectively background-less scenario the resulting excesses follow a Gaussian distribution with mean $\mu_{\text{id}} = 5.15$ and standard deviation $\sigma_{\text{id}} = 1$. The latter comes from the standard deviation of the grand spectra, the former corresponds to the parameters of the simulated signals, which are chosen to yield this significant excess. When instead the SG filter is employed, the mean of the observed signal strength decreases down to $\mu_{\text{f}} = 4.39$. This can be attributed to the axion lineshape filter used in the construction of the grand spectrum not taking into account signal shape alteration induced by the SG filter. However, not only the mean of the distribution is reduced. The standard deviation of the distribution of the excesses is also smaller than σ_{id} with $\sigma_{\text{f}} = 0.93$. This is due to adjacent points in the processed spectrum becoming correlated when applying SG filters. When the spectrum is then cross-correlated with the axion lineshape to obtain the grand spectrum, which also correlates adjacent bins, the effect of the SG filter becomes apparent and leads to a reduction in the grand spectrum standard deviation.

The SG filter not only affects the signal strength, i.e. the mean of the distribution, but also the noise level, i.e. its standard deviation. The degradation of the SNR due to the SG filter is therefore given by

$$\eta_{\text{SG}} = \frac{\mu_{\text{f}}/\sigma_{\text{f}}}{\mu_{\text{id}}/\sigma_{\text{id}}} . \quad (7.21)$$

As expected, larger window sizes lead to less SG filter induced attenuation. In practice it is therefore desirable to choose the largest possible window size which does not lead

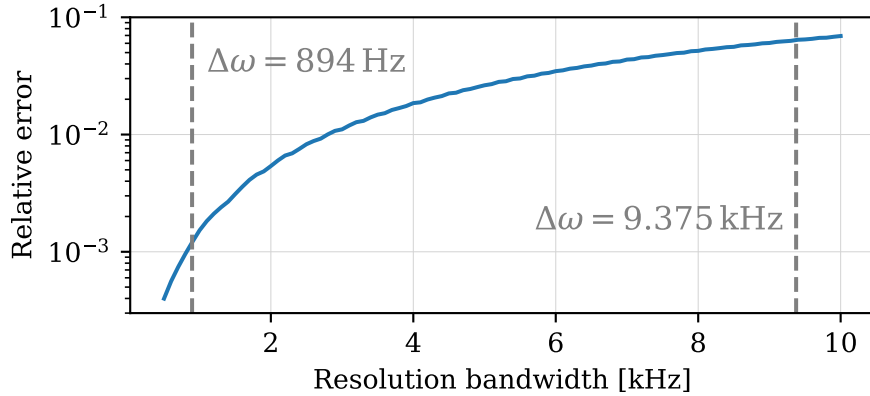


Figure 7.17: Effect of signal misalignment with respect to edges of the frequency bins. Relevant errors only occur for RBWs beyond ~ 3 kHz for an axion signal at ~ 19 GHz. The dashed vertical lines indicate $\Delta\omega = 894$ Hz, the RBW used in the analysis, and $\Delta\omega = 9.375$ kHz, the RBW used in [5].

to residuals of non-axionic RF structures in the grand spectrum. This will be covered in more detail in subsec. 8.2.2. The effect of varying o was not investigated because it is approximately degenerate with w as fig. 7.11 shows.

For the bias-free approach from ch. 6 no signal attenuation is expected, because it is accounted for by the changing lineshape. However it is likely that also the standard deviation of the grand spectrum will not experience attenuation. On the contrary, RF structures that are difficult to remove using the standard approach could increase the standard deviation of the grand spectrum when the effect of the SG filter on the axion lineshape is accounted for, i.e. when the bias-free approach is applied.

7.3.6 Attenuation from misalignment of the axion signal

The SG filter thus attenuates the SNR by correlating nearby bins and altering the expected axion lineshape. Another possible effect impacting the expectation for the axion lineshape is misalignment of the exact frequency of the axion signal with respect to the lowest or highest frequency, also called edges, of the bins.

The axion signal shape in the grand spectrum differs slightly depending on where the frequency corresponding to the axion mass falls within one frequency bin. This effect increases with increasing RBW and does not pose a problem here, as fig. 7.17 shows. For the larger $\Delta\omega$ used in the analysis of [5], however, it had to be considered as a systematic error.

To investigate this effect, datasets with 2×10^4 data points with various different $\Delta\omega$ are constructed containing only a single signal component at the central frequencies and a constant baseline. The signal component of each spectrum is simulated with an axion mass corresponding to a frequency that is offset from the lower edge of one of the bins in equidistant increments between 0 and $\Delta\omega$. These spectra are then processed using a fourth-order SG filter with a filter window width of 1.07 MHz, corresponding to $w = 1201$ data points for the datasets used above and in ch. 8. Every resulting spectrum is then cross-correlated with the axion lineshape in the zero offset case to

obtain mock grand spectra. The maximal excesses in these spectra depends on the offset introduced.

Fig. 7.17 shows the maximal relative difference in grand spectrum excess between any of the offsets and the zero offset case. For the analysis presented in ch. 8 the effect from the misalignment is at the permille level and therefore negligible (left dashed line). The relative error surpasses 1% only for $\Delta\omega \gtrsim 3$ kHz, making it largely irrelevant for RBWs below this level given an axion mass corresponding to ~ 19 GHz. Dark matter searches at these frequencies that use RBWs above 3 kHz may have to take misalignment of the signal with respect to the bins into account as a source of systematic uncertainty, as is done in [5].

7.4 Limit Setting

The procedure above provides all quantities necessary for limit setting. The grand spectrum g_n serves as a test statistic following a known Normal distribution with $\mu_{g_n} = 0$ and $\sigma_g \lesssim 1$ for every frequency bin n that does not contain an axion signal. R_n gauges the values of g_n to the physically meaningful scale of $g_{a\gamma}$ by providing an expectation for μ_{g_n} for a reference axion-photon coupling. SG filter induced correlations lead to a small attenuation in SNR induced by the procedure of the statistical analysis itself. This effect is accounted for by the parameter η_{SG} .

As described in sec. 4.1, a limit is constructed by finding which hypotheses are incompatible with the data. When sampling at random from a Normal distribution with mean μ and standard deviation $\sigma = 1$, the probability of obtaining a value below $\mu - 1.645$ is smaller than 5%. Consequently, locally an observed grand spectrum value of g_n rules out photon couplings corresponding to values larger than $g_n/\sigma_f + 1.645$ at 95% confidence level, where σ_f is the expected grand spectrum standard deviation (see subsec. 7.3.5).

To obtain a 95% confidence limit on the axion-photon coupling one can therefore employ

$$|g_{a\gamma}|_n = |g_{a\gamma}^{\text{ref}}| \sqrt{\frac{g_n/\sigma_f + 1.645}{\eta_{\text{SG}} R_n}} \quad (7.22)$$

where $|g_{a\gamma}^{\text{ref}}|$ must equal the photon coupling used in (7.7). The resulting limit for the example used in this section can be found in fig. 7.18. The median expectation, corresponding to $g_n/\sigma_f = 0$ in (7.22), as well as the boundaries $\pm 1\sigma$, corresponding to $g_n/\sigma_f = \pm 1$, of the limit are shown as dashed and dotted lines, respectively.

When the limit is constructed using either (7.22) or similar methods it may become necessary to slightly adapt the limit for visualisation purposes. The limit shown in fig. 7.18 contains $\sim 6 \times 10^4$ data points, which cannot be individually resolved in the plot. As also demonstrated by the inset, the area at the bottom of the limit with higher colour intensity is actually the limit fluctuating on scales that cannot be resolved in the plot due to the finite width of the lines. By construction, the limit only reaches the dashed line which marks data points exceeding the median sensitivity by 1σ for 16% of the data points. Nevertheless when the full limit is plotted it appears as if the 1σ line is excluded for all data points. This problem persists similarly even if the median expected limit is exceeded by more than one sigma.




Figure 7.18: Simulated limit obtained by applying the analysis pipeline to the mock dataset described in sec. 7.3.

One solution to this visualisation problem is to smooth the limit before plotting, employed e.g. by the authors of [114]. The analysis presented here employs a different approach, showing only limits that do not exceed the median expectation by more than 1σ . This approach is conservative, since it ensures that at every frequency the calculated limit corresponds to a confidence level of at least 95%.

7.5 Grand Spectrum Excesses

So far it was implicitly assumed that no significant excess is observed in the grand spectrum. But what if they are? The purpose of this section is to investigate the case in which an excess in the grand spectrum can be considered significant, collect possible scenarios for excesses and provide courses of action that can be considered.

7.5.1 Significance

Before this can be done a few words are in order on when an excess is significant. Sec. 4.2 showed that a local 5σ excess in the grand spectrum does not correspond to the naively expected global false positive rate of 2.9×10^{-7} and therefore for spectra containing many independent points is not necessarily significant enough to warrant a detection claim. Note that only being interested in positive excesses, the numerical value 2.9×10^{-7} is half of the probability given in sec. 4.2 to obtain an outlier beyond $\pm 5\sigma$.

For the 6×10^4 data points used above a local 5σ excess would have a global significance of only 2.1σ corresponding to a false-positive rate of 1.7% according to eq. (4.1). This calculation, however, ignores correlations between frequency bins, which are introduced by the cross-correlation with the axion lineshape. The correction

factor introduced in eq. (4.2) was calculated for a set of 100 simulated spectra using the same frequencies as the example above ($\Delta\omega = 894$ Hz). The spectra consisted of white noise that was cross-correlated with the axion lineshape. Taking the mean over all 100 correction factors yields an effective number of samples

$$n_{\text{eff}} \approx \frac{n}{47}, \quad (7.23)$$

with $n = 6 \times 10^4$ being the number of correlated data points. This factor crucially depends on the RBW and has a small dependency on frequency due to the increasing width of the axion lineshape with increasing frequency. The latter effect however can safely be ignored for $f = 19 \pm 1$ GHz. Using this new effective number of samples results for a local 5σ excess in a corrected global significance of 3.4σ or false positive rate of 3.7×10^{-4} .

7.5.2 Handling Grand Spectrum Excesses

How could a globally significant excess be distinguished from or shown to be an axion signal?

Firstly, a generic excess in the grand spectrum may be possible to distinguish from an axion signal using the width of the excess. An axion signal in the grand spectrum should have a FWHM similar to the FWHM of the theoretically expected axion lineshape from sec. 2.4, i.e. ~ 26 kHz at a frequency of 19 GHz. Excesses which are significantly more narrow should have already led to even larger excesses in the combined spectrum and can therefore be expected to be caught at this stage of the analysis. A more common problem are excesses which are significantly broader than an expected axion signal. These hint at the presence of RF structures with characteristic scales above the axion FWHM which have not been effectively removed by the SG filter employed in the background removal. The frequency of such a structure coinciding with a sharp peak in the background provides evidence for this hypothesis. Another hint are small frequency shifts between individual datasets, since RF peaks may shift in frequency between different booster configurations while an axion signal would not.

If such broader excesses are present, a smaller SG filter window width or bigger polynomial order can be chosen for the background removal in order to move the pass-band of the filter closer to the axion scale (compare fig. 7.11). A single small axion signal should not significantly affect the standard deviation of the grand spectrum. The same does not need to be true for RF structures, which may lead to multiple, broader excesses. Whether or not a specific set of SG filter parameters may be used for the analysis can therefore be determined by comparing the standard deviation of the grand spectrum σ_g with σ_f from (7.21) which by construction does not contain an effect from residual RF structures. While $\sigma_g > \sigma_f$ can always be expected due to the imperfect filter stop-band, deviations beyond

$$\frac{\sigma_g}{\sigma_f} - 1 \gtrsim 0.02 \quad (7.24)$$

will be deemed unacceptable in ch. 8.

Secondly, a generic excess can differ from an axion signal in various other properties. Potential time variability of an axion signal, either in signal strength, signal width or frequency, is tightly constrained [85]. An excess displaying any kind of variability

can most likely be ruled out as being of axionic origin. Most notably, however, a potential axion signal should scale with B_e^2 . It is therefore desirable to complement any measurement inside a magnetic field with a significant measurement time in the same environment but the magnet turned off or at different values of B_e . If such data are available, an excess can be shown to not be due to an axion signal by calculating the Pearson correlation coefficient between grand spectra at times with and without the external magnetic field.

Lastly, as is also done in most mature haloscope experiments, e.g. [83, 114, 120], a rescan strategy is in order before claiming an axion detection. Due to the narrow-band, scanning approach employed by most haloscope experiments, it is not very time consuming to significantly increase statistics in a frequency region where a persistent excess exists. The development of such a strategy for the MADMAX experiment will be necessary in the long run.

7.6 Summary & Discussion

This chapter provides all necessary ingredients to conduct statistical analysis for the first axion dark matter search with a MADMAX prototype.

Due to the importance of frequency bins in processed and combined spectra being uncorrelated samples from a Gaussian distribution, a new Gaussianity test was introduced. For a normally distributed set of data points with zero mean, taking the squared sum over consecutive positive or negative data points leads to a set of values w following a known distribution. The Runs Gaussianity test leverages this w -distribution in order to investigate the Gaussianity of a dataset. A new metric G based on this test was added to quantify the results. $G \gg 1$ usually indicates correlations between adjacent bins, but there is a non-zero probability of obtaining $G \gg 1$ even for uncorrelated white noise. Future work could therefore focus on improving this metric with respect to its predictability. Results from the Runs Gaussianity test have been shown to be stable for datasets containing at least $N = 10^3$ data points and for a metric G constructed out of at least $K = 50$ ranges. Meeting these criteria, the Runs Gaussianity test displays a better sensitivity than the standard comparison of a histogram of the data with a Normal distribution for a set of example applications.

With the digital notch filter a method was added to remove sinusoidal background components which may be difficult to remove using standard SG filters. It is defined algorithmically by applying a Fourier transform to the dataset, replacing localised peaks in both complex components of the Fourier spectrum with a set of dummy data points and conducting a back-transformation afterwards. The method is effective at removing sinusoidal background components from the raw frequency spectra while only having permille-level impact on potential axion signals.

The procedure to obtain a grand spectrum that takes the expected axion lineshape into account from a set of raw, uncalibrated frequency spectra was described using a set of simulated spectra as an example. It was shown that processed spectra need to be constructed from uncalibrated data since the short measurement time of power calibration measurements would otherwise induce errors in the combined spectrum. Tracing a set of simulated signals through the pipeline reveals the effect of combining processed spectra on the signals as well as a five-fold increase in sensitivity of the grand spectrum compared to the combined spectrum. The simulated signals can

also be used to validate the interpretation of the parameter R_j (see eq. (7.13)) as the expected signal amplitude in frequency bin j . Signal amplitudes are subject to a systematic attenuation due to the effect of the SG filter used. The standard approach to estimate this attenuation is by using a Monte Carlo simulation, which was outlined here and reproduced for one set of SG filter parameters. Another systematic error could in principle arise from misalignment of the axion signal with respect to the frequency bins. This effect, however, needs to be considered only for larger resolution bandwidths of several kHz and is therefore not relevant for this analysis.

