Neutron stars as axion laboratories

Harnessing the power of the magnetosphere

Noordhuis, D.

Publication date
2024

1

Introduction

The nature of dark matter, which is believed to comprise 85% of the total mass content of the Universe today, stands as one of the most intriguing mysteries in contemporary physics. Its existence is proposed to reconcile apparent disparities in observations of galaxies, galaxy clusters, the cosmic microwave background, and the large-scale structure of the Universe, all of which seem to require an unseen form of matter that interacts gravitationally. Evidence pointing to the presence of dark matter is thus abundant, yet it is all of a gravitational nature. Indeed, despite years of dedicated research to demonstrate otherwise, no interaction has been detected between dark matter and one of the other fundamental forces, leaving crucial questions regarding its characterization unanswered. Current models generally posit that dark matter is a new particle, or particles, but defining properties like its mass and coupling to the Standard Model of particle physics are still unknown.

Despite the challenges, however, the quest to unravel the nature of dark matter drives a huge collaborative endeavor encompassing, among other things, terrestrial experiments to identify dark matter in the laboratory, indirect searches that look for dark matter signatures from astrophysical environments, and theoretical work to develop dark matter models consistent with observational criteria. This effort provides, and has provided, one of the most promising avenues for advancing our understanding of physics beyond the Standard Model, as well as shedding light on the dynamics and evolution of our Universe across a wide range of spatial and temporal scales. Hence, the pursuit of identifying dark matter is a profoundly important component of many areas in physics, including particle physics, astrophysics, and cosmology.

In this thesis, we focus on one of the current frontrunners among dark matter candidates, a particle known as the axion, and the accompanying discovery potential provided by neutron star magnetospheres. The present chapter serves to motivate and introduce the rest of the dissertation, and is organized as follows. Section 1.1

gives a brief summary of the historical developments surrounding dark matter, thereby focusing on the evidence accumulated to support its existence. Section 1.2 discusses the likely identity of dark matter as a novel particle, and lists some of the most compelling candidates that have been proposed to date. Then, Section 1.3 hones in on axions and reviews the key concepts underlying these fascinating particles. Section 1.4 focuses on general strategies for axion detection, in particular those exploiting the coupling between axions and photons. Section 1.5 specifically examines axion detection with neutron star magnetospheres, where the extreme environment can dramatically amplify the conventionally feeble axion interactions. Finally, Section 1.6 outlines how the remainder of the thesis is organized. The first two sections of this chapter are largely based on the excellent review provided by [1]; see also [2] for an in-depth historical account on dark matter that takes a more philosophical approach. Great reviews that cover the material of Sections 1.3 and 1.4 in more detail can be found in [3–5]. Equivalent reviews covering the neutron star magnetosphere are given in [6, 7].

1.1 Evidence for dark matter

The notion of conjecturing matter in the Universe that is seemingly absent or invisible has a long-standing history among humanity, dating back even to ancient civilization. Focus in these early times was generally on hypothesizing the existence of new celestial bodies, with the rationale rooted not in scientific observation but in philosophical principles such as the pursuit of symmetry. This is exemplified by the Greek philosopher Philolaus, who proposed that the Earth has a hidden twin, known as the Counter-Earth or Antichton, with both bodies orbiting on opposite sides of some 'Central Fire' (which is a distinct entity from the Sun) [8]. What is more, utilizing a combination of theory and observation to search for missing matter has a similarly rich past. Staying in ancient Greece, the Aristotelian worldview dividing the cosmos into concentric spheres centered around the Earth purported to require no additional constituents. This perspective prevailed until the Renaissance, when astronomers decisively refuted the idea through new technological advancements and observation. In particular Galileo and his invention of the telescope played a significant role. By being able to look at the sky in more detail, he managed to elucidate the Milky Way's composition as individual stars rather than nebulae, observe the previously undetected rings of Saturn, and discover the four largest moons revolving around Jupiter. This series of events demonstrates that accepted models of the Universe and its make-up do not always present the complete picture, and that scientific inquiry coupled with technological progress can work to reveal potentially hidden components. In this sense, it perfectly echoes the essence of contemporary dark matter searches.

While there was no observational evidence available at the time, speculation regarding the existence of what we now refer to as dark matter began to emerge toward the end of the 19th century, with initial attempts at estimating a corresponding abundance appearing shortly after. Notably, in 1904 Lord Kelvin sought to evaluate the total number of stars within a large sphere surrounding the Sun by using their velocity dispersion [9]. In doing so, he theorized that *"Many of our stars, perhaps a great majority of them, may be dark bodies."* While discussing the work of Lord Kelvin, Henri Poincaré appears to have coined the term dark matter (*"matière obscure"* in the original French), and observed that the consistency of the velocity dispersion predicted by Lord Kelvin with observations suggested that visible matter should be more prevalent than dark matter [10]. This sentiment was echoed in subsequent years in work by Ernst Öpik [11], Jacobus Kapteyn [12], James Jeans [13], Bertil Lindblad [14], and Jan Oort [15]. They performed local density calculations in more intricate models of the Milky Way, and similarly concluded that non-luminous matter is likely a subdominant component.

In 1933, it was Fritz Zwicky that begun to reverse the narrative. In a seminal paper, he applied the virial theorem to the Coma cluster in order to estimate the expected velocity dispersion of its constituent galaxies, finding a value of 80 km/s [16]. This stood in stark contrast with the average measured velocity dispersion reported by Edwin Hubble and Milton Humason [17], which was as high as 1000 km/s. To obtain such high values via the virial theorem, the Coma cluster would need to be a lot more massive than observed. Zwicky himself declared that *"If this should be verified, it would lead to the surprising result that dark matter exists in much greater density than luminous matter."* Due to usage of an erroneous Hubble constant, which was believed during this period to be a factor ~ 8.5 times larger than the modern value, Zwicky's estimates were off by more than an order of magnitude. Nevertheless, correcting for this mistake still leaves a discrepancy in the Coma cluster's apparent and inferred velocity dispersion, hinting at the presence of unaccounted-for matter. Several years later, a similar analysis of the Virgo cluster performed by Sinclair Smith reached the same conclusion [18].

The majority of the astrophysical community at the time agreed that the results presented by Zwicky and Smith were puzzling, yet far from everyone subscribed to the notion of dark matter as an explanation. Some in particular were not convinced that the virial theorem could be applied to every galaxy cluster, or argued that the galaxies with the highest observed velocities may not be permanent cluster members [19]. This debate simmered for several decades, until a different piece of evidence gained enough traction to provide resolution: the apparent flatness of galaxy rotation curves.

A galaxy rotation curve relates the circular velocity of stars and gas in a galaxy with their radial distance from the galactic center. Standard Newtonian theory dictates

that this relationship has the form

$$v_c(r) = \sqrt{\frac{GM(r)}{r}}\,, \tag{1.1}$$

where G is the gravitational constant and $M(r)$ is the enclosed mass at distance r. Equation 1.1 clearly shows that beyond the extent of ordinary matter in a galaxy, we expect the circular velocity to decrease as $v_c \propto r^{-1/2}$. This is in accordance with the behavior found in other orbital systems such as stars and their planets. Since galaxies can be treated as a non-interacting collection of stars [20], Equation 1.1 can also be inverted, allowing one to construct a continuous galactic mass distribution by measuring enough stellar circular velocities at enough radii.

The first of such measurements was presented in 1939 by Horace Babcock [21]. By determining the Doppler shifts of spectral lines, he was able to map the rotation curve of the Andromeda Galaxy (M31) out to roughly 20 kpc away from its center. Surprisingly, he found that circular velocities continued to increase at large radii, at odds with the expected behavior from Newtonian theory. Moreover approximating Andromeda's mass distribution, he realized that these results could suggest the presence of a substantial amount of mass in the outer regions of the galaxy, although he did not directly reference a link with dark matter. Unfortunately, subsequent research into this phenomenon was limited until technological developments started to allow for more precise measurements. It was not until decades later, in the late 1960s and early 1970s, that the invention of better spectrographs and the advent of radio astronomy, which allowed for complementary rotation curve surveys through observation of the hydrogen 21 centimeter line, reinvigorated the field. Most famously, Vera Rubin and Kent Ford conducted a spectroscopic survey to obtain an updated rotation curve for the Andromeda Galaxy in 1970 [22]. In the years that followed, a number of other rotation curves, including some derived from radio data, also became available (see *e.g.* [23–25]). A subset of these results is depicted in Figure 1.1, where it can be seen that the rotation curves tend to a flat profile at distances far away from the galactic center. Looking at Equation 1.1, this implies that the enclosed mass continues to scale like $M(r) \propto r$, even in regions well outside of the galactic disk where there is little ordinary matter. This spherical envelope that appears to surround galaxies later came to be known as the dark matter halo.

Like the velocity dispersion calculations in galaxy clusters, galaxy rotation curves thus signified the existence of a potentially very considerable missing mass component. And while these two discrepancies were observed on different scales, the link between them was soon established [26, 27]. In the following years, an increasing amount of rotation curve data was put forward that bolstered the case for dark matter, ultimately cementing its apparent necessity as a major question in physics (see also [28] for a comprehensive review published at the time).

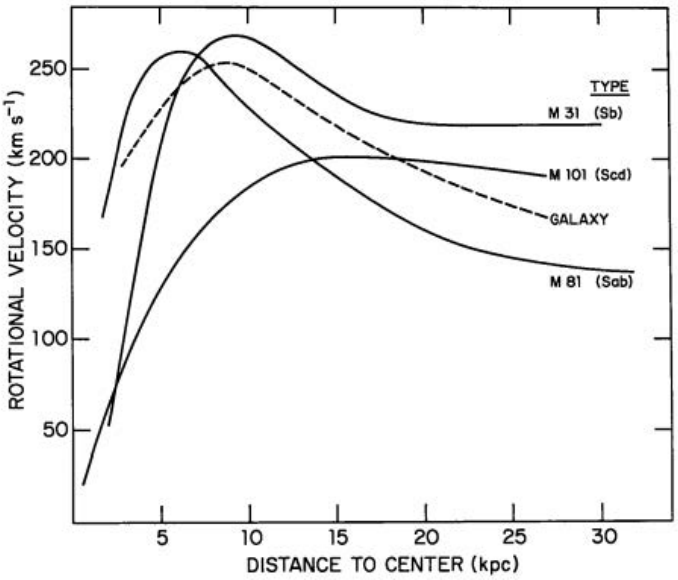

Figure 1.1: The galaxy rotation curves of M31, M81, and M101 (solid lines), as well as the Milky Way (dashed line). This figure is taken from [25], and also includes data from [22–24].

