

ANTOINE GARCON

Searching for Bosonic Dark Matter
with
Nuclear Magnetic Resonance



JOHANNES GUTENBERG
UNIVERSITÄT MAINZ

Submitted on 25 August 2021
for the degree of Doctor of Philosophy

Department of Physics,
Johannes Gutenberg University - Mainz, Germany

Supervisor:
Prof. Dmitry Budker

Declaration of authorship

I, Antoine Garcon, hereby testify that the work presented in this thesis is my own. The work was mainly performed at the Helmholtz Institute - Mainz, under the supervision of Prof. Dmitry Budker. Prof. Stefan Kramer supervised the work presented on the last chapter of this thesis.

The work presented in this thesis mainly consists of the following articles of which I was the lead author:

- Antoine Garcon, Deniz Aybas, John W Blanchard, Gary Centers, Nataniel L Figueroa, Peter W Graham, Derek F Jackson Kimball, Surjeet Rajendran, Marina Gil Sendra, Alexander O Sushkov, Lutz Trahms, Tao Wang, Arne Wickenbrock, Teng Wu, Dmitry Budker. **The Cosmic Axion Spin Precession Experiment (CASPER): a Dark-Matter Search with Nuclear Magnetic Resonance.** *Quantum Science and Technology*, Vol. 3, no. 1, 2017. IOP Publishing.
- Antoine Garcon, John W Blanchard, Gary P Centers, Nataniel L Figueroa, Peter W Graham, Derek F Jackson Kimball, Surjeet Rajendran, Alexander O Sushkov, Yevgeny V Stadnik, Arne Wickenbrock, Teng Wu, Dmitry Budker. **Constraints on Bosonic Dark Matter from Ultralow-Field Nuclear Magnetic Resonance.** *Science Advances*, Vol. 5, no. 10, 2019.
- Antoine Garcon, Julian Vexler, Dmitry Budker, Stefan Kramer. **Deep Neural Networks to Recover Unknown Physical Parameters from Oscillating Time Series.** *preprint arXiv:2101.03850*, 2021.

For simplicity, ease of read, and to avoid repetitions, the articles have been modified or reorganized and some sections have been entirely rewritten.

Abstract

The nature of dark matter, the invisible substance making up over 80% of the matter in the Universe, is one of the most fundamental mysteries of modern physics. Ultralight bosons such as axions, axion-like particles or dark photons could make up most of the dark matter. Couplings between such bosons and nuclear spins may enable their direct detection via nuclear magnetic resonance (NMR) spectroscopy: as nuclear spins move through the galactic dark-matter halo, they couple to dark-matter and behave as if they were in an oscillating magnetic field, generating a dark-matter-driven NMR signal.

In the first chapter of this thesis, we review the predicted couplings of axions and axion-like particles with baryonic matter that enable their detection via NMR. We then describe two measurement schemes being implemented in the Cosmic Axion Spin Precession Experiment (CASPEr), an NMR experiment seeking to detect axion and axion-like particles. The first method, presented in the original CASPEr proposal, consists of a resonant search via continuous-wave NMR spectroscopy. This method offers the highest sensitivity for frequencies ranging from a few Hz to hundreds of MHz, corresponding to masses $m_a \sim 10^{-14}$ – 10^{-6} eV. However, Sub-Hz frequencies are typically difficult to probe with NMR due to the diminishing sensitivity of magnetometers in this region. To circumvent this limitation, we suggest new detection and data processing modalities: a non-resonant frequency-modulation detection scheme, enabling searches from mHz to Hz frequencies ($m_a \sim 10^{-17}$ – 10^{-14} eV).

As a second part of this thesis, we apply the above mentioned non-resonant method and use ultralow-field NMR to probe the axion-fermion “wind” coupling and dark-photon couplings to nuclear spins. No dark matter signal was detected above background, establishing new experimental bounds for dark-matter bosons with masses ranging from 1.8×10^{-16} to 7.8×10^{-14} eV.

In the last chapter of this thesis, we use Deep Neural Networks (DNNs) to disentangle components of oscillating time series, arguably the most

common form of signals acquired during dark-matter searches. We show that the regression and denoising performance is similar to those of least-square curve fittings (LS-fit). We then explore various applications in which we believe our architecture could prove useful for time-series processing, when prior knowledge is incomplete. Because the *Autoencoder* needs no prior information about the physical model, the remaining unknown latent parameters can still be captured, thus making use of partial prior knowledge, while leaving space for data exploration and discoveries.

Acknowledgments

First and foremost I would like to deeply thank my loving parents, stepparents, brother, and stepsister.

My gratitude goes to Evelyne and Christian Clinet, who I have been lucky enough to meet and whose impact on me is still immeasurable.

I would like to thank my friends and teachers from Mainz; Dima Budker whose funny obsession for the the drunken sailor always turned out fruitful, John Blanchard our local sailor and fantastic mentor, and Stefan Kramer for his warm welcome in his group.

Finally, all my gratitude extends to Dionysis Antypas, Gary Centers, Georgios Chatzidrosos, Martin Engler, Mahapan Peter Leyser, and Jim Visschers for being fantastic friends.

Antoine Garcon,
Paris,
July 2021.

Related publications

The following is the list of publications in which I participated, some of which are still in peer review process.

- Antoine Garcon, Deniz Aybas, John W Blanchard, Gary Centers, Nataniel L Figueroa, Peter W Graham, Derek F Jackson Kimball, Surjeet Rajendran, Marina Gil Sendra, Alexander O Sushkov, Lutz Trahms, Tao Wang, Arne Wickenbrock, Teng Wu, Dmitry Budker. **The Cosmic Axion Spin Precession Experiment (CASPER): a Dark-Matter Search With Nuclear Magnetic Resonance.** *Quantum Science and Technology*, Vol. 3, no. 1, 2017. IOP Publishing.
- Antoine Garcon, John W Blanchard, Gary P Centers, Nataniel L Figueroa, Peter W Graham, Derek F Jackson Kimball, Surjeet Rajendran, Alexander O Sushkov, Yevgeny V Stadnik, Arne Wickenbrock, Teng Wu, Dmitry Budker. **Constraints on Bosonic Dark Matter From Ultralow-Field Nuclear Magnetic Resonance.** *Science Advances*, Vol. 5, no. 10, 2019.
- Gary P Centers, John W Blanchard, Jan Conrad, Nataniel L Figueroa, Antoine Garcon, Alexander V Gramolin, Derek F Jackson Kimball, Matthew Lawson, Bart Pelssers, Joseph A Smiga, Yevgeny Stadnik, Alexander O Sushkov, Arne Wickenbrock, Dmitry Budker, Andrei Derevianko. **Stochastic Amplitude Fluctuations of Bosonic Dark Matter and Revised Constraints on Linear Couplings.** *preprint arXiv:1905.13650*, 2019
- Teng Wu, John W Blanchard, Gary P Centers, Nataniel L Figueroa, Antoine Garcon, Peter W Graham, Derek F Jackson Kimball, Surjeet Rajendran, Yevgeny V Stadnik,

Alexander O Sushkov, Arne Wickenbrock, Dmitry Budker. **Search for Axionlike Dark Matter With a Liquid-State Nuclear Spin Comagnetometer.** *Physical review letters*, Vol. 122, no. 19, 2019. American Physical Society.

- D Antypas, O Tretiak, [A Garcon](#), R Ozeri, G Perez, D Budker. **Scalar Dark Matter in the Radio-Frequency Band: Atomic-Spectroscopy Search Results.** *Physical review letters*, Vol. 123, no. 14, 2019. American Physical Society.
- Derek F Jackson Kimball, S Afach, D Aybas, JW Blanchard, D Budker, G Centers, M Engler, NL Figueroa, [A Garcon](#), PW Graham, H Luo, S Rajendran, MG Sendra, AO Sushkov, T Wang, A Wickenbrock, A Wilzewski, T Wu. **Overview of the Cosmic Axion Spin Precession Experiment (CASPER).** *Book: Microwave Cavities and Detectors for Axion Research*, pages 105-121, 2020. Springer, Cham.
- Dudari B Burueva, James Eills, John W Blanchard, [Antoine Garcon](#), Román Picazo-Frutos, Kirill V Kovtunov, Igor V Koptug, Dmitry Budker. **Chemical Reaction Monitoring Using Zero-Field Nuclear Magnetic Resonance Enables Study of Heterogeneous Samples in Metal Containers.** *Angewandte Chemie International Edition*, Vol. 59, no. 39, 2020.
- [Antoine Garcon](#), Julian Vexler, Dmitry Budker, Stefan Kramer. **Deep Neural Networks to Recover Unknown Physical Parameters from Oscillating Time Series.** *preprint arXiv:2101.03850*, 2021.
- Min Jiang, Haowen Su, [Antoine Garcon](#), Xinhua Peng, Dmitry Budker. **Search for Axion-Like Dark Matter with Spin-Based Amplifiers.** *preprint arXiv:2102.01448*, 2021.
- Dionysios Antypas, Oleg Tretiak, Ke Zhang, [Antoine Garcon](#), Gilad Perez, Mikhail G Kozlov, Stephan Schiller, Dmitry Budke **Probing Fast Oscillating Scalar Dark Matter with Atoms and Molecules.** *Quantum Science and Technology*, Vol. 6, no. 3, 2021. IOP Publishing.

- Deniz Aybas, Janos Adam, Emmy Blumenthal, Alexander V Gramolin, Dorian Johnson, Annalies Kleyheeg, Samer Afach, John W Blanchard, Gary P Centers, Antoine Garcon, Martin Engler, Nataniel L Figueroa, Marina Gil Sendra, Arne Wickenbrock, Matthew Lawson, Tao Wang, Teng Wu, Haosu Luo, Hamdi Mani, Philip Mauskopf, Peter W Graham, Surjeet Rajendran, Derek F Jackson Kimball, Dmitry Budker, Alexander O Sushkov. **Search for Axionlike Dark Matter Using Solid-State Nuclear Magnetic Resonance.** *Physical Review Letters*, Vol. 126, no. 14, 2021. American Physical Society.

Contents

List of figures	xv
1 Introduction	1
2 The axion-induced spin precession scheme	7
2.1 Foreword	7
2.2 Axion- and ALP-induced nuclear spin precession	8
2.2.1 ALP-nucleon coupling - CASPEr-Gradient	8
2.2.2 Axion-gluon coupling - CASPEr-Electric	13
2.3 Experimental sensitivity	15
2.4 Ultralow frequencies: sidebands detection	17
2.5 Discussion	23
3 Experimental constraints on bosonic dark matter	27
3.1 Foreword	27
3.2 Experimental scheme	28
3.2.1 Experimental setup overview	28
3.2.2 Measurement schemes	30
3.3 Bosonic dark-matter fields	34
3.3.1 Dark-matter field properties	34
3.3.2 Dark-matter couplings to nuclear spins	36
3.3.3 Dark-matter signatures in zero- to ultralow-field NMR	37
3.3.4 Bosonic dark-matter effective fields	40
3.4 Results	43
3.4.1 Signal processing: post-processing phase cycling	43
3.4.2 Calibration	45
3.4.3 Search and analysis	46
3.4.4 CASPEr-ZULF search results: constraints on bosonic dark matter	47
3.5 Discussion	48

4	Deep neural networks for physics-signal analysis	53
4.1	Foreword	53
4.2	Data description and generation procedure	59
4.3	DNN architecture and training procedure	62
4.4	Results	69
4.4.1	Post-training performance evaluation	69
4.4.2	DNN-assisted LS-fit	70
4.4.3	Partial information regression and denoising	72
4.5	Discussion	76
5	Summary and outlook	79
5.1	Summary and personal notes	79
5.2	Outlook	80
5.2.1	General improvements	80
5.2.2	Complementary Searches	81
5.2.3	Parahydrogen hyperpolarization	82
5.2.4	Search for axions with spin-based amplifiers	83
5.2.5	Further applications of machine learning	84
	Appendices	89
A	Appendix	89
A.1	Chapter 3	89
A.1.1	Dark matter effective fields	89
A.1.2	Signal processing: post-processing phase-cycling	95
A.1.3	Sideband amplitude determination	97
A.1.4	Calibration: signal scaling versus bosonic-field amplitude and frequency	98
A.1.5	Detection threshold determination	98
A.1.6	Coherent averaging: signal scaling with integration time	104
A.1.7	Bandwidth: accessible bosonic mass range	104
A.1.8	Dark-matter-field directionality	107
A.2	Chapter 4	113
A.2.1	DNN architecture implementation	113
A.3	Supplementary figures	116

List of Figures

2.1	ALP-nucleon coupling parameter space	10
2.2	Resonant and sidebands experimental schemes	14
2.3	Resonant versus sidebands detection	24
3.1	Nuclear spin energy levels and NMR spectra of ^{13}C -formic acid measured in three different field conditions	30
3.2	CASPEr-ZULF signal acquisition and processing scheme	38
3.3	ALP wind-nucleon linear coupling exclusions	40
3.4	ALP wind-nucleon quadratic coupling exclusions	49
3.5	Dark photon-nucleon dEDM coupling exclusions	50
4.1	Examples of frequency-modulated sine wave (FM) synthetic time series	58
4.2	Unified DNN architecture and loss description	63
4.3	FM-sine waves test-data prediction errors	66
4.4	Comparison of DNN post-training performance to LS-fits	68
4.5	DNN latent-parameters predictions used as initial guesses for DNN-assisted fits	71
4.6	Performance comparison of the specialized DNN	73
A.1	CASPEr-ZULF calibration data	99
A.2	Example of empirical white-noise power-spectral-density cu- mulative distribution function	103
A.3	Sidebands amplitude A_s , detection threshold A_{th} at 90% confidence level (90% CL)	105
A.4	Sideband amplitude scaling with calibration-field frequency	106
A.5	Example of DNN prediction using a noisy AM-sine wave input	116
A.6	Example of sine wave and AM-sine wave samples	117
A.7	Comparison of DNN post-training performance to LS-fits	118
A.8	Validation loss during training for various values of β	119

Introduction

In 1945 Edward M. Purcell measured the first radio-frequency absorption from nuclear magnetic moments in paraffin [1]. In the following months, Felix Bloch observed nuclear spin precession in water [2]. Subsequently, Bloch's nuclear magnetic resonance (NMR) techniques showed that electrons provide magnetic shielding to the nucleus. The resulting change in magnetic-resonance frequency, known as the chemical shift, provides information on the electronic environment of nuclear spins. The discovery of this phenomenon enabled NMR-based chemical analysis [3] and the field of NMR quickly grew to become a dominant tool in analytical chemistry, medicine and structural biology. NMR also remains at the forefront of fundamental physics, in fields such as materials science, precision magnetometry, quantum control [4, 5] and in searches for exotic forces.

Prior to Purcell's work, the existence of dark matter was postulated in 1933 by Fritz Zwicky to explain the dynamics of galaxies within galaxy clusters. Zwicky discovered that the amount of visible matter in the clusters could not account for the galaxies' velocities and postulated the presence of some invisible, "dunkle" (dark) matter [6]. Armed only with astrophysical observations of the early 20th century and the Virial theorem, Zwicky postulated that if dark matter exists it should be "in much greater density than luminous matter". Modern observations now show that dark matter

composes more than 80% of the matter content of the Universe [7], possibly confirming Zwicky's discovery.

Elucidating the nature of dark matter will profoundly impact our understanding of cosmology, astrophysics, and particle physics, providing insights into the evolution of the Universe and potentially uncovering new physical laws and fundamental forces beyond the Standard Model [8]. While the observational evidences for dark matter are still derived from its gravitational effects at the galactic scale and larger, the key to solving the mystery of its nature lies in directly measuring non-gravitational interactions of dark matter with Standard Model particles and fields. To date, experimental efforts to directly detect dark matter have largely focused on Weakly Interacting Massive Particles (WIMPs), with masses between 10 and 1000 GeV [9, 10]. Despite considerable efforts, there have been no conclusive signs of WIMP interactions with ordinary matter. A multitude of other particles were introduced as possible candidates, but as of today, none of them have been detected and the nature of dark matter remains unknown.

The Standard Model predicts that the strong force could violate the charge conjugation parity symmetry (CP-symmetry). This effect has never been observed and experimental measurements constrain strong-CP violation to an extremely low value [11]. The need for this fine tuning is known as the strong CP problem. In 1977, Roberto Peccei and Helen Quinn introduced a mechanism potentially solving the strong CP problem [12], from which the axion naturally emerges as a bosonic particle [13, 14].

The existence of axions may help to answer other open questions in particle physics such as why the strong force respects the combined charge-

conjugation and parity-inversion (CP) symmetry to such a high degree [15], the relative weakness of the gravitational interaction [16], and how to unify the theories of quantum mechanics and general relativity [17].

In addition to Peccei and Quinn theory, a wide variety of theories predict the existence of new spin-0 bosons such as axion-like particles (ALPs) as well as spin-1 bosons such as dark photons [7, 18]. The absence of evidence for WIMPs has reinvigorated efforts to search for these ultralight bosonic fields, as dark matter candidates [19], composed of bosons with masses smaller than a few eV. Indeed, axions and ALPs should interact only weakly with particles of the Standard Model, making them “dark” and they could account for all of the dark matter density [20–22]. In addition, axions and ALPs could form structures similar to the dark-matter galactic clusters [23, 24].

In addition to interacting via gravity, axions and ALPs are predicted to have a weak coupling to the electromagnetic field, enabling their conversion to photons via the inverse-Primakoff effect [25, 26]. Past and current axion searches largely focus on detecting photons produced by this coupling. These searches include helioscopes, detectors pointed at the Sun, such as the “CERN Axion Solar Telescope” (CAST [27]). Helioscope detection could happen when axions produced in the Sun are converted back to photons in a strong laboratory magnetic field [28].

“Light-Shining-Through-Walls” experiments, such as the “Any Light Particle Search” (ALPS [29]), do not rely on astrophysical-axion sources. These experiments seek to convert laser-sourced photons into axions in a strong magnetic field. Subsequently, axions travel through a wall and are converted back to photons via the same mechanism.

Other searches include haloscopes, aimed at detecting local axions in the Milky Way’s dark-matter halo. These experiments include microwave cavity-enhancement methods, such as the “Axion Dark Matter Experiment” (ADMX [30, 31]), which could convert local axions to photons in a high-Q cavity. More information on the wide array of experimental searches for axion and ALP dark matter can be found in Ref. [32].

While all the aforementioned experiments seek to detect ALPs through their conversion to photons, the possibility of direct ALP detection via their couplings to nucleons and gluons has been recently proposed [33, 34]. These couplings give rise to oscillatory pseudo-magnetic interactions with the dark-matter ALP field. This dark-matter field oscillates at the ALP Compton frequency, which is proportional to the ALP mass. Therefore, another method to search for dark-matter bosonic fields was recently proposed: dark-matter-driven spin-precession [35–37], detected via NMR techniques [36–40]. These concepts were recently applied to data measuring the permanent electric dipole moment of the neutron and successfully constrained ALP dark matter with masses $\lesssim 10^{-17}$ eV [40].

The Cosmic Axion Spin Precession Experiment (CASPER [33]) is a haloscope, seeking to detect the NMR signal induced by these couplings. The CASPER collaboration is composed of two main groups, each searching for axions and ALPs via different couplings: CASPER-Gradient (formally CASPER-Wind) is based on the pseudo-magnetic coupling of ALPs with nucleons and CASPER-Electric is sensitive to the axion-gluon coupling. The two experiments rely on different couplings but are otherwise similar, in the sense that they both measure ALP-induced nuclear-spin precession.

Efforts within this program utilizing zero- to ultralow-field (ZULF) NMR spectroscopy [41] are collectively referred to as CASPEr-ZULF.

In this thesis, we first review the predicted couplings of axions and ALPs with baryonic matter that enable their detection via NMR. We then describe two measurement schemes being implemented in CASPEr. The first method, presented in the original CASPEr proposal, consists of a resonant search via continuous-wave NMR spectroscopy. This method offers the highest sensitivity for frequencies ranging from a few Hz to hundreds of MHz, corresponding to masses $m_a \sim 10^{-14}$ – 10^{-6} eV. However, Sub-Hz frequencies are typically difficult to probe with NMR due to the diminishing sensitivity of magnetometers in this region. To circumvent this limitation, we suggest new detection and data processing modalities. We describe a non-resonant frequency-modulation detection scheme, enabling searches from mHz to Hz frequencies ($m_a \sim 10^{-17}$ – 10^{-14} eV), extending the detection bandwidth by three decades.

As a second part of this thesis, we apply the above mentioned non-resonant method and use ultralow-field NMR to probe the axion-fermion “wind” coupling and dark-photon couplings to nuclear spins. No dark-matter signal was detected above background, establishing new experimental bounds for dark-matter bosons with masses ranging from 1.8×10^{-16} to 7.8×10^{-14} eV. While presenting the experiment, we formally introduce the data analysis modality, in addition to reviewing the experimental scheme and the theoretical effects of bosonic dark matter on nuclear spins.

Lastly, we propose to employ Deep Neural Networks (DNNs) to analyze

time-series data emerging from physics experiments. DNNs are widely used in pattern-recognition tasks for which a human-comprehensible, quantitative description of the data-generating process, cannot be obtained. While doing so, DNNs often produce an abstract (entangled and non-interpretable) representation of the data-generating process. This may be one of the reasons why DNNs are not yet used extensively in physics-experiment signal processing: physicists generally require their analyses to yield quantitative information about the system they study. We find this to be a missed opportunity for the physics community.

In the final chapter of this thesis, we propose to use DNNs to disentangle components of oscillating time series. To this aim, we design and train a DNN on synthetic oscillating time series to perform two tasks: a *regression* of the signal latent parameters and *signal denoising* by an *Autoencoder*-like architecture. We show that the regression and denoising performance is similar to those of least-square curve fittings (LS-fit) with true latent-parameters initial guesses, in spite of the DNN needing no initial guesses at all. We then explore various applications in which we believe our architecture could prove useful for time-series processing, when prior knowledge is incomplete. As an example, we employ the DNN as a preprocessing tool to inform LS-fits when initial guesses are unknown. Moreover, we show that the regression can be performed on some latent parameters, while ignoring the existence of others. Because the *Autoencoder* needs no prior information about the physical model, the remaining unknown latent parameters can still be captured, thus making use of partial prior knowledge, while leaving space for data exploration and discoveries.

The axion spin precession scheme

2.1 Foreword

This chapter is mainly composed of the contents of the following article:

Antoine Garcon, Deniz Aybas, John W Blanchard, Gary Centers, Nataniel L Figueroa, Peter W Graham, Derek F Jackson Kimball, Surjeet Rajendran, Marina Gil Sendra, Alexander O Sushkov, Lutz Trahms, Tao Wang, Arne Wickenbrock, Teng Wu, Dmitry Budker. **The Cosmic Axion Spin Precession Experiment (CASPER): a Dark-Matter Search with Nuclear Magnetic Resonance.** *Quantum Science and Technology*, Vol. 3, no. 1, 2017. IOP Publishing.

In this chapter, we briefly review how the ALP-nucleon and axion-gluon couplings could induce an NMR signal. We later focus on describing two NMR measurement schemes implemented in both CASPER-Gradient and CASPER-Electric. The first method, presented in the original CASPER proposal [33], consists of a resonant search via continuous-wave NMR spectroscopy (CW-NMR). This method offers the highest sensitivity for frequencies ranging from a few Hz to hundreds of MHz. Sub-Hz frequencies are typically difficult to probe with NMR due to the diminishing sensitivity

of magnetometers in this region. We then propose a non-resonant frequency-modulation scheme that may circumvent this limitation.

2.2 Axion- and ALP-induced nuclear spin precession

2.2.1 ALP-nucleon coupling - CASPEr-Gradient

CASPEr-Gradient is a haloscope searching for ALPs in the Milky Way's dark-matter halo via their pseudo-magnetic coupling to nucleons, referred as the ALP-nucleon coupling. As the Earth moves through the galactic ALPs, this coupling gives rise to an interaction between the nuclear spins and the spatial gradient of the scalar ALP field [34]. The Hamiltonian of the interaction written in Natural Units takes the form:

$$H_{aNN} = g_{aNN} \sqrt{2\rho_{DM}} \cos(m_a t) \vec{v} \cdot \vec{\sigma}_N, \quad (2.1)$$

where $\vec{\sigma}_N$ is the nuclear-spin operator, $v \sim 10^{-3}$ is the velocity of the Earth relative to the galactic ALPs, $\rho_{DM} \sim 0.4 \text{ GeV/cm}^3$ is the local dark-matter density [42] and g_{aNN} is the coupling strength in GeV^{-1} . The ALP mass, m_a , usually given in electron-volts, can also be expressed in units of frequency, more relevant for an NMR discussion. The Compton frequency associated to the axion and ALP mass is given by: $\omega_a = m_a c^2 / \hbar$, where c is the speed of light in vacuum and \hbar is the reduced Planck constant. For the rest of the discussion, we set $\hbar = c = 1$.

The coupling in Eq. (2.1) is the inner product of an oscillating vector

field with the nuclear-spin operator. Therefore Eq. (2.1) can be rewritten as an interaction between spins and an oscillating pseudo-magnetic field:

$$H_{aNN} = \gamma \vec{B}_{ALP} \cdot \vec{\sigma}_N, \quad (2.2)$$

where γ is the gyromagnetic ratio of the nuclear spin and we have identified the ALP-induced pseudo-magnetic field known as the ‘‘ALP wind’’:

$$\vec{B}_{ALP} = g_{aNN} \frac{\sqrt{2\rho_{DM}}}{\gamma} \cos(\omega_a t) \vec{v}. \quad (2.3)$$

Equation (2.3) can be understood as follows: as nuclear spins move with velocity \vec{v} through the galactic dark-matter halo, they behave as if they were in an oscillating-magnetic field \vec{B}_{ALP} of frequency ω_a , oriented along \vec{v} . As ρ_{DM} and \vec{v} are determined by astrophysical observations, the only free parameters are the ALP frequency (or equivalently, the ALP mass) and the coupling constant, which define the two-dimensional parameter space of the ALP-nucleon coupling shown in Fig. 2.1. Thus the measured amplitude of \vec{B}_{ALP} probes the value of g_{aNN} . Considering the coupling constant range of interest in Fig. 2.1 ($g_{aNN} \sim 10^{-3}$ – 10^{-23} GeV $^{-1}$) and the ^{129}Xe nuclear gyromagnetic ratio ($\gamma \sim 11.777$ MHz/T) yields an ALP-wind amplitude spanning $|\vec{B}_{ALP}| \sim 10^{-10}$ – 10^{-30} T. In order for an experiment targeting the ALP wind to surpass existing astrophysical and laboratory constraints on g_{aNN} , the experiment must be sensitive to ultralow magnetic fields.

The oscillating nature of the ALP wind suggests a magnetic-resonance-based detection method. In the following discussion we briefly explain why NMR techniques, specifically continuous-wave NMR, are particularly well-




Figure 2.1: ALP-nucleon coupling parameter space: coupling strength g_{aNN} versus ALP mass m_a . The purple line represents the mass-coupling parameter space corresponding to the QCD axion proposed to solve the strong CP problem [34]. The darker purple region of the line shows where the QCD axion could be all of the dark matter. The red line is the projected sensitivity of CASPER-Gradient using hyperpolarized ^{129}Xe . The blue line is the sensitivity using hyperpolarized ^3He during a future upgrade of the experiment. The dashed lines are the limits from magnetization noise for ^{129}Xe (red) and ^3He (blue). The ADMX region shows the mass range already excluded (dark blue) or that will be covered (light blue) by ADMX (probing the axion-photon coupling). The green region is excluded by observations on Supernovae SN1987A [43, 44]. The blue region is excluded by searches for new spin-dependent forces. Figure adapted with permission from Ref. [45].

suited for this application.

Consider a collection of nuclear spins with gyromagnetic ratio γ , immersed in a static magnetic field \vec{B}_0 (the leading field; see Fig. 2.2-a), oriented along the z -axis. The spins orient along the leading field to produce a bulk magnetization along the z -axis. We now introduce an oscillating-magnetic field $\vec{B}_{xy}(t)$ oriented in the transverse xy -plane. If the magnitude of the leading field is such that the Larmor frequency $\gamma|\vec{B}_0|$ is equal to the oscillating field frequency, a resonance occurs. The magnetization responds by building up a transverse component, \vec{M}_{xy} .

Subsequently, the transverse magnetization undergoes a precession about \vec{B}_0 , at the Larmor frequency. This oscillation creates a time-varying magnetic field that can be picked-up via magnetometers, producing the NMR signal. The spectrum exhibits a Lorentzian-shaped peak at the Larmor frequency. The search for the resonance is done by varying the magnitude of the leading field and monitoring the transverse magnetization.

This protocol is the continuous-wave NMR experiment introduced by Bloch's first nuclear induction experiment [2]. The term continuous-wave refers to the fact that the oscillating field is continuously applied to the sample. Frequencies are probed sequentially by varying the leading-field magnitude or the oscillating-field frequency.

The fact that the ALP wind has an unknown frequency and cannot be "switched off" suggests a similar experimental scheme. CASPER-Gradient is effectively a CW-NMR experiment in which the transverse component of \vec{B}_{ALP} relative to \vec{B}_0 is analogous to the oscillating transverse field: $\vec{B}_{ALP,xy} \longleftrightarrow \vec{B}_{xy}$. The resonance is reached when the Larmor frequency

is equal to the ALP frequency ($\gamma|\vec{B}_0| = \omega_a$). The experimental scheme is represented in Fig. 2.2-a.