Alternatively to the construction of the grand spectrum, Bayesian inference methods could be investigated by future work. They are potentially more informative, especially for wider or narrower axion lineshapes. Such differences can be handled by appropriate prior choices. No matter if Bayesian or frequentist methods are used, the signal model should be altered to account for the effect of the SG filter as demonstrated in ch. 6. This has the potential to slightly increase the sensitivity of the procedure, even though the exact effect needs to be investigated for every dataset individually.

If the grand spectrum does not display statistically significant excesses, a limit can be constructed. There are multiple possible ways to obtain a confidence limit in the context of axion haloscopes. One of them is to find for every individual frequency bin an axion-photon coupling large enough to be incompatible with the measured value. For datasets containing many bins visualising such a limit on a logarithmic $g_{a\gamma}$ axis poses some difficulty, which can be alleviated by conservatively rejecting limits exceeding the median expectation significantly.

If the spectrum does display a significant excess, it is crucial to properly calculate its statistical significance taking the number of spectral bins and their correlations into account. Strategies to verify the axionic origin of an excess include investigating the width of the excess in the grand spectrum or its scaling with B_e . Ultimately, however, no detection of axion dark matter should be claimed before extensive rescans are conducted. A concrete strategy for handling excesses is essential and should be put in place before the next MADMAX physics measurements.

Some of the methods presented in this chapter have already been employed for constraining dark photon dark matter in [5].

8

First Results from a MADMAX Axion DM Search

For the MADMAX experiment a dipole magnet with a large open bore is needed. One magnet fulfilling these requirements is the MORPURGO magnet located in the H8 SPS beam-line at the CERN north area. Generous support was given to the MADMAX collaboration in using this magnet over several periods of general CERN shut-downs. During the latest of these in early 2024 physics data was taken that is analysed in the present chapter. This data constitutes the first axion dark matter search by the MADMAX collaboration.

Equipped with the methods introduced in the previous chapter, the data can be used to construct world-leading limits on dark matter axions in regions around axion masses of $76.56\text{-}76.82\ \mu\text{eV}$ and $79.31\text{-}79.53\ \mu\text{eV}$. Crucially important for the limit setting is the determination procedure of the boost factor β^2 , that is outlined in subsec. 3.3.4 and subject to future publication.

In order to construct a limit it is furthermore necessary to ensure Gaussianity of processed and combined spectra. As already hinted at, this requires the removal of several related sinusoidal background components using the digital notch filter (DNF) introduced in the previous chapter.

The second filter used in the processing of the spectra for removal of background components is a Savitzky-Golay (SG) filter. This filter has two approximately degenerate free parameters, which determine the stop-band and pass-band of the filter. The choice of these parameters hinges on the necessity to remove all background structures from the spectra in the dataset, while keeping attenuation of the axion signal to a minimum.

Before a limit can be constructed, systematic uncertainties in system temperature and boost factor need to be accounted for. Their origin has been outlined already in




Figure 8.1: Time evolution of the available magnetic field strength. Shaded regions indicate measurement times for individual datasets considered in the analysis, which are labelled by their respective boost factor maxima.

sec. 3.3, but they now need to be properly combined and translated into a correction of the limit.

The chapter starts with an overview over the data taking and the datasets obtained in sec. 8.1. Sec. 8.2 then provides details on all necessary data processing steps which have not yet been covered in the previous chapter. This includes the choice of exact digital notch and Savitzky-Golay filter parameters, proving the Gaussian white noise behaviour of the processed and combined spectra for the aforementioned axion search measurements as well as the handling of systematic uncertainties. The final axion/axion-like particle dark matter limits are presented in sec. 8.3 before sec. 8.4 concludes the chapter.

8.1 Details of the Data Taking

The MADMAX CB200 setup as described in sec. 3.3 was used for measurements inside the MORPURGO magnet at CERN between 21st of February 2024 and 12th of March 2024. The setup was located at 46.26° N / 6.06° E, with its optical axis oriented horizontally at an azimuth angle of 42° relative to north. Fig. 8.1 details the magnetic field provided by MORPURGO during the measurement campaign. At first the target field strength of 1.58 T was reached, after several problems were encountered with the cryogenic system of the magnet this however was reduced to 1.0 T.

Two significant time periods could not be used in the analysis: On February, 22nd, a test measurement was started using a different receiver chain. Additionally, the booster peak in the spectrum of the measurement started on February, 29th, displayed a significant variation over time. This could be the result of a ferromagnetic screw left on the booster, leading to small distance variations between mirror and first disk after the magnet was turned on. This data could theoretically be analysed, but

starting date	$f(\max \beta^2)$ [GHz]	measurement time [h]	median B field [T]
Feb, 21st	18.531	13.5	1.58
Mar, 02nd	18.543	73.75	1.01
Mar, 05th	18.557	159.0	1.01
Feb, 27th	19.196	40.75	1.27
Feb, 23rd	19.215	61.25	1.58

Table 8.1: Summary of the datasets used for limit setting.

would require many individual boost factor estimations, making it unreasonably time consuming with little added information.

In total five different datasets are used in the analysis described here, corresponding to two configurations, i.e. sets of spacings between the disks, with boost factor peaks around 18.55 GHz and 19.20 GHz, in the following called configuration 1 and 2, respectively. They are summarised in tab. 8.1. The dielectric disks inside the booster are kept at distances (from the taper-adjacent disk to the mirror) of 12.52 mm, 12.25 mm and 8.38 mm in the former case, and 11.89 mm, 12.25 mm and 8.02 mm in the latter (compare fig. 3.5).

These two configurations were further fine tuned by applying pressure to the mirror with a metal rod changing the distance between mirror and first disk by several μm . The altered booster state shifts the booster peak, i.e. the structure in the raw power spectrum associated with a large boost factor, by $\mathcal{O}(10 \text{ MHz})$. The five datasets consist of a total of 1393 spectra with a resolution bandwidth of 8.94 Hz, each containing data integrated over 15 minutes of measurement time. These spectra look almost identical for a given dataset, with T_{sys} variations of $< 10 \text{ K}$ due to the gain drift mentioned in subsec. 3.3.3 and the frequency of the booster peak shifting by $\mathcal{O}(0.1 \text{ MHz})$. Fig. 8.2 therefore shows one representative spectrum for each dataset. Variations between individual spectra are accounted for as systematic uncertainties.

It is clearly visible that the pressure applied to the mirror leads to a small change in its position which causes relevant frequency shifts of $\mathcal{O}(10 \text{ MHz})$ in the boost peak. At the boost peak a strong electric field is expected between mirror and first disk, therefore it is more sensitive to changes in mirror position than other spectral features. In fact, as mentioned in subsec. 3.3.4, most of these other features can not be modelled by the 1D procedure used to obtain the boost factor, leading to less reliable estimates at frequencies close to or containing features due to 3D effects. This issue is accounted for by limiting the analysis to the frequency region around the boost peak that has been identified to be free of such features in preparation of the data-taking. The frequency region of the analysis is indicated as the shaded grey area in fig. 8.2.

Boost factors are obtained using the Monte-Carlo procedure detailed in subsec. 3.3.4. They are shown in fig. 8.3 for the five datasets with mean boost factor values in darker and 1σ uncertainty intervals in lighter colours. In total 100 samples are generated to estimate the uncertainties. The samples approximately follow Normal distributions as the histograms of boost factor samples at the β^2 peak for all datasets show in fig. 8.4.

Additionally, the spectra were rebinned to $\Delta\omega = 894 \text{ Hz}$ before processing by taking the average of 100 adjacent bins. The new $\Delta\omega$ leads to more reasonably sized data-




Figure 8.2: Example spectra for the five datasets used in the analysis. Data regions used for limit setting are shaded in grey. A clear shift is visible in the boost peaks for the individual booster states, while other spectral features remain relatively unchanged.

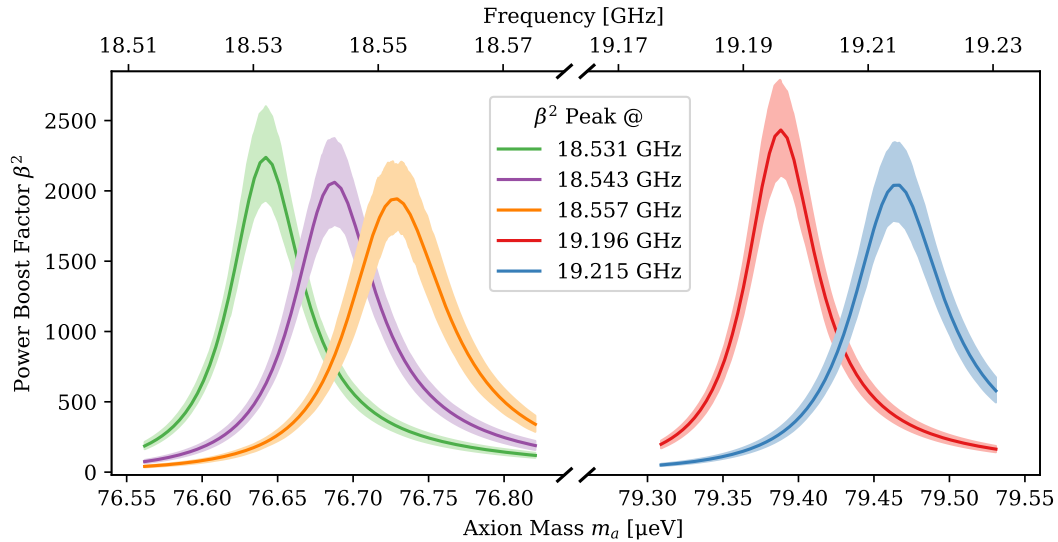


Figure 8.3: Boost factor distributions used in the analysis for the five datasets. Lines denote mean β^2 and shaded regions give $\pm 1\sigma$ intervals.




Figure 8.4: Distribution of boost factor samples at the maximum of the boost factor peaks. For all five datasets the distribution is approximately Gaussian.

sets while still being significantly smaller than the expected axion FWHM of 26 kHz. Therefore no reduction in sensitivity is expected from this rebinning.

The datasets started at Feb, 21st, Feb, 23rd and Feb, 27th with peak β^2 at 18.531 GHz, 19.215 GHz and 19.196 GHz, respectively, have dedicated power calibration measurements available that were taken immediately before the data taking. For both the Mar, 2nd and Mar, 5th measurements with peak β^2 at 18.543 GHz and 18.557 GHz the same power calibration from Feb, 29th was used to avoid having to ramp down the magnet. This is possible, since no changes on the setup occurred between Feb, 29th and Mar, 12th except for adjusting the pressure on the mirror, which could be performed with the magnet on.

8.2 Data Processing

8.2.1 Digital Notch Filtering

In sec. 7.2 a method called digital notch filtering was introduced, which was shown to be effective at removing sinusoidal backgrounds in frequency spectra by replacing the corresponding bins in the Fourier transformation. From fig. 8.5 it can be gleaned that such a background is present in the 2024 data. The Fourier transform of a single spectrum from the dataset with its β^2 peak at 19.215 GHz displays two clear peaks at $\Delta t = 109.22 \mu\text{s}$ and twice that value.¹ Three more peaks at multiples of $109.22 \mu\text{s}$ exist in the data, but are not visible in fig. 8.5 due to the small integration time of the single spectrum. The significant increase of the Fourier spectrum at small Δt and the statistical fluctuations at larger Δt correspond to the correlated background and uncorrelated noise components, respectively, and need to be impacted as little as possible by the digital notch filter.

The Fourier peaks are removed as the first step of the data processing, even before processed spectra are calculated. Five regions around $n \times 109.22 \mu\text{s}$ with $n \in [1, 5]$

¹Note that the frequency spectrum transformed here is real-valued, meaning that phase information is lost and the units of the Fourier transform do not correspond to physical time.




Figure 8.5: Absolute value of the Fourier transform of a single spectrum of the dataset with peak β^2 at 19.215 GHz before and after applying the digital notch filter. The spectrum before filtering (blue) coincides perfectly with the spectrum after filtering (orange) except for two clearly visible peaks at multiples of $109.22 \mu\text{s}$. The inset explicitly shows the replacement of the peak at lower Δt in a window of $0.1 \mu\text{s}$.

are removed with widths of $0.1 \mu\text{s}$ each using the replacement method introduced in sec. 7.2. Fig. 8.8 shows that removing five regions is indeed necessary.

Another perspective on the sinusoidal background component is provided by investigating correlations between individual spectra. The procedure introduced in sec. 7.3 assumes that Gaussian spectra which are stacked to yield the combined spectrum are uncorrelated. This implies that when averaging over a certain number N of spectra with equal standard deviation σ , the resulting spectrum has a standard deviation of σ/\sqrt{N} . Using

$$\sigma_p^2 = (t_{\text{int}}\Delta\omega)^{-1}, \quad (8.1)$$

with integration time t_{int} , for a single processed spectrum one can predict the relationship between t_{int} and σ^{-2} for an experiment with zero dead-time, i.e. t_{int} equals the actual time of the measurement. In other words, if the receiver chain is calibrated and σ_p is known, the integration time can be used for validation. The t_{int} expectation is shown as the dashed black line in fig. 8.6. An experiment with significant dead time would still display a linear relationship, albeit with smaller slope. Positive curvature as displayed by the dataset with peak β^2 at 18.557 GHz without digital notch filtering (blue) indicates correlation between the processed spectra, here caused by the aforementioned sinusoidal background component in the frequency spectra. This dataset is shown here for having the longest measurement time, the same effect however persists similarly for all datasets. With the digital notch filter applied (orange in fig. 8.6), the time evolution of σ^{-2} perfectly matches the ideal case, indicating the absence of further correlations and no significant dead-time.

This perspective is especially valuable because it is agnostic of the type of correlation. In the present example the correlating component removed by the digital notch filter is localised in the Fourier spectrum and therefore easily visible. In case this would not hold, the σ^{-2} evolution would nevertheless provide hints of problems in the dataset.

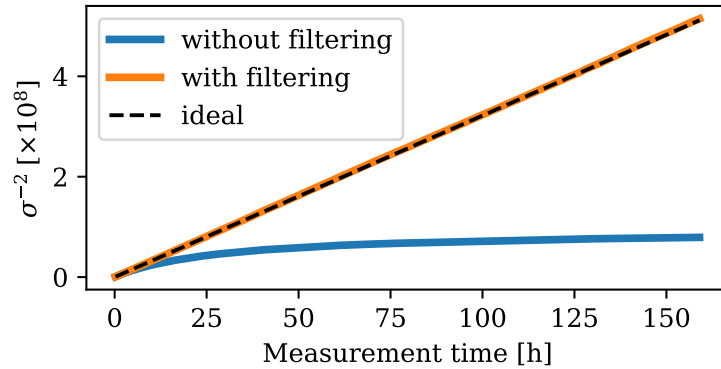


Figure 8.6: Evolution of the standard deviation of averaged processed spectra with respect to measurement time for the dataset with its β^2 peak at 18.557 GHz. If individual spectra are uncorrelated and the experiment has negligible dead-time the evolution should follow the case marked as “ideal” (dashed black). Without applying the digital notch filter a significant deviation is visible (blue), which vanishes when the filter is applied (orange).