Today, many modern astrophysical and cosmological observations have further reinforced the presented historical evidence. Perhaps most compelling are measurements of the temperature anisotropies in the cosmic microwave background (CMB) [31, 32]. The CMB is a relic from the early Universe, consisting of a faint glow of radiation that permeates space in all directions. It is nearly isotropic and uniform, up to small inhomogeneities that reveal crucial information about the structure, composition, and evolution of the Universe. The distribution of these inhomogeneities, also called anisotropies, can be represented via the use of an angular power spectrum. The peaks and valleys in this power spectrum describe the variations in CMB temperature as a function of angular scale, with their exact locations and heights serving as powerful probes of various cosmological parameters. In particular, the first three peaks respectively provide insight into the total energy density, the total baryonic (*i.e.* ordinary) matter density, and the relative dark matter density content of the Universe. This can be understood in general terms (see *e.g.* [33] for a more complete description) by recognizing that before the CMB was emitted, photons and baryonic matter were tightly coupled through Compton and Coulomb scattering processes. CMB anisotropies primarily arise from density perturbations within this photon-baryon fluid. When an overdense region forms, matter gravitationally attracts more matter toward it, while radiation instead exerts an outward pressure. These opposing effects give rise to oscillations, known as baryon acoustic oscillations (BAOs), which dictate how density

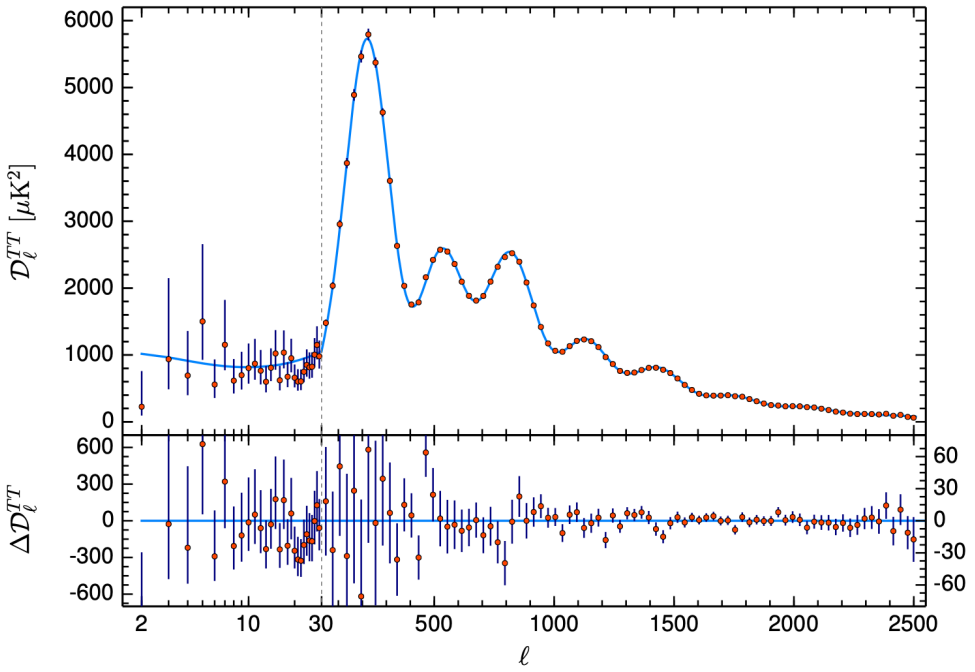

Figure 1.2: The angular power spectrum of CMB temperature anisotropies, as reported by the Planck collaboration [29]. The measured power spectrum is depicted as red dots with error bars, and plotted against angular scale in the form of multipole moment ℓ. The data is compared to a baseline ΛCDM cosmological model, illustrated in blue, which contains dark energy Λ (comprising $\gtrsim 68.3\%$ of the total energy density), dark matter (comprising $\sim 26.5\%$ of the total energy density), and ordinary matter (comprising $\sim 4.9\%$ of the total energy density) – note that dark energy will not be discussed in this thesis, but more information can for example be found in the review provided by [30]. The residuals with respect to the baseline model are presented in the bottom panel.

perturbations evolve[1]. In this sense one can understand that the features of the CMB power spectrum reflect how strongly matter is collapsing into, and bouncing out of, density perturbations under the influence of BAOs. Moreover, dark matter is never coupled to photons and therefore only acts to enhance overdensities, which leads to the dampening of BAOs. This too is captured within the CMB power spectrum. The latest angular power spectrum measured by the Planck satellite, which was the most recent survey aimed at mapping CMB anisotropies, is shown in Figure 1.2. The agreement between theory and observation is incredibly striking, and suggests that our

[1]For completeness, we note that these primordial density perturbations act as seeds for galaxy formation. Consequently, BAOs can also be observed in the distribution of matter at late cosmological times.

Universe contains dark matter in a quantity over five times as abundant as baryonic matter [29]. Other modern findings similarly offer strong support in favor of dark matter. Examples include, but are certainly not limited to, observations of merging clusters [34, 35], large-scale structure [36–38], weak lensing [39], and strong lensing [40, 41]. We emphasize that these pieces of evidence cover nearly the whole age of the Universe (from the emission of the CMB to today), as well as galactic to cosmological scales, and furthermore all predict a dark matter abundance consistent with the CMB result.

Finally, it is important to acknowledge that there also exist alternative explanations for the discussed observational phenomena. These appear in the form of modifications to our accepted theories of gravitation. In the early 1900s, the anomalous precession of Mercury's perihelion had a proposed solution in the shape of an unobserved perturbing planet (dubbed Vulcan). Yet unlike with Galileo's discoveries, there was no new matter to be found here, and it was instead the theory of general relativity that eventually provided the answer. The same could be the case with dark matter. The most well-established alternatives are those of MOdified Newtonian Dynamics (MOND) [42–44] and its associated theories, such as tensor–vector–scalar gravity [45]. These theories have been successful in describing flat galaxy rotation curves, but can, for instance, not yet explain the observed dynamics in galaxy clusters (such as the large velocity dispersions of constituent galaxies) and the anisotropies in the CMB [46–49]. In general, alternative ideas suffer from the vast amount of independent data available working in favor of the dark matter hypothesis – currently, no new theory has been proposed that can fully account for everything.

1.2 The nature of dark matter

The preceding section establishes that there is a large body of evidence pointing toward the probable existence of dark matter. Crucially, however, this conclusion is derived entirely from gravitational probes. Every form of matter interacts via gravity, and therefore such investigations reveal very little about its composition. When early adopters such as Zwicky, Rubin, and Ford published their results, it was generally assumed that dark matter consisted of baryonic entities like stars or gas clouds that were simply too faint to be detected. This class of candidates was later coined MAssive Compact Halo Objects (MACHOs), and at that point also included black holes. Today, the accepted view has instead shifted to an explanation in the shape of one or more non-Standard Model[2] particles. This change in perspective is primarily a consequence of particle dark matter continuing to be consistent with every piece of available data,

[2]The Standard Model of particle physics is a theoretical framework that describes all known elementary particles and their interactions through three of the four fundamental forces (gravity is not included) [50].

while other potential explanations have been gradually ruled out over time[3]. However, it is noteworthy that the past decades also witnessed a sociological shift which may have further encouraged this viewpoint, as an ever-increasing assimilation of particle physics ideas into astrophysics and cosmology have allowed research into a potential particle nature of dark matter to initiate and flourish. Unfortunately, despite narrowing the search, the precise identity of dark matter remains an open question.

There does exist a widespread consensus regarding key properties that any viable particle dark matter candidate must possess. These include:

- Feeble coupling to the Standard Model. If this were not the case, the observed effects today would have extended beyond just gravitational influence. As physicists, we do hope there is some degree of interaction with Standard Model particles to allow for easier dark matter detection, but this need not be true.
- A production mechanism capable of generating an early Universe abundance consistent with the presented CMB observations [29].
- Stability on timescales at least on the order of the age of the Universe, which is $\sim 13.8\,\mathrm{Gyr}$ [29]. This is a simple consequence of the fact that our observations suggest that dark matter was present both before the formation of the CMB (in the early Universe) as well as today.

In addition to the above, the prevailing viewpoint asserts that dark matter is cold, which refers to its constituent particles moving slowly with respect to the speed of light during the era of structure formation [57]. This is in contrast to hot dark matter, whereby the constituent particles are instead relativistic [58]. The distinction is important, as these two types have very different implications for the growth of structure in our Universe. Cold dark matter leads to hierarchical (often referred to as 'bottom-up') structure formation, whereby small initial overdensities seed larger objects that can continue to grow through mergers. Conversely, hot dark matter leads to 'top-down' structure formation, whereby large superclusters form first and subsequently fragment into smaller objects. Numerical simulations swiftly established that only the bottom-up approach leads to the distribution of galaxies and their clusters observed at present, largely excluding hot dark matter as a result [59]. There also exists a middle ground called warm dark matter. This is an alternative hypothesis originally aimed at reconciling some potential shortcomings of the cold dark matter paradigm when comparing numerical simulations to observational data, specifically concerning low-mass galaxies. Today, it is believed that these concerns may also be

[3] In particular, gravitational microlensing surveys [51–53], Big Bang nucleosynthesis [54, 55], and the aforementioned CMB power spectrum [29] have heavily constrained the allowed abundance of baryonic matter, and thereby MACHOs, in the Universe. There is however one related non-particle dark matter candidate that still remains viable; these are black holes that formed prior to Big Bang nucleosynthesis, referred to as primordial black holes [56].