This ALP-induced NMR signal is characterized with two relevant coherence times defining the linewidth of the resonance. The ALP wind oscillates with a temporal coherence inversely proportional to its frequency: $\tau_{ALP} \sim 10^6/\omega_a$ [34]. This effect is modelled by assuming that $\vec{B}_{ALP}(t)$ in Eq. (2.3) acquires a random phase after each time interval τ_{ALP} [33]. In addition, the transverse magnetization decoheres and decays exponentially with a characteristic time T_2 [46]. For the sake of an NMR-based discussion, we will now assume that the signal coherence time is limited by $\tau = \min\{T_2, \tau_{ALP}\}$. As a result the expected linewidth of the resonance becomes $\delta\nu = 1/\pi\tau$ [46]¹.

The allowed values of the ALP mass span many orders of magnitude, yielding a large frequency bandwidth to explore (see Fig. 2.1). Conveniently, the NMR techniques used in CASPER-Gradient are broadly tunable, with the upper bound of the scanned region limited by the achievable magnetic-field strength. In addition, in order to avoid signal broadening that would reduce overall sensitivity, \vec{B}_0 must remain spatially homogeneous over the sample region. These requirements are readily met by superconducting NMR magnets up to fields of about 20 T. Such capabilities enable the detection of ALPs with corresponding frequencies ranging from a few Hz to hundreds of MHz ($m_a \sim 10^{-14}$ – 10^{-6} eV). As such, CASPER-Gradient is a broadband search for light ALP dark matter and is complementary to many other experiments typically looking at higher mass ranges (e.g ADMX searches

¹Recent theories suggest that the signal should not necessarily be Lorentzian-shaped but could be asymmetric, reflecting the fact that the ALPs energy cannot be smaller than $m_a c^2$ (for a stationary ALP) and is higher by $m_a v^2/2$ for a moving ALP [47]. However the discussion remains equivalent.

for axions of mass $m_a \geq 10^{-6}$ eV [32]).

2.2.2 Axion-gluon coupling - CASPEr-Electric

The second CASPEr experiment, CASPEr-Electric, relies on the coupling between CP-solving axions and gluons. This coupling induces an oscillating nucleon electric dipole moment (EDM) [48]:

$$d_n = g_d \frac{\sqrt{2\rho_{DM}}}{\omega_a} \cos(\omega_a t), \quad (2.4)$$

where g_d is the strength of the axion-gluon coupling in GeV^{-2} . As in the case of CASPEr-Gradient, the only two free parameters are g_d and ω_a , giving rise to another two-dimensional parameter space to explore.

The major experimental difference between CASPEr-Electric and CASPEr-Gradient is that CASPEr-Electric makes use of a static electric field applied perpendicularly to the leading magnetic field. As in CASPEr-Gradient, $|\vec{B}_0|$ is tuned to scan for resonance. If the resonance condition is met, the axion-induced EDM oscillates at the Larmor frequency. The interaction between the EDM and the static electric field causes spins to rotate away from the direction of \vec{B}_0 . This produces a non-zero oscillating transverse magnetization, $\vec{M}_{xy}(t)$. Subsequently, $\vec{M}_{xy}(t)$ undergoes precession about \vec{B}_0 and induces the NMR signal.

CASPEr-Electric and CASPEr-Gradient detection schemes are similar in the sense that the effects of both, an oscillating EDM in the presence of a static transverse electric field and an oscillating ALP wind, are analogous to weak magnetic fields, oscillating at the axion or ALP frequencies. Both CASPEr-Gradient and CASPEr-Electric can be described as CW-NMR

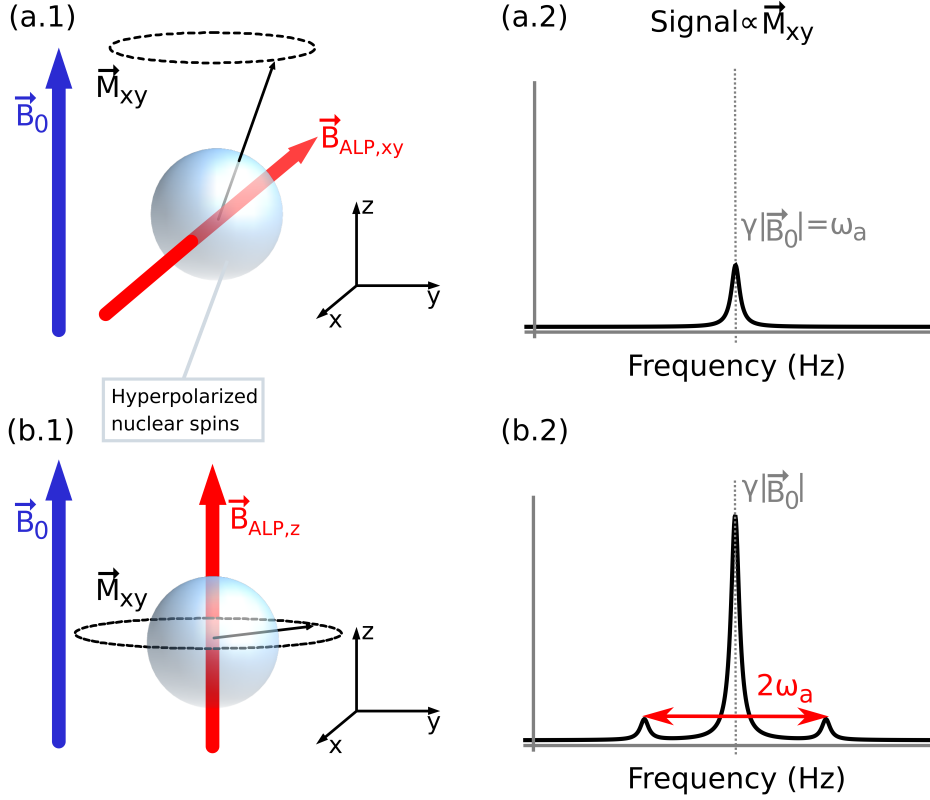


Figure 2.2: Schematic representation of the resonant (CASPER) and sidebands (CASPER, SILFIA) experimental schemes. The hyperpolarized ^{129}Xe sample is immersed in the leading field \vec{B}_0 produced by a tunable NMR magnet. The magnetometer is sensitive to the transverse magnetization \vec{M}_{xy} . The black arrow represents the instantaneous total magnetization. (a.1) Resonant scheme: at resonance ($\gamma|\vec{B}_0| = \omega_a$), the transverse component of the ALP wind, $\vec{B}_{ALP,xy}$, tilts the sample's magnetization which acquires a non-zero component on the xy -plane: \vec{M}_{xy} . \vec{M}_{xy} precesses about \vec{B}_0 at the Larmor frequency. (a.2) The resonant signal is a low amplitude Lorentzian-shaped peak at the ALP frequency. (b.1) Sideband scheme: subsequently to a $\pi/2$ magnetic pulse, the magnetization is on the xy -plane. \vec{M}_{xy} precesses at the Larmor frequency. The longitudinal component of the ALP wind, $\vec{B}_{ALP,z}$, induces modulation of the Larmor frequency. (b.2) The signal of the non-resonant scheme exhibits a carrier frequency ($\gamma|\vec{B}_0|$) and a set of sidebands at $\gamma|\vec{B}_0| \pm \omega_a$. The amplitude of the carrier frequency is large because the full magnetization is rotated by the $\pi/2$ pulse. The sidebands amplitude is expected to be small.

experiments. Although we now focus on the sensitivity of CASPER-Gradient, the discussion for CASPER-Electric is analogous. We note that CASPER-Gradient is not sensitive only to axions and ALPs but could detect any light particle coupling to nuclear spins, in particular, hidden photons [49, 50].

2.3 Experimental sensitivity

Here we discuss the physical parameters affecting the sensitivity of the experiment. Later on, these results are compared to the ones obtained for a non-resonant detection scheme, aiming at probing lower frequencies. During this discussion, we consider the signal on resonance and limit the integration time to $\tau < \min\{\tau_{ALP}, T_2\}$, during which there is coherent averaging of the signal. At resonance, the transverse magnetization increases as if under the action of a low-amplitude rf-field [46]:

$$|\vec{M}_{xy}(\tau)| \propto \gamma \rho P \sin(\gamma |\vec{B}_{ALP,xy}| \tau) \quad \Rightarrow \quad |\vec{M}_{xy}(\tau)| \propto \gamma^2 \rho P |\vec{B}_{ALP,xy}| \tau, \quad (2.5)$$

where ρ is the spin density of the sample given in cm^{-3} and $P \in [0, 1]$ is the dimensionless polarization factor. The NMR signal can be written in terms of the transverse magnetization:²

$$|S(t)| \propto |\vec{M}_{xy}(\tau)| \quad \Rightarrow \quad |S(t)| \propto \gamma^2 \rho P |\vec{B}_{ALP,xy}| \tau. \quad (2.6)$$

To make an estimation on the signal-amplitude threshold A_{min} , for which

²We note that the signal in Eq. (2.6) appears to scale as γ^2 . However, this quadratic scaling is only an artefact arising from the definition of \vec{B}_{ALP} which exhibit a $1/\gamma$ dependence. As such, the signal remain linear in γ .

a event is detected, we use the model proposed in Ref. [33] and assume that the noise is dominated by the magnetometer white noise. Where the white-noise spectral density, $\sqrt{S(f)}$, usually given in fT/ $\sqrt{\text{Hz}}$, is experimentally determined and depends on the frequency probed, f . Limiting the integration time to τ , the signal is coherently averaged and remains constant while the white-noise spectral density decreases as $1/\sqrt{\tau}$ [51]. An event is detected if the signal amplitude is higher than the white noise spectral density after τ , yielding $A_{min} \sim \sqrt{S(f)/\tau}$. The signal-to-noise ratio at resonance, SNR_{RES} , is obtained by comparing the signal amplitude from Eq. (2.6) to A_{min} :

$$SNR_{RES} := |S(\tau)|/A_{min} \quad (2.7)$$

$$\propto \gamma^2 \rho P |\vec{B}_{ALP,xy}| \sqrt{S(f)}^{-1} \tau^{3/2}. \quad (2.8)$$

The SNR of the resonant signal increases as $\tau^{3/2}$. The reasons are as follows: 1) The integration takes place in a time window during which the signal stays phase coherent, thus scaling the SNR as $\tau^{1/2}$ 2) When the resonance condition is satisfied, $|\vec{M}_{xy}(t)|$ increases linearly with time. Hence, the signal is linearly amplified by τ .

Because of this, the SNR greatly benefits from the use of samples with long spin coherence times T_2 . Typical values of T_2 for liquid ^{129}Xe are on the order of 10 to 1000 seconds [52, 53], making xenon an attractive sample. Liquid xenon also provides high spin density, and can be polarized above 50% via spin-exchange optical pumping methods [54], increasing the sensitivity by a factor of at least 10^5 compared to thermal polarization (usually on the order of parts-per-million).

The choice of magnetometer determines the value of $\sqrt{S(f)}$ and is

constrained by the frequency of the expected signal. In the $\omega_a \sim 10\text{--}10^6$ Hz region, the best sensitivities are achieved by SQUIDs. Accounting for all experimental parameters (ALP-wind coherence time, sample geometry, density, and polarization), such SQUIDs would allow CASPER to reach unconstrained regions of the parameters space (see Fig. 2.1; details are given in Ref. [33]).

Frequencies above 2 MHz are usually considered to be the sensitivity cross-over between SQUIDs and inductive pick-up coils. To probe the 2–200 MHz region CASPER will enter its phase II, in which the magnetometer is switched from a SQUID to inductive pick-coils. Atomic magnetometers may also be used at low frequencies, in particular with zero and ultralow magnetic fields as suggested in Ref. [55].

2.4 Ultralow frequencies: sidebands detection

It turns out to be difficult to probe frequencies below ~ 10 Hz with the previously described resonant scheme. Indeed, SQUIDs lose sensitivity below their characteristic “1/f-knee frequency”, typically on the order of a few Hz [56]. To overcome this, one can use a non-resonant measurement protocol that could be implemented as a low-frequency extension to CASPER. The method consists in measuring sidebands induced by modulation of the Larmor frequency and removes the need to scan for resonance. This experimental scheme, represented in Fig. 2.2-b, was first introduced by the “Sideband in Larmor Frequency Induced by Axions” experiment (SILFIA).

In this procedure, the hyperpolarized ^{129}Xe sample is immersed in a leading magnetic field, \vec{B}_0 , oriented along the z -axis and the bulk magnetization is also initially along the z -axis. Prior to the acquisition, a $\pi/2$ magnetic pulse is applied to the sample. After the pulse, the magnetization is in the transverse plane: $M_z \xrightarrow{\pi/2} M_{xy}$. Under the action of \vec{B}_0 , \vec{M}_{xy} precesses about the z -axis at the Larmor frequency $\gamma|\vec{B}_0|$. During the precession, a transient signal $S(t)$ is acquired for a time $\tau = \min\{\tau_{ALP}, T_2\}$. Recalling that $\tau_{ALP} \sim 10^6/\omega_a$ and $T_2 \sim 10 - 1000$ s, then in this low-frequency regime $\tau_{ALP} > T_2$, and the transient-signal coherence time becomes $\tau = T_2$. Once the coherence time T_2 is reached, the transverse magnetization has decayed and the sample is switched for a new one. A $\pi/2$ pulse is applied again and the next transient acquisition takes place.

This measurement method differs from the previous one in the sense that it does not require the ALP wind to tilt the sample magnetization. Following a resonant $\pi/2$ pulse, the magnetization is always in the xy -plane, producing a signal oscillating at the Larmor frequency. The detection involves measuring modulation of the Larmor frequency, induced by the ALP-wind, similarly to the AC-Zeeman effect [57]. In contrast to the resonant scheme, sidebands are induced by the longitudinal component of \vec{B}_{ALP} relative to \vec{B}_0 : $\vec{B}_{ALP,z}$ (see Fig. 2.2-b). The frequency-modulated signal takes the form:

$$S(t) \propto \gamma\rho P e^{-i\gamma|\vec{B}_0|t} e^{-i\gamma|\vec{B}_{ALP,z}|\sin(\omega_a t)/\omega_a}. \quad (2.9)$$

This expression can be written in terms of the Bessel functions of the first

kind J_k :

$$S(t) \propto \gamma \rho \mathbf{P} \sum_{k=-\infty}^{\infty} J_k \left(\frac{\gamma |\vec{B}_{ALP,z}|}{\omega_a} \right) e^{-i\gamma |\vec{B}_0| t} e^{ik\omega_a t}. \quad (2.10)$$

The spectrum of such a frequency modulated signal exhibits a large central peak at the Larmor frequency and sidebands located at $\gamma |\vec{B}_0| \pm k\omega_a$, where $k = 1, 2, 3, \dots$. The k^{th} sideband's amplitude is given by [58]:

$$|S_k(t)| \propto \gamma \rho \mathbf{P} J_k(I) \Rightarrow |S_k(t)| \propto \gamma \rho \mathbf{P} \left(\frac{I^k}{2^k k!} - \frac{I^{k+2}}{2^{k+2} (k+1)!} + \dots \right), \quad (2.11)$$

where $I = \gamma |\vec{B}_{ALP,z}| / \omega_a$ is the modulation index. Recalling that $|\vec{B}_{ALP}| \sim 10^{-10} - 10^{-30}$ T, we see that $I \ll 1$. Thus, the signal can be approximated by its carrier and the first set of sidebands ($k = \pm 1$) arising at frequencies $\gamma |\vec{B}_0| \pm \omega_a$ [59]. Expanding J_1 to first order, the sideband signal can be approximated by:

$$S_{SB}(t) \propto \gamma^2 \rho \mathbf{P} \frac{|\vec{B}_{ALP,z}|}{\omega_a} e^{-i\gamma |\vec{B}_0| t} e^{\pm i\omega_a t}. \quad (2.12)$$

The Fourier transform of $S(t)$ yields a spectrum presenting a central peak at $\gamma |\vec{B}_0|$, surrounded by two sidebands located at $\gamma |\vec{B}_0| \pm \omega_a$. The amplitude of the central peak is large because the full magnetization is rotated by the $\pi/2$ pulse. As $\gamma |\vec{B}_{ALP,z}| / \omega_a \ll 1$, the sidebands amplitude is expected to be small.

Synchronizing the pulse with a stable clock enables recovery of initial phases of each transient signal and allows coherent averaging as long as

the acquisition time is smaller than τ_{ALP} [60, 61]. Thus we can coherently average the transient signals into sets, \bar{s} , each of length T_2 . Here \bar{s} , is the average over $n = \tau_{ALP}/T_2$ transient signals: $\bar{s} = \frac{1}{n} \sum_{i=1}^n S_i(t)$. Coherent averaging of n transient signals decreases the expectation value of the white noise as $1/\sqrt{n}$ [51] while keeping the signal amplitude constant, as in the resonant case. Considering a white-noise spectral density of $\sqrt{S(f)}$, coherent averaging of the n transient signals yields a signal threshold of:

$$A_{min,n} \sim \sqrt{S(f)} (nT_2)^{-1/2}. \quad (2.13)$$

We recall that the sidebands are always located around the Larmor frequency, therefore this scheme removes the need to search for resonance. $|\vec{B}_0|$ is arbitrary adjusted such that the detection can be done in a region where the SQUID sensitivity is optimum, regardless of the ALP frequency: $\sqrt{S(f)} := \sqrt{S_{opt}}$. We can then rewrite Eq. (2.13) setting the SQUID noise spectral density to $\sqrt{S_{opt}}$:

$$A_{min,n} \sim \sqrt{S_{opt}} (nT_2)^{-1/2}. \quad (2.14)$$

The acquisition and coherent averaging are repeated until the desired total integration time T_{tot} is reached. This yields N independent averaged sets $\{\bar{s}_1, \bar{s}_2, \dots, \bar{s}_N\}$ each of length T_2 and of signal threshold $A_{min,n}$. Here $N = T_{tot}/\tau_{ALP}$, is the number of independent sets. Each set is measured for a time τ_{ALP} . We recall that after τ_{ALP} , the ALP-wind phase changes by an unknown amount. As a result, the N sets have uncorrelated phases and averaging them would decrease the signal as well as the noise. However,

some additional signal-processing techniques can be employed to improve the detection threshold, namely, incoherent averaging of the sets \bar{s}_i .

In practice, incoherent averaging corresponds to averaging the sets $\{\bar{s}_1, \bar{s}_2, \dots, \bar{s}_N\}$ in the frequency domain by averaging their power spectral densities: $\overline{\text{PSD}} = \frac{1}{N} \sum_{i=1}^N \text{PSD}(\bar{s}_i)$. Such processing does not reduce the noise mean value, but only the noise-power standard deviation [62]. After the incoherent averaging sequence, the signal threshold becomes (details are given in the appendix of Ref. [33]):

$$A_{min,N} \sim A_{min,n} N^{-1/4} \Rightarrow A_{min,N} \sim \sqrt{S_{opt}} (\tau_{ALP} T_{tot})^{-1/4}. \quad (2.15)$$

The signal threshold scales as $t^{-1/2}$ as long as the signal is phase coherent ($t < \tau_{ALP}$), then scales as $t^{-1/4}$ once the ALP coherence time is reached ($t > \tau_{ALP}$). Recalling that $\tau_{ALP} \sim 10^6/\omega_a$ the signal threshold in Eq. (2.15) becomes:

$$A_{min,N} \sim 10^{-3/2} \sqrt{S_{opt}} (T_{tot} \omega_a)^{-1/4}. \quad (2.16)$$

After the coherent and incoherent averaging sequences, the signal-to-noise ratio $SNR_{sidebands}$ is obtained by comparing the signal amplitude from Eq. (2.12) to $A_{min,N}$:

$$SNR_{SB} := |S_{SB}(t)|/A_{min,N} \quad (2.17)$$

$$\sim 10^{3/2} \gamma^2 \rho \mathbf{P} |\vec{B}_{ALP,z}| \sqrt{S_{opt}}^{-1} T_{tot}^{1/4} \omega_a^{-5/4}. \quad (2.18)$$

SNR_{SB} does not scale with T_2 . Indeed, the transient signals can be

coherently averaged until τ_{ALP} is reached, making T_2 irrelevant (ignoring the duty cycle). Two factors contribute to the gains in sensitivity in the low-frequency ALPs region; 1) Lower ALPs frequencies imply longer ALP-wind coherence time, increasing the time during which the signal can be coherently averaged 2) The amplitude of the transient signal is determined by the modulation index of the first Bessel function. As such, low frequency modulations produce higher sidebands amplitudes.

The relative sensitivity of the resonant and sideband schemes is determined by the ratio of equations (2.18) and (2.8). This ratio is computed by assuming identical samples and equal longitudinal and transverse components of \vec{B}_{ALP} . We impose identical total integration time, T_{tot} , for both the resonant and sidebands measurement scheme. T_{tot} is calculated by assuming coherent averaging for a time $\tau = T_2$ in each frequency bin during the resonant search. Ignoring duty cycle, the total integration time during the resonant search becomes $T_{tot} = \frac{|\Omega|}{\delta\nu} T_2$. Where Ω is the frequency range of interest and $\delta\nu \sim 1/\pi T_2$ is the linewidth of the resonance. The relative sensitivity of the resonant and sideband scheme becomes:

$$SNR_{SB}/SNR_{RES} \sim 10^{3/2} |\Omega|^{1/4} \pi^{-1/4} \frac{\sqrt{S(f)}}{\sqrt{S_{opt}}} \omega_a^{-5/4} T_2^{-1}. \quad (2.19)$$

We now set the frequency range to $\Omega = 0-1$ kHz and assume that the SQUID white-noise level above the knee frequency, $f_0 = 2$ Hz, is approximately constant $\sqrt{S}(f > f_0) = \sqrt{S_{opt}} = 0.9$ fT/ $\sqrt{\text{Hz}}$. Below f_0 , the SQUID noise is set to $\sqrt{S}(f < f_0) \sim 10$ fT/ $\sqrt{\text{Hz}}$ (extracted from Ref. [63]). Values of T_2 higher than 100 seconds impose a total integration time longer

than 1 year to probe the region Ω via the resonant scheme and are ignored. Eq. (2.19) is plotted against ω_a and T_2 in Fig. 2.3.

As shown in Fig. 2.3, the sidebands scheme is beneficial for ALP frequencies below 80 Hz and if long resonant integration time in each frequency bin cannot be achieved (either due to low T_2 or to some constraint on T_{tot}). We recall that the sidebands are located around the central peak at $\gamma|\vec{B}_0| \pm \omega_a$. Thus if the ALP frequency is lower than the linewidth $1/\pi T_2$, the sidebands are inside the central peak and cannot be resolved. This case corresponds to the excluded, red region of the Fig. 2.3 and represents the current lower bound of the CASPEr experiment. Considering a realistic T_2 of 100 seconds, yields a lower limit in the mHz range. This sideband-detection scheme enables detection of ALPs with masses in the $m_a \sim 10^{-17}$ – 10^{-14} eV region, increasing the bandwidth of the experiment by three orders of magnitude and allowing the CASPEr detection region to convene with ultracold neutrons experiments [64].

2.5 Discussion

Axions and ALPs are well-motivated dark-matter candidates; in addition, the QCD axions provides a solution to the strong CP problem. The discovery of such particles would shed light on many fundamental questions in modern physics, offering a glimpse of physics beyond the Standard Model.

The oscillatory pseudo-magnetic couplings between axion/ALPs and matter open the possibility of direct dark-matter detection via NMR techniques. When nuclear spins couple to ALPs in the Milky Way dark-matter halo, the spins behave as if they were in an oscillating magnetic field.

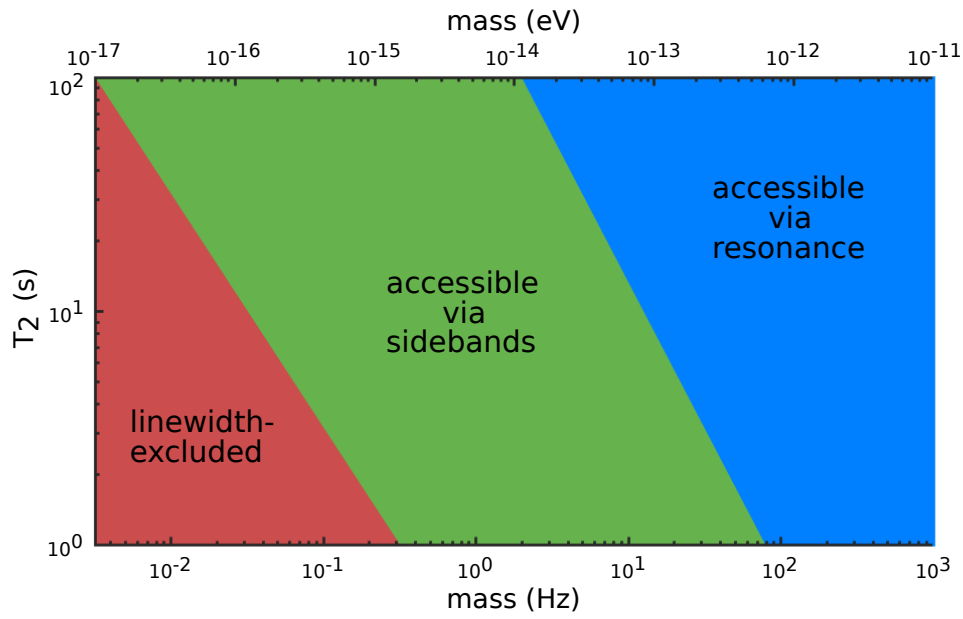


Figure 2.3: Resonant versus sidebands detection: sensitivity cut-off regions for equal total integration time and sample parameters (spin density, geometry). Extracted from Eq. (2.19). Blue area: sensitivity is higher for resonant detection. Green area: region enabled by sidebands detection. Red area: excluded by the linewidth $1/\pi T_2$. We assume that the SQUID white-noise level above the knee frequency, $f_0 = 2$ Hz, is approximately constant $\sqrt{S}(f > f_0) \sim 0.9$ fT/ $\sqrt{\text{Hz}}$. Below f_0 , SQUID noise is set to $\sqrt{S}(f < f_0) \sim 10$ fT/ $\sqrt{\text{Hz}}$ (extracted from Ref. [63]; low- T_c W9L-18D9 SQUID - PTB). Values of T_2 higher than 100 seconds impose a total integration time longer than 1 year to probe the region of interest via the resonant scheme and are ignored.

The axion-gluon coupling can induce an oscillating nucleon electric-dipole moment. CASPEr-Gradient and CASPEr-Electric seek to measure the NMR signals induced by the ALP wind and the nucleon EDM, respectively.

The original CASPEr experiment is based on resonant search via CW-NMR. This method enables the search for axions and ALPs at frequencies ranging from a few Hz to a few hundred MHz ($m_a \sim 10^{-14} - 10^{-6}$ eV). Probing lower frequencies using this approach suffers in sensitivity due to limitations of the magnetometers. Therefore non-resonant detection of ALP-induced sidebands around the Larmor frequency could be beneficial. This detection scheme allows probing of the axion and ALP parameter space in the mHz to Hz region ($m_a \sim 10^{-17} - 10^{-14}$ eV), thus increasing the bandwidth of CASPEr by three decades. Sideband-based ALP searches introduced by the SILFIA experiment using hyperpolarized noble gases in the gas phase and SQUID detection, are already in progress at the The National Metrology Institute of Germany (PTB).

Only a few experiments are tuned to the mass range accessible by CASPEr, even though the presence of ALPs at these frequencies is well-motivated. This makes CASPEr complementary to other searches, which typically look at lower or higher frequencies.

Experimental constraints on bosonic dark matter

3.1 Foreword

This chapter is mainly composed of the contents of the following article:

Antoine Garcon, John W Blanchard, Gary P Centers, Nataniel L Figueroa, Peter W Graham, Derek F Jackson Kimball, Surjeet Rajendran, Alexander O Sushkov, Yevgeny V Stadnik, Arne Wickenbrock, Teng Wu, Dmitry Budker. **Constraints on Bosonic Dark Matter from Ultralow-Field Nuclear Magnetic Resonance.** *Science Advances*, Vol. 5, no. 10, 2019.

In this chapter, we describe an application of the ALP-induced sidebands measurement scheme proposed in Chap. 2. This experiment was carried via ZULF NMR on a sample of ^{13}C -enriched formic acid. We first describe the ZULF NMR experimental setup and review the three relevant magnetic regimes in which the ZULF experiments can be carried. We then discuss in more details the properties of bosonic dark-matter fields and their possible effect on the NMR sample. In particular we show that axion, ALP and dark-photon fields, can be detected by examining ZULF NMR spectra in search for sidebands around the formic acid resonance lines. The subsequent sections

describe the measurement scheme, the data processing techniques, and calibration employed during this search. Finally we present new laboratory results for bosonic dark matter, complementing astrophysical constraints obtained from supernovae 1987A¹ [44].

3.2 Experimental scheme

3.2.1 Experimental setup overview

The experimental setup used in this experiment is the one described in Ref. [4]. Additional descriptions of similar ZULF NMR setups can be found in Refs [67, 68].

The NMR sample consists of $\sim 100 \mu\text{L}$ of liquid ^{13}C -formic acid ($^{13}\text{CHOOH}$) obtained from ISOTECH Stable Isotopes (Millipore Sigma), degassed by several freeze-pump-thaw cycles under vacuum, and flame-sealed in a standard 5 mm glass NMR tube.

