UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Fisica e Astronomia "Galileo Galilei"

Master's Degree in Physics

Final Dissertation

Kaluza–Klein Axions: From Theoretical Formulation to Solar Bounds through Data Analysis

Thesis supervisor

Prof. Francesco D'Eramo

Thesis co-supervisor

Dr. Stefano Piacentini

Candidate

Francesco Peroni

Academic Year 2025/2026

This thesis investigates constraints on axion-like particles (ALPs) and Kaluza–Klein axions through the study of the solar axion basin.

The work begins by reviewing the theoretical motivation for axions, namely the strong CP problem. Possible alternative solutions to this problem are discussed, together with several axion models. Existing bounds and phenomenological arguments constraining the ALP parameter space are then examined, with particular focus on the $(m_a, g_{a\gamma\gamma})$ plane.

The framework of extra-dimensional theories is subsequently introduced, along with the emergence of axions in this context, referred to as Kaluza–Klein (KK) axions. After this theoretical overview, the analysis focuses on axion production in the Sun. The three main production channels are reviewed, and the corresponding energy spectra and fluxes are derived as functions of the axion energy. The number density of gravitationally trapped axions in the solar gravitational potential is then estimated, allowing the associated decay rate to be computed. These results are applied both to a generic ALP model and to KK axions.

Finally, the data collected by LIME, a CYGNO collaboration experiment located at the Gran Sasso National Laboratories, are used to constrain ALP and KK axion models. A dedicated data analysis is performed, including a study of the efficiency of the selection cuts applied to the signal. From this analysis, a sensitivity projection is derived for both generic ALPs and Kaluza–Klein axions.

Contents

Chapter 1

The Strong CP Problem

In this chapter we introduce the strong CP problem. A simple way to state the problem is the following: why is the electric dipole moment of the neutron so small? A little more precisely, although the neutron is electrically neutral, it is composed of charged constituents: one up quark, with electric charge $+2/3$, and two down quarks, each with electric charge $-1/3$. At a qualitative level, one may estimate the electric dipole moment as

$$\vec{d} = \sum_{i \in \text{quarks}} q_i \vec{r}_i \tag{1.1}$$

where Eq. (1.1) should not be interpreted too literally, since quarks cannot be sharply localized inside the neutron. Nevertheless, this expression is useful for building some intuition about the expected size of the neutron electric dipole moment. Taking the characteristic neutron size to be $r_n \sim 1/m_\pi$, one finds

$$|\vec{d}| \simeq 10^{-13} \times \sqrt{1 - \cos\theta} \text{ e cm} \tag{1.2}$$

where θ is simply a geometric angle, as illustrated in Eq. (1.1). Therefore, on dimensional grounds, one would naturally expect a neutron EDM of order 10^{-13} e cm.

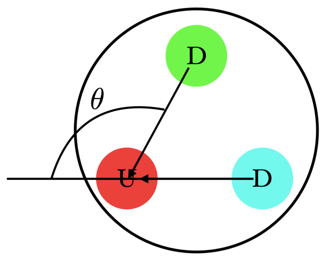

Figure 1.1: Toy model of the neutron, figure taken from [1]

Since the electric dipole moment is a vector, it must be aligned with the only intrinsic vector associated with the neutron, namely its spin. Experimentally, the neutron EDM can be probed by applying parallel electric and magnetic fields, which induce Larmor precession with frequency

$$\nu_\pm = 2\,|\mu B \pm dE|\,. \tag{1.3}$$

Here, ν is the Larmor frequency, E and B are the electric and magnetic fields, and d and μ are respectively the electric and magnetic dipole moments. After an evolution time t, the fields are switched off and the neutron spin state is measured, allowing one to determine ν_+. Repeating the same measurement with antiparallel electric and magnetic fields gives ν_-. The difference between these two frequencies can then be used to constrain the neutron EDM. The current best experimental bound from Refs. [2, 3] is:

$$|\vec{d}| < 10^{-26} \text{ e cm.} \tag{1.4}$$

From a look to Eq. (1.2) it is clear that this implies $\theta < 10^{-13}$. The strong CP problem can therefore be phrased as the question of why the parameter θ, which, as we will show, controls the CP violation in the strong interactions, is so small.

Of course, this qualitative argument must be replaced by a more quantitative and field-theoretic formulation. We begin by briefly reviewing Yang–Mills theories, which provide the general framework for most of the fundamental interactions in Nature: the Standard Model itself is a Yang–Mills theory with a specific matter content. We then focus on Quantum Chromodynamics, the Yang–Mills theory associated with the strong interactions and embedded in the Standard Model. QCD has several distinctive features, and it will be useful to discuss some equivalent ways of formulating this theory. Finally, after introducing the necessary technical ingredients, we analyze the vacuum structure of QCD and show how it gives rise to the strong CP problem in a more formal way.

1.1 Yang-Mills Theories

We first briefly introduce Yang–Mills theories and then turn to quantum chromodynamics (QCD). These topics are covered in essentially every textbook on quantum field theory; see, for example Ref. [4] and Ref. [5]. This introductory discussion is also useful for fixing the conventions that will be used later. The simplest example of a Yang–Mills theory is given by the well-known case of QED. One can derive the QED Lagrangian for a Dirac fermion by requiring local gauge invariance under

$$\psi(x) \to e^{i\alpha(x)}\psi(x)$$

and noticing that this requires the introduction of a new field A_μ in the derivative term, with transformation law

$$A_\mu \to A_\mu - \frac{1}{e}\partial_\mu\alpha$$

where e is the electric charge. It is also instructive to follow a different approach, discussed for example in Ref. [5], which makes the connection between gauge fields and connections more transparent. The reason why one needs to modify the ordinary derivative is that the derivative is formally defined as the limit

$$n^\mu\partial_\mu\psi(x) = \lim_{\epsilon\to 0}\frac{1}{\epsilon}\left[\psi(x+\epsilon n) - \psi(x)\right].$$

However, in the presence of a local phase symmetry, the two fields being subtracted transform independently, and therefore $\partial_\mu\psi$ does not have a simple transformation law. To solve this problem, one can define a scalar quantity $U(y,x)$ depending on two spacetime points x and y, with transformation law

$$U(y,x) \to e^{i\alpha(y)}U(y,x)e^{-i\alpha(x)}.$$

Moreover, as expected, we require that $U(x,x) = 1$. In this way $\psi(y)$ and $U(y,x)\psi(x)$ actually have the same transformation law, and we can build a sensible derivative:

$$n^\mu D_\mu\psi(x) = \lim_{\epsilon\to 0}\frac{1}{\epsilon}\left[\psi(x+\epsilon n) - U(x+\epsilon n, x)\psi(x)\right]$$

Of course in this limit we are actually evaluating $U(y,x)$ at infinitesimally separated points. Quite generally, one can choose $U(x,y)$ to be continuous in its arguments; it can therefore be expanded as:

$$U(x+\epsilon n, x) = 1 - ie n^\mu A_\mu(x) + \mathcal{O}(\epsilon^2)$$

Where we extracted an arbitrary constant e and A_μ is defined as $A_\mu = \frac{\partial U(y,x)}{\partial y^\mu}\Big|_{y=x}$. From this, one obtains the usual definition of the covariant derivative:

$$D_\mu\psi(x) = (\partial_\mu + ieA_\mu(x))\psi(x) \tag{1.5}$$

and sees that the gauge field A_μ can be interpreted as a connection.
In 1954, Yang and Mills Ref. [6] extended the gauge principle to non-Abelian symmetries. In particular, they proposed that the previous construction could be generalized from local phase rotations to

local transformations under a generic continuous symmetry group. The group we are interested in is the group of unitary operators under which the theory is required to be left invariant. Unitarity is needed in order to preserve probabilities or matrix elements.
We focus on groups that can be parameterized continuously by a finite set of parameters and so that contain elements arbitrarily close to the identity, in fact the neutral element belongs to a group by definition, and we are requiring continuity. This type of groups are called Lie groups and can be viewed as differential manifolds. A generic group element g can be written in the form:

$$g(\alpha) = 1 + i\alpha_a T^a$$

where T^a are the generators of the group. These must be Hermitian in order to ensure unitarity. Since we are interested in groups admitting finite-dimensional unitary representations, we restrict our attention to compact Lie groups, that is, groups that are compact finite-dimensional manifolds. The commutator of two generators must itself be a linear combination of generators, since the T^a span the space of infinitesimal group transformations. One therefore has

$$\left[T^a, T^b\right] = if^{abc}T^c,$$

where the coefficients f^{abc} are called the structure constants and define the Lie algebra. If there exists a generator T^a that commutes with all the other generators, then it generates an independent continuous Abelian group $U(1)$.
If no such generator exists, the group is called semisimple. Furthermore, if the generators cannot be split into two mutually commuting subsets, then the algebra (and the corresponding group) is said to be simple. An important result is that a compact Lie group can be decomposed into a direct product of non-Abelian simple groups and Abelian factors. For simplicity, we therefore focus on a simple group. There are not many types of simple groups, in fact only $SU(N)$, $SO(N)$, $Sp(N)$ and the exceptional groups can be groups at the same time simple, compact and with a unitary representation. We focus on $SU(N)$ for some integer N, the set of unitary matrices $N \times N$ satisfying $\det U = 1$, The number of generators is easily shown to be $N^2 - 1$.
Let us now consider an n-plet of Dirac fermions transforming in some representation of a simple group. The transformation law is

$$\psi(x) \to e^{i\alpha^a(x)t_a}\psi(x) = V(x)\psi(x),$$

where t_a are the generators of the Lie algebra in the given n-dimensional representation. In this case, $V(x)$ is an arbitrary element of the corresponding representation of the gauge group, that is, as discussed above, an arbitrary special unitary matrix. Once again, one needs to define a covariant derivative. Proceeding through the introduction of $U(y,x)$, this object is now promoted to an $n \times n$ unitary matrix, with transformation law

$$U(y,x) \to V(y)\, U(y,x)\, V(x)^\dagger.$$

For infinitesimally separated points, $y = x + \epsilon n$, the most general expansion of a unitary matrix is

$$U(x + \epsilon n, x) = 1 + ig\epsilon n^\mu A^a_\mu t_a + \mathcal{O}(\epsilon^2). \tag{1.6}$$

We immediately understand that the number of gauge fields must be equal to the number of generators of the group. From this last expression we easily recover the general form of the covariant derivative:

$$D_\mu = \partial_\mu - igA^a_\mu t^a \tag{1.7}$$

one can use Eq. (1.1) applied to Eq. (1.6) to find the transformation law for A^a_μ:

$$1 + ig\epsilon n^\mu A^a_\mu t^a \to V(x+\epsilon n)\left(1 + ig\epsilon n^\mu A^a_\mu t^a\right)V^\dagger(x)$$

And using:

$$V(x+\epsilon n)V^\dagger(x) = \Big[\left(1 + \epsilon n^\mu \frac{\partial}{\partial x^\mu}\right)V(x)\Big]V^\dagger(x) = 1 - \epsilon n^\mu V(x)\frac{\partial}{\partial x^\mu}V^\dagger(x)$$

One finds:

$$A^a_\mu t^a \to V(x)\left(A^a_\mu t^a + \frac{i}{g}\partial_\mu\right)V^\dagger(x) \tag{1.8}$$

Or, for infinitesimal transformations:

$$A^a_\mu \to A^a_\mu + \frac{1}{g}\partial_\mu \alpha^a + f^{abc}A^b_\mu \alpha^c$$

Of course this implies that the covariant derivative has the same transformation law of ψ, and therefore:

$$\big[D_\mu, D_\nu\big]\psi(x) \to V(x)\big[D_\mu, D_\nu\big]\psi(x)$$

It is useful to define also the field-strength tensor $F^a_{\mu\nu}$ through the relation:

$$\big[D_\mu, D_\nu\big] = -igF^a_{\mu\nu}t^a, \qquad F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + gf^{abc}A^b_\mu A^c_\nu \tag{1.9}$$

1.2 Quantum Chromodynamics

Quantum Chromodynamics (QCD) emerged from the need to provide a coherent quantum-field-theoretic description of the strong interaction. In the 1950s and early 1960s, a rapidly growing “zoo” of strongly interacting particles was discovered, first in cosmic-ray experiments and later in accelerator data. These particles, collectively known as hadrons, displayed a rich spectroscopy but resisted any simple unified interpretation. Unlike other successful quantum field theories, where the fields appearing in the Lagrangian could be more directly associated with the observed particle spectrum, the large number of hadronic states suggested that hadrons might not be elementary.

An essential step toward the modern picture was the development of the SU(3) flavor classification known as the Eightfold Way, introduced independently by Gell-Mann and Ne'eman in 1961, see Refs. [7, 8]. This scheme provided a remarkably successful organization of mesons and baryons into multiplets and led to the prediction of the Ω^- baryon. A further decisive advance came in 1964, when Gell-Mann and, independently, Zweig proposed that hadrons are composite objects built from more fundamental constituents, later called quarks in Refs. [9, 10].

Despite its phenomenological success, the quark model initially faced a serious conceptual difficulty. Some baryons appeared to require three quarks of the same flavor in a spatially symmetric ground state, implying a completely symmetric overall wavefunction for identical spin-$\frac{1}{2}$ fermions. This would be incompatible with the Pauli exclusion principle unless quarks possessed an additional internal degree of freedom Ref. [11]. In 1965, several authors proposed ways to resolve this problem by introducing a new hidden three-valued internal quantum number in Refs. [12, 13]. Gell-Mann later coined the term *color* for this new degree of freedom, which is one of the fundamental ingredients of QCD.

Nevertheless, a clear explanation of why quarks should be confined was still missing. Important progress came in 1970 with Khriplovich in Ref. [14] and in 1971 with 't Hooft in Ref. [15], who computed the one-loop beta function of Yang–Mills theories and found that, under suitable conditions, it is negative. In fact, in four spacetime dimensions asymptotic freedom occurs only in non-Abelian gauge theories. For a generic non-Abelian gauge theory, they obtained

$$\beta \equiv \frac{d\ln g}{d\ln\mu} = -\frac{g^2}{(4\pi)^2}\left(\frac{11}{3}C_2(G) - \frac{4}{3}n_f C(r)\right). \tag{1.10}$$

Here n_f is the number of fermion species in the representation r, and $C(r)$ is defined through the orthogonality relation

$$\mathrm{tr}[t_a t_b] = C(r)\,\delta_{ab}\,,$$

which can always be achieved by an appropriate choice of basis for the generators. For groups such as $SU(N)$, one often adopts the conventional normalization $C(r) = \frac{1}{2}$ in the fundamental representation. The coefficient $C_2(G)$ is the quadratic Casimir in the adjoint representation; more generally, for a representation r the quadratic Casimir $C_2(r)$ is defined by

$$t^a_r t^a_r = C_2(r)\cdot\mathbf{1}\,.$$

One can show that for the fundamental representation of $SU(N)$, $C_2(N) = \frac{N^2-1}{2N}$, while for the adjoint representation $C_2(G) = N$.
Building on these results, in 1973 Gross and Wilczek in Ref. [16] and, independently, Politzer in Ref. [17] identified the phenomenon of asymptotic freedom and clarified its physical significance: at short distances the interaction becomes weak, whereas at long distances it becomes strong.
In the same year, Fritzsch, Gell-Mann, and Leutwyler in Ref. [18] proposed a simple quantum field theory incorporating all the observed key ingredients. Introducing an $SU(3)_c$ gauge group and quarks u, d, and s transforming as color triplets, they obtained a model with an approximate global $SU(3)$ flavor symmetry that could reproduce the hadron spectrum. At the same time, the strong-coupling behavior of $SU(3)_c$ at long distances provides a natural mechanism to confine quarks and gluons into color-neutral bound states—namely, the observed hadrons. In this case the beta function reduces to

$$\beta(g) = -\frac{g^3}{(4\pi)^2}\left(\frac{11}{3}N - \frac{2}{3}n_f\right) < 0\,. \tag{1.11}$$

The associated Lagrangian is:

$$\mathcal{L} = i\bar{q}_{Li}\gamma^\mu D_\mu q_{Li} + i\bar{q}_{Ri}\gamma^\mu D_\mu q_{Ri} - \bar{q}_{Li}M_{ij}q_{Rj} - \bar{q}_{Rj}M^*_{ij}q_{Li} - \frac{1}{4}G^A_{\mu\nu}G^{\mu\nu,A} + \theta\frac{g^2}{64\pi^2}G^A_{\mu\nu}G^A_{\alpha\beta}\epsilon^{\mu\nu\alpha\beta} \tag{1.12}$$

Where we have divided left and right-handed fermions, D_μ is the covariant derivative, defined as usual:

$$(D_\mu q)_{ai} = (\delta_{ab}\partial_\mu - igA^A_\mu T^A_{ab})q_{bi}$$

The i index is in the flavor space, the a, b are color indices, and A is the index for the 8 gauge fields (gluons). We wrote the most general Lagrangian with all the ingredients said above, not assuming any other symmetry. Usually when writing the QED Lagrangian it is done assuming symmetry under parity, eliminating terms such as $c\epsilon F\tilde{F}$, (that however as we will see in the next sections are a total derivative) in this case we cannot work under this assumption, and therefore we must include also the last term in Eq. (1.12).
This is the term that opens the strong CP problem. Before really moving to the strong CP problem we present the symmetries of Eq. (1.12). Besides Poincaré and gauge symmetry, we have also some symmetries in flavor space: in fact, in absence of masses we have symmetry under independent flavor rotations of the left and right handed fermions:

$$q_{Li} \to U_{Lij}q_{Lj}, \qquad q_{Ri} \to U_{Rij}q_{Rj},$$

Where $U_{L,R}$ are $N_f \times N_f$ unitary matrices. This regime is actually particularly important since as we will see u, d and s quarks have masses well below Λ_χ, the scale at which the perturbative regime is lost, which is at about ~ 200 MeV.
The symmetry in this case can act independently for left and right chiralities, since they are not mixed by the kinetic term. We therefore call "chiral", the symmetry group $U(N_f)_L \times U(N_f)_R$, or equivalently $SU(N_f)_L \times SU(N_f)_R \times U(1)_B \times U(1)_A$ the last two abelian groups are given by the transformation relations:

$$\begin{pmatrix} q_L \\ q_R \end{pmatrix} \xrightarrow{U(1)_B} \begin{pmatrix} e^{i\alpha_B}q_L \\ e^{i\alpha_B}q_R \end{pmatrix} \qquad\qquad \begin{pmatrix} q_L \\ q_R \end{pmatrix} \xrightarrow{U(1)_A} \begin{pmatrix} e^{-i\alpha_A}q_L \\ e^{i\alpha_A}q_R \end{pmatrix}$$

However, even in the massless quark limit, the last symmetry $U(1)_A$ is anomalous as we will show.
Let us mention that if, instead, we assume $M_{ij} = m_i\delta_{ij}$ then the symmetry group becomes $SU(N_f)_V \times U(1)_B$, and in this case we have to rotate left-handed fields in the same way as right-handed fields. This regime is important after the formation of quark condensates, where as we will see, the quark bilinear takes a vacuum expectation value.
Finally if after a bi-unitary transformation we arrive at a generic diagonal form for the mass matrix $M_{ij} = m_i\delta_{ij}$ the symmetry in flavor space is given by:

$$q_{Li} \to e^{i\alpha_i}q_{Li} \qquad q_{Ri} \to e^{i\alpha_i}q_{Ri}$$

That is the symmetry given by the group $U(1)_{V_1} \times \cdots U(1)_{V_{N_f}}$.

1.2.1 The Chiral Lagrangian

We now present an alternative way to describe QCD. In fact below a certain scale, Λ_χ the relevant degrees of freedom are not those in Eq. (1.12), in fact QCD becomes non-perturbative, and we should use other degrees of freedom. Let us start by considering at energy scales slightly below Λ_χ, where most of the fields can be safely integrated out, with the exception of only the u, d, and s quarks. Moreover, since $m_{u,d} \ll m_s \lesssim \Lambda_\chi$, the quark mass matrix can be treated as a perturbation, and we may first work in the massless limit.
As discussed before, in this limit the theory has the full chiral symmetry

$$SU(N_f)_L \times SU(N_f)_R \times U(1)_B \times U(1)_A,$$

As we had lowered the energy, perturbation theory breaks down, and a quark bilinear condensate proportional to the identity in flavor space is formed:

$$\langle \bar{q}_{Rj} q_{Li} \rangle = \langle \Phi_{ij} \rangle \propto \delta_{ij} \tag{1.13}$$

This condensate (and in particular its vacuum expectation value) is the order parameter responsible for the breaking of the full chiral symmetry, assuming that the bilinear forms well above the scale set by m_s. It is quite easy to construct a toy model in which such a vacuum expectation value is generated dynamically, for instance the Nambu–Jona-Lasinio model in Ref. [19]. The operator in Eq. (1.13) is the first gauge-invariant term which breaks chiral symmetry. To see explicitly how Eq. (1.13) transforms under chiral rotations, we write

$$\Phi_{ij} = \bar{q}_{Rj} q_{Li} \to L_{ii'} R^*_{jj'} \bar{q}_{Rj'} q_{Li'} = \left(L \Phi R^\dagger \right)_{ij}.$$

Thus the chiral symmetry is spontaneously broken according to

$$SU(N_f)_L \times SU(N_f)_R \times U(1)_B \times U(1)_A \to SU(N_f)_V \times U(1)_B.$$

By Goldstone's theorem, we expect one Nambu–Goldstone boson for each broken generator. Therefore, in this case, we would expect $8+1$ Nambu–Goldstone bosons. Since real-world QCD contains massive quarks, these states are only pseudo-Nambu–Goldstone bosons, because they acquire small masses. However, experimentally one observes only 8 light pseudoscalar particles, and not 9. This is known as the $U(1)$ problem, which we will discuss later. We anticipate that its resolution is related to the fact that the Noether current associated with the $U_A(1)$ symmetry is anomalous, as we will show. In particular this symmetry is already strongly broken when the quark condensate forms, and it should not be regarded as an exact symmetry of the theory. As mentioned above, QCD in the infrared is strongly coupled, and therefore the spectrum of physical states cannot be determined perturbatively. Instead, hadrons are formed, and the relevant degrees of freedom are hadronic rather than quark fields. We must therefore construct an equivalent QFT description in terms of hadrons. This is usually called a dual description. Since the number of hadronic states is extremely large, such a description is useful only in the low-energy regime, where the heavy hadrons can be integrated out and only the Nambu–Goldstone bosons remain dynamical.
We can describe the Nambu–Goldstone bosons using the Callan–Coleman–Wess–Zumino formalism. Starting from a group G spontaneously broken into a subgroup H, the commutation relations can be written as

$$\begin{aligned} \left[T^a, T^b \right] &= i f^{abc} T^c \\ \left[T^a, X^{\hat{b}} \right] &= i f^{a\hat{b}\hat{c}} X^{\hat{c}} \\ \left[X^{\hat{a}}, X^{\hat{b}} \right] &= i f^{\hat{a}\hat{b}c} T^c + i f^{\hat{a}\hat{b}\hat{c}} X^{\hat{c}} \end{aligned} \tag{1.14}$$

where T^a are the unbroken generators of G, while $X^{\hat{a}}$ are the broken generators. The first relation in Eq. (1.14) follows from the fact that H is a group. The second relation follows from the antisymmetry of the structure constants, a property which can always be guaranteed for compact groups, while the third relation is completely general. We define the Nambu–Goldstone matrix as an element $U \in G/H$:

$$U(\Pi) = e^{i\Pi_{\hat{a}} X^{\hat{a}}}, \tag{1.15}$$

where $\Pi_{\hat{a}}$ denotes the set of Nambu–Goldstone bosons, each associated with a spontaneously broken generator $X^{\hat{a}}$. One can then derive the transformation rule for $U(\Pi)$. Since $U(\Pi)$ is an element of G, also $gU(\Pi)$ belongs to G. Therefore, one can write

$$U(\Pi) \to gU(\Pi) = U(\Pi')h,$$

where $h \in H$. It can be shown that this decomposition is unique. Therefore,

$$U(\Pi) \to U(\Pi') = gU(\Pi)h^\dagger. \tag{1.16}$$

We are interested in the particular case of symmetric cosets, for which

$$\left[X^{\hat{a}}, X^{\hat{b}}\right] = if^{\hat{a}\hat{b}c}T^c. \tag{1.17}$$

In this case, the algebra Eq. (1.14) is invariant under the automorphism $\mathcal{P}(T) = T$, $\mathcal{P}(X) = -X$. Starting from $gU = U'h$ and acting with $\mathcal{P}$ on both sides, we obtain

$$\mathcal{P}(g)U^\dagger = U'^\dagger h.$$

Therefore,

$$U'^\dagger = \mathcal{P}(g)U^\dagger h^\dagger \qquad \Longrightarrow \qquad U' = hU\mathcal{P}(g)^\dagger.$$

Combining this relation with $U' = gUh^\dagger$, one can construct a new object with a simple transformation law:

$$\Sigma = UU \to g\Sigma\mathcal{P}(g)^\dagger. \tag{1.18}$$

This is precisely the case relevant for the spontaneous symmetry breaking pattern

$$SU(N)_R \times SU(N)_L \to SU(N)_V.$$

Indeed,

$$\left[T_L^a, T_L^b\right] = if^{abc}T_L^c, \qquad \left[T_R^a, T_R^b\right] = if^{abc}T_R^c, \qquad \left[T_L^a, T_R^b\right] = 0. \tag{1.19}$$

Defining

$$T^a = T_L^a + T_R^a, \qquad X^a = T_L^a - T_R^a,$$

we find

$$\left[T^a, T^b\right] = if^{abc}T^c, \qquad \left[X^a, X^b\right] = if^{abc}T^c, \qquad \left[T^a, X^b\right] = if^{abc}X^c. \tag{1.20}$$

A representation of the generators can be easily constructed. Since T_L and T_R commute with each other, they act on orthogonal spaces:

$$T_L^a = \begin{pmatrix} T^a & 0 \\ 0 & 0 \end{pmatrix}, \qquad T_R^a = \begin{pmatrix} 0 & 0 \\ 0 & T^a \end{pmatrix}, \tag{1.21}$$

where T^a is a generator of $SU(N)$. A generic element of the group G is then given by

$$g = \begin{pmatrix} L & 0 \\ 0 & R \end{pmatrix}, \tag{1.22}$$

while an element h of the unbroken subgroup H is

$$h = \begin{pmatrix} V & 0 \\ 0 & V \end{pmatrix}. \tag{1.23}$$

The Nambu–Goldstone matrix is therefore

$$U = e^{i\Pi^a X_a} = \begin{pmatrix} \xi & 0 \\ 0 & \xi^\dagger \end{pmatrix}. \tag{1.24}$$

We can finally define the matrix Σ, given by $\Sigma = \xi\xi$. From the transformation law $U \to gUh^\dagger$, one obtains $\xi \to L\xi V^\dagger$, and therefore

$$\Sigma \to L\Sigma R^\dagger.$$

Going back to the chiral Lagrangian, we define

$$\Sigma = e^{i\lambda_a \Pi_a/f}. \tag{1.25}$$

Here λ_a are the Gell-Mann matrices. The matrix $\Pi = \lambda_a\Pi_a$ is traceless and Hermitian, and has mass dimension one, while f is the pion decay constant. The lowest-order term that we can write that is invariant under chiral symmetry is

$$\mathcal{L}_\chi = \frac{f^2}{4}\mathrm{Tr}\left[\partial_\mu \Sigma \partial^\mu \Sigma^\dagger\right]. \tag{1.26}$$

With this normalization, the Nambu–Goldstone bosons have canonically normalized kinetic terms. There are different ways to establish the mapping between hadrons and quark bilinears. One possibility is to use the Noether currents associated with $SU(3)_L$ and $SU(3)_R$. Starting from Eq. (1.12), and using the standard definition of the Noether current,

$$S\left[e^{i\alpha_a T^a} q, G\right] - S\left[q, G\right] = -\int d^4x \partial_\mu \alpha_a J^{a\mu}, \tag{1.27}$$

we can write

$$\begin{aligned} J_L^{a\mu} =& \bar{q}_{Li}\gamma^\mu T^a_{ij} q_{Lj}, \\ J_R^{a\mu} =& \bar{q}_{Ri}\gamma^\mu T^a_{ij} q_{Rj}. \end{aligned} \tag{1.28}$$

On the other hand, in the chiral effective theory one finds

$$\begin{aligned} J_L^{a\mu} =& i\frac{f^2}{4}\mathrm{Tr}\left[\left(\partial^\mu\Sigma\Sigma^\dagger - \Sigma\partial^\mu\Sigma^\dagger\right)T^a\right] \\ =& -\frac{f}{2}\mathrm{Tr}\left[\partial^\mu \Pi T^a\right], \\ J_R^{a\mu} =& i\frac{f^2}{4}\mathrm{Tr}\left[\left(\partial^\mu\Sigma^\dagger\Sigma - \Sigma^\dagger\partial^\mu\Sigma\right)T^a\right] \\ =& +\frac{f}{2}\mathrm{Tr}\left[\partial^\mu \Pi T^a\right]. \end{aligned} \tag{1.29}$$

This means that the vector current $J_V^{a\mu}$ has no term linear in Π, while the axial current satisfies $J_A^{a\mu} = -f\partial^\mu\Pi$. By matching the Noether currents in the fundamental and effective descriptions, one finds, for example,

$$\begin{aligned} if\partial^\mu \pi_0 &= \frac{1}{\sqrt{2}}(\bar{u}\gamma^\mu\gamma_5 u - \bar{d}\gamma^\mu\gamma_5 d), \\ if\partial^\mu \pi_+ &= \bar{d}\gamma^\mu\gamma_5 u, \end{aligned} \tag{1.30}$$

and analogous relations for the other mesons. We can now include the effects of the quark mass matrix. To do so, we use the spurion method. The idea is to observe that, if the mass matrix M transformed as

$$M \to LMR^\dagger,$$

then Eq. (1.12) would remain formally invariant under the full chiral symmetry. We must therefore add a term to the chiral Lagrangian which remains invariant under the transformation law $\Sigma \to L\Sigma R^\dagger$, provided that M transforms as above. The leading term of this kind is

$$\mathcal{L}_M = \frac{Bf^2}{2}\mathrm{Tr}[\Sigma M^\dagger + \Sigma^\dagger M]. \tag{1.31}$$

We will use these considerations later to derive the axion mass in the framework of the Peccei–Quinn mechanism.

1.2.2 Anomalies

Before turning to the strong CP problem in detail, we need to introduce one last ingredient: anomalous symmetries. This will also be essential for understanding the resolution of the $U(1)$ problem.
There are several ways to derive anomalies. For example, by regularizing linearly divergent triangle Feynman diagrams, one finds that preserving gauge invariance necessarily breaks axial current conservation. Here we follow instead Fujikawa's approach, in which the anomaly arises from the non-invariance of the path-integral measure under a symmetry transformation.
Let us first recall the standard Ward identity. Suppose that a field transformation $\phi \to \phi'$ changes the action according to Eq. (1.27), and assume for the moment that the path-integral measure is invariant, $\mathcal{D}\phi' = \mathcal{D}\phi$. Then

$$\begin{aligned} Z &= \int \mathcal{D}\phi\, e^{iS[\phi]} = \int \mathcal{D}\phi'\, e^{iS[\phi']} \\ &= \int \mathcal{D}\phi\, \exp\left\{ iS[\phi] - i\int d^4x\, \partial_\mu\alpha\, J^\mu \right\} \\ &= \int \mathcal{D}\phi\, e^{iS[\phi]} \left[1 - i\int d^4x\, \partial_\mu\alpha\, J^\mu \right]. \end{aligned} \tag{1.32}$$

Where in the second equality we Taylor expanded the action. Integrating by parts and using the arbitrariness of $\alpha(x)$, one obtains the usual conservation law for the Noether current,

$$\partial_\mu J^\mu = 0.$$

However, if the same transformation changes the measure according to

$$\mathcal{D}\phi' = \mathcal{D}\phi\, \exp\left\{ i\int d^4x\, \alpha(x)\mathcal{A}(x) \right\},$$

then the same argument gives instead

$$\partial_\mu J^\mu = \mathcal{A}. \tag{1.33}$$

Thus a current associated with a classical symmetry of the Lagrangian need not be conserved in the quantum theory. In this case the symmetry is said to be anomalous, and Eq. (1.33) is the anomalous Ward identity.
Let us now compute the anomaly following Fujikawa's method. We consider a local chiral transformation of the fermion fields,

$$U(x) = e^{i\alpha^a(x)t_a\gamma_5}, \qquad t_a^\dagger = t_a, \qquad [\gamma^\mu, t_a] = 0. \tag{1.34}$$

The corresponding transformation of $\bar\psi$ involves

$$\bar U(x) = i\gamma^0 e^{-i\alpha^a(x)t_a\gamma_5} i\gamma^0 = e^{+i\alpha^a(x)t_a\gamma_5}(i\gamma^0)^2 = e^{i\alpha^a(x)t_a\gamma_5} = U(x). \tag{1.35}$$

therefore the fermionic path-integral measure transforms with the Jacobian

$$\bar{\mathcal{U}}(x) = \mathcal{U}(x) \Longrightarrow (\mathrm{Det}\,\mathcal{U})^{-1}\left(\mathrm{Det}\,\bar{\mathcal{U}}\right)^{-1} = (\mathrm{Det}\,\mathcal{U})^{-2}. \tag{1.36}$$

This factor can be different from 1 and we must compute it. We start from

$$\mathrm{Tr}\log\mathcal{U} = \int d^4x\, \langle x|\mathrm{tr}\log\mathcal{U}|x\rangle = \int d^4x\, \delta^{(4)}(0)\, i\epsilon^\alpha(x)\mathrm{tr}\left(t_\alpha\gamma_5\right), \tag{1.37}$$

where Tr denotes the functional trace, while tr denotes the trace over internal and Dirac indices. Hence

$$(\mathrm{Det}\,\mathcal{U})^{-2} = e^{-2\mathrm{Tr}\log\mathcal{U}} = e^{i\int d^4x\, \epsilon^\alpha(x)a_\alpha(x)},$$

with

$$a_\alpha(x) = -2\delta^{(4)}(0)\,\mathrm{tr}\left(t_\alpha\gamma_5\right).$$

This formal expression is ill-defined: the trace of γ_5 over Dirac indices vanishes, whereas $\delta^{(4)}(0)$ is divergent. Therefore a regulator is needed. Using

$$\delta^{(4)}(0) = \langle x|x\rangle = \int d^4p\, \langle x|p\rangle\langle p|x\rangle = \int \frac{d^4p}{(2\pi)^4} e^{ip(x-y)}\Big|_{x=y},$$

we regularize the trace by introducing a gauge-covariant cutoff. We note that crucially we choose a gauge invariant regularization, since we want to preserve this symmetry. This is very clear if one would have chosen the alternative route: trying to regularize triangular Feynman diagrams one finds that it is not possible to preserve both gauge and chiral symmetry. We compute

$$\int d^4x\, \epsilon^\alpha a_\alpha(x) = -2 \lim_{\Lambda\to\infty} \mathrm{Tr}\,\mathfrak{T}_\Lambda, \qquad \mathfrak{T}_\Lambda = \epsilon^\alpha(x)\gamma_5 t_\alpha f\left((i\gamma^\mu D_\mu/\Lambda)^2\right), \tag{1.38}$$

where $f(s)$ is a smooth function satisfying $f(0) = 1$, $f(\infty) = 0$, and $sf'(s) = 0$ at both $s = 0$ and $s = \infty$. We then have

$$\begin{aligned}
\mathrm{Tr}\,\mathfrak{T}_\Lambda &= \int d^4x\, \mathrm{tr}\left\langle x\left|\epsilon^\alpha(x)\gamma_5 t_\alpha f\left((i\gamma^\mu D_\mu/\Lambda)^2\right)\right|x\right\rangle \\
&= \int d^4x\, \epsilon^\alpha(x) \int d^4p\, \langle x|p\rangle\, \mathrm{tr}\left[\gamma_5 t_\alpha \left\langle p\left|f\left((i\gamma^\mu D_\mu/\Lambda)^2\right)\right|x\right\rangle\right] \\
&= \int d^4x\, \epsilon^\alpha(x) \int \frac{d^4p}{(2\pi)^4} e^{ipx} \mathrm{tr}\left[\gamma_5 t_\alpha f\left(-\frac{1}{\Lambda^2}\left[\gamma^\mu\left(\frac{\partial}{\partial x^\mu} - iA^{\mathcal{R}}_\mu\right)\right]^2\right)\right] e^{-ipx}.
\end{aligned} \tag{1.39}$$

Using

$$f\left(-\frac{1}{\Lambda^2}\left[\gamma^\mu\left(\frac{\partial}{\partial x^\mu} - iA^{\mathcal{R}}_\mu\right)\right]^2\right) e^{-ipx} = e^{-ipx} f\left(-\frac{1}{\Lambda^2}\left[\gamma^\mu\left(\frac{\partial}{\partial x^\mu} - ip_\mu - iA^{\mathcal{R}}_\mu\right)\right]^2\right), \tag{1.40}$$

we obtain

$$\mathrm{Tr}\,\mathfrak{T}_\Lambda = \int d^4x\, \epsilon^\alpha(x) \int \frac{d^4p}{(2\pi)^4} \mathrm{tr}\left[\gamma_5 t_\alpha f\left(-\frac{1}{\Lambda^2}[-i\not{p} + \not{D}]^2\right)\right].$$

Changing variables $p = \Lambda q$, this becomes

$$\mathrm{Tr}\,\mathfrak{T}_\Lambda = \int d^4x\, \epsilon^\alpha(x)\Lambda^4 \int \frac{d^4q}{(2\pi)^4} \mathrm{tr}\left[\gamma_5 t_\alpha f\left([-i\not{q} + \not{D}/\Lambda]^2\right)\right]. \tag{1.41}$$

We now expand in powers of $1/\Lambda$. Using the Clifford algebra,

$$\left(-i\not{q} + \frac{\not{D}}{\Lambda}\right)^2 = -q^2 - \frac{2iq^\mu D_\mu}{\Lambda} + \frac{\not{D}^2}{\Lambda^2}.$$

A non-vanishing trace with γ_5 in four dimensions requires exactly four gamma matrices. Moreover, a finite contribution in the limit $\Lambda \to \infty$ comes from the term quadratic in $\not{D}^2/\Lambda^2$. Therefore

$$\lim_{\Lambda\to\infty} \mathrm{Tr}\,\mathfrak{T}_\Lambda = \int d^4x\, \epsilon^\alpha(x) \int \frac{d^4q}{(2\pi)^4} \frac{1}{2} f''(q^2)\mathrm{tr}\left[\gamma_5 t_\alpha (\not{D}^2)^2\right].$$

After a Wick rotation $q_0 \to -iq_0$, and passing to spherical coordinates in Euclidean momentum space, one finds

$$\int \frac{d^4q}{(2\pi)^4} \frac{1}{2} f''(q^2) = \frac{i}{32\pi^4}\mathrm{vol}(S^3) \int_0^{+\infty} dq\, q^3 f''(q^2).$$

Since $\mathrm{vol}(S^3) = 2\pi^2$, and using the change of variables $q^2 = \xi$, we get

$$\int \frac{d^4q}{(2\pi)^4} \frac{1}{2} f''(q^2) = \frac{i}{32\pi^2}\left[\xi f'(\xi)\Big|_0^\infty - \int_0^\infty d\xi\, f'(\xi)\right] = \frac{i}{32\pi^2},$$

where the boundary conditions imposed on f have been used. We now use

$$\not{D}^2 = \left(\frac{1}{2}\{\gamma^\mu,\gamma^\nu\} + \frac{1}{2}[\gamma^\mu,\gamma^\nu]\right) D_\mu D_\nu = D^\mu D_\mu - \frac{i}{4}e[\gamma^\mu,\gamma^\nu]F_{\mu\nu}.$$

Here $\{\gamma^\mu,\gamma^\nu\} = 2g^{\mu\nu}$ and $[D_\mu, D_\nu] = -iF_{\mu\nu}$. Since only the antisymmetric part contributes to the trace with γ_5, and using

$$\mathrm{tr}_D\left(\gamma_5\gamma^\mu\gamma^\nu\gamma^\sigma\gamma^\rho\right) = 4i\epsilon^{\mu\nu\sigma\rho},$$

one obtains

$$\mathrm{tr}\left[\gamma_5 t_\alpha(\not{D}^2)^2\right] = \left(\frac{i}{4}\right)^2 e^2 \mathrm{tr}_D\left[\gamma_5[\gamma^\mu,\gamma^\nu][\gamma^\sigma,\gamma^\rho]\right]\mathrm{tr}_{\mathcal{R}}\left[t_\alpha F_{\mu\nu}F_{\sigma\rho}\right] = -ie^2\epsilon^{\mu\nu\rho\sigma}\mathrm{tr}_{\mathcal{R}}\left[t_\alpha F_{\mu\nu}F_{\sigma\rho}\right]. \quad (1.42)$$

therefore,

$$\lim_{\Lambda\to\infty}\mathrm{Tr}\,\mathcal{T}_\Lambda = \frac{e^2}{32\pi^2}\int d^4x\,\epsilon^\alpha(x)\epsilon^{\mu\nu\rho\sigma}\mathrm{tr}_{\mathcal{R}}\left[t_\alpha F_{\mu\nu}F_{\sigma\rho}\right],$$

and the anomaly is

$$\mathcal{A}_\alpha(x) = -\frac{e^2}{16\pi^2}\epsilon^{\mu\nu\rho\sigma}\mathrm{tr}_{\mathcal{R}}\left[t_\alpha F_{\mu\nu}(x)F_{\sigma\rho}(x)\right].$$

The trace depends on the representation $\mathcal{R}$ of the gauge group. We can now apply this result to QCD. In the non-Abelian case the field strength carries gauge indices, and the corresponding expression becomes

$$\mathcal{A}_\alpha(x) = -\frac{g^2}{16\pi^2}\epsilon^{\mu\nu\rho\sigma}F^a_{\mu\nu}(x)F^b_{\sigma\rho}(x)\mathrm{tr}_{\mathcal{R}}\left[t_\alpha\mathcal{T}_a\mathcal{T}_b\right], \quad (1.43)$$

where $\mathcal{T}_a$ are the generators of the gauge group in the representation $\mathcal{R}$. For $SU(3)$ in the fundamental representation one has $\mathcal{T}_a = \lambda_a/2$, with λ_a the Gell-Mann matrices.
In QCD, in addition to gauge transformations, one can consider flavor transformations. In particular,

$$\psi \to e^{i\alpha_a T^a}\psi, \qquad \psi \to e^{i\gamma_5\alpha_a T^a}\psi. \quad (1.44)$$

The first transformation is vector-like, while the second is axial. A computation analogous to the one above shows that the vector current is not anomalous, whereas the axial current satisfies

$$\partial_\mu\left(\bar{\psi}\gamma^\mu\gamma^5 T^a\psi\right) = -\frac{g^2}{32\pi^2}\mathrm{Tr}\left[T^a\{\mathcal{T}_A,\mathcal{T}_B\}\right]F^A_{\mu\nu}F^B_{\sigma\rho}\epsilon^{\mu\nu\rho\sigma}. \quad (1.45)$$

Here the trace is over both flavor and color indices. Explicitly, for the flavor generators we take

$$\begin{aligned} U_A(1): &\qquad T^0 = \delta_{ij}\delta_{ab}, \\ \text{axial non-Abelian:} &\qquad T^a = \frac{1}{2}\lambda^a_{ij}\delta_{ab}, \qquad a = 1,\ldots,8, \end{aligned} \quad (1.46)$$

while the color (gauge) generators are

$$\mathcal{T}^A = \delta_{ij}\frac{1}{2}\lambda^A_{ab}, \qquad A = 1,\ldots,8. \quad (1.47)$$

We can now determine which axial symmetries are anomalous. Since flavor and color generators act on different spaces, they commute. For the non-Abelian axial generators,

$$\mathrm{Tr}\left[T^a\{\mathcal{T}^A,\mathcal{T}^B\}\right] = \mathrm{Tr}_{\text{flavour}}[T^a]\,\mathrm{Tr}_{\text{color}}\left[\{\mathcal{T}^A,\mathcal{T}^B\}\right] = \mathrm{Tr}_{\text{flavour}}[T^a]\,\delta^{AB} = 0. \quad (1.48)$$

In the last step we used the fact that the Gell-Mann matrices are traceless. therefore the non-Abelian axial symmetries are not anomalous with respect to QCD. On the other hand, the axial $U(1)_A$ generator is proportional to the identity in flavor space, and hence

$$\mathrm{Tr}\left[T^0\{\mathcal{T}^A,\mathcal{T}^B\}\right] = \mathrm{Tr}_{\text{flavour}}[T^0]\,\mathrm{Tr}_{\text{color}}\left[\{\mathcal{T}^A,\mathcal{T}^B\}\right] = N_f\,\delta^{AB}. \quad (1.49)$$

Thus Eq. (1.45) becomes

$$\partial_\mu \left(\bar{\psi}\gamma^\mu\gamma^5\psi\right) = -\frac{g^2}{32\pi^2} N_f \, \epsilon^{\mu\nu\alpha\beta} G^A_{\mu\nu} G^A_{\alpha\beta}. \tag{1.50}$$

This shows that the classical $U_A(1)$ symmetry is explicitly broken by quantum effects. Around the hadronic scale, where the strong coupling is large, this breaking is not negligible. As a consequence, one does not expect an additional light Goldstone boson associated with $U_A(1)$, in agreement with the observed meson spectrum.
Other gauge interactions, such as electromagnetic and weak interactions, may also contribute to anomalies of axial currents. In those cases, however, the strong coupling g is replaced by e or g_w, and the corresponding effects are much smaller at low energies.
The anomaly of the axial current has another important consequence, directly related to the strong CP problem. Consider the QCD path integral starting from Eq. (1.12), including the most general θ-term:

$$\begin{aligned} Z[M,\theta] &= \int \mathcal{D}[q,\bar{q},G]\, e^{i\int d^4x \left[\bar{q} i \not{D} q - (\bar{q}_L M q_R + \text{h.c.}) - \frac{1}{4} G^A G^A + \theta \frac{g^2}{64\pi^2} \epsilon G^A G^A\right]} \\ &= \int \mathcal{D}[q',\bar{q}',G]\, e^{i\int d^4x \left[\bar{q}' i \not{D} q' - (\bar{q}'_L M q'_R + \text{h.c.}) - \frac{1}{4} G^A G^A + \theta \frac{g^2}{64\pi^2} \epsilon G^A G^A\right]} \\ &= \int \mathcal{D}[q,\bar{q},G]\, e^{i\int d^4x \left[\bar{q} i \not{D} q - (\bar{q}_L e^{2i\alpha} M q_R + \text{h.c.}) - \frac{1}{4} G^A G^A + \theta \frac{g^2}{64\pi^2} \epsilon G^A G^A\right] - i\int d^4x\, \alpha \mathcal{A}} \\ &= \int \mathcal{D}[q,\bar{q},G]\, e^{i\int d^4x \left[\bar{q} i \not{D} q - (\bar{q}_L e^{2i\alpha} M q_R + \text{h.c.}) - \frac{1}{4} G^A G^A + (\theta + 2N_f\alpha) \frac{g^2}{64\pi^2} \epsilon G^A G^A\right]} \\ &= Z[M e^{2i\alpha}, \theta + 2N_f\alpha]. \end{aligned} \tag{1.51}$$

We therefore have the equivalence relation

$$(M,\theta) \sim (M e^{2i\alpha}, \theta + 2N_f\alpha).$$

If θ were not a physical parameter, the phase of the quark mass matrix would also be unphysical. Nevertheless it can be easily shown that this is not the case, so that θ must be treated as a physical parameter.
The equivalence relation in Eq. (1.51) shows that θ and the phase of the quark mass matrix are not separately physical. With the conventions used above, the invariant combination is

$$\bar{\theta} = \theta - \text{Arg}\left[\det M\right]. \tag{1.52}$$

Indeed, under the chiral $U_A(1)$ transformation,

$$\text{Arg}\left[\det M\right] \to \text{Arg}\left[\det M\right] + 2N_f\alpha, \qquad \theta \to \theta + 2N_f\alpha,$$

so that $\bar{\theta}$ is invariant. In the next section we will see how this parameter receives also another contribution related to the non-trivial topological structure of the QCD vacuum.