At the level of the combined spectrum the effect of the Fourier peak on the spectrum itself becomes obvious. Due to the significantly lower uncorrelated noise component compared to a single processed spectrum, the oscillation corresponding to the Fourier peaks visible in figs. 8.5 and 8.8 can easily be seen in a zoomed in version of the combined spectra for both booster configurations (fig. 8.7). Here as well the filter demonstrates its effectiveness at removing the oscillation, since the combined spectra after filtering look like white noise at least at first glance. This is investigated thoroughly in sec. 8.3.

Applying a Fourier transform to the combined spectra before and after digital notch filtering again demonstrates the effectiveness of the digital notch filter at removing




Figure 8.7: Excerpt of combined spectra for both booster configurations with (orange) and without (blue) applying the digital notch filter. The dashed vertical lines indicate the absolute frequencies 18.53 GHz and 19.19 GHz of the spectrum pieces shown.




Figure 8.8: Absolute value of the Fourier transform of the combined spectrum for booster configuration 1 before and after applying the digital notch filter. The inset again shows how exactly the digital notch filter affects the region around the peak at $109.22 \mu\text{s}$.

the Fourier peaks. The result is shown in fig. 8.8 for the booster configuration 1, but equivalently holds for the other configuration. Additionally, note how the SG filter successfully removed the peak at $\Delta t = 0\text{s}$, corresponding to the large-scale background component.

This subsection demonstrates that a combination of digital notch and SG filtering successfully removes all correlations present between individual raw spectra. No remaining correlation can be seen for measurement times of at least 159 hours, potentially more. This means the receiver system is effective for the task of investigating axion dark matter in the range around 19 GHz. Investigating the physical reason for the Fourier peaks with a characteristic timescale of $\Delta t = 109.22 \mu\text{s}$ is nevertheless imperative. The goal is to avoid their presence altogether in the raw spectra of future measurements. This is actively being pursued at the time of writing.

8.2.2 Obtaining SG Parameters

There are two parts to obtaining suitable parameters for the SG filter used to process the spectra. Firstly, as described in subsec. 7.3.5 the SNR attenuation induced by a specific set of filter parameters has to be simulated. Secondly, one must investigate if the residual spectrum contains any non-axionic deviations from Gaussianity, e.g. using eq. (7.24). It is sufficient to treat only one of the two filter parameters, polynomial order o and filter window width w , as free since they are approximately degenerate (see sec. 7.3.5).

To investigate the SNR attenuation, the procedure described in subsec. 7.3.5 was applied separately to all five datasets for a range of reasonable window widths w , always assuming $o = 4$. The attenuation should depend only on the RBW and w , so all datasets should yield the same results. This is approximately the case as can be seen in fig. 8.9, the datasets show systematic percent-level variation. The calculation




Figure 8.9: Evolution of η_{SG} with filter window size for all five datasets. The dashed vertical line indicates the filter window size used for the final limit.

using the dataset with max β^2 at 19.196 GHz produces the smallest η_{SG} for all window sizes. This systematic shift may be explainable by all window sizes using identical values for ideal mean μ_{id} and ideal standard deviation σ_{id} , which are also subject to uncertainty. To be conservative, in the following the lowest value for η_{SG} of any of the datasets was used for a given filter window. The same simulation also yields the expected standard deviation of the grand spectrum σ_f when a specific SG filter is used. It increases with increasing w .

According to fig. 8.9, larger window sizes are clearly preferable due to their increased η_{SG} , i.e. lower SNR attenuation. However, it needs to be ensured that no residual structures, i.e. small excesses due to features in the background, remain in combined or grand spectrum that could be mistakenly associated with an axion signal. To choose an appropriate filter window width, the Gaussianity of the grand spectra for both booster configurations is investigated.

Fig. 8.10 presents grand spectra and histograms thereof for a variety of SG filter window sizes and $o = 4$. With the exception of configuration 2 for $w = 1501$ all shown grand spectra display approximately Gaussian behaviour as demonstrated by the histograms. The standard deviation of a best fit Normal distribution on the histograms is $\sigma_g < 1$ in all but the aforementioned case. This is expected from the interplay between SG filter and the cross-correlation with the axion lineshape as already mentioned in subsec. 7.3.5.

The effect, however, cannot fully explain the increase of σ_g with w seen in fig. 8.10. The actual standard deviation of the grand spectra, σ_g , should for both configurations be close to the theoretical expectation σ_f obtained in the SNR attenuation simulation, which by construction does not include residual structures using the SG filters discussed here. Fig. 8.11 shows how σ_g , obtained from the data, deviates from σ_f , obtained from simulation, for larger w . Both booster configurations show some increase beyond the expected value of $\sigma_g/\sigma_f - 1 = 0$, however much larger values are obtained for configuration 2 up to a relative deviation of almost 10% for $w = 1501$.

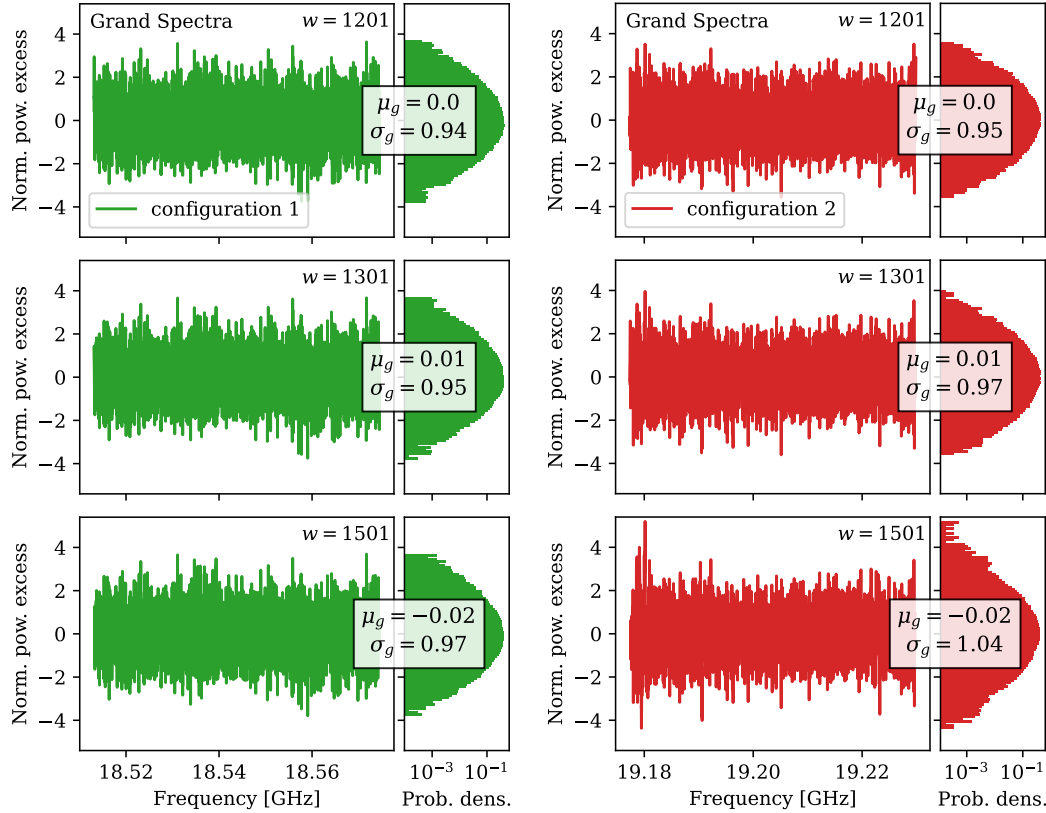


Figure 8.10: Grand spectra and histograms thereof for both booster configurations using $\sigma = 4$ and various SG filter widths w . Mean μ_g and standard deviation σ_g are given for all grand spectra shown.

When increasing w it is inevitable that grand spectra display statistically significant excesses at locations of not fully removed RF structures. Configuration 2 for $w = 1501$ clearly displays such an excess at a frequency of 19.180 GHz with a local significance of 5.45σ . It is significantly attenuated when using smaller w , contrary to an excess of axionic origin the shrinking of which would be governed by the factor $\eta_{\text{SG}}\sigma_f$ depending on w . The excess has been thoroughly investigated and seen to be significantly broader than the FWHM of the axion lineshape. For these two reasons an interpretation of this excess in terms of an axion signal can be ruled out.

An excess at the same frequency is still visible when using $w = 1301$ at 4.19σ local significance. Therefore for the final limit a fourth-order SG filter with window width $w = 1201$ data points was chosen, leading to $\eta_{\text{SG}} = 0.917$ and $\sigma_f = 0.938$.

8.2.3 Gaussianity of Processed and Combined Spectra

In the previous subsection the grand spectra were used to inform the choice of w . As mentioned before, it is vital also for the combined spectra to not include residual structures for this choice. Unlike grand spectra for which adjacent bins get correlated by the cross-correlation with the axion lineshape, combined spectra are expected to




Figure 8.11: Relative deviation between expected grand spectrum standard deviation σ_f and the value obtained from the data, σ_g , for the two booster configurations. Booster configuration 2 displays deviations $> 2\%$ for filter window widths $w > 1201$ indicating the presence of not fully removed RF structures in the spectrum.

display pure white noise behaviour, i.e. no correlation between adjacent bins should exist. Therefore the Runs Gaussianity test introduced in sec. 7.1 can be applied.

The result can be seen in fig. 8.12 (left). The combined spectrum for configuration 2, corresponding to higher frequencies around 19.2 GHz, is perfectly compatible with pure Gaussian white noise, quantified by the metric $G_2 = 0.8 \in [0.7, 3.5]$ being within the interval defined in sec. 7.1. Configuration 1 displays one significant excess leading to $G_1 = 126 \gg 3.5$. The excess is shown in fig. 8.12 (right). It has a width significantly smaller than the axion FWHM, making it impossible to arise from homogeneous axion dark matter. It also does not seem to be a transient event, since it does not coincide with particularly significant excesses in any of the processed spectra. Contamination by a small structure at radio frequencies could be a possible explanation, even though a purely statistical excess cannot be ruled out at this point.

Investigating all processed spectra individually, only one significant transient event was found. The spectrum taken on Feb, 28th at 12:40 PM displays an increase of $\sim 5\%$ in its standard deviation compared to the theoretical expectation due to several excesses between 19.22 GHz and 19.23 GHz. Fig. 8.13 plots all frequency bins of all processed spectra for the individual dataset in a histogram. There this transient event can be seen as the increased probability density for power excesses $> 5\sigma$ for the dataset with $\max \beta^2$ at 19.196 GHz compared to other datasets. With this exception all datasets perfectly follow the expected Gaussian distribution.

Fig. 8.14 shows the distribution of the metric G for all processed spectra individually. The distributions match the theoretical expectation presented in sec. 7.1 well. There are three outliers with $G > 20$: One is in the dataset with $\max \beta^2$ at 18.557 GHz, however it does not coincide with the excess in the combined spectrum for configuration 1 discussed earlier. Two more can be found in the dataset with $\max \beta^2 = 19.196$ GHz and can be attributed to the aforementioned transient event.

All datasets display no relevant deviations from Gaussian behaviour and can therefore be used for limit setting without further processing.




Figure 8.12: **Left:** Runs Gaussianity test for combined spectra of both booster configurations. The Runs weights of both configurations are compatible with Gaussian behaviour except for one excess in configuration 1. **Right:** Zoom into the combined spectrum of configuration 1 at the frequency f_0 of the aforementioned excess.




Figure 8.13: Histogram of all frequency bins of all processed spectra contained in the five datasets.




Figure 8.14: Distribution of metric G for all processed spectra contained in the five datasets. Overall, the distributions follow the expectation well with three exceptions having $G > 20$.

8.2.4 Systematic Uncertainties

Before proceeding to the limit setting, however, the impact of systematic uncertainties on the limit has to be understood.

The limit in frequency bin j is given by eq. (7.22) as

$$|g_{a\gamma}|_j = |g_{a\gamma}^{\text{ref}}| \sqrt{\frac{g_j/\sigma_f + 1.645}{\eta_{\text{SG}} R_j}}.$$

In this formula only R_j , which encodes the SNR expectation for a reference axion-photon coupling in every frequency bin, is subject to significant systematic uncertainties. Those are: Gain drifts of the receiver chain, uncertainty in the power calibration factor and boost factor uncertainty. Their assessment is outlined in subsecs. 3.3.3 and 3.3.4. Fig. 8.15 shows all three systematic uncertainties for the five datasets individually. Each individual uncertainty has comparable magnitudes for all datasets. Uncertainty in the power boost factor is largest (top panel) followed by uncertainty in the power calibration factor (middle panel). Note that two datasets have the same calibration uncertainty because no separate power calibrations were possible (compare also sec. 8.1). The gain drift is defined as variation of the mean $m_i = \frac{1}{N} \sum_j b_{ij}$ between different spectra i , where the sum is over frequency index $j \in [1, N]$. b_{ij} again are SG filtered, uncalibrated spectra. Gain drift only plays a subdominant role (lower panel).

To combine the three sources of uncertainty and properly translate them to an effect on the exclusion limit, the path towards R_j has to be retraced. R_j is defined in (7.13) in terms of $\sigma_{c,j}$, the standard deviation of individual frequency bins of the combined spectrum before normalisation. $\sigma_{c,j}$ as defined in (7.10) in turn depends on a sum over the weights w_{ij} used to combine spectra. These weights are proportional to the inverse square of raw spectra and power calibration factor, as well as the square of the axion signal expectation, which itself is proportional to β_{ij}^2 (see eq. (7.7)). Because power calibrations should be considered correlated between individual datasets, while




Figure 8.15: Leading systematic uncertainties for the five datasets as a function of frequency. From top to bottom: Power boost factor uncertainty, calibration factor uncertainty, gain drift.

other systematic uncertainties are not, it is necessary to divide their contributions to the overall uncertainties. No other quantities used in the calculation of w_{ij} are subject to significant systematic uncertainties, therefore the relative uncertainty from the two different types of systematics can be calculated as

$$\frac{\delta w_{ij}}{w_{ij}} \Big|_{\text{no } Pf} = 2 \sqrt{\left(\frac{\delta \beta_{ij}^2}{\beta_{ij}^2}\right)^2 + \left(\frac{\delta m}{m}\right)^2}, \quad \frac{\delta w_{ij}}{w_{ij}} \Big|_{\text{just } Pf} = 2 \left(\frac{\delta P_{ij}^f}{P_{ij}^f}\right), \quad (8.2)$$

where again P_{ij}^f is the power calibration factor and $m = \frac{1}{M_k} \sum_i m_i$ for M_k spectra indexed by i in a dataset k .

Errors on spectra in a given dataset are considered to be perfectly correlated. This is inevitable for β_{ij}^2 and P_{ij}^f , since for each dataset only one estimate is available for both quantities, i.e. $\beta_{ij}^2 = \beta_j^2$ and $P_{ij}^f = P_j^f$ for all spectra in one dataset but not across different datasets. Similarly the gain drift was calculated by considering all spectra for a given dataset and is therefore unaffected by a sum over the spectra i . Thus, one can drop the index i and introduce the dataset index k such that $\delta w_{ij}/w_{ij} \equiv \delta w/w|^{(kj)}$ for all spectra i in dataset k .