Prior to every transient acquisition, the sample is thermally polarized for ~ 30 s in a 1.8 T permanent magnet, after which the NMR tube is pneumatically shuttled into the zero-field region. During the shuttling, the sample goes through a magnetic field of a guiding solenoid wrapped around a fiberglass tube, ensuring adiabatic transport into the shielded region [41].

The zero-field region is achieved using a four-layer μ -metal and ferrite shield (MS-1F shield produced by Twinleaf LLC, 10^6 magnetic shielding factor). In this region, the residual magnetic field is $\lesssim 10$ pT. Three orthogonal pairs of Helmholtz coils which receive current from an amplified

¹We note that the constraints based on SN1987A data continue to be reexamined; see e.g. Refs. [65, 66].

source (AE Techron 7224-P) are wound on the holder. This allows for the application of AC and DC magnetic fields along three orthogonal directions. This set of coils is used to apply acquisition-starting pulses. A separate set of coils mounted on the innermost shield is used to apply the shimming, leading DC-, calibration AC-, and benchmark AC-fields.

After shuttling, the guiding magnetic field is turned off, a magnetic pulse (corresponding to a π rotation of the ^{13}C spin) is applied to initiate magnetization evolution.

In the zero-field region, the sample stops ~ 1 mm above the spin-exchange-relaxation-free magnetometer's (SERF [69, 70]) rubidium-vapor cell (cell produced by Twinleaf with 500 torr of nitrogen buffer gas).

In order for the magnetometer to operate in SERF regime, the rubidium cell is maintained at 180 C° by means of a resistive heater. The 795 nm circularly-polarized pump beam (from a Toptica Photonics diode laser system, locked to the D1 Rubidium line), propagates along the y-axis.

The linearly polarized probe beam (from another Toptica Photonics diode laser system), propagating along the x-axis, is blue-detuned 10 GHz from the center of the rubidium D2 line to probe the rubidium atoms' polarization. The polarimeter's analog output signal is digitized with a 24 bit acquisition card (NI 9239, National Instrument) at a 5 kHz sampling rate. Synchronization of the experiments and control of shuttling, magnetic pulses, and data acquisition is accomplished with a LabVIEW program.

Following each transient acquisition, the sample is shuttled back into the polarizing magnet and the experiment is repeated. In order to increase the SNR, the transient signals are averaged using the phase-cycling technique

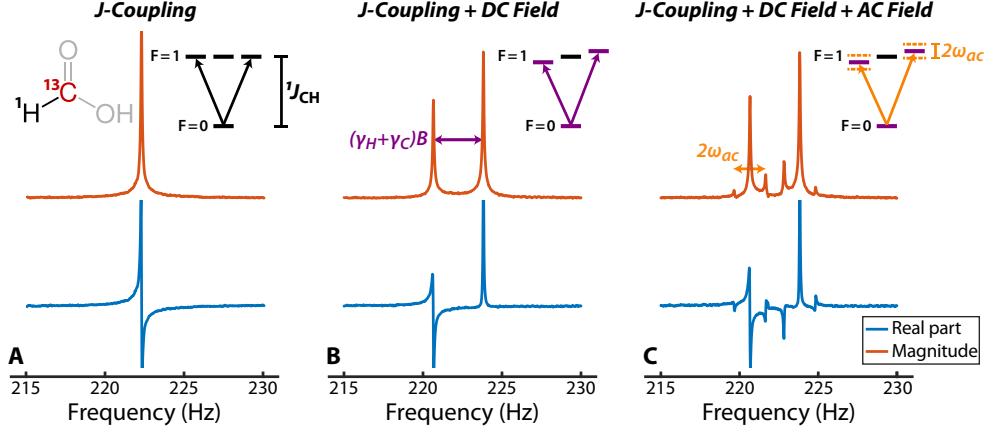


Figure 3.1: Nuclear spin energy levels and NMR spectra of ^{13}C -formic acid measured in three different field conditions. **(a)** At zero magnetic field, the $F = 1$ levels are degenerate, resulting in a spectrum exhibiting a single peak at the J -coupling frequency. **(b)** In the presence of a DC-magnetic field $B_z \approx 50$ nT, the $m_F = \pm 1$ degeneracy is lifted. The spectrum exhibits two split J -resonances. The splitting is equal to $\hbar m_F B_z (\gamma_C + \gamma_H)$. The asymmetry of the resonances is due to the influence of the applied field on the response characteristics of the atomic magnetometer. **(c)** The addition of an oscillating magnetic field along B_z modulates the $m_F = \pm 1$ energy levels, resulting in sidebands located at $J/2\pi \pm B_z(\gamma_C + \gamma_H)/2 \pm \omega_{AC}$ with amplitude proportional to the modulation index: $A_s \propto B_{AC}(\gamma_C + \gamma_H)/2\omega_{AC}$.

described later in Sec. 3.4.1.

3.2.2 Measurement schemes

The experimental setup described above enables measurements to be done in three notable magnetic conditions. The “zero-field” regime corresponds to the situation during which no external field is applied to the sample which evolution then is dominated by the internal Hamiltonian. The “ultralow-field” regime is when a weak DC magnetic field is applied to the sample. In this case, the Zeeman Hamiltonian can be considered as a small perturbation of the internal Hamiltonian. Finally, we refer to the “frequency-modulation”

regime, the situation during which a weak DC magnetic field is applied in addition to an oscillating magnetic field. In this section, we briefly review the physics of these magnetic conditions applied to the polarized sample of formic acid.

Zero-field regime

^{13}C -formic acid possesses an electron-mediated spin-spin coupling between the ^{13}C and its neighbouring ^1H , referred to as J -coupling. In isotropic liquids, rotational motion of the molecules averages the J -coupling Hamiltonian to its isotropic form:

$$\mathcal{H}_J = \hbar J_{\text{CH}} \mathbf{I} \cdot \mathbf{S} , \quad (3.1)$$

where $J_{\text{CH}}/2\pi \approx 221$ Hz for formic acid, \mathbf{I} and \mathbf{S} are the nuclear spin-1/2 operators for ^1H and ^{13}C , respectively.

Prior to the acquisition, the sample is shuttled to the magnetically shielded region where the residual magnetic field is on the order of 10 pT. In such a low magnetic field, the Zeeman interaction is on the order of 0.1 mHz and is negligible compared to the J -coupling. For isotropic fluids in zero-field, other interactions typically relevant in NMR spectroscopy, such as the short-range dipole-dipole couplings, are averaged-out by molecular tumbling and are also negligible. As a result, the sample's evolution in zero-field is dominated by the J -coupling Hamiltonian given by Eq. (3.1).

Following shuttling, the guiding field is turned off suddenly and a magnetic field pulse with area $\gamma_C B_x t_{\text{pulse}} = \pi$ is applied to the sample, where γ_C is the gyromagnetic ratio of the ^{13}C nuclear spin. As a result, the ^{13}C and

^1H nuclear spins are in a superposition of the triplet $|F=1; m_F=\pm 1, 0\rangle$ and singlet $|F=0; m_F=0\rangle$ states which evolves under the J -coupling Hamiltonian. The unperturbed energy levels are given by:

$$E(F)/\hbar = \frac{1}{2}J_{\text{CH}}[F(F+1) - I(I+1) - S(S+1)], \quad (3.2)$$

where $I = S = 1/2$ for ^{13}C and ^1H and $F = 0, 1$. Then the singlet and triplet states energy levels are respectively:

$$\begin{aligned} E(F=0, m_F=0)/\hbar &= -\frac{3}{4}J_{\text{CH}}, \\ E(F=1, m_F=0, \pm 1)/\hbar &= \frac{1}{4}J_{\text{CH}}. \end{aligned} \quad (3.3)$$

The selection rules for transition between the singlet and triplet states, $\Delta F = 0, \pm 1$ and $\Delta m_F = \pm 1$, arise because the observable is the magnetization along the y -axis. The corresponding spectrum exhibits a single peak centred at $[E(F=1, \pm 1) - E(F=0)]/\hbar = J \approx 2\pi \cdot 221 \text{ Hz}$ (see Fig. 3.1.a).

Ultralow-field regime

The experimental setup allows application of a DC-magnetic field, \mathbf{B}_z , along the z -axis via the Helmholtz coils surrounding the sample. In such conditions, the Hamiltonian under which the formic acid molecules evolve becomes:

$$\mathcal{H} = \mathcal{H}_J + \mathcal{H}_Z \quad (3.4)$$

$$= \mathcal{H}_J - \hbar(\gamma_{\text{C}}S_z + \gamma_{\text{H}}I_z)B_z, \quad (3.5)$$

where $\gamma_C/2\pi = 10.70 \text{ MHz}\cdot\text{T}^{-1}$ and $\gamma_H/2\pi = 42.57 \text{ MHz}\cdot\text{T}^{-1}$ are the gyromagnetic ratios of the ^{13}C and ^1H nuclear spins, respectively. A field of about 50 nT yields a Zeeman interaction, H_Z , on the order of $\gamma_C B_z/2\pi \approx \gamma_H B_z/2\pi \approx 0.1 \text{ Hz}$, which remains much weaker than the J -coupling and can thus be treated as a perturbation. To first order in B_z , the eigenstates are those of the unperturbed Hamiltonian, \mathcal{H}_J , and perturbed energies can be read from the diagonal matrix elements of the Zeeman perturbation:

$$\Delta E(F, m_F)/\hbar = -\langle F, m_F | (\gamma_C S_z + \gamma_H I_z) B_z | F, m_F \rangle, \quad (3.6)$$

yielding the perturbed energy levels:

$$\begin{aligned} E(F=0, m_F=0)/\hbar &= -\frac{3}{4} J_{\text{CH}}, \\ E(F=1, m_F=0, \pm 1)/\hbar &= \frac{J_{\text{CH}}}{4} - m_F \frac{\gamma_C + \gamma_H}{2} B_z. \end{aligned}$$

Thus, to first order, the effect of such a field is to break the degeneracy of the triplet state due to the now non-negligible Zeeman splitting. Recalling the selection rules $\Delta F = 0, \pm 1$ and $\Delta m_F = \pm 1$, the magnetometer measures oscillations at two different frequencies between the $|F=0; m_F=0\rangle$ and $|F=1; m_F=\pm 1\rangle$ states. As a result, the single J -coupling line is split into a doublet and the spectrum exhibits two lines located at $J/2\pi \pm B_z \frac{\gamma_C + \gamma_H}{4\pi}$ (see Fig. 3.1.b).

Frequency-modulation regime

In this regime, we apply an oscillating magnetic field $\mathbf{B}_{\text{AC}} = B_{\text{AC}} \cos(\omega t) \cdot \hat{\mathbf{e}}_z$ in addition to the DC-field along z-axis. Under these conditions, the formic

acid molecules evolve under the following Hamiltonian:

$$\mathcal{H} = \mathcal{H}_J + \mathcal{H}_Z + \mathcal{H}_{AC} \quad (3.7)$$

$$= \mathcal{H}_J + \mathcal{H}_Z - \hbar(\gamma_C S_z + \gamma_H I_z) B_{AC} \cos(\omega t) . \quad (3.8)$$

In the case where the modulation index, $\frac{B_{AC}(\gamma_C + \gamma_H)}{2\omega}$, is much smaller than one, the resulting spectrum is composed of the J -coupling doublet peaks, located at $J_{CH}/2\pi \pm B_z \frac{\gamma_C + \gamma_H}{4\pi}$ each exhibiting a set of sidebands located at frequencies $J_{CH}/2\pi \pm B_z \frac{\gamma_C + \gamma_H}{4\pi} \pm \omega/2\pi$. Such a spectrum is shown in Fig. 3.1.c, alongside the corresponding energy levels. The sidebands' amplitude A_s , is proportional to the modulation index [71]. This behavior will later be experimentally confirmed by varying the amplitude and frequency of a calibration AC-field while a DC-field was continuously applied to the sample.

3.3 Bosonic dark-matter fields

3.3.1 Dark-matter field properties

If dark matter predominantly consists of particles with masses $m_{DM} \lesssim 10$ eV, making up the totality of the average local dark-matter density, then they must be bosons with a large mode occupation number. It would be impossible for fermions with such low masses to account for the observed galactic dark matter density, since the Pauli exclusion principle prevents them from having the required mode occupation.

In this scenario, axion and ALP bosonic dark matter is well described by

a classical field $a(t)$, oscillating at the Compton frequency ($\omega_{\text{DM}} \approx m_{\text{DM}}c^2/\hbar$) [72–74]:

$$a(t) \approx a_0 \cos(\omega_{\text{DM}}t) , \quad (3.9)$$

where a_0 is the amplitude of the bosonic field.

The temporal coherence of the bosonic field is limited by the relative motion of the Earth through random spatial fluctuations of the field. The characteristic coherence time τ_{DM} , during which the bosonic dark-matter fields remains phase coherent, corresponds to $\sim 10^6$ periods of oscillation of the fields [36].

The amplitude a_0 can be estimated by assuming that the field energy density constitutes the totality of the average local dark-matter energy density ($\rho_{\text{DM}} \approx 0.4 \text{ GeV}/\text{cm}^3$ [75]). Then a_0 is related to the dark-matter density through:

$$\rho_{\text{DM}} = \frac{1}{2} \frac{c^2}{\hbar^2} m_{\text{DM}}^2 a_0^2 . \quad (3.10)$$

We note that a_0 is expected to fluctuate with relative amplitude of order one due to self-interference of the field. This effect should induce stochastic variations of a_0 , and great care must be employed when computing the actual amplitude of ALP field. As this effect is beyond the scope of this thesis we assume a_0 to be constant for the rest of this discussion. We later account for these variations by using a Monte-Carlo method. We invite the readers to consult Ref. [76] which provide a more complete framework regarding this effect.

3.3.2 Dark-matter couplings to nuclear spins

CASPEr-Gradient is sensitive to any field such that its interaction with nuclear spins can be written in the form:

$$\mathcal{H}_{\text{DM}} = -\hbar g_{\text{N}} \mathbf{I}_{\text{N}} \cdot \mathbf{D}, \quad (3.11)$$

where g_{N} is a coupling constant that parametrizes the coupling of the effective field, \mathbf{D} , to nuclear spins represented by the operator \mathbf{I}_{N} . In analogy with the Zeeman interaction, $\mathcal{H}_{\text{Zeeman}} = -\hbar \gamma_{\text{N}} \mathbf{I}_{\text{N}} \cdot \mathbf{B}$, such an effective field \mathbf{D} may be thought of as a pseudo-magnetic field interacting with nuclear spins, where the nuclear gyromagnetic ratio, γ_{N} , is replaced by the coupling constant, g_{N} .

For clarity, we focus this discussion on the the so-called “axion wind interaction” with effective field $\mathbf{D}_{\text{wind}} = -\nabla a(\mathbf{r}, t)$ and coupling constant g_{aNN} . A number of other possible couplings between nuclear spins and bosonic dark matter fields take a form similar to Eq. (3.11). These include couplings to the “dark” electric (with coupling constant g_{dEDM} and effective field \mathbf{D}_{dEDM}) and magnetic (with coupling constant g_{dMDM} and effective field \mathbf{D}_{dMDM}) fields mediated by spin-1 bosons such as dark photons [19, 77] or a quadratic “wind” coupling to an ALP field (with coupling constant g_{quad} and effective field \mathbf{D}_{quad}) [78]. These are discussed later.

As the Earth orbits around the Sun (itself moving towards the Cygnus constellation at velocity, \mathbf{v} , comparable to the local galactic virial velocity $\sim 10^{-3}c$), it moves through the galactic dark-matter halo and an interaction between axions and ALPs with a given nucleon N , can arise. Assuming that ALPs make up all of the dark matter energy density, ρ_{DM} , and that the dominant interaction with nucleon spins is linear in $\nabla a(\mathbf{r}, t)$, Eqs. (3.9) and

(3.10) can be used to write the the effective field as

$$\mathbf{D}_{\text{wind}}(t) = -\sqrt{2\hbar c\rho_{\text{DM}}} \sin(\omega_{\text{DM}}t) \mathbf{v}. \quad (3.12)$$

Then, given the local galactic virial velocity, \mathbf{v} , only two free parameters remain in the Hamiltonian in Eq. (3.11): the coupling constant, g_{aNN} , and the field's oscillation frequency, ω_{DM} , fixed by the boson mass. A CASPER search consists of probing this parameter space over a bosonic mass range defined by the bandwidth of the experiment. In order to calibrate the experiment, we apply known magnetic fields and the experiment's sensitivity to magnetic fields directly translates to sensitivity to the coupling constant. If no ALP field is detected, upper bounds on the coupling constant can be determined based on the overall sensitivity of the experiment.

3.3.3 Dark-matter signatures in zero- to ultralow-field NMR

The sample - ^{13}C -formic acid, effectively a two-spin ^1H - ^{13}C system - is pre-polarized in a 1.8 T permanent magnet and pneumatically shuttled to a magnetically-shielded environment for magnetization evolution and detection.

The spin Hamiltonian describing the system is

$$\begin{aligned} \mathcal{H} = & \hbar J_{\text{CH}} \mathbf{I} \cdot \mathbf{S} - \hbar (\gamma_{\text{H}} \mathbf{I} + \gamma_{\text{C}} \mathbf{S}) \cdot \mathbf{B} \\ & + \hbar \left(g_{\text{app}} \mathbf{I} - \frac{1}{3} g_{\text{ann}} \mathbf{S} \right) \cdot \mathbf{D}_{\text{wind}}(t), \end{aligned} \quad (3.13)$$

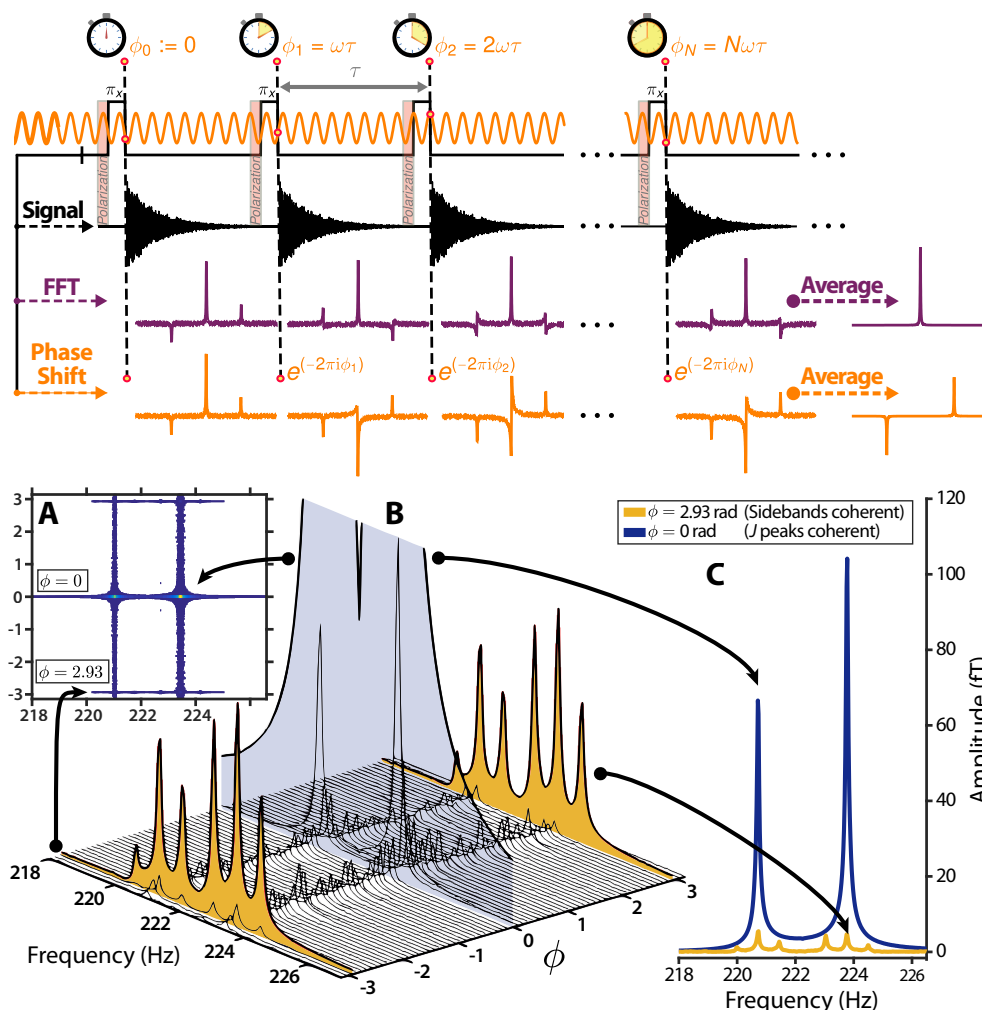


Figure 3.2: **Top:** signal acquisition scheme with simulated spectra. After polarization, each transient acquisition starts following a magnetic π -pulse (corresponding to a 180 degrees flip of the ^{13}C spin along any direction). The external AC-magnetic field's phase varies between transient acquisitions (orange). As a result, the sidebands generally possess different phases in each transient spectrum. Averaging the transients yields a spectrum in which the sidebands are destructively averaged out (purple). Shifting each transient by a phase equal to the external field's accumulated phase restores the sidebands' phase coherence, yielding a spectrum with high signal-to-noise ratio sidebands (orange). For clarity, only one of the two Zeeman-split J -coupling peaks and its two sidebands are shown. **Bottom:** result of the phase-shifting procedure for actual data. **(a)** Transients are averaged using 2001 phase increments and stacked into a two-dimensional plot. **(b)** Side view of **(a)**, sidebands are rescaled by a factor 10 for clarity. **(c)** Averaging with $\phi = 0$ rad corresponds to averaging the transients without phase shift; sidebands are averaged out and carrier peaks appear with maximum amplitude. When the optimal phase (for $\omega/2\pi = 0.73$ Hz, $\phi = 2.93$ rad) is approached, sidebands appear. These spectra were acquired in an experiment during which the AC-field frequency and amplitude were set to 0.73 Hz and 0.24 nT. Transient acquisitions of 30 s were repeated 850 times with a time interval between each transient of $\tau = 61$ s.

where the electron-mediated spin-spin coupling $J_{\text{CH}}/2\pi \approx 221$ Hz for formic acid and \mathbf{I} and \mathbf{S} are the nuclear-spin operators for ^1H and ^{13}C , respectively. Additionally, \mathbf{B} is an applied magnetic field, g_{app} is the ALP-proton coupling strength, and g_{ann} is the ALP-neutron coupling strength. We assume $g_{\text{app}} = g_{\text{ann}} = g_{\text{aNN}}$ [40].

In the absence of external fields ($\mathbf{B} = \mathbf{D}_{\text{wind}}(t) = 0$), the nuclear spin energy eigenstates are a singlet with total angular momentum $F = 0$ and three degenerate triplet states with $F = 1$, separated by $\hbar J_{\text{CH}}$. The observable in our experiment is the y magnetization, leading to selection rules $\Delta F = 0, \pm 1$ and $\Delta m_F = \pm 1$, as in Ref. [79]. The zero-field spectrum thus consists of a single Lorentzian located at $J_{\text{CH}}/2\pi \approx 221$ Hz, as shown in Fig. 3.1(a).

In the presence of a static magnetic field, $\mathbf{B} = B_z \hat{\mathbf{e}}_z$, applied along z , the $m_F = 0$ states are unaffected, while the $m_F = \pm 1$ triplet states' degeneracy is lifted. The corresponding spectrum exhibits two peaks at $J_{\text{CH}}/2\pi \pm B_z(\gamma_{\text{H}} + \gamma_{\text{C}})/2\pi$, as shown in Fig. 3.1(b).

So long as $|J_{\text{CH}}| \gg |\gamma_{\text{H}} B| \gg |g_{\text{aNN}} D_{\text{wind}}|$ the $m_F = 0$ states are unaffected, and the $m_F = \pm 1$ states are shifted by

$$\Delta E(m_F = \pm 1)(t) = \mp \frac{\hbar}{2} B_z (\gamma_{\text{H}} + \gamma_{\text{C}}) \pm \frac{\hbar}{2} \frac{2}{3} g_{\text{aNN}} D_z(t), \quad (3.14)$$

where $D_z(t)$ is the projection of $\mathbf{D}_{\text{wind}}(t)$ along the axis of the applied magnetic field. The time dependence of $\mathbf{D}_{\text{wind}}(t)$ leads to an oscillatory modulation of the $m_F = \pm 1$ energy levels, giving rise to sidebands around the J -coupling doublet as shown in Fig. 3.1(c). The sidebands are separated from the carrier peaks by $\pm \omega_{\text{DM}}/2\pi$ and have an amplitude proportional to




Figure 3.3: ALP wind-nucleon linear coupling parameter space. The CASPER-ZULF region is excluded by this work (90% confidence level) using a thermally polarized sample (data averaged over 850 transient acquisitions of 30 s each). The “New Force” region is excluded by searches for new spin-dependent forces [80]. The SN1987A region represents existing limits from supernova SN1987A cooling [44, 81]. The ν_n/ν_{Hg} region is excluded by measurements of the ratio of neutron and ^{199}Hg Larmor precession frequencies [40]. The dashed line corresponds to the sensitivity of a planned second phase of CASPER-ZULF, with a projected $\sim 10^5$ factor increase in sensitivity and the bandwidth extended towards lower frequencies by using a comagnetometer technique [82] and longer integration times.

the modulation index ($g_{\text{aNN}}/\omega_{\text{DM}}$).

As a result, dark-matter fields with sufficiently strong coupling to nuclear spins can then be detected by searching for frequency-modulation-induced sidebands in the well-defined ZULF NMR spectrum of formic acid.

3.3.4 Bosonic dark-matter effective fields

In order to predict the amplitude of the formic acid response to bosonic dark-matter fields, one must compute the equivalent magnetic field which would

induce an identical response. We refer to these as dark-matter “effective fields”. In the following, we show the effective fields of the ALP-wind linear and quadratic [78, 83] couplings, in addition to those of the dark-photon couplings. The full derivation of these expressions is given in A.1.1 of the Appendix.

In the case of the ALP-wind linear coupling, the field acting on the ^1H - ^{13}C spins induces an energy shift equal to the one produced by a magnetic field with amplitude

$$B_{\text{ALP},z}^{\text{eff}}(t) = -\frac{2}{3} \frac{g_{\text{aNN}}}{\gamma_{\text{H}} + \gamma_{\text{C}}} \sqrt{2\hbar c \rho_{\text{DM}}} \times \sin(\omega_{\text{DM}} t - \mathbf{k} \cdot \mathbf{r} + \phi) \mathbf{v} \cdot \hat{\mathbf{e}}_z, \quad (3.15)$$

where $\omega_{\text{DM}} \approx m_{\text{DM}} c^2 / \hbar$ is the ALP Compton frequency, $\mathbf{k} \approx m_{\text{DM}} \mathbf{v} / \hbar$ is the wave-vector (\mathbf{v} is the relative velocity), m_{DM} is the rest mass of the ALP, ϕ is an unknown phase, and $\hat{\mathbf{e}}_z$ is the axis along which the leading DC-magnetic field is applied.

It is theoretically possible that interaction of nuclear spins with ∇a can be suppressed [78, 83], in which case the dominant axion wind interaction, referred to as the quadratic wind coupling, is related to ∇a^2 . In the case of the ALP-wind quadratic coupling the equivalent magnetic field amplitude is:

$$B_{\text{quad},z}^{\text{eff}}(t) = -\frac{4}{3} \hbar c^2 \frac{g_{\text{quad}}^2}{\gamma_{\text{H}} + \gamma_{\text{C}}} \frac{\rho_{\text{DM}}}{\omega_{\text{DM}}} \times \sin(2\omega_{\text{DM}} t - 2\mathbf{k} \cdot \mathbf{r} + \phi) \mathbf{v} \cdot \hat{\mathbf{e}}_z. \quad (3.16)$$

where g_{quad} , having dimensions of inverse energy, parameterizes the ALP quadratic coupling strength to nuclear spins.

There are two possible interactions of dark photons with nuclear spins that can be detected with CASPER-ZULF: the coupling of the dark electric

field to the dark electric dipole moment (dEDM) and the coupling of the dark magnetic field to the dark magnetic dipole moment (dMDM). The equivalent magnetic field amplitudes are:

$$B_{\text{dEDM}}^{\text{eff}}(t) = \frac{2}{3} \frac{g_{\text{dEDM}}}{\gamma_{\text{C}} + \gamma_{\text{H}}} \sqrt{2\hbar c^3 \rho_{\text{DM}}} \times \cos(\omega_{\text{DM}} t + \phi) \boldsymbol{\epsilon} \cdot \hat{\boldsymbol{e}}_z, \quad (3.17)$$

and

$$B_{\text{dMDM}}^{\text{eff}}(t) = \frac{2}{3} \frac{g_{\text{dMDM}}}{\gamma_{\text{C}} + \gamma_{\text{H}}} \frac{v}{c} \sqrt{2\hbar c^3 \rho_{\text{DM}}} \times \cos(\omega_{\text{DM}} t + \phi) \boldsymbol{\epsilon} \cdot \hat{\boldsymbol{e}}_z, \quad (3.18)$$

with coupling constants g_{dEDM} and g_{dMDM} (having dimensions of inverse energy) and dark photon field polarization $\boldsymbol{\epsilon}$.