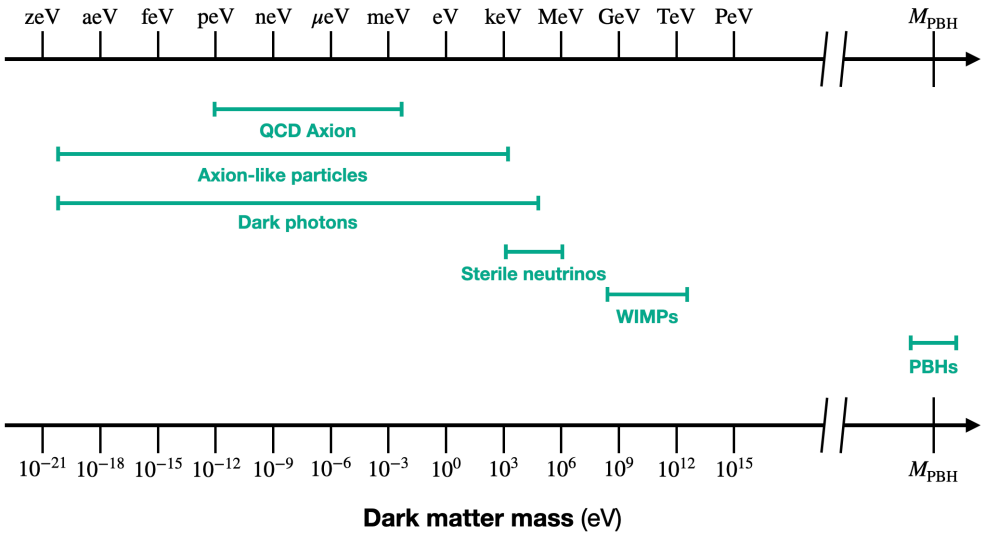

Figure 1.3: Synopsis of the primary candidates for dark matter, shown along with their eligible mass ranges. The mass of primordial black holes can range from scales comparable to asteroids to those of supermassive black holes; however, not all potential masses result in significant densities contributing to the overall matter content of the Universe, see for example [67, 68].

addressed by properly incorporating baryonic matter into simulations, yet warm dark matter remains an intriguing possibility (see [60–62] for examples of recent reviews).

Under the assumption of cold dark matter, one can furthermore derive a general lower limit on the mass of dark matter particles based on the fact that interference effects will erase density variations on scales below their de Broglie wavelength. Accordingly, this de Broglie wavelength cannot be larger than the smallest observable structures in the Universe that contain dark matter (*i.e.* dwarf galaxies). This correspondingly places a lower bound on the dark matter mass at the level of $m_{\mathrm{DM}} \gtrsim 10^{-21} - 10^{-19}\,\mathrm{eV}$ [63–65], depending on precise methodology. In the case of fermionic dark matter, the Pauli exclusion principle further restricts the number of dark matter particles that can occupy dwarf galaxies, resulting in a much more stringent bound of $m_{\mathrm{DM,fer}} \gtrsim 130\,\mathrm{eV}$ [66].

Historically, the first particle dark matter candidate that people turned toward was the neutrino. This is hardly surprising, as neutrinos are the only Standard Model particles that possess most of the essential attributes discussed above. Neutrinos are however an example of hot dark matter, and were therefore quickly excluded by virtue of their effects on structure formation. It thus became evident that a solution would have to be found outside of the Standard Model, and no shortage of candidates has since emerged – an overview of the best-studied ones can be found in Figure 1.3.

The preferred particle candidates are generally motivated through their ability to address other outstanding issues in physics, in addition to acting as dark matter. Premier examples in this regard are sterile neutrinos, weakly interacting massive particles (WIMPs), and axions. Sterile neutrinos are quantitatively similar to standard neutrinos, but do not interact through the electroweak force [69–71]. Besides being a potential dark matter candidate, their main motivation lies in providing a credible mechanism via which standard neutrinos can attain their mass. WIMPs constitute a broad class of particles that find their origin in the Higgs mass hierarchy problem, which suggests that new physics may be present at electroweak energy scales [72–74]. This can for example be realized through supersymmetry [75–79], although non-supersymmetric WIMP candidates also exist. WIMPs gained particular popularity as a result of the so-called 'WIMP miracle', which refers to the observation that thermally produced[4] WIMPs with masses on the electroweak scale naturally lead to the correct dark matter abundance observed today. Despite their attractive qualities, however, the WIMP paradigm currently faces uncertainty as the simplest models are increasingly being ruled out by experiment [81–84]. Axions are light bosonic particles originally conceived to solve the strong CP problem in quantum chromodynamics, which is the theory describing strong interactions [85–88]; later it was realized that they are also an ideal cold dark matter candidate [89–91]. Axions form the primary focus of this dissertation, and we now proceed to exploring their origin, motivation, and detection in the upcoming sections.

1.3 Axions

The Standard Model is a remarkable theory that has served as the backbone of particle and astroparticle physics since its inception. Nevertheless, despite its successes, it is also recognized to be incomplete, with much effort in these fields going out to addressing shortcomings such as the Standard Model not containing any viable dark matter candidate [50]. Of particular importance are those extensions that are capable of tackling multiple problems in tandem. The axion is such an extension, potentially killing two very significant birds with one stone. The first is the addition of dark matter to the Standard Model, and the second is solving the strong CP problem (moreover, axions remain interesting particles even if they ultimately fail to resolve either of these problems, see Section 1.3.4). The strong CP problem functioned as the initial motivation for axions and has its own roots in the so-called $U(1)_A$ problem – we begin this section by exploring these two issues.

[4] Many WIMPs are expected to be thermally produced in the early Universe, through a mechanism known as 'freeze-out' (see *e.g.* [80]).

1.3.1 The $U(1)_A$ and strong CP problems

When the gauge theory of strong interactions, known as quantum chromodynamics (QCD), was being developed, it encountered a troubling complication that came to be known as the $U(1)_A$ problem [92]. Its origin can be understood by considering the QCD Lagrangian, which is given by

$$\mathcal{L}_{\text{QCD}} = \sum_q \bar{q}(i\gamma^\mu D_\mu - m_q e^{i\theta_q \gamma^5})q - \frac{1}{4} G^a_{\mu\nu} G^{a\,\mu\nu} \,. \tag{1.2}$$

Here q represents the quark fields, m_q the quark masses with complex phases θ_q, and the sum is over the different quark flavors (the quark color indices are omitted for brevity). The covariant derivative D_μ couples the quark fields to the gluon fields, which are indexed by a, with the latter also appearing in the gluon field strength tensor $G^a_{\mu\nu}$. In the limit of vanishing quark masses, $m_q \to 0$, this Lagrangian permits a global chiral symmetry of the form $U(N)_V \times U(N)_A$, where N is the total number of quark flavors and V (A) denotes a vector (axial) transformation (note that this should not be confused with the local $SU(3)$ gauge symmetry that defines QCD). Of course, none of the six currently discovered quarks are massless, and this symmetry is therefore not exactly realized in nature. Nonetheless, QCD is expected to be approximately invariant under $U(2)_V \times U(2)_A = SU(2)_V \times SU(2)_A \times U(1)_V \times U(1)_A$. This is a result of the up and down quarks being very light, with masses well below the QCD scale ($\Lambda_{\text{QCD}} \simeq 200\,\text{MeV}$), meaning that they can be reasonably treated as near massless[5].

It turns out that the approximate chiral flavor symmetry $SU(2)_V \times SU(2)_A$ is actually spontaneously broken down by the QCD vacuum to just $SU(2)_V$ [93], which can then be identified with isospin symmetry. As is required, this spontaneous breaking leads to the appearance of three accompanying pseudo-Nambu-Goldstone bosons [94, 95]. These are termed 'pseudo' because the original symmetry is not exact, resulting in Nambu-Goldstone bosons that are not precisely massless, but remain light. The pions are the obvious candidates for fulfilling this role. The second vector component, $U(1)_V$, is a respected symmetry corresponding to baryon number conservation. The second axial component has to be broken similarly to the first, as an unbroken $U(1)_A$ would imply the existence of several particles that we do not observe. This means that we expect a fourth low-mass pseudo-Nambu-Goldstone boson to emerge, akin to the three pions. The absence of such a particle is what constituted the $U(1)_A$ problem.

This conundrum persisted until the mid-1970s, when the discovery of instantons [96] and the accompanying non-trivial structure of the QCD vacuum [97, 98] provided a pathway toward a solution. It was Gerard 't Hooft who subsequently showed that

[5] The strange quark is in fact also relatively light and could be included in this discussion, but this is not necessary for our purposes.

$U(1)_A$ symmetry is anomalous[6] within the more complicated vacuum, and therefore not a true symmetry of quantized QCD [102]. Consequently, it does not undergo spontaneous breaking, thus explaining why there is no fourth pseudo-Nambu-Goldstone boson. While the $U(1)_A$ problem was thereby resolved, unfortunately, in its aftermath a new issue arose.

Associated with the new QCD vacuum there appears an additional term in the Lagrangian of the form[7]

$$\mathcal{L}_\theta = \theta \frac{g_s^2}{32\pi^2} G^a_{\mu\nu} \tilde{G}^{a\,\mu\nu} \,, \tag{1.3}$$

where g_s is the QCD coupling constant, $\tilde{G}^{a\,\mu\nu} = \epsilon^{\mu\nu\alpha\beta} G^a_{\alpha\beta}/2$ is the dual of the gluon field strength tensor, and θ is a phase parameter characterizing the new vacuum. This term changes sign under the combined action of charge conjugation and parity reversal, making it CP violating. Our original Lagrangian, Equation 1.2, also exhibits CP violation through the complex quark phases. In the classical theory this can be removed by performing a chiral rotation, *i.e.*

$$q \to e^{i\alpha_q \gamma^5/2} q \,, \qquad \bar{q} \to \bar{q} e^{i\alpha_q \gamma^5/2} \,, \tag{1.4}$$

and afterward setting $\alpha_q = -\theta_q$. The presence of Equation 1.3 in the quantized theory, however, implies that any rotation of the θ_q can be absorbed via a redefinition of θ [97, 98]. Hence, even when removing CP violation from the quark mass terms, Equation 1.3 will remain CP violating (and vice versa). Inclusion of this term thus transforms QCD into a manifestly CP violating theory.

Within the broader context of the Standard Model, in particular also including electroweak interactions, the phase θ is replaced by $\bar{\theta} = \theta + \arg\det M_u M_d$ when applying the above chiral rotations. Here M_u and M_d denote the up and down quark mass matrices. In this way, all CP violation within the QCD sector is captured by the $G\tilde{G}$ term, with the degree of CP violation being measured solely by $\bar{\theta}$. An observational consequence of this parameter is that it induces an electric dipole moment for the neutron, which is calculated to be [105]

$$d_n = (2.4 \pm 1.0) \times 10^{-16}\, \bar{\theta}\, e\,\mathrm{cm} \,, \tag{1.5}$$

where e is the elementary charge. The current experimental upper bound on the neutron electric dipole moment is $|d_n| < 1.8 \times 10^{-26}\, e\,\mathrm{cm}$ (90% confidence level) [106], implying

$$|\bar{\theta}| \lesssim 10^{-10} \,. \tag{1.6}$$

[6] We call a symmetry anomalous when it is respected within a classical theory, but breaks when generalizing to the full quantum theory [99–101].