This is true for both uncertainty from power calibration factor as well as other uncertainties. The difference lies in the way different datasets have to be combined. Here systematic uncertainties from the power calibration factor are treated as perfectly correlated and other uncertainties as perfectly uncorrelated, leading to

$$\frac{\delta \sigma_{c,j}}{\sigma_{c,j}} \Big|_{\text{no } Pf} = \frac{1}{2} \sqrt{\sum_k \left(\frac{\delta w}{w} \Big|_{\text{no } Pf}^{(kj)}\right)^2}, \quad (8.3)$$

$$\frac{\delta \sigma_{c,j}}{\sigma_{c,j}} \Big|_{\text{just } Pf} = \frac{1}{2} \sum_k \frac{\delta w}{w} \Big|_{\text{just } Pf}^{(kj)}, \quad (8.4)$$

where the sum is over the datasets k belonging to one of the booster configurations. Cross-correlating with the axion lineshape again leaves the error unchanged, because errors on individual frequency bins are approximately equal at least for the relevant frequency differences. Therefore the error in R_j is calculated as

$$\frac{\delta R_j}{R_j} = \sqrt{\left(\frac{\delta \sigma_{c,j}}{\sigma_{c,j}} \Big|_{\text{no } Pf}\right)^2 + \left(\frac{\delta \sigma_{c,j}}{\sigma_{c,j}} \Big|_{\text{just } Pf}\right)^2}. \quad (8.5)$$

The limit on $g_{a\gamma}$ is $\propto R_j^{-1/2}$, therefore the limit including systematic uncertainty is obtained via

$$|g_{a\gamma}|_j \rightarrow |g_{a\gamma}|_j \times \left(1 + \frac{1}{2} \frac{\delta R_j}{R_j}\right), \quad (8.6)$$

where $|g_{a\gamma}|_j$ denotes the limit in frequency bin j without systematic uncertainties.

The resulting value for the total uncertainty $\frac{1}{2} \frac{\delta R_j}{R_j}$ for both booster configurations is shown in fig. 8.16. It is between 5% to 10% for almost the full frequency range, indicating only minor suppression in terms of the final limit.

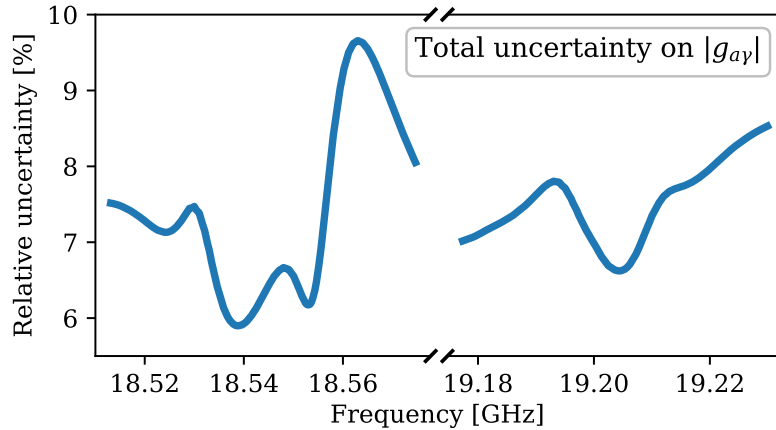


Figure 8.16: Total systematic uncertainty on $|g_{a\gamma}|$ for both booster configurations as a function of frequency.

8.3 Axion Limits

The previous sections discussed in detail how the datasets are obtained and processed. The final grand spectra are shown in fig. 8.17 (yellow) together with the expected signal amplitude for a reference $\mathcal{C}_{a\gamma} = 2 \times 10^3$ over the frequency range (dashed grey line). The signal expectation can be divided into the individual contributions from different datasets (coloured dashed lines). The multimodal shape of the expected signal amplitude is explained by these individual contributions. Their shape is dominated by the frequency dependence of the power boost factors shown in fig. 8.3. The system temperature also plays a role through its influence on the noise. The interplay between these two contributions can lead to slightly multimodal structures even for a single dataset as can be seen for the dataset with max β^2 at 18.531 GHz.

As shown in subsec. 8.2.3, the grand spectra do not display significant deviations from Gaussianity with standard deviations of $\sigma_1 = 0.94$ and $\sigma_2 = 0.95$. The maximal excesses in the grand spectra for booster configurations 1 and 2 are 3.88σ and 3.75σ local significance. The probability for a local excess of at least 3.88σ in both grand spectra combined can be determined to be $p \approx 0.13$ using eq. (4.1) with an effective number of samples governed by (7.23). Existing excesses can therefore be explained by a noise-only hypothesis. Since no evidence for a DM signal was found, a limit is constructed.

The obtained limit as a function of frequency/ axion mass for both configurations can be seen in fig. 8.18. Its shape corresponds to the expected signal amplitudes shown in fig. 8.17 and is determined by the frequency dependent boost factor and system noise temperature. As explained in sec. 7.4 limits exceeding the median sensitivity by more than 1σ are not shown for better intelligibility of the plot. Existing limits from the CAST helioscope [186] are surpassed by up to a factor of 4 if a local density of dark matter axions of $\rho_a = 0.3 \text{ GeV}/\text{cm}^3$ is assumed.

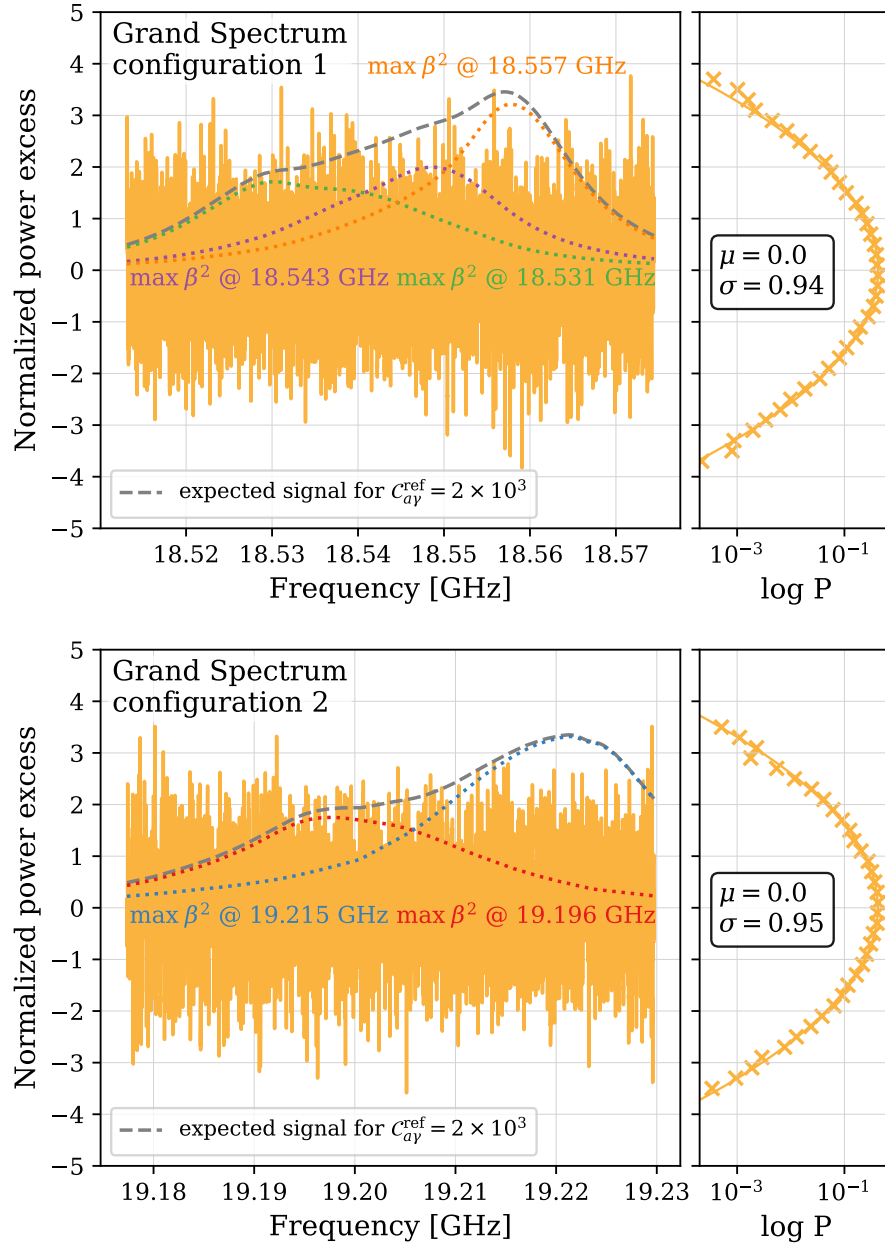


Figure 8.17: Grand spectra for both booster configurations (yellow) as well as expected signal amplitudes for all five datasets individually (dotted) and combined (dashed). Right panels additionally show histograms of the normalised power excesses in both grand spectra as well as mean μ and standard deviation σ of best-fit Gaussian distributions.




Figure 8.18: 95% confidence limit of the 2024 dark matter axion search using the CB 200 MADMAX prototype setup inside the MORPURGO magnet at CERN assuming a local axion dark matter density of $\rho_a = 0.3 \text{ GeV/cm}^3$. The median sensitivity and central one sigma band for expected statistical fluctuation of the limit are shown as dashed and dotted black lines. For comparison the most recent CAST helioscope limit [186] in the frequency range is given (grey).

8.4 Summary & Discussion

The measurements with a MADMAX prototype at CERN in 2024 resulted in world-leading limits on the axion-photon coupling in two mass ranges around $76 \mu\text{eV}$ and $80 \mu\text{eV}$. These are the first limits on axion dark matter published by the MADMAX collaboration.

Five datasets using two different configurations of the CB 200 booster have been used in the analysis. Parts of the datasets with difficult to model boost factor estimates were removed. Here, future work could investigate an empirical and objective way to more appropriately model the reduction in sensitivity for frequency regions where the boost factor is difficult to estimate.

The datasets include a set of peaks at positions in the Fourier spectrum corresponding to timescales of multiples of $109 \mu\text{s}$. They can effectively be removed using the digital notch filter presented in sec. 7.2. The physical reason behind these background components is actively investigated with the goal of avoiding them in future measurements.

The data does not contain other background structures with fluctuations on scales comparable to the axion FWHM, making it possible to apply a fourth-order SG filter with a window width $> 1 \text{ MHz}$ leading to relatively small SNR attenuation of $\eta_{\text{SG}} = 0.92$. The resulting processed and combined spectra display almost perfect Gaussianity, a key feature for being able to extract sensitive limits. From systematic

uncertainties a relative worsening of the limit of 5% to 10% is incurred. Assuming a local axion dark matter density of $\rho_a = 0.3 \text{ GeV/cm}^3$, the obtained limits exceed the previous most stringent limit by up to a factor of four.

The limits were obtained with only two weeks of measurement time using a simple 20 cm setup with three disks inside of a magnet with $B_e \sim 1 \text{ T}$. This highlights the substantial discovery potential of the MADMAX concept, particularly if future designs can achieve magnetic fields of $B_e \sim 10 \text{ T}$ with multiple 100 cm diameter disks leading to a boost factor of up to $\beta^2 = 10^4$. A full-scale MADMAX setup could make a significant portion of the QCD axion parameter space for masses above $m_a \sim 40 \mu\text{eV}$ accessible for experimental exploration.

Conclusion & Outlook

Axion research is entering a golden decade with first experiments reaching sensitivity to minimal QCD axion models and many more promising efforts on the way, including MADMAX.

The work described in this thesis contributes to the axion search program by providing a competitive statistical analysis pipeline for first MADMAX prototype measurements. It combines standard practices in the field, e.g. to properly weight and sum individual power spectra, with new methods applied to axion dark matter data for the first time, such as the Runs Gaussianity test or the digital notch filter.

The Runs Gaussianity test is a new, sensitive way to investigate a set of data points regarding any kind of deviations from white noise. Being based on the expected distribution of consecutive excesses, it is agnostic of the shape of the deviation making it well-suited for the hunt for residual background structures.

The digital notch filter became necessary due to a sinusoidal background component, which could not be removed using standard methods. By replacing relevant parts of the Fourier transform of the spectrum with values informed by the rest of the spectrum, it efficiently removes such components without significant influence on potential axion signals.

Thanks to the digital notch filter, the statistical analysis pipeline is able to remove backgrounds from individual measurements. It then combines them and, since no relevant deviations from Gaussianity are found, cross-correlates the result with the expected axion lineshape. No significant evidence for dark matter axions is found at axion masses in the ranges of $(76.56 \text{ to } 76.82) \mu\text{eV}$ and $(79.31 \text{ to } 79.53) \mu\text{eV}$. Signal attenuation due to the background removal procedure as well as systematic uncertainties are taken into account to yield upper limits on the axion-photon coupling at 95% confidence level. The limits improve upon existing constraints by up to a factor of four, if a local axion density of $\rho_a = 0.3 \text{ GeV}/\text{cm}^3$ is assumed.

This outcome was achieved with just two weeks of measurement time using a compact 20 cm setup containing three disks within a $B_e \sim 1 \text{ T}$ magnet, underscoring the significant discovery potential of the MADMAX concept. Plans are underway to drastically increase the axion-photon conversion rate by scaling up both the size and number of disks. Furthermore, efforts are being made to lower the system temperature through the use of cryogenic quantum sensors. Over the coming years, extended measurements are planned using the infrastructure from CERN until the ultimate goal of a dedicated magnet for MADMAX with $B_e \sim 10 \text{ T}$ can be realised. A setup consisting of multiple 100 cm diameter disks and reaching boost factors of $\beta^2 \sim 10^4$ could enable the experiment to probe a substantial portion of the QCD axion parameter space at high axion masses. The statistical analysis pipeline described in this thesis could be

applied during all these stages of development with only minor modifications. In fact, it has already been adapted for use in a dark photon dark matter search with an open MADMAX setup [5].

Standard haloscope analysis typically relies on linear filters for background removal, which distort the shape of potential axion signal. If these alteration are not accounted for, they may reduce sensitivity, particularly in a Bayesian framework where prior knowledge of signal parameters is explicitly considered. The bias-free signal estimation approach presented in this thesis may alleviate this issue by integrating the signal shape modification caused by the filter directly into the signal model.

With the maturing of the field of axion research it is only a matter of time until first analyses divert from the prevalent frequentist paradigm and apply Bayesian methods to axion measurements. First efforts to this effect are underway [120] by investigating updates in the prior belief. Using Bayesian inference in an axion context is not conceptually difficult as shown by its application to a set of simulated MADMAX-like datasets in this thesis. It was primarily hampered by the absence of usable prior distributions for general QCD axions beyond minimal models.

The work described in this thesis presents such a comprehensive prior distribution for the first time. It systematically investigates the axion-photon couplings of DFSZ-type axion models with up to nine Higgs doublets. By combining the results with a pre-existing KSVZ-type model catalogue, a joint theoretical prior for both major types of invisible axion models can be constructed.