The experimental sensitivity to real magnetic fields then directly translates to sensitivity to the coupling constants g_{aNN} , g_{quad} , g_{dEDM} , and g_{dMDM} . Inverting equations (3.15)–(3.18) yields the corresponding conversion factors from magnetic field to the dark-matter coupling constants:

$$\delta g_{\text{aNN}}(\omega) \approx \left[1.3 \times 10^8 \frac{\text{GeV}^{-1}}{\text{T}} \right] \delta B(\omega), \quad (3.19)$$

$$\delta g_{\text{quad}}(\omega) \approx \left[190 \frac{\text{GeV}^{-1}}{\sqrt{\text{T}} \sqrt{\text{rad/s}}} \right] \sqrt{\omega \cdot \delta B(\omega)}, \quad (3.20)$$

$$\delta g_{\text{dEDM}}(\omega) \approx \left[1.3 \times 10^5 \frac{\text{GeV}^{-1}}{\text{T}} \right] \delta B(\omega), \quad (3.21)$$

$$\delta g_{\text{dMDM}}(\omega) \approx \left[1.3 \times 10^8 \frac{\text{GeV}^{-1}}{\text{T}} \right] \delta B(\omega). \quad (3.22)$$

Here we have used $\gamma_{\text{C}}/2\pi = 10.70 \text{ MHz} \cdot \text{T}^{-1}$ and $\gamma_{\text{H}}/2\pi = 42.57 \text{ MHz} \cdot \text{T}^{-1}$, $v \approx 10^{-3} c$ and $\rho_{\text{DM}} \approx 0.4 \text{ GeV}/\text{cm}^3$.

3.4 Results

3.4.1 Signal processing: post-processing phase cycling

The expected dark matter coherence time (~ 14 hours for a particle with 19 Hz Compton frequency) is much longer than the nuclear spin coherence time in ^{13}C -formic acid (~ 10 s). Taking advantage of this mismatch, we introduce the post-processing phase-cycling technique shown in Fig. 3.2, which consists of incrementally phase shifting the transient spectra and subsequently averaging them together. If the phase increment matches the phase accumulated by the oscillating field between each transient acquisition, the sidebands add constructively. This allows coherent averaging of the complex spectra, such that the signal-to-noise ratio scales as $N^{1/2}$, where N is the number of transients. Because the dark-matter Compton frequency is unknown, it is necessary to repeat this operation for a large number of different phase increments (at least many as the number of transient acquisitions).

For each transient acquisition, the sample is prepared in the same initial state, which determines the phase of the J -coupling peaks. When averaging the transients together, the J -coupling peaks' amplitude and phase remain constant, while the uncorrelated noise is averaged away, thus increasing the SNR as the square root of the total integration time T_{tot} , i.e. $\text{SNR}(T_{\text{tot}}) \propto T_{\text{tot}}^{1/2}$.

However, the dark-matter-related information resides not in the J -coupling peaks, but in their sidebands. The external bosonic field oscillates at an unknown frequency and its phase at the beginning of each transient

acquisition is unknown. This phase directly translates into the phase of the sidebands in the transient spectra: while the phase of the J -coupling peaks is identical from one transient acquisition to another, the phase of the sidebands varies. As a result, naively averaging the transient spectra averages the sidebands away, thus removing the dark-matter-related information from the resulting spectrum.

Here we use a post-processing phase-cycling technique which enables coherent averaging of the spectra in the frequency domain, even for transient signals for which no obvious experimental phase-locking can be achieved due to the unknown frequency of the signal. The method is similar to acquisition techniques in which an external clock is used to register the times of the transient acquisitions and post-processing phase shifting of the transient signals is employed to recover the external field's phase [84, 85].

The method relies on the fact that the bosonic field's phase at the beginning of each transient acquisition is unknown but not random. Indeed we recall that the bosonic fields remain phase coherent for $\sim 10^6$ oscillations, which for frequencies below 19 Hz is longer than the total integration time (14 hours). Thus, precise knowledge of the transient-signal acquisition starting times enables recovery of the phase of the bosonic field.

A full description of this averaging method is given in the appendix A.1.2. Each transient spectrum is incrementally phase shifted prior to averaging. If the phase shift is equal to the phase accumulated by the bosonic field between two transient acquisitions, then the phase stability of the J -coupling peaks is shifted to their sidebands which can thus be coherently averaged.

Considering that the frequency of the bosonic field is unknown, the correct

phase shift is also unknown. Thus, we repeat the operation for 2001 different phase increments between $[-\pi, \pi]$, yielding 2001 averaged spectra, one of which being averaged with the phase increment such that the sidebands are coherent.

To demonstrate the viability of this method, a small magnetic field was applied with amplitude 0.24 nT oscillating at 0.73 Hz, to simulate a dark-matter field. Using this processing technique, the SNR of the sidebands scales as $\text{SNR}(T_{\text{tot}}) \propto T_{\text{tot}}^{1/2}$ (see Fig. A.3 in the appendix), as expected during a coherent averaging procedure. This is a dramatic improvement over the alternative power-spectrum averaging (typically implemented for sets of incoherent spectra), which would yield a $T_{\text{tot}}^{1/4}$ scaling.

3.4.2 Calibration

We remind the reader that the energy shifts produced by $D_z(t)$ are equivalent to those produced by a real magnetic field with amplitude

$$B_{\text{wind}}(t) = \frac{2}{3} \frac{g_{\text{aNN}}}{\gamma_{\text{H}} + \gamma_{\text{C}}} \cdot D_z(t) \propto \frac{g_{\text{aNN}}}{\gamma_{\text{H}} + \gamma_{\text{C}}} \sqrt{\rho_{\text{DM}}} \sin(\omega_{\text{DM}} t) \mathbf{v} \cdot \hat{\mathbf{e}}_z. \quad (3.23)$$

Similar relationships for dark-photon and quadratic-wind couplings were provided by Eqs. (3.16)–(3.18).

Based on Eq. (3.23), the sensitivity of the experiment to dark matter was calibrated by applying a real oscillating magnetic field of known amplitude and frequency and measuring the amplitude of the resulting sidebands in the coherently averaged spectrum.

The experimental sensitivity is defined by the ability to observe dark-matter-induced sidebands above the magnetometer noise floor. In order

to show that the sideband amplitude scales with the modulation index, B/ω , and does not present unusual scalings due to experimental errors, two calibration experiments were performed. A first calibration was performed by varying the amplitude of the AC-field from 95 to 310 pT while holding the frequency constant at $\omega = 2\pi \times 0.73$ Hz. Then the amplitude was held at 160 pT while varying the frequency, from 0.45 to 1.7 Hz. The results of this experiment are shown in Fig. A.1 in the appendix. Similar experiments were performed to determine the minimum detectable frequency (see A.1.7 in the appendix). Based on this calibration, we can extrapolate the expected sidebands amplitude, $A_s(B, \omega)$, for any field of amplitude B and frequency ω :

$$A_s(B, \omega) = 5.53 \times 10^{-6} \text{ rad.s}^{-1} \times \frac{B[\text{T}]}{\omega[\text{rad.s}^{-1}]}. \quad (3.24)$$

Then the magnetometer noise level determines the smallest detectable driving field, which is then converted to dark-matter coupling bounds via eqs. (3.19)-(3.22).

3.4.3 Search and analysis

The dark matter search data were acquired and processed as described above, but without a calibration AC-magnetic field applied.

For each Compton frequency, the appropriate phase increment is computed, which identifies the corresponding coherently averaged spectrum to be analyzed. The noise in the spectrum defines a detection threshold at the 90% confidence level (further details in A.1.5 in the appendix). When the signal amplitude at the given frequency is below the threshold, we set limits

on the dark matter couplings to nuclear spins at levels determined by the calibration and effective-field conversion factors (see Sec. 3.3.4). If the signal is above the threshold, a more stringent analysis is performed by fitting the coherently averaged spectrum to a four-sideband model. When the fit rules out detection, the threshold level is again used to set limits.

In case of an apparent detection, further repeat measurements would need to be performed to confirm that the signal is persistent and exhibits expected sidereal and annual variations.

3.4.4 CASPEr-ZULF search results: constraints on bosonic dark matter

The results of the CASPEr-ZULF search for axionlike particles are given in Fig. 3.3. The frequencies presenting sharp losses in sensitivities at 0.21, 1.69, and 2.16 Hz were the ones for which the nearest optimal phase increment was close to zero, thus presenting maximal-amplitude J -coupling peaks, raising the detection threshold (see discussion in 3.5 in the appendix). The red-shaded area labeled “CASPEr-ZULF” corresponds to upper bounds on nuclear-spin couplings to dark matter consisting of ALPs at the 90% confidence level. This represents our current sensitivity limitation after 850 30-second transient acquisitions using samples thermally polarized at ~ 1.8 T. The “CASPEr-ZULF Phase II” line corresponds to the projected sensitivity of a future iteration of this work that will use a more sensitive magnetometry scheme to measure a larger sample with enhanced (non-equilibrium) nuclear spin polarization.

Figures 3.4 and 3.5 show the search results for the ALP quadratic

interaction and dark photon interactions, respectively. No signal consistent with axion, ALP or dark-photon fields have been observed in the red-shaded areas. The two different limits given in Fig. 3.5 were obtained using the same data set but analyzed by assuming two orthogonal initial polarizations of the dark-photon field (see A.1.8 in the appendix).

In all cases, the search bandwidth was limited from below by the finite linewidth of the J -resonance peaks, preventing to resolve sidebands at frequencies lower than 45 mHz. Due to the finite coherence time of the dark-matter fields (corresponding to $\sim 10^6$ oscillations), the bandwidth's upper limit (19 Hz), is the highest frequency which can be coherently averaged after 14 hours of integration time. The sensitivity fall off is due to the sidebands' amplitude scaling as the modulation index. Further details are given in A.1.4 and A.1.7 of in the appendix.

3.5 Discussion

This work constitutes demonstration of a dark-matter search utilizing NMR techniques with a coupled heteronuclear spin system. The results provide new laboratory-based upper bounds for bosonic dark matter with masses ranging from 1.86×10^{-16} to 7.85×10^{-14} eV, complementing astrophysical bounds obtained from supernova SN1987A [44, 81].

Our data analysis provides a method to perform coherent averaging of the bosonic-field-induced transient signals. This method should prove useful for other experiments seeking to measure external fields of unknown frequency using a detector with a comparatively short coherence time. Conveniently, this phase-cycling approach also suppresses the carrier-frequency signals,




Figure 3.4: ALP wind-nucleon quadratic coupling parameter space. The CASPER-ZULF region is excluded by this work (90% confidence level) using a thermally polarized sample (data averaged over 850 transient acquisitions of 30 s each). Other regions of this figure are defined in the caption of Fig. 3.3.

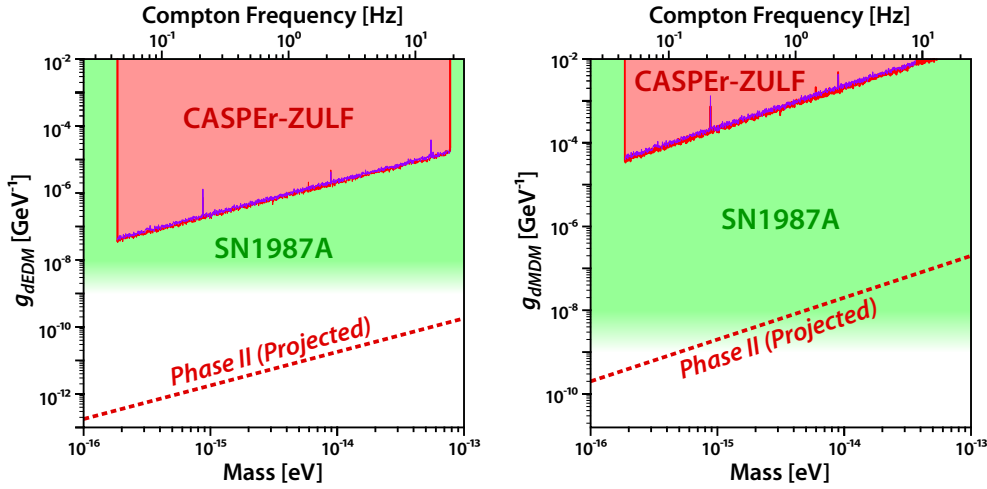


Figure 3.5: Left: dark photon-nucleon dEDM coupling parameter space. The SN1987A region represents existing limits for ALPs from supernova SN1987A cooling [44, 81] adjusted to constrain dark photons as discussed in Ref. [86]. Right: dark photon-nucleon dMDM coupling parameter space. The CASPER-ZULF regions are excluded by this work (90% confidence level) using a thermally polarized sample (data averaged over 850 transient acquisitions of 30 s each). The red and purple lines correspond to the case where the dark-photon field polarization is along the \hat{e}_1 and \hat{e}_3 axes of the non-rotating Celestial frame, respectively (see A.1.8 in the appendix). The dashed lines correspond to the sensitivity of a planned second phase of CASPER-ZULF, with a projected $\sim 10^5$ factor increase in sensitivity.

which would otherwise increase the detection threshold via spectral leakage. As this method is applied during post-processing, it does not require modification of the experiments provided that the data to be analyzed have been time stamped.

Deep neural network for physics-signal analysis

4.1 Foreword

This chapter is mainly composed of the contents of the following article:

Antoine Garcon, Julian Vexler, Dmitry Budker, Stefan Kramer.
**Deep Neural Networks to Recover Unknown Physical
Parameters from Oscillating Time Series.** *preprint*
arXiv:2101.03850, 2021.

Deep neural networks (DNNs) have been successfully used in a wide variety of tasks, such as regression, classification (e.g. in image or speech recognition [87, 88]), and time-series analysis. They are known for being able to construct useful higher-level features from lower-level features in many applications. However, these feature representations frequently remain incomprehensible to humans. This property is one of the reasons why DNNs are not more widely used in physics, in which the approach to data exploration is usually drastically different.

Most systems studied in physics are well described by physical models, generally referred to as *equations of motion*. The experimental data are analysed with respect to a particular model. When doing so, the equations of motion are analytically or numerically solved, yielding a theoretical

description of the data-generating process. The resulting model generally includes a set of mathematical variables that can be adjusted to span the data. The true values of these variables are generally unknown and must be recovered. For that reason, we refer to them as *latent parameters*. The true latent parameters are approximated by comparing the data to the model, typically by fitting the model to the data. With this in mind, the ability of DNNs to find abstract representations of the data features rather than a quantitative generating process is generally seen as a limitation rather than an advantage by physicists. For that reason, DNNs are still often viewed as black boxes in physics and started to be used in the field only in recent years [89].

We find this to be a missed opportunity for the physics community. With physical models at hand, one can generate arbitrarily large volumes of synthetic data to train the DNNs, and later process real-world signals [90]. This circumvents many challenges of supervised learning during which DNNs are trained with data for which the true latent parameters (labeled data) need to be known. Making full use of this possibility, DNNs were recently trained on synthetic nuclear magnetic resonance (NMR) spectroscopic data, simulated by accurate physical models [91]. The large amount of labeled data generated this way enables convergence of the DNN, which is then used to process real NMR data with great accuracy. A similar approach, that is, starting training with synthetic data and continuing with real-world data, has become popular in robotics and autonomous driving.

Moreover, extensive work was done in order to disentangle and make sense of DNN representations. A notable example is that of the β -

variational autoencoder architecture [92]. Correlation loss penalties can also be used during DNN training, without prior knowledge of the data-generating process [93, 94]. These methods consist of penalizing the DNN if its feature representation becomes entangled during training. While doing so, the DNN is encouraged to produce an efficient or disentangled feature representation. While disentangled, the representations achieved through these methods are not readily interpretable and usually require further analysis.

Nonetheless, DNNs are being increasingly used in physics data processing, in particular for signal classification – during which unusual datasets are flagged for further analysis. It was shown that *Autoencoders* can effectively be trained on Large Hadron Collider particle-jet data to detect events or anomalies [95]. In this instance, the DNN is successfully able to increase the events' signal-to-noise ratio by a factor 6. Other searches in high-energy physics, including [96, 97], have recently been performed also with the aim of detecting data displacement from a null-hypothesis (no anomalies). All these searches seek to perform data analyses in a model-independent setting, that is, with minimal prior information or bias. More recently, DNNs have been applied to time-series processing in nano-NMR [98]. In nano-NMR settings, the noise model is complex and noise overpowers the weak signals, rendering standard data analyses inefficient. The DNN was tasked to classify signals (i.e. discriminating two frequencies) and outperformed full-Bayesian methods.

While often achieving great successes, to our knowledge most applications of DNNs in physics are geared toward classification problems. In addition, DNNs are still rarely employed for time-series analyses, although they are

the most common form of data acquired during physics experiments. In this chapter, we propose to use a DNN to disentangle components of monochromatic, amplitude- and frequency-modulated sine waves (AM/FM-sine waves respectively), arguably the most prevalent forms of time-domain signals in physics. The method yields similar performance as more standard analyses such as least-square curve fittings (LS-fits), during which the data-generating process is assumed to be known and a least-squares regression is performed to predict the signal’s latent parameters.

LS-fits, however, require the user to input latent-parameters initial guesses prior to regression. These initial guesses are the prior estimation of the true latent parameters and provide a starting point for the LS-fit gradient descent. The trained DNN, however, needs no initial guesses, thus requiring less prior information about the data-generating process. Indeed, we show that, precisely because DNNs find abstract data representations, they can be used in settings when prior knowledge exists, but is not complete, as it is particularly the case in “new-physics” searches [99], thus leaving space for data exploration and discoveries.

The first part of this chapter describes the synthetic data that we generate and use throughout this work, i.e. monochromatic, AM- and FM-sine waves time series, and their relevance to real-world physics experiments. We then describe our DNN architecture, which incorporates two tasks: A *Regressor* DNN ¹ performs a regression of the signal’s latent parameters that are known to be present in the data-generating process. In addition, an *Autoencoder* [100] denoises the signals by learning an approximation of the

¹We note here that throughout the paper the term *Regressor* is employed as opposed to a *Classifier*, not as an independent variable in a regression. Therefore, *Regressor* refers to the DNN predicting the signal’s latent parameters.

unknown latent parameters. As a benchmarking method, we evaluate the DNN by comparing its performance to an LS-fit with true initial guesses.

We later employ the DNN in realistic settings, when prior knowledge about the data-generating process is incomplete: LS-fit fidelity is typically highly sensitive to initial guesses, thus requiring the user to perform preprocessing work or to possess prior information in order to perform optimally. As a first application, we show that the DNN can be used to predict initial guesses for the model fit evaluation. While consistently converging to optimal solutions, the technique circumvents the usual difficulties arising from fitting signals, such as the need for initial-guesses exploration.

Next, we show that the DNN can be used when the user ignores if the time-series are monochromatic-, AM- or FM-sine waves, but still wishes to recover their main frequency component. In such settings, the user is generally required to repeat the analysis by exploring the space of data-generating processes and initial guesses. Using our architecture enables the user to input only the known information when performing the analysis. That is, the *Regressor* is tasked to recover the user-expected latent parameters while ignoring the existence of others. Because the *Autoencoder* needs no prior information, it is still able to capture unknown information.

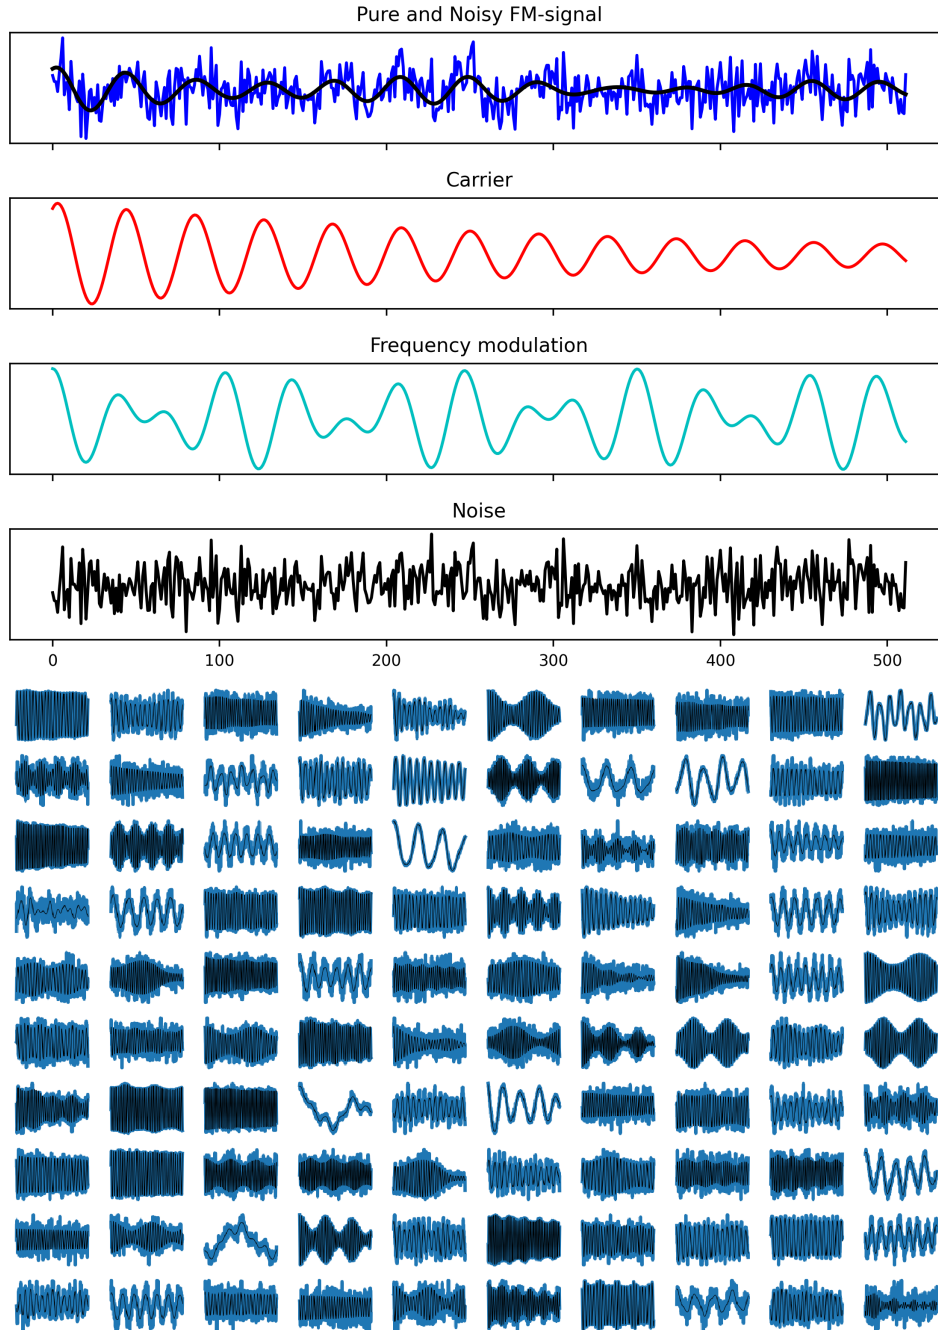


Figure 4.1: Examples of frequency-modulated sine wave (FM) synthetic time series. **Top:** pure and noisy FM-sine wave decomposition. Gaussian noise and frequency modulation are linearly added to decaying sine wave carrier. **Bottom:** random selection of noisy-input (blue) and pure-target (black) samples, illustrating the effect of the random latent parameter selection.

4.2 Data description and generation procedure

The time series studied here are exponentially decaying monochromatic, FM- and AM-sine waves. Gaussian noise is linearly added to the pure signals. An example of FM-signal is shown in Fig. 4.1 (top) alongside its sub-components (decaying-sinewave carrier, frequency-modulation signal, and noise).

Decaying monochromatic-sine waves appear and are prevalent in all fields of physics. They arise from solving the equations of motion of the two-level quantum system, or of the classical harmonic oscillator; to which a multitude of other physical systems can be mathematically reduced to. Notorious examples include the spin-1/2 particle in a DC magnetic field, the orbital motion of planets, or RLC circuits. In information theory, the two-level quantum system also provides a complete description of the qbit. Frequency and amplitude modulation generally arise from external factors such as oscillating magnetic or electric fields applied by the experimenters. Amplitude and frequency modulation of a carrier frequency are also the most common scheme of information communication links. Some form of Gaussian noise, while not necessarily always dominant, is in general present in any real-world signal. The statistical Gaussian noise formalism provides an accurate description of electronic thermal-noise, quantum shot noise, black-body radiation, and of White noise in general.

The time series are generated by propagating the time, t , from 0 to 511 (with length $T = 512$) in 1 s increments, and using the following formula:

$$\begin{aligned}
f_{Mono}(t) &= \cos(2\pi F_c t + \phi) \cdot e^{-t/\tau}, \\
f_{AM}(t) &= \cos(2\pi F_c t + \phi) \cdot e^{-t/\tau} \cdot (1 + I_m \cos(2\pi F_m t)), \\
f_{FM}(t) &= \cos\left(2\pi F_c t + \phi + \frac{0.01 I_m}{F_m} \cos(2\pi F_m t)\right) \cdot e^{-t/\tau},
\end{aligned}$$

where F_c and ϕ are the sine wave carrier frequency and phase, respectively. F_m and I_m are the modulation frequency and amplitude. Noise is linearly added to the signals after being sampled from the Gaussian distribution with zero mean and standard deviation σ .

Before each sample generation, the latent parameters are randomly and uniformly sampled within the following ranges:

$$\begin{aligned}
F_c &\in [10/T, 1/8], & \phi &\in [0, 2\pi], \\
F_m &\in [1/T, 1/16], & I_m &\in [0, 1], \\
\tau &\in [0.2T, 8T], & \sigma &\in [0, 1].
\end{aligned}$$

The range of F_c ensures the carrier frequency remains well within the Fourier and Nyquist limits. The modulation amplitude range ensures the majority of the signal's power remains in its first sidebands and carrier.

Noise is linearly added to the pure signals by sampling the Gaussian distribution with zero mean and standard deviation σ .

Most DNN implementations generally require input and target data to be normalized such as to avoid exploding and vanishing gradients during training [101, 102]. All signals and latent parameters are normalized to lie

within the 0-to-1 range prior to the application of the DNN. The phase ϕ is mapped to two separate parameters, $\phi \rightarrow \left\{ \frac{\sin(\phi)+1}{2}; \frac{\cos(\phi)+1}{2} \right\}$, such as to account for phase periodicity during loss computation, while keeping both targets properly normalized. All other latent parameters are normalized using their respective range. Prior to adding the noise, we normalize the pure signals such that the resulting noisy signals remains within the $[0, 1]$ range with mean 0.5. This normalization is performed identically for all signals.

Despite requiring only 6 latent parameters to generate the samples, these ranges enable a wide scope of functions to be realized. AM/FM-signals with minimum I_m reduce to decaying monochromatic-sine waves and reach 100% modulation with maximum I_m . The coherence time range is wide enough to span underdamped signals up to virtually non-decaying signals. These latent parameter ranges are wide enough such that they would encompass many foreseeable real-world signals. A random selection of FM-signals with and without noise is shown in Fig. 4.1 (bottom), illustrating the richness of the data in a more qualitative manner.

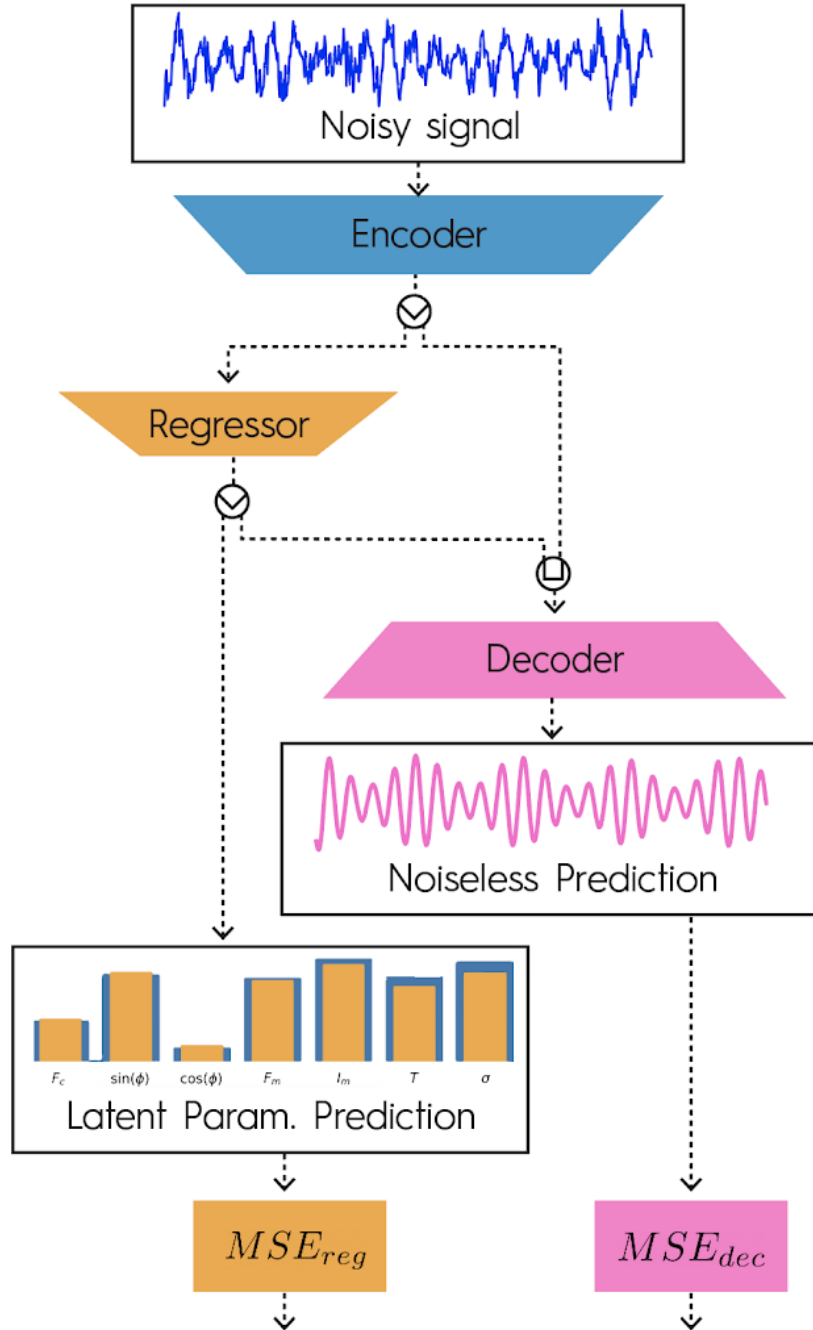
The choice of studying monochromatic, AM-, and FM-sine waves is not only motivated by their richness and prevalence in real-world physics experiments. Indeed, despite originating from different physical models and having different mathematical descriptions, the time series share similar visual features. As a result, within some range of parameters, even expert users could mistake the three generating processes. This is especially the case for weak modulations in the presence of noise, for which visual discrimination in time- or frequency-domain (inspecting the spectrum) may

be impossible. For all the reasons cited above, monochromatic-, AM- or FM-sine waves appear as good representative signals on which to perform our study. Nevertheless, the methods presented in this here can be applied to other types of signals as well.