1.2.3 Back to the Chiral Lagrangian

Before moving on, let us return to the chiral Lagrangian formalism and include the effect of the anomalous $U_A(1)$ symmetry. From the previous discussion, we learned that an axial rotation of the quark fields,

$$q \to e^{i\gamma_5\alpha} q,$$

does not leave the quantum theory invariant, because the path-integral measure is anomalous. Equivalently, such a transformation shifts the coefficient of the topological operator $G_{\mu\nu}\tilde{G}^{\mu\nu}$. In the two-flavour theory, and with the conventions adopted above, this effect can be described by treating θ as a spurion. Namely, we assign to θ a transformation law such that the anomalous transformation is formally promoted to a symmetry:

$$u \to e^{i\gamma_5\alpha}u, \qquad d \to e^{i\gamma_5\alpha}d, \qquad \theta \to \theta - 2\alpha. \tag{1.53}$$

The precise numerical coefficient in the shift of θ depends on the normalization chosen for the axial rotation and for the operator $G\tilde{G}$. What is important is that the anomalous axial rotation can be compensated by a shift of θ.
There are important differences between treating θ as a spurion and treating the quark mass matrix M as a spurion. The mass matrix transforms multiplicatively under chiral transformations, while θ transforms non-linearly: it shifts. In order to write objects that transform linearly, one may instead use $e^{i\theta}$. Moreover, in the limit $M \to 0$ the ordinary chiral symmetry is restored, whereas there is no analogous limit for θ: the $U_A(1)$ symmetry is anomalous from the beginning, and setting $\theta = 0$ does not restore it as an exact quantum symmetry. This is also reflected in the fact that $|e^{i\theta}| = 1$ for any value of θ.
Because θ, a parameter of the theory, transforms under the anomalous symmetry, the singlet pseudoscalar field associated with $U_A(1)$, which we denote by η, acquires a mass even in the limit in which the quark masses vanish. In the chiral Lagrangian, where the degrees of freedom are bosonic and anomalies are encoded in effective terms, we can implement the anomalous symmetry through the formal transformation rules

$$\eta \to \eta + \sqrt{2}\alpha f_\pi, \qquad \theta \to \theta - 2\alpha, \qquad M \to M e^{-i\alpha}. \tag{1.54}$$

These transformation rules allow us to construct an effective Lagrangian that correctly reproduces the anomalous breaking of $U_A(1)$.
Again, we must build a Lagrangian that is invariant under Eq. (1.54). We therefore add to the chiral Lagrangian a term that depends on the determinant of the meson field,

$$\mathcal{L}_\chi = \frac{f_\pi^2}{4}\mathrm{Tr}\left[\partial_\mu\Sigma\partial^\mu\Sigma^\dagger\right] + \frac{Bf_\pi^2}{2}\mathrm{Tr}\left[\Sigma M^\dagger + \Sigma^\dagger M\right] + \frac{C}{2}f_\pi^4 \det U + \frac{C^*}{2}f_\pi^4 \det U^\dagger. \tag{1.55}$$

In order for this term to reproduce the θ-dependence of the anomalous theory, we take $C = |C|e^{i\theta}$. Expanding the determinant term around the singlet direction gives a mass term for η,

$$\mathcal{L}_\eta = \frac{1}{2}m_\eta^2\left(\eta - \frac{\theta f_\pi}{\sqrt{2}}\right)^2, \tag{1.56}$$

with $m_\eta^2 = 2|C|f_\pi^2$. Thus the η field becomes massive even in the chiral limit, as expected from the anomalous breaking of $U_A(1)$.
At energies below the η mass, we may integrate out the singlet field. Using its vacuum expectation value, the remaining $SU(2)$ matrix can be written schematically as

$$U = e^{i\theta/2}\exp\left(i\sigma_i\pi^i/f_\pi\right).$$

The mass term for the neutral pion sector then becomes

$$\mathcal{L}_{mass} = Bf_\pi^2\left[m_u\cos\left(\pi_0/f_\pi + \bar{\theta}/2\right) + m_d\cos\left(\pi_0/f_\pi - \bar{\theta}/2\right)\right], \tag{1.57}$$

where

$$\bar{\theta} = \theta + \theta_u + \theta_d.$$

Here θ_u and θ_d denote the phases of the up- and down-quark masses. The parameter B can be chosen real, since the only CP-violating phase in the QCD Lagrangian is the invariant combination $\bar{\theta}$. The minimum of the potential is obtained from the condition

$$\frac{\partial V}{\partial \pi_0} = 0,$$

which gives

$$\tan\frac{\pi_0}{f_\pi} = \frac{m_u - m_d}{m_u + m_d}\tan\frac{\bar{\theta}}{2}, \qquad V = -m_\pi^2 f_\pi^2 \sqrt{1 - \frac{4m_u m_d}{(m_u+m_d)^2}\sin^2\frac{\bar{\theta}}{2}}. \tag{1.58}$$

Expanding around the minimum, one finds that the neutral pion mass depends on $\bar{\theta}$, while the charged pion mass is

$$m_{\pi^0}^2(\bar{\theta}) = B\sqrt{m_u^2 + m_d^2 + 2m_u m_d\cos\bar{\theta}}, \qquad m_{\pi^\pm}^2 = B(m_u + m_d). \tag{1.59}$$

For $\bar{\theta} = 0$, the usual leading-order chiral perturbation theory result is recovered,

$$m_{\pi^0}^2 = m_{\pi^\pm}^2 = B(m_u + m_d).$$

Experimentally, the neutral and charged pion masses are very close, up to electromagnetic and isospin-breaking corrections. This provides a first indication that $\bar{\theta}$ must be very small.

1.3 QCD vacuum

1.3.1 Path Integral Formulation of Yang–Mills Theories

We now turn to the topological structure of the QCD vacuum. As a first step, following Ref. [20], we review the path-integral formulation of a pure Yang–Mills theory with gauge group $SU(N)$. This first step helps to clarify how to compute things in a single topological sector, where gauge transformation are always connected to the identity. We start from the pure Yang–Mills Lagrangian

$$\mathcal{L} = -\frac{1}{4}\sum_a F_{\mu\nu}^a F_a^{\mu\nu}. \tag{1.60}$$

It is convenient to work in temporal gauge, $A_0^a = 0$. In this gauge the dynamical variables are the spatial components A_i^a. The canonical momenta conjugate to A_i^a are

$$\Pi_i^a = F_{0i}^a \equiv E_i^a, \tag{1.61}$$

and, in temporal gauge, this reduces to $E_i^a = \dot{A}_i^a$. The Hamiltonian density then takes the familiar form

$$\mathcal{H} = \frac{1}{2}\left[(E_i^a)^2 + (B_i^a)^2\right], \tag{1.62}$$

where the non-Abelian chromomagnetic field is defined by

$$B_i^a = \frac{1}{2}\epsilon_{ijk}F_{jk}^a = \frac{1}{2}\epsilon_{ijk}\left(\partial_j A_k^a - \partial_k A_j^a + f^{abc}A_j^b A_k^c\right). \tag{1.63}$$

The Yang–Mills equations of motion are $D_\mu F^{\mu\nu} = 0$. In particular choosing the Lorentz index $\nu = 0$ gives Gauss' law,

$$D_i E_i = 0, \tag{1.64}$$

or, in components,

$$C^a(A) \equiv (D_i E_i)^a = \partial_i \dot{A}_i^a + f^{abc}A_i^b \dot{A}_i^c = 0. \tag{1.65}$$

This equation is a constraint rather than a dynamical equation. Indeed, $C^a(A)$ generates time-independent gauge transformations, while the Hamiltonian is gauge invariant. After our gauge fixing, the restriction to time-independent transformations is necessary because we have fixed the gauge $A_0^a = 0$, and a generic time-dependent gauge transformation would not preserve this condition. For an infinitesimal gauge transformation with parameter $\lambda^a(\mathbf{x})$, the gauge field transforms as

$$A_\mu^a \to A_\mu^a + (D_\mu\lambda)^a, \qquad \delta A_\mu^a = (D_\mu\lambda)^a. \tag{1.66}$$

The corresponding generator is

$$Q_\lambda = \int d^3x\, \Pi_i^a(\mathbf{x})(D_i\lambda)^a(\mathbf{x}), \tag{1.67}$$

where only the spatial components appear because A_0^a is not a dynamical variable in temporal gauge. Using the canonical Poisson brackets

$$\left\{A_i^a(\mathbf{x}), \Pi_j^b(\mathbf{y})\right\} = \delta^{ab}\delta_{ij}\delta^{(3)}(\mathbf{x}-\mathbf{y}), \tag{1.68}$$

one immediately finds, as expected

$$\{A_i^a(\mathbf{x}), Q_\lambda\} = (D_i\lambda)^a(\mathbf{x}), \tag{1.69}$$

Integrating Eq. (1.67) by parts, and assuming that $\lambda^a(\mathbf{x})$ vanishes sufficiently fast at spatial infinity, one obtains

$$Q_\lambda = -\int d^3x\, \lambda^a(\mathbf{x})(D_i\Pi_i)^a(\mathbf{x}). \tag{1.70}$$

Since $\Pi_i^a = E_i^a$, the quantity multiplying λ^a is precisely Gauss' constraint. Therefore physical states must satisfy

$$(D_iE_i)^a|\psi\rangle = 0. \tag{1.71}$$

Equivalently, physical states are invariant under gauge transformations connected to the identity and trivial at spatial infinity:

$$e^{iQ_\lambda}|\psi\rangle = |\psi\rangle, \qquad \lambda^a(\mathbf{x}) \xrightarrow[|\mathbf{x}|\to\infty]{} 0. \tag{1.72}$$

It is useful to rewrite this statement in terms of finite gauge transformations. Defining $A_i = A_i^a T^a$, a time-independent gauge transformation acts as

$$A_i \to UA_iU^{-1} + i\,U\partial_iU^{-1}, \tag{1.73}$$

up to the conventional placement of the gauge coupling in the definition of A_i and of the covariant derivative. Field eigenstates $|A_i\rangle$, however, are not themselves gauge invariant. A gauge-invariant state can be formally obtained by averaging over the gauge orbit,

$$|A_i\rangle_{\text{phys}} = \int [D\lambda]\, e^{iQ_\lambda}|A_i\rangle. \tag{1.74}$$

Since λ is an integration variable, a further gauge transformation only amounts to a change of variables in the integral, leaving the physical state unchanged.
To parametrize the physical configuration space, one may choose a representative on each gauge orbit. For example, one can impose another gauge condition, the transverse gauge fixing:

$$\nabla \cdot A^{\text{tr}} = 0. \tag{1.75}$$

Then a generic field configuration can be written as

$$A_i(\mathbf{x},t) = U(\mathbf{x})A_i^{\text{tr}}(\mathbf{x},t)U^{-1}(\mathbf{x}) + i\,U(\mathbf{x})\partial_iU^{-1}(\mathbf{x}), \qquad \nabla\cdot A^{\text{tr}} = 0. \tag{1.76}$$

Thus the full functional integration over gauge fields can be understood as an integration over physical representatives, such as A_i^{tr}, together with an integration over gauge transformations U. The latter integration only counts gauge redundancy and must be treated appropriately in the path integral.
In this way, the temporal-gauge Hamiltonian formulation makes explicit two important points. First, the condition $A_0 = 0$ can be imposed while leaving a residual class of time-independent gauge transformations. Second, Gauss' law appears as a constraint on physical states, enforcing their invariance under gauge transformations connected to the identity and trivial at spatial infinity. This will be the starting point for understanding the non-trivial vacuum structure of Yang–Mills theory.

1.3.2 Winding Number

We can now study the structure of the Yang–Mills vacuum and the stationary points of the action. Useful references for the following discussion are Refs. [21, 22, 23, 24]. Working in Euclidean space, the requirement of finite action implies that the field strength must vanish at infinity,

$$F^a_{\mu\nu} \xrightarrow[|x|\to\infty]{} 0. \tag{1.77}$$

therefore the gauge potential must approach a pure gauge configuration at infinity. In this limit, Eq. (1.76) reduces to

$$A_\mu(x) \xrightarrow[|x|\to\infty]{} i\, U(x)\partial_\mu U^{-1}(x), \tag{1.78}$$

up to the conventional placement of the gauge coupling in the definition of A_μ. In this way, any finite-action configuration defines a map from the boundary of Euclidean spacetime into the gauge group. Since the boundary of compactified $\mathbb{R}^4$ is a three-sphere, we have

$$U : S^3_\infty \longrightarrow SU(3). \tag{1.79}$$

In other words, we have compactified the boundary of $\mathbb{R}^4$ to S^3_∞. In Minkowski space one can formulate an analogous discussion in terms of a cylinder $S^2 \times \mathbb{R}$, but for the present purpose the Euclidean formulation is the most convenient one.
Such maps are classified by homotopy theory. Two maps are said to belong to the same homotopy class if they can be continuously deformed into each other. Before discussing the formal classification, it is useful to keep the physical picture in mind. As we have said, at infinity the field strength vanishes, and therefore the gauge field is gauge-equivalent to $A_\mu = 0$. Thus the configuration approaches the vacuum. Nevertheless, these asymptotic pure-gauge configurations can still differ by their global topological properties. In the path integral formulation one usually accounts only for local gauge transformation (i.e. those connected to the identity) and if we have different topological structures that cannot be continuously deformed into each other one should add the different contributions of each topological sectors. This different topological sectors are classified by the homotopy class of the map $U : S^3_\infty \to SU(3)$. The relevant result is, indicating with $\pi_n(G)$ the homotopy class for a group G in S^n that $\pi_3(SU(3)) = \mathbb{Z}$, and more generally

$$\pi_3(SU(N)) = \mathbb{Z}, \qquad N \geq 2. \tag{1.80}$$

therefore each homotopy class is labelled by an integer. A simple way to visualize this classification is to consider the lower-dimensional analogy of maps from a circle into $U(1)$, namely $\pi_1(U(1))$. Such maps can be written as

$$g(\chi) = e^{i\phi(\chi)} \in U(1), \qquad \chi \in [0, 2\pi]. \tag{1.81}$$

They are classified by the number of times the phase $\phi(\chi)$ winds around the target circle as χ goes once around the domain circle. This integer is the winding number. In Yang–Mills theory the same idea appears in the classification of maps $S^3 \to SU(N)$.

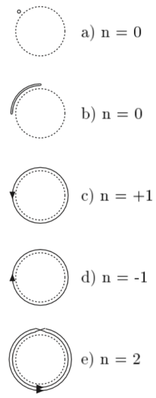

Figure 1.2: Illustration of maps $S^1 \to U(1)$ with different winding numbers. figure taken from Ref. [21]

The integer which labels the homotopy class of a finite-action Yang–Mills configuration is also called Pontryagin index. It is given by

$$\nu[A] = \frac{g^2}{64\pi^2}\epsilon^{\alpha\beta\gamma\delta}\int d^4x\, F^a_{\alpha\beta}F^a_{\gamma\delta} = \frac{g^2}{32\pi^2}\int d^4x\, \mathrm{Tr}\, F_{\mu\nu}\widetilde{F}_{\mu\nu}. \tag{1.82}$$

Here as usual $\widetilde{F}_{\mu\nu} = \frac{1}{2}\epsilon_{\mu\nu\rho\sigma}F_{\rho\sigma}$ is the dual field strength. In Euclidean space we use the convention $\epsilon_{1234} = +1$, while in Minkowski space we take $\epsilon_{0123} = +1$. The normalization in Eq. (1.82) assumes $\mathrm{Tr}(T^aT^b) = \frac{1}{2}\delta^{ab}$. We now show that Eq. (1.82) is indeed a topological quantity. As expected for a quantity which depends only on the homotopy class, it is invariant under smooth local variations of the gauge field. Let us consider a variation $\delta A^a_\mu(x)$. We have

$$\begin{aligned} \frac{\delta A^a_\mu(x)}{\delta A^b_\nu(y)} &= \delta^{ab}\delta^\nu_\mu\delta^{(4)}(x-y), \\ \frac{\delta F^a_{\alpha\beta}(x)}{\delta A^b_\gamma(y)} &= \left(D^{ab}_\alpha\delta^\gamma_\beta - D^{ab}_\beta\delta^\gamma_\alpha\right)_x \delta^{(4)}(x-y). \end{aligned} \tag{1.83}$$

Here D_α is the covariant derivative in the adjoint representation,

$$D^{ab}_\alpha = \delta^{ab}\partial_\alpha + f^{acb}A^c_\alpha, \tag{1.84}$$

up to the convention used for the gauge coupling. Equivalently, one may write $D_\alpha = \partial_\alpha - igA^c_\alpha t^c$, with $(t^c)_{ab} = if^{acb}$. The first equality is simply the definition of the functional derivative, while the second follows from the non-Abelian expression for the field strength. Using these relations, we find

$$\begin{aligned} \frac{\delta\nu[A]}{\delta A^b_\gamma(y)} &= \frac{g^2}{64\pi^2}\epsilon^{\alpha\beta\rho\sigma}\, 2\int d^4x\, \frac{\delta F^a_{\alpha\beta}(x)}{\delta A^b_\gamma(y)}F^a_{\rho\sigma}(x) \\ &= \frac{g^2}{32\pi^2}\epsilon^{\alpha\beta\rho\sigma}\int d^4x\, \left(D^{ab}_\alpha\delta^\gamma_\beta - D^{ab}_\beta\delta^\gamma_\alpha\right)_x \delta^{(4)}(x-y)F^a_{\rho\sigma}(x) \\ &= \frac{g^2}{16\pi^2}\epsilon^{\gamma\mu\alpha\beta}D^{ab}_\mu F^a_{\alpha\beta}(y) \\ &= 0. \end{aligned} \tag{1.85}$$

In the last step we used the Bianchi identity, $\epsilon^{\gamma\mu\alpha\beta}D_\mu F_{\alpha\beta} = D_\mu\widetilde{F}^{\mu\gamma} = 0$.
therefore $\nu[A]$ does not change under smooth local deformations of the gauge field. This proves that

$\nu[A]$ is a topological invariant. As we will also show, for finite-action configurations it is an integer and coincides with the winding number of the map $U : S^3_\infty \to SU(3)$. Hence $\nu[A]$ labels the homotopy class of the gauge configuration.

At this point it is useful to rewrite the winding number in a form which makes its boundary nature explicit. In fact, the topological charge density can be written as a total derivative,

$$\begin{aligned}\frac{g^2}{64\pi^2}\epsilon^{\mu\nu\alpha\beta}F^a_{\mu\nu}F^a_{\alpha\beta} &= \partial_\mu K^\mu, \\ K^\mu &= \frac{g^2}{32\pi^2}\epsilon^{\mu\alpha\beta\gamma}A^a_\alpha\left(F^a_{\beta\gamma}-\frac{1}{3}gf^{abc}A^b_\beta A^c_\gamma\right).\end{aligned} \tag{1.86}$$

We have included the normalization factor directly in the definition of K^μ; different conventions in the literature sometimes define the Chern–Simons current without this prefactor.
Note that the fact that $F\widetilde{F}$ can be written as a total derivative is natural in some sense, since $\nu[A]$ is a topological invariant, it should depend only on the behavior of the gauge field at infinity, namely on the boundary of the configuration. Indeed, using Gauss' theorem one obtains

$$\nu[A]_R = \int_{S^3_R} d\sigma_\mu\, K^\mu, \tag{1.87}$$

and the winding number is obtained by taking the limit $R \to \infty$.
Another important remark is that the current K^μ is not gauge invariant by itself. Under a finite gauge transformation U, it transforms as

$$\begin{aligned}K^\mu \to K'^\mu = K^\mu &+ \frac{1}{24\pi^2}\mathrm{Tr}\,\epsilon^{\mu\alpha\beta\gamma}\left(U^{-1}\partial_\alpha U\right)\left(U^{-1}\partial_\beta U\right)\left(U^{-1}\partial_\gamma U\right) \\ &+ \frac{ig}{8\pi^2}\partial_\beta \mathrm{Tr}\,\epsilon^{\mu\alpha\beta\gamma}\partial_\alpha U\, U^{-1}A_\gamma.\end{aligned} \tag{1.88}$$

The last term is itself a total derivative on the boundary, while the second term is the winding-number density of the gauge transformation U. This makes explicit the fact that the integral of K^μ over S^3_∞ measures the homotopy class of the map $U : S^3_\infty \to SU(N)$.
The limit $R \to \infty$ is well defined for finite-action configurations. Indeed, in Euclidean space one has the pointwise inequality

$$-F^a_{\mu\nu}F^{a\mu\nu} \le F^a_{\mu\nu}\widetilde{F}^{a\mu\nu} \le F^a_{\mu\nu}F^{a\mu\nu}. \tag{1.89}$$

therefore, if the Euclidean action is finite, then the integral of $F\widetilde{F}$ is absolutely convergent. this gives

$$S_E[A] \ge \frac{8\pi^2}{g^2}\,|\nu[A]|. \tag{1.90}$$

And the bound is saturated if and only if

$$F^a_{\mu\nu} = \pm\widetilde{F}^a_{\mu\nu}. \tag{1.91}$$

Configurations satisfying the plus sign are called self-dual, while those satisfying the minus sign are called anti-self-dual. The Yang–Mills equations of motion are

$$\frac{\delta S_E[A]}{\delta A_\beta} = -D_\alpha F_{\alpha\beta} = 0. \tag{1.92}$$

Self-dual and anti-self-dual configurations automatically satisfy these equations, since

$$D_\alpha F_{\alpha\beta} = \pm D_\alpha \widetilde{F}_{\alpha\beta} = 0 \tag{1.93}$$

by the Bianchi identity. However, they are not the only possible stationary points of the Euclidean Yang–Mills action. Rather, they are the absolute minima of the action in a fixed topological sector, because they saturate the bound

$$S_E[A] \ge \frac{8\pi^2}{g^2}|\nu[A]|.$$

The general construction of all self-dual instanton solutions is provided by the ADHM in Ref. [25]. The simplest example of instanton solution is obtained for the gauge group $SU(2)$ and winding number $\nu = 1$. It is the instanton found in Ref. [26], also referred to as BPST solution:

$$A_\mu(x) = \frac{\rho^2 (x-x_0)_\lambda}{(x-x_0)^2\left[(x-x_0)^2+\rho^2\right]} \eta^i_{\mu\lambda} \left(g\sigma^i g^{-1}\right). \tag{1.94}$$

Here σ^i are the Pauli matrices, ρ is the instanton size, x_0 is the position of the instanton, and $g \in SU(2)$ is a global gauge orientation. The symbols $\eta^i_{\mu\nu}$ are the self-dual 't Hooft symbols, which intertwine the spacetime indices μ, ν with the $SU(2)$ gauge index i. In the conventions used here, they can be written as

$$\eta^1 = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix}, \qquad \eta^2 = \begin{pmatrix} 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix}, \qquad \eta^3 = \begin{pmatrix} 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{pmatrix}.$$

The corresponding anti-instanton is obtained by replacing the self-dual symbols $\eta^i_{\mu\nu}$ with the anti-self-dual symbols $\bar{\eta}^i_{\mu\nu}$.

We just mention that the BPST solution contains a finite number of parameters, called collective coordinates. These include the position x_0, the size ρ, and the global gauge orientation. More generally, for an $SU(N)$ instanton with topological charge ν, the dimension of the instanton moduli space (that is the parameter space that defines an instanton solution) is $4N|\nu|$ see e.g. Ref. [21].

In the path integral, one should integrate over these collective coordinates in the semiclassical approximation.

However for self-dual or anti-self-dual at configurations at classical level, the action is fixed entirely by the topological charge:

$$S_E[A] = \frac{8\pi^2}{g^2}|\nu[A]|, \qquad \nu[A] \in \mathbb{Z}. \tag{1.95}$$

This will be enough to carry on our discussion. From Eq. (1.95) is clear that Yang–Mills theory possesses infinitely many classical finite-action configurations, one for each topological sector. Physically, instantons describe tunneling events between vacua with different winding numbers. Slightly more precisely, an instanton solution with topological charge ν interpolates between vacua whose winding numbers differ by ν. Since both the initial and final configurations are pure gauge, while the instanton has non-zero field strength in the bulk of Euclidean spacetime, the process can be interpreted as tunneling between distinct classical minima. If one goes carefully in the imaginary time formulation, one can see that in this case instanton solutions are just the usual classical motion.

The non-perturbative nature of instantons is already visible from Eq. (1.95). The semiclassical contribution of a configuration with action S_E to the Euclidean path integral is weighted by

$$e^{-S_E} = \exp\left(-\frac{8\pi^2}{g^2}|\nu|\right), \tag{1.96}$$

which is non-analytic in g^2 around $g = 0$. Therefore such effects cannot be obtained at any finite order in ordinary perturbation theory.

Very briefly, and as we have anticipated, the path integral should then be understood as a sum over contributions from all topological sectors. In each sector, one expands semiclassically around the corresponding classical configurations, including instantons, anti-instantons, and multi-instanton configurations. The latter can be understood, in a dilute-gas approximation, as configurations made of several instantons and anti-instantons that are well separated in Euclidean spacetime see Ref. [20]. However, this does not mean that instantons exhaust all possible stationary points of the Euclidean Yang–Mills action. Rather, they are the relevant leading semiclassical saddles because they are finite-action configurations of minimal action in each topological sector. Other stationary points, when present, have larger action or may be singular, and are not included in the leading instanton approximation.

Merons are singular configurations with fractional topological charge and do not have finite Euclidean action when considered in isolation. Sphalerons, on the other hand, are unstable static saddle points of the energy functional, not regular finite-action Euclidean solutions. For this reason, in the standard semiclassical treatment of pure Yang–Mills theory on $\mathbb{R}^4$, the relevant leading finite-action saddles are instantons and anti-instantons.

Some Properties

Following Ref. [22], we now show explicitly that the topological charge $\nu[A]$ is an integer. For simplicity we specialize to the gauge group $SU(2)$. The result generalizes to $SU(N)$, with $N \geq 2$, since $\pi_3(SU(N)) = \mathbb{Z}$. As discussed above, finite Euclidean action implies that the field strength vanishes sufficiently fast at infinity. More precisely, for the class of smooth configurations we are considering, one may require

$$x^2 F_{\mu\nu}(x) \xrightarrow[|x|\to\infty]{} 0 \tag{1.97}$$

in almost every direction. Hence the gauge field approaches a pure gauge configuration:

$$gA^a_\mu T^a \xrightarrow[|x|\to\infty]{} i\, G^{-1}\partial_\mu G, \tag{1.98}$$

where $G(x) \in SU(2)$. With respect to Eq. (1.88), we have changed notation in order to reserve G for the group element defining the asymptotic pure gauge. Using Eq. (1.98) in the expression for the Chern–Simons current, one finds that, at infinity,

$$K_\mu \to \frac{1}{24\pi^2}\epsilon_{\mu\alpha\beta\gamma}\mathrm{Tr}\left[(G^{-1}\partial_\alpha G)(G^{-1}\partial_\beta G)(G^{-1}\partial_\gamma G)\right]. \tag{1.99}$$

therefore the winding number can be written as the surface integral

$$\nu = \frac{1}{24\pi^2}\int_{S^3_R} d\sigma_\mu\, \epsilon_{\mu\alpha\beta\gamma}\mathrm{Tr}\left[(G^{-1}\partial_\alpha G)(G^{-1}\partial_\beta G)(G^{-1}\partial_\gamma G)\right], \tag{1.100}$$

with the integer value obtained in the limit $R \to \infty$. We now make the geometrical meaning of Eq. (1.100) explicit. Since we are working with $SU(2)$, every group element can be parametrized as

$$G = V_4\,\mathbb{I} - i\,\vec{\sigma}\cdot\vec{V}, \tag{1.101}$$

where $\vec{\sigma}$ are the Pauli matrices. The conditions $G^\dagger G = \mathbb{I}$ and $\det G = 1$ imply

$$V_4^2 + \vec{V}^{\,2} = 1. \tag{1.102}$$

Thus the four real variables $V_\rho = (V_1, V_2, V_3, V_4)$ parametrize a unit three-sphere. This is the familiar identification $SU(2) \simeq S^3$. Using this parametrization, Eq. (1.100) becomes

$$\nu = \frac{1}{12\pi^2}\int_{S^3_R} d\sigma(x)\,\epsilon_{\mu\alpha\beta\gamma}\hat{x}_\mu\,\epsilon_{\rho\tau\eta\delta}V_\rho\partial_\alpha V_\tau\partial_\beta V_\eta\partial_\gamma V_\delta. \tag{1.103}$$

This expression has a simple interpretation: it is the Jacobian of the map from the sphere at infinity in x-space to the group manifold $SU(2) \simeq S^3$ in V-space. To see this more explicitly, note that $d\sigma_\mu(x) = \hat{x}_\mu d\sigma(x)$, so that

$$d\sigma_\mu(x)\epsilon_{\mu\alpha\beta\gamma} = d\sigma(x)\,\hat{x}_\mu\epsilon_{\mu\alpha\beta\gamma}. \tag{1.104}$$

This object selects the three tangent directions to S^3_R in x-space. Similarly, since $V_\rho V^\rho = 1$, the vector V_ρ is normal to the target sphere S^3 in V-space, and $\epsilon_{\rho\tau\eta\delta}V_\rho$ selects the corresponding oriented volume form on the target S^3. To see explicitly how Eq. (1.100) becomes the Jacobian of the map between the two three-spheres, it is useful to rewrite the integrand in terms of the coordinates V_ρ. We first use

$$G^{-1}\partial_\mu G = -(\partial_\mu G^{-1})G, \tag{1.105}$$

which follows from $\partial_\mu(G^{-1}G) = 0$. Then Eq. (1.100) can be written as

$$\begin{aligned}\nu &= \frac{1}{24\pi^2}\int_{S_R^3} d\sigma_\mu\, \epsilon_{\mu\alpha\beta\gamma}\mathrm{Tr}\left[G^{-1}\partial_\alpha G\,\partial_\beta G^{-1}\,\partial_\gamma G\right] \\ &= \frac{1}{24\pi^2}\int_{S_R^3} d\sigma_\mu\, \epsilon_{\mu\alpha\beta\gamma}V^\rho\partial_\alpha V^\tau\partial_\beta V^\eta\partial_\gamma V^\delta\,\mathrm{Tr}\left[\sigma_\rho^\dagger\sigma_\tau\sigma_\eta^\dagger\sigma_\delta\right].\end{aligned} \tag{1.106}$$

Here we have introduced the compact notation

$$\sigma^\rho = (\mathbb{I}, -i\sigma^i), \qquad \sigma^{\rho\dagger} = (\mathbb{I}, +i\sigma^i), \tag{1.107}$$

so that

$$G = V^\rho\sigma_\rho, \qquad G^{-1} = G^\dagger = V^\rho\sigma_\rho^\dagger. \tag{1.108}$$

The trace appearing in Eq. (1.106) selects the completely antisymmetric part of the product of four V's. Indeed, any contribution in which two of the internal indices coincide gives a term containing a symmetric product such as

$$\partial_\alpha V^\lambda\,\partial_\beta V^\lambda, \tag{1.109}$$

which vanishes after contraction with the antisymmetric tensor $\epsilon_{\mu\alpha\beta\gamma}$. Thus only the part of the trace which is antisymmetric in the four indices ρ,τ,η,δ contributes. Using the Pauli-matrix identities

$$\mathrm{Tr}(\sigma_i\sigma_j\sigma_k) = 2i\epsilon_{ijk}, \qquad \mathrm{Tr}(\mathbb{I}) = 2, \tag{1.110}$$

one finds

$$\mathrm{Tr}\left[\sigma_\rho^\dagger\sigma_\tau\sigma_\eta^\dagger\sigma_\delta\right]_{\text{antisym}} = 2\,\epsilon_{\rho\tau\eta\delta}. \tag{1.111}$$

therefore Eq. (1.106) reduces to

$$\nu = \frac{1}{12\pi^2}\int_{S_R^3} d\sigma_\mu\,\epsilon_{\mu\alpha\beta\gamma}\epsilon_{\rho\tau\eta\delta}V^\rho\partial_\alpha V^\tau\partial_\beta V^\eta\partial_\gamma V^\delta. \tag{1.112}$$

Writing $d\sigma_\mu = \hat{x}_\mu d\sigma(x)$, this becomes

$$\nu = \frac{1}{12\pi^2}\int_{S_R^3} d\sigma(x)\,\hat{x}_\mu\epsilon_{\mu\alpha\beta\gamma}\epsilon_{\rho\tau\eta\delta}V^\rho\partial_\alpha V^\tau\partial_\beta V^\eta\partial_\gamma V^\delta. \tag{1.113}$$

This is precisely the structure of the Jacobian of the map $V : S_x^3 \longrightarrow S_V^3$. Equivalently,

$$d\sigma(V) = \frac{1}{6}\hat{x}_\mu\epsilon_{\mu\alpha\beta\gamma}\epsilon_{\rho\tau\eta\delta}V_\rho\partial_\alpha V_\tau\partial_\beta V_\eta\partial_\gamma V_\delta\,d\sigma(x). \tag{1.114}$$

Substituting this into Eq. (1.113), we obtain

$$\nu = \frac{1}{2\pi^2}\int_{\text{mult}} d\sigma(V). \tag{1.115}$$

Here the notation $\int_{\text{mult}}$ indicates that the target sphere may be covered more than once by the map $V : S_x^3 \to S_V^3$. Since the volume of the unit three-sphere is $\mathrm{Vol}(S^3) = 2\pi^2$, the integral gives an integer $\nu \in \mathbb{Z}$. More formally, ν is the degree of the map $V : S_x^3 \to S_V^3$. For a regular point of the target sphere, the degree is obtained by summing over all its preimages, counting each preimage with a sign determined by whether the map preserves or reverses orientation. This signed number is an integer and is independent of the chosen regular point. Therefore the winding number ν is an integer topological invariant.

For a gauge group $SU(N)$ with $N > 2$, non-trivial configurations can already be obtained by embedding an $SU(2)$ subgroup into $SU(N)$. More generally, the result is guaranteed by

$$\pi_3(SU(N)) = \mathbb{Z}, \qquad N \geq 2, \tag{1.116}$$

so the integer winding number persists for all $SU(N)$ Yang–Mills theories.

The Index Theorem

The fact that $\nu[A]$ is an integer can also be understood from the Atiyah–Singer index theorem. This theorem states that in the presence of an instanton background with winding number k, each Weyl fermion picks $2kT(R)$ zero modes, where $T(R)$ depends on the specific group representation, in the fundamental representation $T(R) = 1/2$. The index theorem links the index of the Euclidean Dirac operator in that background see Refs. [25, 27], with the winding number. Let us consider the Euclidean Dirac operator

$$i\not{D}_E \varphi_n = \lambda_n \varphi_n, \tag{1.117}$$

where $i\not{D}_E$ is Hermitian, so that all eigenvalues λ_n are real. Moreover, since $\{i\not{D}_E, \gamma_5\} = 0$, if φ_n is an eigenfunction with eigenvalue λ_n, then $\gamma_5\varphi_n$ is an eigenfunction with eigenvalue $-\lambda_n$. Therefore as already seen in Sec. (1.2.2) all non-zero eigenvalues come in pairs of opposite chirality and do not contribute to the index.
The only possible unpaired states are the zero modes. In the zero-mode subspace, $\lambda_n = 0$, one can diagonalize γ_5. We denote by n_+ the number of zero modes with positive chirality, $\gamma_5\varphi_u = +\varphi_u$, and by n_- the number of zero modes with negative chirality, $\gamma_5\varphi_v = -\varphi_v$. Using a regulator f, with $f(0) = 1$ and $f(\infty) = 0$, one can write the index as

$$\mathrm{ind}(i\not{D}_E) = \lim_{\Lambda\to\infty} \mathrm{Tr}\left[\gamma_5 f\left(-\frac{\not{D}_E^2}{\Lambda^2}\right)\right] = n_+ - n_-. \tag{1.118}$$

On the other hand, evaluating the same regulated trace locally gives

$$\mathrm{ind}(i\not{D}_E) = \frac{g^2}{32\pi^2}\int d^4x\, \mathrm{Tr}\, F_{\mu\nu}\widetilde{F}_{\mu\nu} = \nu[A]. \tag{1.119}$$

Thus

$$n_+ - n_- = \nu[A]. \tag{1.120}$$

This is the Atiyah–Singer index theorem applied to the Euclidean Dirac operator in a Yang–Mills background. Since n_+ and n_- are integers, this also shows that the topological charge $\nu[A]$ is an integer.
We now show explicitly how, for a self-dual or anti-self-dual background, only one chirality can have zero modes. Let ψ_0 be a zero mode of the Euclidean Dirac operator,

$$i\not{D}_E\psi_0 = 0. \tag{1.121}$$

Then, of course,

$$(i\not{D}_E)^2\psi_0 = 0. \tag{1.122}$$

Using Euclidean gamma matrices, one finds

$$(i\not{D}_E)^2 = -D_\mu D_\mu - \frac{g}{2}\sigma_{\mu\nu}F_{\mu\nu}, \tag{1.123}$$

where $\sigma_{\mu\nu} = \frac{1}{2}[\gamma_\mu, \gamma_\nu]$. We decompose the zero mode into its two chiral components, and use the useful identity:

$$\sigma_{\mu\nu} = -\gamma_5\widetilde{\sigma}_{\mu\nu}, \qquad \widetilde{\sigma}_{\mu\nu} = \frac{1}{2}\epsilon_{\mu\nu\rho\sigma}\sigma_{\rho\sigma}. \tag{1.124}$$

Now suppose that the gauge field is self-dual, Then

$$\sigma_{\mu\nu}F_{\mu\nu} = \sigma_{\mu\nu}\widetilde{F}_{\mu\nu} = \widetilde{\sigma}_{\mu\nu}F_{\mu\nu} = -\gamma_5\sigma_{\mu\nu}F_{\mu\nu}. \tag{1.125}$$

Multiplying this equation by the chiral projectors, one obtains

$$\sigma_{\mu\nu}F_{\mu\nu}P_+ = \sigma_{\mu\nu}F_{\mu\nu}\frac{1+\gamma_5}{2} = 0. \tag{1.126}$$

Thus, for a self-dual field, the spin term in Eq. (1.123) does not act on the positive-chirality component. Projecting Eq. (1.123) on $\psi_{0,+}$, we get

$$0 = (i\not{D}_E)^2\psi_{0,+} = -D_\mu D_\mu \psi_{0,+}. \tag{1.127}$$

For a positive definite operator one cannot have a vanishing eigenvalue, if the eigenvector is different from zero. Therefore

$$\psi_{0,+} = 0, \qquad n_+ = 0. \Rightarrow -n_- = \nu[A]. \tag{1.128}$$

Thus, in a self-dual background, only negative-chirality zero modes can exist. Analogously, in the anti-self background only positive-chirality zero modes can exist.

1.3.3 The QCD Vacuum in Minkowski spacetime

The discussion above was formulated in Euclidean spacetime. There, finite-action gauge configurations approach pure gauge configurations at infinity, so that one may compactify $\mathbb{R}^4$ to S^4. The corresponding topological charge is an integer and is associated with the homotopy class of the gauge transformation at infinity. We now reformulate the same physics in Minkowski spacetime, following Ref. [22], in order to understand how the non-trivial vacuum structure of Yang–Mills theory leads to a new term in the QCD Lagrangian. We work again in temporal gauge,

$$A_0 = 0, \tag{1.129}$$

so that the residual gauge transformations are time-independent. Consider the cylindrical surface $\sigma(r, T)$ in Minkowski spacetime, with spatial radius r and time extension $2T$,

$$\sigma(r, T) = \{x^\mu : |\mathbf{x}| = r \text{ or } t = \pm T\}\,. \tag{1.130}$$

As in the Euclidean formulation, we restrict the path integral to gauge fields which approach pure gauge configurations at the boundary. Thus, on the boundary of the cylinder, one may write

$$-ig\, A_i^a T^a \longrightarrow \begin{cases} G_m^{-1}(\mathbf{x})\,\partial_i G_m(\mathbf{x}), & t \to -\infty, \\ G_n^{-1}(\mathbf{x})\,\partial_i G_n(\mathbf{x}), & t \to +\infty, \\ g^{-1}(\Omega)\,\partial_i g(\Omega), & |\mathbf{x}| \to \infty. \end{cases} \tag{1.131}$$

Here G_m and G_n are time-independent gauge functions defined on the initial and final constant-time hypersurfaces, while Ω denotes the angular coordinates at spatial infinity. Consistency on the lateral surface requires G_m and G_n to approach the same angular function $g(\Omega)$ as $|\mathbf{x}| \to \infty$. As we have seen, since the topological charge density is a total derivative,

$$\frac{g^2}{64\pi^2}\epsilon^{\mu\nu\rho\sigma} F^a_{\mu\nu} F^a_{\rho\sigma} = \partial_\mu K^\mu, \tag{1.132}$$

the topological charge can be written as a surface integral,

$$\nu = \lim_{r,T\to\infty} \int_{\sigma(r,T)} d\sigma_\mu\, K^\mu(t, \mathbf{x}). \tag{1.133}$$

For a time-independent gauge transformation G, we define its winding number by

$$K[G] = -\frac{1}{24\pi^2}\int d^3x\, \mathrm{Tr}\left[\epsilon_{ijk} G^{-1}\partial_i G\, G^{-1}\partial_j G\, G^{-1}\partial_k G\right]. \tag{1.134}$$

The labels m and n are then defined by

$$K[G_m] = m, \qquad K[G_n] = n. \tag{1.135}$$

In the limit $r \to \infty$, the lateral contribution to Eq. (1.133) vanishes, and the topological charge reduces to the difference between the Chern–Simons numbers on the final and initial time slices:

$$\nu = \int d^3x \, \left[K^0(t = +\infty, \mathbf{x}) - K^0(t = -\infty, \mathbf{x})\right]. \tag{1.136}$$

With the convention used above, this gives

$$\nu = n - m. \tag{1.137}$$

The opposite sign may appear if the orientation of the cylindrical surface is chosen differently; the invariant statement is that ν measures the change in winding number between the final and initial asymptotic vacua. It is useful to remove the common angular dependence at spatial infinity by a gauge transformation,

$$G_m(\mathbf{x}) \longrightarrow G'_m(\mathbf{x}) = G_m(\mathbf{x})\, g^{-1}(\Omega), \tag{1.138}$$

and similarly for G_n. Then $G'_m \to \mathbb{I}$ as $|\mathbf{x}| \to \infty$, so each constant-time hypersurface may be compactified as

$$\mathbb{R}^3 \cup \{\infty\} \simeq S^3. \tag{1.139}$$

The asymptotic vacua are therefore classified by maps

$$S^3 \longrightarrow SU(3), \tag{1.140}$$

which are characterized by

$$\pi_3(SU(3)) = \mathbb{Z}. \tag{1.141}$$

Thus the integers m and n label the topological sectors of the initial and final pure gauge configurations. Although the individual representatives used to describe these sectors are gauge dependent, the difference $\nu = n - m$ is gauge invariant.

In terms of Hilbert space, let $|m\rangle$ denote an equivalence class of states associated with a pure-gauge vacuum of winding number m, that we will also denote as topological sector. Configurations related by small gauge transformations, namely transformations continuously connected to the identity, belong to the same equivalence class. By contrast, large gauge transformations have non-trivial winding number and map one sector into another. Thus the states $|m\rangle$ form a convenient basis of asymptotic vacua, but they are not themselves invariant under all gauge transformations.

Perturbative fluctuations preserve the winding number and therefore contribute to transitions of the form

$$|m\rangle_- \longrightarrow |m\rangle_+. \tag{1.142}$$

Non-perturbative configurations, such as instantons, instead connect vacua with different winding numbers. A configuration with topological charge ν mediates

$$|m\rangle_- \longrightarrow |m+\nu\rangle_+. \tag{1.143}$$

Let $\widetilde{G}$ be a large gauge transformation with unit winding number,

$$K[\widetilde{G}] = 1, \qquad \widetilde{G}(\mathbf{x}) \to \mathbb{I} \quad \text{for} \quad |\mathbf{x}| \to \infty. \tag{1.144}$$

Then $\widetilde{G}$ maps a representative of the sector m into one of the sector $m+1$. Denoting by U the unitary operator implementing this large gauge transformation, we have

$$U|m\rangle = |m+1\rangle. \tag{1.145}$$

Equivalently, the winding number transforms as

$$U^{-1}\left(\int d^3x\, K^0(x)\right) U = \int d^3x\, K^0(x) + 1, \tag{1.146}$$

so that

$$\int d^3x\, K^0(t = \pm\infty, \mathbf{x})\, |m\rangle_\pm = m\, |m\rangle_\pm. \tag{1.147}$$

The basis $\{|m\rangle_\pm\}$ may be chosen orthonormal separately in the in- and out-Hilbert spaces,

$$_\pm\langle m|n\rangle_\pm = \delta_{mn}, \tag{1.148}$$

but as we have seen transition amplitudes between in- and out-states need not be diagonal in m. Physical states must correspond to eigenstates of gauge transformation, working in the basis $\{|m\rangle\}$ we are already sure that this is the case for small gauge transformation, therefore it is enough to find eigenstates of U. Since U is unitary this means that a physical state must satisfy an equation of the type:

$$U|\text{phys}\rangle = e^{-i\theta}|\text{phys}\rangle \tag{1.149}$$

Where θ at this point is just an arbitrary phase. We can easily find the expression for the eigenstate with eigenvalue $e^{-i\theta}$ in terms of the pure-gauge states, and it is given by:

$$|\theta\rangle_\pm = \sum_m e^{im\theta}|m\rangle_\pm \to U|\theta\rangle_\pm = e^{-i\theta}|\theta\rangle_\pm \tag{1.150}$$

Where the last step can be verified using Eq. (1.145). For this eigenstates we have that the normalization is the following:

$$_\pm\langle\theta' \mid \theta\rangle_\pm = \sum_{m,n} e^{i\theta m - i\theta' n}\delta_{mn} = \sum_m e^{im(\theta-\theta')} = 2\pi\delta(\theta-\theta') \tag{1.151}$$

We have then built an asymptotic vacuum state (being a superposition of vacua in the basis $\{|m\rangle\}$) that is also gauge invariant, meaning that we have the physical vacuum. Since a QFT is defined to have only one vacuum state, we can split the $V-$space in different H_θ sectors, each of them describing physical states with a different $|\theta\rangle$. The condition on physical observables $O_k = U^{-1}O_kU$ and Eq. (1.145) imply that the T-product on some observables depends only on $\nu = m - n$:

$$T_+\langle m|\Pi_k O_k|n\rangle_- = T\langle\Pi_k O_k\rangle_\nu \tag{1.152}$$

then we can write:

$$T_+\langle\theta'|\Pi_k O_k|\theta\rangle_- = \sum_{m,n} e^{i(m\theta - n\theta')}T_+\langle n|\Pi_k O_k|m\rangle_- = 2\pi\delta(\theta-\theta')\sum_\nu e^{-i\nu\theta}T\langle\Pi_k O_k\rangle_\nu \tag{1.153}$$

After dividing by the vacuum-to-vacuum amplitude, the normalized Green function in a fixed θ-sector is therefore

$$T\langle\text{vac}|\prod_k O_k|\text{vac}\rangle_\theta = \frac{\sum_\nu e^{-i\nu\theta}\,\langle\prod_k O_k\rangle_\nu}{\sum_\nu e^{-i\nu\theta}\langle 1\rangle_\nu}. \tag{1.154}$$

In the Euclidean path-integral formulation this corresponds to

$$\sum_\nu e^{-i\nu\theta}\int [dA]_\nu \prod_k O_k\, e^{-S_E[A]} = \sum_\nu \int [dA]_\nu \prod_k O_k\, e^{-S_E[A] - i\theta\nu}. \tag{1.155}$$

Here $[dA]_\nu$ denotes integration over gauge configurations with topological charge ν. Since

$$\nu = \frac{g^2}{32\pi^2}\int d^4x\, \mathrm{Tr}\, G_{\mu\nu}\widetilde{G}^{\mu\nu}, \tag{1.156}$$

the phase associated with the θ-vacuum may be equivalently incorporated by adding the term

$$\mathcal{L}_\theta = \theta\, \frac{g^2}{32\pi^2}\mathrm{Tr}\, G_{\mu\nu}\widetilde{G}^{\mu\nu} \tag{1.157}$$

to the QCD Lagrangian. Thus working in a theory with a non-trivial θ-vacuum is equivalent to working with a trivial vacuum performing the path integral over only small gauge transformations, but with the modified Lagrangian density,

$$\mathcal{L}'_{\text{QCD}} = \mathcal{L}_{\text{QCD}} + \theta\, \frac{g^2}{32\pi^2} \text{Tr}\, G_{\mu\nu}\widetilde{G}^{\mu\nu}. \tag{1.158}$$

This term adds with the phase arising from the quark mass matrix. As we have seen, starting from the general quark mass term

$$\mathcal{L}_{\text{mass}} = -\bar{q}_{Li} M_{ij} q_{Rj} - \bar{q}_{Ri} M^\dagger_{ij} q_{Lj}. \tag{1.159}$$

The complex matrix M can be diagonalized by independent unitary transformations acting on left- and right-handed quarks. As we have already seen, in general this procedure involves an anomalous axial rotation, which shifts the θ-angle. We have reached the heart of the strong CP problem, that is since θ receives so different, and independent contributions, why it is so small.
Before elaborating further, we delve to some results that will be useful for the next chapters. The first is given by the vacuum energy density, that having found the physical vacuum we can now compute starting from:

$$\langle\theta'|e^{-HT}|\theta\rangle = 2\pi\delta(\theta-\theta')\sum_\nu e^{-i\theta\nu}\int \mathcal{D}A_\nu e^{-S_E} \tag{1.160}$$

In the limit $T \to \infty$, the left-hand side is dominated by the lowest-energy state in the sector with fixed θ, since the contributions from the excited Hamiltonian eigenstates are exponentially suppressed:

$$\langle\theta|e^{-HT}|\theta\rangle = \sum_n e^{-E_nT}\langle\theta|n\rangle\langle n|\theta\rangle \longrightarrow e^{-E(\theta)T}|\langle\theta|\theta\rangle|^2 = e^{-\mathcal{E}(\theta)V} \tag{1.161}$$

Here $\mathcal{E}(\theta)$ is the vacuum energy density and V_4 is the Euclidean spacetime volume. Therefore

$$\mathcal{E}(\theta) = -\frac{1}{V}\log\left(\sum_\nu e^{-i\theta\nu}\int \mathcal{D}A_\nu e^{-S_E}\right) = -\frac{1}{V}\log Z(\theta) \tag{1.162}$$

A crude approximation can be obtained by keeping only the perturbative sector $\nu = 0$ and the one-instanton sectors $|\nu| = 1$. Since for a configuration with topological charge ν one has

$$S_E \geq \frac{8\pi^2}{g^2}|\nu|, \tag{1.163}$$

higher-winding sectors are exponentially suppressed in the weak-coupling regime. Thus

$$\mathcal{E}(\theta) = -\frac{1}{V}\log Z_0 - \frac{1}{V}\log\left(1 + K_+ e^{-8\pi^2/g^2}e^{i\theta} + K_- e^{-i\theta}e^{-8\pi^2/g^2}\right) \tag{1.164}$$

Where $K_+ = K_-$ is given integrating the gauge fields around the (anti-) istanton solutions weighted against the perturbative regime. Then, assuming to work in the perturbative (weak coupling) regime we find the vacuum energy density:

$$\mathcal{E}(\theta) - \mathcal{E}(0) \simeq -2Ke^{\frac{-8\pi^2}{g^2}}\cos\theta. \tag{1.165}$$

The second result is that we can relate the winding number of a certain theory with chiral violation, in fact:

$$\Delta Q = \int d^4x\; \partial_\mu J_5^\mu(x) = n_f \frac{g^2}{16\pi^2}\int d^4x\; \text{Tr}\; G^{\mu\nu}\tilde{G}_{\mu\nu} = 2n_f\nu \tag{1.166}$$

Then if $\nu = 0$ (i.e. we are in a regime dominated by perturbation theory) the charge associated to J_5^μ is conserved, while in $\nu \neq 0$, the Noether current is not conserved. It was in Ref. [28] shown that the non-perturbative contributions are actually relevant in QCD even if we do not show here the details, meaning that we have no suppressed $\Delta Q \neq 0$, so that Q, that here we would identify to be the chirality, must change by multiples of $2n_f$. At this point let us mention that in Ref. [29] and Ref. [28] t'Hooft showed that the fermionic zero modes of the full path integral formulation in the form

$$Z = e^{-\frac{8\pi^2}{g^2}} e^{i\theta} \int [d\delta A][dq][d\bar{q}] e^{i\delta^2 S} \tag{1.167}$$

once a source term $\eta, \bar{\eta}$ is included, generate an effective Lagrangian, called 't Hooft vertex, that is schematically given by

$$\mathcal{L}_{eff} = \Pi_i^{N_f} (\bar{\psi}_i \omega)(\bar{\omega}\psi_i) \tag{1.168}$$

where ω is some fixed Dirac spinor. This interaction makes explicit that instantons violate the anomalous axial symmetry.
We now go back to the two different contributions given to $\bar{\theta}$, that we showed are given by the argument of the quark mass matrix, and by the topological structure of the vacuum QCD. These contributions add up, to see this explicitly consider the current J_5^μ. Although J_5^μ is gauge invariant, it is not conserved. One can define instead a conserved current

$$\widetilde{J}_5^\mu = J_5^\mu - n_f K^\mu, \tag{1.169}$$

whose associated charge is

$$\widetilde{Q}_5 = \int d^3x\, \widetilde{J}_5^0. \tag{1.170}$$

This charge generates axial transformations, but it is not invariant under large gauge transformations. In particular, for the unit-winding transformation U,

$$U^{-1} \widetilde{Q}_5 U = \widetilde{Q}_5 - n_f. \tag{1.171}$$

Using $U|\theta\rangle = e^{-i\theta}|\theta\rangle$, one finds

$$e^{i\alpha \widetilde{Q}_5} |\theta\rangle = |\theta + n_f \alpha\rangle. \tag{1.172}$$

Finally let us mention that the term given in the operator in Eq. (1.157) induces a neutron electric dipole moment. This can be seen schematically by performing an anomalous chiral rotation which removes the $\bar{\theta} G\widetilde{G}$ term and transfers CP violation into the quark mass sector. For illustration, consider two light flavours with a common mass m_q.

$$\begin{pmatrix} u \\ d \end{pmatrix} \to \exp\left[i\, \frac{\bar{\theta}\gamma_5}{2} \right] \begin{pmatrix} u \\ d \end{pmatrix} \tag{1.173}$$

To first order in $\bar{\theta}$, the mass term becomes

$$\mathcal{L}_{\text{mass}} \supset -m_q(\bar{u}u + \bar{d}d) - i m_q \bar{\theta}(\bar{u}\gamma_5 u + \bar{d}\gamma_5 d). \tag{1.174}$$

The second term is CP-violating and induces CP-odd hadronic interactions. These interactions in turn generate an electric dipole moment for the neutron, usually estimated as

$$d_n \sim 10^{-16}\, \bar{\theta}\, e\, \text{cm}, \tag{1.175}$$

up to an order-one coefficient depending on the hadronic calculation. Since experimentally d_n is extremely small, one obtains the bound

$$|\bar{\theta}| \lesssim 10^{-10}. \tag{1.176}$$

The strong CP problem is precisely the absence of an explanation, within ordinary QCD, for why this parameter should be so tiny.