Even in the large set of DFSZ-type models analysed, anomaly ratios of the minimal DFSZ model are still favoured. In contrast to the minimal models, no significant differences between the anomaly ratio distributions for general DFSZ- and KSVZ-type models are found. A subset of DFSZ-type models with domain wall number of unity was found and could be considered phenomenologically preferred. These models additionally display photon couplings, which are on average almost one order of magnitude larger than for the full set of models, making them more easily accessible for experiments.

The investigations could be extended into multiple directions, such as better estimating anomaly ratio distributions for eight or more Higgs doublets and accounting for massive neutrinos. However, mathematical considerations show that the resulting distributions would likely be similar.

The ultimate goal of the extensive axion search program cannot be the setting of incrementally tighter limits, but to eventually find the axion and in doing so solve two of the greatest mysteries in physics. If an axion signal is found, the methods presented here are necessary for a bias-free estimation of its signal parameters. This thesis therefore not only helps to shrink the viable axion parameter space by analysing data from a MADMAX prototype setup, but also contributes by providing methods applicable to measurements beyond the MADMAX prototype and towards axion detection.

List of Figures

1.1	Anomaly ratio distributions for a single heavy quark in any representation and the complete set of preferred models with up to 28 heavy quarks.	14
2.1	Schematic of topological defects in 2D after explicit breaking of the PQ symmetry for $N_{\text{DW}} > 1$	18
2.2	The QCD axion parameter space.	20
2.3	Axion signal shape computed from (2.11) for an axion at 19 GHz.	22
3.1	Sketch of axion-induced emission sourced at a dielectric interface in a magnetic field in the classical field picture.	25
3.2	Simulated boost factor curves for two different configurations for a system with 20 dielectric disks.	27
3.3	Baseline design of the MADMAX experiment.	29
3.4	Influence of geometrical inaccuracies on the obtainable power boost factor.	31
3.5	Exploded view of the CB 200 booster.	32
3.6	Example power calibration factor from the calibration measurement of Feb, 21st.	34
3.7	Comparison between TE_{11} mode of a circular waveguide and axion induced electric field.	35
3.8	1D power boost factors.	36
4.1	Example scenario of a frequentist hypothesis test.	38
5.1	Anomaly ratio distributions for DFSZ-type models with 4 Higgs doublets.	54
5.2	Influence of drawing PQ charges from different distributions on the resulting anomaly ratio distribution.	56
5.3	Anomaly ratio distributions for different numbers of Higgs doublets.	57
5.4	Comparison between anomaly ratio distributions for KSVZ-type and DFSZ-type axion models.	62
5.5	Cumulative probability of models with $ \mathcal{C}_{a\gamma} $ larger than the indicated values.	64
5.6	Cumulative probability of models with $ \mathcal{C}_{a\gamma} $ larger than the indicated values for the complete set of DFSZ-type and KSVZ-type models as well as for models with $N_{\text{DW}} = 1$ specifically.	65
5.7	Probability density in $\log \mathcal{C}_{a\gamma} $ -space of the combined DFSZ-type and KSVZ-type “all N_{DW} ”-case.	66

5.8	Preferred $g_{a\gamma}$ parameter space for all non-minimal DFSZ-type as well as preferred KSVZ-type models.	67
6.1	Components of a simulated example dataset.	75
6.2	Distribution of the signal amplitude across the mock datasets, relative to the background and the noise level.	77
6.3	Effect of using a bias-free signal model for the fit on the data.	79
6.4	Validation of the chosen SG filter parameters in terms of background estimation and noise removal.	80
6.5	Reduced chi-square distributions for all mock datasets, from fits using the standard biased method as well as the bias-free method.	81
6.6	Distribution of the posterior quantile bin in which the true signal parameter resides, for the three parameters.	82
6.7	Posterior and prior probability density for signal parameter $\mathcal{C}_{a\gamma}$ for a mock dataset with no simulated axion present.	83
7.1	Exemplary calculation of runs weights.	87
7.2	Probability density distributions for w given reference datasets of various lengths N	88
7.3	Dependence of the metric G on parameters N , the number of data points in the dataset, and K , the number of ranges used.	90
7.4	Three example use cases demonstrating the power of the runs Gaussianity test: Peak search, inserted sine wave and correlated noise.	91
7.5	Fourier transform of a mock frequency spectrum.	93
7.6	Fourier peak of the sinusoidal component before and after the digital notch filter was applied.	95
7.7	Effect of and error from the digital notch filter.	96
7.8	Effect of the digital notch filter on the shape of a potential axion signal.	97
7.9	Influence of amplitude and location of the Fourier peak on signal shape and error for the removal of the sine wave.	98
7.10	Datasets simulated for this section consisting of three components: Background, white noise and equidistant signals.	100
7.11	Transfer functions of different SG filters.	102
7.12	Type-II example spectrum and isolated signal component of the spectrum after processing.	103
7.13	Spectra at intermediate steps towards limit setting compared with simulated axion signal components at each stage.	105
7.14	Expected signal amplitude for the simulated dataset.	107
7.15	Gaussianity of the combined spectrum when accounting for the power calibration before or after constructing the processed spectra.	110
7.16	Result of the SG attenuation simulation.	111
7.17	Effect of signal misalignment with respect to edges of the frequency bins.	112
7.18	Simulated limit obtained by applying the analysis pipeline to the mock dataset.	114
8.1	Time evolution of the available magnetic field strength.	120
8.2	Example spectra for the five datasets used in the analysis.	122

8.3	Boost factor distributions used in the analysis for the five datasets. Lines denote mean β^2 and shaded regions give $\pm 1\sigma$ intervals.	122
8.4	Distribution of boost factor samples at the maximum of the boost factor peaks.	123
8.5	Absolute value of the Fourier transform of a single spectrum before and after applying the digital notch filter.	124
8.6	Evolution of the standard deviation of averaged processed spectra with respect to measurement time.	125
8.7	Excerpt of combined spectra for both booster configurations with and without applying the digital notch filter.	125
8.8	Absolute value of the Fourier transform of the combined spectrum for booster configuration 1 before and after applying the digital notch filter.	126
8.9	Evolution of η_{SG} with filter window size for all five datasets.	127
8.10	Grand spectra and histograms thereof for both booster configurations using various SG filter widths w	128
8.11	Relative deviation between expected grand spectrum standard deviation σ_f and the value obtained from the data, σ_g , for the two booster configurations.	129
8.12	Runs Gaussianity test for combined spectra of both booster configurations.	130
8.13	Histogram of all frequency bins of all processed spectra contained in the five datasets.	130
8.14	Distribution of metric G for all processed spectra contained in the five datasets.	131
8.15	Leading systematic uncertainties for the five datasets as a function of frequency.	132
8.16	Total systematic uncertainty on $ g_{a\gamma} $ for both booster configurations as a function of frequency.	134
8.17	Grand spectra for both booster configurations as well as expected signal amplitudes for all five datasets individually and combined.	135
8.18	95% confidence limit of the 2024 dark matter axion search using the CB200 MADMAX prototype setup inside the MORPURGO magnet at CERN assuming a local axion dark matter density of $\rho_a = 0.3 \text{ GeV}/\text{cm}^3$	136

List of Tables

1.1	Weak hypercharges of SM fermions excluding neutrinos.	10
5.1	Resulting conditions on PQ charges from quadrilinears, constructed from corresponding bilinears.	51
5.2	All possible solutions for PQ charges of Higgs doublets in terms of χ_S and all possible anomaly ratios for the $n_D = 3$ Yukawa sector under consideration.	52
5.3	Important statistics of DFSZ-type models broken down by number of Higgs doublets n_D	58
5.4	$ \mathcal{C}_{a\gamma} $ lower prior limits for selected combinations of models.	68
6.1	Overview over the fixed parameters in the analysis and the priors used for the three non-fixed parameters.	77
6.2	Coverage properties of the Bayesian credible intervals.	83
7.1	Selected properties of the G distributions shown in fig. 7.3.	90
7.2	Variables introduced for different spectra used in the analysis.	100
8.1	Summary of the datasets used for limit setting.	121

Bibliography

- [1] Johannes Diehl and Emmanouil Koutsangelas. ‘Dine-Fischler-Srednicki-Zhitnitsky-type axions and where to find them’. In: *Phys. Rev. D* 107.9 (2023), p. 095020. DOI: 10.1103/PhysRevD.107.095020. arXiv: 2302.04667 [hep-ph].
- [2] Johannes Diehl, Jakob Knollmüller and Oliver Schulz. ‘Bias-free estimation of signals on top of unknown backgrounds’. In: *Nucl. Instrum. Meth. A* 1063 (2024), p. 169259. DOI: 10.1016/j.nima.2024.169259. arXiv: 2306.17667 [astro-ph.IM].
- [3] B. Ary dos Santos Garcia, J. Diehl, D. Leppla-Weber, A. Martini et. al. ‘First search for axion dark matter with a MADMAX prototype’. *in preparation*.
- [4] S. Knirck et al. ‘Simulating MADMAX in 3D: requirements for dielectric axion haloscopes’. In: *JCAP* 10 (2021), p. 034. DOI: 10.1088/1475-7516/2021/10/034. arXiv: 2104.06553 [hep-ex].
- [5] J. Egge et al. ‘First search for dark photon dark matter with a MADMAX prototype’. In: (Aug. 2024). arXiv: 2408.02368 [hep-ex].
- [6] Emmanouil Koutsangelas. ‘How the axion paves the way beyond the standard model’. PhD thesis. Munich U., 2023.
- [7] Luca Di Luzio et al. ‘The landscape of QCD axion models’. In: *Physics Reports* 870 (July 2020), pp. 1–117. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2020.06.002. URL: <http://dx.doi.org/10.1016/j.physrep.2020.06.002>.
- [8] Luca Di Luzio, Federico Mescia and Enrico Nardi. ‘Window for preferred axion models’. In: *Phys. Rev. D* 96.7 (2017), p. 075003. DOI: 10.1103/PhysRevD.96.075003. arXiv: 1705.05370 [hep-ph].
- [9] Vaisakh Plakkot and Sebastian Hoof. ‘Anomaly ratio distributions of hadronic axion models with multiple heavy quarks’. In: *"Phys. Rev. D"* 104.7, 075017 (Oct. 2021), p. 075017. DOI: 10.1103/PhysRevD.104.075017. arXiv: 2107.12378 [hep-ph].
- [10] Kazuo Fujikawa. ‘Path-Integral Measure for Gauge-Invariant Fermion Theories’. In: *Phys. Rev. Lett.* 42 (18 Apr. 1979), pp. 1195–1198. DOI: 10.1103/PhysRevLett.42.1195. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.42.1195>.
- [11] Gerard 't Hooft. ‘Symmetry Breaking Through Bell-Jackiw Anomalies’. In: *Phys. Rev. Lett.* 37 (1976). [226(1976)], pp. 8–11. DOI: 10.1103/PhysRevLett.37.8.

- [12] G. 't Hooft. 'Computation of the quantum effects due to a four-dimensional pseudoparticle'. In: *Phys. Rev. D* 14 (12 Dec. 1976), pp. 3432–3450. DOI: 10.1103/PhysRevD.14.3432. URL: <https://link.aps.org/doi/10.1103/PhysRevD.14.3432>.
- [13] C. Abel et al. 'Measurement of the Permanent Electric Dipole Moment of the Neutron'. In: *Phys. Rev. Lett.* 124.8 (2020), p. 081803. DOI: 10.1103/PhysRevLett.124.081803. arXiv: 2001.11966 [hep-ex].
- [14] B. Graner et al. 'Reduced Limit on the Permanent Electric Dipole Moment of Hg 199'. In: "*Phys. Rev. Lett.*" 116.16, 161601 (Apr. 2016), p. 161601. DOI: 10.1103/PhysRevLett.116.161601. arXiv: 1601.04339 [physics.atom-ph].
- [15] R. J. Crewther et al. 'Chiral Estimate of the Electric Dipole Moment of the Neutron in Quantum Chromodynamics'. In: *Phys. Lett.* 88B (1979). [Erratum: *Phys. Lett.* 91B, 487 (1980)], p. 123. DOI: 10.1016/0370-2693(80)91025-4, 10.1016/0370-2693(79)90128-X.
- [16] Varouzhan Baluni. 'CP-nonconserving effects in quantum chromodynamics'. In: *Phys. Rev. D* 19 (7 Apr. 1979), pp. 2227–2230. DOI: 10.1103/PhysRevD.19.2227. URL: <https://link.aps.org/doi/10.1103/PhysRevD.19.2227>.
- [17] Antonio Pich and Eduardo de Rafael. 'Strong CP-violation in an effective chiral lagrangian approach'. In: *Nuclear Physics B* 367.2 (1991), pp. 313–333. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(91\)90019-T](https://doi.org/10.1016/0550-3213(91)90019-T). URL: <https://www.sciencedirect.com/science/article/pii/055032139190019T>.
- [18] Maxim Pospelov and Adam Ritz. 'Theta vacua, QCD sum rules, and the neutron electric dipole moment'. In: *Nucl. Phys. B* 573 (2000), pp. 177–200. DOI: 10.1016/S0550-3213(99)00817-2. arXiv: hep-ph/9908508.
- [19] Lorenzo Bartolini et al. 'Neutron electric dipole moment from gauge/string duality'. In: *Phys. Rev. Lett.* 118.9 (2017), p. 091601. DOI: 10.1103/PhysRevLett.118.091601. arXiv: 1609.09513 [hep-ph].
- [20] M. Abramczyk et al. 'Lattice calculation of electric dipole moments and form factors of the nucleon'. In: *Phys. Rev. D* 96.1 (2017), p. 014501. DOI: 10.1103/PhysRevD.96.014501. arXiv: 1701.07792 [hep-lat].
- [21] Jack Dragos et al. 'Confirming the Existence of the strong CP Problem in Lattice QCD with the Gradient Flow'. In: *Phys. Rev. C* 103.1 (2021), p. 015202. DOI: 10.1103/PhysRevC.103.015202. arXiv: 1902.03254 [hep-lat].
- [22] Roberto D. Peccei. 'The Strong CP Problem and Axions'. In: *Axions: Theory, Cosmology, and Experimental Searches*. Ed. by Markus Kuster, Georg Raffelt and Berta Beltrán. Springer Berlin Heidelberg, 2008, pp. 3–17. ISBN: 978-3-540-73518-2. DOI: 10.1007/978-3-540-73518-2_1. URL: https://doi.org/10.1007/978-3-540-73518-2_1.
- [23] Javier Redondo. *Lecture on axions, Invisibles16 Summer School*. URL: <https://indico.cern.ch/event/464396/contributions/2217970/>.
- [24] S. Yu. Khlebnikov and M. E. Shaposhnikov. 'Extra space-time dimensions: Towards a solution to the strong CP-problem'. In: *Physics Letters B* 203.1-2 (Mar. 1988), pp. 121–124. DOI: 10.1016/0370-2693(88)91582-1.