4.3 DNN architecture and training procedure

The latent-parameters regression and signal denoising are performed by two separate architectures (Python code implementation is given in the appendix).

Denoising is performed by an *Autoencoder* architecture [100] composed of an *Encoder* followed by a *Decoder*. Noisy signals are first passed through the *Encoder*. The *Encoder* output layer has 64 neurons and thus produces a compressed representation of the input signal. Following this step, the *Encoder* output is passed through the *Decoder*, which decompresses the signal to its original size. This type of [*Encoder-Decoder*] architecture, is widely used, inter alia, for data denoising [103]. As the *Encoder* output dimension is smaller than the dimension of the input data, the *Encoder*'s output layer acts as an information bottleneck, or more specifically dimensionality reduction, thus encouraging the *Autoencoder* to capture relevant latent features while discarding noise or redundant information [100]. Latent-parameters regression is also performed while passing the data through the *Encoder*. The *Encoder* output is then passed through a third DNN referred to as the *Regressor*.



$$Loss = \beta \cdot MSE_{reg} + (1 - \beta) \cdot MSE_{dec}$$

Figure 4.2: Unified DNN architecture and loss description. The *Encoder* produces a reduced representation of the input noisy signals. The *Encoder* output is passed to the *Regressor*, which outputs the latent parameters' prediction. The *Encoder* and *Regressor* outputs are passed to the *Decoder*, which produces a noiseless prediction of the inputs. The *Regressor* and *Decoder* outputs are used to compute the regression and denoising losses, MSE_{reg} and MSE_{dec} , respectively. The loss used during backpropagation is a weighted sum of MSE_{reg} and MSE_{dec} using a bias parameter β .

The *Encoder* composed of 2[*Conv1D-Maxpool*] layers, followed by 2 *Dense* layers. The *Encoder* output layer has 64 neurons. The *Regressor* is composed of 2[*Conv1D – Maxpool*] followed by a *Conv1D* and 4 *Dense* layers. The output dimension of the *Regressor* is adjusted to the number of latent parameters that the *Regressor* is tasked to detect. The *Decoder* is composed of 1 *Dense-3[Conv1D-Maxpool-Upsampling]* layers, followed by a single *Conv1D* layer. The *Decoder* consists of a concatenation of the *Regressor* and *Encoder* outputs.

All activation functions are rectified linear units, with the exception of the *Regressor* and *Decoder* outputs, which are linear and sigmoid function, respectively.

After refining the base *Encoder*, *Regressor* and *Decoder*, we unify the three architectures into a single DNN as depicted in Fig. 4.2 such that the *Regressor* and *Decoder* share the same *Encoder*. As such, we pass the *Encoder* output to the *Regressor*, which predicts the signal’s latent parameters. The *Decoder* input then consists of a concatenation of the *Regressor* and *Encoder* outputs.

The latent parameter regression and signal-denoising losses are computed simultaneously (MSE_{reg} and MSE_{dec} , respectively). The loss used during backpropagation is computed as a weighted sum of MSE_{reg} and MSE_{dec} as follows:

$$\mathcal{L}_{oss} = \beta \cdot MSE_{reg} + (1 - \beta) \cdot MSE_{dec}, \quad \beta \in [0, 1], \quad (4.1)$$

where the hyperparameter β is the bias adjustment between the two tasks.

This architecture presents the advantage of enabling bias control via

a unique hyperparameter. Moreover, both networks are naturally trained at the same time rather than in alternate, thus accelerating training approximately two-fold and enabling high-momentum gradient optimizers.

To illustrate the architecture’s output, we train the DNN on AM-sine waves and show a prediction example in Fig. A.5 of the appendix, alongside the noisy input signal.

During training, the *Regressor’s* target data consists of the latent parameters, and the *Decoder* target data are the noiseless signals. For both, the loss function is the mean squared error (*MSE*). The optimized architectures achieve sufficient performance, while keeping the number of trainable parameters under 1 million, such as to be able to perform training on a modern laptop GPU under 12 hours for a typical training session of 20 training sets of 100’000 samples, over 10 epochs. Due to the number and characteristics of the instances, asymptotic loss is reached within a small number of epochs. In general, increasing the number of training set instances was more beneficial than increasing the number of epochs.

To illustrate the effect of the bias parameter, we train the unified DNN on identical FM-sine waves datasets with varying values of β . For this experiment, training is performed using 12 training sets of 100’000 randomly generated samples for 10 epochs. Because the number of synthetic samples is large and the latent parameters are continuous random variables, overfitting (controlled by a validation set, unseen during training) was never an issue.

The performance of the trained DNN is evaluated using a test set of 100’000 randomly generated FM-samples, which were unseen during training. Figure 4.3 shows the test-sample losses for the denoising (top) and regression




Figure 4.3: FM-sine waves test-data prediction errors: total weighted loss βMSE , Denoising loss MSE_{dec} and Regression loss MSE_{reg} for varying values of β . Setting $\beta = 0$ or $\beta = 1$ fully biases training toward one of the two tasks, preventing the negatively-biased tasks to reach sufficient performance. Middle-range values enable both tasks to be learned simultaneously.

(bottom) tasks after training. Setting $\beta = 0$ fully biases training towards the denoising tasks, which achieves best performance, while the parameter regression yields the worst results; vice versa for $\beta = 1$. This behaviour is also observed in Fig. A.8 in the appendix, which shows the validation losses during training. The training curves show that extremum values of β prevents validation loss improvement of the negatively-biased task. Middle-range values enable both tasks to be learned simultaneously.

We find that the best values of β are those for which the initial β -weighted regression and denoising losses are within the same order of magnitude. As a result, determining a good value for β is a trivial task: A single forward pass is performed to obtain the initial values of MSE_{reg} and MSE_{dec} . Regardless of the type of data (monochromatic-, AM- and FM-samples), DNNs trained with $\beta = 0.001$ achieve good overall performance (lowest weighted total loss) and little bias towards any of the tasks. This value of β is employed throughout the entire chapter. For all that follows, training is always performed using 20 training sets of 100'000 randomly generated samples for 10 epochs. This training is always enough to reach asymptotic loss, while exhibiting no noticeable overfitting. Training can be performed on decaying monochromatic-, AM-, FM-sine waves or a combination of all three processes.

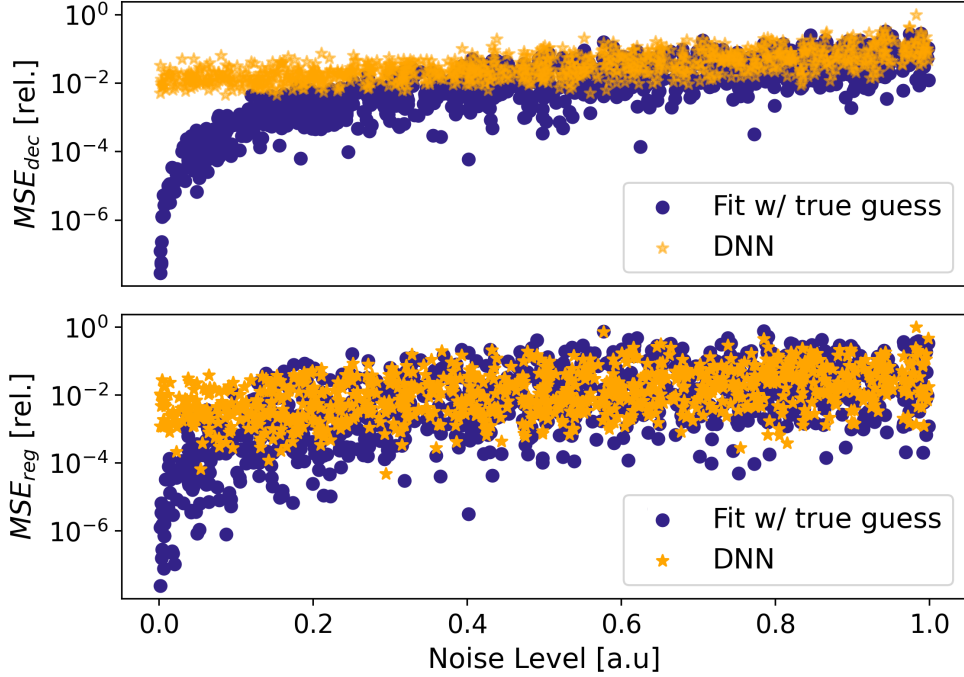


Figure 4.4: Comparison of DNN post-training performance to LS-fits with true latent-parameters initial guesses for 1000 random monochromatic, decaying sine waves from the test set (unseen during training). The denoising (MSE_{dec} , top) and latent-parameters relative regression losses (MSE_{reg} , bottom) are sorted by increasing noise levels. The DNN was trained on monochromatic sine waves samples. MSE_{reg} is the MSE from the true latent parameters to the predicted latent parameters. For the DNN, MSE_{dec} is the MSE from the true noiseless signal to the *Decoder* noiseless-signal prediction. For the LS-fits, MSE_{dec} is computed similarly, but the noiseless-signal prediction is generated by inputting the predicted latent parameters in the noiseless data-generating process. The LS-fit with true initial guesses vastly outperforms the DNN for low-noise signals but both systems reach similar performance for high-noise.

4.4 Results

4.4.1 Post-training performance evaluation

We train the DNN on a random selection of decaying monochromatic sine waves (no modulation). The training, validation, and test samples are generated using random frequency, phase, coherence time, and noise levels. After training, we evaluate the DNN performance by comparing its prediction error to an LS-fit using the Python Scipy library.

When performing the LS-fit, the input data is the noisy signals and the objective function is with respect to the noiseless data-generating process. The LS-fit then produces predictions of the true latent parameters. To this end, the LS-fit requires latent-parameters initial guesses to start the gradient descent. The initial guesses used here are the true latent parameters (i.e. true frequency, phase, and coherence time). After gradient descent, we use the LS-fit outputs to generate a prediction of the noiseless signals. This is done by inputting the LS-fit latent-parameters predictions in the data-generating process. The LS-fit and DNN performance are then compared in two ways: (i) the latent-parameters regression loss is the MSE from the true latent parameters for both the LS-fit and DNN (MSE_{reg}), and (ii) the denoising error is the MSE from the true noiseless signals for both the LS-fit and DNN (MSE_{dec}). Note that this comparison drastically favors the LS-fit, which then constitutes a good benchmark method. Indeed, in any practical applications the true value of the latent parameters are hidden from the user, and LS-fits are employed precisely to approximate them.

A random selection of 1000 noisy signals from the test set is processed

using this method. Figure 4.4 shows the relative MSE_{reg} and MSE_{dec} for both the DNN and LS-fit sorted by noise level (examples of signals with extremum noise levels, alongside LS-fit and DNN predictions are shown in Fig. A.6 in the appendix). A similar evaluation is performed using AM-samples. In this experiment, the DNN is specifically trained on AM-samples. Examples of such samples are given in Fig. A.6 in the appendix. Figure A.7 shows the prediction errors of all samples, for both the DNN and LS-fit sorted by noise levels.

For both monochromatic and AM-signals, the DNN performs generally worse than the LS-fit for low-noise signals. However, the DNN reaches LS-fit performance-level once the noise reaches the top half of the allowed range, while requiring no initial guesses. The latent-parameters regression follows a similar trend. We note that, in general, DNN outputs are less sensitive to noise, and the performance is more consistent throughout both datasets. These results show that our architecture is a good alternative to LS-fits for time-series analysis, as it reaches acceptable performance when benchmarked to standard LS-fits with true guesses, while needing no initial guesses.

4.4.2 DNN-assisted LS-fit

We now wish to apply our DNN in more realistic settings. Fitting oscillating time series using LS-fits is notoriously difficult because the MSE is in general a non-convex function of the latent parameters and possesses numerous local minima. Consequently, the quality of the LS-fit is highly dependent on the initial guesses in addition to the noise. In the previous experiments, LS-fits were only performed as a benchmark method, and the initial guesses were

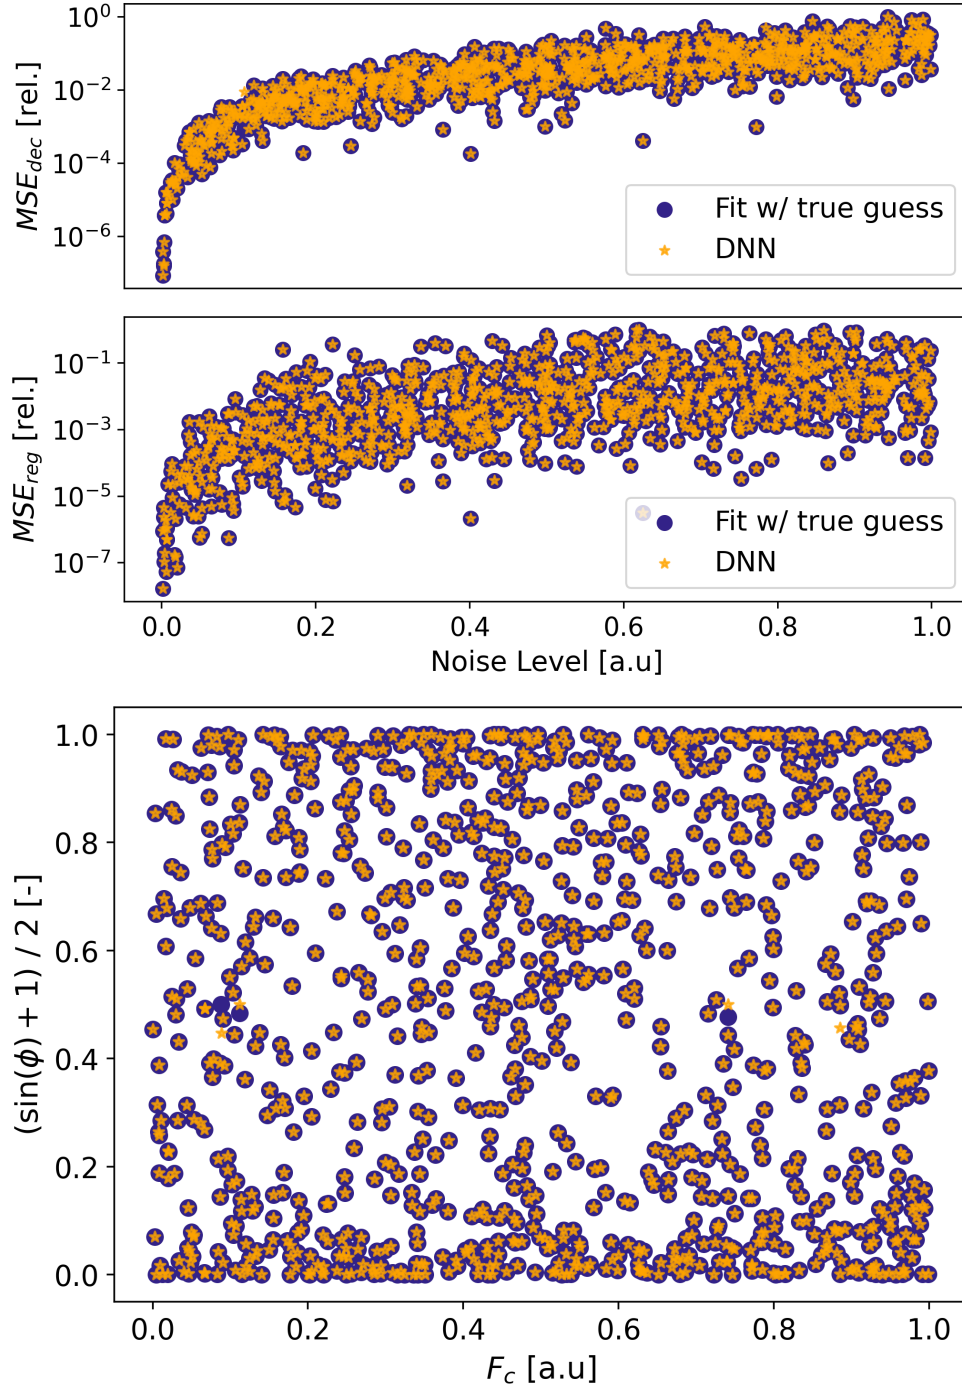


Figure 4.5: DNN latent-parameters predictions used as initial guesses for DNN-assisted fits. Comparison to LS-fit with true initial guesses. **Top:** The denoising (MSE_{dec}) and latent-parameters relative regression losses (MSE_{reg}) are sorted by increasing noise levels. See Fig. 4.4 for MSE_{reg} and MSE_{dec} computation methods. **Bottom:** Phase and carrier frequency predictions for the DNN-assisted fits and LS-fits. Both methods converge to the same losses and predictions for over 99% of the samples. The DNN and data employed here are identical as in Fig. 4.4.

the true latent parameters. In any real-world setting, the user must perform preprocessing work or use prior information to find initial guesses leading to the global minima. We propose to employ the DNN as a preprocessing tool to assist LS-fit in the situation when the user possesses no prior information about the initial guesses and wishes to recover the signal’s latent parameters. The sine wave samples from the previous experiment are fitted while using the DNN latent predictions as initial guesses. Results of this experiment are shown in Fig. 4.5 alongside LS-fits with true initial guesses results.

Because the DNN predictions are always within the vicinity of the true parameters, almost all DNN-assisted LS-fits converge to optimal solutions. In settings when the initial guesses are unknown or samples are numerous, the user can initially train the DNN on synthetic data and use it for DNN-assisted fits. As the latter performs optimally regardless of the noise level, this enables fast and accurate analysis of large datasets by removing the need for initial guesses exploration. This conjunction use of DNNs and LS-fits also enables access to prediction uncertainties, typically in the form of a covariance matrices, which DNNs typically lack [104–106].

4.4.3 Partial information regression and denoising

In the experiments presented above, the data-generating process was assumed to be fully known by the user. The DNN or DNN-assisted LS-fits were employed to recover the signal latent parameters and denoise the signal. We now wish to explore the possibility of employing the DNN in a situation where the data-generating processes to be explored are multi-fold and guesses must be done. This is typically the case in “new-physics searches”

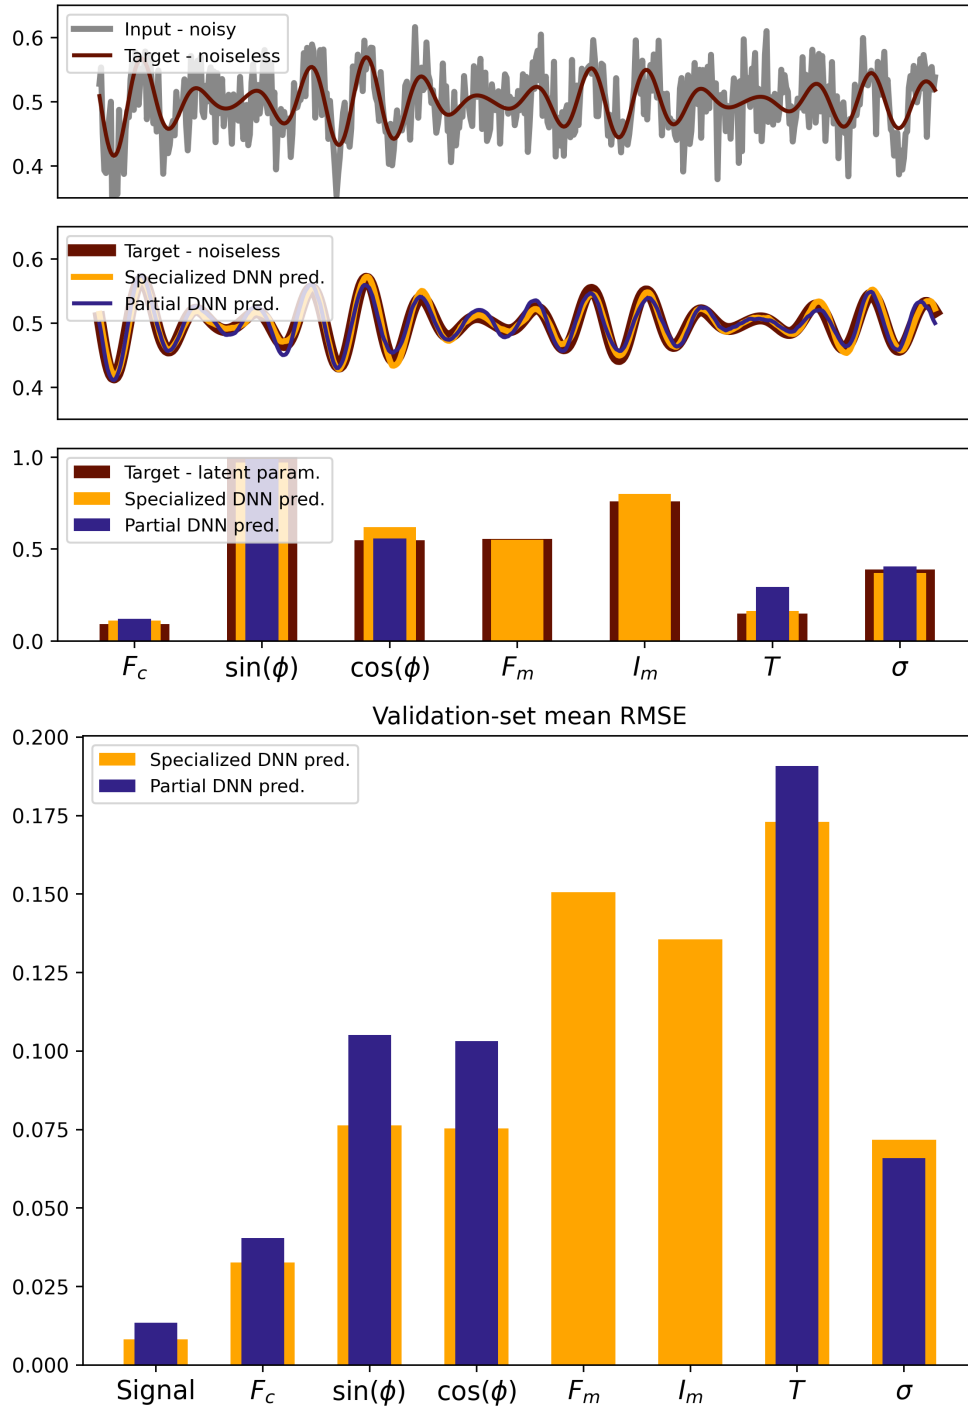


Figure 4.6: Performance comparison of the specialized DNN (trained on AM-sine waves, tasked to denoise signals and recover *all* latent parameters of AM-sine waves) and of the partial DNN (trained on monochromatic, AM- and FM-sine waves, tasked to denoise signals and recover the carrier frequency, phase, coherence time and noise level only). **Top:** Example of noisy input AM-signal, alongside specialized- and partial-DNN denoised and latent predictions. **Bottom:** Individual latent parameters and signal denoising root mean squared error (RMSE), averaged over the whole AM-sinewave test set (100'000 samples) for both DNNs.

experiments [99], during which hypothetical and undiscovered particles may cause signals deviating from the null-hypothesis (i.e. no new particles). As the hypothetical particles are numerous, they may have many potential effects on the signals. We take the situation in which a potential external source could modulate a carrier signal produced by the experiment, as it is sometimes the case for bosonic dark-matter [107].

Specifically, we study the case in which the end-user is aware of the existence of an oscillation in the signal provided by the experimental setup. The user ignores if the signal is monochromatic, amplitude or frequency modulated. Nonetheless, the user wishes to recover the frequency, phase, and coherence time of the expected oscillation.

In this situation, the typical approach is to test all allowed processes by varying the LS-fits objective functions and explore the space of initial guesses for each process. This approach presents a new set of challenges, as this exploration is time consuming and sometimes unrealistic, if the data is too large or if too many processes are to be tested. Moreover, in some situations, all guesses can be wrong.

We show that it is possible to perform the regression and denoising with partial prior information about the physical process producing the data. That is, the DNN is tasked to perform the regression only on the narrow set of latent parameters that exist across all models: frequency, phase, coherence time, and noise level. However, the DNN ignores any form of modulation. This is done by decreasing the number of neurons in the *Regressor*'s output layer. The DNN is then trained on signals from every explored model (monochromatic, AM and FM). We now refer to this DNN as the *partial*

DNN (ignoring the existence of particular modulation type).

After training, we compare the performance of the partial DNN to a specialized DNN, trained specifically on AM signals, which performs a regression of all latent parameters. Figure 4.6 shows the MSE_{reg} and MSE_{dec} averaged over the AM-sine wave test set (100'000 samples) for both the AM-specialized DNN and the partial DNN. Because the partial DNN is trained on a wider variety of data, its prediction accuracy is, on average, lower than the specialized DNN. However, both DNNs performance remain close and the partial DNN is sufficiently accurate to perform the previously described experiments.

Using this method, the user's prior information is encoded into the *Regressor* architecture and training data. The *Regressor* then captures the expected latent parameters co-existing across the entire training set. The *Encoder* and *Decoder* remain unchanged and are still able to capture unknown latent parameters by reproducing noiseless signals. As a result, prior to the analysis, the user need not be fully aware of the data generating physical model but can instead train the DNN on a wider class of models (in this instance, the DNN was trained assuming the presence of any type of modulation). This feature is of particular importance in exploration of data obtained from new physics searches (such as dark-matter searches), in which the multitude of allowed physical models enables various forms of emerging signatures within the data. Thus, in addition to removing the need to iteratively explore models, this method enables weaker and partial prior information to be employed, while leaving space for signal exploration and unexpected discoveries.

4.5 Discussion

We have presented an efficient DNN that combines the denoising of times series and regression of their latent parameters. The DNN was trained and evaluated on synthetic monochromatic, frequency- and amplitude-modulated decaying sine waves with Gaussian noise; some of the most prevalent forms of signals acquired in physics.

For high-noise signals, the DNN reaches same levels of precision as an LS-fit with true initial guesses, in spite of the DNN needing no guesses at all. In addition, the architecture requires no hyperparameter fine tuning to perform consistently. Moreover, because large volumes of synthetic training data can be generated, the DNN is quickly adaptable to a broad range of physical signals. This makes our architecture a good alternative to LS-fits for analysing large volumes of data, when fitting individual signals requires too much computation or user time.

The DNN architecture is flexible and can accommodate for various levels of user prior information. First, the DNN was used to assist LS-fits and predict initial guesses, unknown by the user. In this situation, DNN-assisted LS-fits consistently converge to the optimal solutions. Moreover, the regression task can be adapted to accommodate for partial prior information about the data-generating process. The known latent parameters are encoded in the *Regressor* and training data, while the *Decoder* helps the *Encoder* to still capture unknown signal features, thus leaving space for data exploration and discoveries.

Because training is done on arbitrarily large volumes of synthetic data, raw performance could be improved by increasing the number of trainable

parameters such as adding more layers or neurons, without too much concern for overfitting.

Therefore, we believe this architecture is readily applicable to existing physics experiments, in particular bosonic dark-matter searches [107–109], in which large quantities of data are to be analyzed with partial prior information.

Summary and outlook

5.1 Summary and personal notes

My work in Prof. Dmitry Budker's group started by catching up with my colleagues and slowly understanding how axions and ALPs were expected to generate CASPER's NMR signal. While doing so, we realized that for subHz frequencies, CASPER's original resonant CW-NMR measurement scheme was not optimal, and that a non-resonant sideband search would yield higher sensitivities. This part of our work is described in the second chapter of this thesis.

Later on, my work focused on applying the sideband scheme not on the main CASPER experimental setup but on the ZULF NMR setup available at the time. This implementation led us to the understanding of the effect of bosonic dark-matter fields on small organic molecules such as formic acid (effectively a two-spins system), rather than on pure xenon spins as was originally proposed.