Chapter 2

Solutions to the Strong CP Problem

After introducing the strong CP problem in the previous chapter, we now discuss some possible solutions. Let us first continue the simplified argument presented in the introduction of chapter 1. One possible way to forbid a neutron electric dipole moment is to require parity to be an exact symmetry of Nature. Under parity, polar vectors change sign, whereas axial vectors do not. Therefore,

$$\vec{x} \to -\vec{x}, \qquad \vec{d} \to -\vec{d}, \qquad \vec{s} \to \vec{s}, \tag{2.1}$$

where $\vec{s}$ and $\vec{d}$ denote, respectively, the spin and the electric dipole moment of the neutron.
If the neutron had a non-vanishing electric dipole moment, rotational symmetry would require $\vec{d}$ to be aligned or anti-aligned with the spin, namely $\hat{d} = \pm\hat{s}$. Under parity, because $\hat{d}$ is an axial vector, this configuration is mapped into $\hat{d} = \mp\hat{s}$. However, since there is only one neutron state, an exact parity symmetry would require the neutron to be mapped into itself. The only way for both configurations to describe the same physical state is therefore

$$\vec{d} = 0. \tag{2.2}$$

Thus, exact parity would solve the strong CP problem by forbidding the neutron electric dipole moment. However, this possibility is not realized in Nature: parity is violated in the Standard Model: the weak interaction acts only on left-handed fermion doublets; the minimal Standard Model contains no right-handed neutrino.
A similar argument can be made by imposing time-reversal invariance. Under time reversal one has

$$\vec{d} \to \vec{d}, \qquad \vec{s} \to -\vec{s}. \tag{2.3}$$

Hence a configuration with $\hat{d} = \pm\hat{s}$ is again mapped into one with $\hat{d} = \mp\hat{s}$, implying that an exact time-reversal symmetry would also force the neutron electric dipole moment to vanish. Since, by the CPT theorem, assuming CPT invariance, time reversal violation is equivalent to CP violation, this would remove the strong CP-violating term. Nevertheless, this is not a satisfactory explanation within the Standard Model, where CP violation is already present through the complex phase of the CKM matrix.

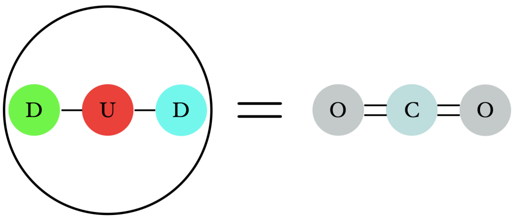

Figure 2.1: The CO_2 atom, the electrostatic potential is minimized with $\theta = 0$. Figure taken from Ref. [30]

Finally, one can consider the axion solution. The basic idea is to promote the θ angle to a dynamical field. Its effective potential is generated by QCD non-perturbative effects and is minimized at a CP-conserving value, dynamically driving the physical value of θ to zero. At a classical level, this can be understood by analogy with a system whose configuration is fixed by energy minimization. For instance, in a CO_2 molecule the negatively charged oxygen atoms align with the carbon atom because this configuration minimizes the electrostatic energy. In the axion case, the dynamics of the axion field plays an analogous role: the vacuum relaxes to the minimum of the potential, where the effective strong CP phase vanishes.
In the following we present in a more quantitative way the solutions to the Strong CP problem, focusing on the Peccei-Quinn solution and on axion models.

2.1 Some Possibilities

We now briefly present some alternatives to the Peccei–Quinn mechanism and to the axion solution.

2.1.1 Massless Up Quark

This solution was first proposed by 't Hooft in Ref. [28]. The simplest way to remove the strong CP phase is to assume that one quark, for instance the up quark, is massless. In this case one can perform a chiral rotation of the u-quark field. Since the left- and right-handed components are not coupled by a mass term, this transformation does not change the physical content of the theory, while it shifts the parameter $\bar{\theta}$. Therefore, one can always choose the chiral rotation in such a way that $\bar{\theta} = 0$
In this limit, $\bar{\theta}$ is not a physical parameter.
Although this solution is conceptually very simple, it is experimentally ruled out. The current value quoted by the Particle Data Group is

$$m_u = 2.16 \pm 0.07 \text{ MeV}, \tag{2.4}$$

which is clearly incompatible with a massless up quark. See Ref. [31].

2.1.2 Parity Symmetry

Another possible solution, already mentioned above, is based on parity. This idea was first discussed in Refs. [32, 33]. There are different realizations of parity solutions to the strong CP problem. Very schematically, Ref. [33] presents a possible ultraviolet completion in which the weak gauge sector is doubled. One introduces a right-handed weak group $SU(2)_R$, possibly accompanied by a $U(1)_{Y_R}$, and a scalar doublet H_R, in analogy with the Standard Model Higgs doublet H_L. The color gauge group, however, remains common to the two sectors and is still $SU(3)_c$. Under parity, the two sectors are exchanged according to

$$SU(2)_L \longleftrightarrow SU(2)_R, \qquad Q_L \longleftrightarrow Q_R^\dagger, \qquad H_L \longleftrightarrow H_R^\dagger, \qquad L_L \longleftrightarrow L_R^\dagger. \tag{2.5}$$

Parity invariance implies that the Yukawa matrices are Hermitian. Moreover, because there are two weak gauge groups, the vacuum expectation values of both H_L and H_R can be chosen to be real. As a consequence, the physical strong CP phase vanishes. Indeed, the bare QCD angle θ, associated with the common $SU(3)_c$ gauge group, is forbidden by parity, while the contribution coming from the quark mass matrices is real:

$$\arg\det M_q = 0. \tag{2.6}$$

Thus one obtains at tree level:

$$\bar{\theta} = \theta - \arg\det M_q = 0 \tag{2.7}$$

Parity must eventually be spontaneously broken in order to recover the observed chiral structure of the Standard Model at low energies. This can be achieved by assuming

$$\langle H_R \rangle = \lambda \langle H_L \rangle, \qquad \lambda \neq 1. \tag{2.8}$$

In this picture, the ordinary Standard Model sector has the observed masses and interactions, while the mirror sector has masses rescaled by the real factor λ. In addition, weak interactions act on right-handed fields in the mirror sector, rather than on left-handed fields. This scaling property follows from the fact that the Higgs vacuum expectation value is the only fundamental mass scale in the electroweak sector of the Standard Model.
At tree level, these models give $\bar{\theta} = 0$. However, radiative corrections can regenerate a non-zero value. It was shown in Ref. [34] that, in the Standard Model, finite contributions to $\bar{\theta}$ arise only at three loops. These corrections are related to the renormalization of the quark mass matrices. More generally, the running of $\bar{\theta}$ can be studied by treating it as the coefficient of an effective CP-odd coupling of the form $\bar{f}\gamma_5 f$. If a high-energy mechanism sets $\bar{\theta}$ to zero in the ultraviolet, the resulting Standard Model contribution provides a useful estimate of the induced value at low energies. Parity models provide precisely such a mechanism. As argued in Ref. [33], in these models, the additional corrections generated by the renormalization-group flow are even smaller.
A minimal realization of these models can be described by the effective Yukawa Lagrangian

$$\mathcal{L}_Y = -y_u \frac{Q_L H_L Q_R H_R}{\Lambda_u} - y_d \frac{Q_L H_L^\dagger Q_R H_R^\dagger}{\Lambda_d} + \text{h.c.} \tag{2.9}$$

Here Q_R denotes the right-handed Standard Model quarks, while H_R is the $SU(2)_R$ doublet introduced above. After H_R acquires a vacuum expectation value v_R, the effective Standard Model Yukawa matrices are

$$Y_u = y_u \frac{v_R}{\Lambda_u}, \qquad Y_d = y_d \frac{v_R}{\Lambda_d}. \tag{2.10}$$

By parity, the matrices y_u and y_d are Hermitian, and therefore the effective Yukawa matrices are Hermitian as well. Their determinants are then real, so that the same argument applies: the contribution $\arg\det M_q$ vanishes, and one obtains $\bar{\theta} = 0$ at tree level.

2.1.3 CP Symmetry

These models were first introduced by Nelson in Ref. [35] and were later generalized by Barr in Ref. [36]. They were originally proposed in the context of grand unified theories, where the gauge interactions are embedded into a single larger gauge group. In the ultraviolet, CP is assumed to be an exact symmetry. Under suitable assumptions, in particular on the allowed Yukawa couplings between fermions belonging to different representations of the large gauge group, one can again get $\arg\det M = 0$. Since CP forbids the bare QCD vacuum angle, this implies

$$\bar{\theta} = \theta - \arg\det M = 0 \tag{2.11}$$

at tree level. At the same time, CP can be spontaneously broken, allowing for a non-trivial complex phase in the CKM matrix. A simple toy model illustrating this mechanism is described by the Lagrangian

$$\mathcal{L} = \mu q q^c + Y_{ij} H Q_i d_j^c + A^{ia} \eta^a d_i^c q + \text{h.c.} \tag{2.12}$$

Here q and q^c are additional vector-like quarks, while η^a denotes a set of complex scalar fields whose vacuum expectation values can be complex. The corresponding down-type mass matrix has the schematic form

$$M = \begin{pmatrix} \mu & A\langle\eta\rangle \\ 0 & m_d \end{pmatrix}. \tag{2.13}$$

Although this matrix contains complex entries, its determinant is simply

$$\det M = \mu \det m_d. \tag{2.14}$$

Therefore, if μ and $\det m_d$ are real, as enforced by the original CP symmetry, one obtains $\arg\det M = 0$. This argument is exact at tree level. Radiative corrections can in general regenerate a non-zero value of $\bar{\theta}$, but with suitable model-building assumptions it can remain sufficiently small. Compared with parity-based solutions, CP-based solutions are typically more fragile. In order to obtain a large

CKM phase while keeping $\bar{\theta}$ small, one often needs non-trivial hierarchies among scales and a restricted structure of the allowed couplings. Otherwise, dangerous contributions to $\bar{\theta}$ may already appear at tree level. Furthermore, it has been noted that spontaneous symmetry breaking of CP symmetry in the early Universe, leads to domains with different CP values, separated by domain walls carrying a large energy density. Since this energy dissipates very slowly, for VEVs of the order of the weak scale the present-day domain-wall energy density would exceed the critical density of the Universe, see Ref. [37]. Many different realizations of this idea have been proposed in the literature; see, for instance, Refs. [38, 39, 40, 41].

2.2 The Peccei–Quinn Solution

The Peccei–Quinn solution provides a dynamical mechanism by which the physical parameter $\bar{\theta}$ is driven to zero in the true vacuum of the theory. In Refs. [42, 43], Peccei and Quinn considered models in which at least one fermion obtains its mass from the vacuum expectation value of a color-singlet scalar field. After integrating out the gauge fields and the fermions, the effective potential for the scalar field is such that, for a suitable range of parameters, its minimum corresponds to $\bar{\theta} = 0$. We first illustrate the mechanism for a single-flavour toy model; the generalization to more flavours is straightforward. The Lagrangian is

$$\mathcal{L} = -\frac{1}{4}F^a_{\mu\nu}F^{\mu\nu}_a + i\bar{\psi}\not{D}\psi + \bar{\psi}\left(G\varphi\frac{1+\gamma_5}{2} + G^*\varphi^*\frac{1-\gamma_5}{2}\right)\psi - |\partial_\mu\varphi|^2 - \mu^2|\varphi|^2 - h|\varphi|^4. \tag{2.15}$$

This Lagrangian is formally invariant under a new assumed symmetry, the Peccei–Quinn symmetry $U(1)_{\mathrm{PQ}}$,

$$\psi \longrightarrow e^{i\omega\gamma_5}\psi, \qquad \varphi \longrightarrow e^{-2i\omega}\varphi. \tag{2.16}$$

However, the corresponding chiral current is anomalous. Under a Peccei–Quinn transformation, the fermionic measure in the path integral acquires a non-trivial Jacobian, which produces the shift

$$\delta S_{\text{eff}} = -i\,2\omega\,\frac{g^2}{32\pi^2}\int d^4x\, F^a_{\mu\nu}\widetilde{F}^{a\mu\nu}. \tag{2.17}$$

Equivalently, the θ parameter is redefined as $\theta \longrightarrow \theta - 2\omega$. This shows that different values of θ are related by a field redefinition. In the case in which the fermion is massless, or in which the mass term transforms in such a way as to remain invariant under the chiral rotation (that is the general case in which we are working in) θ is not a physical parameter.

We still have to show, however, that the theory is CP conserving in the true vacuum of the scalar φ. If $\mu^2 > 0$, then $\langle\varphi\rangle = 0$, and the quark remains massless; in this case one recovers the massless-quark solution discussed above, and the theory is CP invariant. In the more general case, one must require that, after setting $\theta = 0$, the fermion mass term is real.

To see this more explicitly, let us start from the path-integral formulation. As discussed in Chapter (1), the generating functional in the θ vacuum can be written as

$$Z[J,J^*] = \sum_q e^{iq\theta}\int (dA_\mu)_q \int d\psi\, d\bar{\psi}\int d\varphi\, d\varphi^*\ \exp\left\{-\int d^4x\ \left[\mathcal{L}(A_\mu,\psi,\varphi) + J^*\varphi^* + J\varphi\right]\right\}. \tag{2.18}$$

The integer q labels the topological sector over which the gauge fields are integrated. The scalar vacuum expectation value is defined by

$$\langle\varphi\rangle = \frac{1}{Z[J,J^*]}\frac{\delta Z[J,J^*]}{\delta J}\bigg|_{J=J^*=0} = \lambda e^{i\beta}. \tag{2.19}$$

We now integrate over the fermionic and gauge fields. The index theorem tells us that the topological charge q counts the difference between the numbers of right- and left-handed zero modes of the Dirac operator. As seen in (1.3.2) for $q > 0$ the zero modes are purely right-handed, while for $q < 0$ they

are purely left-handed . The non-zero modes, instead, come in pairs of opposite chirality.
At fixed gauge-field background, the fermionic integration gives a determinant. In general in fact,

$$\int \mathcal{D}\bar{\psi}\,\mathcal{D}\psi\ \exp\left\{-\int d^4x\,\bar{\psi}M(x)\psi\right\} = \det M. \tag{2.20}$$

In the present case,

$$M(x) = i\not{D} + G\varphi\frac{1+\gamma_5}{2} + G^*\varphi^*\frac{1-\gamma_5}{2}. \tag{2.21}$$

In order to use the index theorem, it is convenient to expand the fermion fields in eigenfunctions of $i\not{D}$ rather than directly in eigenfunctions of $M(x)$, and to work in Euclidean signature. The determinant then separates into a zero-mode contribution and a non-zero-mode contribution. Let us first consider the sector $q = m > 0$. If ϕ_n are the right-handed zero modes and $\psi = \sum_n c_n\phi_n$, then, restricting to the zero-mode subspace,

$$\begin{aligned}\int d^4x\,\mathcal{L} &= \int d^4x\,\bar{\psi}\left(G\varphi\frac{1+\gamma_5}{2} + G^*\varphi^*\frac{1-\gamma_5}{2}\right)\psi = \\ &= \sum_{n,r=1}^{m}\bar{c}_r c_n\int d^4x\,G\varphi\,\bar{\phi}_r\phi_n = \sum_{n=1}^{m}\bar{c}_n c_n\,G\varphi,\end{aligned} \tag{2.22}$$

where in the last step we used the orthonormality of the zero modes and assumed φ to be approximately constant. Therefore, in the zero-mode subspace, the determinant gives a factor $(G\varphi)^m$. Similarly, in the sector $q = -m < 0$, the zero modes are left-handed and one obtains $(G^*\varphi^*)^m$.
Let us now consider the non-zero modes. Since these modes do not have definite chirality, it is useful to work in the two-dimensional subspace spanned by ϕ_n and $\gamma_5\phi_n$. If $i\not{D}\phi_n = \lambda_n\phi_n$, then $i\not{D}\gamma_5\phi_n = -\lambda_n\gamma_5\phi_n$. In this basis the operator M takes the form

$$M_n = \begin{pmatrix}\lambda_n + \dfrac{G\varphi + G^*\varphi^*}{2} & \dfrac{G\varphi - G^*\varphi^*}{2} \\ \dfrac{G\varphi - G^*\varphi^*}{2} & -\lambda_n + \dfrac{G\varphi + G^*\varphi^*}{2}\end{pmatrix}. \tag{2.23}$$

Its determinant depends only on the modulus of the scalar field: $\det M_n = |G\varphi|^2 - \lambda_n^2$. Thus the determinant over the non-zero modes has the schematic form

$$\prod_{\lambda_n\neq 0}\left(|G\varphi|^2 - \lambda_n^2\right), \tag{2.24}$$

and depends only on $|\varphi|$. This property is preserved after integrating over the gauge fields. We denote by $A_m(|\varphi|^2)$ the contribution from the non-zero modes and the gauge-field integration in the sector of topological charge m. The coefficients in opposite topological sectors are related by CP. Indeed, schematically:

$$Z_q = \int_q \mathcal{D}\Phi\,e^{-S[\Phi]} = \int_{-q}\mathcal{D}\Phi^{CP}\,e^{-S[\Phi^{CP}]} = \int_{-q}\mathcal{D}\Phi\,e^{-S[\Phi]^*} = Z^*_{-q}, \tag{2.25}$$

where $\Phi = \{A_\mu, \bar{\psi}, \psi, \varphi\}$ and the CP transformation maps the sector q into the sector $-q$. We used the fact that, in the absence of the θ term, QCD is CP invariant, while the scalar field transforms as $\varphi \to \varphi^*$. It follows that

$$A^*_m = A_{-m}. \tag{2.26}$$

Therefore the generating functional can be written schematically as

$$Z[J,J^*] = \int \mathcal{D}\varphi\,\mathcal{D}\varphi^*\left[A_0(|\varphi|^2) + \sum_{n>0}A_n(|\varphi|^2)\left(G\varphi e^{i\theta}\right)^n + A^*_n(|\varphi|^2)\left(G^*\varphi^* e^{-i\theta}\right)^n\right]. \tag{2.27}$$

Here A_0 denotes the perturbative contribution. More generally, when the scalar field depends on spacetime, (in fact in some steps we assumed φ to be constant) the coefficients A_n are non-local functionals of $\varphi\varphi^*$, schematically of the form

$$A_n[\varphi\varphi^*] = \sum_m \prod_{i=1}^{m} \int d^4x_i\, d^4y_i\, \varphi(x_i)\varphi^*(y_i)\, c_n^m(x_i, y_i). \tag{2.28}$$

We now define $\alpha = \arg\left(G\langle\varphi\rangle e^{i\theta}\right)$ and expand the scalar field around its vacuum expectation value as

$$\varphi = e^{i\beta}\left(\lambda + \rho + i\sigma\right). \tag{2.29}$$

Then

$$A_n(|\varphi|^2)\left(G\varphi e^{i\theta}\right)^n = A_n(|\varphi|^2)e^{in\alpha}|G|^n\left(\lambda + \rho + i\sigma\right)^n. \tag{2.30}$$

As a result, Eq. (2.27) can be rewritten in the schematic form

$$Z[J, J^*] = \int d\rho\, d\sigma\, \left[A_0(\rho, \sigma^2) + \sum_{n>0}\left(F_n(\rho, \sigma^2)\cos n\alpha - \sigma\, G_n(\rho, \sigma^2)\sin n\alpha\right)\right]. \tag{2.31}$$

Here F_n and σG_n denote, respectively, the real and imaginary parts of $A_n|G|^n(\lambda + \rho + i\sigma)^n$. Their dependence on σ is fixed by the fact that A_n depends only on $|\varphi|^2$: the real part is even in σ, while the imaginary part is odd in σ. Therefore all terms odd in σ vanish after integration over σ. By definition of the vacuum, one must impose

$$\begin{aligned} \langle\rho(\alpha)\rangle &= \int d\rho\, d\sigma\, \rho\left[A_0(\rho, \sigma^2) + \sum_{n>0} F_n(\rho, \sigma^2)\cos n\alpha\right] = 0, \\ \langle\sigma(\alpha)\rangle &= -\int d\rho\, d\sigma\, \sigma^2 \sum_{n>0} G_n(\rho, \sigma^2)\sin n\alpha = 0. \end{aligned} \tag{2.32}$$

If, for a certain value of α, the first condition is not satisfied, one can redefine $\lambda(\alpha)$ so that $\langle\rho(\alpha)\rangle = 0$. The same argument cannot be applied to the second condition, because α and β are not independent. In general, the second condition is satisfied only for $\alpha = 0, \pi$.

The true vacuum is determined by the effective potential. At linear order in G, its form is:

$$V_{\text{eff}}(\varphi) = \mu^2|\varphi|^2 + h|\varphi|^4 - K|G\varphi|\cos\alpha, \tag{2.33}$$

where the last term is generated by integrating out the fermionic and gauge fields. The coefficient K can be written schematically as

$$K = \frac{\int (dA_\mu)_1\, d\psi\, d\bar{\psi} \int d^4x\, \bar{\psi}(x)\frac{1+\gamma_5}{2}\psi(x)\, \exp\left[-\int d^4x'\left(-\frac{1}{4}F^a_{\mu\nu}F_a^{\mu\nu} + i\bar{\psi}\not{D}\psi\right)\right]}{\int (dA_\mu)_0\, d\psi\, d\bar{\psi}\, \exp\left[-\int d^4x'\left(-\frac{1}{4}F^a_{\mu\nu}F_a^{\mu\nu} + i\bar{\psi}\not{D}\psi\right)\right]}, \tag{2.34}$$

and is real and positive. We have found out that the effective potential has no $\mathrm{U}_{PQ}(1)$ symmetry and it explains for example the lack of a goldstone boson after (for real QCD the lack of the η pseudo-Goldstone boson).

Although the expression above was derived only to linear order in G, it is sufficient to show how the mechanism works. For a suitable range of parameters, the minimum is obtained at $\alpha = 0$ (note in fact that α depends on $\langle\varphi\rangle$, and so it depends on the parameters of the potential). More generally, the Vafa–Witten theorem (see Ref. [44]) guarantees (actually, under certain assumptions, satisfied by QCD) that the minimum of the potential is located at the CP-conserving value $\alpha = 0$. From the definition of α, it is clear that if $\alpha = 0$, one can perform a chiral rotation of the fermion fields such that the mass term generated by $\langle\varphi\rangle$ becomes real, while at the same time the vacuum angle is shifted as $\theta \longrightarrow \theta - \theta = 0$. Thus the Peccei–Quinn mechanism provides a dynamical relaxation of the strong CP phase to zero. We now turn to an important prediction of this mechanism: the existence of a new pseudoscalar particle, the axion.

2.3 The Axion Model

We have seen that after integrating out the fermionic fields, QCD non-perturbative effects generate an effective potential for φ that in particular depends on the phase of its vacuum expectation value (Eq. (2.33)) and that explicitly breaks $U(1)_{\rm PQ}$. Since the Peccei-Quinn symmetry is spontaneously broken, we expect a pseudo-Nambu–Goldstone boson associated with its breaking. This particle is the axion.
From the derivation of the previous section, one can then already identify the axion field, the axion is simply the angular part of φ, or equivalently what we called in the previous discussion σ, renormalized by an energy scale f_{PQ} which dictates the scale of the symmetry breaking of the complex scalar field φ and has no associated potential. It is mostly a matter of convention whether one calls "axion" the full angular phase of φ, or only the fluctuation around its vacuum expectation value. We start with the first convention and later redefine the axion as the fluctuation around the minimum.
However, the Peccei–Quinn mechanism can be reformulated as follows. One extends the Standard Model by adding new Yukawa interactions, involving additional scalars singlets and fermions, in such a way that the full theory is invariant under a global $U(1)_{\rm PQ}$ symmetry. When this symmetry is spontaneously broken, the axion appears as the corresponding Nambu–Goldstone boson. Its transformation law then must be:

$$a(x) \longrightarrow a(x) + \alpha f_a, \tag{2.35}$$

where f_a is the order parameter of the symmetry breaking and is called the axion decay constant. Note that this procedure is quite general, and therefore in the definition of the axion field we do not want to stick just to the model presented in the previous section. More generally, the relevant terms of our theory become:

$$\mathcal{L}_{\rm tot} = \mathcal{L}_{\rm SM} + \bar{\theta}\,\frac{g^2}{32\pi^2} G^a_{\mu\nu}\widetilde{G}^{\mu\nu}_a - \frac{1}{2}\partial_\mu a\,\partial^\mu a + \mathcal{L}_{\rm int}(\partial_\mu a, \Psi) + \xi\,\frac{a}{f_a}\frac{g^2}{32\pi^2} G^a_{\mu\nu}\widetilde{G}^{\mu\nu}_a. \tag{2.36}$$

Here ξ is the anomaly coefficient. We give motivation of Eq. (2.36). The first two terms are just the usual SM Lagrangian, with the addition of the $\bar{\theta}$ term. Then we have the usual kinetic term (that for example arises from the kinetic term of φ) then we add an interaction term with fermions, with a derivative coupling given the Goldstone boson nature of the axion, and finally we have the last term that couples the axion with gluons. This term is necessary, since under a $U(1)_{\rm PQ}$ transformation, the anomalous variation of the fermionic measure shifts the vacuum angle as

$$\bar{\theta} \longrightarrow \bar{\theta} - \xi\alpha, \tag{2.37}$$

while the axion shifts as

$$a(x) \longrightarrow a(x) + \alpha f_a. \tag{2.38}$$

Therefore the combination

$$\bar{\theta}_{\rm eff} = \bar{\theta} + \xi\frac{a}{f_a} \tag{2.39}$$

is invariant, and the Lagrangian in Eq. (2.36) is left unchanged, as it should, given that before considering the effective potential corrections, our theory must be PQ invariant. The Peccei–Quinn mechanism states that the axion dynamically relaxes the effective strong CP phase to zero (this follows from the definition of α in the previous section) so:

$$\langle\bar{\theta}_{\rm eff}\rangle = \bar{\theta} + \xi\frac{\langle a\rangle}{f_a} = 0, \tag{2.40}$$

or equivalently $\langle a\rangle = -\bar{\theta}\frac{f_a}{\xi}$. This is the same condition denoted by $\alpha = 0$ in the previous section. The same conclusion follows from the θ-dependence of the vacuum energy, discussed in Eq. (1.165): the effect of the axion, as easily seen in Eq. (2.36) is to simply replace:

$$\bar{\theta} \longrightarrow \bar{\theta} + \xi\frac{a}{f_a}. \tag{2.41}$$

The cosine structure obtained in the one-instanton approximation then confirms that the minimum is reached for $\langle a \rangle = -\bar{\theta} f_a/\xi$. We now define $\bar{a}$ as the fluctuation around this vacuum expectation value. The axion mass is obtained from the second derivative of the vacuum energy with respect to the axion field. Equivalently,

$$m_a^2 = \frac{\delta^2}{\delta \bar{a}^2} \log Z(\bar{\theta}) \bigg|_{\bar{a}=0} = \frac{\xi^2}{f_a^2} \frac{\delta^2}{\delta \bar{\theta}^2} \log Z(\bar{\theta}) \bigg|_{\bar{\theta}=0} = \frac{\xi^2}{f_a^2} \chi_{\text{top}}. \tag{2.42}$$

In the last step we used the fact that the axion potential is given by the QCD vacuum energy density $\mathcal{E}(\bar{\theta})$, evaluated after the replacement of Eq. (2.41). The QCD topological susceptibility is defined as

$$\chi_{\text{top}} = \int d^4x \left\langle \frac{g^2}{32\pi^2} G^a_{\mu\nu} \widetilde{G}^{a\mu\nu}(x) \frac{g^2}{32\pi^2} G^b_{\rho\sigma} \widetilde{G}^{b\rho\sigma}(0) \right\rangle. \tag{2.43}$$

This expression follows by treating $\bar{\theta}$ as a source for the topological charge density. Using translational invariance, one has

$$\frac{1}{V} \int d^4x\, d^4y\, G\widetilde{G}(x)\, G\widetilde{G}(y) = \int d^4z\, G\widetilde{G}(z)\, G\widetilde{G}(0), \tag{2.44}$$

where $z = x - y$. The quantity χ_{top} is particularly useful because it can be computed using chiral Lagrangian techniques as well as non-perturbative lattice methods. Most importantly, it fixes the characteristic scaling of the axion mass with the decay constant:

$$m_a^2 f_a^2 = \xi^2 \chi_{\text{top}}. \tag{2.45}$$

Equivalently, if one absorbs the anomaly coefficient into the definition of f_a, this relation is often written as

$$m_a^2 f_a^2 = \chi_{\text{top}}. \tag{2.46}$$

We can find the axion mass also using the chiral Lagrangian, again using Eq. (2.41) in Eq. (1.58), one finds

$$V = -m_\pi^2 f_\pi^2 \sqrt{1 - \frac{4 m_u m_d}{(m_u + m_d)^2} \sin^2\left(\frac{\bar{\theta}}{2} + \frac{\xi a}{2 f_a}\right)}. \tag{2.47}$$

As expected $\langle a \rangle = -\bar{\theta} \frac{f_a}{\xi}$. Expanding the potential around the minimum gives

$$m_a^2 = \frac{\xi^2}{f_a^2} \frac{m_u m_d}{(m_u + m_d)^2} m_\pi^2 f_\pi^2. \tag{2.48}$$

If the anomaly coefficient is absorbed into the definition of the axion decay constant, the standard expression becomes

$$m_a = \frac{\sqrt{m_u m_d}}{m_u + m_d} \frac{m_\pi f_\pi}{f_a} \simeq 5.7\, \mu\text{eV} \left(\frac{10^{12}\,\text{GeV}}{f_a}\right). \tag{2.49}$$

We will return to this expression in the next chapter, where we analyze the experimental bounds on axions. We now move to a general issue associated with axion models.

2.3.1 The Axion Quality Problem

The axion quality problem arises from the fact that the Peccei–Quinn symmetry is not an exact symmetry of the theory, but an anomalous global symmetry. In other words, the strong CP problem is solved by a dynamical mechanism which relies on a symmetry that is expected to be explicitly broken. This lack of an exact symmetry leads to two related issues.

First, from an effective field theory point of view, once our Lagrangian contains a coupling of the form $(\bar{\theta} + a/f_a)G\widetilde{G}$, one should in principle include all other operators compatible with the exact symmetries of the theory. If no exact symmetry forbids them, Peccei–Quinn-violating operators are generically allowed.

Second, quantum gravity is widely believed to break global symmetries. This is easily understood by

considering black-holes. Therefore, even if a global Peccei–Quinn symmetry is imposed at the classical level, Planck-scale effects can explicitly break it. In both cases, the axion then receives an additional contribution to its potential, whose minimum is not in general aligned with the CP-conserving point. This could reintroduce the strong CP problem.
For example, Ref. [45] found that gravitationally induced operators which explicitly violate the Peccei–Quinn symmetry must have dimension ten or higher (the current bound is dimension 14 [1]) in order for the axion to provide a natural solution to the strong CP problem. Thus we have to gauge PQ symmetry, or assume some additional exact symmetry, to protect the PQ solution. We just mention that gravitational corrections may also have implications for the axion domain-wall problem. The effect of gravity on the Peccei–Quinn solution can be illustrated by considering the effective Lagrangian

$$\mathcal{L} = |\partial_\mu\phi|^2 - \lambda\left(|\phi|^2 - f_a^2\right)^2 + \Lambda^4\cos\frac{a}{f_a} + \frac{|\phi|^{n+3}\left(ge^{-i\delta}\phi + ge^{i\delta}\phi^*\right)}{2M_{\rm Pl}^n}. \tag{2.50}$$

Here

$$\phi = (f_a + \widetilde{\phi})e^{ia/f_a}, \qquad \langle\widetilde{\phi}\rangle = 0, \tag{2.51}$$

and $\Lambda \simeq 100$ MeV parametrizes the QCD contribution to the axion potential. The last term is a dimension-$(n+4)$ operator and represents the effect of quantum gravity. In the absence of this term, the minimum would be located at $a = 0$. The phase δ is defined relative to the phase of the quark mass matrix and has no reason to be small. For small a, the potential can be expanded as

$$V(a) = \frac{1}{2}\Lambda^4\frac{a^2}{f_a^2} + \frac{1}{2}g\frac{f_a^{n+4}}{M_{\rm Pl}^n}\left[\left(\frac{a}{f_a}\right)^2\cos\delta - 2\frac{a}{f_a}\sin\delta\right]. \tag{2.52}$$

Defining

$$m_*^2 = g\frac{f_a^{n+2}}{M_{\rm Pl}^n}\cos\delta, \qquad m_{\rm phys}^2 = m_0^2 + m_*^2, \qquad m_0^2 = \frac{\Lambda^4}{f_a^2}, \tag{2.53}$$

one obtains

$$V(a) = \frac{1}{2}m_{\rm phys}^2 a^2 - m_*^2 f_a a\tan\delta. \tag{2.54}$$

The minimum is therefore shifted away from the CP-conserving value, giving

$$\langle\bar{\theta}\rangle = \frac{m_*^2\tan\delta}{m_0^2 + m_*^2}. \tag{2.55}$$

Requiring the induced value of $\bar{\theta}$ to satisfy the experimental bound then implies roughly

$$\frac{m_*^2}{m_0^2} < 10^{-9} \Rightarrow \frac{gf_a^{n+4}}{\Lambda^4 M_{\rm Pl}^n} < 10^{-9}, \tag{2.56}$$

Using $M_{\rm Pl} \simeq 10^{19}$ GeV, and assuming $\delta = \mathcal{O}(1)$, this condition can be written as

$$\log_{10} f_a \lesssim \frac{19n - 13 - \log_{10} g}{n+4}. \tag{2.57}$$

Combining this with a typical bound $f_a \gtrsim 10^{10}$ GeV, one finds that for natural couplings, $g \sim 1$, one must require $n \geq 6$. Therefore the leading PQ-violating operators must have dimension at least ten. Smaller values of n are possible only if the coefficient g is highly suppressed, for example $|g| \lesssim 10^{-8}$ for $n = 5$ and $|g| \lesssim 10^{-43}$ for $n = 1$.
Similar, and in some cases even stronger, conclusions were found in Ref. [46, 45]. Furthermore, the authors of Ref. [46] stress that it is not consistent to assume that all low-dimensional PQ-violating operators are absent while only a very high-dimensional operator has an $\mathcal{O}(1)$ coefficient, unless an additional symmetry forbids the lower-order terms. The reason is that, once a PQ-violating operator such as $|\phi|^{2m}\phi^n$ is present with an unsuppressed coupling, loop contractions can generate lower-dimensional PQ-violating operators. The inverse powers of $M_{\rm Pl}$ in the original operator are then compensated by

ultraviolet-divergent loop integrals cut off at the Planck scale, so the resulting lower-dimensional operators are not additionally suppressed. More generally, two operators violating the global charge by n and n' can induce an operator violating it by $|n - n'|$. Therefore Planck-scale PQ-violation can be consistently suppressed only if all dangerous lower-dimensional operators are exactly forbidden, for instance by gauge or discrete symmetries, or if all symmetry-violating couplings are themselves extremely small.

Several mechanisms have been proposed to address the axion quality problem. They can be broadly grouped into models in which $U(1)_{\mathrm{PQ}}$ arises as an accidental symmetry, models in which it originates from higher-dimensional gauge symmetries, and constructions motivated by string and more broadly extra dimensional theories. Here we focus on the first two classes. Although these constructions provide possible ultraviolet completions of the Peccei–Quinn mechanism, they are generally much less economical than the low-energy axion effective field theory. In this sense, the axion quality problem represents the hidden complication behind the otherwise elegant low-energy solution.
An example of an accidental Peccei–Quinn symmetry is given in Ref. [47]. The construction introduces four sets of quarks $Q_{1,2,3,4}$ and two new confining gauge groups, $SU(N)$ and $SU(M)$. With a suitable choice of representations, the groups $SU(N) \times SU(M)$ confine and produce a new pseudo-Nambu–Goldstone boson which plays the role of the axion. If M is chosen sufficiently large, higher-dimensional PQ-violating operators are suppressed enough for the axion to continue solving the strong CP problem.
The extra-dimensional approach addresses the problem in a different way. The idea is to promote the axion shift symmetry to a gauge symmetry in the ultraviolet, so that it is protected against Planck-scale corrections. This can be realized, for example, by identifying the axion with the fifth component A_5 of a higher-dimensional gauge field. At high energies, the axion is then protected by five-dimensional gauge invariance. Another possibility is to gauge $U(1)_{\mathrm{PQ}}$ in the bulk. Finally, these five-dimensional solutions can often be dimensionally deconstructed into purely four-dimensional models, although the higher-dimensional picture provides the simplest "UV" motivation for their special structure.

2.3.2 Axion Models

We now present some explicit axion models, starting with the Peccei–Quinn–Weinberg–Wilczek model. In this original realization, the Peccei–Quinn symmetry is broken at the electroweak scale. As a consequence, the axion is relatively heavy and strongly coupled, and in fact this model was ruled out long ago, see Ref. [48] for a modern review.
We then introduce a broader class of models known as *invisible axion* models. In these constructions, a new energy scale $f_{\mathrm{PQ}} \gg v$ is introduced, where v is the electroweak scale. Since the axion mass and its couplings scale inversely with f_{PQ}, the axion becomes lighter and much more weakly coupled. This new scale is often associated with very high-energy physics, for instance with the grand-unification scale.
The two main classes of invisible axion models are the KSVZ (Kim–Shifman–Vainshtein–Zakharov) model and the DFSZ (Dine–Fischler–Srednicki–Zhitnitsky) model. In KSVZ models one introduces new heavy exotic quarks, together with a complex scalar field whose angular excitation is identified with the axion. In DFSZ models, instead, no exotic quarks are required: the Standard Model scalar sector is enlarged by introducing two Higgs doublets, together with the same kind of complex scalar field. In both cases, the vacuum expectation value of this scalar field sets the Peccei–Quinn scale f_{PQ}.
A common feature of axion models is the structure of the Yukawa sector. After Peccei–Quinn symmetry breaking, Yukawa interactions typically induce couplings between the axion and fermions. These couplings can be written in two equivalent forms. The first form is obtained by expanding the scalar field around its vacuum expectation value, for example

$$\Phi = f_{\mathrm{PQ}} e^{ia/f_{\mathrm{PQ}}}. \tag{2.58}$$

This gives non-derivative pseudoscalar interactions such as

$$\mathcal{L}_{\mathrm{int}} = -i\frac{m}{f_{\mathrm{PQ}}} a\,\bar{\psi}\gamma_5\psi + \frac{m}{2f_{\mathrm{PQ}}^2} a^2\,\bar{\psi}\psi + \cdots\,. \tag{2.59}$$

Equivalently, one can remove the axion phase from the mass term by a chiral field redefinition. Starting from

$$m\bar{\psi}e^{i\gamma_5 a/f_{\mathrm{PQ}}}\psi, \tag{2.60}$$

one performs the rotation $\psi \longrightarrow e^{-i\gamma_5 a/(2f_{\mathrm{PQ}})}\psi$. The axion dependence is then shifted from the mass term to the fermion kinetic term, giving the derivative interaction

$$\mathcal{L}_{\mathrm{int}} = \frac{1}{2f_{\mathrm{PQ}}}\bar{\psi}\gamma_\mu\gamma_5\psi\,\partial^\mu a. \tag{2.61}$$

The two descriptions are physically equivalent. The non-derivative form is often useful for perturbative expansions in powers of a/f_{PQ}, while the derivative form makes the pseudo-Nambu–Goldstone nature of the axion more transparent.

The PQWW Axion Model

If all Standard Model fermions carry PQ charge, one has to introduce two Higgs doublets. This is necessary in order to absorb independent chiral transformations of the up- and down-type quarks, and of the charged leptons. The Yukawa sector is

$$\mathcal{L}_{\mathrm{Yukawa}} = Y^u_{ij}\bar{Q}_{Li}\Phi_1 u_{Rj} + Y^d_{ij}\bar{Q}_{Li}\Phi_2 d_{Rj} + Y^\ell_{ij}\bar{L}_{Li}\Phi_1\ell_{Rj} + \mathrm{h.c.} \tag{2.62}$$

Here Q_L and L_L are the left-handed quark and lepton $SU(2)_L$ doublets, while Φ_1 and Φ_2 are two scalar $SU(2)_L$ doublets. Their hypercharges are fixed by gauge invariance:

$$Y_{\Phi_1} = -\frac{1}{2}, \qquad Y_{\Phi_2} = +\frac{1}{2}. \tag{2.63}$$

Using $Q = T_3 + Y$, one obtains the component decomposition

$$\Phi_1 = \begin{pmatrix}\phi_0\\ \phi_-\end{pmatrix}, \qquad \Phi_2 = \begin{pmatrix}\phi'_+\\ \phi'_0\end{pmatrix}. \tag{2.64}$$

The vacuum expectation values are chosen so as to generate fermion masses,

$$\langle\Phi_1\rangle = \begin{pmatrix}v_1\\ 0\end{pmatrix}, \qquad \langle\Phi_2\rangle = \begin{pmatrix}0\\ v_2\end{pmatrix}, \qquad v_F = \sqrt{v_1^2 + v_2^2}. \tag{2.65}$$

Most excitations around these vacua can be written as $SU(2)_L$ gauge transformations. The axion is the neutral phase excitation orthogonal to the Goldstone mode eaten by the electroweak gauge bosons. After gauging away the unphysical excitations, the scalar doublets can be parametrized as

$$\Phi_1 = e^{iax/v_F}\begin{pmatrix}v_1\\ 0\end{pmatrix}, \qquad \Phi_2 = e^{ia/(xv_F)}\begin{pmatrix}0\\ v_2\end{pmatrix}, \tag{2.66}$$

where $x = \frac{v_2}{v_1}$. Since under $U(1)_{\mathrm{PQ}}$ the axion shifts as $a \longrightarrow a + \alpha v_F$, and choosing a basis in which the left-handed fermions are neutral under $U(1)_{\mathrm{PQ}}$, the right-handed quarks transform as

$$u_{Rj} \longrightarrow e^{-i\alpha x}u_{Rj}, \qquad d_{Rj} \longrightarrow e^{-i\alpha/x}d_{Rj}. \tag{2.67}$$

The charged leptons transform analogously to the up-type quarks in the Yukawa assignment written above. The Peccei–Quinn current can be obtained in the usual way from Noether's theorem for internal symmetries:

$$J^\mu_{\mathrm{PQ}} = \sum_\varphi \frac{\partial\mathcal{L}}{\partial(\partial_\mu\varphi)}\delta\varphi. \tag{2.68}$$

In the present basis it takes the schematic form

$$J^\mu_{\mathrm{PQ}} = -v_F\partial^\mu a + x\sum_i \bar{u}_i\gamma^\mu u_i + \frac{1}{x}\sum_i \bar{d}_i\gamma^\mu d_i + x\sum_i \bar{\ell}_i\gamma^\mu\ell_i. \tag{2.69}$$

From the QCD anomaly of this current one obtains as anomaly coefficient

$$\xi = N_c \left(x + \frac{1}{x} \right) . \tag{2.70}$$

To compute the axion mass, we use the chiral Lagrangian. For simplicity we keep only the two light quarks, u and d. The kinetic term is

$$\mathcal{L}_\chi = -\frac{f_\pi^2}{4} \, \mathrm{Tr} \left[\partial_\mu \Sigma \, \partial^\mu \Sigma^\dagger \right] , \tag{2.71}$$

where

$$\Sigma = \exp \left[\frac{i}{f_\pi} \left(\tau^i \pi_i + \eta \right) \right] . \tag{2.72}$$

Here τ_i are the Pauli matrices and η is the would-be Goldstone boson associated with the anomalous axial $U(1)$ symmetry in the two-flavor theory. The mass term is written as

$$\mathcal{L}_M = \frac{1}{2} \left(f_\pi m_\pi^0 \right)^2 \mathrm{Tr} \left[\Sigma A M + \left(\Sigma A M \right)^\dagger \right] , \tag{2.73}$$

with

$$A = \begin{pmatrix} e^{-iax/v_F} & 0 \\ 0 & e^{-ia/(xv_F)} \end{pmatrix} , \qquad M = \begin{pmatrix} \dfrac{m_u}{m_u + m_d} & 0 \\ 0 & \dfrac{m_d}{m_u + m_d} \end{pmatrix} . \tag{2.74}$$

The invariance of $\mathcal{L}_M$ under $U(1)_{\mathrm{PQ}}$ requires

$$\Sigma \longrightarrow \Sigma \begin{pmatrix} e^{i\alpha x} & 0 \\ 0 & e^{i\alpha/x} \end{pmatrix} . \tag{2.75}$$

In the definition of M and in the prefactor of $\mathcal{L}_M$ we have absorbed the constant B introduced in Sec. (1.2.1) through

$$m_\pi^2 = B(m_u + m_d). \tag{2.76}$$

Thus the expression above is the same chiral mass term discussed previously in Sec. (1.2.1), with the only difference that the quark mass matrix now contains the axion phase. Expanding to quadratic order in the neutral fields gives

$$\mathcal{L}_{\mathrm{mass},\, \pi^0, \eta, a} = -\frac{m_\pi^2}{2} \left[\frac{m_u}{m_u + m_d} \left(\pi^0 + \eta - \frac{x f_\pi}{v_F} a \right)^2 + \frac{m_d}{m_u + m_d} \left(\pi^0 - \eta - \frac{f_\pi}{x v_F} a \right)^2 \right] . \tag{2.77}$$

At this stage the η would have a mass of the same order as the pion, which would reproduce the usual $U(1)_A$ problem. Moreover, after diagonalization the axion would remain massless, since there are only two independent quadratic combinations for three neutral fields. Both issues are resolved by including the anomalous contribution. In the absence of the axion, the anomaly can be modeled by adding

$$\mathcal{L}_{\mathrm{anom}} = -\frac{C}{2} \left(\theta - i \ln \det \Sigma \right)^2 . \tag{2.78}$$

For $\theta = 0$ this term gives a mass to the η field. The constant C can therefore be traded for the physical η mass, so that schematically

$$\mathcal{L}_{\mathrm{anom}} = -\frac{1}{2} m_\eta^2 \eta^2 . \tag{2.79}$$

When the axion is included, the anomaly effectively shifts the vacuum angle. One obtains

$$\mathcal{L}_{\mathrm{anom}} = -\frac{1}{2} m_\eta^2 \left[\eta + \frac{f_\pi}{v_F} \frac{(N_g - 1)}{2} \left(x + \frac{1}{x} \right) a \right]^2 . \tag{2.80}$$

The coefficient multiplying the axion reflects the relative strength of the couplings of the axion and of the η to $G\widetilde{G}$. Since the couplings of the axion to the light quarks are already included in Eq. (2.77),

one must subtract the first generation, which explains the factor $N_g - 1$.
The anomalous term completes the quadratic Lagrangian in the neutral sector. After diagonalization, one obtains three massive eigenstates: the neutral pion, the heavier η, and the axion. We do not present the full diagonalization here, but quote the final results:

$$\begin{aligned} m_a &= \lambda_m\, m_\pi \frac{f_\pi}{v_F} \frac{\sqrt{m_u m_d}}{m_u + m_d}, \qquad \lambda_m = N_g \left(x + \frac{1}{x}\right), \\ \xi_{a\pi} &= \lambda_\pi \frac{f_\pi}{v_F}, \qquad \lambda_\pi = \frac{1}{2}\left[\left(x - \frac{1}{x}\right) - N_g \left(x + \frac{1}{x}\right)\frac{m_d - m_u}{m_u + m_d}\right], \\ \xi_{a\eta} &= \lambda_\eta \frac{f_\eta}{v_F}, \qquad \lambda_\eta = \frac{1}{2}(1 - N_g)\left(x + \frac{1}{x}\right). \end{aligned} \tag{2.81}$$

All axion models are also characterized by their coupling to two photons. We write the corresponding interaction as

$$\mathcal{L}_{a\gamma\gamma} = \frac{\alpha}{4\pi} K_{a\gamma\gamma} \frac{a}{f_a} F_{\mu\nu} \widetilde{F}^{\mu\nu}. \tag{2.82}$$

Starting from the electromagnetic anomaly of the Peccei–Quinn current,

$$\partial_\mu J^\mu_{\mathrm{PQ}} = \frac{\alpha}{4\pi} \xi_{a\gamma\gamma} F_{\mu\nu} \widetilde{F}^{\mu\nu}, \tag{2.83}$$

In fact one sees that the effective Lagrangian must contain a term of the form Eq. (2.82), in order to reproduce the anomalous variation under a Peccei–Quinn transformation. Including the charged leptons, the anomaly coefficient for the PQWW model is

$$\begin{aligned} \xi_{a\gamma\gamma} &= N_g \left[\left(\frac{2}{3}\right)^2 3x + \left(\left(-\frac{1}{3}\right)^2 3 + (-1)^2\right)\frac{1}{x}\right] \\ &= \frac{4}{3} N_g \left(x + \frac{1}{x}\right). \end{aligned} \tag{2.84}$$

The coefficient $K_{a\gamma\gamma}$ differs from $\xi_{a\gamma\gamma}$ for two reasons. First, the contribution of the lightest quark generation has to be subtracted, since the interactions between the axion and the light quarks are already included in the chiral Lagrangian. Second, the axion mixes with the neutral mesons π^0 and η, which themselves couple to two photons. After taking these effects into account, one finds

$$K_{a\gamma\gamma} = N_g \left(x + \frac{1}{x}\right) \frac{m_u}{m_u + m_d}. \tag{2.85}$$

As we said in the PQWW model the Peccei–Quinn scale is fixed by the electroweak scale,

$$v_F = (\sqrt{2} G_F)^{-1/2} \simeq 250 \text{ GeV}. \tag{2.86}$$

The axion is therefore not invisible: its mass and couplings are too large. Parametrically one finds from Eq. (2.81) $m_a \sim 100$ keV. Moreover, the model predicts rare meson decays such as

$$\mathrm{BR}(K^+ \to \pi^+ a) = 3 \times 10^{-5} \left(x + \frac{1}{x}\right)^2. \tag{2.87}$$

This is incompatible with the experimental bound on invisible kaon decays, see Ref. [49],

$$\mathrm{BR}(K^+ \to \pi^+ + \text{“nothing”}) < 10^{-8}. \tag{2.88}$$

Thus the original PQWW axion is experimentally excluded. This motivates the construction of *invisible* axion models. The key point is that the solution of the strong CP problem does not depend on the Peccei–Quinn scale being tied to the electroweak scale. One can instead introduce a much larger scale f_a, associated schematically with an angular mode of the form e^{ia/f_a}. For $f_a \gg (\sqrt{2} G_F)^{-1/2}$, the axion becomes very light,

$$m_a \sim \frac{\Lambda^2_{\mathrm{QCD}}}{f_a}, \tag{2.89}$$

and its interactions become very weak.