- [25] G. Schierholz. ‘Towards a dynamical solution of the strong CP problem’. In: *Nucl. Phys. B Proc. Suppl.* 37.1 (1994). Ed. by Y. Kuno and Y. Okada, pp. 203–210. DOI: 10.1016/0920-5632(94)90751-X. arXiv: hep-lat/9403012.
- [26] M. A. B. Beg and H. -S. Tsao. ‘Strong P, T Noninvariances in a Superweak Theory’. In: *Phys. Rev. Lett.* 41 (1978), p. 278. DOI: 10.1103/PhysRevLett.41.278.
- [27] Howard Georgi. ‘A Model of Soft CP Violation’. In: *Hadronic J.* 1 (1978), p. 155.
- [28] Rabindra N. Mohapatra and G. Senjanovic. ‘Natural Suppression of Strong p and t Noninvariance’. In: *Phys. Lett. B* 79 (1978), pp. 283–286. DOI: 10.1016/0370-2693(78)90243-5.
- [29] S. M. Barr. ‘Solving the Strong CP Problem without the Peccei-Quinn Symmetry’. In: *Phys. Rev. Lett.* 53 (4 July 1984), pp. 329–332. DOI: 10.1103/PhysRevLett.53.329. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.53.329>.
- [30] Ann Nelson. ‘Naturally weak CP violation’. In: *Physics Letters B* 136.5 (1984), pp. 387–391. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(84\)92025-2](https://doi.org/10.1016/0370-2693(84)92025-2). URL: <https://www.sciencedirect.com/science/article/pii/0370269384920252>.
- [31] R. D. Peccei and Helen R. Quinn. ‘CP Conservation in the Presence of Pseudoparticles’. In: *Phys. Rev. Lett.* 38 (25 June 1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.38.1440>.
- [32] R. D. Peccei and Helen R. Quinn. ‘Constraints imposed by CP conservation in the presence of pseudoparticles’. In: *Phys. Rev. D* 16 (6 Sept. 1977), pp. 1791–1797. DOI: 10.1103/PhysRevD.16.1791. URL: <https://link.aps.org/doi/10.1103/PhysRevD.16.1791>.
- [33] Cumrun Vafa and Edward Witten. ‘Parity Conservation in Quantum Chromodynamics’. In: *Phys. Rev. Lett.* 53 (6 Aug. 1984), pp. 535–536. DOI: 10.1103/PhysRevLett.53.535. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.53.535>.
- [34] Steven Weinberg. ‘A New Light Boson?’ In: *Phys. Rev. Lett.* 40 (4 Jan. 1978), pp. 223–226. DOI: 10.1103/PhysRevLett.40.223. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.40.223>.
- [35] Frank Wilczek. ‘Problem of Strong P and T Invariance in the Presence of Instantons’. In: *Phys. Rev. Lett.* 40 (1978), pp. 279–282. DOI: 10.1103/PhysRevLett.40.279.
- [36] Giovanni Grilli di Cortona et al. ‘The QCD axion, precisely’. In: *JHEP* 01 (2016), p. 034. DOI: 10.1007/JHEP01(2016)034. arXiv: 1511.02867 [hep-ph].
- [37] Adel Bilal. ‘Lectures on Anomalies’. In: (Feb. 2008). arXiv: 0802.0634 [hep-th].

- [38] T. W. Donnelly et al. ‘Do axions exist?’ In: *Phys. Rev. D* 18 (5 Sept. 1978), pp. 1607–1620. DOI: 10.1103/PhysRevD.18.1607. URL: <https://link.aps.org/doi/10.1103/PhysRevD.18.1607>.
- [39] Lawrence J. Hall and Mark B. Wise. ‘Flavor changing Higgs boson couplings’. In: *Nuclear Physics B* 187.3 (1981), pp. 397–408. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(81\)90469-7](https://doi.org/10.1016/0550-3213(81)90469-7). URL: <https://www.sciencedirect.com/science/article/pii/0550321381904697>.
- [40] Frank Wilczek. ‘Decays of Heavy Vector Mesons into Higgs Particles’. In: *Phys. Rev. Lett.* 39 (21 Nov. 1977), pp. 1304–1306. DOI: 10.1103/PhysRevLett.39.1304. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.39.1304>.
- [41] William A. Bardeen, R.D. Peccei and T. Yanagida. ‘Constraints on variant axion models’. In: *Nuclear Physics B* 279.3 (1987), pp. 401–428. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(87\)90003-4](https://doi.org/10.1016/0550-3213(87)90003-4). URL: <https://www.sciencedirect.com/science/article/pii/0550321387900034>.
- [42] S. Anderson et al. ‘Improved upper limits on the FCNC decays $B \rightarrow K\ell^+\ell^-$ and $B \rightarrow K^*(892)\ell^+\ell^-$ ’. In: *Phys. Rev. Lett.* 87 (2001), p. 181803. DOI: 10.1103/PhysRevLett.87.181803. arXiv: hep-ex/0106060.
- [43] A. Tumasyan et al. ‘Search for Flavor-Changing Neutral Current Interactions of the Top Quark and Higgs Boson in Final States with Two Photons in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV’. In: *Phys. Rev. Lett.* 129.3 (2022), p. 032001. DOI: 10.1103/PhysRevLett.129.032001. arXiv: 2111.02219 [hep-ex].
- [44] Emmanuel A. Paschos. ‘Diagonal neutral currents’. In: *Phys. Rev. D* 15 (7 Apr. 1977), pp. 1966–1972. DOI: 10.1103/PhysRevD.15.1966. URL: <https://link.aps.org/doi/10.1103/PhysRevD.15.1966>.
- [45] Sheldon L. Glashow and Steven Weinberg. ‘Natural Conservation Laws for Neutral Currents’. In: *Phys. Rev. D* 15 (1977), p. 1958. DOI: 10.1103/PhysRevD.15.1958.
- [46] Mark Srednicki. ‘Axion couplings to matter: (I). CP-conserving parts’. In: *Nuclear Physics B* 260.3 (1985), pp. 689–700. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(85\)90054-9](https://doi.org/10.1016/0550-3213(85)90054-9). URL: <https://www.sciencedirect.com/science/article/pii/0550321385900549>.
- [47] Luca Di Luzio, Federico Mescia and Enrico Nardi. ‘Redefining the Axion Window’. In: *Phys. Rev. Lett.* 118.3 (2017), p. 031801. DOI: 10.1103/PhysRevLett.118.031801. arXiv: 1610.07593 [hep-ph].
- [48] Masahiro Kawasaki, Kazunori Kohri and Takeo Moroi. ‘Big-bang nucleosynthesis and hadronic decay of long-lived massive particles’. In: *Phys. Rev. D* 71 (8 Apr. 2005), p. 083502. DOI: 10.1103/PhysRevD.71.083502. URL: <https://link.aps.org/doi/10.1103/PhysRevD.71.083502>.
- [49] J. Chluba. ‘Distinguishing different scenarios of early energy release with spectral distortions of the cosmic microwave background’. In: *Monthly Notices of the Royal Astronomical Society* 436.3 (Oct. 2013), pp. 2232–2243. ISSN: 0035-8711. DOI: 10.1093/mnras/stt1733. eprint: <https://academic.o>

- up.com/mnras/article-pdf/436/3/2232/4081410/stt1733.pdf. URL: <https://doi.org/10.1093/mnras/stt1733>.
- [50] M. Ackermann et al. ‘Fermi LAT search for dark matter in gamma-ray lines and the inclusive photon spectrum’. In: *Phys. Rev. D* 86 (2 July 2012), p. 022002. DOI: 10.1103/PhysRevD.86.022002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.86.022002>.
- [51] N. Aghanim et al. ‘Planck 2018 results. VI. Cosmological parameters’. In: (2018). arXiv: 1807.06209 [astro-ph.CO].
- [52] Y. Akrami et al. ‘Planck 2018 results. X. Constraints on inflation’. In: *Astron. Astrophys.* 641 (2020), A10. DOI: 10.1051/0004-6361/201833887. arXiv: 1807.06211 [astro-ph.CO].
- [53] Mark P Hertzberg, Max Tegmark and Frank Wilczek. ‘Axion Cosmology and the Energy Scale of Inflation’. In: *Phys. Rev. D* 78 (2008), p. 083507. DOI: 10.1103/PhysRevD.78.083507. arXiv: 0807.1726 [astro-ph].
- [54] Peter Graf and Frank Daniel Steffen. ‘Thermal axion production in the primordial quark-gluon plasma’. In: *Phys. Rev. D* 83 (2011), p. 075011. DOI: 10.1103/PhysRevD.83.075011. arXiv: 1008.4528 [hep-ph].
- [55] Edward W. Kolb and Michael S. Turner. *The Early Universe*. Vol. 69. 1990. ISBN: 978-0-201-62674-2. DOI: 10.1201/9780429492860.
- [56] David J. E. Marsh. ‘Axion Cosmology’. In: *Phys. Rept.* 643 (2016), pp. 1–79. DOI: 10.1016/j.physrep.2016.06.005. arXiv: 1510.07633 [astro-ph.CO].
- [57] Sz. Borsanyi et al. ‘Calculation of the axion mass based on high-temperature lattice quantum chromodynamics’. In: *Nature* 539.7627 (2016), pp. 69–71. DOI: 10.1038/nature20115. arXiv: 1606.07494 [hep-lat].
- [58] Claudio Bonati et al. ‘Axion phenomenology and θ -dependence from $N_f = 2+1$ lattice QCD’. in: *JHEP* 03 (2016), p. 155. DOI: 10.1007/JHEP03(2016)155. arXiv: 1512.06746 [hep-lat].
- [59] Peter Petreczky, Hans-Peter Schadler and Sayantan Sharma. ‘The topological susceptibility in finite temperature QCD and axion cosmology’. In: *Phys. Lett. B* 762 (2016), pp. 498–505. DOI: 10.1016/j.physletb.2016.09.063. arXiv: 1606.03145 [hep-lat].
- [60] Claudio Bonati et al. ‘Topology in full QCD at high temperature: a multicanonical approach’. In: *JHEP* 11 (2018), p. 170. DOI: 10.1007/JHEP11(2018)170. arXiv: 1807.07954 [hep-lat].
- [61] Florian Burger et al. ‘Chiral observables and topology in hot QCD with two families of quarks’. In: *Phys. Rev. D* 98.9 (2018), p. 094501. DOI: 10.1103/PhysRevD.98.094501. arXiv: 1805.06001 [hep-lat].
- [62] Patrick Fox, Aaron Pierce and Scott D. Thomas. ‘Probing a QCD string axion with precision cosmological measurements’. In: (Sept. 2004). arXiv: hep-th/0409059.
- [63] Olivier Wantz and E. P. S. Shellard. ‘Axion cosmology revisited’. In: *Phys. Rev. D* 82 (12 Dec. 2010), p. 123508. DOI: 10.1103/PhysRevD.82.123508. URL: <https://link.aps.org/doi/10.1103/PhysRevD.82.123508>.

- [64] P. Sikivie. ‘Axions, Domain Walls, and the Early Universe’. In: *Phys. Rev. Lett.* 48 (17 Apr. 1982), pp. 1156–1159. DOI: 10.1103/PhysRevLett.48.1156. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.48.1156>.
- [65] A. Vilenkin and E. P. S. Shellard. *Cosmic Strings and Other Topological Defects*. Cambridge University Press, July 2000. ISBN: 978-0-521-65476-0.
- [66] R. A. Battye and E. P. S. Shellard. ‘Global string radiation’. In: *Nucl. Phys. B* 423 (1994), pp. 260–304. DOI: 10.1016/0550-3213(94)90573-8. arXiv: astro-ph/9311017.
- [67] Masahide Yamaguchi, Jun’ichi Yokoyama and M. Kawasaki. ‘Evolution of a global string network in a matter dominated universe’. In: *Phys. Rev. D* 61 (2000), p. 061301. DOI: 10.1103/PhysRevD.61.061301. arXiv: hep-ph/9910352.
- [68] Takashi Hiramatsu et al. ‘Production of dark matter axions from collapse of string-wall systems’. In: *Phys. Rev. D* 85 (2012). [Erratum: Phys.Rev.D 86, 089902 (2012)], p. 105020. DOI: 10.1103/PhysRevD.85.105020. arXiv: 1202.5851 [hep-ph].
- [69] Leesa Fleury and Guy D. Moore. ‘Axion dark matter: strings and their cores’. In: *JCAP* 01 (2016), p. 004. DOI: 10.1088/1475-7516/2016/01/004. arXiv: 1509.00026 [hep-ph].
- [70] Vincent B. Klaer and Guy D. Moore. ‘How to simulate global cosmic strings with large string tension’. In: *JCAP* 10 (2017), p. 043. DOI: 10.1088/1475-7516/2017/10/043. arXiv: 1707.05566 [hep-ph].
- [71] Marco Gorghetto, Edward Hardy and Giovanni Villadoro. ‘Axions from Strings: the Attractive Solution’. In: *JHEP* 07 (2018), p. 151. DOI: 10.1007/JHEP07(2018)151. arXiv: 1806.04677 [hep-ph].
- [72] Alejandro Vaquero, Javier Redondo and Julia Stadler. ‘Early seeds of axion miniclusters’. In: *JCAP* 04 (2019), p. 012. DOI: 10.1088/1475-7516/2019/04/012. arXiv: 1809.09241 [astro-ph.CO].
- [73] Malte Buschmann, Joshua W. Foster and Benjamin R. Safdi. ‘Early-Universe Simulations of the Cosmological Axion’. In: *Phys. Rev. Lett.* 124.16 (2020), p. 161103. DOI: 10.1103/PhysRevLett.124.161103. arXiv: 1906.00967 [astro-ph.CO].
- [74] Ken’ichi Saikawa et al. ‘Spectrum of global string networks and the axion dark matter mass’. In: (Jan. 2024). arXiv: 2401.17253 [hep-ph].
- [75] Kari Enqvist et al. ‘A novel way to determine the scale of inflation’. In: *JCAP* 02 (2018), p. 006. DOI: 10.1088/1475-7516/2018/02/006. arXiv: 1711.07344 [astro-ph.CO].
- [76] Georg G. Raffelt. ‘Astrophysical axion bounds’. In: *Axions: Theory, Cosmology, and Experimental Searches*. Ed. by Markus Kuster, Georg Raffelt and Berta Beltran. Vol. 741. Springer Berlin Heidelberg, 2008, pp. 51–71. DOI: 10.1007/978-3-540-73518-2_3. arXiv: hep-ph/0611350.
- [77] S. J. Asztalos et al. ‘SQUID-Based Microwave Cavity Search for Dark-Matter Axions’. In: *Phys. Rev. Lett.* 104.4, 041301 (Jan. 2010), p. 041301. DOI: 10.1103/PhysRevLett.104.041301. arXiv: 0910.5914 [astro-ph.CO].