While performing the experiments, we subsequently found a data-processing scheme enabling restoration of the sidebands phase coherence, even in conditions in which the ability to do so was not obvious. This experiment provided new experimental bounds on the existence of ALPs

and dark-photons. This part of our work was described in the third chapter of this thesis.

Following up on these experiments, my work focused on analyzing data from complementary experiments, published elsewhere [110–112]. This finally led me to concurrently join Prof. Stefan Kramer’s group who enabled me to explore applications of machine learning, in particular deep learning to physics-signal processing. A first approach in applying neural networks to disentangle noisy time series is presented in the final chapter of this thesis.

5.2 Outlook

5.2.1 General improvements

Future experiments will be carried out with increased integration time. We recall that the bosonic dark-matter fields are coherent for a time τ_{DM} on the order of 10^6 periods of oscillation. The phase-cycling procedure depicted in Fig. 3.2 is valid for data sets with total time less than τ_{DM} . For integration times longer than the coherence time of the bosonic field $T_{\text{tot}} > \tau_{\text{DM}}$. In such a situation, one could for example coherently average the data in sets of duration τ_{DM} using the phase-cycling procedure. This yields $T_{\text{tot}}/\tau_{\text{DM}}$ sets of coherently averaged data. To profit from longer integration time and further increase the SNR, these sets can be incoherently averaged by averaging their PSDs, yielding an overall SNR scaling as $(T_{\text{tot}}\tau_{\text{DM}})^{1/4}$ (see Supplementary Materials in Ref. [36]).

5.2.2 Complementary Searches

In order to increase the bandwidth of the experiment (see A.1.7), we propose complementary measurement procedures. As the amplitude of the sidebands scales as $1/\omega_{\text{DM}}$, the sensitivity of the experiment decreases for higher frequencies. To probe frequencies ranging from ~ 20 to 500 Hz (corresponding to bosonic masses of $\sim 8 \times 10^{-14}$ to 2×10^{-12} eV), it was shown in Ref. [71] that a resonant detection method will be more sensitive than the current frequency-modulation-induced sidebands measurement scheme. Resonant AC fields can induce phase shifts in the J -coupling peaks [113]; cosmic fields can induce the same effect. By gradually varying the magnitude of a leading magnetic field, one can tune the splitting of the J -coupling multiplets to match the dark-matter field frequency. Such a resonance would manifest itself by shifting the phase of the J -coupling peak.

For frequencies below ~ 45 mHz (corresponding to bosonic masses $\lesssim 2 \times 10^{-16}$ eV), the sidebands are located inside of the J -coupling peaks and the experimental sensitivity drops rapidly. This represents the lower limit of the bandwidth accessible by the frequency-modulation-induced sidebands measurement scheme presented in this work. To probe down to arbitrarily low frequencies, another measurement scheme has been implemented based on a single-component liquid-state nuclear-spin comagnetometer [114]. Further details and results of this scheme are presented elsewhere [82].

5.2.3 Parahydrogen hyperpolarization

In order to search a greater region of the bosonic dark-matter parameter spaces, several sensitivity-enhancing improvements are planned for the next phase of the experiment. In this work, nuclear spin polarization was achieved by allowing the sample to equilibrate in a 1.8-T permanent magnet, which yields a ^1H polarization $\lesssim 10^{-5}$. For the next phase of the experiment, a substantial sensitivity improvement can be obtained by using so-called “hyperpolarization” methods to achieve much higher, non-equilibrium nuclear-spin polarization. Current efforts are focused on the implementation of non-hydrogenative parahydrogen-induced polarization (NH-PHIP)¹ [115]. Signal enhancement via NH-PHIP has been demonstrated at zero field [116] and after optimization is expected to increase nuclear spin polarization levels to at least 1%. Because parahydrogen can be flowed continuously into the sample, a steady-state polarization enhancement can be achieved [117], improving the experimental duty cycle.

Additional sensitivity enhancement will be provided by magnetometer improvements and use of a larger sample. In the experiments reported here, only about 50 μL of the sample contributed to the signal, which was detected from below with an atomic magnetometer with a noise floor around 10 fT/ $\sqrt{\text{Hz}}$. With a larger ($\gtrsim 1$ mL) sample hyperpolarized via NH-PHIP detected via a gradiometric magnetometer array with optimized geometry and sensitivity below 1 fT/ $\sqrt{\text{Hz}}$, we anticipate an improvement by $\gtrsim 10^5$ relative to the results presented here.

¹Non-hydrogenative parahydrogen-induced polarization methods are often referred to with the acronym SABRE, for Signal Amplification by Reversible Exchange of parahydrogen.

Following the publication of the work presented in Chap.3 we started the design and construction of our own parahydrogen generator. The generator is currently fully functional and allows for high volumes of parahydrogen to be produced, stored and flowed through the NMR samples. While the generator is used for other scientific work [118], no dark-matter search has been conducted with it yet. It is my hope that a future group of students will do so.

5.2.4 Search for axions with spin-based amplifiers

After the publication of our results, a collaboration was started with Prof. Xinhua Peng and Dr. Min Jiang. Indeed, their MASER setup located at the University of Science and Technology of China (Hefei) presented great potential towards more sensitive axion searches. We remind the reader that the search presented in this thesis was performed using thermally polarized spins (formic acid sample) located a few mm away from the magnetometer's vapor cell. In Prof. Xinhua Peng setup, ^{129}Xe spins are trapped within the magnetometer vapor cells. As a result, the xenon vapor is hyperpolarized via spin-exchange collisions with the optically pumped ^{87}Rb atoms. Moreover, the lower spin density of the xenon vapor (versus liquid-state spins) is largely compensated by Fermi-contact enhancement. Following up on this idea, a resonant search was successfully carried out to completion [112], constraining the presence of axions in a much larger portion of the parameter space and with a ~ 4 orders of magnitude sensitivity improvement.

The same setup can be operated as a Floquet-state MASER [119, 120] by introducing a feedback magnetic field oscillating at the xenon spins Larmor

frequency. The coherence time of this self-oscillating nuclear spin resonator was shown to reach $\sim 10^4$ s. As a result, combining xenon hyperpolarization with this feedback system could conceivably offer further improvement at frequencies around 0.1 mHz. Work to perform this non-resonant sideband search is ongoing.

5.2.5 Further applications of machine learning

Further work to explore application of machine learning and DNNs in physics data processing is ongoing. The DNN architecture presented in this thesis could be augmented by adding an upstream classifier DNN-module, which could identify the type of signals being analyzed. Classified signals could then be processed via specialized versions of our architecture, trained on the corresponding type of signals.

Time-domain oscillations generally appear as peaks or peak multiplets in frequency-domain spectra. Frequency, amplitude, and phase information is then localized to narrow regions of the spectral data. For that reason, we believe further improvements could be attained by making use of frequency-domain information. We suggest to use Fourier transforms or power spectra as DNN inputs, in addition to the raw time series.

The proposed DNN architecture can be used to detect and approximate hidden features in time series data. The *Regressor* outputs a prediction of prior known parameters, but real signals could still contain unknown latent variables. These hidden latent variables can be detected and approximated by our DNN, as it also incorporates an *Autoencoder*-like structure. As such, the bottleneck layer contains a feature representation of the time series, used

by the *Decoder* to recreate the original signal. This bottleneck layer will be further investigated, in order to detect and specify hidden latent parameters.

We remain aware that in physics data analysis, a sole estimation of latent parameters often provides insufficient information. Standard analysis usually requires a quantitative estimation of the prediction uncertainty, often represented as error bars or confidence intervals. In LS-fits, this uncertainty is naturally obtained by maximizing the fit likelihood under the assumption of Gaussian distributed latent variables [121]. Despite extensive efforts, DNNs still lack the capacity for reliable uncertainty evaluation [104–106]. While the DNN-assisted LS-fit method presented above partially solves this issue, more work needs to be done in this area to further generalize DNN usage in physics signal processing.

Appendices

A.1 Chapter 3

A.1.1 Dark matter effective fields

In the following, we explicitly state the Hamiltonians describing the interaction of ALPs and dark photons with a single nuclear spin-1/2 having gyromagnetic ratio γ_N . By analogy to the Zeeman interaction, we then derive the expressions of the corresponding pseudo-magnetic fields acting on a two-spin system such as formic acid. In the following, all expressions are given in SI units.

ALP Wind linear coupling to nuclear spins

For a single nuclear spin, axions and ALPs generate a pseudo-magnetic field through what is known as the “axion wind interaction”, described by the nonrelativistic Hamiltonian [19, 35, 37],

$$\mathcal{H}_{\text{wind}} = \sqrt{\hbar^3 c^3} g_{\text{aNN}} \nabla a(\mathbf{r}, t) \cdot \mathbf{I}_N \quad (\text{A.1})$$

$$= -\hbar g_{\text{aNN}} \mathbf{D}_{\text{wind}} \cdot \mathbf{I}_N , \quad (\text{A.2})$$

where:

$$\mathbf{D}_{\text{wind}}(t) = -\sqrt{\hbar c^3} \nabla a(\mathbf{r}, t) , \quad (\text{A.3})$$

is the ALP effective field, acting on the nuclear spin, and \mathbf{I}_N is the spin operator in units of \hbar .

The ALP field, $a(\mathbf{r}, t)$, can be written as:

$$a(\mathbf{r}, t) = a_0 \cos(\omega_{\text{DM}} t - \mathbf{k} \cdot \mathbf{r} + \phi) , \quad (\text{A.4})$$

where $\omega_{\text{DM}} \approx m_{\text{DM}} c^2 / \hbar$ is the ALP Compton frequency, $\mathbf{k} \approx m_{\text{DM}} \mathbf{v} / \hbar$ is its wave-vector with \mathbf{v} , the average velocity of the particles in the laboratory frame, m_{DM} is the rest mass of the ALP particle and ϕ is an unknown phase. As the value of ϕ has no incidence on the measurement we set its value to zero for the rest of this discussion. Differentiating Eq. (A.4) yields:

$$\nabla a(\mathbf{r}, t) = \frac{m_{\text{DM}} a_0}{\hbar} \sin(\omega_{\text{DM}} t - \mathbf{k} \cdot \mathbf{r}) \mathbf{v} , \quad (\text{A.5})$$

Recalling that a_0 is related to the dark-matter density through

$$\rho_{\text{DM}} = \frac{1}{2} \frac{c^2}{\hbar^2} m_{\text{DM}}^2 a_0^2 , \quad (\text{A.6})$$

yields:

$$\mathbf{D}_{\text{wind}}(t) = -\sqrt{2\hbar c \rho_{\text{DM}}} \times \sin(\omega_{\text{DM}} t - \mathbf{k} \cdot \mathbf{r}) \mathbf{v} .$$

We now assume that the ALP field is acting on the nucleons of the coupled ^{13}C and ^1H nuclear spins while a weak DC magnetic field is applied to the

sample. The axion wind Hamiltonian now becomes

$$\mathcal{H}_{\text{wind}} = -\hbar \mathbf{D}_{\text{wind}} \cdot \left(g_{\text{app}} \mathbf{I} - \frac{1}{3} g_{\text{ann}} \mathbf{S} \right), \quad (\text{A.7})$$

where, g_{app} is the ALP-proton coupling strength, g_{ann} is the ALP-neutron coupling strength. To first order, the $m_F = 0$ states are unaffected, and the $m_F = \pm 1$ states are shifted by

$$\Delta E(m_F = \pm 1)(t) = \pm \frac{\hbar}{2} \frac{2}{3} g_{\text{aNN}} D_z(t), \quad (\text{A.8})$$

where we have assumed that $g_{\text{app}} = g_{\text{ann}} = g_{\text{aNN}}$, and $D_z(t)$ is the projection of $\mathbf{D}_{\text{wind}}(t)$ along the axis of the applied magnetic field. Thus the perturbation induced by $D_z(t)$ takes a similar form to that of a magnetic field $B_z(t)$ (see Eq. (3.7)):

$$\begin{aligned} B_z(t) &= \frac{2}{3} \frac{g_{\text{aNN}}}{\gamma_{\text{H}} + \gamma_{\text{C}}} D_z(t) \\ &= -\frac{2}{3} \frac{g_{\text{aNN}}}{\gamma_{\text{H}} + \gamma_{\text{C}}} \sqrt{2\hbar c \rho_{\text{DM}}} \sin(\omega_{\text{DM}} t - \mathbf{k} \cdot \mathbf{r}) \mathbf{v} \cdot \hat{\mathbf{e}}_z. \end{aligned} \quad (\text{A.9})$$

ALP Wind quadratic coupling to nuclear spins

It is theoretically possible that interaction of nuclear spins with ∇a can be suppressed [78, 83], in which case the dominant axion wind interaction, referred to as the quadratic wind coupling, is related to ∇a^2 :

$$\mathcal{H}_{\text{quad}} = \hbar^2 c^2 (g_{\text{quad}})^2 \nabla a^2(\mathbf{r}, t) \cdot \mathbf{I}_{\text{N}} \quad (\text{A.10})$$

$$= -\hbar (g_{\text{quad}})^2 \mathbf{D}_{\text{quad}} \cdot \mathbf{I}_{\text{N}}, \quad (\text{A.11})$$

where g_{quad} , having dimensions of inverse energy, parameterizes the ALP quadratic coupling strength to nuclear spins and where we have introduced the effective field due to the quadratic axion coupling:

$$\mathbf{D}_{\text{quad}} = -\hbar c^2 \nabla a^2(\mathbf{r}, t) . \quad (\text{A.12})$$

Assuming the form of $a(r, t)$ given by Eq. (A.4):

$$\nabla a^2(\mathbf{r}, t) = a_0^2 \mathbf{k} \sin(2\omega_{\text{DM}}t - 2\mathbf{k} \cdot \mathbf{r}) , \quad (\text{A.13})$$

we obtain for the effective field:

$$\mathbf{D}_{\text{quad}}(t) = -2\hbar c^2 \frac{\rho_{\text{DM}}}{\omega_{\text{DM}}} \sin(2\omega_{\text{DM}}t - 2\mathbf{k} \cdot \mathbf{r}) \mathbf{v} . \quad (\text{A.14})$$

As for the linear coupling, the effect of \mathbf{D}_{quad} acting on a single spin is equivalent to that of a magnetic field, whose amplitude in this case is given by:

$$\begin{aligned} B_{\text{quad},z}(t) &= \frac{g_{\text{quad}}^2}{\gamma_{\text{N}}} D_{\text{quad},z}(t) \\ &= -2\hbar c^2 \frac{g_{\text{quad}}^2}{\gamma_{\text{N}}} \frac{\rho_{\text{DM}}}{\omega_{\text{DM}}} \times \sin(2\omega_{\text{DM}}t - 2\mathbf{k} \cdot \mathbf{r}) \mathbf{v} \cdot \hat{\mathbf{e}}_z . \end{aligned} \quad (\text{A.15})$$

Specifically for formic acid, assuming equal coupling of the axion field to the proton and the neutron, we have

$$B_{\text{quad},z}(t) = -\frac{4}{3} \hbar c^2 \frac{g_{\text{quad}}^2}{\gamma_{\text{H}} + \gamma_{\text{C}}} \frac{\rho_{\text{DM}}}{\omega_{\text{DM}}} \times \sin(2\omega_{\text{DM}}t - 2\mathbf{k} \cdot \mathbf{r}) \mathbf{v} \cdot \hat{\mathbf{e}}_z . \quad (\text{A.16})$$

We note that for the quadratic coupling case, the modulation frequency is twice the Compton frequency of the ALPs. Therefore, an experiment sensitive to a certain range of axion masses for the linear coupling is, at the same time, probing quadratic couplings of axions of half the mass.

Dark-photon couplings to nuclear spins

There are two possible interactions of dark photons with nuclear spins that can be detected with CASPER-ZULF: the coupling of the dark electric field to the dark electric dipole moment (dEDM) and the coupling of the dark magnetic field to the dark magnetic dipole moment (dMDM).

Assuming the dark photons make up the totality of the dark matter energy density, ρ_{DM} , the energy stored in the dark electric field, $\mathbf{E}_{\text{dEDM}}(t)$, is set equal to ρ_{DM} , in a form analogous to that of the usual electromagnetism:

$$\rho_{\text{DM}} = \frac{\epsilon_0}{2} |\mathbf{E}_{\text{dEDM}}|^2, \quad (\text{A.17})$$

where ϵ_0 is permittivity of free space. The dark magnetic field is given by:

$$|\mathbf{B}_{\text{dMDM}}| \approx \frac{v}{c} |\mathbf{E}_{\text{dEDM}}|. \quad (\text{A.18})$$

The Hamiltonians describing the dark EDM and dark MDM are respectively given by:

$$\mathcal{H}_{\text{dEDM}} = g_{\text{dEDM}} \sqrt{\frac{\hbar^3 c^3}{\pi}} \mathbf{E}_{\text{dEDM}}(t) \cdot \mathbf{I}_{\text{N}}, \quad (\text{A.19})$$

and:

$$\mathcal{H}_{\text{dMDM}} = g_{\text{dMDM}} \sqrt{\frac{\hbar^3 c^3}{\pi}} \mathbf{B}_{\text{dMDM}}(t) \cdot \mathbf{I}_N, \quad (\text{A.20})$$

where g_{dEDM} and g_{dMDM} (of dimension of inverse energy) parametrize the strength of the dark photon couplings to the EDM and MDM, respectively.

Similarly to axions and ALPs, the above Hamiltonians are expressed by analogy to the Zeeman interaction by introducing the following magnetic fields:

$$B_{\text{dEDM}}(t) = \frac{g_{\text{dEDM}}}{\gamma_N} \sqrt{2\hbar c^3 \rho_{\text{DM}}} \cos(\omega_{\text{DM}} t) \hat{\mathbf{e}} \cdot \hat{\mathbf{e}}_z, \quad (\text{A.21})$$

and

$$B_{\text{dMDM}}(t) = \frac{g_{\text{dMDM}} v}{\gamma_N c} \sqrt{2\hbar c^3 \rho_{\text{DM}}} \cos(\omega_{\text{DM}} t) \hat{\mathbf{e}} \cdot \hat{\mathbf{e}}_z, \quad (\text{A.22})$$

where $\hat{\mathbf{e}}$ is the bosonic field's polarization vector. Assuming equal coupling to the proton and neutron, yields for a two-spin (^1H - ^{13}C) system:

$$B_{\text{dEDM}}(t) = \frac{2}{3} \frac{g_{\text{dEDM}}}{\gamma_C + \gamma_H} \sqrt{2\hbar c^3 \rho_{\text{DM}}} \cos(\omega_{\text{DM}} t) \hat{\mathbf{e}} \cdot \hat{\mathbf{e}}_z, \quad (\text{A.23})$$

and

$$B_{\text{dMDM}}(t) = \frac{2}{3} \frac{g_{\text{dMDM}} v}{\gamma_C + \gamma_H c} \sqrt{2\hbar c^3 \rho_{\text{DM}}} \cos(\omega_{\text{DM}} t) \hat{\mathbf{e}} \cdot \hat{\mathbf{e}}_z, \quad (\text{A.24})$$

Note that for ordinary photons, the coupling to the EDM is strongly suppressed since the electromagnetic interaction respects the CP symmetry. However, the dark photon need not respect the CP symmetry, in which case there would be no associated suppression [86].

A.1.2 Signal processing: post-processing phase-cycling

The averaging method is illustrated in Fig. 3.2 in the publication main text. Subsequent transient acquisitions are separated by a time interval τ . The n^{th} transient acquisition yields a transient spectrum denoted FFT_n . The operation is repeated N times, yielding a set of N transient spectra: $\{\text{FFT}_n \text{ where } n = 1, 2 \dots N\}$. Once every FFT_n has been acquired, a phase shift, ϕ , incremented by the current acquisition number n , is applied to each FFT_n :

$$\text{FFT}_n^\phi := \text{FFT}_n \times \exp(-in\phi) . \quad (\text{A.25})$$

The N transient phase-shifted FFT_n^ϕ are then averaged:

$$\langle \text{FFT}^\phi \rangle := \frac{1}{N} \times \sum_{n=1}^N \text{FFT}_n^\phi . \quad (\text{A.26})$$

If the phase shift matches the accumulated phase of the dark-matter field between each transient acquisition,

$$\phi := \int_0^\tau \omega_{\text{DM}} dt \quad (\text{modulo } 2\pi), \quad (\text{A.27})$$

where ω_{DM} is the oscillation frequency of the dark-matter field, the sidebands will be averaged coherently. Generally, the ‘‘carrier’’ J -coupling peaks will be averaged to zero, with the exception of cases where $\omega_{\text{DM}} = 2\pi/\tau$. Because the sidebands are coherently averaged, their SNR scales as $\text{SNR}(T_{\text{tot}}) \propto T_{\text{tot}}^{1/2}$ while the J -coupling peaks are averaged away along with the uncorrelated

noise.

The data shown in Fig. 3.2 of the main text were acquired in an experiment during which the AC-field frequency was set to 0.73 Hz. The time interval between each measurement was set to $\tau = 61$ s (including 30 s of acquisition and 31 s for polarization, shuttling and pulsing, combined). The optimal phase shift is computed using Eq. (A.27), yielding $\phi \approx 2.93$ rad. Note that as the frequency of the bosonic fields is unknown, this optimal phase shift cannot be computed during the actual dark-matter-search run. However, this optimal phase shift must necessarily lie within the $[-\pi, \pi]$ interval. In practice we numerically try 2001 phase increments within $[-\pi, \pi]$: $\{\phi_k = k\pi/1000 \text{ where } k = -1000, -999, \dots, 1000\}$. This method produces a set of 2001 phase-shifted averaged FFTs: $\langle \text{FFT}^{\phi_k} \rangle$.

Figure 3.2 illustrates the result of the phase-shifting processing method for an experiment where the transient acquisition is repeated 850 times. The $\langle \text{FFT}^{\phi_k} \rangle$ are ordered by phase increment ϕ_k , and stacked in a two-dimensional plot. This plot is then examined and searched for sidebands. The phase increment of $\phi = 0$ rad corresponds to averaging without applying any phase shift. In this case, the averaging is optimized for the J -resonances and sidebands cannot be seen (i.e., they are destructively averaged). However, a specific phase increment of $\phi = 2.93$ rad (highlighted in orange), yields maximum sideband amplitude while the J -coupling peaks are averaged out. This phase value (and its negative counterpart) corresponds to the optimal value for which the benchmark-field coherence is restored and yields maximum SNR, matching the computed value using Eq. (A.27).

A.1.3 Sideband amplitude determination

A benchmark sideband measurement was performed using a test field of amplitude $B_0 = 0.24$ nT, oscillating at the frequency $\omega_0/2\pi = 0.73$ Hz. The transient acquisition data, consisting of 30 s time-series sampled at 5 kHz, are Fast-Fourier Transformed. The transient FFTs are multiplied by a calibration function expressing the magnetometer sensitivity within the 0 to 500 Hz bandwidth. The calibration function is obtained by measuring the magnetometer response to an applied frequency-varying magnetic field of fixed amplitude.

Following the calibration, the baseline of the FFTs is fit and subtracted. The transient FFTs are then averaged using the phase-shifting scheme described in the previous section.

The next operation consists in extracting the sideband amplitudes from the averaged spectrum. To this end, six complex Lorentzians are fit to the averaged spectrum (two for the J -peaks, each possessing two sidebands):

$$L(\nu) := \frac{\alpha}{\Gamma - i(\nu - \nu_c)} \exp(i\phi) , \quad (\text{A.28})$$

where α , Γ , ν_c and ϕ are the amplitudes, linewidths, center frequencies and phases of the Lorentzians, respectively, and are real, freely fitted parameters. The benchmark sideband amplitude, centered at ν_c is given by $A_{s_0} = |L(\nu_c)|$.

A.1.4 Calibration: signal scaling versus bosonic-field amplitude and frequency

In order to experimentally verify that the sideband amplitude scales with the modulation index, $A_s(B, \omega) \propto B/\omega$, two calibration experiments were performed.

A first calibration was performed by varying the amplitude of the AC-field with constant frequency. The results of this experiment are shown in Fig. A.1 and show the linear dependence of the sideband amplitude on the amplitude of the AC-field. Each data point corresponds to the measured sideband amplitude (as defined in the previous section) for AC-calibration fields of amplitude varying from 0.095 to 0.31 nT at fixed frequency 0.73 Hz.

The second calibration was performed by varying the frequency of the AC-calibration field from 0.45 to 1.7 Hz with constant amplitude of 0.16 nT. The results of this calibration (Fig. A.1) show the $1/\omega$ dependence of the sideband amplitude with the frequency of the AC-field.

Both calibration experiments were carried out with 2100 seconds total integration time for each data point (corresponding to 70 transient acquisitions of 30 seconds).

A.1.5 Detection threshold determination

The detection threshold, A_{th} , is defined as the required signal amplitude such that we can claim detection with a 90% confidence in the case: $A_s(B, \omega) > A_{\text{th}}$, where $A_s(B, \omega)$ is the measured signal as described in Eq. (3.24). To define the detection threshold, we use the standard p -value test described in Ref. [122]. To this aim we consider a complex spectrum containing noise

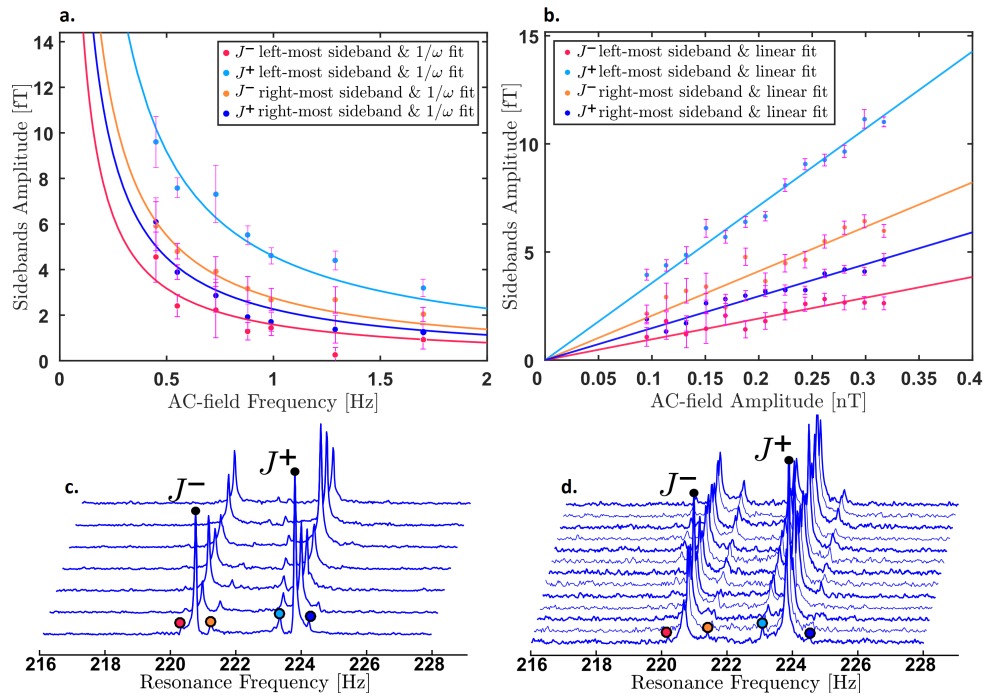


Figure A.1: Calibration data. **(a) & (b)** : Sideband amplitude scaling with external field amplitude and frequency. **(a) & (c)** : Calibration-field set to 0.16 nT and frequency, $\omega/2\pi$, varied from 0.45 to 1.7 Hz. The sideband amplitude follows a $1/\omega$ dependence on the calibration-field frequency ω . **(b) & (d)** : Calibration-field frequency set to 0.73 Hz and amplitude varied from 0.095 to 0.31 nT. The sideband amplitude follows a linear dependence on the calibration-field amplitude. **(a) & (b)** : Each point is the amplitude of the sideband lorentzian fit. Data are averaged over 70 scans of 30 s using the optimal phase increment corresponding to the probed frequency. Error bars show the $1-\sigma$ confidence interval obtained from the fitting procedure.

only:

$$\text{FFT}(\omega) = \chi_1(\omega) + i\chi_2(\omega) , \quad (\text{A.29})$$

where $\chi_1(\omega)$ and $\chi_2(\omega)$ are the real and imaginary part of the FFT, respectively, and $\omega/2\pi$ is the probed frequency. The signal detection threshold can be defined by studying the power spectral density (PSD) of this noisy spectrum:

$$\text{PSD}(\omega) \propto \chi_1(\omega)^2 + \chi_2(\omega)^2 . \quad (\text{A.30})$$

To this aim we use the result that if the Fourier coefficients $\chi_1(\omega)$ and $\chi_2(\omega)$ are normally distributed with zero mean, variance σ^2 and are frequency independent (consistent with white noise), then the PSD cumulative distribution function (CDF) follows an exponential distribution [122]. Hence, at any frequency $\omega/2\pi$, the probability that the measured power P is smaller than z_0 is:

$$\Pr\{P(\omega) < z_0\} = 1 - e^{-z_0/2\sigma^2}, \quad \forall \omega . \quad (\text{A.31})$$

In practice, the value of σ^2 is obtained by fitting the evaluating the PSD CDF of a noise-only spectrum (in the present case via the MATLAB Kaplan-Meier estimator). We then fit the PSD CDF to the exponential distribution given in Eq. (A.31). An illustration of this procedure is shown in Fig. A.2.