The DFSZ Axion Model

The Dine–Fischler–Srednicki–Zhitnitsky (DFSZ) model, first proposed in Refs. [50, 51] provide a well-motivated realization of the Peccei–Quinn mechanism in which the axion is "invisible", namely very light and very weakly coupled. In this framework, the Standard Model scalar sector is extended by two Higgs doublets, Φ_1 and Φ_2, together with a complex scalar singlet S, all charged under a global $U(1)_{\rm PQ}$ symmetry. For the Yukawa sector we have,

$$\mathcal{L}_{\rm Yukawa} = Y^u_{ij}\,\bar{Q}_{Li}\,\widetilde{\Phi}_1\,u_{Rj} + Y^d_{ij}\,\bar{Q}_{Li}\,\Phi_2\,d_{Rj} + Y^\ell_{ij}\,\bar{L}_{Li}\,\Phi_2\,\ell_{Rj} + \text{h.c.}, \tag{2.90}$$

where $\widetilde{\Phi}_1 = i\sigma_2\Phi_1^*$, ensuring gauge invariance and consistency with the PQ charge assignment. The scalar potential is constructed to be invariant under $U(1)_{\rm PQ}$ and contains interactions linking the singlet and doublet sectors, enforcing a non-trivial PQ charge assignment and connecting the singlet sector, where the high PQ scale is generated, to the electroweak Higgs sector. The PQ symmetry is spontaneously broken when the scalar fields acquire vacuum expectation values,

$$\langle S\rangle = \frac{v_s}{\sqrt{2}}, \qquad \langle\Phi_1\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v_1\end{pmatrix}, \qquad \langle\Phi_2\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v_2\end{pmatrix}, \tag{2.91}$$

We define $v = \sqrt{v_1^2 + v_2^2}$ setting the electroweak scale, and therefore the Higgs sector is basically left untouched. After PQ and electroweak symmetry breaking, the fields can be parametrized as

$$S = \frac{1}{\sqrt{2}}(v_s + \rho_s)e^{ia_s/v_s}, \qquad \Phi_i = \frac{1}{\sqrt{2}}e^{ia_i/v_i}\begin{pmatrix}0\\ v_i + h_i\end{pmatrix}, \qquad i = 1, 2. \tag{2.92}$$

As usual the physical axion is the Goldstone mode associated with the broken $U(1)_{\rm PQ}$ symmetry, orthogonal to the neutral Goldstone boson eaten by the Z boson. In the invisible limit $v_s \gg v$, it is mostly aligned with the scalar phase,

$$a \simeq a_s + \mathcal{O}\left(\frac{v}{v_s}\right), \qquad f_a \simeq v_s. \tag{2.93}$$

The radial excitation of the singlet acquires a mass of order the PQ scale. From the singlet part of the scalar potential,

$$V(S) \supset m_S^2|S|^2 + \lambda_S|S|^4, \tag{2.94}$$

and one finds, after imposing the minimum condition,

$$m_{\rho_s}^2 \simeq 2\lambda_S v_s^2. \tag{2.95}$$

And we can safely integrate out of the theory the radial excitation. In the limit $v_s \gg v$, the DFSZ model realizes an invisible axion because the axion decay constant is set by the large singlet vacuum expectation value, rather than by the electroweak scale.

At low energies, using the same chiral Lagrangian methods discussed in the previous section, one finds in the two-flavor approximation for the axion mass

$$m_a^2 = \frac{m_u m_d}{(m_u + m_d)^2}\,\frac{m_\pi^2 f_\pi^2}{f_a^2}. \tag{2.96}$$

More precisely, including the strange quark and using current algebra (for a treatment of this topic see Refs. [52, 53, 54, 55, 56]) one obtains

$$m_a = \frac{m_\pi f_\pi}{f_a}\left[\frac{z}{(1 + z + w)(1 + z)}\right]^{1/2}, \qquad z \equiv \frac{m_u}{m_d}, \qquad w \equiv \frac{m_u}{m_s}. \tag{2.97}$$

Numerically this gives the standard relation

$$m_a \simeq 5.7\ \mu\text{eV}\left(\frac{10^{12}\ \text{GeV}}{f_a}\right). \tag{2.98}$$

The couplings of the axion to gauge fields are fixed by the PQ anomalies. The effective low-energy Lagrangian necessarily contains the gluonic interaction

$$\mathcal{L}_{aGG} = \frac{\alpha_s}{8\pi}\frac{a}{f_a}G^a_{\mu\nu}\widetilde{G}^{a\mu\nu}, \tag{2.99}$$

which is responsible for the solution of the strong CP problem. We repeat an argument that was already presented with slightly different notation. Note, repeating what was already said above, that we expect after a chiral rotation of angle α sends $\theta \to \theta - N\alpha$, being N the anomaly coefficient, while we know that $a \to a + f_{PQ}\alpha$, so we that $f_{PQ} = Nf_a$. The QCD potential is periodic with period $2\pi f_a$, whereas the interpretation of the axion as the phase of the PQ field requires periodicity $2\pi f_{\text{PQ}}$. Therefore N must be an integer. This integer is sometimes also referred as the domain-wall number: the theory contains N degenerate vacua, each satisfying $\bar{\theta}_{\text{eff}} = 0$ and therefore solving the strong CP problem. This can be easily understood once one considers the scalar field to have a phase in the form $\Phi_i \sim e^{ia/f_{PQ}}$.
To compute the anomaly coefficient for QCD it is enough to consider a chiral rotation, each fermion will have its own PQ charge:

$$f \longrightarrow e^{-iX_f\gamma^5 a/(2f_a)}f. \tag{2.100}$$

then $N = \sum_q T(r_q)X_q$, where $T(r_q) = 1/2$ for quarks in the fundamental representation of $SU(3)_c$. The axion coupling to photons can be derived in analogous way, computing the electromagnetic anomaly

$$\Delta\mathcal{L} = \frac{\alpha}{8\pi}\frac{a}{f_a}\left[\sum_f N_c^f Q_f^2 X_f\right]F_{\mu\nu}\widetilde{F}^{\mu\nu}. \tag{2.101}$$

The quantity in square brackets is the electromagnetic anomaly coefficient E, again given the shift of the axion field it is easy to understand that in the UV the effective Lagrangian is:

$$\mathcal{L}^{\text{UV}}_{a\gamma\gamma} = \frac{\alpha}{8\pi}\frac{E}{N}\frac{a}{f_a}F_{\mu\nu}\widetilde{F}^{\mu\nu}. \tag{2.102}$$

At low energies, however, the physical axion also mixes with neutral mesons, especially with π^0 and η. Since these mesons couple anomalously to two photons, this mixing induces an additional model-independent contribution to the axion-photon coupling. Using current algebra again, Refs. [54, 57], the low-energy interaction can be written as

$$\mathcal{L}_{a\gamma\gamma} = -\frac{1}{4}g_{a\gamma\gamma}aF_{\mu\nu}\widetilde{F}^{\mu\nu} = g_{a\gamma\gamma}\,\vec{E}\cdot\vec{B}, \tag{2.103}$$

with

$$g_{a\gamma\gamma} = \frac{\alpha}{2\pi f_a}\left[\frac{E}{N} - \frac{2}{3}\frac{4+z+w}{1+z+w}\right]. \tag{2.104}$$

Finally the fermionic-axion couplings are of order

$$g_{af} = \frac{X_j m_j}{f_{\text{PQ}}}. \tag{2.105}$$

In the DFSZ model the coefficients X_f are determined by the PQ charges of the Higgs doublets and, after imposing orthogonality with the Goldstone boson eaten by the Z, depend on the ratio

$$\tan\beta = \frac{v_2}{v_1}. \tag{2.106}$$

For example, in a common DFSZ convention one finds, up to an overall sign,

$$X_u = \frac{1}{3}\cos^2\beta, \qquad X_d = \frac{1}{3}\sin^2\beta, \qquad X_e = \frac{1}{3}\sin^2\beta. \tag{2.107}$$

After QCD confinement, one should describe the axion couplings in terms of nucleons rather than quarks. For the DFSZ model the resulting nucleon couplings see Refs. [54, 57, 58, 59], can be written as

$$X_p = -0.10 - 0.45\cos^2\beta, \qquad X_n = -0.18 + 0.39\cos^2\beta. \tag{2.108}$$

The KSVZ Model

The Kim–Shifman–Vainshtein–Zakharov (KSVZ) axion model, introduced in Refs. [60, 61] is the prototype of hadronic invisible axion models. Its main idea is to introduce at least one new heavy vector-like quark Ψ, charged under $SU(3)_c$, together with a complex scalar field Φ which is a singlet under the electroweak gauge group $SU(2)_L \times U(1)_Y$. The scalar field carries Peccei–Quinn charge and its vacuum expectation value sets the PQ scale. Since the Standard Model fermions do not carry PQ charge in the minimal realization of this model, the axion has no tree-level couplings to ordinary quarks and leptons. The relevant part of the Lagrangian can be written as

$$\mathcal{L} = i\bar{\Psi}\not{D}\Psi + \partial_\mu\Phi\,\partial^\mu\Phi^\dagger - V(|\Phi|) - \left[Y\Phi\,\bar{\Psi}_L\Psi_R + \text{h.c.}\right], \tag{2.109}$$

where Ψ is vector-like with respect to the Standard Model gauge group. The scalar potential is taken as before to be

$$V(\Phi) = \lambda\left(|\Phi|^2 - \frac{f_{\rm PQ}^2}{2}\right)^2, \tag{2.110}$$

so that the minimum satisfies $|\langle\Phi\rangle| = \frac{f_{\rm PQ}}{\sqrt{2}}$. This Lagrangian is invariant under the global Peccei–Quinn transformation

$$\Phi \longrightarrow e^{i\alpha}\Phi, \qquad \Psi \longrightarrow e^{-i\gamma_5\alpha/2}\Psi. \tag{2.111}$$

After spontaneous PQ breaking, the scalar field can be parametrized as

$$\Phi = \frac{f_{\rm PQ} + \rho}{\sqrt{2}} e^{ia/f_{\rm PQ}}, \tag{2.112}$$

where ρ is the heavy radial mode and a is the axion. The heavy quark mass is then $m_\Psi = \frac{Y f_{\rm PQ}}{\sqrt{2}}$. At energies much below the PQ scale, as in the models presented above, the radial excitation can be integrated out. Under the PQ symmetry the axion shifts as always as $a \longrightarrow a + \alpha f_{\rm PQ}$. And again we have the usual PQ solution to the strong CP problem, as usual the interaction Lagrangian is:

$$\mathcal{L}_{aGG} = \frac{\alpha_s}{8\pi}\frac{a}{f_a} G^a_{\mu\nu}\widetilde{G}^{a\mu\nu}. \tag{2.113}$$

In the simplest KSVZ model with a single heavy quark in the fundamental representation of $SU(3)_c$, the domain-wall number is $N = 1$, so that $f_a = f_{\rm PQ}$. More generally, if several heavy quarks or different color representations are introduced, the anomaly coefficient changes and $f_a = f_{\rm PQ}/N$. The axion-photon coupling has the same general form discussed above in Eq. (2.104) where E and N are respectively the electromagnetic and color anomaly coefficients of the Peccei–Quinn current that depend on the electroweak quantum numbers of the heavy quark. The second term in Eq. (2.104) is model independent and comes from axion mixing with the neutral mesons. Note that if the heavy quark is electrically neutral, then $E = 0$ and the photon coupling is entirely induced by the model-independent meson mixing contribution. If instead the heavy quark has electric charge Q, one finds $E/N = 6Q^2$. Thus different choices of the heavy quark electric charge lead to different KSVZ variants, with the same solution to the strong CP problem but different phenomenology in photon-coupling searches. A characteristic feature of minimal KSVZ models is that the Standard Model fermions are neutral under $U(1)_{\rm PQ}$. As a consequence, the axion has no tree-level coupling to any leptons at the PQ scale. This is why KSVZ models are often called *hadronic* axion models.
Couplings to ordinary matter are nevertheless generated at low energies through axion–meson mixing and through radiative effects. For the nucleon couplings one has (see Refs. [54, 57, 58, 59]):

$$C_p \simeq -0.39, \qquad C_n \simeq -0.04, \tag{2.114}$$

up to the usual uncertainties of order ~ 0.01 associated with the nucleon matrix elements.

Chapter 3

Axion phenomenology, cosmology and experimental searches

In the first two chapters we introduced the strong CP problem and discussed some of its possible solutions, with particular emphasis on the Peccei–Quinn mechanism. We then showed that this solution predicts the existence of a new pseudoscalar particle, the axion. As discussed, the phenomenology of the axion is not unique: there are different axion models with sufficient freedom to obtain different, model-dependent predictions.
This chapter is devoted to the search for axions and axion-like particles (ALPs). Their phenomenology is rich and can be explored in a variety of physical settings. We begin by discussing the cosmology of axions and ALPs, showing how these particles can provide a viable dark matter candidate. We then review the main constraints on axions and ALPs arising from astrophysical arguments, and finally present the current experimental strategies aimed at their detection.

3.1 Axion Cosmology

One of the most appealing features of the axion is that it can also account for the dark matter abundance of the Universe, and therefore, addressing two fundamental problems at the same time: the strong CP problem and the nature of dark matter, see for example Refs. [62, 63, 64]. Axions can be produced in the early Universe through different mechanisms, the most important of which is the misalignment mechanism, which can be studied in different cosmological settings. In this section, we review how the axion relic abundance can be estimated in these scenarios.
The misalignment mechanism in fact can operate when the Peccei–Quinn (PQ) symmetry is broken during inflation and then is not restored afterwards. In this case, inflation selects a nearly homogeneous initial value of the axion field across the observable Universe. Conversely, if the PQ symmetry is restored either during or after inflation, different causally disconnected regions acquire different initial values of the axion field. This leads to the formation of topological defects, such as axion strings and domain walls, whose subsequent evolution and decay can provide an additional contribution to the axion abundance

3.1.1 ALP Production

The scalar misalignment mechanism describes the production of dark matter from a scalar field initially displaced from the minimum of its potential. We consider a homogeneous scalar field ϕ with potential $V(\phi) = \frac{1}{2}m_\phi^2\phi^2$ and initial value $\phi_i \neq 0$. The action is

$$S = \int d^4x\sqrt{-g}\left[-\frac{1}{2}\partial_\mu\phi\,\partial^\mu\phi - \frac{1}{2}m_\phi^2\phi^2\right], \tag{3.1}$$

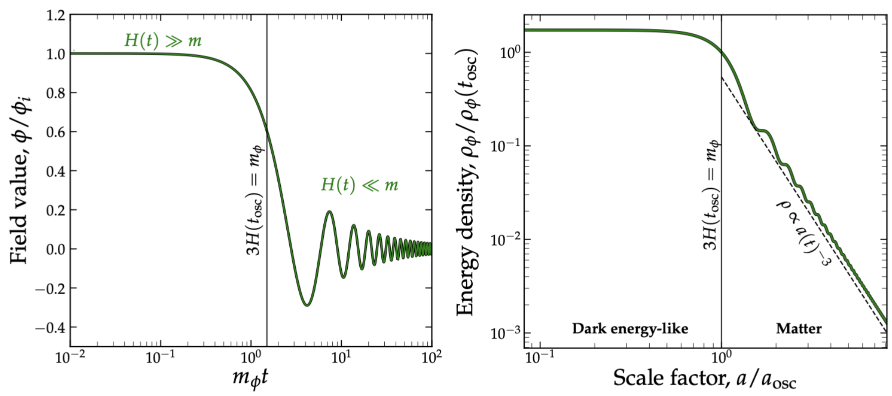

Figure 3.1: Evolution of the scalar field and of the energy density associated with the scalar field ϕ. Figure taken from Ref. [65].

where we assume an FRW metric, $g_{\mu\nu} = \mathrm{diag}(-1, a^2, a^2, a^2)$. Varying the action gives the Klein–Gordon equation,

$$\Box\phi - m_\phi^2\phi = 0. \tag{3.2}$$

For homogeneous modes, neglecting spatial derivatives, this becomes

$$\ddot{\phi} + 3H(t)\dot{\phi} + m_\phi^2\phi = 0. \tag{3.3}$$

Thus, the field behaves as a damped harmonic oscillator.
We now assume that these processes take place before matter-radiation equality, so that the Universe is radiation dominated and $H = 1/2t$. With initial conditions $\phi(0) = \phi_i$ and $\dot{\phi}(0) = 0$, the solution is

$$\phi(t) = \phi_i \left(\frac{2}{m_\phi t}\right)^{1/4} \Gamma\left(\frac{5}{4}\right) J_{1/4}(m_\phi t), \tag{3.4}$$

where $J_{1/4}$ is a Bessel function. A typical profile of the solution is shown in Fig. (3.1). Physically, at early times the Hubble friction term dominates and the field remains approximately frozen. Oscillations begin when $3H(T_{\rm osc}) \simeq m_\phi$.

At late times, for $m_\phi t \gg 1$, the friction term becomes negligible and the solution can be written approximately as $\phi(t) \simeq \phi_{\rm env}(t)\cos(m_\phi t)$, where $\phi_{\rm env}(t)$ is a slowly decreasing envelope. The energy density stored in the scalar field is

$$\rho_\phi = \frac{1}{2}\dot{\phi}^2 + \frac{1}{2}m_\phi^2\phi^2. \tag{3.5}$$

Averaged over many oscillations, this density scales as $\rho_\phi \propto a^{-3}$, which is the same scaling as non-relativistic matter. Therefore, after oscillations begin, the scalar field behaves as dark matter.
The present-day abundance is obtained by evolving the density from the onset of oscillations to today:

$$\rho_\phi(T_0) = \rho_\phi(T_{\rm osc})\frac{g_{*s}(T_0)}{g_{*s}(T_{\rm osc})}\left(\frac{T_0}{T_{\rm osc}}\right)^3. \tag{3.6}$$

This equation follows from the matter-like scaling of the energy density, together with entropy conservation, $S \sim g_{*s}(T)T^3a^3 = \text{const.}$. Using $\rho_\phi(T_{\rm osc}) \simeq \frac{1}{2}m_\phi^2\phi_i^2$ and the Friedmann equation at $T_{\rm osc}$, one obtains approximately

$$\Omega_\phi h^2 \simeq 0.12 \left(\frac{\phi_i}{4.7\times 10^{16}\,\mathrm{GeV}}\right)^2 \left(\frac{m_\phi}{10^{-21}\,\mathrm{eV}}\right)^{1/2}. \tag{3.7}$$

This relation shows how the scalar mass m_ϕ and the initial displacement ϕ_i must be chosen in order for the field to reproduce the observed dark matter abundance. The numerical values used here are typical benchmark values, with $\Omega_{\rm DM} \simeq 0.12$.

3.1.2 Axion Production

We now repeat the same arguments of the previous section for the QCD axion. The main difference is a feature that has not been made explicit so far: the previous discussion was performed at zero temperature. However, the axion mass is temperature-dependent. This is because the topological susceptibility, introduced in Eq. (2.43), can in general depend on the temperature. Therefore, Eq. (3.3) becomes

$$\ddot{\theta} + 3H(t)\dot{\theta} + m_\theta(T)^2 \sin\theta = 0. \tag{3.8}$$

In writing Eq. (3.8), we should remember that in the ALP case the particle was defined as a small perturbation around the minimum $\theta = 0$. During the cosmological evolution, however, the initial displacement of the QCD axion does not necessarily have to be small. For this reason, the Taylor expansion used to identify the axion mass around the minimum is not always the most appropriate description.

The full temperature dependence of the axion mass has been studied using analytic techniques with different regimes of validity, such as the dilute instanton gas approximation shown in Ref. [66] and the interacting instanton-liquid model in Ref [67]. It has also been computed numerically using lattice QCD in Ref. [68]. The result that we will use is that, for $T > \Lambda_{QCD}$, $m_a(T) \approx m_a(\Lambda_{QCD}/T)^{n/2}$.

To derive Ω_a, we start from the axion number density. Indeed, since the axion mass depends on temperature, the comoving energy density is not conserved, whereas the comoving number density is conserved. Therefore,

$$n_a(T_0) = n_a(T_{\rm osc}) \left(\frac{a_0}{a_{\rm osc}}\right)^{-3}. \tag{3.9}$$

At the onset of oscillations, we can approximate $\theta(T_{\rm osc}) \simeq \theta_i$ and $\dot{\theta}(T_{\rm osc}) \simeq 0$. Then,

$$n_a(T_{\rm osc}) = \frac{\rho_a(T_{\rm osc})}{m_a(T_{\rm osc})} \simeq \frac{1}{2} m_a(T_{\rm osc}) f_a^2 \theta_i^2. \tag{3.10}$$

It follows that the present axion energy density is

$$\rho_a(T_0) \simeq \frac{1}{2}\theta_i^2 f_a^2 m_a m_a(T_{\rm osc}) \left(\frac{a_0}{a_{\rm osc}}\right)^{-3}. \tag{3.11}$$

Again, as in the previous subsection, we use entropy conservation,

$$g_{*,s}(T) a^3 T^3 = \text{const.}, \tag{3.12}$$

to rewrite the dilution factor in terms of the temperature and the relativistic degrees of freedom. This gives

$$\rho_a(T_0) \simeq \frac{1}{2}\theta_i^2 f_a^2 m_a m_a(T_{\rm osc}) \left(\frac{g_{*,s}(T_0) T_0^3}{g_{*,s}(T_{\rm osc}) T_{\rm osc}^3}\right). \tag{3.13}$$

The temperature $T_{\rm osc}$ is defined by the condition $3H(T_{\rm osc}) = m_a(T_{\rm osc})$. For $T > T_{\rm QCD}$, we take $m_a(T) = m_a \left(\frac{T_{\rm QCD}}{T}\right)^n$, with typically $n \simeq 8$. Again, following the same steps as before and using the expression for the Hubble parameter during radiation domination, one finds approximately

$$T_{\rm osc} = \left[\left(\frac{10}{\pi^2 g_*(T_{\rm osc})}\right)^{1/2} m_a M_{\rm Pl}\right]^{\frac{2}{n+4}} T_{\rm QCD}^{\frac{n}{n+4}}. \tag{3.14}$$

In practice, one finds $T_{\rm osc} \sim 1\,\text{GeV}$, so that one may use $g_*(T_{\rm osc}) \simeq g_{*,s}(T_{\rm osc}) \simeq 61.75$, and $g_{*,s}(T_0) = 3.91$.

Dividing by the critical density, the relic abundance is

$$\Omega_a h^2 = \frac{\theta_i^2 f_a^2 m_a m_a(T_{\rm osc})}{6 H_0^2 M_{\rm Pl}^2} \left(\frac{g_{*,s}(T_0) T_0^3}{g_{*,s}(T_{\rm osc}) T_{\rm osc}^3}\right). \tag{3.15}$$

Although this expression seems to depend on m_a, f_a, and θ_i, for the QCD axion the zero-temperature mass is fixed by

$$m_a = (5.70 \pm 0.07)\,\mu\text{eV}\left(\frac{10^{12}\,\text{GeV}}{f_a}\right). \tag{3.16}$$

Thus, f_a can be eliminated in favour of m_a. Since $T_{\text{osc}} \propto m_a^{1/(n+4)}$, one obtains the scaling $\Omega_a h^2 \propto \theta_i^2 m_a^{-\frac{n+6}{n+4}}$. For $n = 8$,

$$\Omega_a h^2 \propto \theta_i^2 m_a^{-7/6}. \tag{3.17}$$

A more accurate numerical treatment, including the full temperature dependence of the axion mass and of the relativistic degrees of freedom, gives instead

$$\Omega_a h^2 \simeq 0.12\left(\frac{\theta_i}{2.155}\right)^2\left(\frac{28\,\mu\text{eV}}{m_a}\right)^{1.16}. \tag{3.18}$$

This result shows that lighter axions give a larger relic abundance. At first, this may seem counter-intuitive, since one might expect heavier particles to contribute more to the energy density. The important point is that the quantity fixed at T_{osc} is the number density. If the axion is lighter, the condition $3H(T_{\text{osc}}) = m_a(T_{\text{osc}})$ is satisfied later, so the axion number density is less diluted by the expansion of the Universe. As a consequence, lighter axions can lead to a larger present-day abundance.

There is, however, an important cosmological subtlety hidden in the choice of the initial misalignment angle θ_i: the PQ symmetry may be broken either before or after inflation. In the following, we discuss these two different possibilities.

3.1.3 Pre-inflationary Scenario

Let us first consider the case in which the Peccei–Quinn symmetry is broken before, or during, inflation and is not restored afterwards. In this scenario, inflation stretches a single causal patch to a size much larger than the observable Universe. Therefore the axion field is homogenized, and our whole observable Universe is described by one single initial misalignment angle θ_i. Following Ref. [63].
This makes the pre-inflationary scenario simple, but also less predictive. The relic abundance is not obtained by averaging over many possible values of θ_i, but depends on the particular value selected in our patch. For a generic initial angle $\theta_i \sim \mathcal{O}(1)$, requiring the axion to reproduce the observed dark matter abundance selects a typical QCD axion mass range. Smaller masses would usually overproduce dark matter, unless θ_i is tuned to be small. Conversely, larger masses can be obtained if the field starts close to the top of the potential, $\theta_i \simeq \pi$, where anharmonic corrections become important. Fig. (3.2) shows the initial angle θ_i needed to saturate the dark matter density for fixed QCD axion mass.

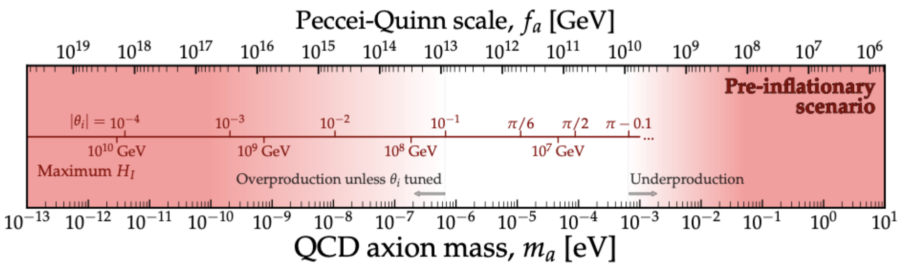

Figure 3.2: Figure taken from [65]

A 95 % credible interval is given as $0.12\ \mu\text{eV} < m_a < 150\ \mu\text{eV}$ assuming a uniform prior on θ_i .
A useful feature of this scenario is that topological defects, such as axion strings and domain walls, are inflated away. Thus the abundance is dominated by the misalignment mechanism alone. However,

the fact that the axion exists during inflation also leads to an important constraint. Like any light scalar field in de Sitter space, the axion receives quantum fluctuations of order

$$\delta\theta \simeq \frac{H_I}{2\pi f_a}, \tag{3.19}$$

where H_I is the Hubble scale during inflation. These fluctuations are later converted into fluctuations of the axion number density, producing isocurvature perturbations, see Refs. [69, 70, 71]. Cosmological perturbations can be divided into adiabatic and isocurvature modes. In adiabatic perturbations all components fluctuate together, so that the ratios between matter, radiation and the other species remain fixed. In isocurvature perturbations, instead, the total curvature perturbation is unchanged, but the relative abundances of the different components fluctuate. Since the CMB is observed to be mostly adiabatic, the axion's isocurvature contribution must be very small.
This gives an upper bound on the scale of inflation. In terms of the axion parameters, one finds approximately

$$H_I \lesssim 2.8 \times 10^7\,\mathrm{GeV}\,\theta_i \left(\frac{f_a}{10^{12}\,\mathrm{GeV}}\right). \tag{3.20}$$

Thus, if the PQ symmetry is broken before inflation, high-scale inflation is strongly constrained. If one further assumes that axions make up all of the dark matter, then θ_i and f_a are not independent, since the relic abundance fixes their combination. Using this relation, the isocurvature constraint can be rewritten as a bound involving only the axion mass and the inflationary scale. For $n = 8$, one obtains

$$H_I \lesssim 8.8 \times 10^8\,\mathrm{GeV} \left(\frac{1\,\mathrm{neV}}{m_a}\right)^{5/9}. \tag{3.21}$$

This expresses the consistency condition between pre-inflationary axion dark matter and inflation. Fig. (3.2) shows also the maximum for H_I at fixed axion mass.
Therefore the pre-inflationary scenario is viable, but it ties the axion parameter space to the scale of inflation. A future detection of primordial tensor modes, corresponding to a large value of H_I, could exclude large portions of this scenario. In the absence of such a detection, however, pre-inflationary axion dark matter remains possible, with its main uncertainty encoded in the unknown initial value of θ_i.

3.1.4 Post-Inflationary Scenario

We now consider the opposite possibility: the Peccei–Quinn symmetry is broken after inflation. In this case, inflation does not homogenize the axion field, and different causally disconnected regions of the Universe start with different initial values of the misalignment angle θ_i. As the horizon grows, all these regions eventually enter our observable Universe, so the relic abundance is not determined by one arbitrary value of θ_i, but by an average over many patches.
This makes the post-inflationary scenario more predictive than the pre-inflationary one. Once the average over θ_i is performed, including the anharmonic corrections to the axion potential, one obtains an effective value $\theta_i^{\mathrm{eff}} \simeq 2.155$. Using this value in the relic abundance formula (3.18) (so assuming that axions make up all the dark matter content) gives a preferred axion mass around $m_a \sim 28\,\mu\mathrm{eV}$, up to theoretical uncertainties related, for example, to the temperature dependence of the topological susceptibility.
However, this simple estimate is not the full story. In deriving the misalignment abundance we treated the axion field as spatially homogeneous, keeping only its zero-momentum mode. In the post-inflationary scenario this approximation is not fully justified, because the axion field is initially inhomogeneous. The equation of motion should therefore include spatial gradients,

$$\ddot{\theta} + 3H\dot{\theta} - \frac{1}{a^2}\nabla^2\theta + m_a^2\theta = 0. \tag{3.22}$$

These gradient terms are precisely what lead to additional phenomena, namely the formation of topological defects.

Cosmic Strings

When the PQ symmetry breaks after inflation, the axion field randomly chooses different phases in different causally disconnected regions. Since the axion angle is periodic, there can be closed paths in space along which the field winds by 2π. At the center of such a winding the phase is not well-defined, so the full PQ field must go to zero and the symmetry is locally restored, as in the Mexican-hat potential. In three dimensions, these singular points form one-dimensional objects: global cosmic strings as described in Ref. [72].
These strings carry energy and form a network in the early Universe. As the Universe expands, the network evolves and radiates axions. Therefore, in the post-inflationary scenario, the final axion abundance receives not only the usual misalignment contribution, but also an additional contribution from the decay of the string network, see Refs. [72, 73]. The precise size of this contribution is difficult to determine and usually requires numerical simulations, which is one of the main sources of uncertainty in the post-inflationary prediction Ref. [74].

Domain Walls

When the temperature approaches the QCD scale, the axion acquires its potential and the field starts to roll toward its minimum. At this stage another type of topological defect can appear: domain walls. These are two-dimensional surfaces separating regions where the axion field relaxes toward different equivalent minima of the potential, see Ref. [75]. The fate of these walls depends on the domain wall number N_{DW}, which, as anticipated in Sec. (2.3.2), counts the number of distinct vacua of the axion potential. For $N_{\text{DW}} = 1$ there is only one true vacuum, so the walls bounded by strings are unstable and the network eventually collapses, releasing its energy into axions. This case is cosmologically acceptable, although it still affects the final relic abundance.
For $N_{\text{DW}} > 1$, instead, the axion potential has several degenerate vacua. In this case the walls separating different vacua cannot simply disappear, because no vacuum is preferred over the others. The wall network can then become stable, and its energy density is diluted very slowly compared with matter and radiation. As a consequence, it may come to dominate the energy density of the Universe, leading to a cosmology in clear conflict with observations. This is the domain wall problem (Refs. [75, 76]). Viable post-inflationary axion models must therefore either have $N_{\text{DW}} = 1$ or include some mechanism that makes the walls unstable before they dominate the cosmic energy density.
In summary, the post-inflationary scenario is more predictive because the initial misalignment angle is averaged over many patches. However, it is also more complicated because spatial gradients cannot be ignored: cosmic strings and domain walls form and contribute to the final axion abundance. Furthermore, other new features emerge in this context: we only mention that these topological defects can lead to highly inhomogeneous distributions, such as axion miniclusters and axion stars. For this reason, the preferred axion mass in this scenario is tied not only to the misalignment mechanism, but also to the dynamics of topological defects.

3.2 ALP Cosmological Bounds

It is useful to summarize the cosmological constraints (and their preferred axion parameters) of axion-like particles in the two-dimensional parameter space spanned by the ALP mass m_a and the inverse decay constant $1/f_a$. This is the parameter space shown in Fig. (3.3). Since the interaction strengths of an ALP are typically suppressed by the scale f_a, moving upward in the plot corresponds to stronger couplings, while moving downward corresponds to more weakly coupled ALPs. Unlike the QCD axion, for a generic ALP the mass m_a and the scale f_a are independent parameters. Therefore the allowed regions are not restricted to a single QCD-axion line, but may fill extended bands.

While for QCD axion appears in Fig. (3.3) as a special line in the $(m_a, 1/f_a)$ plane, because for the QCD axion the mass and decay constant are related. The yellow part of the line corresponds to the pre-inflationary misalignment scenario, where θ_i can be chosen freely. The red segments correspond to post-inflationary predictions, including the contribution from topological defects. We do not repeat the cosmological bounds presented above, we see that different cosmological assumption identify

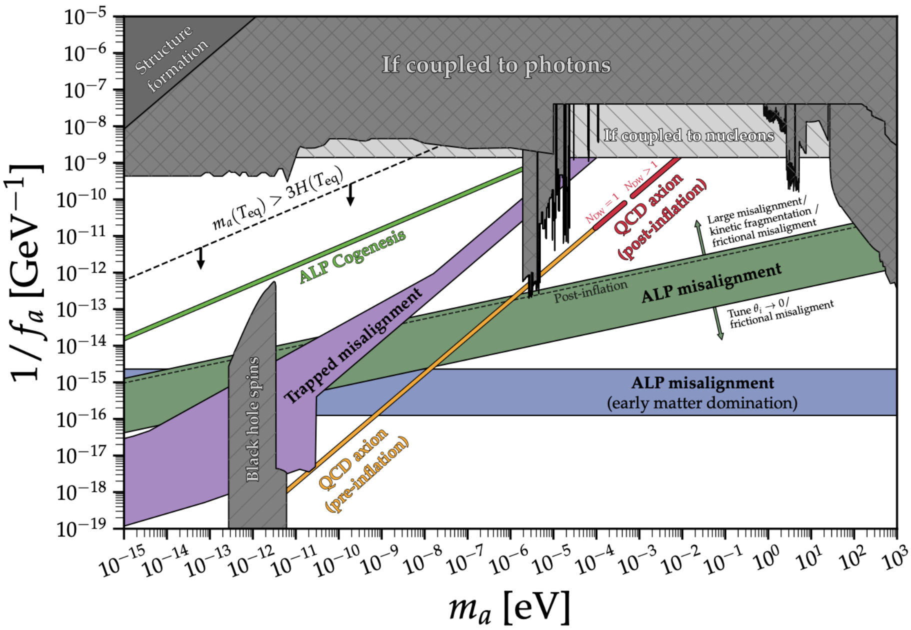

Figure 3.3: Excluded regions in the ALP parameter space, assuming different cosmological scenarios. Figure taken from [65]

different regions in the parameter space. However, we mention that further cosmological constraints can be applied to axions. For example from structure formation we have that if the ALP is extremely light, its de Broglie wavelength can become astrophysically large. In this regime, gradient pressure suppresses the growth of structure below a characteristic scale, modifying the matter power spectrum. This gives a lower bound on the ALP mass if the ALP makes up all of the dark matter. Very light ALPs are therefore constrained by observations of large-scale structure and small-scale structure, in a way similar to fuzzy dark matter bounds.

The figure (3.3) also shows how non-standard cosmological histories can modify the preferred ALP dark matter region. For example, an early matter-dominated era before BBN dilutes the ALP abundance through entropy injection. As a result, larger values of f_a can be consistent with the observed dark matter density, shifting the viable band downward in the $(m_a, 1/f_a)$ plane. Conversely, a period of kination, where the energy density is dominated by the kinetic energy of a scalar field and redshifts faster than radiation, changes the time at which ALP oscillations begin and can shift the preferred region in the opposite direction.

There are also mechanisms that modify the standard misalignment picture itself. In kinetic misalignment, the ALP field starts with a large initial velocity rather than from rest. The field can then roll over many periods of its potential before becoming trapped, delaying the onset of ordinary oscillations. This changes the final relic abundance and can populate regions above the standard ALP misalignment band. In trapped misalignment, the ALP initially oscillates around a false or temporary minimum before eventually settling into the true minimum. The delay enhances the abundance and therefore changes the relation between m_a and f_a required to obtain the observed dark matter density.

Finally, one should distinguish these cosmological bounds from the grey regions in Fig. (3.3). The grey regions are not cosmological abundance constraints: they usually assume that the ALP has a coupling either to photons or to nucleons, and therefore combine astrophysical, laboratory and cosmological information. For instance, sufficiently large photon couplings are constrained by stellar cooling, helioscope searches, supernova bounds and cosmological considerations. We present this type of arguments in the following sections. Similarly, nucleon couplings are constrained by stellar evolution, supernova physics and fifth-force or spin-dependent searches. These bounds are model dependent, because their position depends on which coupling is assumed to dominate. By contrast, the colored bands in the figure mainly indicate regions where different cosmological production mechanisms can generate the observed dark matter abundance.

3.2.1 Thermal Relic Axions

The bounds discussed so far rely on the assumption that axions account for the cold dark matter abundance. There is, however, another cosmological constraint on the QCD axion which is independent of this assumption. It comes from the fact that axions are thermally produced in the hot plasma of the early Universe (Refs. [77, 78]).
Since the QCD axion couples to Standard Model particles, thermal processes can produce a population of axions in the primordial bath. If the interaction rate is larger than the Hubble expansion rate, axions are kept in thermal equilibrium. As the Universe expands and cools, the interaction rate eventually becomes inefficient and axions decouple from the plasma. After decoupling, this population survives as a relic cosmic axion background (Refs. [77, 79]).
The cosmological effect of this thermal population depends on whether the axions are relativistic or non-relativistic at the epochs probed by observations. If thermal axions decouple early enough to be relativistic, they behave as hot dark matter. In this case they free-stream out of small-scale overdensities and suppress the matter power spectrum, in the same way as massive neutrinos, see Refs. [79, 80]. Large-scale structure observations strongly constrain this possibility, excluding thermal QCD axions with masses of order the eV scale, see Refs. [79, 80].

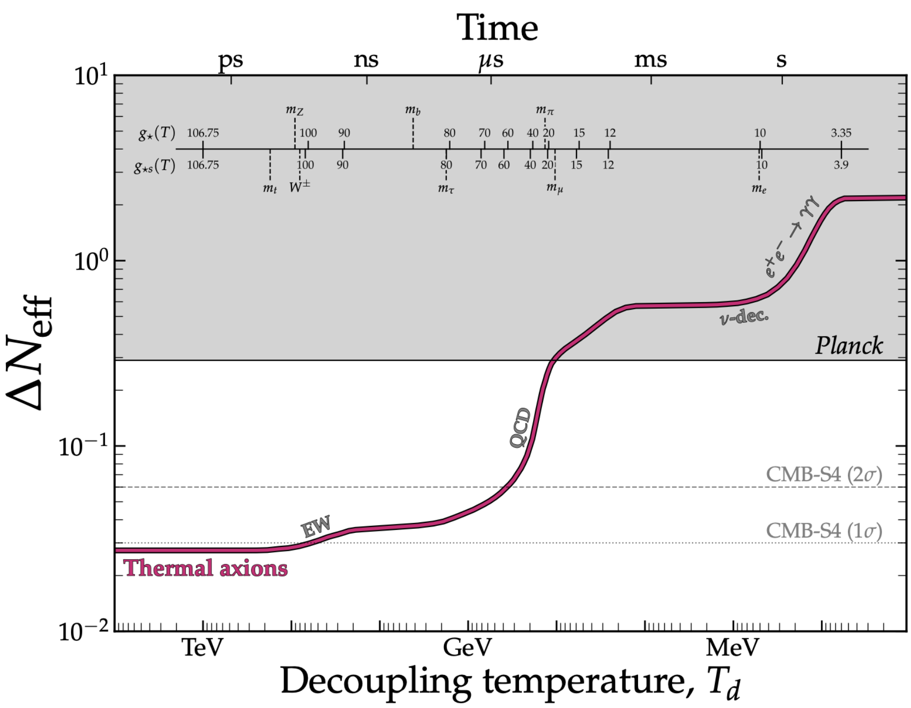

Figure 3.4: Upper limit on $\Delta N_{\rm eff}$ as a function of the axion decoupling temperature. Figure taken from Ref. [65]

If instead axions remain relativistic for most of the relevant cosmological history, without becoming non-relativistic, they contribute to the radiation content of the Universe as dark radiation. By dark radiation one means any relativistic energy density beyond the Standard Model photons and neutrinos. Its effect is usually described through the effective number of relativistic species $N_{\rm eff}$, see Refs. [81, 82], defined by writing the total radiation density after electron-positron annihilation as

$$\rho_{\rm rad} = \rho_\gamma + \rho_\nu + \rho_{\rm DR} = \rho_\gamma \left[1 + \frac{7}{8}\left(\frac{4}{11}\right)^{4/3} N_{\rm eff}\right]. \tag{3.23}$$

In the Standard Model one has

$$N_{\rm eff}^{\rm SM} \simeq 3.044, \tag{3.24}$$

where the small deviation from 3 comes from non-instantaneous neutrino decoupling and finite temperature effects, Ref. [83]. Any extra relativistic component is then parametrized as

$$N_{\text{eff}} = N_{\text{eff}}^{\text{SM}} + \Delta N_{\text{eff}}. \tag{3.25}$$

For a thermal axion background, the contribution to ΔN_{eff} is obtained by comparing the axion energy density with the energy density of one neutrino species. For a real scalar boson,

$$\rho_a = \frac{\pi^2}{30} T_a^4, \tag{3.26}$$

while one neutrino species, with two internal degrees of freedom, contributes

$$\rho_{\nu_i} = \frac{7}{8}\frac{\pi^2}{15} T_\nu^4. \tag{3.27}$$

Therefore a relativistic thermal axion gives

$$\Delta N_{\text{eff}} = \frac{\rho_a}{\rho_{\nu_i}} = \frac{4}{7}\left(\frac{T_a}{T_\nu}\right)^4. \tag{3.28}$$

Thus the relevant quantity we need to compute is the ratio between the axion temperature and the neutrino temperature.
This ratio is fixed by entropy conservation after axion decoupling. If axions decouple at a temperature T_d, they no longer share the entropy released by the annihilation of Standard Model particles. Consequently, the photon and neutrino baths are heated relative to the axion bath. One finds

$$\frac{T_a}{T_\gamma} = \left(\frac{g_{*s}(T_0)}{g_{*s}(T_d)}\right)^{1/3}, \tag{3.29}$$

where g_{*s} is the effective number of relativistic degrees of freedom contributing to the entropy density. Since after electron-positron annihilation

$$\frac{T_\nu}{T_\gamma} = \left(\frac{4}{11}\right)^{1/3}, \tag{3.30}$$

the axion contribution weighted against neutrinos can be written as

$$\Delta N_{\text{eff}} = \frac{4}{7}\left[\frac{11}{4}\frac{g_{*s}(T_0)}{g_{*s}(T_d)}\right]^{4/3}. \tag{3.31}$$

Equivalently, using the Standard Model entropy degrees of freedom, this is often written as

$$\Delta N_{\text{eff}} \simeq 0.027\left(\frac{106.75}{g_{*s}(T_d)}\right)^{4/3}. \tag{3.32}$$

This expression shows the main physical point: the earlier axions decouple, the larger $g_{*s}(T_d)$ is, and the smaller their contribution to dark radiation becomes, Refs. [82, 84]. Conversely, if axions decouple late, their temperature is closer to the neutrino temperature and ΔN_{eff} is larger.
Current CMB and large-scale structure data are consistent with the Standard Model value of N_{eff}, leaving little room for extra relativistic degrees of freedom, see Ref. [85]. The reason why this becomes an upper bound on the QCD axion mass is tied to entropy injection in the thermal bath. Once axions decouple, they no longer interact efficiently with the Standard Model plasma and therefore do not receive the entropy released when massive species become non-relativistic and annihilate. In particular, if axions decouple before electron-positron annihilation, the entropy released by e^+e^- annihilation heats the photon bath, while the axion bath remains colder.
More generally, the earlier axions decouple, the larger the number of Standard Model degrees of freedom that subsequently annihilate into the thermal bath. This increases the entropy dilution of

the axion background and lowers the ratio T_a/T_γ, or equivalently T_a/T_ν. As a result, early decoupling gives a smaller contribution to dark radiation,

$$\Delta N_{\text{eff}} = \frac{4}{7}\left(\frac{T_a}{T_\nu}\right)^4 . \tag{3.33}$$

This is why small values of m_a, corresponding to large f_a and weaker interactions, are safe: the axion decouples early, remains colder than the photon and neutrino baths, and contributes only a small amount to ΔN_{eff}.
Conversely, heavier QCD axions have smaller f_a and therefore stronger interactions with the thermal plasma. They remain in equilibrium for longer and decouple later, after fewer Standard Model degrees of freedom have annihilated. The entropy dilution is then less efficient, the axion temperature is closer to the neutrino temperature, and the resulting ΔN_{eff} is larger. Requiring this extra radiation to be compatible with cosmological bounds therefore excludes axions that are too strongly coupled, or equivalently too heavy. Present analyses typically give sub-eV limits, of order

$$m_a \lesssim 0.1\text{–}0.3\,\text{eV}, \tag{3.34}$$

with the precise value depending on the axion model, the thermal production rates and the cosmological datasets used in Refs. [86, 87, 88].

It is important to stress that the decoupling temperature corresponding to the present bound lies around the QCD crossover, $T_d \sim \Lambda_{\text{QCD}} \sim 100\text{–}200\,\text{MeV}$.
This is precisely the regime where the computation of the production rate is most delicate, see Ref. [89]. At temperatures well above the QCD transition, the relevant degrees of freedom are quarks and gluons, and axion production is described using the perturbative coupling of the axion to QCD. At temperatures below the transition, instead, the appropriate degrees of freedom are hadrons, and the dominant processes are described by the chiral effective Lagrangian, for example through axion-pion interactions, see Refs. [79, 89]. Therefore, a reliable determination of the thermal axion bound requires matching, or at least consistently interpolating, between the perturbative QCD description at high temperatures and the chiral description at low temperatures.
Future CMB experiments will improve the sensitivity to small values of ΔN_{eff}, as explained in Ref. [84, 90]. This is particularly interesting because even a very early decoupled thermal relic gives a non-zero minimal contribution to dark radiation. A future detection of $\Delta N_{\text{eff}} > 0$ would therefore be a possible hint of extra light relics that contributed to the early Universe thermal bath, including but not limited to axions, while the absence of such a signal would push thermal axion bounds towards smaller masses and weaker couplings.
Another possible production regime is freeze-in. In this case the axion interactions are too weak to bring the axion population into thermal equilibrium with the Standard Model bath, so there is no subsequent freeze-out temperature in the usual sense, see Ref. [91, 92]. Nevertheless, axions can still be slowly produced by rare processes in the plasma. The condition is schematically

$$\Gamma(T) < H(T) \tag{3.35}$$

at all relevant temperatures. Since the axion abundance never reaches the equilibrium value, the resulting contribution to dark radiation is typically smaller than in the freeze-out case. Moreover, the axion distribution need not be exactly thermal, so it is more appropriate to define its contribution as

$$\Delta N_{\text{eff}} = \frac{\rho_a}{\rho_{\nu_i}}, \tag{3.36}$$

where ρ_{ν_i} is the energy density of one neutrino species. Unlike the freeze-out case, the freeze-in abundance can depend on the details of the early thermal history, in particular on the maximum temperature reached after reheating, since production may be dominated at the highest available temperatures, see Ref. [93].