- [78] P. Brun et al. ‘A new experimental approach to probe QCD axion dark matter in the mass range above $40 \mu\text{eV}$ ’. in: *Eur. Phys. J. C* 79.3 (2019), p. 186. DOI: 10.1140/epjc/s10052-019-6683-x. arXiv: 1901.07401 [physics.ins-det].
- [79] Sebastian Hoof et al. ‘Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT’. in: *JHEP* 03 (2019). [Erratum: *JHEP* 11, 099 (2019)], p. 191. DOI: 10.1007/JHEP03(2019)191. arXiv: 1810.07192 [hep-ph].
- [80] Luca Visinelli and Javier Redondo. ‘Axion Miniclusters in Modified Cosmological Histories’. In: *Phys. Rev. D* 101.2 (2020), p. 023008. DOI: 10.1103/PhysRevD.101.023008. arXiv: 1808.01879 [astro-ph.CO].
- [81] S. Beurthey et al. ‘MADMAX Status Report’. In: (Mar. 2020). arXiv: 2003.10894 [physics.ins-det].
- [82] Jihn E. Kim and Se-Jin Kim. ‘“Invisible” QCD axion rolling through the QCD phase transition’. In: *Phys. Lett. B* 783 (2018), pp. 357–363. DOI: 10.1016/j.physletb.2018.07.020. arXiv: 1804.05173 [hep-ph].
- [83] B. M. Brubaker et al. ‘HAYSTAC axion search analysis procedure’. In: *Phys. Rev. D* 96.12 (2017), p. 123008. DOI: 10.1103/PhysRevD.96.123008. arXiv: 1706.08388 [astro-ph.IM].
- [84] Jo Bovy et al. ‘The Milky Way’s circular-velocity curve between 4 and 14 kpc from APOGEE data’. In: *The Astrophysical Journal* 759.2 (2012), p. 131.
- [85] Stefan Knirck et al. ‘Directional axion detection’. In: *JCAP* 11 (2018), p. 051. DOI: 10.1088/1475-7516/2018/11/051. arXiv: 1806.05927 [astro-ph.CO].
- [86] Khyati Malhan, Rodrigo A Ibata and Nicolas F Martin. ‘Measuring the Sun’s velocity using Gaia EDR3 observations of Stellar Streams’. In: (2020). arXiv: 2012.05271.
- [87] *Nasa Earth Fact Sheet*. <https://nssdc.gsfc.nasa.gov/planetary/factsheet/earthfact.html>. Mar. 2024.
- [88] *Harry Ransom Centre Website*. <https://www.hrc.utexas.edu/niepc-heliograph/>. July 2024.
- [89] V. Anastassopoulos et al. ‘New CAST Limit on the Axion-Photon Interaction’. In: *Nature Phys.* 13 (2017), pp. 584–590. DOI: 10.1038/nphys4109. arXiv: 1705.02290 [hep-ex].
- [90] E. Armengaud et al. ‘Physics potential of the International Axion Observatory (IAXO)’. in: *JCAP* 2019.6, 047 (June 2019), p. 047. DOI: 10.1088/1475-7516/2019/06/047. arXiv: 1904.09155 [hep-ph].
- [91] Katharina-Sophie Isleif. ‘The Any Light Particle Search Experiment at DESY’. in: *Moscow Univ. Phys. Bull.* 77.2 (2022), pp. 120–125. DOI: 10.3103/S002713492202045X. arXiv: 2202.07306 [hep-ex].
- [92] C. Bartram et al. ‘Search for Invisible Axion Dark Matter in the $3.3\text{--}4.2 \mu\text{eV}$ Mass Range’. In: *Phys. Rev. Lett.* 127.26 (2021), p. 261803. DOI: 10.1103/PhysRevLett.127.261803. arXiv: 2110.06096 [hep-ex].

- [93] L. Zhong et al. ‘Results from phase 1 of the HAYSTAC microwave cavity axion experiment’. In: *Phys. Rev. D* 97.9 (2018), p. 092001. DOI: 10.1103/PhysRevD.97.092001. arXiv: 1803.03690 [hep-ex].
- [94] Alexander J. Millar et al. ‘Dielectric Haloscopes to Search for Axion Dark Matter: Theoretical Foundations’. In: *JCAP* 01 (2017), p. 061. DOI: 10.1088/1475-7516/2017/01/061. arXiv: 1612.07057 [hep-ph].
- [95] P. Sikivie. ‘Experimental Tests of the “Invisible” Axion’. In: *Phys. Rev. Letters* 51.16 (Oct. 1983), pp. 1415–1417. DOI: 10.1103/PhysRevLett.51.1415.
- [96] Frank Wilczek. ‘Two applications of axion electrodynamics’. In: *Phys. Rev. Lett.* 58 (18 May 1987), pp. 1799–1802. DOI: 10.1103/PhysRevLett.58.1799. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.58.1799>.
- [97] Stefan Paul Nikolas Knirck. ‘How To Search for Axion Dark Matter with MADMAX (Magnetized Disk and Mirror Axion eXperiment)’. PhD thesis. Munich, Tech. U., 2020.
- [98] Jesse Liu et al. ‘Broadband Solenoidal Haloscope for Terahertz Axion Detection’. In: *Phys. Rev. Lett.* 128.13 (2022), p. 131801. DOI: 10.1103/PhysRevLett.128.131801. arXiv: 2111.12103 [physics.ins-det].
- [99] Stefan Knirck et al. ‘First Results from a Broadband Search for Dark Photon Dark Matter in the 44 to 52 μeV Range with a Coaxial Dish Antenna’. In: *Phys. Rev. Lett.* 132.13 (2024), p. 131004. DOI: 10.1103/PhysRevLett.132.131004. arXiv: 2310.13891 [hep-ex].
- [100] S. Kotaka et al. ‘Search for Dark Photon Dark Matter in the Mass Range 74 – –110 μeV with a Cryogenic Millimeter-Wave Receiver’. In: *Phys. Rev. Lett.* 130 (7 Feb. 2023), p. 071805. DOI: 10.1103/PhysRevLett.130.071805. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.071805>.
- [101] Dieter Horns et al. ‘Searching for WISPy Cold Dark Matter with a Dish Antenna’. In: *JCAP* 04 (2013), p. 016. DOI: 10.1088/1475-7516/2013/04/016. arXiv: 1212.2970 [hep-ph].
- [102] Jacob Egge. ‘Axion haloscope signal power from reciprocity’. In: *JCAP* 04 (2023), p. 064. DOI: 10.1088/1475-7516/2023/04/064. arXiv: 2211.11503 [hep-ph].
- [103] C.W. Steele. ‘A Nonresonant Perturbation Theory’. In: *IEEE Transactions on Microwave Theory and Techniques* 14.2 (1966), pp. 70–74. DOI: 10.1109/TMTT.1966.1126168.
- [104] J. Egge et al. ‘Experimental determination of axion signal power of dish antennas and dielectric haloscopes using the reciprocity approach’. In: *JCAP* 04 (2024), p. 005. DOI: 10.1088/1475-7516/2024/04/005. arXiv: 2311.13359 [hep-ex].
- [105] Javier De Miguel. ‘A dark matter telescope probing the 6 to 60 GHz band’. In: *JCAP* 04 (2021), p. 075. DOI: 10.1088/1475-7516/2021/04/075. arXiv: 2003.06874 [physics.ins-det].

- [106] Javier De Miguel et al. ‘Discovery prospects with the Dark-photons & Axion-like particles Interferometer’. In: *Phys. Rev. D* 109.6 (2024), p. 062002. DOI: 10.1103/PhysRevD.109.062002. arXiv: 2303.03997 [hep-ph].
- [107] Gray Rybka et al. ‘Search for dark matter axions with the Orpheus experiment’. In: *Phys. Rev. D* 91.1 (2015), p. 011701. DOI: 10.1103/PhysRevD.91.011701. arXiv: 1403.3121 [physics.ins-det].
- [108] R. Cervantes et al. ‘ADMX-Orpheus first search for 70 μeV dark photon dark matter: Detailed design, operations, and analysis’. In: *Phys. Rev. D* 106.10 (2022), p. 102002. DOI: 10.1103/PhysRevD.106.102002. arXiv: 2204.09475 [hep-ex].
- [109] Masha Baryakhtar, Junwu Huang and Robert Lasenby. ‘Axion and hidden photon dark matter detection with multilayer optical haloscopes’. In: *Phys. Rev. D* 98.3 (2018), p. 035006. DOI: 10.1103/PhysRevD.98.035006. arXiv: 1803.11455 [hep-ph].
- [110] A. Torre and L. Quettier. ‘Editorial of the MADMAX Special Section’. In: *IEEE Transactions on Applied Superconductivity* 33.7 (Oct. 2023), p. 1.
- [111] B. Ary Dos Santos Garcia et al. ‘First mechanical realization of a tunable dielectric haloscope for the MADMAX axion search experiment’. In: (July 2024). arXiv: 2407.10716 [physics.ins-det].
- [112] David M. Pozar. *Microwave Engineering, 4th edition*. Wiley, 2012.
- [113] Helio Migon and Dani Gamerman. *Statistical Inference: An Integrated Approach*. Hodder Arnold, 1999.
- [114] C. Bartram et al. ‘Axion dark matter experiment: Run 1B analysis details’. In: *Phys. Rev. D* 103.3 (2021), p. 032002. DOI: 10.1103/PhysRevD.103.032002. arXiv: 2010.06183 [astro-ph.CO].
- [115] Sara Algeri et al. ‘On methods for correcting for the look-elsewhere effect in searches for new physics’. In: *JINST* 11.12 (2016), P12010. DOI: 10.1088/1748-0221/11/12/P12010. arXiv: 1602.03765 [physics.data-an].
- [116] Charles J. Geyer. *Handbook of markov chain monte carlo*. Chapman & Hall, 2011. Chap. Introduction to markov chain monte carlo. ISBN: ISBN 9781420079418.
- [117] Asimina Arvanitaki et al. ‘String Axiverse’. In: *Phys. Rev. D* 81 (2010), p. 123530. DOI: 10.1103/PhysRevD.81.123530. arXiv: 0905.4720 [hep-th].
- [118] Ciaran A. J. O’Hare. ‘Cosmology of axion dark matter’. In: *PoS COSMICWISPerS* (2024), p. 040. DOI: 10.22323/1.454.0040. arXiv: 2403.17697 [hep-ph].
- [119] P. Astone and G. D’Agostini. ‘Inferring the intensity of Poisson processes at the limit of the detector sensitivity (with a case study on gravitational wave burst search)’. In: (Sept. 1999). arXiv: hep-ex/9909047.
- [120] D. A. Palken et al. ‘Improved analysis framework for axion dark matter searches’. In: *Phys. Rev. D* 101.12 (2020), p. 123011. DOI: 10.1103/PhysRevD.101.123011. arXiv: 2003.08510 [astro-ph.IM].

- [121] Anne Ernst, Andreas Ringwald and Carlos Tamarit. ‘Axion Predictions in $SO(10) \times U(1)_{PQ}$ Models’. In: *JHEP* 02 (2018), p. 103. DOI: 10.1007/JHEP02(2018)103. arXiv: 1801.04906 [hep-ph].
- [122] Marco Farina et al. ‘The photo-philic QCD axion’. In: *JHEP* 01 (2017), p. 095. DOI: 10.1007/JHEP01(2017)095. arXiv: 1611.09855 [hep-ph].
- [123] G. R. Dvali and Goran Senjanovic. ‘Is there a domain wall problem?’ In: *Phys. Rev. Lett.* 74 (1995), pp. 5178–5181. DOI: 10.1103/PhysRevLett.74.5178. arXiv: hep-ph/9501387.
- [124] G. Lazarides and Q. Shafi. ‘Axion models with no domain wall problem’. In: *Physics Letters B* 115.1 (1982), pp. 21–25. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(82\)90506-8](https://doi.org/10.1016/0370-2693(82)90506-8). URL: <https://www.sciencedirect.com/science/article/pii/0370269382905068>.
- [125] Antonio Pich. ‘Flavour Physics and CP Violation’. In: *6th CERN-Latin-American School of High-Energy Physics*. 2013, pp. 119–144. DOI: 10.5170/CERN-2013-003.119. arXiv: 1112.4094 [hep-ph].
- [126] Igor P. Ivanov. ‘Building and testing models with extended Higgs sectors’. In: *Prog. Part. Nucl. Phys.* 95 (2017), pp. 160–208. DOI: 10.1016/j.ppnp.2017.03.001. arXiv: 1702.03776 [hep-ph].
- [127] M. Ya. Gogberashvili and G. R. Dvali. ‘Hierarchy at Yukawa constants and K_0 anti- K_0 , B_0 anti- B_0 oscillations in the model with two Higgs doublets. (In Russian)’. In: *Sov. J. Nucl. Phys.* 53 (1991), pp. 491–492.
- [128] Antonio Pich and Paula Tuzon. ‘Yukawa Alignment in the Two-Higgs-Doublet Model’. In: *Phys. Rev. D* 80 (2009), p. 091702. DOI: 10.1103/PhysRevD.80.091702. arXiv: 0908.1554 [hep-ph].
- [129] Ivo de Medeiros Varzielas and Jim Talbert. ‘FCNC-free multi-Higgs-doublet models from broken family symmetries’. In: *Phys. Lett. B* 800 (2020), p. 135091. DOI: 10.1016/j.physletb.2019.135091. arXiv: 1908.10979 [hep-ph].
- [130] T. P. Cheng and Marc Sher. ‘Mass-matrix ansatz and flavor nonconservation in models with multiple Higgs doublets’. In: *Phys. Rev. D* 35 (11 June 1987), pp. 3484–3491. DOI: 10.1103/PhysRevD.35.3484. URL: <https://link.aps.org/doi/10.1103/PhysRevD.35.3484>.
- [131] Sergio Carrolo et al. ‘Symmetry and decoupling in multi-Higgs boson models’. In: *Phys. Rev. D* 103 (7 Apr. 2021), p. 075026. DOI: 10.1103/PhysRevD.103.075026. URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.075026>.
- [132] Jeff Bezanson et al. ‘Julia: A Fast Dynamic Language for Technical Computing’. In: *CoRR* abs/1209.5145 (2012). arXiv: 1209.5145. URL: <http://arxiv.org/abs/1209.5145>.
- [133] *StaticArrays.jl*. <https://juliaarrays.github.io/StaticArrays.jl/>. Version v1.5.6. Aug. 2022.
- [134] Karl Pearson. ‘Mathematical Contributions to the Theory of Evolution. XIX. Second Supplement to a Memoir on Skew Variation’. In: *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a*