We now define a signal power z_0 such that, if at the probed frequency $\omega/2\pi$, the power of a spectrum containing both the signal and noise, $P(\omega)$, is greater than z_0 , a detection can be claimed with a false-alarm probability

p_0 (a false alarm corresponds to the case where the measured power $P(\omega)$ is due to noise fluctuations). This yields:

$$p_0 := \Pr\{P(\omega) > z_0\} \quad (\text{A.32})$$

$$= 1 - \Pr\{P(\omega) < z_0\} \quad (\text{A.33})$$

$$= e^{-z_0/2\sigma^2} . \quad (\text{A.34})$$

Solving the above equation for z_0 yields:

$$z_0 = -2\sigma^2 \ln(p_0) . \quad (\text{A.35})$$

Setting now the false-alarm probability to $p_0 = 0.10$ (corresponding to a confidence level of 90%) we find the signal power threshold in terms of the noise fluctuation:

$$z_0 \approx 4.6\sigma^2 . \quad (\text{A.36})$$

As our analysis is done on the spectrum amplitude, defined as the square-root of the PSD, we define the detection threshold as:

$$A_{\text{th}} = \sqrt{z_0} \quad (\text{A.37})$$

$$\approx 2.14\sigma . \quad (\text{A.38})$$

Any signal with amplitude $A_s(B, \omega) > A_{\text{th}}$ is inspected for detection. The

case where no such signal is observed, yields the upper-bound exclusion line:

$$A_s(B, \omega) = A_{\text{th}} . \quad (\text{A.39})$$

Using Eq. (3.24) this limit can be expressed in term of magnetic field versus frequency:

$$B_{\text{th}}(\omega) = \frac{A_{\text{th}} B_0}{A_{s_0} \omega} . \quad (\text{A.40})$$

Such analysis is done in a frequency-dependent manner by computing the phase increment corresponding to the studied frequency. The averaged spectrum obtained using this phase increment is then studied as explained above to define the signal-detection threshold at ω . The results of this analysis are given in Fig. A.2.c, where Eq. (A.40) is evaluated for every frequency within the bandwidth.

The exclusion plots shown in Figs. 3.3, 3.4 and 3.5 in the publication main text, are obtained by converting ω to bosonic masses in eV and $B_{\text{th}}(\omega)$ to coupling constants in GeV^{-1} using Eqs. (3.19), (3.20), (3.21) and (3.22) in the publication main text.

The exclusion lines show the 2.14σ confidence level of the experiment (90% CL). Such frequency-dependent analysis yields higher detection threshold for frequencies for which the phase increment is close to 0 rad. Indeed, the spectra obtained from the 0 radian phase increment contains maximum amplitude J -couplings, which then raise the detection threshold to a higher value due to spectral leakage. This effect manifests itself in the exclusion lines by sporadic and sharp increases of the detection threshold.

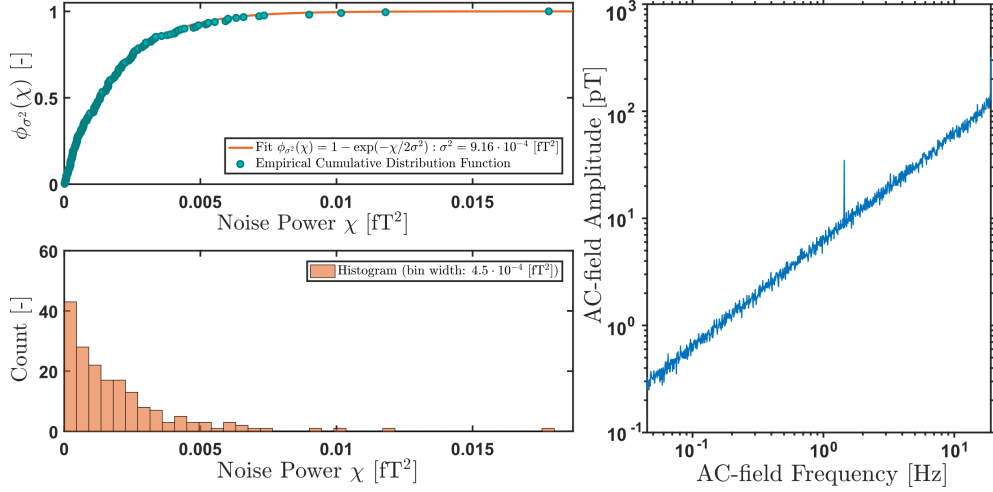


Figure A.2: (a) : Example of empirical white-noise power-spectral-density cumulative distribution function (CDF) and fit to the exponential distribution $\phi_{\sigma^2}\{\chi\} = 1 - e^{-\chi/2\sigma^2}$. The fit yields a white-noise variance $\sigma^2 \approx 9.16 \times 10^{-4}$ [fT²]. (b) : White-noise power histogram. (a), (b) : Non-corrected for misalignment (see A.1.8), noisy-data analysed around 1.5 Hz in a 5 Hz frequency bin using the optimal phase increment corresponding to the 1.5 Hz frequency. (c) : Amplitude of AC magnetic field yielding a signal above the detection threshold for frequencies within the bandwidth 0.045 to 19 Hz. Based on the calibration data (see section A.1.4) and power threshold evaluation. Total integration time of 25500 s (corresponding to 850 transient acquisitions each 30 seconds long, with a duty cycle of 50% imposed by the polarization step between each acquisition).

We note that in the literature, similar analyses are done by applying a statistical penalty for searching multiple frequencies within a bandwidth of interest to account for the look-elsewhere effect (corresponding to raising Eq. (A.31) to the power of the number of frequencies probed). Our analysis does not require such a penalty. Indeed, the characteristic signature of the bosonic fields, a set of four sidebands around the J -coupling frequencies, allows one to differentiate false alarms from true detection events.

A.1.6 Coherent averaging: signal scaling with integration time

Figure A.3 shows the sidebands amplitude A_s , detection threshold A_{th} , and signal-to-noise ratio, $SNR := A_s/A_{th}$, scaling with the total integration time using the phase-shifting procedure described in A.1.2.

During this experiment, the AC-field is maintained at a frequency $\omega_{AC}/2\pi = 0.73$ Hz and amplitude $B_{AC} = 0.24$ nT. The experiment is repeated with increased number of transient acquisitions each of length 30 s. All values are extracted using the averaged spectrum obtained from the optimal phase increment corresponding to ω_{AC} .

Figure A.3 shows that the sideband amplitude remains constant while the signal detection-thresholds decreases as $A_{th}(T_{tot}) \propto T_{tot}^{-1/2}$ as expected while performing a coherent averaging. The SNR scales accordingly as $SNR \propto T_{tot}^{1/2}$.

A.1.7 Bandwidth: accessible bosonic mass range

The bandwidth of the experiment is limited from below by the linewidth of the J -coupling peaks. For bosonic fields with frequencies lower than 45 mHz (corresponding to bosons with mass $\lesssim 1.86 \times 10^{-16}$ eV), the corresponding sidebands are located inside the J -resonance and cannot be resolved. As a result, we limit the lower end of the search bandwidth to 45 mHz, below which calibration could not be achieved (see Fig. A.4). In principle, the bandwidth extends to frequencies up to ~ 500 Hz, corresponding to the maximum detectable frequency of the magnetometer (see Materials and Methods Sec. 3.2). However, the data are acquired during ~ 14 hours (corresponding




Figure A.3: Sidebands amplitude A_s , detection threshold A_{th} at 90% confidence level (90% CL), and signal-to-noise ratio A_s/A_{th} (SNR) versus integration time T_{tot} , using the phase shifting processing method. A_s remains constant, A_{th} scales as $T_{tot}^{-1/2}$ during the averaging, consistent with coherent averaging. Thus the SNR scales as $T_{tot}^{1/2}$. During this experiment, the AC-field is maintained at a frequency $\omega_{AC}/2\pi = 0.73$ Hz and amplitude $B_{AC} = 0.24$ nT. Noisy data analysed around 0.73 Hz in a 5 Hz frequency bin. Data is averaged over 850 scans of 30 seconds using the optimal phase increment corresponding to the 0.73 Hz frequency. Error bars show the 1- σ confidence interval obtained from the fitting procedure.




Figure A.4: Sideband amplitude scaling with calibration-field frequency. Calibration-field set to 0.8 nT and frequency, $\omega/2\pi$, varied from 0.045 to 0.088 Hz. The sideband amplitude follows a $1/\omega$ dependence with the calibration-field frequency $\omega/2\pi$. Each point is the amplitude of the sideband lorentzian fit. Data are averaged over 70 scans of 30 s using the optimal phase increment corresponding to the probed frequency. Error bars show the $1\text{-}\sigma$ confidence interval obtained from the fitting procedure.

to 850 transient acquisitions each 30 seconds long, with a duty cycle of 50% imposed by the polarization step between each acquisition). The 850 transient signals are coherently averaged for as long as the bosonic field remains phase-coherent. The coherence time of the bosonic fields is on the order of 10^6 periods of oscillation, which for fields of ~ 19 Hz corresponds to ~ 14 hours. This limits the search bandwidth to fields of frequencies below 19 Hz (corresponding to bosons with mass lower than 7.8×10^{-14} eV).

A.1.8 Dark-matter-field directionality

Daily and annual modulations

We note that the ALP field described in Eq. (A.4), seen in the Earth's reference frame, has annual and daily modulations due to the motion of the Earth in the Solar System. These modulations are expected to **1-** modulate the orientation of the experiment's sensitive axis with respect to the dark-matter wind, thus modulating the overall experimental sensitivity and **2-** change the lineshape of the experimental power spectra from pure Lorentzians to skewed Lorentzians. These effects are addressed and accounted for during the data processing by weighting transient spectra during the averaging sequence using the following method.

For axions and ALPs, the experiment is sensitive to the projection of the bosonic pseudo-magnetic field gradient, $\nabla a(\mathbf{r}, t)$, onto the direction of the leading magnetic field. We denote this axis in the laboratory frame by the unit vector $\hat{\mathbf{e}}_z$. As a result of the yearly revolution of the Earth around the Sun and the daily revolution of the Earth, the alignment of $\hat{\mathbf{e}}_z$ with $\nabla a(\mathbf{r}, t)$ varies over time. This effect can increase or totally suppress the expected signal depending on when transient acquisitions happen and must be included in the analysis to avoid overestimating the experiment's sensitivity.

To this aim, we use the analysis proposed in Ref. [123], originally applied to nuclear recoil in WIMP-nucleus scattering but which directly translates to this sensitive-axis misalignment. Detailed and comprehensive calculations are given in Ref. [123], here we show how these results are applied to our current signal analysis. We recall that the ALP scalar-field gradient can be

written as (ignoring its initial phase):

$$\nabla a(\mathbf{r}, t) = a_0 \cos(\omega_{\text{DM}}t - \mathbf{k} \cdot \mathbf{r} - \pi/2) \cdot \mathbf{k} , \quad (\text{A.41})$$

where $\omega_{\text{DB}} = \sqrt{(m_{\text{DM}}c^2)^2 + (m_{\text{DM}}vc)^2}/\hbar$ is the bosonic particle de Broglie frequency, $\mathbf{k} = m_{\text{DM}}\mathbf{v}/\hbar$ is its wave-vector with \mathbf{v} , the average velocity of the particles in the laboratory frame and m_{DM} is the mass of the particle. When shifting from the galactic-center to the laboratory frame, \mathbf{v} can be written as the following vectorial sum of velocities:

$$\mathbf{v}(t) = \mathbf{v}_{\text{DM}} - \mathbf{v}_{\text{SG}} - \mathbf{v}_{\text{ES}}(t) - \mathbf{v}_{\text{LE}}(t) , \quad (\text{A.42})$$

where \mathbf{v}_{DM} is the velocity of the local bosonic field with respect to the center of the galaxy and is taken to be zero on average due to the isotropic structure of the dark-matter halo, \mathbf{v}_{SG} is the velocity of the Sun with respect to the center of the galaxy, i.e. \mathbf{v}_{SG} is the velocity of the Sun towards the Cygnus constellation and is approximately constant, $\mathbf{v}_{\text{ES}}(t)$ is the velocity of the center of the Earth with respect to the sun and $\mathbf{v}_{\text{LE}}(t)$ is the velocity of the laboratory with respect to center of the Earth. The expected signal amplitude S , is proportional to the projection of the field's gradient onto the

sensitive axis:

$$S(t) \propto \nabla a(\mathbf{r}, t) \cdot \hat{\mathbf{e}}_z \quad (\text{A.43})$$

$$\propto \frac{m_{\text{DM}} a_0}{\hbar} \cos(\omega_{\text{DB}} t - \mathbf{k} \cdot \mathbf{r}) \mathbf{v}(t) \cdot \hat{\mathbf{e}}_z \quad (\text{A.44})$$

$$\propto \frac{m_{\text{DM}} a_0}{\hbar} \cos(\omega_{\text{DB}} t - \mathbf{k} \cdot \mathbf{r}) \dots \quad (\text{A.45})$$

$$\times \left(\mathbf{v}_{\text{SG}} \cdot \hat{\mathbf{e}}_z + \mathbf{v}_{\text{ES}}(t) \cdot \hat{\mathbf{e}}_z + \mathbf{v}_{\text{LE}}(t) \cdot \hat{\mathbf{e}}_z \right).$$

During the calibration and benchmark acquisitions, an oscillating-magnetic field is continuously applied along the $\hat{\mathbf{e}}_z$ direction. After averaging the benchmark sideband amplitude A_{s_0} is obtained from the fitting procedure. To avoid overestimating the expected signal amplitude and account for the previously discussed astrophysical motions, we weigh the transient spectra by the factor $\frac{|\mathbf{v}(t_n) \cdot \hat{\mathbf{e}}_z|}{|v|} \approx \frac{|\mathbf{v}(t_n) \cdot \hat{\mathbf{e}}_z|}{10^{-3}c}$ before including them in the average spectra, where t_n is the starting time of the n^{th} transient acquisition. This replicates the daily and yearly amplitude modulations of the signal induced by the motion of the Earth in the Solar System such as to avoid overestimating the expected signal amplitude from the bosonic fields.

During the search acquisitions, no oscillating-magnetic field is applied and this misalignment effect is naturally included in the data. Nonetheless we weight the transient spectra by the factor $\frac{\mathbf{v}(t_n) \cdot \hat{\mathbf{e}}_z}{|v|}$ before averaging them. Indeed, a transient spectrum acquired when $\mathbf{v}(t_n) \cdot \hat{\mathbf{e}}_z = 0$, cannot possibly include any signal and thus contains only noise. The effect of weighting such a transient spectrum by this factor is to prevent the addition of noise in the averaged spectrum. On the contrary, spectra acquired when $\mathbf{v}(t_n) \cdot \hat{\mathbf{e}}_z$ is maximal possess higher weight to account for the fact that the signal should

be maximal during the acquisition.

For dark photons, the experiment is sensitive to the projection of the field's polarization, denoted by the unit-vector $\hat{\epsilon}$, onto the leading field direction \hat{e}_z . We note that, at any point in space and time, $\hat{\epsilon}$ can take any direction and is uniformly distributed along the unit-3D-sphere. Moreover, $\hat{\epsilon}$ is assumed to be constant during one coherence time, which is always longer than the total measurement time. Therefore the experiment probes a unique value of $\hat{\epsilon}$, modulated by the annual and daily motions of the sensitive axis of detection.

The analysis is performed assigning three equally probable orientations $\hat{\epsilon}_i$; and by time propagating $\hat{\epsilon}_i \cdot \hat{e}_z(t_0)$. This quantity is then used as a weight during the averaging of transient signal similar to the ALP case, thus yielding three corresponding limits.

To obtain limits comparable with other experiments, the three $\hat{\epsilon}_i$, are given in the non-rotating Celestial Frame. The axis of this frame are practically fixed on any relevant time scale, since it takes the Solar System ~ 250 million years to complete one orbit about the galactic centre, so different experiments can be compared easily in this case (even if they are performed many years apart in time).

The three assigned directions are the usual orthonormal axis of the non-rotating Celestial Frame. The first assigned orientation, $\hat{\epsilon}_1$, lies in the Celestial Equator Plane and points towards the Sun at the Vernal Equinox. The second orientation, $\hat{\epsilon}_2$, points towards the North Celestial Pole (i.e. along the axis of rotation of the Earth). In this case the sensitivity of the experiment is null, as \hat{e}_z is pointing East in the laboratory frame and

remains perpendicular to $\hat{\epsilon}_2$ at any point in time. The third orientation is the following vector product $\hat{\epsilon}_3 = \hat{\epsilon}_1 \times \hat{\epsilon}_2$, and lies in the Celestial Equator Plane. A complete mathematical description and schematic of this coordinate system can be found in Ref. [123].

Search data time-stamp

The search data are composed of 850 transient scans, each 30 s long, spaced by a polarization time of 31 s. The first transient scan was acquired on the 28th of July 2018 at 10:00:06 (CET). Data acquired in Mainz, Germany.

Expected sideband lineshape

The oscillating bosonic fields described in the Earth frame also exhibit annual and daily modulations due to the presence of the modulated wave-vector \mathbf{k} , and de Broglie frequency ω_{DB} in the arguments of the cosine function. The effect of such modulations is to modify the lineshape of the power spectra from the expected Lorentzian signal. The first modulation due to the motion of the Earth with respect to the galactic center enters in the de Broglie frequency $\omega_{\text{DB}} = \sqrt{(m_{\text{DM}}c^2)^2 + (m_{\text{DM}}vc)^2}/\hbar$, where \mathbf{v} is modulated as in Eq. (A.42). The amplitude of the velocities in Eq. (A.42) are $|\mathbf{v}_{\text{DM}}| \approx 0$, $|\mathbf{v}_{\text{SG}}| \approx 220 \text{ km}\cdot\text{s}^{-1} \approx 10^{-3}c$, $|\mathbf{v}_{\text{ES}}(t)| \approx 30 \text{ km/s}$ and $|\mathbf{v}_{\text{LE}}(t)| \approx 0.5 \text{ km/s}$ [123]. Those components are therefore negligible compared to the rest mass-energy of the particle, so the usual approximation $\omega_{\text{DB}} \approx m_{\text{DM}}c^2/\hbar$ is used. The second modulation appears in the spatial argument of the cosine function, $\mathbf{k} \cdot \mathbf{r}$. This component is dominated by the velocity of the Sun towards the Cygnus constellation, $|\mathbf{v}_{\text{SG}}|$, which is

taken as constant both in amplitude and direction. Therefore, we neglect other components of the wave-vector. These two approximations yield a non-modulated cosine-form gradient, thus yielding a pure Lorentzian power-spectral line.

A.2 Chapter 4

A.2.1 DNN architecture implementation

The pseudo-code architectures of the *Encoder*, *Regressor*, *Decoder*, and unified DNN, as implemented in Python (Keras - Tensorflow), are given below in addition to the custom weighted-loss function.

```

bottleneck_dim = 64 # Output dimension of the Encoder
latent_dim      = 7  # Output dimension of the Regressor (# latent parameters)
signal_length   = 512

## ENCODER subDNN
input_shape = (signal_length,1)
i = Input(shape=input_shape)
x = Conv1D(64, kernel_size=64, activation='relu', padding='same')(i)
x = MaxPooling1D(4, padding='same')(x)
x = Conv1D(64, kernel_size=32, activation='relu', padding='same')(x)
x = MaxPooling1D(4, padding='same')(x)
x = Flatten()(x)
x = Dense(128, activation='relu', kernel_initializer='he_uniform')(x)
bottleneck = Dense(64, activation='relu', kernel_initializer='he_uniform')(x)

encoder = Model(i, bottleneck)

## REGRESSOR subDNN
input_shape = (bottleneck_dim,1)
r_i = Input(shape=input_shape)
x = Reshape((bottleneck_dim,1))(r_i)
x = Conv1D(64, kernel_size=64, activation='relu', padding='same')(x)
x = MaxPooling1D(4, padding='same')(x)

```

```
x = Conv1D(64, kernel_size=32, activation='relu', padding='same')(x)
x = MaxPooling1D(4, padding='same')(x)
x = Conv1D(64, kernel_size=32, activation='relu', padding='same')(x)
x = Flatten()(x)
x = Dense(256, activation='relu', kernel_initializer='he_uniform')(x)
x = Dense(128, activation='relu', kernel_initializer='he_uniform')(x)
x = Dense(64, activation='relu', kernel_initializer='he_uniform')(x)
latent = Dense(latent_dim)(x)

regressor = Model(r_i, latent)

## DECODER subDNN
input_shape = (bottleneck_dim+latent_dim,1)
d_i = Input(shape=input_shape)
x = Dense(128, activation='relu', kernel_initializer='he_uniform')(d_i)
x = Reshape((128,1))(x)
x = Conv1D(64, kernel_size=32, activation='relu', padding='same')(x)
x = MaxPooling1D(2, padding='same')(x)
x = UpSampling1D(4)(x)
x = Conv1D(64, kernel_size=32, activation='relu', padding='same')(x)
x = MaxPooling1D(2, padding='same')(x)
x = UpSampling1D(4)(x)
x = Conv1D(64, kernel_size=32, activation='relu', padding='same')(x)
x = MaxPooling1D(2, padding='same')(x)
x = UpSampling1D(2)(x)
decoded = Conv1D(1, kernel_size=32, activation='sigmoid', padding='same')(x)

decoder = Model(d_i, decoded)

## Unified DNN model
concat = Concatenate()( [encoder(i), regressor(encoder(i))] )
```

```
ae_outputs = decoder(concat)
flatten_ae_outputs = Reshape((signal_length,))(ae_outputs)
concat2 = Concatenate()([flatten_ae_outputs, regressor(encoder(i)) ])
DNN_outputs = concat2
DNN = Model(i, DNN_outputs)

## CustomLoss function
def customLoss(yTrue, yPred):
    latentSize = 7
    SignalSize = 512
    beta = 0.001
    mseSignal = K.square(yTrue[:, 0:SignalSize] - yPred[:, 0:SignalSize])
    mseSignal = K.abs(mseSignal)
    mseSignal = K.sum(mseSignal, axis=-1)
    mseSignal = mseSignal/SignalSize

    mseLatent = K.square(yTrue[:, SignalSize:] - yPred[:, SignalSize:])
    mseLatent = K.abs(mseLatent)
    mseLatent = K.sum(mseLatent, axis=-1)
    mseLatent = mseLatent/latentSize

    weighted_mse = (1-beta)*mseSignal + beta*mseLatent

    return weighted_mse

DNN.compile(optimizer = 'ADAM', loss = customLoss)
DNN.summary()
```

A.3 Supplementary figures

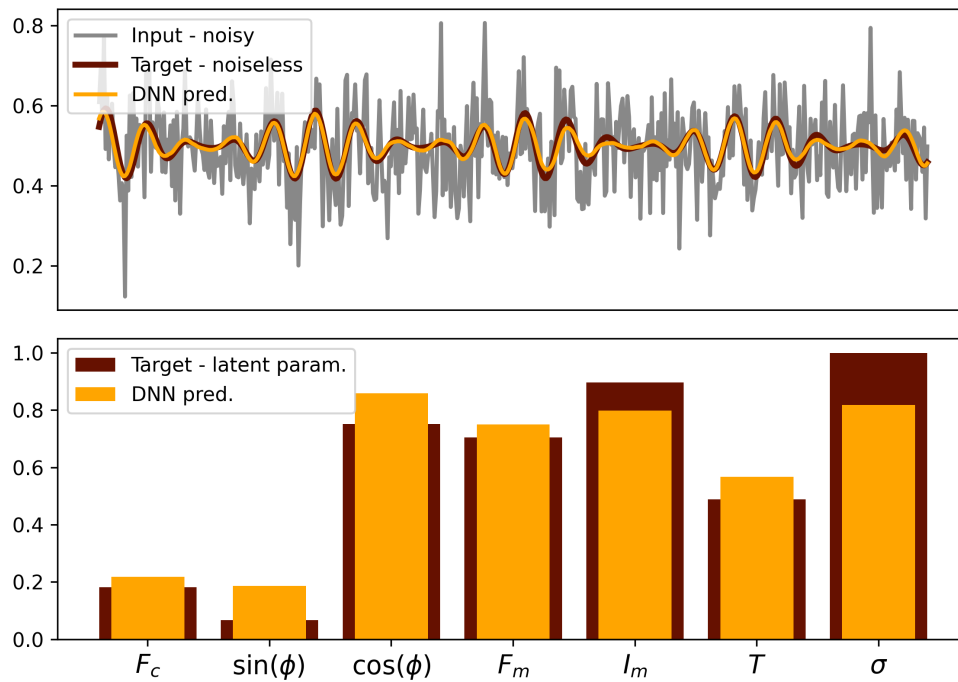


Figure A.5: Example of DNN prediction using a noisy AM-sine wave input. The DNN was trained only on AM-sine waves samples. **Top:** Noisy input and *Decoder* denoised prediction. **Bottom:** True latent parameter targets and *Regressor* prediction. The phase ϕ is mapped to two separate latent variables to accommodate for phase periodicity during loss computation.

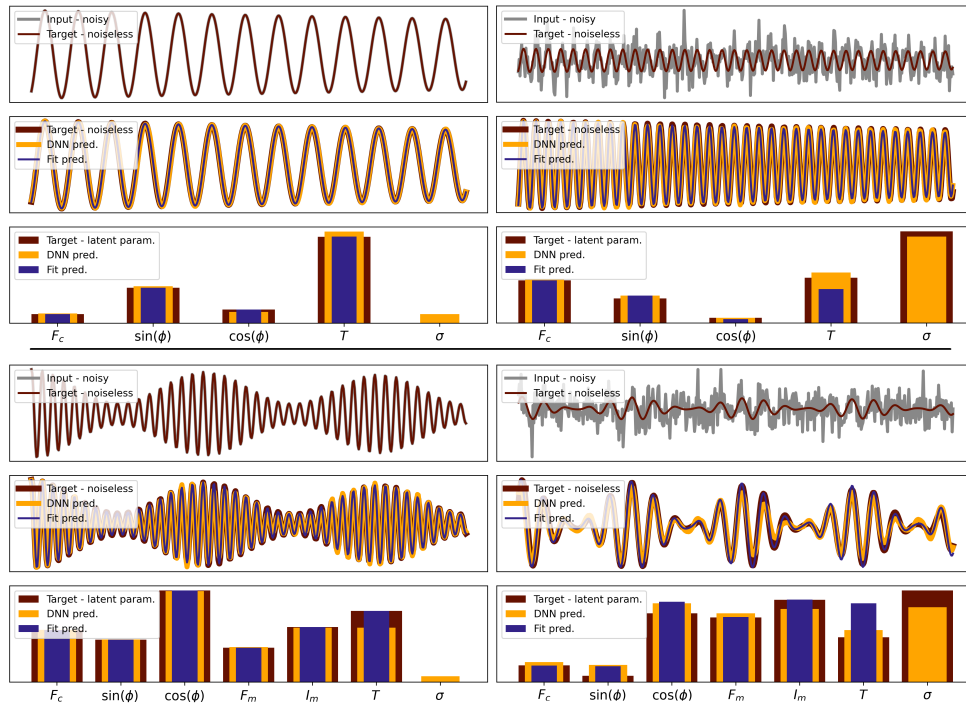


Figure A.6: Example of minimum (left) and maximum noise (right) sine wave (top) and AM-sine wave (bottom) samples. DNN and LS-fits denoised and latent parameters predictions are shown below the respective inputs.




Figure A.7: Comparison of DNN post-training performance to LS-fits with true latent-parameters initial guesses for 1000 random AM-sine wave from the test set (unseen during training). The denoising (MSE_{dec} , top) and latent-parameters relative regression losses (MSE_{reg} , bottom) are sorted by increasing noise levels. The denoising (MSE_{dec} , top) and latent-parameters regression losses (MSE_{reg} , bottom) are sorted by increasing noise levels. The DNN was trained only on AM-sine waves samples. See Fig. 4.4 for MSE_{reg} and MSE_{dec} computation methods. The LS-fit with true initial guesses vastly outperforms the DNN for low-noise signals but both systems reach similar performance for high-noise.