3.2.2 Lower Bounds from Fuzzy Dark Matter

For very small axion masses, the axion field can behave as fuzzy dark matter. In this regime, as anticipated above, the de Broglie wavelength becomes comparable to astrophysical scales,

$$\lambda_{\rm dB} \sim \frac{1}{m_a v}, \tag{3.37}$$

where v is the typical velocity of dark matter inside a halo. If m_a is too small, wave effects become important and the axion field cannot cluster efficiently below a characteristic length scale. This produces a suppression of small-scale structure and leads to lower bounds on the axion mass.
The same effect appears in the matter power spectrum as a cutoff at large wavenumber k. In terms of structure formation, this means that the abundance of low-mass halos and subhalos is reduced compared with the standard cold dark matter prediction. Measurements of galaxy abundances, halo mass functions and subhalo populations can therefore be used to constrain ultralight axions. Note however that this argument can be applied to any particle responsible for DM. These probes typically give bounds around

$$m_a \gtrsim 10^{-22}\text{–}10^{-21}\,\text{eV}, \tag{3.38}$$

although the precise value depends on the dataset and on the modeling of nonlinear structure formation, Refs. [94, 95].
One of the strongest probes is the Lyman-α forest. The absorption lines in the spectra of distant quasars trace the distribution of neutral hydrogen along the line of sight and therefore give access to the matter power spectrum on small scales. Since fuzzy dark matter suppresses precisely these scales, Lyman-α data can set significantly stronger lower bounds, reaching approximately

$$m_a \gtrsim 10^{-20}\,\text{eV}. \tag{3.39}$$

This number should not be interpreted as completely model independent, since it depends on the thermal history of the intergalactic medium, reionization, and hydrodynamical modeling, see Ref. [96].
Constraints can also be obtained from larger cosmological scales, for example using the CMB. These bounds are usually weaker, roughly

$$m_a \gtrsim 10^{-24}\,\text{eV}, \tag{3.40}$$

but they are theoretically cleaner because they mostly rely on linear perturbation theory rather than on nonlinear structure formation and baryonic effects Ref. [97].
Fuzzy dark matter is also phenomenologically interesting because it can soften some of the small-scale tensions of cold dark matter. In particular, wave pressure naturally leads to the formation of solitonic cores in the centres of halos, producing shallower density profiles and potentially addressing the core–cusp problem, see Ref. [98]. However, there is a tension between this motivation and the observational constraints: the masses light enough to produce sizeable cores are often close to, or below, the values excluded by Lyman-α forest data.
For sufficiently large masses, the de Broglie wavelength becomes too small to affect the observed structure formation, and the model approaches ordinary cold dark matter. Future probes such as galaxy surveys, gravitational lensing, 21-cm cosmology and high-precision CMB measurements may push the lower bound to larger masses, potentially approaching

$$m_a \sim 10^{-18}\,\text{eV}. \tag{3.41}$$

A compact review of fuzzy, scalar-field and wave dark matter phenomenology can be found in Ref. [99].

3.3 Astrophysical Bounds

Already in the previous sections we have encountered constraints on axions which rely not only on cosmology, but also on astrophysical arguments. In this section we briefly review some of the most important astrophysical bounds on axions and axion-like particles. We mainly follow Ref. [100]. Our goal is not to provide a complete review of all existing limits, but rather to introduce the physical

mechanisms behind the most relevant bounds and to highlight the role played by astrophysics.
Astrophysical systems are particularly powerful laboratories for weakly coupled light particles. If axions are produced in stellar interiors, they can escape almost freely because of their extremely small interaction rates. In this way, they provide an additional channel for energy loss. The existence of such a new cooling mechanism modifies the structure and the evolution of stars, supernovae, neutron stars and black holes and can therefore be constrained by comparison with observations.

3.3.1 Solar bounds

The Sun is the closest and best studied star, and hence represents an interesting source of axions. The solar interior is described by Standard Solar Models (SSMs), which reproduce the present solar luminosity, radius, age and surface composition by evolving the Sun from its formation to the present time. These models include nuclear reactions, radiative transport, convection, diffusion and the equation of state of the solar plasma. Their reliability is tested by helioseismology and by the measurement of solar neutrino fluxes, see Refs. [101, 102].
In the solar core, axions can be produced through several processes, depending on their couplings to Standard Model particles. We will compute the production rate inside the Sun for the case of axions coupled only to photons for different processes in chapter (5). For axion-like particles dominantly coupled to photons, the most important production channel is the Primakoff process (assuming we work with very light axions),

$$\gamma + Ze \to Ze + a, \tag{3.42}$$

in which a thermal photon is converted into an axion in the electric field of charged particles in the plasma. The relevant interaction is

$$\mathcal{L}_{a\gamma} = -\frac{1}{4} g_{a\gamma\gamma} a F_{\mu\nu} \tilde{F}^{\mu\nu} = g_{a\gamma\gamma} a \, \mathbf{E} \cdot \mathbf{B}. \tag{3.43}$$

A very simple estimate of the solar bound can be obtained from an energy-loss argument. If the Sun emitted too much energy in the form of axions, its nuclear burning rate would have to be larger in order to account for the observed photon luminosity. This would modify the solar structure and shorten the solar lifetime. A conservative requirement is therefore $L_a \lesssim L_\odot$, where L_a is the axion luminosity and $L_\odot$ is the observed solar photon luminosity, Ref. [101]. For Primakoff production, it is customary to define $g_{10} \equiv \frac{g_{a\gamma\gamma}}{10^{-10}\,\mathrm{GeV}^{-1}}$. A standard solar model calculation gives approximately

$$L_a \simeq 1.85 \times 10^{-3} \, g_{10}^2 \, L_\odot. \tag{3.44}$$

Imposing the conservative condition $L_a \lesssim L_\odot$ therefore gives $g_{10} \lesssim 23$, or equivalently

$$g_{a\gamma\gamma} \lesssim 2.3 \times 10^{-9}\,\mathrm{GeV}^{-1}. \tag{3.45}$$

This estimate is useful because it shows how stellar energy-loss arguments constrain very weakly coupled particles. However, it is not the strongest solar bound.
A more refined constraint comes from the fact that an additional energy-loss channel changes the solar core temperature. Since the flux of high-energy solar neutrinos, in particular those from ^{8}B, is very sensitive to the central temperature, solar neutrino measurements provide a stronger probe of axion emission, see Refs. [103, 102]. The axion luminosity induced by Primakoff production leads approximately to an enhancement of the ^{8}B neutrino flux of the form

$$\frac{\Phi(^8\mathrm{B})}{\Phi_{\mathrm{SSM}}(^8\mathrm{B})} \simeq \left(1 + \frac{L_a}{L_\odot}\right)^{4.4}. \tag{3.46}$$

Requiring consistency with the measured solar neutrino flux gives a stronger bound, typically quoted as

$$g_{a\gamma\gamma} \lesssim 6 \times 10^{-10}\,\mathrm{GeV}^{-1}. \tag{3.47}$$

As shown in Fig. (3.5) this bound is subject to the details of the solar model, in particular to the assumed solar metallicity, but it illustrates the strong constraining power of solar observations [102].

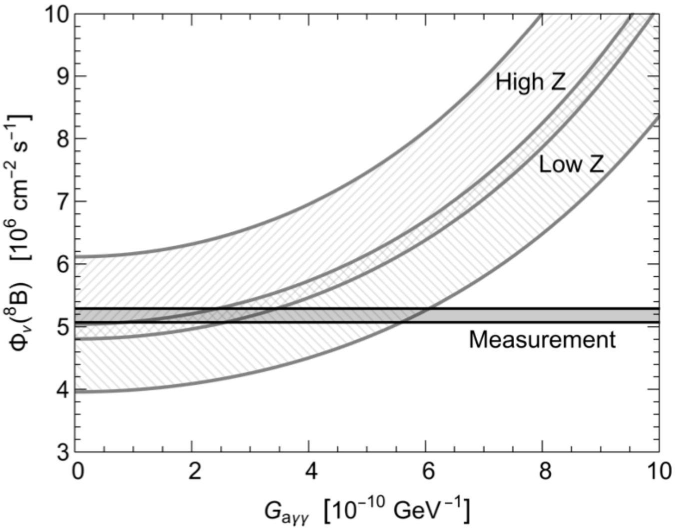

Figure 3.5: We see how the neutrino flux changes as a function of the $g_{a\gamma\gamma}$ for different solar models: Z represent the metallicity of different solar models. Figure taken from Ref. [100]

Trapping Regime

The bounds discussed above assume that axions, once produced in the solar core, escape freely from the Sun. This is the usual situation for sufficiently small couplings, since the axion mean free path is much larger than the solar radius. In this free-streaming regime, axions act as an efficient additional energy-loss channel and the corresponding luminosity can be constrained by solar evolution and solar neutrino data.
However, for very large values of the axion coupling this assumption eventually breaks down. If the interaction rate becomes large enough, axions can be reabsorbed or scattered before leaving the Sun. This effect is further enhanced once one considers also axion-electron couplings. In this case their mean free path becomes shorter than the solar radius and axions no longer provide an efficient mechanism for energy loss. This is known as the trapping regime. Qualitatively, the same coupling that increases the production rate also increases the absorption rate. Therefore, beyond a certain coupling strength, the simple energy-loss argument cannot be applied. For Primakoff-produced axions, the relevant mean free path is approximately the inverse of the Primakoff conversion rate. In the solar core, where the temperature is of order $T_c \simeq 1.3\,\text{keV}$, (we will see the temperature profile inside the Sun later on) and the Debye–Hückel screening scale is roughly $k_s \simeq 9\,\text{keV}$, that is the screening scale of the solar plasma, whose inverse gives the typical distance over which electric fields are screened by the surrounding charged particles. Again we will come back to this later on.
One finds that axions with energies of a few keV would be trapped only for extremely large photon couplings, roughly

$$g_{a\gamma\gamma} \gtrsim 10^{-3}\,\text{GeV}^{-1}. \tag{3.48}$$

This value is many orders of magnitude above the couplings relevant for ordinary QCD axions and for the solar bounds discussed above. Moreover, if axions were trapped efficiently, they would contribute to radiative energy transport inside the Sun. In order not to significantly modify the solar structure, their mean free path should then be shorter than the photon mean free path, which is of order several centimeters in the solar core.
The trapping argument is therefore mainly important conceptually. It shows that stellar energy-loss bounds do not necessarily exclude arbitrarily large couplings: they constrain the free-streaming region, while at very large couplings the new particle may become trapped. For ordinary solar axions coupled through the Primakoff process, however, the trapping regime lies at couplings so large that it is not relevant for the phenomenology considered here.
There is an additional subtlety for QCD axions. In this case the axion mass and the photon coupling

are not independent, since one can write schematically $g_{a\gamma\gamma} = \frac{C_{a\gamma}}{2\pi f_a}$, and $m_a \propto \frac{1}{f_a}$. Thus, very large values of $g_{a\gamma\gamma}$ correspond to very large axion masses. When the axion mass becomes larger than the solar temperature, $m_a \gg T$, thermal production is exponentially suppressed by the Boltzmann factor. In this heavy-mass regime, as we will see more quantitatively in chapter (5), the dominant production channel is no longer standard Primakoff conversion, but rather photon coalescence,

$$\gamma + \gamma \to a. \tag{3.49}$$

The resulting production rate is suppressed as $Q_a \propto e^{-m_a/T}$, so that sufficiently heavy axions are again difficult to produce in the solar plasma. If produced, such heavy axions may decay into two photons $a \to \gamma + \gamma$, possibly generating a flux of hard solar X-rays see Refs. [104, 105, 106, 107]. The absence of such an anomalous X-ray flux provides another constraint on heavy axions, as shown in Refs [105, 106]. However another important argument can be applied to solar axions, i.e. the solar basin argument, that we will present in more detail in chapter (5).

Axion–Electron Coupling

So far we have mainly discussed solar axions produced through the Primakoff process, which is controlled by the axion–photon coupling. However, in several axion models the axion can also have a direct coupling to electrons. This is the case, for instance, of the DFSZ-type models.
As anticipated, when the axion couples directly to electrons, the solar axion flux receives additional contributions. The most important production mechanisms are usually referred to as ABC processes, namely Atomic recombination, de-excitation, Bremsstrahlung and Compton scattering. Schematically, these include the processes: $\gamma + e \to e + a$, for Compton-like production, $e + I \to e + I + a$, for electron-ion bremsstrahlung, and bound-bound processes $I^* \to I + a$.
These atomic processes are particularly important because the solar plasma is not made only of free electrons and fully ionized nuclei: heavier elements can still have bound electrons in some regions of the Sun. For this reason, a precise calculation of the ABC flux requires detailed solar opacity tables and a careful treatment of atomic physics. The resulting solar axion spectrum is different from the Primakoff one. Primakoff axions are produced mainly by thermal photons with energies of a few keV and have an average energy around 4 keV. By contrast, the ABC flux is softer, with a typical average energy around $\simeq 2.3$ keV.
Moreover, because atomic transitions contribute to the emission, the spectrum can contain several line-like features in addition to the continuous component. Defining $g_{11} \equiv \frac{g_{ae}}{10^{-11}}$. A solar model calculation gives approximately, from Ref. [108]:

$$L_a \simeq g_{11}^2\, 7.15 \times 10^{-2} L_\odot. \tag{3.50}$$

As in the Primakoff case, the conservative requirement that $L_a \lesssim 0.04 L_\odot$, gives:

$$g_{ae} \lesssim 0.75 \times 10^{-11}. \tag{3.51}$$

For DFSZ axions this can be translated into a lower bound on the axion decay constant and an upper on the axion mass. In the usual normalization one finds approximately

$$f_a \gtrsim 1.1 \times 10^7\,\text{GeV}\,(2\sin^2\beta), \qquad m_a \lesssim \frac{0.50\,\text{eV}}{2\sin^2\beta}. \tag{3.52}$$

Here β is the usual DFSZ angle defined in chapter (2). There is an important difference with respect to the pure Primakoff case. If g_{ae} is non-vanishing, the total solar axion flux is not only controlled by $g_{a\gamma\gamma}$, but also by g_{ae}. In DFSZ-like models the ABC flux can even dominate over the Primakoff flux. The signal may receive a contribution proportional to

$$R_{\text{ABC}} \propto g_{ae}^2 g_{a\gamma\gamma}^2, \tag{3.53}$$

This is why the interpretation of helioscope bounds becomes more model-dependent when the electron coupling is present.

Solar axions coupled to electrons can also be searched for in low-threshold direct-detection experiments. In this case axions from the Sun can be absorbed or scattered in the detector, producing electronic recoils with keV-scale energies. Experiments originally designed for dark matter searches have therefore become sensitive to solar axions as well. In particular, the XENONnT experiment, Ref. [109] currently provides one of the strongest bounds on the axion–electron coupling from the solar ABC flux,

$$g_{ae} < 1.9 \times 10^{-12} \qquad (90\%\,\text{C.L.}). \tag{3.54}$$

In DFSZ models this corresponds approximately to

$$f_a \gtrsim 4.5 \times 10^7\,\text{GeV}\,(2\sin^2\beta), \qquad m_a \lesssim \frac{0.13\,\text{eV}}{2\sin^2\beta}. \tag{3.55}$$

Finally we only mention that in principle one should consider la also axion-nuclear couplings for the solar production, nevertheless they are strongly model dependent and do not provide competitive bounds to the ones presented in the previous discussion.

Solar Basins

We conclude this section on solar axion bounds by anticipating one of the main topics of the following chapters: the possibility that axions or axion-like particles produced in the Sun do not necessarily escape to infinity, but remain gravitationally trapped in bound orbits around the Sun. This population of trapped particles is usually referred to as a *solar basin*.
The basic idea is that particles produced in the solar interior with velocity larger than the solar escape velocity leave the Sun and contribute to the usual solar axion flux at Earth. However, a fraction of the particles can be produced with sufficiently small kinetic energy that their velocity is below the escape velocity. These particles remain gravitationally bound to the Sun and can accumulate over the solar lifetime. The effect is most relevant for particles with masses comparable to the temperature of the stellar interior. In the case of the Sun, this corresponds to the keV mass range, since the core temperature is of order $T_c \sim$ keV. In fact for much lighter particles the typical thermal velocities are relativistic and the probability of gravitational capture is strongly suppressed.
This mechanism was first proposed in Ref. [110] and it is particularly interesting for models containing a tower of massive states, i.e. Kaluza–Klein axions. In these scenarios, the higher-dimensional axion appears in four dimensions as a set of modes with different masses. Some of these modes can naturally lie in the keV range, where solar production is efficient and gravitational trapping can become relevant. Even if each individual mode is very weakly coupled, the presence of many accessible states can enhance the total population accumulated around the Sun. We will focus on this models and extra dimensional theories in the next chapter.
A solar basin of keV-mass axions or ALPs could be searched for in low-background direct-detection experiments, where the absorption or scattering of these particles would produce electronic signals in the keV energy range. Depending on the dominant coupling, the relevant detection channels may involve the axion–electron coupling g_{ae} or the axion–photon coupling $g_{a\gamma\gamma}$. In addition, if the particles are unstable, their radiative decay into two photons could generate an X-ray signal from the quiet Sun or from the solar surroundings see Refs. [111, 112, 106].

3.3.2 Axion–photon Conversion in Neutron-Star Magnetospheres

In a magnetic field, axions and photons can mix through the axion–photon coupling $g_{a\gamma\gamma}$. This leads to the possibility of axion–photon conversion in astrophysical environments with large magnetic fields and macroscopic length scales. The relevant quantity controlling the efficiency of the conversion is roughly the product BL, where B is the magnetic field strength and L is the size of the conversion region. Therefore, objects such as pulsars, magnetars and large-scale intergalactic magnetic fields can provide interesting probes of axion-like particles.
However, the conversion is efficient only if the axion and photon remain coherent during propagation. For non-negligible axion masses, the momentum mismatch between the two states suppresses the

conversion probability. For very light axions, a typical bound reached is $g_{a\gamma\gamma} < 10^{-12}\mathrm{GeV}^{-1}$. This is the reason why many astrophysical bounds on $g_{a\gamma\gamma}$, although very strong for generic ALPs, do not usually reach the QCD axion line. In this sense, helioscopes such as CAST and future experiments like IAXO or BabyIAXO remain among the most promising approaches to probe realistic QCD axion models.

Neutron stars are especially interesting because of their extremely large magnetic fields, typically of order $B \gtrsim 10^{12}\,\mathrm{G}$. If dark-matter axions fall into the magnetosphere of a neutron star, they can encounter a region where the plasma frequency matches the axion mass, see Ref. [113]. In this case, axion–photon conversion can become resonant, producing a narrow radio line. This possibility has motivated several searches using radio-telescope observations, for example in Ref. [114].

Independently of the dark-matter axion scenario, neutron stars themselves may also produce axions in their magnetospheres, especially in polar-cap regions where large values of $\mathbf{E}\cdot\mathbf{B}$ can act as an efficient source. Some of these axions may remain gravitationally bound to the neutron star, forming dense axion clouds, see Ref. [115]. Such clouds could accumulate over the lifetime of the star and potentially lead to observable electromagnetic signals.

3.3.3 Globular-Cluster Stars

Globular clusters provide some of the most important astrophysical constraints on axion couplings. They are gravitationally bound systems containing a large number of stars which formed approximately at the same time and with similar chemical composition. For this reason, the stars in a globular cluster mainly differ by their mass, making these systems very useful laboratories for testing stellar evolution.

The evolutionary stage of the stars in a globular cluster can be studied through a color–magnitude diagram. In this diagram, stars populate characteristic branches corresponding to different phases of their life. Since axion emission would provide an additional energy-loss channel, it can modify stellar evolution and therefore change the relative population of stars along these different branches or the brightness of specific evolutionary features.

Tip of the Red Giant Branch and Axions

Low-mass stars, such as those found in old globular clusters, spend most of their life on the main sequence (MS), where they burn hydrogen into helium in their cores. Once the central hydrogen is exhausted, the star leaves the main sequence and enters the subgiant branch (SGB) see figures (3.6) and (3.7). In this phase the core is mostly made of inert helium, while hydrogen burning continues in a shell around it. The star then ascends the red giant branch (RGB). During this phase, the helium core is supported by electron degeneracy pressure, and for this systems the relation between the core mass and radius is approximately $R \sim M^{-1/3}$.

As the mass of the degenerate helium core increases, its radius decreases and its gravitational potential becomes deeper. This drives up the temperature in the hydrogen-burning shell and therefore increases the luminosity of the star. The star continues to climb the RGB until the core becomes hot and dense enough to ignite helium, roughly at $\rho \sim 10^6\,\mathrm{g/cm^3}$ and $T \sim 10^8\,\mathrm{K} \simeq 8.6\,\mathrm{keV}$. For low-mass stars, $M \lesssim 3M_\odot$, helium ignition occurs under degenerate conditions and produces the so-called helium flash. More massive stars ignite helium under non-degenerate conditions.

The degenerate RGB core before helium ignition is essentially a helium white dwarf, except that it is surrounded by an active hydrogen-burning shell. After helium ignition, the core expands to non-degenerate densities of order $10^4\,\mathrm{g/cm^3}$, while its temperature remains of order $10^8\,\mathrm{K}$. The star then settles on the horizontal branch (HB), where it quietly burns helium into carbon and oxygen in the core, while hydrogen burning continues in a shell. The star is now much dimmer than at the TRGB because the larger core radius lowers the temperature of the hydrogen-burning shell, reducing the dominant hydrogen-burning luminosity. Later, after central helium exhaustion, the star moves to the asymptotic giant branch (AGB), where it has an inert carbon–oxygen core surrounded by helium- and hydrogen-burning shells.

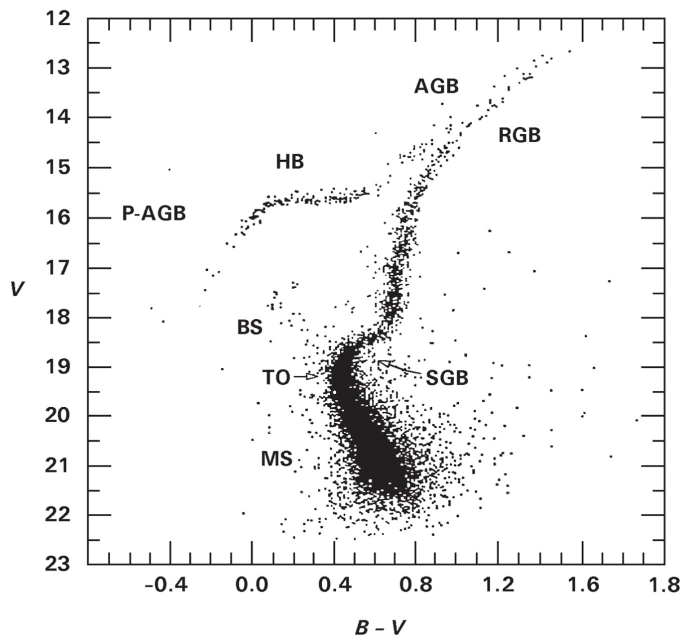

Figure 3.6: Color–magnitude diagram of the globular cluster M5. The main stellar evolutionary phases are visible: the main sequence (MS), the subgiant branch (SGB), the red giant branch (RGB), the horizontal branch (HB), and the asymptotic giant branch (AGB). The tip of the red giant branch (TRGB) corresponds to the brightest point reached by low-mass red giants before helium ignition.

The maximum brightness reached just before helium ignition defines the tip of the red giant branch, or TRGB.

This observable is very sensitive to additional cooling mechanisms. If axions couple to electrons, they can be efficiently produced in the dense red-giant core, mainly through bremsstrahlung processes involving electrons and ions. This extra cooling makes the core lose energy more efficiently. As a consequence, the helium core must grow to a slightly larger mass before reaching the conditions required for helium ignition. The star therefore climbs further up the red giant branch, and the TRGB becomes brighter. The absence of such an anomalous increase in the observed TRGB brightness can be used to constrain the axion–electron coupling. It is customary to introduce the rescaled coupling $g_{13} \equiv \frac{g_{ae}}{10^{-13}}$. The current bound is, from Refs. [116, 117, 118]

$$g_{ae} \lesssim 1 \times 10^{-13},$$

For DFSZ axions, where the axion–electron coupling is present at tree level, this translates into a strong lower bound on the axion decay constant and therefore into an upper bound on the axion mass.

$$g_{ae} = C_e \frac{m_e}{f_a} \simeq 2.8 \times 10^{-14}\, C_e \left(\frac{m_a}{\text{meV}}\right),$$

where C_e is a model-dependent coefficient. In DFSZ models C_e is of order one and depends on the angle β.

At the limit $g_{13} = 1$, the core-averaged axion luminosity is already a fraction of the standard neutrino luminosity. In particular, one finds approximately is roughly one fifth of the standard neutrino luminosity. In this sense, the present TRGB bound is already quite ambitious: it constrains an exotic cooling channel whose contribution is significantly smaller than the ordinary neutrino cooling.

TRGB bounds are among the strongest astrophysical constraints on g_{ae}.

We finally mention that TRGB bounds can also be obtained without relying exclusively on globular clusters. Field halo stars in the outskirts of galaxies also belong to old stellar populations and can show a sharp break in their luminosity function. This break is produced by the TRGB and is commonly used as a distance indicator. In this approach, the TRGB is measured in resolved stellar populations in galactic halos rather than in globular clusters.

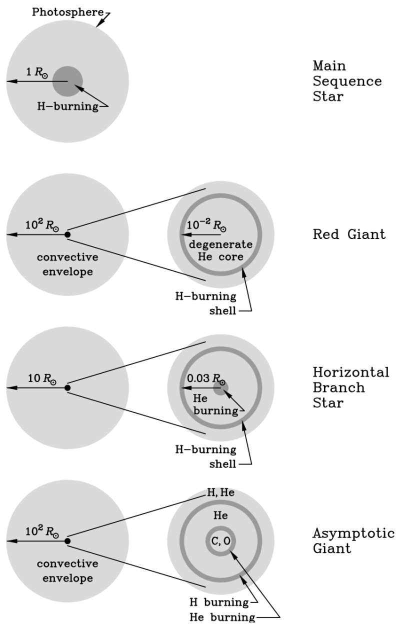

Figure 3.7: Schematic structure of a low-mass star in different evolutionary phases. During the main sequence the star burns hydrogen in its core. After core hydrogen exhaustion, hydrogen burning continues in a shell around an inert helium core, and the star moves through the SGB and RGB phases. After helium ignition, the star becomes an HB star, burning helium in the core and hydrogen in a shell. Later, during the AGB phase, the star develops an inert carbon–oxygen core surrounded by helium- and hydrogen-burning shells.

This method is important because one of the main systematic uncertainties in the globular-cluster analysis is the determination of the cluster distances. A wrong distance calibration would directly affect the inferred absolute luminosity of the TRGB, and therefore the axion bound. Using halo field stars provides an independent way to calibrate the TRGB and reduces the dependence on globular-cluster distance determinations.

Helium-burning Lifetime and Axions

Horizontal-branch (HB) stars burn helium in a non-degenerate core with a typical mass of order $0.5\,M_\odot$. The core temperature and density are approximately $T \sim 10^8\,\mathrm{K}$, $\rho \sim 10^4\,\mathrm{g\,cm^{-3}}$. In these conditions, axions or ALPs coupled to photons can be efficiently produced through mainly the Primakoff process $\gamma + Ze \to a + Ze$, providing an additional cooling channel. As a consequence, helium would be consumed more rapidly and the HB lifetime would be shortened. The Primakoff energy-loss rate per unit mass scales approximately as $\epsilon_a \propto \frac{T^7}{\rho}\,G_{10}^2$, as done in Ref. [119] and for typical HB-core conditions, the core-averaged axion energy-loss rate is of order

$$\epsilon_a \simeq 30\,G_{10}^2\,\mathrm{erg\,g^{-1}\,s^{-1}}.$$

This should be compared with the standard helium-burning energy generation rate, roughly

$$\epsilon_{\mathrm{He}} \simeq 80\,\mathrm{erg\,g^{-1}\,s^{-1}}.$$

Therefore, for $G_{10} \sim 1$, axion cooling would reduce the HB lifetime by a factor of order

$$\frac{80}{80 + 30G_{10}^2},$$

corresponding to a reduction of about 30% for $G_{10} = 1$. This effect can be constrained observationally by comparing the number of HB stars with the number of red-giant branch (RGB) stars brighter than the HB. The relevant observable is the so-called R parameter, $R \equiv \frac{N_{\mathrm{HB}}}{N_{\mathrm{RGB}}} = \frac{\tau_{\mathrm{HB}}}{\tau_{\mathrm{RGB}}}$. Since axion cooling shortens τ_{HB}, a large axion–photon coupling would reduce R. The observed value of R in globular clusters is compatible with standard stellar evolution, leaving little room for an additional Primakoff cooling channel. This leads to the bound, Ref. [120]

$$g_{a\gamma\gamma} < 0.65 \times 10^{-10}\,\mathrm{GeV}^{-1} \qquad (95\%\ \mathrm{C.L.}).$$

For KSVZ axions, this can be translated into

$$f_a > 3.4 \times 10^{7}\,\mathrm{GeV}, \qquad m_a < 0.17\,\mathrm{eV}.$$

For DFSZ models the axion also couples to electrons, and this additional interaction would significantly affect RGB evolution. As a result, the R parameter would no longer isolate the axion–photon interaction alone. For DFSZ axions, constraints from the tip of the red-giant branch provide the more appropriate and typically stronger bound.

Bounds from the Asymptotic Giant Branch

After helium is exhausted in the core of a horizontal-branch star, the star develops a degenerate carbon–oxygen core. Helium burning continues in a shell around this core, while hydrogen burning proceeds in a second, outer shell. As the carbon–oxygen core grows in mass, it contracts and the helium-burning shell becomes hotter and more efficient. The star then ascends the asymptotic giant branch (AGB). This phase occurs for stars with $M \lesssim 8\,M_\odot$, which do not ignite the next nuclear-burning stage, but instead lose their envelope and end their evolution as white dwarfs. AGB evolution is interesting but complicated in that the burning shells can interact and the occurrence of thermal pulses. New channels of energy loss, such as axion emission, can modify this evolution. In particular, excessive cooling may affect the final white-dwarf mass, the number of luminous AGB stars, and the chemical composition of the material brought to the surface by dredge-up episodes. For QCD axions, many of these effects in intermediate-mass stars are marginal once present bounds are imposed. However, for more general ALPs, hotter stellar environments can probe larger ALP masses and can exclude regions of the $G_{a\gamma\gamma}$–m_a plane that are not accessible with low-temperature systems.

The most relevant constraint for low-mass stars in globular clusters comes from the AGB clump. After the end of core helium burning, the star does not immediately climb the AGB. Instead, it remains for some time at a luminosity of roughly a few times the HB luminosity. During this stage, low-mass AGB stars spend most of their AGB lifetime in a narrow luminosity region known as the AGB clump. The subsequent ascent along the AGB is much faster and ends when the envelope is lost and the star becomes a white dwarf.

The key point is that helium shell burning in the AGB clump occurs at a higher temperature than core helium burning during the HB phase. Therefore, energy losses from neutrinos or axions are more efficient. Axion emission shortens both the HB and AGB-clump lifetimes, but it affects the AGB clump more strongly. As a result, the ratio

$$R_2 \equiv \frac{N_{\mathrm{AGB}}}{N_{\mathrm{HB}}}$$

decreases in the presence of Primakoff axion emission. Since number counts trace evolutionary lifetimes, R_2 provides a sensitive probe of additional energy losses in helium-burning stars.

Observationally, HST photometry of globular clusters finds a narrow AGB clump and a value of R_2 consistent with standard stellar evolution. Since axion cooling would reduce the AGB-clump lifetime and hence lower R_2, the absence of such a deficit constrains the axion–photon coupling. A conservative recent analysis gives, Ref. [121]

$$g_{a\gamma\gamma} < 0.47 \times 10^{-10}\,\mathrm{GeV}^{-1} \qquad (95\%\ \mathrm{C.L.}).$$

For KSVZ axions, this corresponds to

$$f_a > 4.7 \times 10^{7}\,\mathrm{GeV}, \qquad m_a < 0.12\,\mathrm{eV}.$$

As in the HB case, this bound should not be directly translated to DFSZ axions in a self-consistent way.

3.3.4 White Dwarf and Axions

As seen in the previous section, stars with masses $M \lesssim 8\,M_\odot$ after losing most of their envelope, end their evolution as carbon–oxygen white dwarfs (WD), with typical masses in the range $M_{\rm WD} \sim 0.5 - 1.4\,M_\odot$. WDs are also a powerful laboratory to test new physics.
The subsequent evolution of a white dwarf is essentially a cooling process. Since nuclear burning has ceased, the star gradually loses its residual thermal energy.
Axions again can provide an additional cooling channel in white dwarfs. In this environment, the dominant production mechanism is electron bremsstrahlung, $e + Ze \to e + Ze + a$, rather than the Primakoff process. Therefore, white-dwarf cooling mainly constrains the axion–electron coupling g_{aee}.
The effect of axion cooling can be searched for in two main ways. The first is through the white-dwarf luminosity function, namely the number density of white dwarfs as a function of their bolometric magnitude (that is the magnitude over all possible wavelengths). The second is through pulsating white dwarfs, where an anomalous cooling rate can appear as a drift in the pulsation period.

White-Dwarfs Function Luminosity

The white-dwarf luminosity function (WDLF) describes the density of white dwarfs in the Galactic neighborhood as a function of their brightness, usually expressed in terms of the absolute bolometric magnitude $M_{\rm bol}$. Assuming a constant birthrate, the slope of the WDLF directly reflects the cooling speed: if cooling is slower, more white dwarfs accumulate in a given magnitude interval.
At early times, when white dwarfs are still hot, energy loss is dominated by neutrino emission, mainly through plasmon decay, $\gamma \to \nu\bar{\nu}$.
At later times, photon emission from the surface becomes the dominant cooling mechanism. In the absence of axions, and for a simple carbon–oxygen white dwarf model, the intermediate-luminosity part of the WDLF is well described by the Mestel cooling law,

$$\log\left(\frac{dN}{dM_{\rm bol}}\right) = \frac{2}{7}M_{\rm bol} - 6.84 + \log B_3,$$

where B_3 is the white-dwarf birthrate normalized to $10^{-3}\,{\rm pc}^{-3}\,{\rm Gyr}^{-1}$. Axions modify both the normalization and the shape of the WDLF. An additional axion luminosity L_a increases the total cooling rate, so that white dwarfs spend less time in a given luminosity interval. This changes the predicted number of white dwarfs at different magnitudes. Comparing stellar-evolution models including axion emission with the observed WDLF therefore gives a bound on g_{aee}.
A recent self-consistent evolutionary analysis, comparing theoretical models with the WDLF of the Galactic disk obtained from SDSS, Ref. [122] and SuperCOSMOS data, Ref. [123], found the bound

$$g_{aee} < 2.8 \times 10^{-13}$$

at a nominal 99% confidence level. For DFSZ axions, the original WDLF limit can be translated into

$$f_a > 3 \times 10^8\,{\rm GeV}, \qquad m_a < \frac{19\,{\rm meV}}{2\sin^2\beta}.$$

Pulsating White Dwarfs

White dwarfs can also become pulsationally unstable when they cross the instability strip in the Hertzsprung–Russell diagram. The instability strip is a quasi-vertical region in the Hertzsprung–Russell diagram: stars contracting increase their opacity, this leads to an increase in pressure and a subsequent increase in their radius, bringing a more transparent stellar medium. When WD enter this regime they become ZZ Ceti stars. Their pulsation periods are extremely stable and can be used as stellar clocks.
As a white dwarf cools, its internal structure slowly changes, and the pulsation period drifts with time. The rate of period change therefore provides a direct measure of the cooling speed. If an additional cooling channel such as axion emission is present, the star cools faster and the observed period drift

can be larger than expected from standard cooling alone.
A particularly important example is the pulsating white dwarf G117-B15A. Its measured period drift is larger than the value expected from standard white-dwarf cooling. If interpreted entirely as axion bremsstrahlung, it would require $g_{aee} = (5.66 \pm 0.57) \times 10^{-13}$, from Ref. [124]. Similar values have been inferred from other pulsating white dwarfs.
However, this interpretation is problematic. Such a large value of g_{aee} would imply a strong axion contribution to white-dwarf cooling and should also show up clearly in the WDLF. It is in tension with the WDLF bounds and even more strongly with the TRGB constraint on the axion–electron coupling. Since both the WDLF and TRGB limits rely on less delicate physics than the modeling of pulsating white dwarfs, the apparent cooling hint from ZZ Ceti stars is usually regarded as more likely due to unresolved systematics in pulsation physics rather than evidence for axions.

3.3.5 Neutron-Star and Axions

Neutron stars (NS) provide a natural laboratory for axions because QCD axions couple to nucleons. In this environment, the relevant production mechanism is mainly nucleon bremsstrahlung, $N + N \to N + N + a$, and, depending on the composition of the core, especially $n + n \to n + n + a$. Axion emission would act as an additional cooling channel and would modify the thermal evolution of NS.
The cooling of a neutron star proceeds in different stages. During the first seconds after core collapse, the star cools through an intense neutrino burst. Shortly afterwards, it becomes transparent to neutrinos, but continues to cool by neutrino volume emission. For the first $\sim$ 100 yr the crust thermally relaxes. After this stage, the stellar interior is approximately isothermal and the cooling is controlled by the core. At late times, around $\sim 10^6$ yr, neutrino emission becomes inefficient and photon emission from the surface takes over.
Axion bounds from neutron-star cooling are obtained by comparing theoretical cooling curves with observations of neutron-star ages and surface temperatures. If axion emission is too strong, the star cools too rapidly and becomes colder than observed. This gives constraints on the axion–nucleon couplings, especially the axion–neutron and axion–proton Yukawa couplings,

$$g_{ann} = C_{ann}\frac{m_N}{f_a}, \qquad g_{app} = C_{app}\frac{m_N}{f_a}.$$

These constraints can then be translated into bounds on f_a and m_a once a specific QCD axion model is chosen.

Young Neutron Stars

Young neutron stars, with ages from a few hundred to a few tens of thousands of years, are still in the neutrino-cooling era. A relevant example is given by the neutron star in the supernova remnant Cassiopeia A (Cas A).
Cas A is a very young neutron star, with an age of about $t \simeq 340$ yr. It was initially interesting because its surface temperature appeared to decrease faster than expected from standard cooling processes such as modified Urca emission (that are those processes that produce neutrinos inside NS cores in non superfluid conditions). A possible explanation involved the onset of neutron superfluidity in the core, which enhances neutrino emission through pair breaking and formation.
The relevant pairing channel for neutrons is the triplet $^3P_2$ state. Here the notation means that the two neutrons form Cooper pairs with total spin $S = 1$, orbital angular momentum $L = 1$ corresponding to a P wave, and total angular momentum $J = 2$.
When the temperature drops below the critical temperature for this superfluid transition, neutron Cooper pairs begin to form and break. This produces efficient neutrino emission through the pair-breaking-and-formation process. The same physics also affects axion emission, because axions can couple to the spin fluctuations of nucleons.
Although the originally inferred rapid cooling of Cassiopeia A is now considered likely to be affected by detector-calibration systematics, the source can still be used to place order-of-magnitude constraints on axion cooling. Cooling simulations including axion losses give, for representative QCD axion models, Ref. [125]: $f_a > 0.3 \times 10^8$ GeV for KSVZ axions, and $f_a > 4.5 \times 10^8$ GeV for a representative

DFSZ choice with $\cot\beta = 10$. However, these limits are strongly model-dependent because the axion–neutron coupling can accidentally become small or vanish for particular choices of the QCD axion parameters.

Old Neutron Stars

Older neutron stars, with ages of order $\sim 10^6$ yr, are close to the transition between neutrino-dominated and photon-dominated cooling. In this regime, the neutrino luminosity starts to be suppressed by the decreasing core temperature, while the photon luminosity from the surface becomes increasingly important. Therefore, an additional axion luminosity can have a visible impact on the late-time cooling curve.

The strongest and most reliable neutron-star limits come from old isolated neutron stars, in particular members of the Magnificent Seven together with the pulsar J1605, see Ref. [126]. Their measured luminosities and kinematic ages can be compared with numerical cooling sequences including axion emission. If the axion–nucleon coupling is too large, the predicted luminosity at the measured age becomes too small.

For DFSZ axions, the resulting bound can be written approximately as

$$m_a^{\mathrm{DFSZ}} < \left(33 - 17\sin^2\beta - 1.8\sin^4\beta\right)\,\mathrm{meV}.$$

For the reference value $\sin^2\beta = \frac{1}{2}$, this gives approximately $m_a < 24$ meV. For KSVZ axions, the corresponding conservative limit is

$$m_a < 19\ \mathrm{meV}, \qquad f_a > 3\times 10^8\ \mathrm{GeV}.$$

Overall, neutron-star cooling constrains axion emission by requiring that the additional axion luminosity does not make observed neutron stars colder than they are. Young neutron stars mainly probe the axion–neutron coupling through neutron bremsstrahlung and superfluid effects, while old neutron stars give the most robust QCD-axion bounds after translating the nucleon couplings into model-dependent limits on f_a and m_a.

3.3.6 Supernova Neutrinos

Bounds on axion models can also be derived from supernovae. The traditional supernova bound on axions comes from the duration of the neutrino signal observed from SN 1987A, see Ref. [127]. Neutrinos from this supernova were detected by different experiments, in particular IMB, Kamiokande-II and BUST experiments, mainly through inverse beta decay $\bar{\nu}_e + p \to n + e^+$.

After core collapse, the remnant is a proto-neutron star with nuclear density and temperature of order $T \sim 30$ MeV. In this environment, neutrinos are trapped and diffuse out on a timescale of several seconds. If a new weakly interacting particle, such as an axion, carried away too much energy from the core, the proto-neutron star would cool faster and the neutrino burst would have been shorter.

For QCD axions, the relevant production mechanism is again mainly nucleon bremsstrahlung, $N+N \to N+N+a$, controlled by an effective axion–nucleon coupling g_{aNN}. As in the trapping regime discussed for the solar bounds, the behavior depends on the size of this coupling. There are two regimes in which this argument can be eluded. If g_{aNN} is very small: axion emission is negligible and the neutrino signal is unchanged. For larger couplings, axions are produced efficiently and free-stream out of the core, removing energy and shortening the neutrino burst. If we have even larger couplings, axions become trapped, similarly to neutrinos, and are emitted only from an "axion sphere". In that trapping regime, the effect on the neutrino signal becomes smaller again. However, larger couplings of this type are not necessarily harmless. First, they may play an important role during the SN collapse phase. Second, in the water Cherenkov detectors that registered the SN 1987A neutrinos, these axions would have interacted with oxygen nuclei, leading to the release of $\gamma-$rays and causing too many events.

The most useful regime for the standard SN 1987A cooling bound is therefore the free-streaming

regime. A commonly used analytic criterion is that the axion energy-loss rate per unit mass should not exceed roughly: $\epsilon_a \lesssim 10^{19}\,\mathrm{erg\,g^{-1}\,s^{-1}}$, as done in Ref. [128]. this corresponds to an axion luminosity $L_a \sim \epsilon_a M_{\rm core} \sim 2\times 10^{52}\,\mathrm{erg\,s^{-1}}$. This is comparable to the neutrino luminosity during the cooling phase. If axion losses were much larger than this, they would noticeably shorten the neutrino burst.
To put a bound on g_{aNN} we need to study nuclear scattering. Early estimates modeled the nucleon–nucleon interaction with a one-pion-exchange (OPE) potential however, after including also phenomenological improvements, and information from nucleon–nucleon scattering data, one gets for KSVZ axions the approximate bound, given in Ref. [129]

$$f_a \gtrsim 4\times 10^8\,\mathrm{GeV}, \qquad m_a \lesssim 16\,\mathrm{meV}.$$

For DFSZ axions, the same bound depends on the parameter β. It can be approximated as

$$m_a^{\rm DFSZ} < \left(23 + 6\sin^2\beta - 38\sin^4\beta + 22\sin^6\beta\right)\mathrm{meV}.$$

For the reference choice $\sin^2\beta = \frac{1}{2}$ this gives $m_a < 19\,\mathrm{meV}$.
These bounds could be further improved: in fact other processes could lead to axion emission such as $\pi^- p \to na$, nevertheless the importance of axion production form pion's scattering depends on the conditions found in SNs cores, therefore while one cannot exclude that the emission rate is enhanced by pions, one also cannot rely on them for a trustworthy bound. We only mention that including axions in SN could modify the explosion physics of SN. Furthermore we would expect, in analogy with neutrinos a diffuse SN axion background coming form axions produced in stellar collapse. These axions could then convert in photons in astrophysical magnetic fields leading to interesting signals.

3.3.7 Black-Hole Superradiance

Superradiance is a wave-amplification phenomenon that occurs when a wave interacts with a rotating dissipative system. In the case of a Kerr black hole, an incoming wave can be reflected with a larger amplitude than the one it initially had, extracting energy and angular momentum from the black hole rotation, in a similar way to the Penrose's process see Refs. [130, 131]. The kinematic condition for amplification is

$$\omega < m\Omega_H, \tag{3.56}$$

where ω is the wave frequency, m is the azimuthal angular momentum, and Ω_H is the angular velocity of the black-hole horizon. Physically, the wave is amplified when its angular phase velocity is smaller than the angular velocity of the horizon.
This phenomenon is not exclusive to black holes. A useful analogy is Cherenkov radiation: a charged particle moving through a medium faster than the phase velocity of light in that medium emits electromagnetic radiation. In that process, part of the kinetic energy of the particle is converted into radiation. Similarly, in rotational superradiance, rotational energy is transferred from the rotating object to the wave.
For black holes, the effect becomes especially important for massive bosonic fields, in fact they can form bound states around the black hole, analogous to the orbitals of the hydrogen atom. The region outside the black hole then acts as a kind of gravitational cavity: the wave remains trapped near the black hole and is repeatedly amplified by superradiance. This leads to the exponential growth of a bosonic cloud around the black hole, at the expense of the black hole rotational energy.
The efficiency of the process is controlled by the gravitational analogue of the fine-structure constant,

$$\alpha_g = G_N M_{\rm BH} m_a, \tag{3.57}$$

where $M_{\rm BH}$ is the black-hole mass and m_a is the axion mass. Superradiance is most efficient when the Compton wavelength of the boson is comparable to the gravitational radius of the black hole, $\lambda_C \sim r_g$, Equivalently, a black hole is most sensitive to bosons with masses of order $m_a \sim \frac{1}{r_g} \sim \frac{1}{G_N M_{\rm BH}}$ This immediately shows that black holes with different masses probe different axion-mass ranges: stellar-mass black holes are sensitive to heavier axions, while supermassive black holes probe much lighter axions, giving the possibility to probe GUT and Planckian axions.

The growth of the superradiant cloud extracts spin from the black hole. Therefore, the observation of old, rapidly spinning black holes can be used to exclude bosons whose masses would make superradiance efficient. If such an axion existed and its interactions were sufficiently weak, the bosonic cloud would have grown on astrophysical timescales and spun down the black hole. Observing a black hole that is still rapidly rotating therefore rules out the corresponding range of m_a.
For stellar-mass black holes, spin measurements in X-ray binary systems exclude an axion-mass window of approximately, from Ref. [132]

$$0.3 \times 10^{-12}\,\mathrm{eV} \lesssim m_a \lesssim 7 \times 10^{-12}\,\mathrm{eV}. \tag{3.58}$$

The bound does not depend only on the axion mass, but also on its interactions. If the self-interactions are too large, the axion cloud can collapse or saturate before extracting a significant amount of spin from the black hole. In that case, the superradiance bound can be weakened or may not apply. For very weakly interacting axion-like particles, instead, the growth of the cloud is more efficient and the black-hole spin bounds are more robust.