[7] For more detailed accounts regarding the origin of this term, as well as the solution to the $U(1)_A$ problem, one can for instance consult [103, 104].

This is a remarkably small number. As a free phase parameter, θ could a priori take any value between 0 and 2π. Why does it then assume a value precisely canceling the contribution from the quark mass matrices, which has a completely unrelated origin? This question, which is equivalent to asking why QCD seems to preserve CP, despite appearing as though it should not, is referred to as the strong CP problem.

1.3.2 Solving the strong CP problem with axions

There is no apparent Standard Model or anthropic explanation for the smallness of $\bar{\theta}$ [107], suggesting that new physics may be required to solve this puzzle. Perhaps the most elegant solution in this regard is that of Roberto Peccei and Helen Quinn, who proposed the Peccei-Quinn (PQ) mechanism [85, 86][8]. Simply put, this involves the introduction of a new global $U(1)_{\mathrm{PQ}}$ symmetry that is spontaneously broken and thereby allows $\bar{\theta}$ to dynamically settle to zero. Frank Wilczek and Steven Weinberg separately pointed out that this broken symmetry then also necessitates the existence of a Nambu-Goldstone boson, which Wilczek named the axion [87, 88].

In more elaborate terms, the PQ mechanism introduces a pseudo-scalar field $a(x)$, hereby referred to as the axion field, that adheres to a global shift symmetry of the form $a(x) \to a(x) + \alpha f_a$ (*i.e.* a $U(1)_{\mathrm{PQ}}$ symmetry). The energy scale f_a, which is closely related to the $U(1)_{\mathrm{PQ}}$ symmetry breaking scale, is termed the axion decay constant. A further defining quality of this $U(1)_{\mathrm{PQ}}$ symmetry is that it is anomalous under QCD. Analogous to the previously discussed $U(1)_A$ symmetry, therefore, it generates a term in the Lagrangian given by

$$\mathcal{L}_a \supset -\frac{a}{f_a}\frac{g_s^2}{32\pi^2}G^a_{\mu\nu}\tilde{G}^{a\,\mu\nu}\,. \tag{1.7}$$

This converts the total prefactor of the $G\tilde{G}$ term into a dynamical field $\bar{\theta}_{\mathrm{eff}} = \bar{\theta} - a/f_a$. The anomaly moreover gives axions a mass $m_a \simeq \Lambda^2_{\mathrm{QCD}}/f_a$ (hence axions, like pions, are actually pseudo-Nambu-Goldstone bosons), which manifests in the Lagrangian as the lowest-order term in the expansion of an axion potential $V(a)$. Given sufficient time, the axion field will evolve and settle into the minimum of this potential. With the use of chiral perturbation techniques it can be shown that the corresponding vacuum expectation value equals $v_a = \bar{\theta} f_a$ [112, 113], coinciding exactly with the value required to set $\bar{\theta}_{\mathrm{eff}} = 0$. This same conclusion can also be derived by applying the Vafa-Witten theorem [114]. Including the axion field into QCD thus naturally relaxes $\bar{\theta}_{\mathrm{eff}}$ to its CP conserving value, and in doing so resolves the strong CP problem.

[8] Other possible solutions to the strong CP problem also exist, such as the Nelson-Barr mechanism [108, 109] and a massless up quark [110]. Nevertheless, these ideas are generally less well-motivated, with the latter for example being heavily challenged by lattice QCD simulations [111].

1.3.3 Invisible axions

In the original implementation of axions, called the Peccei-Quinn-Weinberg-Wilczek (PQWW) model, a $U(1)_{\mathrm{PQ}}$ symmetry appears through the introduction of a second Higgs doublet that is added to the Standard Model in addition to the existing Higgs boson [85–88]. Importantly, this fixes the axion decay constant in the PQWW model to be near the electroweak scale, $f_a \simeq 250\,\mathrm{GeV}$. Such values were quickly ruled out, however, as they would lead to axion masses exceeding experimental constraints [115–117]. In order to avoid these bounds, f_a needed to be significantly larger than initially anticipated and, consequently, axions needed to be lighter and more feebly interacting (as both the axion mass and the coupling strengths between axions and the Standard Model scale inversely with f_a). The result of this realization was the emergence of 'invisible' axion models, wherein the axion decay constant became a free parameter (*i.e.* a parameter that can take on any value). The first and most famous examples of such models are the Kim-Shifman-Vainshtein-Zakharov (KSVZ) axion [118, 119] and the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) axion [120, 121].

The KSVZ axion model is arguably the most illustrative of the two, and we will quickly examine it to see how some of the concepts discussed in the previous subsection materialize. In this model, a new quark ψ and a new complex scalar ϕ (which is a Standard Model singlet) are added to our theory. This manifests as the Lagrangian

$$\mathcal{L}_{\mathrm{KSVZ}} = i\bar{\psi}\gamma^\mu D_\mu \psi + \partial^\mu \phi^\dagger \partial_\mu \phi - V(|\phi|) + (y\bar{\psi}_L \phi \psi_R + \mathrm{h.c.})\,, \tag{1.8}$$

where we have included kinetic terms, a potential term for ϕ, and a Yukawa interaction term with strength y. The left- and right-handed projections are defined in the standard way as $\psi_L \equiv \frac{1}{2}(1-\gamma^5)\psi$ and $\psi_R \equiv \frac{1}{2}(1+\gamma^5)\psi$. This Lagrangian accommodates a $U(1)_{\mathrm{PQ}}$ symmetry described by

$$\phi \to e^{i\alpha}\phi\,, \qquad \psi_L \to e^{i\alpha/2}\psi_L\,, \qquad \psi_R \to e^{-i\alpha/2}\psi_R\,, \tag{1.9}$$

which is spontaneously broken through the potential (note that this is not the axion potential)

$$V(|\phi|) \propto \left(|\phi|^2 - \frac{v_\phi^2}{2}\right)^2\,. \tag{1.10}$$

The absolute minimum of the potential lies at $|\phi| = v_\phi/\sqrt{2}$, with v_ϕ being the vacuum expectation value of ϕ, and it is occupied by an infinite number of states. This becomes evident when decomposing the scalar field into polar coordinates

$$\phi = \frac{v_\phi + \rho}{\sqrt{2}} e^{ia/v_\phi}\,. \tag{1.11}$$

This system thus adheres to a $U(1)_{\mathrm{PQ}}$ symmetry that is spontaneously broken by the vacuum expectation value v_ϕ. The angular mode a can be identified with the pseudo-scalar axion (*i.e.* the corresponding Nambu-Goldstone boson). The radial mode ρ represents a different, very heavy particle that is generally integrated out of the low-energy theory. Doing so results in the Lagrangian

$$\mathcal{L}_{\mathrm{KSVZ}} = i\bar{\psi}\gamma^\mu D_\mu \psi + \frac{1}{2}(\partial_\mu a)^2 + m_\psi \bar{\psi} e^{i a \gamma^5 / v_\phi}\psi + \dots \,. \tag{1.12}$$

Here, we have defined $m_\psi \equiv y v_\phi/\sqrt{2}$ and rewrote the Yukawa interaction in terms of ψ to make the connection with Equation 1.2. The PQ symmetry transformations become

$$a \to a + \alpha v_\phi \,, \qquad \psi_L \to e^{i\alpha/2}\psi_L \,, \qquad \psi_R \to e^{-i\alpha/2}\psi_R \,, \tag{1.13}$$

and we remind the reader that these transformations are anomalous under QCD. Accordingly, chiral rotations on the mass term in Equation 1.12 give rise to terms $\propto G\tilde{G}$, *i.e.* Equation 1.7, analogous to what happens when applying chiral rotations to the quark mass terms in the original QCD Lagrangian. It is also clear from Equation 1.13 that we can identify v_ϕ with the axion decay constant f_a in this scenario, but this correspondence is not exact for every axion model.

The defining characteristics of invisible axions make them excellent dark matter candidates, as they are feebly interacting and stable on cosmological timescales. There moreover exist production methods capable of supplying ample axions to account for the dark matter in the Universe, the most prominent of which are the misalignment mechanism [89–91] and decay via topological defects [122, 123]. These production methods are non-thermal[9], meaning axions are created with minimal velocity dispersion. They can thus be classified as cold dark matter, consistent with our expectations from structure formation and the ΛCDM model. The misalignment mechanism is an unavoidable consequence of a scalar field being present in an expanding Universe. This can be illustrated by looking at the cosmological equation of motion governing the axion field, given by [89–91]

$$\frac{d^2\phi_a}{dt^2} + 3H\frac{d\phi_a}{dt} + m_a^2\phi_a = 0 \,, \tag{1.14}$$

where $\phi_a = a/f_a$, H is the Hubble rate, and we have chosen a fixed axion potential $V(a) = m_a^2 a^2/2$. The second term in this equation accounts for the expansion of the Universe, which acts as a friction term for the axion field. Consequently, we find that in the early Universe, when $H \gg m_a$, the axion field retains a constant amplitude due to the friction term dominating its evolution. The misalignment mechanism relies

[9] Thermal production of axions is also a possibility, and may generate a non-negligible hot axion population for axion masses $m_a \gtrsim 0.1\,\mathrm{eV}$ [124–127].