- Mathematical or Physical Character* 216 (1916), pp. 429–457. ISSN: 02643952. URL: <http://www.jstor.org/stable/91092> (visited on 31/01/2023).
- [135] J.V. Sloan et al. ‘Limits on axion–photon coupling or on local axion density: Dependence on models of the Milky Way’s dark halo’. In: *Physics of the Dark Universe* 14 (2016), pp. 95–102. ISSN: 2212-6864. DOI: <https://doi.org/10.1016/j.dark.2016.09.003>. URL: <https://www.sciencedirect.com/science/article/pii/S2212686416300504>.
- [136] V. Anastassopoulos et al. ‘Towards a medium-scale axion helioscope and haloscope’. In: *Journal of Instrumentation* 12 (June 2017). DOI: [10.1088/1748-0221/12/11/P11019](https://doi.org/10.1088/1748-0221/12/11/P11019).
- [137] S. Scopel. ‘Particle Dark Matter Candidates’. In: *Journal of Physics: Conference Series* 120 (July 2008), p. 042003. DOI: [10.1088/1742-6596/120/4/042003](https://doi.org/10.1088/1742-6596/120/4/042003).
- [138] I. Shilon et al. ‘Conceptual Design of a New Large Superconducting Toroid for IAXO, the New International AXion Observatory’. In: *IEEE Trans. Appl. Supercond.* 23.3 (2013). Ed. by Steve Gourlay, Gian Luca Sabbi and Lance Cooley, p. 4500604. DOI: [10.1109/TASC.2013.2251052](https://doi.org/10.1109/TASC.2013.2251052). arXiv: 1212.4633 [physics.ins-det].
- [139] Ben T. McAllister et al. ‘The ORGAN Experiment: An axion haloscope above 15 GHz’. In: *Phys. Dark Univ.* 18 (2017), pp. 67–72. DOI: [10.1016/j.dark.2017.09.010](https://doi.org/10.1016/j.dark.2017.09.010). arXiv: 1706.00209 [physics.ins-det].
- [140] Junu Jeong et al. ‘Search for Invisible Axion Dark Matter with a Multiple-Cell Haloscope’. In: *Phys. Rev. Lett.* 125.22 (2020), p. 221302. DOI: [10.1103/PhysRevLett.125.221302](https://doi.org/10.1103/PhysRevLett.125.221302). arXiv: 2008.10141 [hep-ex].
- [141] Alexander V. Gramolin et al. ‘Search for axion-like dark matter with ferromagnets’. In: *Nature Phys.* 17.1 (2021), pp. 79–84. DOI: [10.1038/s41567-020-1006-6](https://doi.org/10.1038/s41567-020-1006-6). arXiv: 2003.03348 [hep-ex].
- [142] Chiara P. Salemi et al. ‘Search for Low-Mass Axion Dark Matter with ABRACADABRA-10 cm’. In: *Phys. Rev. Lett.* 127.8 (2021), p. 081801. DOI: [10.1103/PhysRevLett.127.081801](https://doi.org/10.1103/PhysRevLett.127.081801). arXiv: 2102.06722 [hep-ex].
- [143] Jonathan L. Ouellet et al. ‘First Results from ABRACADABRA-10 cm: A Search for Sub- μeV Axion Dark Matter’. In: *Phys. Rev. Lett.* 122.12 (2019), p. 121802. DOI: [10.1103/PhysRevLett.122.121802](https://doi.org/10.1103/PhysRevLett.122.121802). arXiv: 1810.12257 [hep-ex].
- [144] K. M. Backes et al. ‘A quantum-enhanced search for dark matter axions’. In: *Nature* 590.7845 (2021), pp. 238–242. DOI: [10.1038/s41586-021-03226-7](https://doi.org/10.1038/s41586-021-03226-7). arXiv: 2008.01853 [quant-ph].
- [145] C Haggmann et al. ‘Results from a search for cosmic axions’. In: *Physical Review D* 42.4 (1990), p. 1297.
- [146] S DePanfilis et al. ‘Limits on the abundance and coupling of cosmic axions at $4.5 < m_a < 5.0 \mu\text{eV}$ ’. in: *Physical Review Letters* 59.7 (1987), p. 839.
- [147] N. Du et al. ‘A Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment’. In: *Phys. Rev. Lett.* 120.15 (2018), p. 151301. DOI: [10.1103/PhysRevLett.120.151301](https://doi.org/10.1103/PhysRevLett.120.151301). arXiv: 1804.05750 [hep-ex].

- [148] Thierry Grenet et al. ‘The Grenoble Axion Haloscope platform (GrAHal): development plan and first results’. In: (Oct. 2021). arXiv: 2110.14406 [hep-ex].
- [149] T. Braine et al. ‘Extended Search for the Invisible Axion with the Axion Dark Matter Experiment’. In: *Phys. Rev. Lett.* 124.10 (2020), p. 101303. DOI: 10.1103/PhysRevLett.124.101303. arXiv: 1910.08638 [hep-ex].
- [150] Yuta Michimura et al. ‘DANCE: Dark matter Axion search with riNg Cavity Experiment’. In: *J. Phys. Conf. Ser.* 1468.1 (2020). Ed. by Masayuki Nakahata, p. 012032. DOI: 10.1088/1742-6596/1468/1/012032. arXiv: 1911.05196 [physics.ins-det].
- [151] Beatriz Aja et al. ‘The Canfranc Axion Detection Experiment (CADEx): search for axions at 90 GHz with Kinetic Inductance Detectors’. In: *JCAP* 11 (2022), p. 044. DOI: 10.1088/1475-7516/2022/11/044. arXiv: 2206.02980 [hep-ex].
- [152] Alexander J. Millar et al. ‘ALPHA: Searching For Dark Matter with Plasma Haloscopes’. In: (Sept. 2022). arXiv: 2210.00017 [hep-ph].
- [153] Matthew Lawson et al. ‘Tunable axion plasma haloscopes’. In: *Phys. Rev. Lett.* 123.14 (2019), p. 141802. DOI: 10.1103/PhysRevLett.123.141802. arXiv: 1904.11872 [hep-ph].
- [154] Hsin Chang et al. ‘First Results from the Taiwan Axion Search Experiment with a Haloscope at 19.6 μeV ’. in: *Phys. Rev. Lett.* 129.11 (2022), p. 111802. DOI: 10.1103/PhysRevLett.129.111802. arXiv: 2205.05574 [hep-ex].
- [155] D. Alesini et al. ‘Search for Galactic axions with a high-Q dielectric cavity’. In: *Phys. Rev. D* 106.5, 052007 (Sept. 2022), p. 052007. DOI: 10.1103/PhysRevD.106.052007. arXiv: 2208.12670 [hep-ex].
- [156] A. Álvarez Melcón et al. ‘First results of the CAST-RADES haloscope search for axions at 34.67 μeV ’. in: *JHEP* 21 (2020), p. 075. DOI: 10.1007/JHEP10(2021)075. arXiv: 2104.13798 [hep-ex].
- [157] Aaron P. Quiskamp et al. ‘Direct Search for Dark Matter Axions Excluding ALP Cogenesis in the 63-67 micro-eV Range, with The ORGAN Experiment’. In: *arXiv e-prints*, arXiv:2203.12152 (Mar. 2022), arXiv:2203.12152. DOI: 10.48550/arXiv.2203.12152. arXiv: 2203.12152 [hep-ex].
- [158] M. J. Jewell et al. ‘New results from HAYSTAC’s phase II operation with a squeezed state receiver’. In: *Phys. Rev. D* 107.7 (2023), p. 072007. DOI: 10.1103/PhysRevD.107.072007. arXiv: 2301.09721 [hep-ex].
- [159] S. Lee et al. ‘Axion Dark Matter Search around 6.7 μeV ’. in: *Phys. Rev. Lett.* 124.10, 101802 (Mar. 2020), p. 101802. DOI: 10.1103/PhysRevLett.124.101802. arXiv: 2001.05102 [hep-ex].
- [160] Ohjoon Kwon et al. ‘First Results from an Axion Haloscope at CAPP around 10.7 μeV ’. in: *Phys. Rev. Lett.* 126.19, 191802 (May 2021), p. 191802. DOI: 10.1103/PhysRevLett.126.191802. arXiv: 2012.10764 [hep-ex].
- [161] Youngjae Lee et al. ‘Searching for Invisible Axion Dark Matter with an 18 T Magnet Haloscope’. In: *Phys. Rev. Lett.* 128.24, 241805 (June 2022),

- p. 241805. DOI: 10.1103/PhysRevLett.128.241805. arXiv: 2206.08845 [hep-ex].
- [162] Jinsu Kim et al. ‘Near-Quantum-Noise Axion Dark Matter Search at CAPP around $9.5 \mu\text{eV}$ ’. in: *arXiv e-prints*, arXiv:2207.13597 (July 2022), arXiv:2207.13597. DOI: 10.48550/arXiv.2207.13597. arXiv: 2207.13597 [hep-ex].
- [163] Andrew K. Yi et al. ‘Axion Dark Matter Search around $4.55 \mu\text{eV}$ with Dine-Fischler-Srednicki-Zhitnitskii Sensitivity’. In: *arXiv e-prints*, arXiv:2210.10961 (Oct. 2022), arXiv:2210.10961. DOI: 10.48550/arXiv.2210.10961. arXiv: 2210.10961 [hep-ex].
- [164] Asher Berlin et al. ‘Heterodyne broadband detection of axion dark matter’. In: *Phys. Rev. D* 104.11 (2021), p. L111701. DOI: 10.1103/PhysRevD.104.L11701. arXiv: 2007.15656 [hep-ph].
- [165] L. Brouwer et al. ‘Projected sensitivity of DMRadio-m3: A search for the QCD axion below $1 \mu\text{eV}$ ’. in: *Phys. Rev. D* 106.10 (2022), p. 103008. DOI: 10.1103/PhysRevD.106.103008. arXiv: 2204.13781 [hep-ex].
- [166] David Alesini et al. ‘The KLASH Proposal’. In: (July 2017). arXiv: 1707.06010 [physics.ins-det].
- [167] N. Crisosto et al. ‘ADMX SLIC: Results from a Superconducting LC Circuit Investigating Cold Axions’. In: *Phys. Rev. Lett.* 124.24 (2020), p. 241101. DOI: 10.1103/PhysRevLett.124.241101. arXiv: 1911.05772 [astro-ph.CO].
- [168] Zhongyue Zhang, Dieter Horns and Oindrila Ghosh. ‘Search for dark matter with an LC circuit’. In: *Phys. Rev. D* 106.2 (2022), p. 023003. DOI: 10.1103/PhysRevD.106.023003. arXiv: 2111.04541 [hep-ex].
- [169] Koji Nagano et al. ‘Axion Dark Matter Search with Interferometric Gravitational Wave Detectors’. In: *Phys. Rev. Lett.* 123.11 (2019), p. 111301. DOI: 10.1103/PhysRevLett.123.111301. arXiv: 1903.02017 [hep-ph].
- [170] Hongwan Liu et al. ‘Searching for Axion Dark Matter with Birefringent Cavities’. In: *Phys. Rev. D* 100.2 (2019), p. 023548. DOI: 10.1103/PhysRevD.100.023548. arXiv: 1809.01656 [hep-ph].
- [171] Jan Schütte-Engel et al. ‘Axion quasiparticles for axion dark matter detection’. In: *JCAP* 08 (2021), p. 066. DOI: 10.1088/1475-7516/2021/08/066. arXiv: 2102.05366 [hep-ph].
- [172] David J. E. Marsh et al. ‘Proposal to Detect Dark Matter using Axionic Topological Antiferromagnets’. In: *Phys. Rev. Lett.* 123.12 (2019), p. 121601. DOI: 10.1103/PhysRevLett.123.121601. arXiv: 1807.08810 [hep-ph].
- [173] Jack A. Devlin et al. ‘Constraints on the Coupling between Axionlike Dark Matter and Photons Using an Antiproton Superconducting Tuned Detection Circuit in a Cryogenic Penning Trap’. In: *Phys. Rev. Lett.* 126.4 (2021), p. 041301. DOI: 10.1103/PhysRevLett.126.041301. arXiv: 2101.11290 [astro-ph.CO].
- [174] N. Crescini et al. ‘Axion search with a quantum-limited ferromagnetic haloscope’. In: *Phys. Rev. Lett.* 124.17 (2020), p. 171801. DOI: 10.1103/PhysRevLett.124.171801. arXiv: 2001.08940 [hep-ex].

- [175] D. Alesini et al. ‘Galactic axions search with a superconducting resonant cavity’. In: *Phys. Rev. D* 99.10 (2019), p. 101101. DOI: 10.1103/PhysRevD.99.101101. arXiv: 1903.06547 [physics.ins-det].
- [176] D. Alesini et al. ‘Search for invisible axion dark matter of mass $m_a = 43 \mu\text{eV}$ with the QUAX- $a\gamma$ experiment’. In: *Phys. Rev. D* 103.10 (2021), p. 102004. DOI: 10.1103/PhysRevD.103.102004. arXiv: 2012.09498 [hep-ex].
- [177] Catriona A. Thomson et al. ‘Upconversion Loop Oscillator Axion Detection Experiment: A Precision Frequency Interferometric Axion Dark Matter Search with a Cylindrical Microwave Cavity’. In: *Phys. Rev. Lett.* 126.8 (2021). [Erratum: *Phys.Rev.Lett.* 127, 019901 (2021)], p. 081803. DOI: 10.1103/PhysRevLett.127.019901. arXiv: 1912.07751 [hep-ex].
- [178] Ciaran O’Hare. *cajohare/AxionLimits: AxionLimits*. <https://cajohare.github.io/AxionLimits/>. Version v1.0. Jan. 2023. DOI: 10.5281/zenodo.3932430.
- [179] Frederik Beaujean, Allen Caldwell and Olaf Reimann. ‘Is the bump significant? An axion-search example’. In: *Eur. Phys. J. C* 78.9 (2018), p. 793. DOI: 10.1140/epjc/s10052-018-6217-y. arXiv: 1710.06642 [hep-ex].
- [180] A. Savitzky and M. J. E. Golay. ‘Smoothing and differentiation of data by simplified least squares procedures’. In: *Analytical Chemistry* 36 (Jan. 1964), pp. 1627–1639. DOI: 10.1021/ac60214a047.
- [181] Johannes Buchner. ‘A statistical test for Nested Sampling algorithms’. In: *Statistics and Computing* 26.1-2 (Jan. 2016), pp. 383–392. DOI: 10.1007/s11222-014-9512-y.
- [182] Johannes Buchner. ‘Collaborative Nested Sampling: Big Data versus Complex Physical Models’. In: *Publications of the Astronomical Society of the Pacific* 131.1004 (Oct. 2019), p. 108005. DOI: 10.1088/1538-3873/aae7fc.
- [183] Johannes Buchner. ‘UltraNest - a robust, general purpose Bayesian inference engine’. In: *The Journal of Open Source Software* 6.60, 3001 (Apr. 2021), p. 3001. DOI: 10.21105/joss.03001. arXiv: 2101.09604 [stat.CO].
- [184] Oliver Schulz et al. ‘BAT.jl: A Julia-Based Tool for Bayesian Inference’. In: *SN Computer Science* 2.3 (Apr. 2021), p. 210. ISSN: 2661-8907. DOI: 10.1007/s42979-021-00626-4. URL: <https://doi.org/10.1007/s42979-021-00626-4>.
- [185] Frederik Beaujean and Allen Caldwell. ‘A test statistic for weighted runs’. In: *Journal of Statistical Planning and Inference* 141.11 (2011), pp. 3437–3446. ISSN: 0378-3758. DOI: <https://doi.org/10.1016/j.jspi.2011.04.022>. URL: <https://www.sciencedirect.com/science/article/pii/S0378375811001935>.
- [186] K. Altenmüller et al. ‘A new upper limit on the axion-photon coupling with an extended CAST run with a Xe-based Micromegas detector’. In: (June 2024). arXiv: 2406.16840 [hep-ex].