3.4 Direct searches for Axions

Before discussing direct searches for axions and axion-like particles, it is useful to clarify how axion dark matter should be treated in terrestrial experiments. In the mass range relevant for many direct-detection searches, axions are extremely light and have a very large occupation number per Compton wavelength. For this reason, the axion dark-matter background can be described with its classical solution.
At finite volume $V = L^3$, the Hamiltonian of a free scalar field can be written as

$$H = \sum_n \omega_n a_n^\dagger a_n, \tag{3.59}$$

where $a_n^\dagger$ and a_n are the creation and annihilation operators of the mode n, and ω_n is its energy. The field is decomposed as

$$\phi(x,t) = \sum_n \frac{1}{L^{3/2}} \frac{1}{\sqrt{2\omega_n}} \left(a_n e^{ip_n\cdot x - i\omega_n t} + a_n^\dagger e^{-ip_n\cdot x + i\omega_n t} \right). \tag{3.60}$$

A classical field corresponds to a state with a very large number of particles in the relevant modes. More precisely, we can consider a coherent state for each mode, that is an eigenstate of $a_n^\dagger$ for every n:

$$|N\rangle = \sum_n e^{-N_n/2}\, e^{\sqrt{N_n} a_n^\dagger} |0\rangle, \tag{3.61}$$

where N_n is the average occupation number of the mode n. In such a state, it is easy to show

$$\langle N|H|N\rangle = \sum_n \omega_n N_n. \tag{3.62}$$

One can then compute $\langle\phi\rangle$ and $\langle\phi^2\rangle$. Because particles come in waves, the energy at a fixed position is not constant in time. Let us define:

$$\langle H\rangle_T = \frac{1}{T}\int dt H(t) = \int d^3v f(v)\omega_v \bar{n} \tag{3.63}$$

$\bar{n}$ is the number density and $f(v)$ the distribution velocity. It is then easy to see that:

$$\langle H\rangle_T = \frac{1}{V}\sum_n \omega_n N_n. \qquad N_n = f(v)\bar{n}\left(\frac{2\pi}{m_a}\right)^3 \tag{3.64}$$

We now expand around a random position x0 and a time t0. This means that all of the relative phases between the sums will be basically random, for simplicity we assume $f(v)$ to be isotropic.

$$\langle\phi\rangle = \sum_{i_v\, i_\theta\, i_\phi} \left(\frac{2\pi}{mL}\right)^{3/2} \sqrt{\frac{2\bar{n} f_{i_v}}{\omega_{i_v}}} \cos\left(\omega_{i_v} t + \phi_{i_v} + \phi_{i_\theta} + \phi_{i_\phi}\right) \tag{3.65}$$

The velocity distribution has a characteristic scale v_0, therefore we can discretize the modulus of the velocity in bins with width such that Eq. (3.65) remains constant, inside this interval we assume that the sum over the velocity and its angles are completely random. Between velocity v and $v + \Delta v$ the number of random phases that we are adding is:

$$N = 4\pi v^2 \Delta v \left(\frac{Lm}{2\pi}\right)^3 \tag{3.66}$$

Summing over this N random phases one gets a factor proportional to $N^{1/2}$. We get:

$$\begin{aligned} \langle\phi\rangle =& \sum_{a_v} \left(\frac{2\pi}{m_a L}\right)^{3/2} \sqrt{\frac{2\bar{n} f_{a_v}}{m_a}} N^{1/2} \cos\left(\omega_{a_v} t + \phi_{a_v}\right) = \\ \simeq& \frac{\sqrt{\bar{\rho}}}{m} \sum_{a_v} \sqrt{4\pi v^2 f(v) \Delta v} \cos\left(\omega_{a_v} t + \phi_{a_v}\right) \end{aligned} \tag{3.67}$$

Where a_v is the index for the modulus of the velocity, and we assumed that ω_v is well indentified with the axion mass. In an analogous way one can compute also $\langle\phi^2\rangle$. The final results are:

$$\begin{aligned} \langle\phi\rangle &= \frac{\sqrt{\bar{\rho}}}{m} \sum_{a_v} \sqrt{4\pi v^2 f(v) \Delta v}\, \alpha_{a_v} \cos\left(\omega_{a_v} t + \phi_{a_v}\right) \\ \langle\phi^2\rangle &= \frac{\bar{\rho}}{m^2} \sum_{a_v} 4\pi v^2 f(v) \Delta v \cos^2\left(\omega_{a_v} t + \phi_{a_v}\right) \end{aligned} \tag{3.68}$$

These sums converge independently of the value of Δv as long as Δv is small. The result is that the axion field is well approximated by $a(t) = \frac{\sqrt{2\rho}}{m} \cos mt$ as long as coherence is still present, that for intervals of time $\tau \sim \frac{4\pi}{mv^2}$. Another important preliminary ingredient is given by the interaction Lagrangian:

$$\mathcal{L}_{a\gamma\gamma} = -\frac{1}{4} g_{a\gamma\gamma}\; a F^{\mu\nu} \tilde{F}_{\mu\nu} = -a\; g_{a\gamma\gamma}\; E \cdot B \tag{3.69}$$

and the Maxwell's equations become:

$$\partial_\mu F^{\mu\nu} = J^\nu - g_{a\gamma\gamma} (\partial_\mu a) \tilde{F}^{\mu\nu} \tag{3.70}$$

in components:

$$\begin{aligned} &\nabla \cdot E = \rho - g_{a\gamma\gamma} \nabla a \cdot B \\ &\nabla \times B - \frac{\partial E}{\partial t} = J + g_{a\gamma\gamma} (\dot{a} B + \nabla a \times E) \\ &\nabla \cdot B = 0 \\ &\nabla \times E - \frac{\partial B}{\partial t} = 0 \end{aligned} \tag{3.71}$$

Having this equations, we can solve them in the vacuum. It will be useful to study the probability of photon-axion conversion inside an external magnetic field, a relevant situation for different direct axion searches. Suppose that the external magnetic field has the form $B_{ext} = B\hat{x}$, and a photon propagates in the $z-$direction. We fix the radiation gauge $A_0 = 0$ and $\nabla \cdot A = 0$. For a wave propagating along the $z-$direction we write:

$$A(z,t) = A_x(z,t)\hat{x} + A_y(z,t)\hat{y} \tag{3.72}$$

It is easy to understand that having $E = \partial_t A$, the component that couples with the axion field is only A_x, that we rename $A_\parallel$. We can then write the relevant Lagrangian density:

$$\mathcal{L} = \frac{1}{2}\left((\partial_t A_\parallel)^2 - (\partial_z A_\parallel)^2\right) + \frac{1}{2}\left((\partial_t a)^2 - (\partial_z a)^2 - m_a^2\right) - g_{a\gamma\gamma}\ a\ B_{ext}\partial_t A_\parallel \tag{3.73}$$

The equation of motion are:

$$\begin{aligned}(\partial_t^2 - \partial_z^2)A_\parallel &= g_{a\gamma\gamma}B_{ext}\ \partial_t a\\ (\partial_t^2 - \partial_z^2 + m_a^2)a &= -g_{a\gamma\gamma}B_{ext}\ \partial_t A_\parallel\end{aligned} \tag{3.74}$$

Now imposing $A_\parallel(z,t) = A_\parallel(z)e^{-i\omega t}$ and $a_\parallel(z,t) = a_\parallel(z)e^{-i\omega t}$ we get:

$$\begin{aligned}(\omega^2 + \partial_z^2)A_\parallel &= i\omega\ g_{a\gamma\gamma}B_{ext}\ a\\ (\omega^2 + \partial_z^2 - m_a^2)a &= -i\omega g_{a\gamma\gamma}B_{ext}\ A_\parallel\end{aligned} \tag{3.75}$$

We first pass $a \to ia$. Finally we rewrite $A_\parallel = A(z)e^{i\omega z}$ and $a_\parallel = a(z)e^{i\omega z}$ and neglecting second derivative terms (we are assuming that $A(z)$ and $a(z)$ are slowly varying) we arrive at:

$$\begin{aligned}2i\omega\partial_z A_\parallel &= \omega\ g_{a\gamma\gamma}B_{ext}\ a\\ 2i\omega\partial_z a &= m_a^2 a + \omega g_{a\gamma\gamma}B_{ext}\ A_\parallel\end{aligned} \tag{3.76}$$

From which we get:

$$i\frac{d}{dz}\begin{pmatrix}A_\parallel\\ a\end{pmatrix} = \begin{pmatrix}0 & \frac{g_{a\gamma\gamma}B_{ext}}{2}\\ \frac{g_{a\gamma\gamma}B_{ext}}{2} & -\frac{m_a^2}{2\omega}\end{pmatrix}\begin{pmatrix}A_\parallel\\ a\end{pmatrix} \tag{3.77}$$

It is now easy to find the amplitude for the photon-axion conversion, it is enough to find the two eigenstates of the matrix, express the photon as a superposition of the two and evolve the states. The final solution is:

$$P_{\gamma\to a}(L) = \sin^2 2\theta \sin^2\frac{\Delta_{osc}L}{2} \tag{3.78}$$

Where $\sin 2\theta = (g_{a\gamma\gamma}B_{ext})/\Delta_{osc}$ and $\Delta_{osc} = \sqrt{m^4 + g_{a\gamma\gamma}^2 B_{ext}^2}/2\omega$. In the small mixing limit we get:

$$P_{\gamma\to a}(L) = \left(\frac{g_{a\gamma\gamma}B_{ext}}{q}\right)^2 \sin^2\frac{qL}{2} \tag{3.79}$$

$q = \frac{m_a^2}{2\omega}$ is the momentum mismatch between the photon and the axion.

3.4.1 Light-Shining-Through-Walls Experiments

As anticipated above, a laboratory strategy to search for axions and axion-like particles (ALPs) is based on photon regeneration in external magnetic fields, commonly known as a light-shining-through-walls (LSW) experiment. In this setup, a high-intensity photon beam, usually provided by a laser, propagates through a strong transverse magnetic field. In the production region, photons can oscillate into axions through the axion–photon coupling. An opaque wall blocks the ordinary photons, while axions, interacting only very weakly with matter, pass through it. In a second magnetic region, called the regeneration region, axions may reconvert into photons, which can then be detected by low-noise optical sensors. Therefore we are interested in the chain $\gamma \to a \to \gamma$. The basic principle is shown in the figure below. We have already found the probability to have photon-axion conversion in this set up in Eq. (3.79). When $qL \ll 1$, namely when the length of the magnetic region is much shorter than the oscillation length, the conversion remains coherent and the probability reduces to

$$P_{\gamma\to a}(L) \simeq \left(\frac{g_{a\gamma\gamma}BL}{2}\right)^2. \tag{3.80}$$

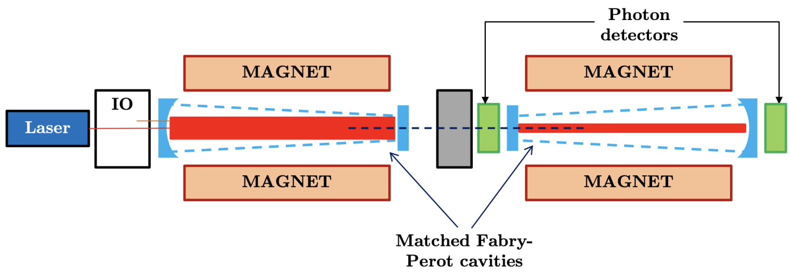

Figure 3.8: A schematic setup of a LSW experiment. Figure taken from Ref. [133].

Since an LSW experiment requires both photon-to-axion conversion and axion-to-photon reconversion, the total regeneration probability is just the square of Eq. (3.80):

$$P_{\gamma\to a\to\gamma} = P_{\gamma\to a}P_{a\to\gamma} \simeq \left(\frac{g_{a\gamma\gamma}BL}{2}\right)^4 .$$

In resonantly enhanced LSW experiments, optical cavities may be placed both in the production and regeneration regions. These Fabry–Perot cavities increase the effective number of photon passes in the magnetic field and enhance the signal by the power build-up factors β_P and β_R, giving

$$P_{\gamma\to a\to\gamma} \simeq \left(\frac{g_{a\gamma\gamma}BL}{2}\right)^4 \beta_P\beta_R.$$

The production cavity increases the number of photons available for conversion, while the regeneration cavity enhances the reconversion of axions into photons at the resonant frequency. In practice, the two cavities must be mode-matched and phase-locked, which represents one of the main experimental challenges.

Several LSW experiments have already set bounds on the axion–photon coupling in the range $g_{a\gamma\gamma} \sim 10^{-6} - 10^{-7}\,\mathrm{GeV}^{-1}$. Important examples include ALPS at DESY and OSQAR at CERN, Refs. [134, 135], which use strong accelerator dipole magnets. More advanced setups, such as ALPS II, see Ref. [136], exploit resonant regeneration and long strings of superconducting magnets to improve the sensitivity by several orders of magnitude, targeting couplings down to $g_{a\gamma\gamma} \lesssim 2\times 10^{-11}\,\mathrm{GeV}^{-1}$ for $m_a < 10^{-4}$eV. Other implementations have also been proposed or performed at microwave Ref. [137], giving similar performances to the optical experiments and X-ray frequencies as in Ref. [138], although optical and microwave LSW experiments currently provide the most competitive laboratory sensitivities, given the low number of photons available for x-ray LSW experiments.

A complementary class of purely laboratory searches is provided by polarization experiments. In this case no wall is used: instead, a linearly polarized laser beam propagates through a strong external magnetic field. Since only the photon polarization component parallel to the magnetic field couples to the axion field, photon–axion oscillations can induce a small depletion of this component. This produces a rotation of the polarization plane, an effect known as dichroism. At the same time, the virtual exchange of axions can generate a phase delay between the parallel and perpendicular polarization components, leading to an ellipticity of the outgoing beam, and to a different refraction index, namely to birefringence. These measurements are also sensitive to the standard QED effect of vacuum magnetic birefringence, described by the Euler–Heisenberg effective interaction. Experiments such as PVLAS, Ref. [139] and future experiments such as BMV, Ref. [140] exploit high-finesse optical cavities and strong magnetic fields to increase the optical path length and improve the sensitivity to very small polarization rotations or ellipticities. Compared with LSW experiments, polarization experiments do not require axion reconversion behind a wall, but their interpretation can be more model-dependent, since the signal appears as a small modification of the optical properties of the vacuum.

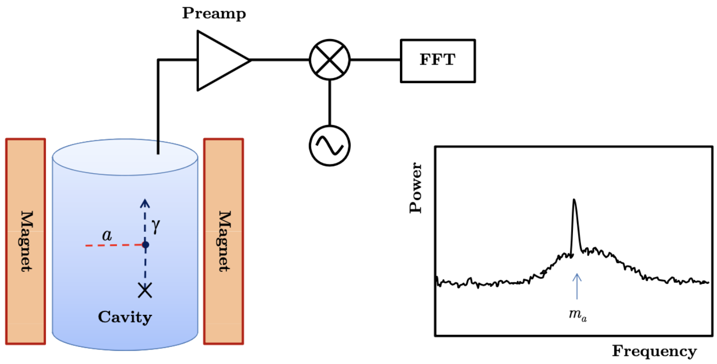

Figure 3.9: Schematic arrangement of an axion haloscope. If the axion has the same frequency of the cavity, the axion will show as a narrow peak in the power spectrum extracted from the cavity.

3.4.2 Axion Haloscopes

As we have seen before, axions can be treated as coherent oscillating fields. Furthermore, if we assume that axions constitute a sizeable fraction of the galactic dark matter halo, they are expected to be present in the laboratory as a non-relativistic field. The local dark matter density is usually taken to be of order $\rho_{\rm DM} \simeq 0.3\,{\rm GeV/cm^3}$, so that, for axion masses in the μeV range, the corresponding number density can be very large. The oscillation frequency is essentially fixed by the axion energy, which is well approximated by the axion mass, with a small correction due to the halo velocity distribution. The velocity spread depends on the galactic halo model used, but it is usually of order

$$\frac{\delta\nu_a}{\nu_a} \sim v^2 \sim 10^{-6}.$$

This means that, in the frequency domain, axions appear as a narrow spectral line, which is an important feature of axion dark matter searches [133].
One of the standard experimental techniques that exploits this framework to search for axions is the haloscope. A conventional haloscope consists of a high-quality-factor resonant cavity placed inside a strong static magnetic field, see Fig. (3.9) for a schematic representation. Through the Primakoff effect, galactic axions can convert into real photons in the magnetic field. If the resonant frequency of the cavity matches the axion oscillation frequency, and therefore the frequency of the photon produced by the axion, the conversion is enhanced and the produced photons populate a resonant electromagnetic mode of the cavity, see . The signal is then extracted through an antenna or a coupled port and amplified by a low-noise radio-frequency detection chain, see Refs. [141, 142].
For a resonant cavity haloscope, the expected signal power can be written schematically as

$$P_s \sim g_{a\gamma\gamma}^2 \rho_a B^2 V C \frac{Q}{m_a},$$

where $g_{a\gamma\gamma}$ is the axion–photon coupling, ρ_a is the local axion density, B is the external magnetic field applied in the cavity, V is the cavity volume, Q is the quality factor of the resonator, and C is a geometric form factor describing the overlap between the cavity mode and the external magnetic field. More explicitly, for a cavity mode with electric field $\mathbf{E}$, the form factor is

$$C = \frac{\left|\int dV\, \mathbf{E}\cdot\mathbf{B}\right|^2}{B^2 V \int dV\, |\mathbf{E}|^2}.$$

Therefore, the most useful modes are those whose electric field is well aligned with the external magnetic field.
Since the axion mass is not known, the cavity frequency must be tunable. This is usually achieved by inserting movable elements that change the geometry of the cavity. The experiment then scans over

axion masses by tuning the cavity step by step and integrating for a certain time at each frequency. A real axion signal would appear as a narrow excess of power above the noise background, with a linewidth determined by the dark matter velocity dispersion.
The sensitivity of a haloscope depends strongly on the main experimental parameters. The figure of merit can be defined as proportional to the time needed to scan a fixed mass range down to a given signal-to-noise ratio. In the absence of systematics, assuming therefore that increasing the exposure time improves the quality of the readout, one has

$$F_{\text{halo}} \propto \rho_a^2 g_{a\gamma\gamma}^4 m_a^2 B^4 V^2 \frac{Q}{T_{\text{sys}}^2},$$

where T_{sys} is the effective noise temperature of the detection chain. This expression shows why large magnetic fields, large volumes, high quality factors, and very low-noise amplifiers are crucial. This is true up to a certain quality factor. Indeed, as discussed above, the axion signal itself has an effective quality factor

$$Q_a \sim \frac{\nu_a}{\delta\nu_a} \sim 10^6.$$

If the cavity quality factor becomes larger than Q_a, part of the axion signal lies outside the cavity bandwidth, and therefore part of the signal is lost. For this reason, the relevant enhancement is effectively controlled by the smaller of the cavity quality factor and the axion quality factor.
The most established haloscope experiment is ADMX, which uses a large microwave cavity inside a strong solenoidal magnetic field and has reached sensitivity to QCD axion models in the μeV mass range, Ref. [141].

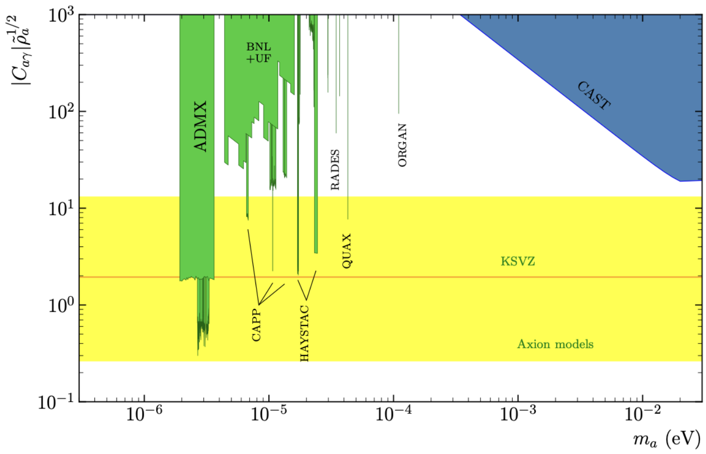

Figure 3.10: ALP parameter space with its exlcuded regions, focusing on haloscopes. The horizontal line corresponds to the QCD axions.

ADMX represents the prototype of the conventional cavity haloscope approach. Other experiments follow the same basic idea but target different mass ranges or use different technological improvements. HAYSTAC, for example, searches at somewhat higher masses with a smaller cavity and a low-noise microwave detection chain, Ref. [142]. Experiments such as CAPP and QUAX also use resonant microwave cavities and low-noise detection chains, exploring other regions of the axion parameter space Refs. [143, 144].
At higher axion masses, the resonant frequency increases and the required cavity volume decreases, which reduces the signal power. This makes high-frequency searches more challenging. Several strategies have been proposed to overcome this limitation, including the use of multiple phase-matched cavities, arrays of smaller resonators, higher-order modes, dielectric structures, or more complex resonant geometries. Examples of this direction include RADES, which studies arrays of coupled cavities, and other high-frequency programs such as ORGAN [145, 146]. At lower axion masses, on the other hand, the required cavity volume becomes very large, motivating proposals based on larger magnets or alternative resonator geometries.

For axion masses well below the μeV scale, conventional microwave cavity haloscopes become less effective. In this low-mass region, an alternative strategy is to search for the tiny oscillating magnetic field induced by the axion dark matter field in the presence of a strong static magnetic field. This is the basic idea of so-called DM-radio experiments.
A typical setup consists of a large external magnet and a pick-up coil placed inside it. The axion field induces a small time-dependent magnetic flux through the coil, which can then be read out with a sensitive magnetometer.
The signal can be resonantly enhanced by coupling the pick-up coil to a tunable LC circuit. Compared with a microwave cavity, an LC resonator is easier to tune at very low frequencies, making this technique particularly suitable for light axions. A broadband, non-resonant mode of operation is also possible, although in that case the signal enhancement is smaller.
Experiments such as ABRACADABRA, SHAFT, BEAST, ADMX SLIC, and BASE explore different implementations of this idea, see Refs. [147, 148, 149, 150]. Current searches are mainly sensitive to axion masses in the approximate range

$$10^{-12}\,\mathrm{eV} \lesssim m_a \lesssim 10^{-7}\,\mathrm{eV}.$$

With sufficiently large magnetic volumes, strong magnetic fields, and low-noise readout, DM-radio techniques could become sensitive to QCD axion dark matter for masses below about $m_a \lesssim 10^{-6}\,\mathrm{eV}$.

Dielectric Haloscopes

As already said, conventional cavity haloscopes become progressively more difficult at high axion masses. For this reason, alternative haloscope concepts have been proposed to extend dark matter axion searches above the mass range most naturally accessible with large microwave cavities.
A particularly interesting approach is the dielectric haloscope, whose main example is the MADMAX proposal [151]. The idea is to place a series of dielectric disks in front of a reflecting mirror, all immersed in a strong external magnetic field. In the presence of the axion dark matter field, the axion–photon coupling induces an effective oscillating electric field proportional to the external magnetic field. When this field crosses an interface between two media with different dielectric constants, the (modified) electromagnetic boundary conditions require the parallel component of the electric field to be continuous. As a consequence, to maintain continuity, electromagnetic waves are emitted at each dielectric interface.
If the positions and thicknesses of the dielectric disks are chosen appropriately, the waves emitted by the different interfaces can interfere constructively in one direction. The dielectric stack therefore acts in a way similar to a phased antenna array: the signal is not enhanced by a single narrow cavity resonance, but by the coherent superposition of radiation emitted at many interfaces. The resulting power enhancement is usually described by a boost factor β^2, so that the emitted power may be written schematically as

$$P_s \sim g_{a\gamma\gamma}^2 \rho_a B^2 A\, \beta^2,$$

where A is the transverse area of the disks and β^2 encodes the coherent enhancement produced by the dielectric structure.
The important difference with respect to a resonant cavity haloscope is that the relevant size is no longer the volume of a microwave cavity tuned to a single mode. Instead, the sensitivity depends mainly on the magnetic field, the transverse area of the dielectric disks, and the achievable boost factor. This makes dielectric haloscopes especially attractive at higher frequencies, where the volume of a conventional cavity would become very small.
Another advantage is that the frequency response can be engineered by changing the separations between the disks. In this way, one can tune the dielectric haloscope to different axion masses and choose a compromise between a large boost factor over a narrow bandwidth or a smaller boost factor over a broader bandwidth. This tunability is essential because, as in all haloscope searches, the axion mass is not known in advance.

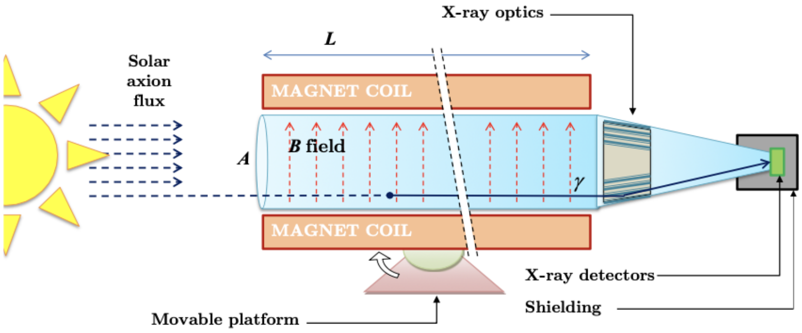

Figure 3.11: Schematic arrangement of a helioscope. Solar axions are converted into photons by the transverse magnetic field inside the bore of a powerful magnet, and then the resulting beam of photons is focused into a low background detector.

3.4.3 Helioscope Searches

The same solar axion flux that leads to energy-loss bounds can also be searched for directly in laboratory experiments. The basic idea is the same as in the LSW experiments: axions produced in the solar core travel to the Earth and are converted back into photons inside a strong transverse magnetic field, see Eq. (3.11). We have seen that for an axion of energy ω propagating through a magnetic field B over a length L, the probability of conversion into a photon is given by Eq. (3.79), however for helioscopes, the regime $qL \ll 1$ is given for axions with masses $m_a < 0.02$ eV, For larger axion masses, the axion and photon waves lose coherence along the magnet length. It is clear why helioscopes, as haloscopes and LSW experiments require very strong magnetic fields and long magnet lengths. In practice, the magnet is mounted on a movable platform in order to track the Sun for the largest possible fraction of the day. Solar axions enter the magnet bore almost parallel to its axis and may be converted into X-ray photons inside the transverse magnetic field. These photons are then searched for at the opposite end of the magnet. Modern helioscope designs use X-ray focusing optics placed after the magnet bore, so that the quasi-parallel photon beam is concentrated onto a small spot on a low-background detector. This is important because the signal is localized in a small region of the detector plane, whereas the background can be measured either outside the signal spot or during periods in which the magnet is not pointing at the Sun. The signal rate scales as $R \propto \Phi_a P_{a\to\gamma}$, since the solar Primakoff axion flux and the conversion probability scale both with $g^2_{a\gamma\gamma}$, we have that $R \propto g^4_{a\gamma\gamma}$.
The performance of a helioscope is therefore not determined only by the magnetic field and length, but also by the magnet aperture, the detector background and the efficiency of the optical system. A useful figure of merit can be written schematically as

$$F_{\text{helio}} \propto B^2 L^2 A \, \frac{\epsilon_o \epsilon_d}{\sqrt{ba}} \, \sqrt{\epsilon_t t}, \tag{3.81}$$

where A is the magnet cross-sectional area, ϵ_o is the throughput of the X-ray optics, ϵ_d is the detector efficiency, b is the normalized detector background, a is the focal spot area, ϵ_t is the fraction of time during which the magnet tracks the Sun, and t is the total data-taking time. This expression shows why next-generation helioscopes focus on large-aperture magnets, efficient X-ray optics, small focal spots and ultra-low-background detectors.
the now-decommissioned CAST experiment at CERN, Ref. [152] found in this regime:

$$g_{a\gamma\gamma} < 0.66 \times 10^{-10}\,\text{GeV}^{-1}, \qquad m_a \lesssim 0.02\,\text{eV}. \tag{3.82}$$

This remains one of the strongest laboratory bounds on the axion–photon coupling in the low-mass region.

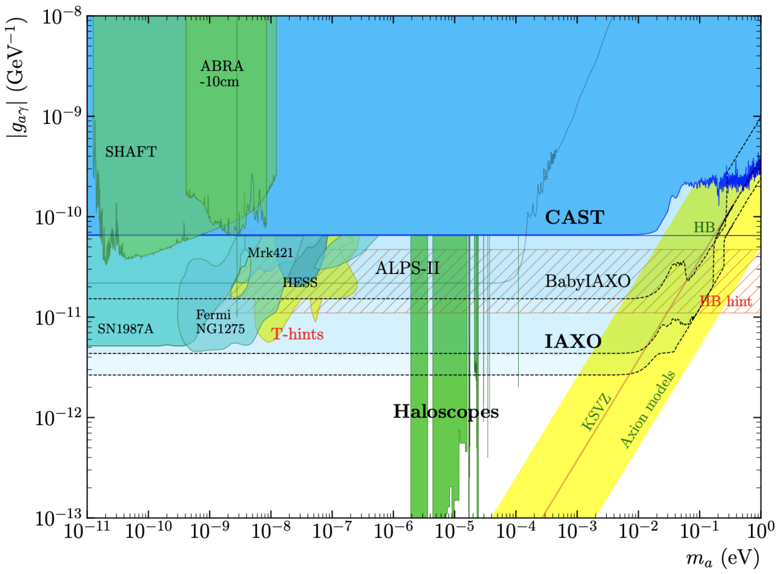

Figure 3.12: Excluded region from current and future experiments in the ALP parameter space, focusing on haloscopes.

To probe higher axion masses, CAST filled the magnet bore with a low-density buffer gas. In a medium, photons acquire an effective mass m_γ, and by tuning the gas pressure one can arrange $m_\gamma \simeq m_a$. This restores coherence for a narrow range of axion masses. Scanning over different gas pressures, therefore allows helioscopes to explore masses that would otherwise be inaccessible in vacuum. CAST also performed runs with helium buffer gas in order to extend the search to larger masses. Using ^{4}He and later ^{3}He, the experiment was able to scan the region up to masses of order $m_a \sim 1\,\text{eV}$. In particular, CAST excludes a part of the QCD axion parameter space in the eV mass range, although the precise limit depends on the axion model and on the value of the model-dependent coefficient $C_{a\gamma}$.
The successor of CAST is the International Axion Observatory, IAXO Ref. [153] a next-generation axion helioscope designed to improve the sensitivity to the axion–photon coupling by more than one order of magnitude. Its design is based on a large superconducting toroidal magnet with several magnet bores, each equipped with dedicated X-ray optics and low-background detectors. The use of a large magnet aperture increases the axion-induced photon flux collected by the experiment, while the optics reduce the effective detector area by focusing the signal onto small spots, thereby improving the signal-to-background ratio.
An intermediate experimental stage, BabyIAXO, is planned as a technically representative version of the full IAXO setup. It is intended to test the magnet, optics and detector technologies while already improving the sensitivity beyond CAST. Besides the standard Primakoff axion search, these experiments may also be sensitive to solar axions produced through axion–electron interactions, which would lead to different spectral features and therefore provide complementary information on the axion couplings. It was shown that for QCD axions they expect to probe $g_{a\gamma\gamma} \sim 10^{-11}\,\text{GeV}^{-1}$, or below, depending on the axion mass range. This big leap in sensitivity is achieved by the realization of a large-scale magnet, as well as by extensive use of X-ray focusing optics and low background detectors. The figure (3.12) summarizes the present experimental exclusions and the projected reach of future helioscopes in the $(m_a, g_{a\gamma\gamma})$ plane. In particular, CAST provides the strongest current helioscope bound at low masses, while BabyIAXO and IAXO are expected to explore a significant part of the parameter space motivated by QCD axion models and by astrophysical hints.

3.5 A General View

We summarize the bounds presented in the sections of this chapter in one plot. However let us remark that we did not review fully all the possible arguments and experiments used to constrain the axion parameter space.

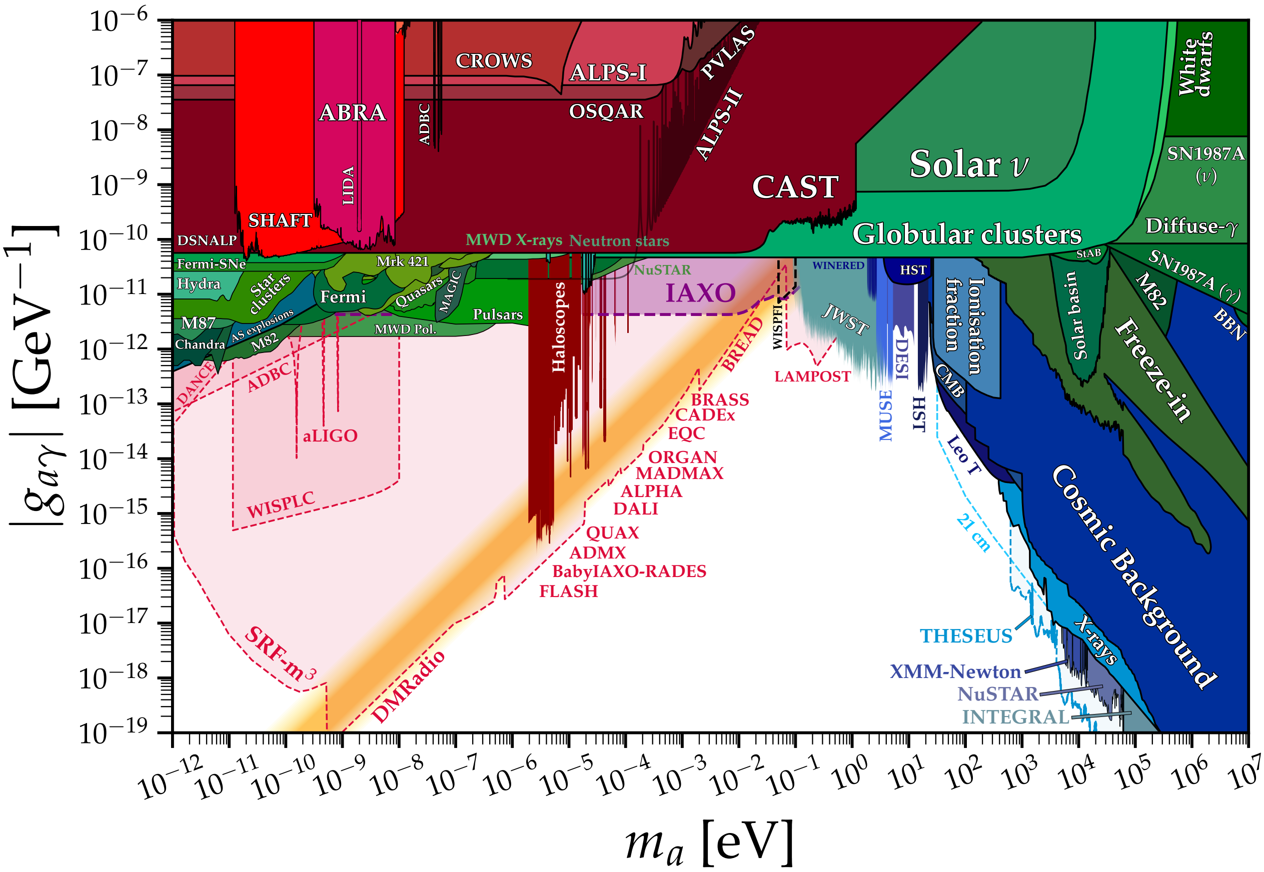

Figure 3.13: ALP parameter space, and its excluded regions using different arguments. We included also projections coming from future experiments, in dashed lines.

We conclude this chapter briefly adding two other bounds present in Fig. (3.13). First the blue region indicated with "cosmic background". The cosmic background (CB) is defined as the isotropic diffuse radiation field with extragalactic origin found across the electromagnetic spectrum. Different astrophysical sources dominate the CB emission at different energies. Assuming that dark matter consists of ALPs with masses on the order of electron volts or higher, an additional contribution to the CB due to their decay into two photons is produced in the early universe, their decay would produce a photon cosmic background that remains undetected, see Ref. [154]. In the shaded regions, there are different projections of future experiments, in particular indicated with "DMRadio" and "SRF-m^3. In fact at very low axion masses, below the usual microwave-cavity range, several new haloscope concepts are now being proposed. DMRadio uses a tunable lumped-element LC resonator placed in a strong static magnetic field. The axion dark matter field induces an effective electromagnetic current, which is read out with superconducting quantum sensors. This approach is designed to probe sub-μeV QCD axion dark matter [155].
SRF-m^3, exploits superconducting radio-frequency cavities. In this case the axion field can induce transitions between two nearly degenerate electromagnetic modes of a high-Q cavity. By tuning the frequency splitting between the modes, the experiment can scan over the axion mass, Ref. [156]. These techniques are complementary to conventional microwave-cavity haloscopes such as ADMX and extend the search toward lower axion masses.

Chapter 4

Axions in Extra-Dimensional Theories

Extra-dimensional theories provide a geometrical way of extending four-dimensional quantum field theory. The basic idea is that spacetime may contain one or more compact spatial dimensions which are not directly accessible at low energies. From the four-dimensional point of view, fields propagating in the higher-dimensional bulk appear as towers of massive states, the so-called Kaluza-Klein modes. Moreover, if some fields are localized on lower-dimensional hypersurfaces, or branes, the resulting low-energy theory can differ substantially from an ordinary four-dimensional theory. These ingredients make extra-dimensional models useful not only as ultraviolet-motivated frameworks, but also as phenomenological tools to address hierarchy problems, flavor hierarchies, chirality, and the structure of beyond-the-Standard-Model sectors.

4.1 Extra-Dimensional Theories in General

The idea of extra dimensions is surprisingly an old idea: For example in 1914 Nordström in Ref. [157] proposed a 5-dimensional theory to simultaneously describe electromagnetism and a scalar version of gravity, while Kaluza (1921) in Ref. [158] and Klein (1926) in Ref. [159] elaborated a 5-dimensional theory, with one spatial dimension compactified on a circle of radius R in an attempt to unify general relativity and electromagnetism. This theory nevertheless has some predictions that are not satisfied, and it is not really a viable model to describe nature.
We can briefly describe the basic Kaluza-Klein theory, the metric in this 5-dimensional theory is given by:

$$G_{MN} = \begin{pmatrix} \bar{g}_{\mu\nu} & \bar{A}_\mu \\ \bar{A}_\mu^T & \bar{g}_{55} \end{pmatrix}$$

with a field redefinition, and assuming that also in in this 5D theory we get the Einstein-Hilbert action one gets:

$$S = -\frac{1}{2}M_5^3 \int d^5x\sqrt{G}R_5[G] = -\frac{1}{2}M_P^2 \int d^4x\sqrt{g}\left[R_r[g] + \frac{1}{4}\phi F_{\mu\nu}F^{\mu\nu} - \frac{1}{6}\frac{\partial_\mu\phi\partial^\mu\phi}{\phi^2}\right] \tag{4.1}$$

Where note that as usual that assuming that the metric is dimensionless $[R_5] = [R_4] = 2$ where $R_{5(4)}$ is the Ricci scalar in 5 (4) dimensions. Therefore in the lhs of Eq. (4.1) M_5^3 is the "Planck" scale in 5 dimensions and M_p is the usual 4-dimensional Planck scale, that we get after integrating over the compactified dimension, and in particular, assuming a radius R for the extra dimension, $M_P^2 = M_5^3(2\pi R)$, we will turn back to this point in the next section.
We have different interesting features from Eq. (4.1): in fact beside having the typical action associated to a gauge theory, we also have a scalar field, called the *radion*, after canonical normalization of the kinetic term, this fields gets the typical couplings of a *dilaton* field. Already in this theory there are other interesting features, that is the Kaluza-Klein decomposition of the fields, that we are going to show below.
We already anticipate some crucial aspects where the old Kaluza-Klein theory fails: in fact, without

making additional assumption, that in the following we will work out, it is not possible to embed chiral fermions in this theory, and therefore we have no hope to describe the SM with a circular extra dimension. After Kaluza and Klein the idea of extra-dimensions was later on developed in the 1970's and 1980's, for theories such as supergravity and superstring theories, in this theories nevertheless extra dimensions were always at the scale of the Planck length. During the 1990's people began to consider the possibility of large extra dimensions, see Refs. [160, 161, 162, 163]. In phenomenology extra-dimensions became particularly popular after 1998, when Arkani-Hamed, Dimopoulos and Dvali (ADD) addressed them to solve the hierarchy problem in Refs. [164, 165]. We briefly present the main idea of this theory. It was always assumed that Nature has two different fundamental scales, that are $m_{EW} \sim 10^3$GeV and $M_P \sim 10^{18}$GeV, this two scales are quite different, and this disparity is, as last analysis, the hierarchy problem (that can of course can be formulated introducing the momentum cut-off for renormalization of the Higgs mass). In Ref. [165] is proposed that actually the energy scale associated at gravity is $\sim m_{EW}$ but, once n extra dimensions are included, its 4-dimensional counterpart is the observed Planck mass. This theory relies on the assumption that gravity, differently from all the other fields, can also propagate in the extra dimensions (this can vary in different models).In this case, new physics will only appear in the gravitational sector, and only when distances as short as the size of the extra dimension are actually reached. Nevertheless it is hard to test gravity at short distances, because of its weakness. We can say that Newton's law has been tested up to 0.05 mm, see Ref. [166], that is therefore the real bound on the size of an extra dimension.
It is easy to find the relation between the Newton constants $G_{N(4+n)}, G_{N(4)}$ (or analogously their corresponding "Planck Masses")

$$
\begin{aligned}
F_{(4+n)}(r) &= G_{N(4+n)}\frac{m_1 m_2}{r^{n+2}} \sim \frac{m_1 m_2}{M_{P(4+n)}^{n+2} r^{n+2}} \\
F_{(4)}(r) &= G_{N(4)}\frac{m_1 m_2}{r^2} \sim \frac{m_1 m_2}{M_P^2 r^2}
\end{aligned}
\tag{4.2}
$$

using Gauss Law and considering a flat space with mirror masses placed periodically in all the new dimensions one can show that:

$$
G_{N(4)} = \frac{S_{n+3}}{4\pi}\frac{G_{N(4+n)}}{V_n} \tag{4.3}
$$

Where $V_n = R^n$ is the volume of the extra dimensions and $S_n = \frac{2\pi^{n/2}}{\Gamma(n/2)}$ is the surface area of the unit sphere in n spatial dimensions. Therefore we get the crucial relation:

$$
M_P^2 \sim M_{P(4+n)}^{n+2}(2\pi R)^n \tag{4.4}
$$

Demanding R to be chosen to reproduce the observed Planck mass and to be equal for all the extra dimensions yields:

$$
R(n) = 2 \times 10^{\frac{31}{n}-16}\text{mm} \times \left(\frac{1\text{TeV}}{M_{(4+n)}}\right)^{1+\frac{2}{n}} \tag{4.5}
$$

Requiring that $M_{P(4+n)} = m_{EW} \sim 1$TeV we have for $n = 1, R \sim 10^{14}$mm implying deviations from Newtonian gravity over solar system distances, so this case is empirically excluded. For all $n \geq 2$, however, the modification of gravity only becomes noticeable at distances smaller than those currently probed by experiment (i.e. at distances $d < 0.05$ mm). For example at $n = 2$,with $M = 10$ TeV $R \sim 0.01$ mm. Different bounds exist for the radii of extra dimensional theories, that we will present in the following. Note that in a sense we moved the problem setting the Planck mass to the problem of setting the geometry (i.e. the number, the radius and the topology) of the extra dimensions. Nevertheless different models were proposed to build dynamical mechanism that bring the radii of extra dimensions to the "right" value, see Ref. [167].

4.1.1 Boundary Conditions

Let us fix first of all some notations: we will consider a five-dimensional spacetime with coordinates

$$x^M = (x^\mu, y), \qquad M, N = 0, 1, 2, 3, 5, \qquad \mu, \nu = 0, 1, 2, 3, \tag{4.6}$$

where x^μ denote the usual four-dimensional spacetime coordinates, while $y \equiv x^5$ parametrizes the extra spatial dimension. In the following we will assume that the fifth dimension is compactified on an interval, $y \in [0, L]$, whose endpoints may be interpreted as four-dimensional boundaries, or branes, located at $y = 0$ and $y = L$. Unless otherwise stated, capital Latin indices run over the full five-dimensional spacetime, while Greek indices run only over the four ordinary dimensions:

$$M, N, \ldots = 0, 1, 2, 3, 5, \qquad \mu, \nu, \ldots = 0, 1, 2, 3. \tag{4.7}$$

We now present the most important aspects of extra-dimensional theories. In particular as anticipated we need to set the topology of extra-dimensional theories. For $n = 1$ a natural choice would be S_1, nevertheless in this case (as in the Kaluza Klein theory) is not possible to build chiral fermions, that are fermions with different quantum numbers between their left right handed parts. For this reason it is necessary to consider compactifications with some sort of (mild) singularities. One example in two dimensions is the *chiral square*, where, differently from the torus adjacent instead of opposite sides are identified, with one extra dimension we have an even simpler possibility, that is the compactification interval, that we will consider in the following. Once we have a dimension that is compactified, something new arises: the theory is not fully defined until one specifies the boundary conditions, in fact in infinite space time this problem is often slipped under the implicit assumption that fields vanish at infinity (even if, sometimes this request can bring rich phenomenology, as we have seen in Eq. (1)). In the interval compactification, nevertheless different boundary conditions are physically allowed. We start from the simple example of a real scalar field:

$$S = \int d^5x \sqrt{g} \left(\frac{1}{2} g_{MN} \partial^M \phi \partial^N \phi - V(\phi) \right) \tag{4.8}$$

where g_{MN} is the background metric. Under a variation $\delta\phi$, the action receives two contributions $\delta S = \delta S_V + \delta S_A$:

$$\begin{aligned} \delta S_V &= - \int d^5x \sqrt{g} \left(\frac{1}{\sqrt{g}} \partial_M \left[\sqrt{g}\, g^{MN}\ \partial_N\ \phi \right] + \frac{\partial V}{\partial \phi} \right) \delta\phi \\ \delta S_A &= \int d^4x\ \sqrt{g}\ g^{5N}\ \partial_N \phi \cdot\ \delta\phi \Big|_{y=0}^{y=L} \end{aligned} \tag{4.9}$$

Where the last term is a surface term arising from integration by parts. As usual we assumed that $\phi \to 0$ as $x^\mu \to \pm\infty$. To derive the equations of motions of ϕ from $\delta S = 0$, we need to set to zero the surface term, that is:

$$-\partial_y \phi\ \delta\ \phi \Big|_{y=L} + \partial_y \phi\ \delta\phi \Big|_{y=0} \tag{4.10}$$

To satisfy Eq. (4.10), one possibility is to impose periodicity $\phi(L) = \phi(0)$ and $\phi'(L) = \phi'(0)$, however as we will soon see when describing fermions this defines a compactificiation on a circle, and therefore a vector-like theory. Other simple possibilities are to impose the Neumann or Dirichlet boundary conditions:

$$\text{Neumann: } \partial_y \phi \Big| = 0 \qquad\qquad \text{Dirichlet: } \phi \Big| = 0 \tag{4.11}$$

We can impose this boundaries in various combinations for example $(+, -) = (\text{N at } y = 0, \text{D at } y = L)$ and the other 3 combinations $(+, +)$; $(-, +)$; $(-, -)$.

4.1.2 Kaluza-Klein Decomposition

Having said that we need to specify the boundary conditions, we study how a field propagating in the bulk (that is, in all dimensions, and specifically in our 5D compactification) can be rewritten as a superposition of 4D fields. This is immediately seen with S_1 extra dimension: imposing periodicity, one can Fourier decompose the field only the extra dimensions, and after integrating over them one gets a 4D action. In general one needs an orthogonal and complete set of functions in the extra dimensions. We analyze the case of a scalar, a fermion and a gauge field.

Scalar Fields

Writing a scalar field in the following way:

$$\Phi(x^\mu, y) = \frac{e^{A(y)}}{\sqrt{L}} \sum_n \phi_n(x_\mu) f_n(y) \tag{4.12}$$

we expect that starting from the action of a free scalar field, we can write it as a :

$$S = \int d^5x \sqrt{g} \left(\frac{1}{2} g_{MN} \partial^M \phi \partial^N \phi - V(\phi) \right) \Rightarrow \sum_n \int d^4x \frac{1}{2} \{ \partial_\mu \phi_n \partial^\mu \phi_n - m_n^2 \phi_n^2 \} \tag{4.13}$$

Some considerations: Eq. (4.12) is the decomposition for a generic metric written in the form:

$$ds^2 = g_{MN} dx^M dx^N = e^{-2A(y)} \eta_{\mu\nu} dx^\mu dx^\nu - dy^2 \tag{4.14}$$

where eventually we can have $A(y) = 0$, and we recover a flat metric (the factor $e^{A(y)}$ in Eq. (4.12) in some way compensates the term $\sqrt{g}$). Furthermore note that in 5D Φ has dimension 3/2, while a generic ϕ_n has dimension 1. To find the values of m_n one needs to find the set $\{\phi_n\}$, in particular one writes the EOM for Φ in terms of Eq. (4.12), this defines an ODE for f_n with some (Neumann or Dirichlet) boundary conditions that has solutions only for some m_n. An important final remark: the KK decomposition is defined in terms of the free action, while interactions are to be treated perturbatively after inserting the previous KK decomposition. For instance:

$$\mathcal{L}_{int} = \int_0^L dy \sqrt{g} \frac{1}{3!} \lambda \Phi^3 = \sum_{n,m,l} \frac{1}{3!} \lambda_{mnl} \phi_m \phi_n \phi_l \tag{4.15}$$

where:

$$\lambda_{mnl} = \frac{\lambda}{L^{3/2}} \int_0^L dy\; e^{-A(y)} f_m f_n f_l \tag{4.16}$$

and f_j are the same functions found from the free scalar decomposition.

Fermions

We have really to define fermions in 5D. In 5d, As in the 4D case, fermions are representations of the Lorentz group. In this case the Clifford algebra nevertheless is different (since the 5D Lorentz group is different from its 4D counterpart) and it contains the γ_5 Dirac matrix as well. While in 4D the smallest representation of the Lorentz group is a two-component Weyl fermion, and one gets four component Dirac fermions only if one requires parity invariance, in 5D due to the fact the Clifford algebra contains γ_5 the smallest irreducible representation is the four component Dirac fermion.
Since a Dirac fermion contains two two-component fermions with opposite chirality, whenever one talks about higher dimensional fermions one has to start out with an intrinsically non-chiral set of 4D fermions. This is solved as we have said multiple times with an appropriate compactification of the extra dimensions. Said that we can proceed with the KK decomposition of fermions. We first present the action and the EOM for a fermion in 5D in a general background and then its KK decomposition. Starting from:

$$S_\Psi = \int d^5x \sqrt{g} \left(\frac{i}{2} \bar{\Psi} e^M_A \Gamma^A D_M \Psi - \frac{i}{2} (D_M \Psi)^\dagger \Gamma^0 e^M_A \Gamma^A \Psi - M(y) \bar{\Psi} \Psi \right) \tag{4.17}$$

Where $\Gamma_A = (\gamma_\mu, -i\gamma_5)$, e_A^M is the fünfbein and D_M is the covariant derivative, $D_M = \partial_M + \frac{1}{8}\omega_{MAB}[\Gamma^A, \Gamma^B]$ (ω are the spin connections of the metric g_{MN}) and we inserted a generic $y-$dependence of the mass from the extra-dimension coordinate. From Eq. (4.17) one gets the EOM for Ψ and the boundary conditions. For the metric given in Eq. (4.14) we get:

$$\left[ie^{A(y)}\partial\!\!\!/ + \left(\partial_5 - \frac{1}{2}A'\right)\gamma_5 - M(y)\right]\Psi = 0 \tag{4.18}$$

and the requirement:

$$\delta\bar{\Psi}\gamma_5\Psi\Big|_{y=L} + \delta\bar{\Psi}\gamma_5\Psi\Big|_{y=0} = 0 \tag{4.19}$$

Writing the boundary conditions in terms of the chiral parts Ψ_L and Ψ_R:

$$(\delta\bar{\Psi}_L\Psi_R + \delta\bar{\Psi}_R\Psi_L)\Big|_{y=0}^{y=L} = 0 \tag{4.20}$$

If we now impose periodic conditions for example on the left-handed part of Ψ, $\Psi_L(y=0) = \Psi_L(y=L)$, we see that we would need the same conditions also on the right-handed part, and therefore there would be nothing that distinguishes between the two chiralities. Instead we can impose:

$$\Psi_L\Big|_y = 0 \qquad \underline{\text{or}} \qquad \Psi_R\Big|_y = 0 \tag{4.21}$$

In this case one can show that imposing both boundary conditions brings $\Psi \equiv 0$ in the y coordinate. Writing also Eq. (4.18) in terms of the chiral parts we get:

$$\begin{aligned} &ie^{A(y)}\partial\!\!\!/\Psi_L + \left[\partial_5 - \frac{1}{2}A' - M(y)\right]\Psi_R = 0 \\ &ie^{A(y)}\partial\!\!\!/\Psi_R + \left[-\partial_5 + \frac{1}{2}A' - M(y)\right]\Psi_L = 0 \end{aligned} \tag{4.22}$$

and we have the two (times two) possibilities for each boundary:

$$\begin{aligned} (-) \equiv \Psi_L\Big|_y = 0 \Rightarrow \partial_5\Psi_R\Big|_y &= \left(\frac{A'}{2} + M(y)\right)\Psi_R\Big|_y, \\ (+) \equiv \Psi_R\Big|_y = 0 \Rightarrow \partial_5\Psi_L\Big|_y &= \left(\frac{A'}{2} - M(y)\right)\Psi_L\Big|_y, \end{aligned} \tag{4.23}$$

This is at the basis of the advertised result: compactification on an interval necessarily leads to boundary conditions that distinguish left from right handed fields, which will allow us to easily embed the SM structure. Similar to the scalar case, specifying the boundary conditions on both boundaries leads to the four possibilities presented above.