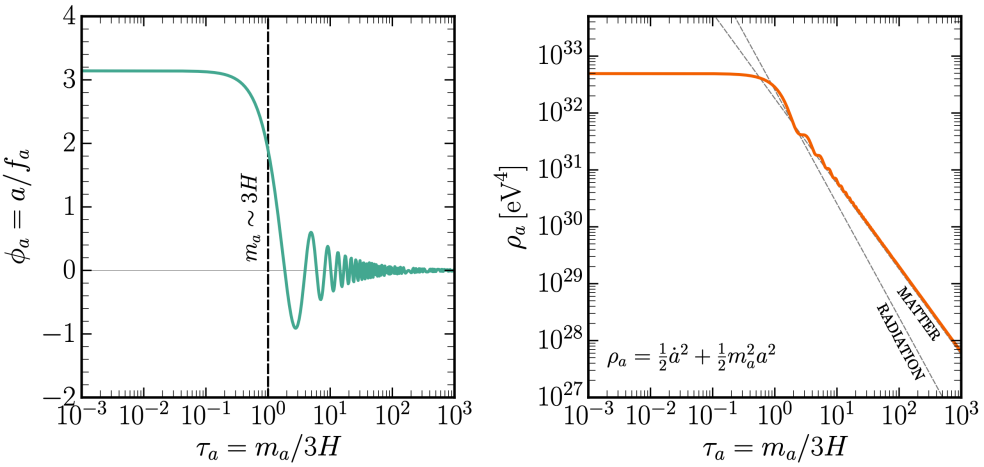

Figure 1.4: The axion misalignment mechanism, originally from [130]. *Left:* Plot of the axion field versus dimensionless time τ_a. The initial value of the axion field (set here to $\phi_i = \pi$) remains frozen until the Hubble rate reaches $H \simeq m_a/3$. At this point, given that its initial value is misaligned from the minimum of the axion potential, the axion field will begin to exhibit damped oscillations. *Right:* Plot of the axion energy density, $\rho_a = (\dot{a}^2 + m_a^2 a^2)/2$, versus dimensionless time τ_a. We also show the energy density evolution of matter ($\rho_m \propto a^{-3}$) and radiation ($\rho_r \propto a^{-4}$). Once oscillations for the axion field initiate, it clearly starts behaving like a component of matter.

on the premise that this initial amplitude typically differs from the minimum of the axion potential. Therefore, once the Hubble rate decreases to a value $H \simeq m_a/3$, the axion field starts to undergo damped oscillations around the potential minimum. Importantly, this gives rise to axions that behave exactly as matter. We portray the axion misalignment mechanism in Figure 1.4. Topological defects are irregular configurations of matter that arise from spontaneously broken symmetries in the early Universe, through a process known as the Kibble mechanism [128, 129]. In the case of a $U(1)_{\rm PQ}$ symmetry they take the shape of axion strings or domain walls whose decay can source a population of axions. Whether they form, however, depends on the time at which the $U(1)_{\rm PQ}$ symmetry gets broken, particularly if this occurs before/during or after inflation.

If the $U(1)_{\rm PQ}$ symmetry was broken post inflation, we would find that the now observable Universe consists of many different random patches containing different initial axion field values, as these patches were out of causal contact at the time of symmetry breaking. In this case, the appropriate quantity to use is the root-mean-square of all possible initial field values, which can be computed, and for example equals $\sqrt{<\phi_i^2>} = \pi/\sqrt{3}$ for a quadratic axion potential [3, 5]. Given that the abundance of axions generated via the misalignment mechanism is determined solely by the initial

axion field and the axion mass (see *e.g.* [5]), there exists in this scenario a unique axion mass that gives rise to the observed dark matter content today. That is, if the misalignment mechanism is the only production method. This is not the case, however, and the inclusion of topological defects notably complicates the discussion, as the amount of axions produced via their decay is relatively hard to quantify. Predictions on the axion mass required to saturate the present dark matter abundance generally lie around $m_a \simeq 10^{-5} - 10^{-3}\,\mathrm{eV}$ [131–136]. Nevertheless, we stress that these calculations currently remain subject to uncertainties.

If the $U(1)_{\mathrm{PQ}}$ symmetry was broken prior to inflation, we would find that inflation expands one of the abovementioned random patches to encompass the entire observable Universe, resulting in a homogeneous initial axion field value everywhere. This also dilutes the network of topological defects, and accordingly, the misalignment mechanism is the sole source of axion production in this scenario. However, unlike in the post-inflationary breaking case, the initial axion field value is now an unknown parameter. As a consequence, we cannot directly predict a single axion mass that reproduces the correct contemporary dark matter abundance. Instead, we need to take into account the uncertainty regarding the initial axion field, which for reasonable assumptions (*i.e.* $\phi_i \in [0.3, 3]$) translates to possible masses in the range $m_a \simeq 10^{-6} - 10^{-4}\,\mathrm{eV}$ [4, 137]. Nonetheless, arguments can also be made for more fine-tuned configurations suggesting alternative axion masses, see [138–140].

1.3.4 Further motivation for axions

In addition to providing a leading solution to both the dark matter problem and the strong CP problem, there are a multitude of other reasons further motivating the pursuit of axions. One example is axiogenesis [141], which is a proposed explanation for the baryon asymmetry in the Universe that involves $U(1)_{\mathrm{PQ}}$ symmetry breaking. Baryon asymmetry is another important phenomenon currently left unexplained by the Standard Model [142]. Axions moreover appear ubiquitously in string theory [143–145] and may be components of grand unified theories [146–148], where they potentially offer a window into high-energy physics. It is important to note, however, that when discussing axions in such contexts, it is generally assumed that they do not require any association with the PQ mechanism. These particles, instead, are light, feebly interacting pseudo-Nambu-Goldstone bosons termed 'axion-like particles' (ALPs) [149, 150]. This differentiates them from the 'QCD axion', which shares the same fundamental properties but furthermore solves the strong CP problem[10].

The focus in this section has until now been on the QCD axion, but ALPs serve as an equally attractive dark matter candidate. Crucially, they have the benefit that their mass is considerably less constrained owing to their distinct origin. As noted

[10] Additionally, string theory ALPs do not generically form topological defects in the early Universe, see [151].

earlier, the QCD axion attains its mass through a chiral anomaly. This yields a precise dependence on the axion decay constant, largely independent of the chosen axion model, which is calculated to be [113]

$$m_a = (5.70 \pm 0.07) \times \left(\frac{10^{12}\,\mathrm{GeV}}{f_a} \right) \mu\mathrm{eV}\,. \tag{1.15}$$

This restrictive relationship coupled with cosmological arguments regarding QCD axion production (via misalignment and topological defects) suggests the likely mass ranges for axion dark matter provided in the previous subsection. ALPs, however, can a priori have any mass, and therefore significantly extend the potential parameter space that axions may inhabit. It is furthermore worth noting that the existence of ALPs does not prohibit the existence of the QCD axion. Indeed, several types of axions may be realized in nature, with in that case the QCD axion potentially just being one of them. Such configurations even arise naturally in string theoretical models [143–145]. Due to their likeness, many axion experiments and other searches make no distinction between ALPs and the QCD axion. The same is true for the material discussed in this thesis. In what follows, therefore, we simply collectively refer to both particle types as axions.

1.4 Axion detection

Having established axions as a promising extension to the Standard Model and a compelling candidate for dark matter, our attention now turns to the question of their detection. Crucially, we remind the reader that a non-gravitational signature is necessary to ascertain whether dark matter is composed of axions. In pursuit of this goal, a plethora of ideas for axion searches have been developed and operationalized over the years, several of which we highlight in this section. Detecting axions, much like any potential dark matter particle, is naturally extremely challenging due to their feeble interactions and the vast range of masses that requires investigation. Nonetheless, we have currently entered a very exciting phase for axion research, with recent times especially seeing considerable effort directed toward devising new search strategies and constructing axion detection experiments.

Axions have the capability to interact with various particles in the Standard Model, including for example photons, electrons, and nucleons. While all of these couplings are subject to investigation, the main emphasis has traditionally been on probing the axion-photon coupling, as this arguably offers the most promising detection opportunities. The work presented here conforms to this approach, and as such, we will focus exclusively on axion interactions with electromagnetism. We review current constraints on these interactions and some of the efforts that led to them, but first, we proceed to a more general discussion of their phenomenology.

1.4.1 The axion-photon coupling

The coupling between axions and electromagnetism is a generic expectation across axion models. It generates several noteworthy effects, all of which can be examined by considering the following Lagrangian [152]

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - A_\mu J^\mu + \frac{1}{2}(\partial_\mu a)^2 - \frac{1}{2}m_a^2 a^2 - \frac{1}{4}g_{a\gamma\gamma}F_{\mu\nu}\tilde{F}^{\mu\nu}a\,. \tag{1.16}$$

This Lagrangian encapsulates both electromagnetism and the axion, along with an interaction between the two that emerges via charged quark loops. The corresponding coupling strength is given by $g_{a\gamma\gamma}$, also called the axion-photon coupling constant. Equation 1.16 furthermore features the electromagnetic field strength tensor $F_{\mu\nu}$ and its dual, the four-current $J^\mu = (\rho, \vec{J})$, the photon field A_μ, and the axion field a. When investigating axions, particularly within the context of dark matter, their number densities are often very high because the particles themselves posses little mass. Consequently, axions tend to exhibit behavior more akin to that of a classical field, rather than acting as a collection of individual particles. Photons operate in the same manner, allowing for their description in the form of electric and magnetic fields. This similarity motivates the inclusion of axions in our existing theory of electrodynamics, where they are treated using a generalization of Maxwell's equations. The relevant formulas can be found by considering the equations of motion resulting from Equation 1.16, and are given by[11]

$$\nabla\cdot\vec{E} = \rho - g_{a\gamma\gamma}\vec{B}\cdot\nabla a\,, \tag{1.17}$$

$$\nabla\cdot\vec{B} = 0\,, \tag{1.18}$$

$$\nabla\times\vec{E} + \partial_t\vec{B} = 0\,, \tag{1.19}$$

$$\nabla\times\vec{B} - \partial_t\vec{E} = \vec{J} + g_{a\gamma\gamma}(\vec{B}\partial_t a - \vec{E}\times\nabla a)\,, \tag{1.20}$$

$$(\partial_t^2 - \nabla^2 + m_a^2)a = g_{a\gamma\gamma}\vec{E}\cdot\vec{B}\,. \tag{1.21}$$

Here $\vec{E}$ and $\vec{B}$ are the electric and magnetic field, respectively, and we have utilized the relation $F_{\mu\nu}\tilde{F}^{\mu\nu} = -4\vec{E}\cdot\vec{B}$.

The above are known as the axion-modified Maxwell's equations, and they unveil different ways in which axions interact with electromagnetism. One can for example infer that the axion enters as an effective four-current, with associated charge density $\rho_a = -g_{a\gamma\gamma}\vec{B}\cdot\nabla a$ and current density $\vec{J_a} = g_{a\gamma\gamma}(\vec{B}\partial_t a - \vec{E}\times\nabla a)$. These contributions indicate that the axion field is able to induce an electric field through the presence of a magnetic field, and conversely, a magnetic field through the presence of either an electric or magnetic field. Similarly, the axion equation of motion, Equation 1.21, shows

[11] Note that in this thesis we work in natural units, where $\epsilon_0 = c = \hbar = 1$. We furthermore use the metric signature $(+,-,-,-)$.