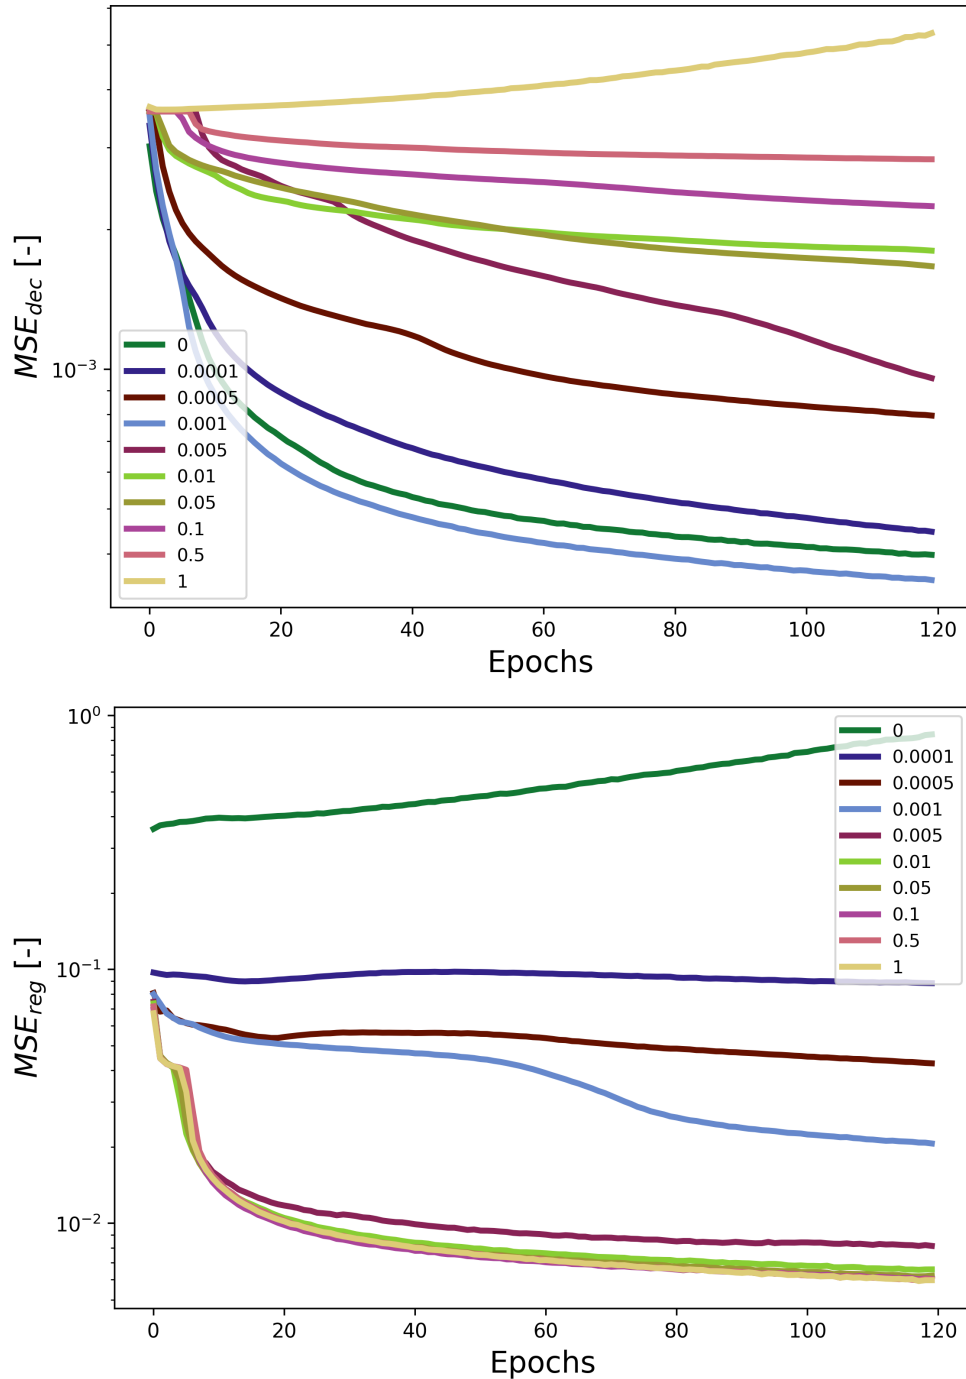


Figure A.8: Validation loss during training for various values of β . **Top:** signal denoising losses. **Bottom:** latent-parameters regression losses. Training is performed on 12 training sets of 100'000 randomly generated FM-sine waves for 10 epochs. For clarity, validation-loss is displayed after each individual training set backpropagation.

Bibliography

- [1] E. Purcell, H. Torrey, and R. Pound. Resonance Absorption by Nuclear Magnetic Moments in a Solid. *Physical Review Letters*, 69(1-2):37–38, 1946.
- [2] F. Bloch, W.W. Hansen, and M. Packard. The nuclear induction experiment. *Physical Review Letters*, 70(7-8):474–485, 1946.
- [3] J. T. Arnold, S. S. Dharmatti, and M. E. Packard. Chemical effects on nuclear induction signals from organic compounds. *The Journal of Chemical Physics*, 19(4):507–507, 1951.
- [4] Min Jiang, Teng Wu, John W. Blanchard, Guanru Feng, Xinhua Peng, and Dmitry Budker. Experimental benchmarking of quantum control in zero-field nuclear magnetic resonance. *Sci. Adv.*, 4, 2018.
- [5] J. Bian, M. Jiang, J. Cui, X. Liu, B. Chen, Y. Ji, B. Zhang, J.W. Blanchard, X. Peng, and J. Du. Universal quantum control in zero-field nuclear magnetic resonance. *Physical Review A*, 95, 2017.
- [6] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta. 6: 110–127*, 1933.
- [7] L. Ackerman, M.R. Buckley, S.M. Carroll, and M. Kamionkowski. Dark matter and dark radiation. *Phys. Rev. D*, 79:023519, 2009.
- [8] G. Bertone and T.M.P. Tait. A new era in the search for dark matter. *Nature*, 562, 2018.
- [9] G. Bertone, D. Hooper, and J. Silk. Particle dark matter: Evidence, candidates and constraints. *Phys. Rep.*, 405:279, 2005.
- [10] J. L. Feng. Dark matter candidates from particle physics and methods of detection. *Annu. Rev. Astron. Astr.*, 48:495, 2010.

-
- [11] C. Amsler et al. Review of Particle Physics. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics*, 667(1-5): 1–6, 2008.
- [12] R. D. Peccei and H.R. Quinn. CP conservation in the presence of pseudoparticles. *Physical Review Letters*, 38(25):1440–1443, 1977.
- [13] S. Weinberg. A new light boson? *Physical Review Letters*, 40(4): 223–226, 1978.
- [14] F. Wilczek. Problem of Strong P and T Invariance in the Presence of Instantons. *Physical Review Letters*, 40(5), 1978.
- [15] R. Peccei and H. Quinn. CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38:1440, 1977.
- [16] P.W. Graham, D.E. Kaplan, and S. Rajendran. Cosmological relaxation of the electroweak scale. *Phys. Rev. Lett.*, 115:221801, 2015.
- [17] P. Svrcek and E. Witten. Axions in string theory. *J. High Energy Phys.*, 06:051, 2006.
- [18] M. S. Safronova, D. Budker, D. DeMille, D. F. Jackson Kimball, A. Derevianko, and C. W. Clark. Search for new physics with atoms and molecules. *Rev. Mod. Phys.*, 90:025008, 2018.
- [19] P. W. Graham, I. G. Irastorza, S. K. Lamoreaux, A. Lindner, and K. A. van Bibber. Experimental searches for the axion and axion-like particles. *Annu. Rev. Nucl. Part. S.*, 65:485, 2015.
- [20] J. Ipser and P. Sikivie. Can Galactic Halos Be Made of Axions? *Physical Review Letters*, 50(12), 1983.
- [21] L.F. Abbott and P. Sikivie. A cosmological bound on the invisible axion. *Physical Review Letters B*, 120, 1983.
- [22] M. Dine and W. Fischler. The not so harmless axion. *Physical Review Letters B*, 120, 1983.
- [23] J. Preskill, M.B. Wise, and F. Wilczek. Cosmology of the invisible axion. *Physics Review Letters B*, 120(1-3):127–132, 1983.
- [24] D.J.E. Marsh. Axion cosmology. *Physics Reports*, 643(2):1–79, 2016.
- [25] M. Srednicki. Axion couplings to matter. (I). CP-conserving parts. *Nuclear Physics, Section B*, 260(3-4):689–700, 1985.

-
- [26] H. Primakoff. Photo-production of neutral mesons in nuclear electric fields and the mean life of the neutral meson. *Phys. Rev.*, 81:899–899, 1951.
- [27] K. Zioutas et al. First results from the cern axion solar telescope. *Physical Review Letters*, 94:121301, 2005.
- [28] G. Raffelt and L. Stodolsky. Mixing of the photon with low-mass particles. *Physical Review Letters D*, 37(5), 1988.
- [29] K. Ehret et al. Resonant laser power build-up in ALPS-A ”light shining through a wall” experiment. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 612(1):83–96, 2009.
- [30] S.J. Asztalos et al. SQUID-based microwave cavity search for dark-matter axions. *Physical Review Letters*, 104(4):1–4, 2010.
- [31] P. Sikivie. Experimental Tests of the Invisible Axion. *Physical Review Letters*, 51(16), 1983.
- [32] P.W. Graham, I.G. Irastorza, S.K. Lamoreaux, A. Lindner, and K.A. Van Bibber. Experimental Searches for the Axion and Axion-Like Particles. *Annual Review of Nuclear and Particle Science*, 65(1):485–514, 2015.
- [33] D. Budker, P.W. Graham, M. Ledbetter, S.Rajendran, and A.O. Sushkov. Proposal for a cosmic axion spin precession experiment (CASPER). *Physical Review Letters X*, 4(2):1–10, 2014.
- [34] Peter W. Graham and Surjeet Rajendran. New observables for direct detection of axion dark matter. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 88:1–16, 2013.
- [35] Y.V. Stadnik and V.V. Flambaum. Axion-induced effects in atoms, molecules, and nuclei: Parity nonconservation, anapole moments, electric dipole moments, and spin-gravity and spin-axion momentum couplings. *Phys. Rev. D*, 89:043522, 2014.
- [36] D. Budker, P. W. Graham, M. Ledbetter, S. Rajendran, and A. O. Sushkov. Proposal for a cosmic axion spin precession experiment (CASPER). *Phys. Rev. X*, 4:021030, 2014.
- [37] P. W. Graham and S. Rajendran. New observables for direct detection of axion dark matter. *Phys. Rev. D*, 88:035023, 2013.

- [38] D. F. Jackson Kimball, S. Afach, D. Aybas, J. W. Blanchard, D. Budker, G. Centers, M. Engler, N. L. Figueroa, A. Garcon, and P. W. Graham. Overview of the cosmic axion spin precession experiment (CASPER). *arXiv:1711.08999*, 2017.
- [39] T. Wang, Derek F. Jackson Kimball, A.O. Sushkov, D. Aybas, J.W. Blanchard, G. Centers, S.R. O' Kelley, A. Wickenbrock, J. Fang, and D. Budker. Application of spin-exchange relaxation-free magnetometry to the cosmic axion spin precession experiment. *Physics of the Dark Universe*, 19:27–35, 2018.
- [40] C. Abel, N. J. Ayres, G. Ban, G. Bison, K. Bodek, V. Bondar, M. Daum, M. Fairbairn, V. V. Flambaum, P. Geltenbort, K. Green, W. C. Griffith, M. van der Grinten, Z. D. Gruji, P. G. Harris, N. Hild, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic, K. Kirch, H.-C. Koch, S. Komposch, P. A. Koss, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemièrè, D. J. E. Marsh, P. Mohanmurthy, A. Mtchedlishvili, M. Musgrave, F. M. Piegsa, G. Pignol, M. Rawlik, D. Rebreyend, D. Ries, S. Roccia, D. Rozpedzik, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, Y. V. Stadnik, A. Weis, E. Wursten, J. Zejma, and G. Zsigmond. Search for axionlike dark matter through nuclear spin precession in electric and magnetic fields. *Phys. Rev. X*, 7:041034, 2017.
- [41] J.W. Blanchard and D. Budker. Zero- to Ultralow-Field NMR. *eMagRes*, pages 1395–1410, 2016.
- [42] R. Catena and P. Ullio. A novel determination of the local dark matter density. *arXiv:0907.0018v2*, 2009.
- [43] G.G. Raffelt. *Astrophysical Axion Bounds*, pages 51–71. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [44] M.I. Vysotsky, Y.B. Zeldovich, M.Y. Khlopov, and V.M. Chechetkin. Some astrophysical limitations on the axion mass. *Journal of Experimental and Theoretical Physics Letters*, 27, 1978.
- [45] P.W. Graham, I.G. Irastorza, S.K. Lamoreaux, A. Lindner, and K.A. Van Bibber. Experimental Searches for the Axion and Axion-Like Particles. *Annual Review of Nuclear and Particle Science*, 2015.
- [46] M.H. Levitt. *Spin Dynamics: Basics of Nuclear Magnetic Resonance - 2nd Ed, ch.2–3, p.34–39*,. Wiley, 2008.

-
- [47] A. Derevianko. Detecting dark matter waves with precision measurement tools. *arXiv:1605.09717v2*, 2016.
- [48] P.W. Graham and S. Rajendran. Axion dark matter detection with cold molecules. *Physical Review Letters D*, 84, 2011.
- [49] J. Jaeckel. A force beyond the Standard Model – Status of the quest for hidden photons. *Frascati Physics Series*, 56, 2012.
- [50] S. Chaudhuri, P.W. Graham, K. Irwin, J. Mardon, S. Rajendran, and Y. Zhao. A radio for hidden-photon dark matter detection. *Physical Review Letters D*, 92, 2015.
- [51] John H. Karl. *An Introduction to Digital Signal Processing*. Academic Press, 1991.
- [52] M.P. Ledbetter and M.V. Romalis. Nonlinear Effects from Dipolar Interactions in Hyperpolarized Liquid ^{129}Xe . *Physical Review Letters*, 89, 2002.
- [53] M.V. Romalis and M.P. Ledbetter. Transverse Spin Relaxation in Liquid ^{129}Xe in the Presence of Large Dipolar Fields. *Physical Review Letters*, 8725(33), 2001.
- [54] B. Driehuys, G.D. Cates, E. Miron, K. Sauer, D.K. Walter, and W. Happer. High-volume production of laser-polarized ^{129}Xe . *Applied Physics Letters*, 69, 1996.
- [55] T. Wang, D.F.J. Kimball, A.O. Sushkov, D. Aybas, J.W. Blanchard, G. Centers, S.R.O. Kelley, J. Fang, and D. Budker. Application of Spin-Exchange Relaxation-Free Magnetometry to the Cosmic Axion Spin Precession Experiment. *arXiv:1701.08082*, 2017.
- [56] D. Robbes. Highly sensitive magnetometers-a review. *Sensors and Actuators, A: Physical*, 129(1-2 SPEC. ISS.):86–93, 2006.
- [57] D. Budker, D. Kimball, and D. DeMille. *Atomic physics: An exploration through problems and solutions, 2nd Ed.*, p.90,95,109,121. Oxford University Press, 2008.
- [58] G. Arfken. *Mathematical Methods for Physicists, 6th Ed.*, Ch.11, p.675–693, volume 40. American Journal of Physics, 1972.
- [59] D. Budker and S. Rochester. *Optically Polarized Atoms, Ch.4*, p.75–81. Oxford University Press, 2008.

- [60] J. M. Boss, K. S. Cujia, J. Zopes, and C. L. Degen. Quantum sensing with arbitrary frequency resolution. *Science*, 356(6340):837–840, 2017.
- [61] D.B. Bucher, D.R. Glenn, J. Lee, M.D. Lukin, and H. Park. High resolution magnetic resonance spectroscopy using solid-state spins. *arXiv:1705.08887v2*, 2017.
- [62] R.G. Lyons. *Understanding Digital Signal Processing, Ch.8, p.320–330*, volume 40. Pearson, 2004.
- [63] D. Drung. High-performance DC SQUID read-out electronics. *Physica C: Superconductivity and its Applications*, 368(1-4):134–140, 2002.
- [64] C. Abel et al. Search for axion-like dark matter through nuclear spin precession in electric and magnetic fields. *arXiv:1708.06367*, 2017.
- [65] K. Blum and D. Kushnir. Neutrino signal of collapse-induced thermonuclear supernovae: The case for prompt black hole formation in SN 1987A. *The Astrophysical Journal*, 828:31, 2016.
- [66] J.H. Chang, R. Essig, and S.D. McDermott. Supernova 1987A constraints on sub-GeV dark sectors, millicharged particles, the QCD axion, and an axion-like particle. *arXiv:1803.00993*, 2018.
- [67] Michael C. D. Tayler, Thomas Theis, Tobias F. Sjolander, John W. Blanchard, Arne Kentner, Szymon Pustelny, Alexander Pines, and Dmitry Budker. Invited review article: Instrumentation for nuclear magnetic resonance in zero and ultralow magnetic field. *Rev. Sci. Instrum.*, 88:091101, 2017.
- [68] M.C.D. Tayler, T.F. Sjolander, A. Pines, and D. Budker. Nuclear magnetic resonance at millitesla fields using a zero-field spectrometer. *J. Magn. Reson.*, 270:35–39, 2016.
- [69] J. C. Allred, R. N. Lyman, T. W. Kornack, and M. V. Romalis. High-sensitivity atomic magnetometer unaffected by spin-exchange relaxation. *Phys. Rev. Lett.*, 89:130801, 2002.
- [70] D. Budker and M. Romalis. Optical magnetometry. *Nat. Phys.*, 3: 227–234, 2007.
- [71] A. Garcon, D. Aybas, J. W. Blanchard, G. Centers, N. L. Figueroa, P. W. Graham, D. F. Jackson Kimball, S. Rajendran, M. G. Sendra, A. O. Sushkov, L. Trahms, T. Wang, A. Wickenbrock, T. Wu, and

- D. Budker. The cosmic axion spin precession experiment (CASPER): a dark-matter search with nuclear magnetic resonance. *Quantum Science and Technology*, 3:014008, 2018.
- [72] J. Preskill, M. B. Wise, and F. Wilczek. Cosmology of the invisible axion. *Phys. Lett. B*, 120:127, 1983.
- [73] M. Dine and W. Fischler. The not-so-harmless axion. *Phys. Lett. B*, 120:137, 1983.
- [74] L.F. Abbott and P. Sikivie. A cosmological bound on the invisible axion. *Phys. Lett. B*, 120:133 – 136, 1983. ISSN 0370-2693.
- [75] R. Catena and P. Ullio. A novel determination of the local dark matter density. *J. Cosmol. Astropart. P.*, 2010:004, 2010. URL <http://stacks.iop.org/1475-7516/2010/i=08/a=004>.
- [76] Gary P. Centers, John W. Blanchard, Jan Conrad, Nataniel L. Figueroa, Antoine Garcon, Alexander V. Gramolin, Derek F. Jackson Kimball, Matthew Lawson, Bart Pelssers, Joseph A. Smiga, Alexander O. Sushkov, Arne Wickenbrock, Dmitry Budker, and Andrei Derevianko. Stochastic amplitude fluctuations of bosonic dark matter and revised constraints on linear couplings. *arXiv:1905.13650*, 2020.
- [77] A.E. Nelson and J. Scholtz. Dark light, dark matter, and the misalignment mechanism. *Phys. Rev. D*, 84:103501, 2011.
- [78] M. Pospelov, S. Pustelny, M. P. Ledbetter, D. F. Jackson Kimball, W. Gawlik, and D. Budker. Detecting domain walls of axionlike models using terrestrial experiments. *Phys. Rev. Lett.*, 110:021803, 2013.
- [79] M. P. Ledbetter, T. Theis, J. W. Blanchard, H. Ring, P. Ganssle, S. Appelt, B. Blümich, A. Pines, and D. Budker. Near-zero-field nuclear magnetic resonance. *Phys. Rev. Lett.*, 107:107601, 2011.
- [80] G. Vasilakis, J. M. Brown, T. W. Kornack, and M. V. Romalis. Limits on new long range nuclear spin-dependent forces set with a $\mathbf{K}-^3\text{He}$ comagnetometer. *Phys. Rev. Lett.*, 103:261801, 2009.
- [81] G.G. Raffelt. *Astrophysical Axion Bounds*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [82] T. Wu, J W Blanchard, G P Centers, N L Figueroa, A Garcon, P. W. Graham, D. F. Jackson Kimball, S Rajendran, Y. V.

- Stadnik, A. O. Sushkov, A Wickenbrock, and D Budker. Search for axionlike dark matter with nuclear spins in a single-component liquid. *arXiv:1901.10843*, 2019.
- [83] K. A. Olive and M. Pospelov. Environmental dependence of masses and coupling constants. *Phys. Rev. D*, 77:043524, 2008.
- [84] B. Blümich, P. Blümli, and J. Jansen. Presentation of sideband envelopes by two-dimensional one-pulse (TOP) spectroscopy. *Solid State Nucl. Mag.*, 1:111 – 113, 1992.
- [85] S. Schmitt, T. Gefen, F.M. Stürner, T. Uden, G. Wolff, C. Müller, J. Scheuer, B. Naydenov, M. Markham, S. Pezzagna, J. Meijer, I. Schwarz, M. Plenio, A. Retzker, L.P. McGuinness, and F. Jelezko. Submillihertz magnetic spectroscopy performed with a nanoscale quantum sensor. *Science*, 356:832–837, 2017.
- [86] P.W. Graham, D.E. Kaplan, J. Mardon, S. Rajendran, W.A. Terrano, L. Trahms, and T. Wilkason. Spin precession experiments for light axionic dark matter. *Phys. Rev. D*, 97:055006, 2018.
- [87] Andrew L. Maas, Peng Qi, Ziang Xie, Awni Y. Hannun, Christopher T. Lengerich, Daniel Jurafsky, and Andrew Y. Ng. Building DNN acoustic models for large vocabulary speech recognition. *Computer Speech and Language*, 41:195–213, 2017.
- [88] Dan Ciregan, Ueli Meier, and Jurgen Schmidhuber. Multi-column deep neural networks for image classification. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, (February):3642–3649, 2012.
- [89] Dan Guest, Kyle Cranmer, and Daniel Whiteson. Deep learning and its application to LHC physics. *Annual Review of Nuclear and Particle Science*, 68:161–181, 2018.
- [90] Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, Varun Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Boochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 969–977, 2018.

-
- [91] Steven G. Worswick, James A. Spencer, Gunnar Jeschke, and Ilya Kuprov. Deep neural network processing of DEER data. *Science Advances*, 4(8):1–18, 2018.
- [92] –VAE: Learning basic visual concepts with a constrained variational framework. *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, pages 1–22, 2017.
- [93] Greg Ver Steeg. Unsupervised Learning via Total Correlation Explanation. *arXiv*, pages 5151–5155, 2017.
- [94] Shuyang Gao, Rob Brekelmans, Greg Ver Steeg, and Aram Galstyan. Auto-encoding total correlation explanation. *AISTATS 2019 - 22nd International Conference on Artificial Intelligence and Statistics*, 2020.
- [95] Marco Farina, Yuichiro Nakai, and David Shih. Searching for new physics with deep autoencoders. *Physical Review D*, 101(7):75021, 2020.
- [96] Mikael Kuusela, Tommi Vatanen, Eric Malmi, Tapani Raiko, Timo Aaltonen, and Yoshikazu Nagai. Semi-supervised anomaly detection - Towards model-independent searches of new physics. *Journal of Physics: Conference Series*, 368(1), 2012.
- [97] Raffaele Tito D’Agnolo and Andrea Wulzer. Learning new physics from a machine. *Physical Review D*, 99(1):1–34, 2019.
- [98] Nati Aharon, Amit Rotem, Liam P. McGuinness, Fedor Jelezko, Alex Retzker, and Zohar Ringel. NV center based nano-NMR enhanced by deep learning. *Scientific Reports*, 9(1):1–11, 2019.
- [99] M. S. Safronova, D. Budker, D. Demille, Derek F. Jackson Kimball, A. Derevianko, and Charles W. Clark. Search for new physics with atoms and molecules. *Reviews of Modern Physics*, 90(2), 2018.
- [100] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- [101] Sepp Hochreiter. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(2):107–116, 1998.

- [102] Boris Hanin. Which neural net architectures give rise to exploding and vanishing gradients? In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 582–591. Curran Associates, Inc., 2018.
- [103] Lovedeep Gondara. Medical Image Denoising Using Convolutional Denoising Autoencoders. *IEEE International Conference on Data Mining Workshops, ICDMW*, 0:241–246, 2016.
- [104] K. S. Kasiviswanathan and K. P. Sudheer. Methods used for quantifying the prediction uncertainty of artificial neural network based hydrologic models. *Stochastic Environmental Research and Risk Assessment*, 31(7):1659–1670, 2017.
- [105] H. M. Dipu Kabir, Abbas Khosravi, Mohammad Anwar Hosen, and Saeid Nahavandi. Neural Network-Based Uncertainty Quantification: A Survey of Methodologies and Applications. *IEEE Access*, 6(c):36218–36234, 2018.
- [106] Yukun Ding, Jinglan Liu, Jinjun Xiong, and Yiyu Shi. Revisiting the evaluation of uncertainty estimation and its application to explore model complexity-uncertainty trade-off. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 2020-June:22–31, 2020.
- [107] Antoine Garcon, John W. Blanchard, Gary P. Centers, Nataniel L. Figueroa, Peter W. Graham, Derek F. Jackson Kimball, Surjeet Rajendran, Alexander O. Sushkov, Yevgeny V. Stadnik, Arne Wickenbrock, Teng Wu, and Dmitry Budker. Constraints on bosonic dark matter from ultralow-field nuclear magnetic resonance. *Science Advances*, 5(10), 2019.
- [108] Hector Masia-Roig, Joseph A. Smiga, Dmitry Budker, Vincent Dumont, Zoran Grujic, Dongok Kim, Derek F. Jackson Kimball, Victor Lebedev, Madeline Monroy, Szymon Pustelny, Theo Scholtes, Perrin C. Segura, Yannis K. Semertzidis, Yun Chang Shin, Jason E. Stalnaker, Ibrahim Sulai, Antoine Weis, and Arne Wickenbrock. Analysis method for detecting topological defect dark matter with a global magnetometer network. *Physics of the Dark Universe*, 28:100494, 2020.
- [109] Teng Wu, John W. Blanchard, Gary P. Centers, Nataniel L. Figueroa, Antoine Garcon, Peter W. Graham, Derek F. Jackson Kimball,

- Surjeet Rajendran, Yevgeny V. Stadnik, Alexander O. Sushkov, Arne Wickenbrock, and Dmitry Budker. Search for axionlike dark matter with a liquid-state nuclear spin comagnetometer. *Phys. Rev. Lett.*, 122:191302, May 2019.
- [110] D. Antypas, O. Tretiak, A. Garcon, R. Ozeri, G. Perez, and D. Budker. Scalar dark matter in the radio-frequency band: Atomic-spectroscopy search results. *Phys. Rev. Lett.*, 123:141102, Oct 2019.
- [111] Dionysios Antypas, Oleg Tretiak, Ke Zhang, Antoine Garcon, Gilad Perez, Mikhail G. Kozlov, Stephan Schiller, and Dmitry Budker. Probing fast oscillating scalar dark matter with atoms and molecules. 6(3):034001, apr 2021.
- [112] Min Jiang, Haowen Su, Antoine Garcon, Xinhua Peng, and Dmitry Budker. Search for axion-like dark matter with spin-based amplifiers, 2021.
- [113] T.F. Sjolander, M.C.D. Tayler, J.P. King, D. Budker, and A. Pines. Transition-Selective Pulses in Zero-Field Nuclear Magnetic Resonance. *J. Phys. Chem. A*, 120:4343–4348, 2016.
- [114] Teng Wu, John W. Blanchard, Derek F. Jackson Kimball, Min Jiang, and Dmitry Budker. Nuclear-spin comagnetometer based on a liquid of identical molecules. *Phys. Rev. Lett.*, 121:023202, 2018.
- [115] Ralph W. Adams, Juan A. Aguilar, Kevin D. Atkinson, Michael J. Cowley, Paul I. P. Elliott, Simon B. Duckett, Gary G. R. Green, Iman G. Khazal, Joaquín López-Serrano, and David C. Williamson. Reversible interactions with para-hydrogen enhance NMR sensitivity by polarization transfer. *Science*, 323:1708–1711, 2009.
- [116] Thomas Theis, Micah P. Ledbetter, Gwendal Kervern, John W. Blanchard, Paul J. Ganssle, Mark C. Butler, Hyun D. Shin, Dmitry Budker, and Alexander Pines. Zero-field NMR enhanced by parahydrogen in reversible exchange. *J. Am. Chem. Soc.*, 134:3987–3990, 2012.
- [117] Jan-Bernd Hövener, Niels Schwaderlapp, Thomas Lickert, Simon B. Duckett, Ryan E Mewis, Louise AR Highton, Stephen M Kenny, Gary GR Green, Dieter Leibfritz, Jan G Korvink, et al. A hyperpolarized equilibrium for magnetic resonance. *Nat. Commun.*, 4:2946, 2013.

-
- [118] Dudari B. Burueva, James Eills, John W. Blanchard, Antoine Garcon, Román Picazo-Frutos, Kirill V. Kovtunov, Igor V. Koptug, and Dmitry Budker. Chemical reaction monitoring using zero-field nuclear magnetic resonance enables study of heterogeneous samples in metal containers. *Angewandte Chemie International Edition*, 59(39):17026–17032, 2020.
- [119] Martin Sufke, Sören Lehmkuhl, Alexander Liebisch, Bernhard Blümich, and Stephan Appelt. Para-hydrogen raser delivers sub-millihertz resolution in nuclear magnetic resonance. *Nat. Phys.*, 13: 568, 2017.
- [120] M. Jiang, H. Li, Z. Zhu, X. Peng, and D. Budker. Floquet-state maser under real-time quantum feedback control. *arXiv:1901.00970*, 2019.
- [121] M. Bishop. *Pattern Recognition and Machine Learning - Chapter.1*. Springer, 2011. ISBN 9780387310732.
- [122] Jeffrey D Scargle. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal*, 263:835–853, 1982.
- [123] A. Bandyopadhyay and D. Majumdar. Diurnal and annual variations of directional detection rates of dark matter. *Astrophysical Journal*, 746, 2012.