We are now ready for the KK decomposition of 5D fermions. We write:

$$\Psi_{L,R}(x^\mu, y) = \frac{e^{\frac{3A(y)}{2}}}{\sqrt{L}}\sum_n \psi_{L,R}^n\, f_{L,R}^n \tag{4.24}$$

as in the scalar case we want to write the action of the fermionic field, in this case Eq. (4.17) as a superposition of KK free wavefunctions, that is:

$$S_\Psi \Rightarrow \sum_n \int d^4x\ \bar{\psi}_n(i\partial\!\!\!/\psi_n - m_n)\psi_n \tag{4.25}$$

therefore writing the EOM of Ψ in terms of (4.24) leads to some ODE for the $\{f_{L,R}^n\}$, that wll lead to the spectrum of $\{m_n\}$:

$$\begin{aligned} \left(\partial_y + M - \frac{1}{2}A'\right) f_L^n &= m_n e^A f_R^n \\ \left(\partial_y - M - \frac{1}{2}A'\right) f_R^n &= -m_n e^A f_L^n \end{aligned} \tag{4.26}$$

and the two boundaries conditions are given as:

$$(-) \equiv f^n_L\Big|_y = 0 \Rightarrow \partial_y f^n_R\Big|_y = \left(\frac{A'}{2} + M(y)\right) f^n_R\Big|_y,$$
$$(+) \equiv f^n_R\Big|_y = 0 \Rightarrow \partial_5 f^n_L\Big|_y = \left(\frac{A'}{2} - M(y)\right) f^n_L\Big|_y, \tag{4.27}$$

Note that therefore we got a chiral theory for every zero mode. An important remark is surely that it is always allowed a massless mode. Indeed setting $m_{n=0} = 0$ the ODE's for $\{f^n_{L,R}\}$ can be integrated immediately to give:

$$f^0_{L,R}(y) = N_{L,R}\exp\left\{\left(\frac{1}{2}A(y) \mp \int_0^y dz M(z)\right)\right\} \tag{4.28}$$

For example in the $(+,+)$ prescription, we must set $N_R = 0$, and similarly for the other cases, thus the 0-mode level is indeed chiral. We note that presumably, the SM fermions could be identified with such 0-modes, in fact there can be additional contributions to the physical masses arising from EWSB. Thus, in realistic models, the "0-modes" may not be strictly massless. Furthermore the chirality of the zero modes is a feature present also in other types of compactification (i.e. those that do not represent orbifold, that for the other modes are *vector-like*) and it is an alternative way to get a chiral theory. It is worth spending some time on this other possibility, that can actually be embedded also in the "segment" compactification. This is given by the localization mechanism: if we have a "bulk" Yukawa coupling of the form $\Phi(y)\bar{\Psi}\Psi$, then if $\Phi(y) \sim 2\mu^2 y$, then the solution for the left-handed zero-mode is:

$$\Psi_L = \frac{\sqrt{\mu}}{(\pi/2)^{1/4}} e^{-\mu^2 y^2} \tag{4.29}$$

That is a way to localize a chiral fermion on a domain wall (in this case at $y = 0$). At this point let us consider the case when we have many fermions in the bulk, Ψ_i, and their action is schematically given by:

$$S = \int d^5x\sqrt{g}\left(\frac{i}{2}\bar{\Psi}_i e^M_A \Gamma^A D_M \Psi_i - \frac{i}{2}(D_M\Psi_i)^\dagger \Gamma^0 e^M_A \Gamma^A \Psi_i - \bar{\Psi}_i(\lambda\Phi(y) - m)_{ij}\Psi_j\right) \tag{4.30}$$

therefore there will be chiral fermion zero modes centered at the zeros of $(\lambda\Phi(y) - m)_{ij}$, and if we again set $\Phi(y) \sim 2\mu^2 y$ then the center of the fermion wave function will be at $y_i = m_i/2\mu^2$ where m_i are the eigenvalue of the m_{ij} matrix. Thus we have localized different fermions in different point of the extra dimensions! This can be used to easily solve two different problem: hierarchy between fermion masses and explain strongly suppressed higher dimensional operators (for example those that would give proton decay). Schematically consider:

$$S = \int d^5x\ \bar{L}\left[i\gamma^M\partial_M + \Phi(y)\right]L + \int d^5x\ \bar{E}^C\left[i\gamma^M\partial_M + \Phi(y) - m\right]E^c + \int d^5x\ \kappa H L^T C_5 E^C \tag{4.31}$$

L and E are for example the left-handed and right-handed lepton, and the last term is the bulk Yukawa coupling term written in a form that is both gauge and Lorentz invariant. Note that the terms $\bar{L}\Phi L$ and $\bar{E}^C\Phi E^C$ do not contribute to the SM fields, if we assume that them are identified with the KK zero-modes. Then the effective 4D Yukawa coupling for the zero modes will be of the form:

$$\int d^4x\ \kappa H\ l e^c \int dy\ \Psi_L(y)\Psi_{e^c}(y - r) \tag{4.32}$$

where l and e^c are the zero modes, Ψ is the wave function of these zero modes, that at the moment we are assuming to be a gaussian centered around two different points (0 and r) since we assume again the linear profile for $\Phi(y)$. After integration over y we get a coupling exponentially suppressed, $\kappa e^{-\mu^2 r^2/2}$. This is an hint of how fermions localized in different points of the extra dimension could naturally generate the fermion mass hierarchy observed in the SM. As last application we can turn as anticipated to proton decay. In this case we assume that leptons and quarks are localized in different points of the extra dimension, A dangerous proton decay operator generated by quantum gravity at the TeV scale would have the form (we still assume that the electroweak scale is the only natural scale, and thus associated also to gravity, following the ADD reasoning):

$$\int d^5x \, \frac{1}{M_{EW}^3}(Q^T C_5 L)(U^{cT} C_5 D^c) \tag{4.33}$$

and again calculating the wave function overlap of the zero modes as before we get that the effective 4D coupling is exponentially suppressed $\sim e^{-\mu^2 r^2}$.

Gauge Fields

We now pass to gauge fields and gauge theories in extra dimensions. We start from:

$$S_A = \int d^5x \sqrt{g} \left[-\frac{1}{4} F^{MN} F_{MN} \right] + S_{GF} \tag{4.34}$$

where the S_{GF} represent the gauge fixing term, chosen to avoid mixing between A_5 and A_μ. We decompose the gauge field in the form:

$$A_{\mu,5}(x^\mu, y) = \frac{1}{\sqrt{L}} \sum_n A^n_{\mu,5}(x^\mu) f^n_{\mu,5}(y) \tag{4.35}$$

From the EOM for $A_{\mu,5}$ one derives an ODE for f_n. Specifically, using the background defined as in Eq. (4.14) one gets:

$$\begin{aligned} &\partial_y \left(e^{-2A(y)} \partial_y f^n_A \right) + m_n^2 f^n_A = 0 \\ &\partial_y^2 \left(e^{-2A(y)} f^n_5 \right) + m_n^2 f^n_5 = 0 \end{aligned} \tag{4.36}$$

Studying the zero-mode solution, for example with $(+,+)$ boundary conditions, i.e. $\partial_y f^n_A \Big| = f^n_5 \Big| = 0$, one gets $f^0_\mu = 1$ and $f^0_5 = 0$. We note that in this case the zero-mode is not localized on a brane, in fact unlike the case of fermion 0-modes (which are controlled by the 5D Dirac mass), there are no free parameters to control the localization of the gauge zero- mode. The fact that the zero-mode of a gauge field is completely flat in the extra dimension ensures the universality of the gauge coupling, schematically:

$$g_4 \sim \frac{g_5}{L} \int dy \, |\psi(y)|^2 f^0_A(y) \Rightarrow g_4^2 = \frac{g_5^2}{L} \tag{4.37}$$

where in the implication we have set $f^0_A(y) = \text{const.}$ and used the normalization of fermions. For a $(-,-)$ boundary condition, that is $\partial_y e^{-2A} f^n_A \Big| = f^n_5 \Big| = 0$, there is no spin-1 zero-mode, but there is a spin-0 zero-mode with $f^0_5 \propto e^{2A}$, In certain constructions, such scalars can be identified with the SM Higgs field. Notice again that, given the background, the profile of such a 4D scalar 0-mode is fixed (no adjustable parameters to control its localization).

4.1.3 Effective Field Theory of a 3D Brane

In the ADD theory all particles of the SM are localized in four dimensions, while it is assumed that gravity can propagate also in the extra dimensions. We will refer to the surface along which some of the particles are localized as "branes", in particular a p-brane is a brane (a surface) with p spatial dimensions. We want to describe this branes with an effective field theory approach, that is we will not care about the high energy theory that will give rise to the brane, we will just assume their existence, therefore our thoery is consistent up to a certain cut-off scale, above which the dynamics that gave rise to the brane as to be taken into account. This is the approach given in Ref. [168].
First of all we fix some notations, that are slighlty different from the ones used above: the 3-brane will be assumed to be R^4, while the compactified extra dimensions are T^n (a n dimensional torus). The coordinates in the bulk are denoted with $X^M, M = 0, 1, \ldots, n+3$ the coordinates of the brane

are denoted by x_μ , $\mu = 0, 1, 2, 3$ while the coordinates along the extra dimensions are denoted by x_m , $m = 4, \ldots, 3+n$. as before the metric is given by $G_{MN}(X)$ and the position of the brane in the extra dimension is given by $Y^M(x)$, function only of the coordinates on the brane, note that these coordinates are n new bosonic degrees of freedom, we will show that these are dynamical fields, given our choice for the action of the 3D brane. The effective field theory that we are building is defined around a vacuum state:

$$G_{MN} = \eta_{MN}; \qquad\qquad Y^M = \delta^M_\mu x^\mu \tag{4.38}$$

That is, we expand about a Minkowski bulk spacetime, with the 3-brane occupying the subspace spanned by the four $X_\mu-$axes, and with the intrinsic 3-brane coordinates, x_μ, agreeing with the bulk coordinates, X_μ.

We now define the bulk action as anticipated above, using usual higher dimensional Einstein-Hilbert, plus a bulk cosmological constant:

$$S = -\int d^{4+n}X\sqrt{|G|}\ (M_*^{n+2}R^{(4+n)} + \Lambda) \tag{4.39}$$

Here M_* is the scale that we have denoted before $M_{(4+n)}$ or M_{EW}. to write the effective action, we need the induced metric on the brane, this is easily found to be:

$$ds^2 = G_{MN}dY^M\ dY^N = G_{MN}\frac{\partial Y^M}{\partial x^\mu}dx^\mu\ \frac{\partial Y^N}{\partial x^\nu}dx^\nu \tag{4.40}$$

Then the induced metric is just:

$$g_{\mu\nu}(x) = G_{MN}(x)\partial_\mu Y(x)^M \partial_\nu Y(x)^N = \eta_{\mu\nu} + \partial_\mu Y^m \partial_\nu Y_m \tag{4.41}$$

To write down a 4D effective action we need to write terms that are invariant under a transformation of the X coordinates, since it corresponds to the usual general covariance of a higher dimensional gravitational theory. Given that the x coordinates are completely nonphysical, since they only provide a way to parameterize the surface $Y(x)$, we need terms that are also invariant under a transformation of the brane coordinates. Practically, this means that one has to contract bulk indices with bulk indices and brane indices with brane indices. The general form of the brane action is:

$$S_{EFT} = \int d^4x\ \sqrt{|g|}\left[-f^4 - R^{(4)} + \frac{g^{\mu\nu}}{2}D_\mu\Phi D_\nu\Phi - V(\Phi) - \frac{g^{\mu\nu}g^{\alpha\beta}}{4}F_{\mu\alpha}F_{\nu\beta} + \ldots\right] \tag{4.42}$$

Where f is the energy density of the brane, or brane tension and we considered a general scalar and gauge field. As we have said $x \to x'(x)$ is an invariance of the Lagrangian, and therefore we can set a gauge fixing condition, which will eliminate the non-physical components of Y^M. Fortunately, it is straightforward to eliminate this redundancy, leaving us with only the physical number of 3-brane degrees of freedom, we set:

$$Y^\mu(x) - x^\mu = 0 \tag{4.43}$$

while the other $Y^m(x)$ are indeed physical. Then, in the path integral formulation, one should integrate only in the fields that satisfy our gauge fixing condition, furthermore the ghost determinant is actually trivial.

Finally, expanding the metric given our gauge choice we get:

$$g_{\mu\nu} = \eta_{\mu\nu} + \partial_\mu Y^m \partial_\nu Y_m + \cdots \Rightarrow \det g = -\partial_\mu Y^m \partial^\mu Y_m \tag{4.44}$$

and therefore that the leading term in the action becomes:

$$S = \int d^4x\ f^4 \partial_\mu Y^m \partial^\mu Y_m \tag{4.45}$$

We can conclude that the canonically normalized field is just $Z^m = fY^m$.

We now study in more detail the graviton modes. First of all we can expand the metric in the following way:

$$G_{MN} = \eta_{MN} + \frac{1}{2M_*^{\frac{n}{2}+1}}h_{MN} \tag{4.46}$$

In this way the graviton has dimension $\frac{n}{2}+1$ that is the right dimension for a bosonic field in $4+n$ dimensions, By definition its coupling with matter is given in terms of the stress-energy tensor of matter:

$$S_{int} = \int d^4x T^{\mu\nu} \frac{h_{\mu\nu}}{M_*^{\frac{n}{2}+1}} \tag{4.47}$$

Thus as expected the graviton couples linearly to $T^{\mu\nu}$ (this follows immediately by the definition of the stress energy tensor and by the fact that we impose the Einstein-Hilbert action in the bulk). Nevertheless we haven't written Eq. (4.47) in terms of KK modes. For extra dimensions compactified on a torus, $h_{\mu\nu}$ is periodic in the extra dimensions:

$$h_{MN}(x,y) = \sum_{k_1=-\infty}^{\infty} \cdots \sum_{k_n=-\infty}^{\infty} \frac{h_{MN}^{\vec{k}}(x)}{\sqrt{V_n}} e^{i\frac{\vec{k}\cdot\vec{y}}{R}} \tag{4.48}$$

Basically we Fourier expanded h_{MN} in the extra dimensions, here $V_n = (2\pi R)^n$ is a normalization factor, then we can rewrite Eq. (4.47):

$$\sum_{\vec{k}} \int d^4x T^{\mu\nu} \frac{1}{M_*^{\frac{n}{2}+1}} \frac{h_{\mu\nu}^{\vec{k}}}{\sqrt{V_n}} = \sum_{\vec{k}} \int d^4x \frac{T^{\mu\nu} h_{\mu\nu}^{\vec{k}}}{M_P} \tag{4.49}$$

Furthermore this KK modes can be decomposed in different parts: start with noting that h_{MN} is a $(n+4)\times(n+4)$ symmetric matrix, so with $(n+4)\cdot(n+5)/2$ components. Then we can use invariance under X coordinate transformation, therefore we can fix $n+4$ conditions with $h_{MN} \to \partial_M \xi_N + \partial_N \xi_M$, for example in the harmoinc gauge $\partial^M h_{MN} = \frac{1}{2}\partial_N h^M_M$, however this condition does not fully specify ξ, since adding ξ'_μ such that ${}'_\mu = 0$ still satify the harmonic gauge, therefore other $n+4$ conditions can be imposed. We conclude that the number of degrees of freedom for the graviton are $(n+4)\cdot(n+1)/2$. This means that in dimension higher than 4 we have degrees of freedom other than the two usual helicity states.

4D Graviton and its KK Modes. The four-dimensional graviton modes are contained in the upper-left 4×4 block of the $(4+n)\times(4+n)$ higher-dimensional graviton matrix:

$$\left(\begin{array}{c|c} G_{\mu\nu}^{\vec{k}} & \\ \hline & \end{array}\right). \tag{4.50}$$

These modes are labeled by the n-dimensional vector $\vec{k}$, which corresponds to the KK numbers along the extra dimensions. The zero mode correspond to the usual graviton, while the other KK modes become massive, "eating" a massless gauge field and a massless scalar.

4D Vectors and their KK Modes. The four-dimensional vector fields arise from the off-diagonal blocks of the higher-dimensional graviton matrix:

$$\left(\begin{array}{c|c} & V_{\mu j}^{\vec{k}} \\ \hline V_{i\nu}^{\vec{k}} & \end{array}\right). \tag{4.51}$$

Naively one could think there there would be n such four dimensional vectors, however we have seen above that the 4D graviton, at fixed $\vec{k}$ has to eat one 4D vector to form a full massive spin-2 field, thus there are only n 1 massive spin 1 particles in 4D. The missing of the last scalar is expressed by the constraint

$$\hat{k}^j V_{\mu j}^{\vec{k}} = 0, \tag{4.52}$$

so that there are only $n-1$ independent vectors.

4D Scalars and their KK Modes. The lower-right $n \times n$ block of the higher-dimensional graviton matrix corresponds to four-dimensional scalar fields:

$$\left(\begin{array}{c|c} & \\ \hline & S_{ij}^{\vec{k}} \end{array} \right). \tag{4.53}$$

Since S_{ij} is symmetric, again naively we would expect $n(n+1)/2$ scalar fields, however the graviton (at fixed $\vec{k}$ eats one scalar, while each vector eats another scalar, for a total of n "eaten" scalars. Therefore we have only $n(n-1)/2$. This can be seen from the constraint:

$$\hat{k}^i S_{ij}^{\vec{k}} = 0. \tag{4.54}$$

It is also convenient to separate the radion as the scalar field associated with the trace $h_j^{\vec{k}j}$. The remaining scalar modes are then chosen to be traceless.
The explicit expressions for the canonically normalized four-dimensional fields are

$$\text{radion:} \qquad H^{\vec{k}} = \frac{1}{\alpha} h_j^{\vec{k}j}, \tag{4.55}$$

$$\text{scalars:} \qquad S_{ij}^{\vec{k}} = h_{ij}^{\vec{k}} - \frac{\alpha}{n-1} \left(\eta_{ij} + \frac{\hat{k}_i \hat{k}_j}{\hat{k}^2} \right) H^{\vec{k}}, \tag{4.56}$$

$$\text{vectors:} \qquad V_{\mu j}^{\vec{k}} = \frac{i}{\sqrt{2}} h_{\mu j}^{\vec{k}}, \tag{4.57}$$

$$\text{gravitons:} \qquad G_{\mu\nu}^{\vec{k}} = h_{\mu\nu}^{\vec{k}} + \frac{\alpha}{3} \left(\eta_{\mu\nu} + \frac{\partial_\mu \partial_\nu}{\hat{k}^2} \right) H^{\vec{k}}. \tag{4.58}$$

Where $\alpha = \sqrt{\frac{3(n-1)}{n+2}}$. The equations of motion in the presence of brane-localized sources are

$$\left(\Box + \hat{k}^2 \right) \begin{pmatrix} G_{\mu\nu}^{\vec{k}} \\ V_{\mu j}^{\vec{k}} \\ S_{ij}^{\vec{k}} \\ H^{\vec{k}} \end{pmatrix} = \begin{pmatrix} \frac{1}{M_P} \left[-T_{\mu\nu} + \left(\eta_{\mu\nu} + \frac{\partial_\mu \partial_\nu}{\hat{k}^2} \right) \frac{T_\rho^\rho}{3} \right] \\ 0 \\ 0 \\ \frac{\alpha}{3 M_P} T_\mu^\mu \end{pmatrix}. \tag{4.59}$$

Therefore, only the four-dimensional graviton, the radion and their KK modes couple directly to brane-localized sources. The vector and traceless scalar modes do not couple directly to matter fields confined on the brane, and are therefore less relevant for processes involving only brane matter.

4.1.4 Phenomenology with Large Extra Dimensions

In this section we present some phenomenology of extra dimensional theories, and how it can be used to set bounds on this models. We follow Ref. [169]. Before presenting the phenomenological features of these models, a first important preliminary, which will turn useful again in the next chapter. In fact a phenomenological feature of large extra dimensional theories is that the mass splitting between KK modes is of the order of the inverse of the compactification radius:

$$\Delta m \sim \frac{1}{R} = M_* \left(\frac{M_*}{M_P} \right)^{\frac{2}{n}} = \left(\frac{M_*}{\text{TeV}} \right)^{\frac{n+2}{2}} 10^{\frac{12n-31}{n}} \text{eV} \tag{4.60}$$

where we have just rewritten Eq. (4.4) using a different notation ($M_* = M_{(4+n)} = M_{EW}$). This implies that for high energy processes in general there is an enormous number of KK modes available. To compute the number of KK modes available between $\vec{k}$ and $\vec{k} + d\vec{k}$ we can write:

$$dN = S_{n-1} k^{n-1} dk \tag{4.61}$$

where S_{n-1} is the usual surface of an n dimensional sphere with unit radius. The actual mass of a given KK mode is on general grounds $m \simeq \frac{|\vec{k}|}{R}$ and therefore:

$$dN = S_{n-1} m^{n-1} R^n dm = S_{n-1} m^{n-1} \frac{M_P^2}{M_*^{n+2}} dm \tag{4.62}$$

and from that the inclusive cross section of a process is just:

$$\frac{d\sigma}{ds} = \int_0^s S_{n-1} m^{n-1} \frac{M_P^2}{M_*^{n+2}} \frac{d\sigma_m}{ds} dm \tag{4.63}$$

For example, if we consider KK gravitons, we expect that $\frac{d\sigma_m}{ds} \simeq \frac{1}{M_p^2}$ and therefore $\frac{d^2\sigma}{dsdm} \sim S_{n-1} \frac{m^{n-1}}{M_*^{n+2}}$. Some of the features of large extra dimensions are:

- *Graviton production at colliders.*

  One of the most characteristic signatures of large extra dimensions is the direct production of KK gravitons in high-energy collisions, as presented in Refs. [170, 171, 172]. Although the coupling of each individual KK graviton to brane-localized matter is suppressed by the four-dimensional Planck scale, the inclusive rate can be sizeable because of the large multiplicity of accessible KK states.
  A typical process is the production of a single KK graviton in association with a Standard Model particle,

  $$e^+e^- \to \gamma G_{\mathrm{KK}}, \qquad pp \to j G_{\mathrm{KK}}, \tag{4.64}$$

  where the graviton escapes and is not detected. Therefore the experimental signature is a photon or a jet recoiling against missing energy. Since the decay width of an individual KK graviton is of order $\Gamma_G \sim \frac{m^3}{M_P^2}$, each mode is extremely long-lived on detector scales. The graviton thus behaves as an invisible weakly coupled particle. These signatures are particularly clean, since they lead to mono-photon or mono-jet events with missing transverse energy.

- *Virtual graviton exchange.*

  Besides direct production, KK gravitons can also contribute virtually to Standard Model scattering processes, see Refs. [170, 172, 173]. For example, the process

  $$e^+e^- \to f\bar{f} \tag{4.65}$$

  receives corrections from the exchange of the whole tower of massive gravitons. Schematically, the amplitude contains terms of the form

  $$\mathcal{A} \sim \frac{1}{M_P^2} \sum_{\vec{k}} \left[T^{\mu\nu} \frac{P_{\mu\nu\alpha\beta}}{s - m_{\vec{k}}^2} T^{\alpha\beta} \right], \tag{4.66}$$

  where $T_{\mu\nu}$ is the energy-momentum tensor of the brane fields and $P_{\mu\nu\alpha\beta}$ is the spin-2 projector. Since the number of KK modes grows rapidly with the energy, the sum over the tower can compensate the Planck-suppressed coupling.
  For $n \geq 2$, however, the sum over virtual KK modes is ultraviolet sensitive. This means that virtual graviton exchange cannot be predicted unambiguously within the low-energy effective theory alone, and depends on the ultraviolet completion of the model. It is nevertheless useful to parametrize these effects through higher-dimensional contact operators suppressed by the fundamental scale M_*.

- *Astrophysical constraints: supernova cooling.*

  Strong bounds on large extra dimensions arise from astrophysics. In particular, KK gravitons can be produced in the core of a supernova and escape into the bulk, carrying away energy. This mechanism is analogous to the usual axion cooling bounds that we have seen in (3), any

new weakly coupled light degree of freedom can provide an additional cooling channel for the supernova core.
The typical temperature in a core-collapse supernova is $T \sim 30\,\text{MeV}$. Since many KK modes with masses below this temperature can be emitted, the total graviton production rate is enhanced by the KK multiplicity. A useful way of estimating the bound is to compare the graviton emission rate with the standard axion cooling bound. Parametrically, the graviton production rate scales as

$$\Gamma_{\text{graviton}} \sim \frac{T^n}{M_*^{n+2}}. \tag{4.67}$$

Requiring that this energy loss does not exceed the observed cooling rate of supernovae gives very strong constraints, especially for $n = 2$ [174, 175]. Typically, one obtains a bound of order $M_* \gtrsim 10 - 100\,\text{TeV}, \qquad n = 2$, while for larger values of n the bound becomes weaker.

- *Cosmological bounds and cooling into the bulk.*

  Cosmology gives another important class of constraints. At sufficiently high temperature, the thermal bath on the brane can efficiently emit gravitons into the bulk. If this process is too efficient, the brane loses a large fraction of its energy density into bulk modes, modifying the standard cosmological history. The ordinary cooling due to the Hubble expansion is parametrically

$$\left.\frac{d\rho}{dt}\right|_{\text{expansion}} \sim -3H\rho \sim -3\frac{T^2}{M_P}\rho, \tag{4.68}$$

  while the energy loss into bulk gravitons scales as

$$\left.\frac{d\rho}{dt}\right|_{\text{evaporation}} \sim -\frac{T^n}{M_*^{n+2}}\rho. \tag{4.69}$$

  Equating the two rates defines a "normalcy temperature" T_*, below which the universe evolves approximately as an ordinary four-dimensional universe:

$$T_* \sim \left(\frac{M_*^{n+2}}{M_P}\right)^{\frac{1}{n+1}}. \tag{4.70}$$

  The reheating temperature after inflation must be below this scale; otherwise the universe would overproduce bulk gravitons. This bound is particularly restrictive for small n, since in that case the normalcy temperature can be close to the temperature required for successful Big Bang nucleosynthesis, see Refs. [165, 176].

- *Black-hole production.*

  A striking possible consequence of large extra dimensions is that the fundamental scale of quantum gravity may be close to the TeV scale. If this is the case, trans-Planckian particle collisions could produce microscopic black holes, as shown in Refs. [177, 178]. The idea is simple. In four dimensions, the Schwarzschild radius associated with a mass M is of order $r_H \sim G_N M$.
  In $4+n$ dimensions, instead, the gravitational potential scales differently, and the horizon radius is actually

$$r_H \sim \frac{1}{M_*}\left(\frac{M}{M_*}\right)^{\frac{1}{n+1}}. \tag{4.71}$$

  Therefore, if the impact parameter of a collision is smaller than r_H, a microscopic black hole may form. The corresponding geometrical cross section is roughly

$$\sigma \sim \pi r_H^2 \sim \frac{1}{M_*^2}\left(\frac{M_{\text{BH}}}{M_*}\right)^{\frac{2}{n+1}}. \tag{4.72}$$

  Such black holes would not be stable. They would decay through Hawking radiation into Standard Model particles and gravitons. The resulting events would be characterized by high multiplicity, large transverse energy, and a nearly thermal distribution of final states.

We summerize the current bound in the table below (credits to Ref. [106]):

n	1	2	3	4	5	6
Pendulum	2.2×10^5	/	/	/	/	/
Colliders	/	2.5×10^4	6.2×10^{-1}	3.1×10^{-3}	1.2×10^{-4}	1.5×10^{-5}
SN 1987A	2.5×10^{12}	4.9×10^3	5.8	1.9×10^{-1}	2.5×10^{-2}	6.1×10^{-3}
Neutron star	2.2×10^5	7.9×10^{-1}	1.3×10^{-2}	1.7×10^{-3}	5.1×10^{-4}	2.2×10^{-4}

Table 4.1: Upper limits on the compactification radius $R_{\max}$, expressed in keV^{-1}, as a function of the number of extra dimensions n. The bounds come from torsion-pendulum tests of Newton's law, see Ref. [179], collider searches by CMS, Ref. [180], and astrophysical constraints from SN 1987A and neutron-star excess heat, Ref. [181].

4.2 Kaluza Klein Axions

We now present axions in the framework of extra dimensions, i.e. Kaluza-Klein axions, following mainly the seminal papers presented in Ref. [182, 104, 183]. We have seen already that the QCD axion solution to the Peccei-Quinn solution introduces a new scale, f_{PQ} that is the energy scale at which the PQ symmetry is spontaneously broken, and that we expect that the mass of the axion is of order Λ^2_{QCD}/f_{PQ}, and has we have already seen also the axion-photon or the axion-fermion coupling are $\propto 1/f_{PQ}$, so that in general more the axion is light, more its coupling are weak. We have already shown that in its original proposal this scale was at the order of the electroweak scale, but that this is already ruled out. The crucial ingredient here is that both the axion's mass and its coupling are determined by essentially a single energy scale f_{PQ}.
In extra dimensional theories, however this is no longer true. Firs of all in our previous discussion we stated that only the electroweak scale is the fundamental scale of Nature, and how the Planck mass is actually just arises in the 4D point of view. However this statement as it is presented has the problem to not explain the intermediate scales f_{PQ}. This issue will be solved basically placing the axion in a number of extra dimensions $\delta < n$. We now present the model in Ref. [182]: Assume there is a complex scalar field that live in the bulk and it is charged under $U_{PQ}(1)$:

$$\phi \to e^{i\alpha}\phi \tag{4.73}$$

and assume it is spontaneously broken by the bulk dynamics, $\langle\phi\rangle = f_{PQ}/\sqrt{2}$, and as usual, neglecting the radial term we write:

$$\phi = \frac{f_{PQ}}{\sqrt{2}} e^{ia/f_{PQ}} \tag{4.74}$$

and under a PQ transformation:

$$a \to a + \alpha f_{PQ} \tag{4.75}$$

Where we assume that α is a constant also in the extra dimensions (i.e. $U_{PQ}(1)$ is a global symmetry also in the extra dimensions). The axion Lagrangian is composed at least by two terms, the first is just the kinetic term, that comes from the kinetic term of ϕ, as done in Ref. [182] and in the previous section, we stick to 5 dimensions:

$$S_K = \int d^4x\ dy M_S\ \partial^M\phi^*\partial_M\phi = \int d^4x\ dy\ M_S\frac{1}{2}\partial_M a\partial^M a \tag{4.76}$$

Then from the chiral anomaly restricted to the brane where our fermionic fields live we have:

$$S = \int d^4x dy\ \frac{\xi}{f_{PQ}}\frac{g^2}{32\pi^2} aF_{\mu\nu}\tilde{F}^{\mu\nu}\ \delta(y) \tag{4.77}$$

note that in odd dimension we can not have local anomalies in the bulk, while we can localize anomalies on the boundaries. For more information on the chiral anomaly in 5 dimensions, using our compactification see Ref. [184], the result is that the bulk "is not chiral" that is it cannot have anomalies, that live on the fixed points (i.e. on the boundary) of our theory. We focus on $y = 0$. This result is somewhat not so solid if one considers other compactifications, see e.g. Ref. [185]. Other useful instructive references are given in Refs. [186, 187]. Our full action is then:

$$S = \int d^4x \; dy \; \left(M_S \frac{1}{2} \partial_M a \partial^M a \; + \frac{\xi}{f_{PQ}} \frac{g^2}{32\pi^2} a F_{\mu\nu} \tilde{F}^{\mu\nu} \; \delta(y) \right) \tag{4.78}$$

The axion can be written as a KK tower:

$$a(x^\mu, y) = \sum_{n=0}^{\infty} a_n(x^\mu) \cos\left(\frac{ny}{R}\right) \tag{4.79}$$

and under a Peccei-Quinn transformation the KK axion modes transforms as:

$$\begin{aligned} a_0 &\to a_0 + \alpha f_{PQ} \\ a_k &\to a_k \end{aligned} \tag{4.80}$$

Inserting the Kaluza-Klein decomposition for the axion, Eq. (4.79) in Eq. (4.78) we arrive at the effective Lagrangian:

$$\mathcal{L}_{eff} = \mathcal{L}_{QCD} + \frac{1}{2} \sum_{n=0}^{\infty} (\partial_\mu a_n)(\partial^\mu a_n) - \frac{1}{2} \sum_{n=1}^{\infty} \frac{n^2}{R^2} a_n^2 + \frac{\xi}{\hat{f}_{PQ}} \frac{g^2}{32\pi^2} \left(\sum_{n=0}^{\infty} r_n a_n \right) F_{\mu\nu} \tilde{F}^{\mu\nu} \tag{4.81}$$

where $r_n = 1$ is $n = 0$ and otherwise $r_n = 1/\sqrt{2}$. Basically this is due to the integral:

$$\int_0^{\pi R} \cos\left(\frac{nx}{R}\right) \cos\left(\frac{mx}{R}\right) \, dx = \begin{cases} 0, & m \neq n, \\ \pi R, & m = n = 0, \\ \dfrac{\pi R}{2}, & m = n \neq 0. \end{cases}$$

After canonical nomralization of the KK modes it is easy to see that $\hat{f}_{PQ} = f_{PQ}(V M_s)^{1/2}$, where V is the volume of our compactified space. For δ extra dimensions this definition generalizes to $\hat{f}_{PQ} \sim f_{PQ}(RM_s)^{\frac{\delta}{2}}$, and we note that the energy scale that sets the coupling between the KK axion modes and the gluons (or photons) is the the volume-renormalized quantity $\hat{f}_{PQ}$. Since we $M_s \gg R^{-1}$ (we have assumed that $M_s \sim 1$ TeV) it follows that $\hat{f}_{PQ} \gg f_{PQ}$. This is very similar to what was proposed for gravity in Ref. [164], in that case we had a very similar relation, $M_P \sim (RM_*)^{n/2} M_*$, therefore since we expect $M_s \sim M_*$ and we know that R has to be adjusted to get the correct Planck mass, then if $\delta = n$ we would have $\hat{f}_{PQ} \sim M_P$, which would overclose the universe. However, as also proposed in Ref. [188], we can set an intermediate set an intermediate $\hat{f}_{PQ}$ if $\delta < n$. In other words, under these assumptions, the axion must be restricted to a subspace of the full higher-dimensional bulk A second important remark is that even if $\hat{f}_{PQ}$ sets the scale for couplings between gauge fields and individual axion modes, the gauge fields couple not to an individual axion mode, but to a linear superposition of them:

$$a' \equiv \frac{1}{\sqrt{N}} \sum_{n=0}^{n_{max}} r_n a_n = \frac{1}{\sqrt{N}} \left(a_0 + \sqrt{2} \sum_{n=1}^{n_{max}} a_n \right) \tag{4.82}$$

where $N = 1 + 2n_{max}$ is a normalization factor, and n_{max} is a cutoff, determined by the relation $n_{max} \sim M_s R$, that is when the KK axion mass is of the order of M_s the scale at which our 5-dimensional description is no longer valid. But we have seen that $RM_s \gg 1$, then, the coupling between KK axion modes and gluons is:

$$\frac{\sqrt{N}}{\hat{f}_{PQ}} a' F_{\mu\nu} \tilde{F}^{\mu\nu} \simeq \frac{1}{f_{PQ}} a' F_{\mu\nu} \tilde{F}^{\mu\nu} \tag{4.83}$$

In other words, the effects of the volume factor in are canceled by the normalization of the Kaluza-Klein linear superposition a. This result furthermore remains true of any number of extra dimensions. This is a serious issue to invisible axion models in extra dimensions, that however, as we will show, can be solved. let us mention that very similar considerations, that we will not repeat, hold for axion-fermions couplings. We now study if axions still solve the strong CP problem even in this extra-dimensional framework. It is enough to consider just a', that is the superposition that KK axion modes that couples to gluons and photons to correctly take into account how the QCD potential depends on $\bar{\theta}$. Then basically as we have seen already in chapter (1), we have that a new term arises $V(\bar{\theta}) \sim -\Lambda^4 \cos\bar{\theta} + a'$, so that the potential associated to the KK axion modes is given by:

$$V(a_n) = \frac{1}{2}\sum_{n=1}^{\infty}\frac{n^2}{R^2}a_n^2 + \frac{g^2}{32\pi^2}\Lambda^4\left[1-\cos\left(\frac{\xi}{\hat{f}_{PQ}}\sum_{n=0}^{\infty}r_n a_n + \bar{\theta}\right)\right] \tag{4.84}$$

This potential is minimized when we look for stationary solutions:

$$\frac{\partial V}{\partial a_n} = \frac{n^2}{R^2}a_n(1-\delta_{n0}) + r_n\frac{\xi}{\hat{f}_{PQ}}\frac{g^2}{32\pi^2}\Lambda^4\sin\left(\frac{\xi}{\hat{f}_{PQ}}\sum_{n=0}^{\infty}r_n a_n + \bar{\theta}\right) = 0 \tag{4.85}$$

This has the solutions:

$$\begin{aligned} a_0 &= \frac{\hat{f}_{PQ}}{\xi}(-\bar{\theta} + m\pi) \\ a_k &= 0 \end{aligned} \tag{4.86}$$

where $m \in 2\mathbb{Z}$. Thus this higher dimensional Peccei-Quinn mechanism continues to solve the strong CP problem, where a_0 is the usual Peccei-Quinn axion. We now study the axion mass matrix:

$$\mathcal{M}_{nn'} = \frac{\partial^2 V}{\partial a_n \partial a_{n'}} = \frac{n^2}{R^2}\delta_{nn'} + \xi^2\frac{g^2}{32\pi^2}\frac{\Lambda^4}{\hat{f}_{PQ}}r_n r_{n'} \tag{4.87}$$

Calling $m_{PQ}^2 = \xi^2\frac{g^2}{32\pi^2}\frac{\Lambda^4}{\hat{f}_{PQ}}$ and $y = 1/(m_{PQ}R)$:

$$\mathcal{M} = m_{PQ}^2(\delta_{nn'}y^2n^2 + r_n r_{n'}) \tag{4.88}$$

It is not so easy to find the eigenvalues of $\mathcal{M}$, however after applying some mathematical theorems, the eigenvalues λ are the solution of the transcendental equation:

$$\frac{\pi\tilde{\lambda}}{y}\cot\left(\frac{\pi\tilde{\lambda}}{y}\right) = \tilde{\lambda} \tag{4.89}$$

that we can also rewrite as:

$$\pi R\lambda\cot(\pi R\lambda) = \frac{\lambda^2}{m_{PQ}^2} \tag{4.90}$$

and $\lambda = \tilde{\lambda}m_{PQ}$. For each eigenvalue λ^2, the corresponding normalized mass eigenstate $\hat{a}_\lambda$ in terms of the KK modes is:

$$\hat{a}_\lambda = \sum_{k=0}^{\infty}U_{\lambda k}a_k, \qquad U_{\lambda k} = \left(\frac{r_k\tilde{\lambda}^2}{\tilde{\lambda}^2 - k^2y^2}\right)\frac{\sqrt{2}}{\tilde{\lambda}}(\tilde{\lambda} + 1 + \pi^2/y^2)^{-1/2} \tag{4.91}$$

Here we present the zero mode solution as a function of y^{-1} measured in units of $1/R$. The horizontal line represent the upper limit on the axion mass, that is $1/2R$. Studying Eq. (4.90) it is clear in fact that the axion mass is always below $1/2R$. The other line represent the case where the axion mass is equal to m_{PQ}. This limit is recovered in the case when $m_{PQ}R \to 0$.

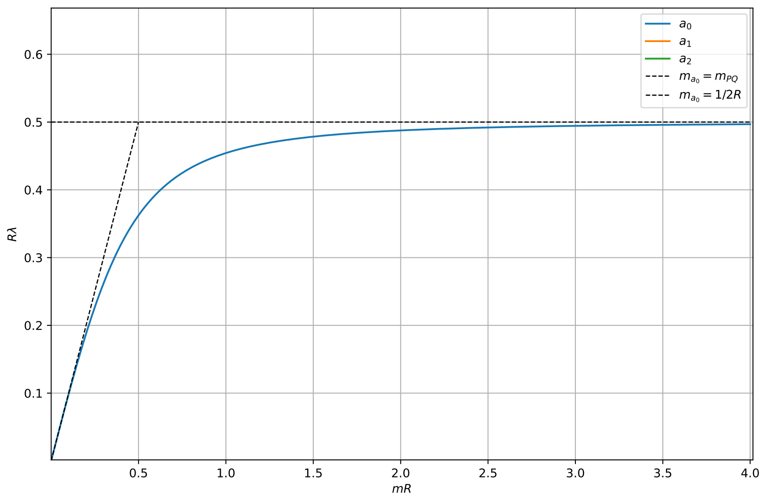

Figure 4.1: Mass of the zeroth axion mode as a function of mR. We reproduced the figure in Ref. [182].

As a result of this unexpected higher-dimensional behavior for the axion mass, we see that when $m_{PQ} > 1/2R$, the Peccei-Quinn scale essentially decouples from the axion mass! once the axion mass constraint is satisfied by an appropriate choice of R, the parameter m_{PQ}, or equivalently the Peccei–Quinn scale f_{PQ}, can still be varied without spoiling the bound, provided that

$$m_{PQ} \gtrsim \frac{1}{2R}.$$

This freedom allows one to suppress the axion couplings to matter, thereby offering a possible mechanism for realizing an invisible axion.

Chapter 5

Axion Production in the Sun and Its Number Density on Earth

In this chapter, we compute the axion production rate inside the Sun, assuming a photon-axion coupling. The relevant processes are three: photon coalescence, the Primakoff effect, and plasmon decay. We then compute their contribution to the axion number density around the Sun, which allows us to determine the axion decay rate on Earth. Finally, we apply our computations to KK axions.

5.1 Photon Coalescence

We start by considering photon coalescence, $\gamma\gamma \to a$. Inside the solar plasma, photons acquire an effective mass, which we denote by ω_P. We have

$$\omega_P^2 = \frac{4\pi\alpha n_e}{m_e}. \tag{5.1}$$

Here, n_e and m_e are the electron number density and the electron mass, respectively. The electron number density is determined by assuming that the Sun is composed only of hydrogen (H1) and helium (He4), and by imposing electric charge neutrality through the expression

$$n_e = \rho\frac{1 + X_{H1}}{2m_p}, \tag{5.2}$$

where ρ is the density inside the Sun, X_{H1} is the hydrogen mass fraction, and m_p is the proton mass. Knowing these quantities as functions of the solar radius, we can plot the effective photon mass as a function of the position inside the Sun. We use the Standard Solar Model (SSM) of Ref. [189].

The effective photon mass is a central quantity in our computations, since, for example, plasmon decay would otherwise be kinematically forbidden. Below, we also show the solar mass fraction constituted only by hydrogen and helium, confirming that our approximation in Eq. (5.2) is well justified.
Let us mention, for future reference, that photons inside a plasma are actually divided into two types: longitudinal and transverse plasmons, each with its own dispersion relation. We will return to this point later, when discussing plasmon decay. For now, we recall that we work only with transverse plasmons, whose dispersion relation is given by

$$\omega^2 = k^2 + \omega_P^2\left(1 + \frac{k^2}{\omega^2}\frac{T}{m_e}\right) \simeq k^2 + \omega_P^2, \tag{5.3}$$

where T is the solar temperature. To show that $T/m_e \ll 1$, and hence that the correction to the usual dispersion relation is negligible, we plot the solar temperature as a function of the solar radius.

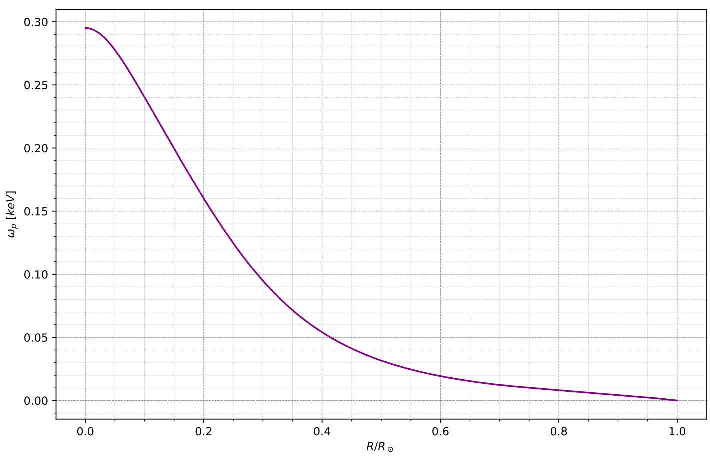

Figure 5.1: Effective photon mass as a function of the radius inside the Sun, based on data from the Saclay model presented in Ref. [189]. The photon mass is of the order of 0.1 keV, which is not completely negligible compared with the typical solar temperature.

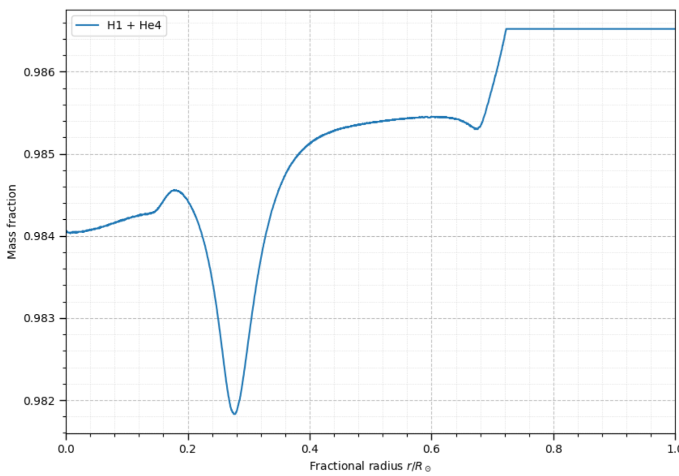

Figure 5.2: Mass fraction of Hydrogen plus Helium. Eq. (5.2) is valid under the assumption that the Sun is composed only of hydrogen and helium; this approximation is valid within 2%.

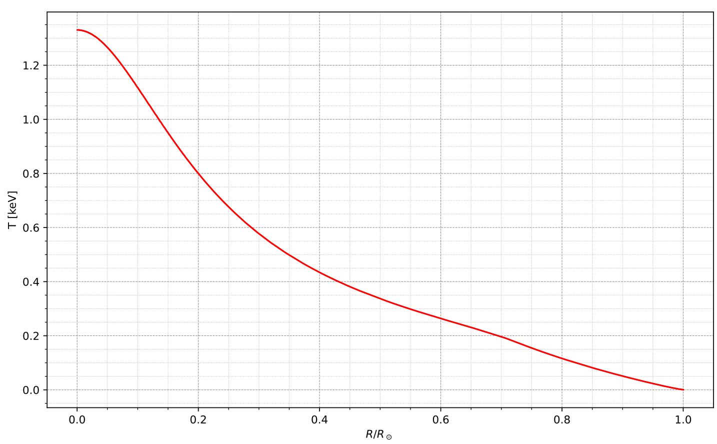

Figure 5.3: Temperature inside the Sun as a function of the radius. A typical temperature is about 0.5 keV, comparable to the effective photon mass.

We are now ready to derive the production rate due to photon coalescence inside the Sun. We start from the usual interaction Lagrangian

$$\mathcal{L} = -\frac{g_{a\gamma\gamma}}{4} a F_{\mu\nu}\tilde{F}^{\mu\nu}, \tag{5.4}$$

from which the Feynman rule is easily obtained:

$$-ig_{a\gamma\gamma}\epsilon^{\mu\nu\sigma\rho}p_{1_\sigma}p_{2_\rho}. \tag{5.5}$$

We study the axion distribution function, $f_a(E)$, through the Boltzmann equation. As usual, $f_a(E)$ is defined through the relation

$$n_a = \int pEf_a(E)\frac{4\pi}{(2\pi)^3}dE. \tag{5.6}$$

From the Boltzmann equation, we have

$$\frac{df_a}{dt}(E) = \frac{1}{2E}\int \frac{d^3k_1}{(2\pi)^3 2\omega_1}\frac{d^3k_2}{(2\pi)^3 2\omega_2}(2\pi)^4\delta^4(k_1+k_2-p)|\mathcal{M}_\gamma|^2 f_1^{eq}(\omega_1)f_2^{eq}(\omega_2)(1+f_a(E)). \tag{5.7}$$

Here, p and E are the momentum and energy of the axion, respectively, while k_1, k_2, ω_1, and ω_2 denote the three-momenta and energies of the two photons.
We can compute the matrix element for inverse axion decay by taking into account the effective photon mass. The averaged matrix element is given by

$$\sum_{pol}|\bar{\mathcal{M}_\gamma}| = \frac{g_{a\gamma}^2}{4}\epsilon^{\mu\nu\sigma\tau}\epsilon_{\mu\nu\alpha\beta}k_{1_\sigma}k_{2_\tau}k_{1_\alpha}k_{1_\beta} = -\frac{g_{a\gamma}^2}{2}\left((k_1k_2)^2 - k_1^2k_2^2)\right) = \frac{g_{a\gamma}^2}{8}m_a^2\left(m_a^2 - 4\omega_P^2\right). \tag{5.8}$$

Using the identity

$$f_1^{eq}(\omega_1)f_2^{eq}(\omega_2)(1+f_a^{eq}(E)) = f_a^{eq}(E)(1+f_1^{eq}(\omega_1))(1+f_2^{eq}(\omega_2)), \tag{5.9}$$

and working in the dilute limit, we can rewrite Eq. (5.7) as

$$\frac{df_a}{dt}(E) \simeq f_a^{eq}(E)\frac{1}{2E}\int \frac{d^3k_1}{(2\pi)^3 2\omega_1}\frac{d^3k_2}{(2\pi)^3 2\omega_2}(2\pi)^4\delta^4(k_1+k_2-p)|\mathcal{M}_\gamma|^2(1+f_1^{eq}(\omega_1)+f_2^{eq}(\omega_2)). \tag{5.10}$$

At this point, it is useful to define

$$\Gamma_{a\gamma\gamma}(T) = \frac{1}{2E}\int \frac{d^3k_1}{(2\pi)^3 2\omega_1}\frac{d^3k_2}{(2\pi)^3 2\omega_2}(2\pi)^4\delta^4(k_1+k_2-p)|\mathcal{M}_\gamma|^2(1+f_1^{eq}(\omega_1)+f_2^{eq}(\omega_2)). \tag{5.11}$$

Since the matrix element in Eq. (5.8) does not depend on the photon momenta, we factor it out and focus on the phase-space integral. The integral over d^3k_2 is trivial because of the Dirac delta, from which we obtain the relations $\mathbf{k}_2 = \mathbf{p} - \mathbf{k}_1$ and $E_2 = \sqrt{E_1^2 + p^2 - 2|k||p|\cos\theta}$.