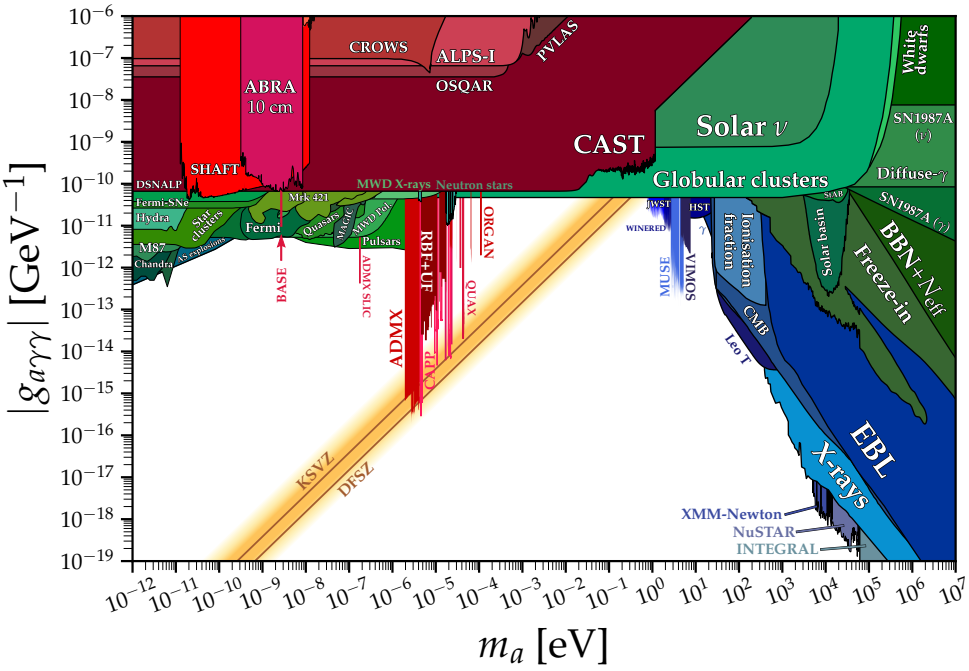

Figure 1.5: An overview of axion-photon parameter space. The colored regions are excluded through various search strategies. These are grouped in three categories: laboratory experiments (red), astrophysical searches (green), and astrophysical searches that have to assume axions are dark matter (blue). The QCD axion band is displayed in yellow. Unexplored parameter space is left as white. This plot is produced by [154], and references for all constraints can be found therein.

that $\vec{E} \cdot \vec{B}$ acts as a source term for axions, meaning that they may be produced via the existence of a non-vanishing $\vec{E} \cdot \vec{B}$ field. The axion-modified Maxwell's equations can furthermore be used to derive an electromagnetic wave equation involving axions, which reveals the potential of axions and photons to undergo mixing while traversing a background electromagnetic field. All of these interactions scale with the strength of the electric and magnetic fields involved. The latter two effects moreover form the basis for the findings presented in this dissertation – we introduce these more comprehensively in Section 1.5. Lastly, we note that there exist more possible interactions between axions and electromagnetism, including for instance axion decay into two photons (see *e.g.* [153] and references therein), but these are better studied by considering the axion-photon coupling vertex separate from Maxwell's equations.

Endeavors to observe axions via their coupling to electromagnetism can be neatly summarized using Figure 1.5[12]. No decisive detection of axions has yet been made, leading searches to produce alternative results in the form of constraints on the axion-photon coupling and axion mass. All presently established constraints are illustrated in color in the figure. The QCD axion band is also shown, representing those combinations of coupling and mass that yield an axion capable of solving the strong CP problem. This is depicted as a linear relationship because conventional QCD axion models predict a proportionality $g_{a\gamma\gamma} \propto m_a \propto f_a^{-1}$. To provide a sense of how the different constraints are determined, we will now describe a selection of the underlying efforts. In doing so, we stay in theme with Figure 1.5 and consider two separate categories: laboratory and astrophysical searches.

1.4.2 Axion detection in the laboratory

There are various strategies that can be employed to detect the axion-photon coupling, each with its own advantages and drawbacks. An obvious approach is to build a detector here on Earth and attempt to measure axions in the laboratory. This has the considerable benefit of allowing one to carefully design an experimental setup and mitigate possible confounding backgrounds. The first axion detectors came online in the late 1980s and early 1990s, and ever-increasing interest has led to a surge in the development of new instruments in recent years. Large-scale coherent magnetic fields are easier to create in the laboratory than large-scale coherent electric fields, and therefore nearly all experiments utilize a background magnetic field to try and generate axion signatures. Nevertheless, despite this similarity, detectors can differ substantially in terms of setup and scale.

The oldest and most developed type of axion detector is known as the axion haloscope. Haloscopes consist of a cavity filled with a magnetic field, typically on the order of $10^4 - 10^5$ G, and are designed to identify the conversion of ambient dark matter axions passing through the apparatus. The cavities are tunable, meaning that they can be reshaped in order to resonate with a particular axion mass. This significantly amplifies the induced signal, but also confines each experimental run to only scan a narrow range of axion masses. Accordingly, haloscopes have the capability to probe relatively low axion-photon couplings, yet they require a long time to cover substantial portions of parameter space. Examples of axion haloscopes are ADMX [160–162], CAPP [163], HAYSTAC [164–166], ORGAN [167], MADMAX [168], and QUAX [169, 170]. Axion helioscopes stand as a second well-established class of axion detection experiments, with the CERN Axion Solar Telescope (CAST) [171, 172] being the current benchmark. CAST employs a decommissioned test magnet from the Large

[12]While this plot showcases all regions of interest for the QCD axion and a large amount for ALPs, additional parameter space does exist for very high-mass axions [155–158] and very low-mass axions (*i.e.* fuzzy dark matter [159]).

Hadron Collider to convert potential axions created by the Sun into X-rays, which can subsequently be observed. The Sun may produce axions when thermal photons interact with the ambient electromagnetic field supplied by charged particles in the solar plasma, through a process known as the Primakoff effect (see Figure 1.7). Notably, as opposed to haloscopes, helioscopes do not operate on the assumption that axions comprise dark matter. They also require no tuning because solar axions are relativistic (in contrast to dark matter axions), which enables the axion and photon four-momentum to naturally align. In return, however, helioscopes have a smaller reach in axion-photon coupling due to the signal requiring two conversions: first, the production of axions in the Sun, and second, their reconversion in the detector, with each process scaling with the coupling squared. Other laboratory searches that examine the interaction between axions and electromagnetism include light-shining-through-walls studies such as ALPS [173, 174] and OSQAR [175], haloscope alternatives operating at lower axion masses such as ABRACADABRA [176, 177] and SHAFT [178], and polarization experiments such as PVLAS [179]. We emphasize that this list is by no means exhaustive – more detailed versions can for example be found in [4, 152].

1.4.3 Axion detection in astrophysics

Laboratory searches provide a direct means of measuring axions as they appear in your experimental setup. Axions can also be detected indirectly, however, through observation of signals generated in astrophysical environments where axions may have played a role. This approach has three important merits. First, the conditions found in astrophysical settings are usually far more extreme than those in controlled laboratories, and therefore able to induce much stronger axion effects. For example, neutron stars can harbor magnetic fields as powerful as 10^{14} G, which is over nine orders of magnitude larger than what we manage to produce in haloscopes and helioscopes. Similarly, experiments conducted on Earth are limited by the local abundance of axions, whereas indirect searches can examine environments where axion densities may be larger. Second, astrophysical efforts often naturally cover a broad array of axion masses while probing relatively low axion-photon couplings. In this way, they are a great complimentary tool to laboratory searches, particularly in serving as a guide for haloscopes. Third, when considering axion dark matter, laboratory experiments are necessarily constrained to probe the product of the local dark matter density times the axion-photon coupling. This is not the case for astrophysical searches, and they can therefore play a key role in helping to break this degeneracy.

There are also difficulties associated with indirect axion detection. Notably, the inability to control astrophysical systems to the same extent as an experiment means that there are generally more uncertainties pertaining to the calculation of anticipated signals. Moreover, while the in situ signals can be very strong, this may be compensated by the large distances that are traversed before reaching Earth. During this journey,

signal power diminishes with distance according to the inverse square law, *i.e.* as $1/d^2$. Note, however, that this decay does not apply to every potential observable, as some may also originate only from axions affecting the local environment (axions may for instance modify the cooling rate of neutron stars, see [180]).

This thesis focuses on axion detection using neutron star magnetospheres, and we devote the next section to discussing the fundamentals of this field specifically. As evidenced by Figure 1.5, there are many other examples of indirect astrophysical searches involving the axion-photon coupling. The variety of methods employed is in fact much greater than in the case of laboratory experiments, and we therefore refrain from giving a summary here that would inevitably be largely incomplete. Instead, we refer the reader to [181] for a recent comprehensive account.

1.5 Axion detection with neutron star magnetospheres

Neutron stars are dense[13] celestial objects, primarily composed of tightly packed neutrons, resulting from the gravitational collapse of massive stars that can no longer support nuclear fusion. They offer one of the most extreme environments known to humanity, and the characteristics of these environments make them incredibly promising arenas for axion detection. The main reasoning behind this is twofold. First, as mentioned in the preceding section, neutron star magnetic fields are enormous, the strongest in the Universe in fact, which dramatically enhances the power of any axion-induced phenomena. Their magnetic fields are furthermore coherent over large length scales[14], also providing a vast region over which the relevant phenomena can occur. Second, neutron stars are surrounded by an ambient plasma known as the magnetosphere, which is highly conducive to certain axion-photon interactions. In particular, the plasma attributes photons with an effective mass that allows for the axion and photon four-momentum to become exactly equal, boosting the probability of mixing between these particles. Moreover, the neutron star magnetosphere is expected to contain regions of non-vanishing $\vec{E} \cdot \vec{B}$ that can source a substantial local population of axions – the axion densities produced in this way can potentially overcome the feeble axion interaction strength that normally troubles detection methods[15]. The core objective of this thesis is to explore the role these effects play in neutron star magnetospheres and how their occurrence can potentially lead to observational

[13] Neutron stars have radii on the order of $\sim 12\,\mathrm{km}$ and masses between $1 - 2\,\mathrm{M}_\odot$ [182].

[14] The coherence length of neutron star magnetic fields can be on the order of kilometers close to the star, and grow to the order of thousands of kilometers in more distant regions.

[15] For completeness, we mention that efforts also exist which try to utilize different interactions, like conversion of thermal axions produced inside of the neutron star [183]. However, the detection prospects for these methods are typically weaker.

consequences. We introduce this specific topic here, starting with a description of the relevant aspects of neutron star magnetospheres.

1.5.1 The neutron star magnetosphere

Neutron stars are known to contain a magnetic field in their interior, believed to be generated via the conservation of magnetic flux during the collapse of the progenitor star. This internal magnetic field has to be continuous at the star's surface, inevitably producing a magnetic field outside of the star as well. The external magnetic field, in turn, necessarily leads to the formation of a magnetosphere. This last fact was first recognized in 1969 by Peter Goldreich and William Julian, who pioneered the model for neutron star magnetospheres [184]. They noticed that a rotating neutron star cannot exist in vacuum, as the rotation of its magnetic field should induce an electric field strong enough to overpower gravity and strip charges from the neutron star surface. These charges are confined to move along magnetic field lines and therefore pulled into corotation with the star, which can be sustained for field lines contained inside the light cylinder (defined as the radial distance where corotating particles reach the speed of light). This plasma is what constitutes the neutron star magnetosphere. Within the light cylinder, the dynamics of the magnetosphere are dominated by the powerful electromagnetic field, resulting in it being accurately described as a force-free system. This indicates that the Lorentz force acting on plasma particles is negligible, which has the well-known consequence that the plasma will move to screen the electric field lying parallel to the magnetic field, *i.e.* $\vec{E} \cdot \vec{B} = 0$. The corresponding plasma charge density is given by the Goldreich-Julian (GJ) density, which is derived by finding the minimum charge density allowed by a stable, self-consistent solution to Maxwell's equations in the presence of a strong rotating magnetic field, and has the form [184]

$$\rho_{\mathrm{GJ}} = \frac{2\vec{\Omega}_{\mathrm{NS}} \cdot \vec{B}_{\mathrm{NS}}}{1 - \Omega_{\mathrm{NS}}^2 r^2 \sin^2\theta} \,. \tag{1.22}$$

Here r and θ are standard spherical coordinates, Ω_{NS} is the neutron star rotational frequency, and B_{NS} is the neutron star magnetic field. We note that the denominator in Equation 1.22 contains a relativistic correction that only becomes significant either near the light cylinder or when considering millisecond pulsars. In this dissertation, we mainly study the vicinity of slower-rotating neutron stars, where this correction is instead close to unity.

The force-free GJ configuration serves as an excellent baseline, but we know that it cannot hold true everywhere in the magnetosphere. If this were the case, particle acceleration would be impossible, which is a necessary ingredient to account for phenomena such as the observed electromagnetic radiation emanating from pulsars. A force-free magnetosphere is furthermore inherently unstable, as there is no mechanism

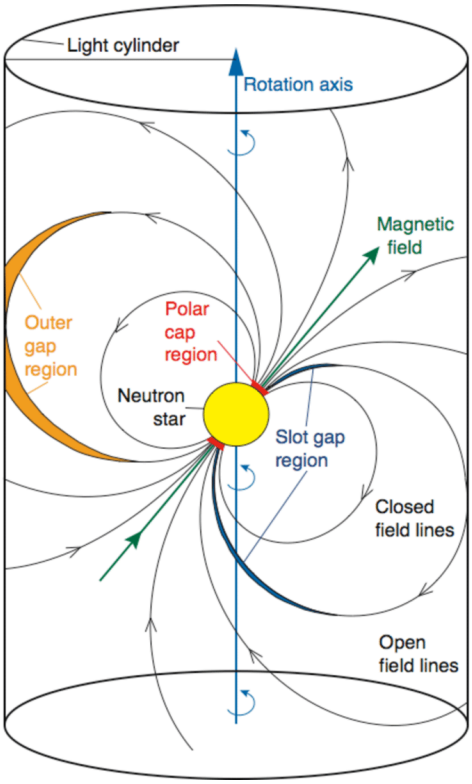

Figure 1.6: A simplified two-dimensional picture of the neutron star magnetosphere, adapted from [185]. The depicted neutron star is an oblique rotator, meaning that its magnetic and rotation axis are not aligned – the angle between these two axes is known as the misalignment angle. The image shows the location of the polar-cap gaps, along with the location of two other potential vacuum gap regions (the slot gap and outer gap). The light cylinder is positioned at $R_{\rm LC} = 1/\Omega_{\rm NS}$, which for the neutron stars considered in this thesis corresponds to roughly $10^3 - 10^5$ star radii.

capable of sustaining plasma production. For these reasons, it is anticipated that there exist localized areas within the magnetosphere that allow for non-zero $\vec{E} \cdot \vec{B}$ fields, thereby enabling particle acceleration and production [186, 187]. These areas are characterized by a lower density of plasma, and therefore referred to as vacuum gaps. There are several locations where vacuum gaps can potentially arise; we focus on the most well-studied one, which is above the two magnetic poles of the neutron star [187]. To understand the appearance of these so-called polar-cap gaps, it is important to note that the light cylinder effectively splits the magnetosphere into two zones, differentiated by featuring either open or closed magnetic field lines (see

Figure 1.6). The open zones are anchored to the magnetic poles, and these regions may consequently experience a reduction in density as plasma exits the magnetosphere along open magnetic field lines. As long as the current supplied by these charges is furthermore insufficient to screen the parallel electric field, polar-cap gaps form as a result (see *e.g.* [7]). Importantly, this is a cyclic process. Recent numerical simulations have established that once the polar-cap gaps are depleted, the resulting unscreened parallel electric field will initiate a pair production cascade that refills the gap with plasma [188–191]. The newly created plasma proceeds to once again screen the electric field, shutting down pair production and allowing the gap to empty, thereby restarting the cycle. This procedure offers a continuous means of generating plasma, and is moreover a proposed explanation for the origin of pulsar emission [189, 192]. Additionally, and of particular relevance to us, it results in a time-oscillating $\vec{E} \cdot \vec{B}$ field that is capable of creating axions.

A schematic of the neutron star magnetosphere is provided in Figure 1.6. Throughout this thesis, we adopt the GJ model as our benchmark, and additionally consider the existence of vacuum gaps allowing for axion production in Chapters 3 and 4. We furthermore approximate the magnetosphere as comprising only electron-positron pair plasma, and the magnetic field as being dipolar. It is worth mentioning that there is a possibility of realistic magnetospheres having features beyond this model, including:

- Areas with larger charge multiplicities, defined by $M \equiv n_e/n_{\rm GJ}$, where n_e and $n_{\rm GJ}$ denote the electron and GJ number density, respectively.
- The presence of ions. Their qualitative impact, however, remains negligible as long as their number density is less than one part per thousand.
- Non-dipolar magnetic field configurations. Specifically, the assumption of a dipolar magnetic field may not be valid close to the neutron star surface, where higher order multipoles can contribute, and outside of the light cylinder.

The extent to which these features occur is currently unknown. They are therefore not taken into account in the work presented here, with the intention of addressing them in the future.

1.5.2 Resonant axion-photon mixing and axion sourcing in vacuum gaps

Any particle that has a two-photon vertex may be generated by a photon moving inside of an external electromagnetic field. The external field in this case supplies a virtual photon to facilitate the interaction. The reverse process, whereby the particle in question instead creates a photon, is also a possibility. In the context of axions, this mechanism is referred to as the Primakoff effect [193], see Figure 1.7. As we will demonstrate explicitly in Chapter 5, axions and photons actually undergo mixing

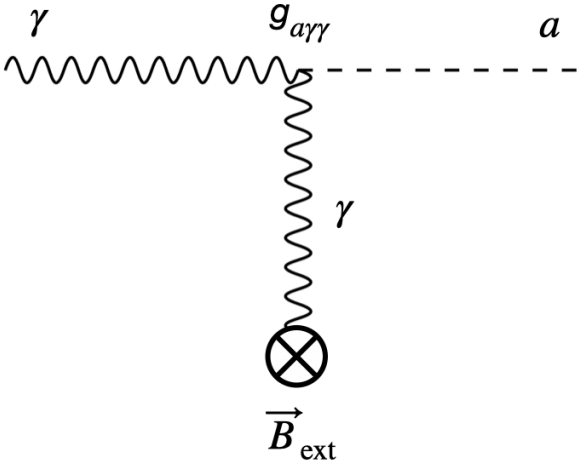

Figure 1.7: Feynman diagram representing the Primakoff effect for axions. The external field is chosen to be magnetic.

rather than straightforward conversion, analogous to how neutrinos exhibit oscillations from one type to another [194, 195]. We show in Chapter 5 that the equations governing this mixing process can be computed from the axion-modified Maxwell's equations. Importantly, axion-photon mixing can be particularly efficient when it takes place within a plasma like that found in the magnetosphere. Here, photons acquire an effective mass through plasma oscillations [196], allowing for the possibility of their four-momentum overlapping with axions. If this indeed occurs, the mixing process is enhanced, giving rise to what is known as resonant axion-photon mixing. We note that this operates on the same principle of aligning four-momenta as axion helioscopes. As previously mentioned, conversion in the case of helioscopes is not resonant, however, as alignment is instead achieved as a result of axions being relativistic, which minimizes the relevance of the axion mass.

The concept of resonant axion-photon mixing in neutron star magnetospheres was introduced alongside the broader notion of axion-photon mixing itself, in an influential paper authored by Georg Raffelt and Leo Stodolsky [195]. They exclusively studied the case of ultra-relativistic axions, however, and ultimately concluded that these could not lead to strong resonant effects. Many years later, Maxim Pshirkov and Sergei Popov revisited the issue, instead focusing on what turned out to be a more interesting source: dark matter axions [197]. Their proposal can be explained in relatively simple terms. As a neutron star travels through its respective galaxy, axions from the dark matter halo will continuously fall toward the star, thereby traversing the magnetosphere. By virtue of being cold dark matter, these axions are at most semi-relativistic, and along their path they may encounter areas where resonant mixing

into photons is possible. In a fully general context, this happens when the momenta of the two particles coincide, *i.e.* $k_a \simeq k_\gamma$. In the case of non- or semi-relativistic axions, however, this condition simplifies to $m_a \simeq \omega_p$. Here, ω_p denotes the plasma frequency, which is the quantity that determines the effective photon mass. In a cold[16] plasma it is given by

$$\omega_p = \sqrt{\sum_i \frac{4\pi\alpha n_i}{m_i}}\,, \tag{1.23}$$

where α is the fine-structure constant and the sum is over plasma species i with number density n_i and mass m_i. For an electron-positron pair plasma, this reduces to $\omega_p = \sqrt{4\pi\alpha n_e/m_e}$. If one sets n_e equal to the GJ density, this reveals that typical axion masses able to resonantly convert in neutron star magnetospheres are on the order of $m_a \simeq 10^{-9} - 10^{-4}\,\text{eV}$. As outlined in Section 1.3, this is a well-motivated mass range for axion searches. The corresponding photons have frequencies between $\sim 1\,\text{MHz} - 100\,\text{GHz}$, and potential signals due to resonant axion-photon mixing will thus reside in the radio band. Moreover, energy and momentum are conserved during the conversion, meaning that all photons produced through non-relativistic resonant axion-photon mixing have an energy roughly equal to the axion mass. The above mechanism is thus expected to yield a monochromatic line centered at the axion mass, which Pshirkov and Popov showed is possibly observable.

A narrow spectral line of this nature would provide a striking signature for the existence of axion dark matter, and recent years have seen much effort devoted to its characterization and detection [198–208]. Initial calculations showed promise, but failed to take into account many of the uncertainties regarding, for example, the neutron star magnetosphere. It was therefore unclear whether these results were truly accurate. We discuss this further in Chapter 2, where we also provide a framework for dealing with a substantial amount of these uncertainties. This framework allowed for the first robust signal calculations within this search strategy, and has since been used to set leading constraints on the axion-photon coupling [209] (shown with the label 'Neutron stars' in Figure 1.5). It furthermore inspired several pieces of follow-up work [210, 211], including an investigation into gravitationally bound structures of dark matter axions called axion miniclusters (which we do not discuss further here, but present an interesting possibility regarding the distribution of axions) [212].

Rather than exploiting ambient axion dark matter, Reference [213] identified that axions can also be produced locally, in the vacuum gap regions of the neutron star magnetosphere. As we have examined more thoroughly in the previous subsection, vacuum gaps are defined by the presence of a non-zero, oscillating $\vec{E} \cdot \vec{B}$ field, which

[16] When treating resonant axion-photon mixing, we will consider the neutron star magnetosphere to be a cold (*i.e.* non-relativistic) plasma. This is expected to be valid for the closed field line regions that corotate with the star. The open field line regions may also contain relativistic plasma, but these only occupy a relatively small volume.

appears as a source term in the axion equation of motion (see Equation 1.21). Axions are in this case created through an interaction resembling the one depicted in Figure 1.7, but with the incoming photon also being virtual and supplied by the external electric field. Reference [213] showed that the density of axions produced via this mechanism can be enormous, often surpassing the dark matter density in the neutron star vicinity. These axions moreover have a rich phenomenology, with, for example, an interesting dichotomy of the population into relativistic axions that escape the neutron star and non-relativistic axions that are gravitationally captured to form bound states. We explore these possibilities in Chapters 3 and 4, respectively, where we also perform the first-ever computations of the various signatures arising from resonant mixing of axions produced in vacuum gaps. Just like the monochromatic line due to dark matter axions discussed above, these are anticipated to yield significant potential for axion detection.

1.6 Summary (thesis organization)

In the realm of physics, the term 'matter' is used to collectively refer to every physical object that we can observe, from the tiniest atoms to the most massive stars. The building blocks of matter are elementary (*i.e.* indivisible) particles, such as quarks and electrons, that interact through the four known fundamental forces: gravity, electromagnetism, the weak nuclear force, and the strong nuclear force. This network of particles and their interactions is encapsulated in a theory known as the Standard Model of particle physics, which is firmly supported by both theory and experimental evidence. However, despite its accomplishments, the Standard Model is not a complete theory, as it fails to explain several major phenomena.

One such phenomenon is the implication of missing mass in the Universe. According to abundant gravitational evidence, approximately 85% of the total matter content of the Universe is not directly observed. This unseen matter, termed 'dark matter', exerts gravitational force on visible matter and light but does not radiate or interact in any other way that we can currently detect. Though seemingly invisible, dark matter plays a crucial role in shaping the structure of the cosmos, significantly influencing the formation and behavior of galaxies and galaxy clusters. Understanding the nature of dark matter remains one of the foremost challenges in modern astrophysics and cosmology, with researchers employing many different approaches to unveil its elusive properties.

The majority of models postulate that dark matter is an undiscovered particle, characterized by being feebly interacting, produced in the early Universe, and long-lived. Among the numerous proposed candidates, axions stand out as especially well-motivated. These hypothetical particles naturally emerge in extensions of the Standard Model and, in addition to being an ideal dark matter candidate, can address

other significant outstanding issues in physics. For example, axions were initially proposed to resolve the so-called 'strong CP problem' plaguing the theory of strong interactions, which involves the question of why the strong force seems to unnecessarily preserve the combined symmetries of charge conjugation and parity.

Unfortunately, detecting axions, as with any dark matter candidate, is a formidable challenge. Despite this, recent years have witnessed a surge of interest in axion research. Physicists are increasingly developing and building Earth-based detectors aimed at measuring axions in the laboratory, as well as formulating new strategies to indirectly observe axions by analyzing their impact on astrophysical environments. A prime example in the latter category is the modern effort to study axion-photon interactions in the environment surrounding neutron stars, known as the magnetosphere, where the ambient plasma and powerful magnetic field can provide the necessary conditions to overcome the feeble axion interaction strength.

In this dissertation, entitled *Neutron stars as axion laboratories: Harnessing the power of the magnetosphere*, we aim to comprehend the intricate interplay between big, in the form of neutron stars, and small, in the form of axions. As mentioned above, axions have emerged as a leading candidate for beyond the Standard Model physics, and neutron stars with their accompanying magnetospheres present ideal settings for studying these elusive particles. Two mechanisms have been identified as particularly promising: resonant axion-photon mixing and axion production via vacuum gaps – these accordingly form the foundation on which this thesis is built. Resonant axion-photon mixing refers to the conversion of axions into photons, and vice versa, at points where the four-momenta of the two particles coincide. This is facilitated within the magnetosphere by the combination of plasma and a background magnetic field. Axion production via vacuum gaps refers to the creation of axions in localized regions of the magnetosphere, ideally suited for this purpose due to the presence of oscillating electromagnetic fields. The research presented here represents a step forward in understanding the phenomenology and impact of these mechanisms, and thereby in elucidating the combined system of axions and neutron star magnetospheres. Ultimately, I hope that this may also serve as progress toward the discovery of axions.

In **Chapter 2**, we introduce an end-to-end analysis framework capable of calculating the photon flux produced via resonant conversion of dark matter axions in neutron star magnetospheres. This framework employs an advanced auto-differentiation ray-tracing algorithm, enabling precise tracking of axions and photons as they travel. Crucially, we take into account the non-trivial nature of the magnetospheric plasma, which is paramount in determining both the photon trajectories and the resonant conversion points. Through this approach, we address critical uncertainties regarding the impact of the magnetosphere on signal computations, helping to considerably improve their reliability in the process. Our results indicate that, given realistic parameter choices, resonant mixing of dark matter axions generates a signal character-

ized by strong anisotropy, time dependence, and a larger line width than previously anticipated.

In **Chapter 3**, we shift our attention toward the analysis of axions sourced locally within neutron star vacuum gaps. We develop a pipeline that allows for the computation of initial axion spectra, the subsequent trajectory of these axions through the magnetosphere, their resonant mixing with photons, and the resulting radio emission. The first of these steps is performed in two ways, using either a self-developed semi-analytic model or particle-in-cell simulations from a prior study. These two methods demonstrate great agreement, underscoring the robustness of our approach to modeling uncertainties. The final radio signal is determined utilizing an adapted version of the ray-tracing algorithm introduced in Chapter 2. Concentrating solely on the relativistic axion population, we show that, contrary to dark matter axions, these generate a broadband flux when resonantly mixing. We finally compare our signal predictions to observations of 27 nearby pulsars, setting the strongest current constraints on the axion-photon coupling in the mass range $m_a \simeq 10^{-8} - 10^{-5}\,\mathrm{eV}$. These limits are labeled 'Pulsars' in Figure 1.5.

In **Chapter 4**, we explore the consequences of the large non-relativistic axion population that is also produced via vacuum gaps. Such axions are unable to escape a neutron star's gravitational pull, causing them to gradually accumulate around the star to form an 'axion cloud'. The existence of axion clouds requires no assumptions beyond an axion in the appropriate mass range that couples to electromagnetism, and therefore they are expected to be a common occurrence if the axion mass is roughly on the order of $m_a \simeq 10^{-9} - 10^{-4}\,\mathrm{eV}$, even when the axion-photon coupling is small. We study the formation, characteristics, and evolution of axion clouds, showing their capacity to achieve enormous densities over wide regions of parameter space. We furthermore analyze possible energy dissipation mechanisms, identifying resonant axion-photon mixing as the primary candidate for generating observable effects. This work represents the first investigation into this category of axion clouds, inaugurating a new interdisciplinary field that harbors great potential for axion research.

Finally, in **Chapter 5**, we diverge from our previous broad perspective and instead zoom in on axion-photon mixing specifically. While this process is relatively well-understood in idealized scenarios, it poses a challenging task to completely solve the axion-photon mixing equations in highly non-uniform backgrounds like the neutron star magnetosphere. There have been multiple theoretical efforts focused on this issue in recent years, but these have yielded differing results despite ostensibly relying on the same approximations. To address this discrepancy, we treat the axion-photon mixing problem with numerical methods, finding agreement with only one of the published theoretical solutions. This work, while not considering highly non-uniform backgrounds yet, serves as the starting point for a broader investigation into axion-photon mixing in complex astrophysical settings, which will play a crucial role in advancing the accuracy of signal calculations for indirect axion searches.