We therefore have

$$\Gamma_{a\gamma\gamma}(T) \sim \frac{1}{2E}\int \frac{d^3k_1}{(2\pi)^2 2\omega_1 2\omega_2}\delta(\omega_1+\omega_2-E)(1+f_1^{eq}(\omega_1)+f_2^{eq}(\omega_2)). \tag{5.12}$$

The integral over the azimuthal variable ϕ is trivial, giving a factor of 2π, and we can rewrite the expression as

$$\begin{aligned}\Gamma_{a\gamma\gamma}(T) \sim \frac{1}{4\pi E}&\int_0^{+\infty}\frac{k_1^2}{4\,\omega_1\,\omega_2}\int_{-1}^{1}\delta\Big(\sqrt{\omega_1^2+p^2-2|k_1|\,|p|\,\cos\theta}+\omega_1-E\Big)\\ &\times\left(1+f_1^{\mathrm{eq}}(\omega_1)+f_2^{\mathrm{eq}}(\omega_2)\right)d\cos\theta.\end{aligned} \tag{5.13}$$

The Dirac delta ensures that, for each kinematically allowed value of ω_1, the corresponding angle θ is selected in such a way that four-momentum conservation is satisfied. Applying the properties of the Dirac delta, we obtain

$$\begin{aligned}\Gamma_{a\gamma\gamma}(T) \sim \frac{1}{16\pi E}\int_{E_{\min}}^{E_{\max}} & k_1\,\omega_1\,d\omega_1\,\frac{\omega_2}{k_1\,p\,\omega_1\,\omega_2} \\ & \times \left(1+f_1^{\rm eq}(\omega_1)+f_2^{\rm eq}(\omega_2)\right).\end{aligned} \tag{5.14}$$

For each angle, one can find a corresponding value of the energy ω_1, but not all energies are allowed. The maximum and minimum values of the energy correspond to the back-to-back decay configurations. If the axion energy in the solar frame is $E = m_a\gamma$ and therefore $p = m_a\gamma\beta$, where γ is the Lorentz factor, then the maximum and minimum photon energies are, respectively,

$$\frac{m_a\gamma}{2} \pm \frac{\gamma\beta}{2}\sqrt{m_a^2-4\omega_p^2},$$

with the condition $m_a > 2\omega_p$. At first order in ω_p, we approximate the maximum and minimum energies as $(E \pm p)/2$. We are now ready to perform the integral. Since the integration domain is symmetric around $E/2$, and since $\omega_2 = E - \omega_1$, we can write

$$\Gamma_{a\gamma\gamma}(T) \sim \frac{1}{16\pi E p}\int_{\frac{E-p}{2}}^{\frac{E+p}{2}} \left(1+\frac{2}{e^{\omega/T}-1}\right) d\omega. \tag{5.15}$$

To perform the integral, it is sufficient to change variables to $x = \omega/T$. We finally obtain

$$\Gamma_{a\gamma\gamma}(T) \sim \frac{1}{8\pi E}\left(1+\frac{2T}{p}\ln\left(\frac{1-e^{-\frac{E+p}{2T}}}{1-e^{-\frac{E-p}{2T}}}\right)\right). \tag{5.16}$$

Restoring the matrix element, we get

$$\Gamma_{a\gamma\gamma}(T) = \frac{g_{a\gamma\gamma}^2 m_a^3}{64\pi}\frac{(m_a^2-4\omega_P^2)}{m_a^2}\frac{m_a}{E}\left(1+\frac{2T}{p}\ln\left(\frac{1-e^{-\frac{E+p}{2T}}}{1-e^{-\frac{E-p}{2T}}}\right)\right). \tag{5.17}$$

From the axion distribution function, we can obtain the number density per unit energy and time by using the definition of $\Gamma_{a\gamma\gamma}(T)$:

$$\frac{d^2n_a(E)}{dEdt} = \frac{pEf_a^{eq}(E)}{2\pi^2}\Gamma_{a\gamma\gamma}(T). \tag{5.18}$$

We now integrate over the SSM presented in Ref. [189] to obtain the differential solar flux, namely the number of axions produced per unit time and per unit energy:

$$\frac{d\Phi_{coal}}{dE} = \frac{1}{4\pi d^2}\int d^3r\frac{dn_a(E)}{dE} = \frac{R_\odot^3}{d^2}\int \frac{dn_a(E)}{dE}\bar{r}^2 d\bar{r}. \tag{5.19}$$

To express the differential flux in units of cm^{-2} s^{-1} keV^{-1}, we multiply the integral by a factor 2.8×10^{11}. In the following, we present the results for $g_{a\gamma\gamma} = 10^{-12}$ GeV^{-1}. The differential flux for the coalescence process is shown below.

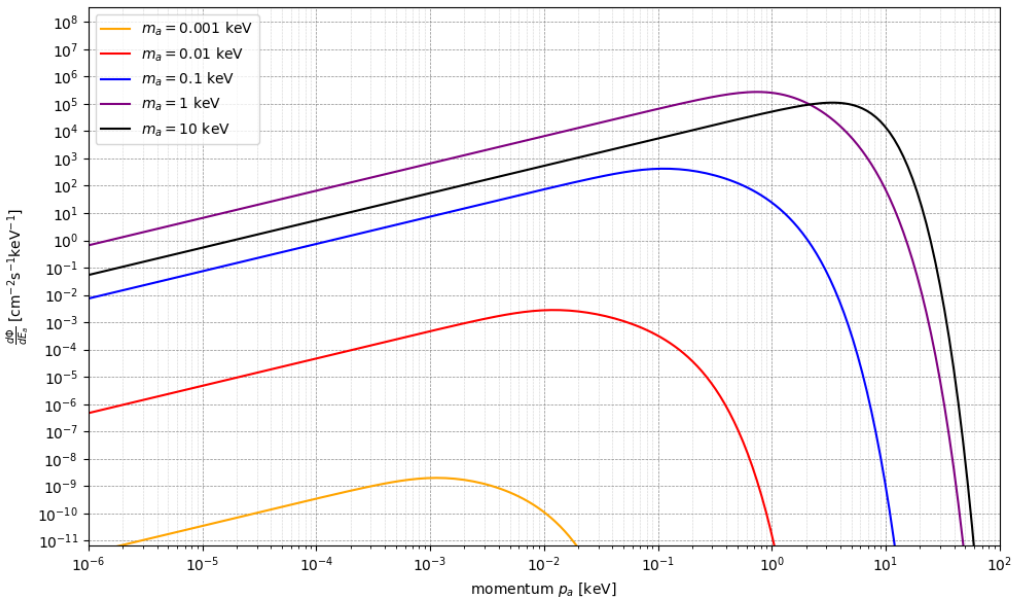

Figure 5.4: Differential flux for photon coalescence, for different values of the axion mass, as a function of the axion momentum p. As we will see, this process is less kinematically suppressed than the other two production processes.

5.2 Primakoff Effect

The Primakoff process, $Ze + \gamma \to Ze + a$, converts photons into axions in the electrostatic field of a nucleus or an electron. To derive the differential flux associated with this process, we can again start from the Boltzmann equation and follow the same arguments used in the photon coalescence case:

$$\frac{df_a}{dt}(E) = \frac{1}{2E}\int \frac{d^3k}{(2\pi)^3 2\omega}(2\pi)\delta(\omega - E)|\mathcal{M}_\gamma|^2 f_\gamma^{eq}(\omega)(1 + f_a(E)). \tag{5.20}$$

Here, the axion has energy E and momentum p, while the photon scattering off the static nucleus has energy ω and momentum k. The Dirac delta enforces energy conservation, and we fix the momentum transferred by the nucleus to the axion-photon system to be $q = k - p$. Following the discussion in Ref. [190], we assume that the recoiling nucleus carries no energy; therefore, $q = (0, \mathbf{q})$. From this, we can write the averaged squared matrix element for a photon interacting with a nucleus of charge Ze as

$$|\bar{\mathcal{M}}| = \frac{g_{a\gamma\gamma}^2 Z^2 e^2}{2}\frac{|\mathbf{k}\times\mathbf{p}|^2}{q^4}\cdot\frac{q^2}{q^2+\kappa^2}. \tag{5.21}$$

In general, one should sum over the contributions from all nuclear species. In Eq. (5.21), κ is the Debye screening scale, defined as

$$\kappa^2 = \frac{4\pi\alpha}{T}\frac{\rho}{m_u}\left(Y_e + \sum_j Z_j^2 Y_j\right). \tag{5.22}$$

Here, Y_e is the number of electrons per baryon, while Y_j is the number of nuclei of species j per baryon in the medium. The Debye screening length is the typical length scale of an electromagnetic interaction inside the plasma. Since the photon acquires an effective mass in a plasma, the potential of a point-like charge is modified according to the substitution $\frac{1}{r} \to \frac{e^{-\kappa r}}{r}$. We show below the Debye screening scale as a function of the solar radius, together with its ratio to the temperature:

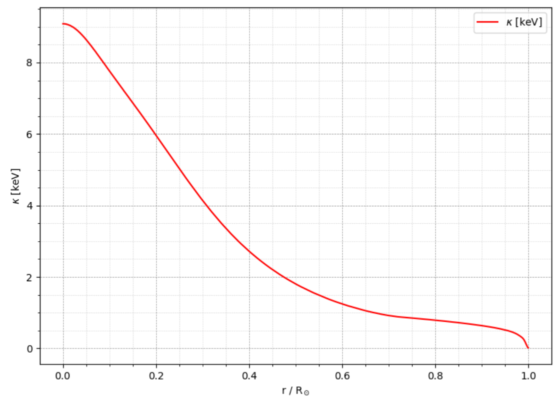

Figure 5.5: Debye screening scale as a function of the solar radius.

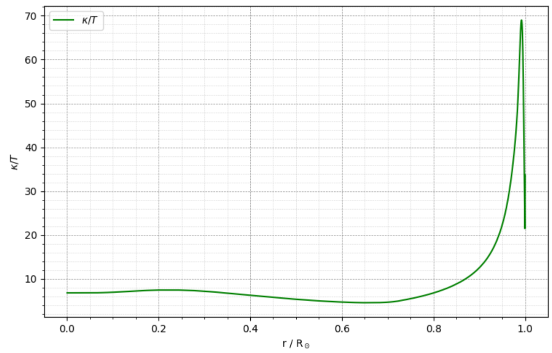

Figure 5.6: κ/T is ~ 7 over most of the solar volume [190].

Choosing $p^\mu = (E, 0, 0, p)$ for the axion four-momentum, we can explicitly write the integral over θ. We also perform the trivial integral over ϕ, which gives a factor of 2π:

$$\begin{aligned}\frac{df_a}{dt}(E) =& \frac{g_{a\gamma\gamma}^2 Z^2 e^2}{16\pi E}\int \frac{k^2}{\omega}\,dk\,\delta\left(\sqrt{k^2+\omega_P^2}-E\right) f_\gamma^{eq}(\omega)(1+f_a(E))\\ &\times \int_{-1}^{1} d\cos\theta\,\frac{k^2p^2(1-\cos^2\theta)}{(p^2+k^2-2pk\cos\theta)(p^2+k^2-2pk\cos\theta+\kappa^2)}.\end{aligned} \tag{5.23}$$

The angular integral can be performed analytically, giving the same result as in Ref. [104]. We again include the factor 2π from the integration over the azimuthal angle ϕ, and perform the integration over k by applying the properties of the Dirac delta:

$$\begin{aligned}\frac{df_a}{dt}(E) = \frac{g^2Z^2e^2}{32\pi}\frac{k}{E}&\left[\frac{\left((k+p)^2+\kappa^2\right)\left((k-p)^2+\kappa^2\right)}{4kp\kappa^2}\ln\left(\frac{(k+p)^2+\kappa^2}{(k-p)^2+\kappa^2}\right)\right.\\ &\left.-\frac{(k-p)^2}{4kp\kappa^2}\ln\left(\frac{(k+p)^2}{(k-p)^2}\right)-1\right] f_\gamma^{eq}(\omega)(1+f_a(E)).\end{aligned} \tag{5.24}$$

We now restore the contribution from each nuclear species and obtain

$$\begin{aligned}\frac{df_a}{dt}(E) = \frac{g_{a\gamma\gamma}^2\kappa^2 T}{32\pi}\frac{k}{E}&\left[\frac{\left((k+p)^2+\kappa^2\right)\left((k-p)^2+\kappa^2\right)}{4kp\kappa^2}\ln\left(\frac{(k+p)^2+\kappa^2}{(k-p)^2+\kappa^2}\right)\right.\\ &\left.-\frac{(k-p)^2}{4kp\kappa^2}\ln\left(\frac{(k+p)^2}{(k-p)^2}\right)-1\right]\cdot\frac{2}{e^{\omega/T}-1}\cdot\frac{e^{E/T}}{e^{E/T}-1}.\end{aligned} \tag{5.25}$$

In the last two factors, we have explicitly written the thermal distributions; the factor of two comes from the two photon polarizations. Schematically, Eq. (5.25) can be written as

$$\frac{df_a}{dt}(E) = \frac{g_{a\gamma\gamma}^2\kappa^2 T}{32\pi}\frac{k}{E} f_\gamma^{eq}(\omega)(1+f_a(E))\cdot F_k(p). \tag{5.26}$$

As shown in Ref. [104], the function $F_k(p)$, corresponding to the large square brackets in Eq. (5.25), and it can be approximated in different ways. We denote by $F_k(p)$ the exact expression, by $F_E(p)$ the expression obtained by replacing k with E, and, for small momenta, by

$$F(p) = \frac{8p^2}{3(\kappa^2+m_a^2)}. \tag{5.27}$$

We show below the values of these functions for different radii, comparing the three expressions:

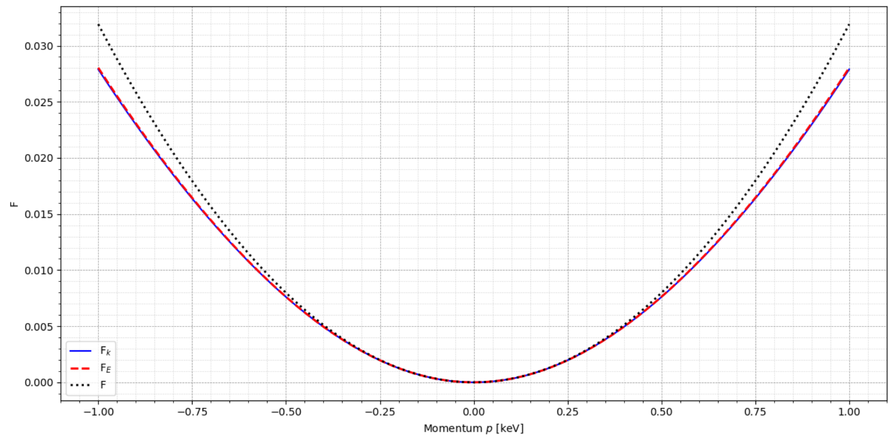

Figure 5.7: Comparison between F_k, F_E, and F for $m_a = 1$ keV. Up to 0.5 keV, all three functions are in excellent agreement. We fix $r = 0$.

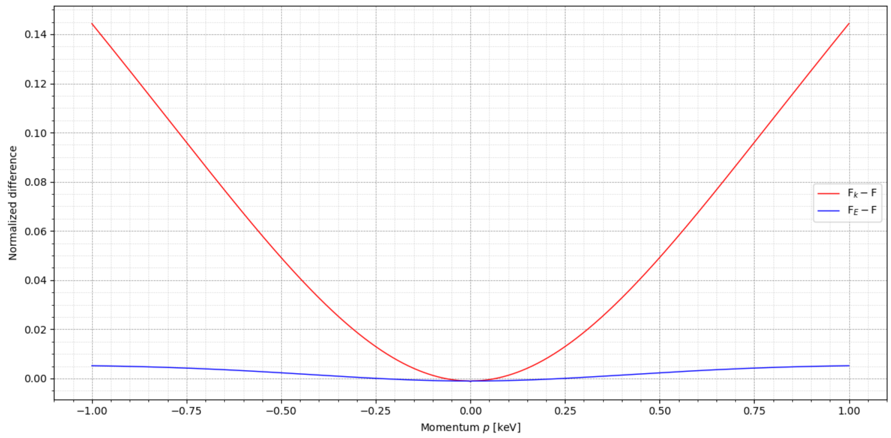

Figure 5.8: Normalized difference between F_E, F_k, and F for $m_a = 1$ keV. The corresponding errors are about 1% and 15%, respectively; however, restricting to momenta below 0.1 keV, the error is below 1% in both cases.

As already done in the photon coalescence case, we can obtain the number of axions produced per unit volume, energy, and time through the equation

$$\frac{d^2 n_a(E)}{dEdt} = \frac{pE}{2\pi^2}\frac{df_a(E)}{dt}. \tag{5.28}$$

Integrating over the same solar model, we obtain the differential axion flux:

$$\frac{d\Phi_{prim}}{dE} = \frac{1}{4\pi d^2}\int d^3r \frac{dn_a(E)}{dE} = \frac{R_\odot^3}{d^2}\int \frac{dn_a(E)}{dE}. \tag{5.29}$$

The results are shown below:

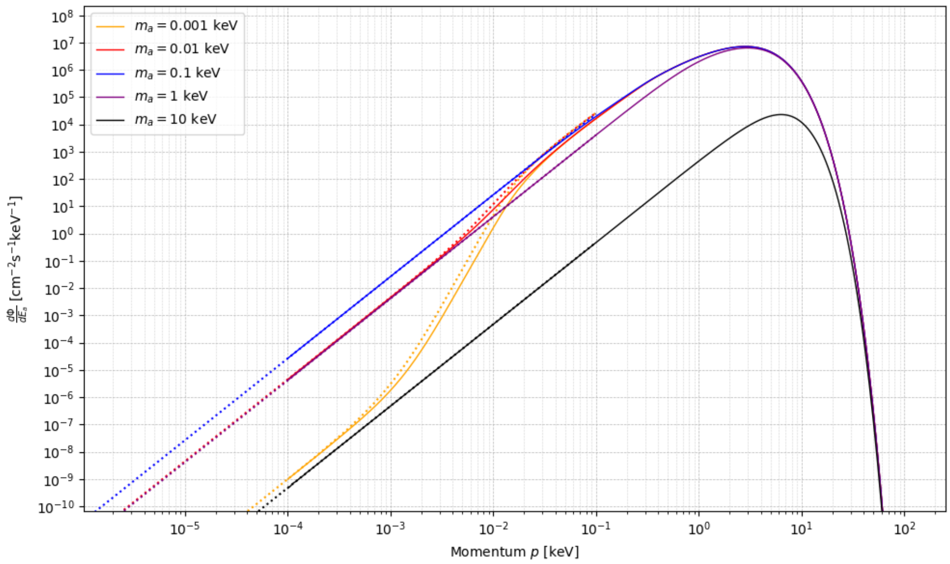

Figure 5.9: Differential flux from the Primakoff process for different axion masses. Note the logarithmic scale and the suppression of the process at low momenta. Due to numerical errors, we could not compute the full expression for the flux $d\Phi/dE$ at low momenta. Therefore, we plot the full expression for momenta between 10^{-4} keV and 10^2 keV, and the approximate expression given in Eq. (5.27) in the interval between 10^{-6} keV and 10^{-1} keV.

5.3 Plasmon Decay

We now study the third and last production mechanism: plasmon decay and coalescence. We want to compute the conversion rate for $\gamma_t \to a$. From now on, we distinguish between the two excitation modes of the electromagnetic field in a plasma: transverse modes, γ_t, and longitudinal modes, γ_l, each with its own dispersion relation:

$$\begin{aligned} \omega^2 &= k^2 + \omega_P^2 \left(1 + \frac{k^2}{\omega^2}\frac{T}{m_e}\right) \simeq k^2 + \omega_P^2, \\ \omega^2 &= \omega_P^2 \left(1 + 3\frac{k^2}{\omega^2}\frac{T}{m_e}\right) \simeq \omega_P^2. \end{aligned} \tag{5.30}$$

For photon coalescence and the Primakoff process, we used, respectively, two transverse photons and one transverse photon. The main reason why we distinguish between the Primakoff process and plasmon decay and coalescence, unlike what is done in other treatments such as Ref. [190], is that in the latter case we do not assume that the energy of the longitudinal photon vanishes. For the Primakoff process, this corresponds to assuming that the recoil energy of the nucleus is negligible, i.e. that $m_Z \gg E$. Furthermore, we associate a thermal Maxwell-Boltzmann distribution to the longitudinal photons; this assumption also changes the final result with respect to the Primakoff process.
As already discussed, from Fig. (5.1), the temperature inside the Sun never exceeds 1.3 keV; therefore, the thermal corrections in Eq. (5.30) are negligible.
As in the previous cases of photon coalescence and the Primakoff process, we start from the Boltzmann equation:

$$\begin{aligned} \frac{df_a}{dt} = \frac{1}{2E}\int & \frac{d^3k_t}{(2\pi)^3 2\omega_t}\frac{d^3k_l}{(2\pi)^3 2\omega_l}(2\pi)^4 Z_l Z_t |\mathcal{M}|^2 \\ & \times f_{\gamma_t}(\omega_T)\left(\delta^4(k_t - k_l - p)(1 + f_{\gamma_l}) + \delta^4(k_t + k_l - p) f_{\gamma_l}\right)(1 + f_a(E)). \end{aligned} \tag{5.31}$$

The first term in brackets is associated with plasmon decay, while the second term corresponds to the coalescence of a transverse and a longitudinal plasmon. Notice the different thermal factor.

The Z's are vertex renormalization factors. In our case, we work in a non-relativistic and non-degenerate plasma, as appropriate for the Sun. Therefore, following Ref. [190], we can approximate $Z_l = \frac{\omega_l^2}{\omega_l^2 - k_l^2}$ and $Z_t = 1$.
For the matrix element, we note that the interaction Lagrangian can be rewritten as

$$\mathcal{L} = -g_{a\gamma\gamma} E \cdot B. \tag{5.32}$$

Since $B = \nabla \times A$, no magnetic field can be associated with the longitudinal mode. Given $E = -\nabla A^0 - \partial_t A$, and using a wave decomposition for the electromagnetic field, we obtain $B \propto k_t \times \epsilon_t$ and $E \propto \omega_l \epsilon_l - k_l \epsilon_l^0$. Following the discussion in Ref. [190], the normalized polarization vector for the longitudinal photon is

$$\epsilon_l = \frac{(k_l^2, \omega_l k_l)}{k_l \sqrt{\omega_l^2 - k_l^2}}.$$

Putting everything together, we can write the averaged matrix element as

$$|\bar{\mathcal{M}}|^2 = \frac{g_{a\gamma\gamma}^2}{2}(\omega_l^2 - k_l^2)|\hat{k}_l \cdot (k_t \times \epsilon_t)|^2 = \frac{g_{a\gamma\gamma}^2}{2}(\omega_l^2 - k_l^2)|\hat{k}_l \times k_t|^2. \tag{5.33}$$

In the last equality, we used $a \cdot (b \times c) = c \cdot (a \times b)$. Following the discussion in Ref. [191], one may also include an additional factor:

$$|\bar{\mathcal{M}}|^2 = \frac{g_{a\gamma\gamma}^2}{2}(\omega_l^2 - k_l^2)|\hat{k}_l \times k_t|^2 \frac{\omega_l^2}{\omega_l^2 + k_l^2}. \tag{5.34}$$

We will present the computation both with and without this thermal factor.

After integrating over $d^3 k_l$, the delta function gives $k_l = k_t \pm p$, where p is the three-momentum of the axion. The sign depends on the process considered. We obtain

$$\begin{aligned} \frac{df_a}{dt} &= \frac{g_{a\gamma\gamma}^2}{16(4\pi^2)E} \int d^3 k_t \frac{\omega_l}{\omega_t} f_{\gamma_t}(\omega_t) k_t^2 p^2 \sin^2\theta (1 + f_a(E)) \\ &\quad \times \left(\frac{(1 + f_{\gamma_l})\delta(\omega_t - \omega_l - E)}{p^2 + k_t^2 - 2pk\cos\theta} + \frac{f_{\gamma_l}\delta(\omega_t + \omega_l - E)}{p^2 + k_t^2 + 2pk\cos\theta} \right). \end{aligned} \tag{5.35}$$

With the factor introduced in Eq. (5.34), the same integral becomes

$$\begin{aligned} \frac{df_a}{dt} &= \frac{g_{a\gamma\gamma}^2}{16(4\pi^2)E} \int d^3 \mathbf{k}_t \, \frac{\omega_l^3}{\omega_t} f_{\gamma_t}(\omega_t)\, k_t^2 p^2 \sin^2\theta \, (1 + f_a(E)) \times \\ &\quad \left[\frac{(1 + f_{\gamma_l})\, \delta(\omega_t - \omega_l - E)}{(p^2 + k_t^2 - 2pk\cos\theta)(p^2 + k_t^2 - 2pk\cos\theta + \omega_l^2)} \right. \\ &\quad \left. + \frac{f_{\gamma_l}\, \delta(\omega_t + \omega_l - E)}{(p^2 + k_t^2 + 2pk\cos\theta)(p^2 + k_t^2 + 2pk\cos\theta + \omega_l^2)} \right]. \end{aligned} \tag{5.36}$$

We explicitly write the angular integral, choosing the z-axis along the direction of $\mathbf{p}$. One angular integration is trivial and gives a factor of 2π. For the term without the additional factor, we have

$$\begin{aligned} \frac{df_a}{dt} &= \frac{g_{a\gamma\gamma}^2}{32\pi E} \int dk_t k_t^2 \frac{\omega_l}{\omega_t} f_{\gamma_t}(\omega_t)\, [\delta(\omega_t - \omega_l - E)(1 + f_{\gamma_l}) + f_{\gamma_l}\delta(\omega_t + \omega_l - E)] \\ &\quad \times \int_{-1}^{1} d\cos\theta \frac{k_t^2 p^2 \sin^2\theta}{p^2 + k_t^2 + 2pk\cos\theta}(1 + f_a(E)). \end{aligned} \tag{5.37}$$

The integral over θ can be performed analytically in both cases. We obtain

$$\begin{aligned} \frac{df_a}{dt} &= \frac{g_{a\gamma\gamma}^2}{32\pi E} \int dk_t k_t^2 \frac{\omega_l}{\omega_t} f_{\gamma_t}(\omega_t)\, [\delta(\omega_t - \omega_l - E)(1 + f_{\gamma_l}) + f_{\gamma_l}\delta(\omega_t + \omega_l - E)] \\ &\quad \times (1 + f_a(E)) \frac{4k_t p(p^2 + k_t^2) - (k_t^2 - p^2)^2 \log\left(\frac{(k_t+p)^2}{(k_t-p)^2}\right)}{8k_t p}. \end{aligned} \tag{5.38}$$

and, when including the additional factor,

$$\begin{aligned}\frac{df_a}{dt} = \frac{g_{a\gamma\gamma}^2}{32\pi E} \int dk_t k_t^2 \frac{\omega_l}{\omega_t} f_{\gamma_t}(\omega_t) \left[\delta(\omega_t - \omega_l - E)(1+f_{\gamma_l}) + f_{\gamma_l}\delta(\omega_t + \omega_l - E)\right] \\ \times (1+f_a(E))\omega_p^2 \Bigg[\frac{\left((k+p)^2+\omega_p^2\right)\left((k-p)^2+\omega_p^2\right)}{4kp\omega_p^2} \ln\left(\frac{(k+p)^2+\omega_p^2}{(k-p)^2+\omega_p^2}\right) \\ - \frac{(k-p)^2}{4kp\omega_p^2} \ln\left(\frac{(k+p)^2}{(k-p)^2}\right) - 1 \Bigg].\end{aligned} \tag{5.39}$$

Using $\omega_l = \omega_p$ and the properties of the Dirac delta, we get

$$\begin{aligned}\frac{df_a}{dt} = \frac{g_{a\gamma\gamma}^2}{32\pi E}\omega_l f_{\gamma_t}(\omega_t)(1+f_a(E)) \Bigg((1+f_{\gamma_l}) \frac{4k_t p(p^2+k_t^2) - (k_t^2-p^2)^2 \log\left(\frac{(k_t+p)^2}{(k_t-p)^2}\right)}{8p} \\ + f_{\gamma_l} \frac{4k_t p(p^2+k_t^2) - (k_t^2-p^2)^2 \log\left(\frac{(k_t+p)^2}{(k_t-p)^2}\right)}{8p} \Bigg),\end{aligned} \tag{5.40}$$

and

$$\begin{aligned}\frac{df_a}{dt} =& \frac{g_{a\gamma\gamma}^2}{32\pi E} k_t \omega_p^3 f_{\gamma_t}(\omega_t)(1+f_a(E)) \\ &\times \Bigg[(1+f_{\gamma_l}) \left(\frac{\left((k+p)^2+\omega_p^2\right)\left((k-p)^2+\omega_p^2\right)}{4kp\omega_p^2} \ln\left(\frac{(k+p)^2+\omega_p^2}{(k-p)^2+\omega_p^2}\right) - \frac{(k-p)^2}{4kp\omega_p^2} \ln\left(\frac{(k+p)^2}{(k-p)^2}\right) - 1 \right) \\ &+ f_{\gamma_l} \left(\frac{\left((k+p)^2+\omega_p^2\right)\left((k-p)^2+\omega_p^2\right)}{4kp\omega_p^2} \ln\left(\frac{(k+p)^2+\omega_p^2}{(k-p)^2+\omega_p^2}\right) - \frac{(k-p)^2}{4kp\omega_p^2} \ln\left(\frac{(k+p)^2}{(k-p)^2}\right) - 1 \right) \Bigg].\end{aligned} \tag{5.41}$$

In both cases, in the first term k_t is fixed by the relation $\omega_t = \omega_p + E$, while in the second term it is fixed by $\omega_t + \omega_p = E$.

We have used the relations

$$\begin{aligned} f_{\gamma_t}(1+f_a)(1+f_{\gamma_l}) &= (1+f_{\gamma_t})f_a f_{\gamma_l}, \\ f_{\gamma_t} f_{\gamma_l}(1+f_a) &= f_a(1+f_{\gamma_t})(1+f_{\gamma_l}). \end{aligned} \tag{5.42}$$

Finally, to obtain the differential flux, we integrate over the Saclay solar model. As before, to express the differential flux in units of cm^{-2} s^{-1} keV^{-1}, we multiply the integral by a factor 2.8×10^{11}. We present below the results both with and without the factor used in Ref. [191], choosing $g_{a\gamma\gamma} = 10^{-12}$ GeV^{-1}:

$$\frac{d\Phi_{plas}}{dE} = \frac{1}{4\pi d^2} \int d^3r \frac{pE}{2\pi^2} \frac{df_a}{dt}(E) = \frac{R_\odot^3}{d^2} \int \frac{dn_a(E)}{dE}. \tag{5.43}$$

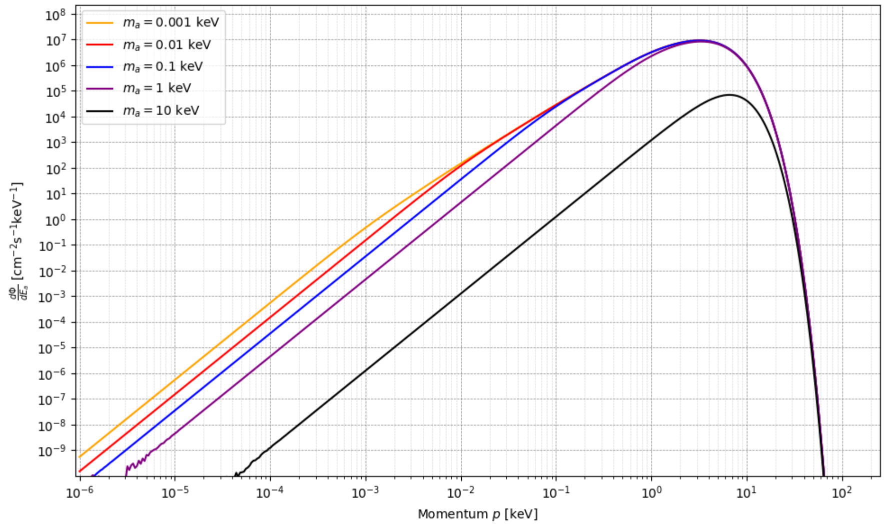

Figure 5.11: Differential flux for the plasmon process without the thermal factor, as in [191].

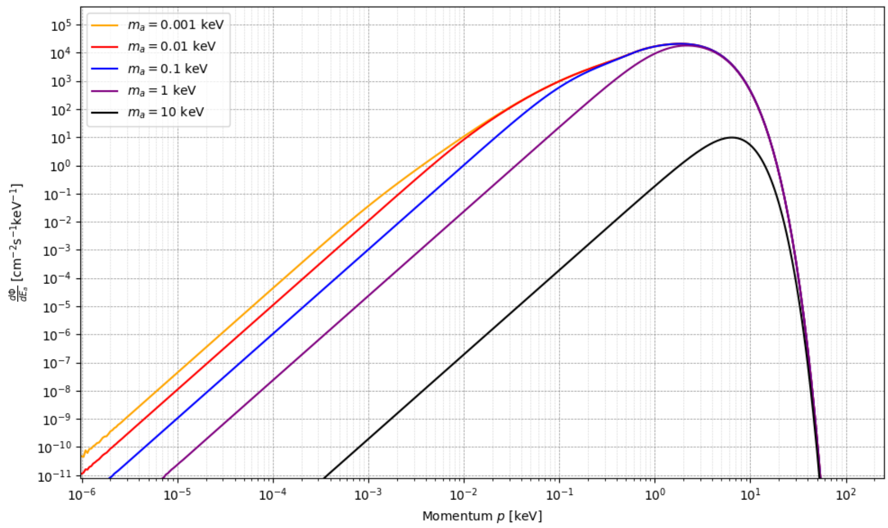

Figure 5.10: Differential flux for different values of the axion mass for the plasmon processes with the thermal factor.

One can easily see that the inclusion of the thermal factor does not drastically change our conclusions. Furthermore, as we will see in the next section, we are interested in the study of solar basins, which are produced by axions with low momenta. In this regime, the dominant production process is photon coalescence.

5.4 Number Density of Trapped Axions

We now focus on computing the axion number density in the vicinity of the Earth. Low-momentum axions are expected to remain gravitationally trapped around the Sun, forming a cosmic axion basin, as shown in Ref. [111]. The long Solar-System timescales over which this production mechanism remains active allow a significant axion number density to accumulate around the Sun. The large astrophysical timescales involved in this scenario can therefore lead to a substantial population of trapped axions. The accumulated trapped population can exceed the instantaneous solar axion population by several orders of magnitude. This argument becomes stronger when it is applied to KK axions, since in this case inside the Sun even more modes can be produced.

We first present a comparison, for different axion masses, of the production rates inside the Sun for the three different processes:

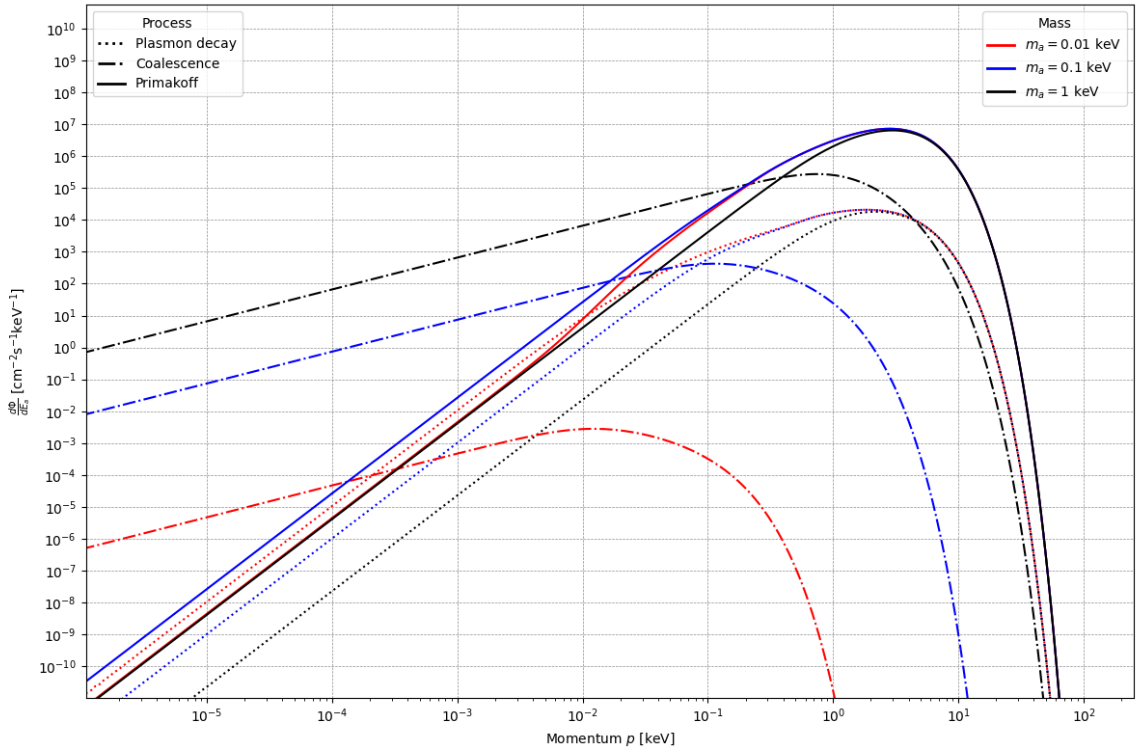

Figure 5.12: Comparison between the three axion production processes inside the Sun for different axion masses.

From the figure, it is clear that, at low values of the axion momentum p, photon coalescence is the dominant production process. We therefore focus on this process.

For an axion to remain gravitationally bound to the Sun, its speed at production must be below the escape velocity at the production point r_0. As a reference, at the solar surface one has $v_{\rm esc}(R_\odot) \simeq 6.0 \times 10^2\,{\rm km\,s^{-1}} \approx 2 \times 10^{-3}\,c$. For a representative axion mass of $1\,$keV, this corresponds to a momentum threshold $p \simeq m_a v_{\rm esc} \approx 2 \times 10^{-3}\,{\rm keV}/c$. As shown in Fig. (5.12), at this momentum, photon coalescence is already the dominant production channel.

To compute the number density, we need to perform several integrals. We first consider the production rate at a distance r for a mode of mass m, keeping in mind that we will later apply the same computation to KK axions:

$$P_{KK}(r,m) = 4\pi \int dr_0 r_0^2 \int d^3p \Gamma_{a\gamma}(T(r_0),p) f_a(E) P_r(r_0,p). \tag{5.44}$$

Here, $P_{KK}(r,m)$ is the production rate of the number density at distance r for axions of mass m, produced per unit time. The inner integral accounts for the contribution from all energies, or more generally from all axion momenta, while the outer integral sums over all possible production positions inside the Sun. We have also introduced a new quantity, $P_r(r_0,p)$, which is the probability for an axion produced at position r_0 with momentum p to be found at distance r.

Of course, Eq. (5.44) can be interpreted as the number of axions produced inside the Sun per unit time, given by the first two integrals, multiplied by the probability of finding the axion at distance r from the Sun. At the end of our computation, we will set $r = 1$ A.U.

For the moment, we focus on a fixed axion mass. We will then apply our computation either to a generic

ALP with fixed mass or to a Kaluza–Klein axion tower. In the latter case, we will need to integrate over all masses, taking into account the corresponding degeneracy, as discussed in Chapter (4).

The probability density $P_r(r_0, p)$ can be obtained by noting that the probability of finding an axion between r and $r + dr$, given its initial conditions, is proportional to the time the axion spends in that radial interval. This time is proportional to the inverse of the radial velocity $v_r(r)$. This reasoning applies only if the axion is in a bound orbit. If the trajectory is unbound, then over sufficiently long times the probability of finding the axion between r and $r + dr$ for any finite r goes to zero. Therefore, we must integrate only over momenta, and correspondingly velocities, that give bound trajectories.

Following Ref. [191], we first compute the gravitational potential energy per unit mass inside the Sun, given a density profile $\rho(r)$:

$$V(r) = -4\pi G R_\odot^2 \left(\int_0^r \frac{r'^2}{r} \rho(r') dr' + \int_r^1 r' \rho(r') dr' \right). \tag{5.45}$$

Here, $R_\odot$ is the solar radius, and all variables inside the integrals are normalized. It is straightforward to show that this potential reproduces the classical Newtonian force law; equivalently, one can derive this result from Poisson's equation.

In Figs. (5.13) and (5.14), we first show the solar density profile, taken from Ref. [189], and then the corresponding gravitational potential:

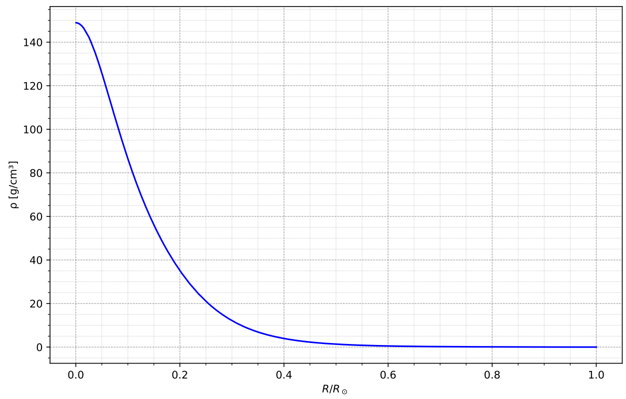

Figure 5.13: Density inside the Sun as a function of the solar radius.

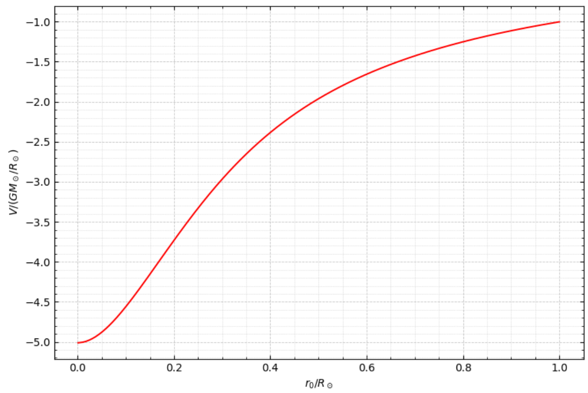

Figure 5.14: Potential inside the Sun, normalized to minus the potential at the solar surface.

Because of the low momenta involved, we perform the computation in the classical non-relativistic limit. Imposing energy conservation, we obtain

$$\frac{1}{2} R_\odot^2 \dot{r}^2 + \frac{R_\odot^2 l^2}{2r^2} + V(r) = \text{const.} \tag{5.46}$$

We used $l = r^2 \dot{\theta}$, and again r denotes the radius in units of the solar radius. Of course, the factor $R_\odot$ can be simplified. From this expression, it is clear that we can write

$$\dot{r} = \sqrt{2(E(r_0, v_0) - U(r))}. \tag{5.47}$$

For $r < 1$, we define

$$U(r) = \frac{l^2}{2r^2} - 4\pi G \left(\int_0^r \frac{r'^2}{r} \rho(r') dr' + \int_r^1 r' \rho(r') dr' \right), \tag{5.48}$$

while for $r > 1$ we have

$$U(r) = \frac{l^2}{2r^2} - \frac{GM}{r}. \tag{5.49}$$

From the discussion above, it follows that

$$P_r(r_0, v_0) = \frac{N}{\sqrt{2(E(r_0, v_0) - U(r))}}. \tag{5.50}$$

Here, N is a normalization factor, which can be written in the form $N = 1/4\pi I$, where I is given by

$$I = \int_{r_{min}}^{r_{max}} \frac{r^3}{(2E(r_0, v_0)r^2 + 2GMr - l^2)^{\frac{1}{2}}}. \tag{5.51}$$

This formula is strictly valid only for $r > 1$. Nevertheless, we are interested in orbits that reach distances much larger than the solar radius, since $R_\odot \simeq 0.005$ A.U.. We therefore approximate the potential inside the Sun by the expression in Eq. (5.49). This has the advantage that N does not need to be computed numerically for each value of r_0, and instead an analytical closed-form expression can be obtained. Moreover, this is an excellent approximation: with this choice of N, the integral over all space for fixed r_0, v_{a0}, and v_{r0} is equal to one up to an error of order $\sim 0.5\%$.

The integral is restricted to the allowed values of r, where r_{min} and r_{max} are obtained, as usual, from the condition that the kinetic energy vanishes. Computing Eq. (5.51) under these assumptions, we obtain

$$I = \frac{GM(5G^2M^2 - 6|E|l^2)}{16\sqrt{2}|E|^{7/2}}\pi. \tag{5.52}$$

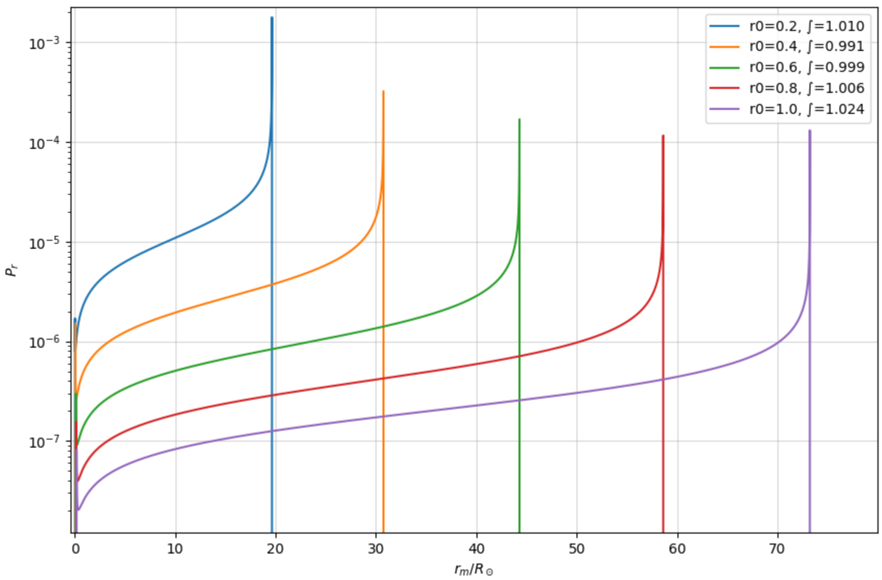

Figure 5.15: Probability for $v_{r0} = 0.9\,v_{esc}$, $v_{a0} = 0.42\,v_{esc}$, and for different values of $r_0 = \{0.2, 0.4, 0.6, 0.8, 1\}$. Similar behavior is observed for different values of v_{r0} and v_{a0}. The distribution has the expected physical interpretation: it is inversely proportional to the radial velocity. Therefore, two peaks appear near the minimum and maximum distances, while the minimum of the distribution, corresponding to the maximum radial velocity, lies near the minimum radius, since the radial velocity increases as the axion moves toward the Sun.

For the orbit to be bound, the energy E must be negative. We therefore have

$$N = \frac{4\sqrt{2}|E|^{7/2}}{GM(5G^2M^2 - 6|E|l^2)\pi^2}. \tag{5.53}$$

We are now ready to compute Eq. (5.44). First, the momentum integration measure, in the non-relativistic regime, becomes $d^3p \simeq m^3v^2 dv d\cos\theta d\phi$. We switch to cylindrical coordinates, choosing the z-direction along the axis connecting the axion production point r_0 and the center of the Sun, while the ρ-direction coincides with the component of the velocity orthogonal to the z-axis, corresponding to $r\dot{\theta}$ in the previous spherical coordinates. There is no velocity component along the ϕ-direction, so the integration over this variable is trivial. We therefore obtain the integration measure $d^3p \simeq 2\pi m^3 v_a dv_a dv_r$, where v_a and v_r are the angular and radial velocities, respectively. Taking into account all factors of π, we obtain

$$P_{KK}(r,m) = \frac{8\pi^2 R_\odot^3}{8\pi^3} \int m^3 v_a dv_a dv_r \int dr_0 r_0^2 \Gamma_{a\gamma}(T(r_0),p,m) f_a(E,m) P_r(r_0,p). \tag{5.54}$$

We now have all the ingredients needed to compute the number density for each axion mode, $n_{KK}(r,m)$:

$$\frac{dn_{KK}(r,m)}{dt} = -\Gamma_0(m) n_{KK}(r,m) + P_{KK}(r,m). \tag{5.55}$$

The solution to this differential equation is easily found

$$n_{KK}(r,m) = \frac{P_{KK}(r,m)}{\Gamma_0(m)}(1-\exp(-\Gamma_0(m)t_\odot)). \tag{5.56}$$

We now need to define $t_\odot$. Several considerations must be taken into account. First, since we have adopted a specific solar model for our integration, we must ensure that this model remains valid over the full integration time. This requirement is essentially satisfied, since the Sun has remained in a quasi-steady state from approximately 50 million years after its formation and, for our purposes, has been equivalent to its current state.

Following Ref. [111], a further consideration comes from the dynamical stability of the Solar System. The Lyapunov time at $R = 1$ A.U. is about $t_L = 5 \times 10^6$ yr. Taking this as the timescale over which an orbit efficiently explores the entire phase space, a conservative estimate is to set the characteristic ejection time from the Solar System equal to twice t_L, giving $t_\odot = 10^7$ yr. Another possible estimate comes from the study of Near-Earth Objects (NEOs), leading to $t_\odot = 10^8$ yr. However, this value is likely underestimated. The dominant contribution to the trapped-axion number density comes from directionally isotropic, almost radial orbits with aphelia close to 1 A.U., at all inclinations i with respect to the ecliptic plane. In contrast, NEOs are preferentially located near the ecliptic plane, $i \simeq 0$, and therefore near unstable regions of phase space, since they are usually resonantly driven away from the asteroid belt or, to a lesser extent, captured through close encounters in the inner Solar System. Furthermore, the NEO population exhibits strong correlations among its orbital parameters; thus, because this sample is biased, it is not a reliable statistical sample from which to infer $t_\odot$.

We also note that, in principle, one should consider the possibility that trapped axions are reabsorbed inside the Sun, but we neglect this effect. As recently found in Ref. [192], $t_\odot = (1.20 \pm 0.09) \times 10^9$ yr.

To compute $P_{KK}(r,m)$, we use the rate obtained for photon coalescence, which we rewrite here for convenience:

$$\Gamma_{a\gamma\gamma}(T) = \frac{g_{a\gamma\gamma}^2 m_a^3}{64\pi} \frac{(m_a^2 - 4\omega_P^2)}{m_a^2} \frac{m_a}{E} \left(1 + \frac{2T}{p} \ln\left(\frac{1-e^{-\frac{E+p}{2T}}}{1-e^{-\frac{E-p}{2T}}}\right)\right). \tag{5.57}$$

Following Ref. [191], we compare it with

$$\Gamma_{a\gamma\gamma}(T) = \frac{g_{a\gamma\gamma}^2 m_a^3}{64\pi} \frac{m_a}{E} = \Gamma_0(m)\frac{m_a}{E}. \tag{5.58}$$

The difference is of order $\mathcal{O}(1)$ for light masses, up to ~ 3 keV. We show below the number density associated with several axion masses:

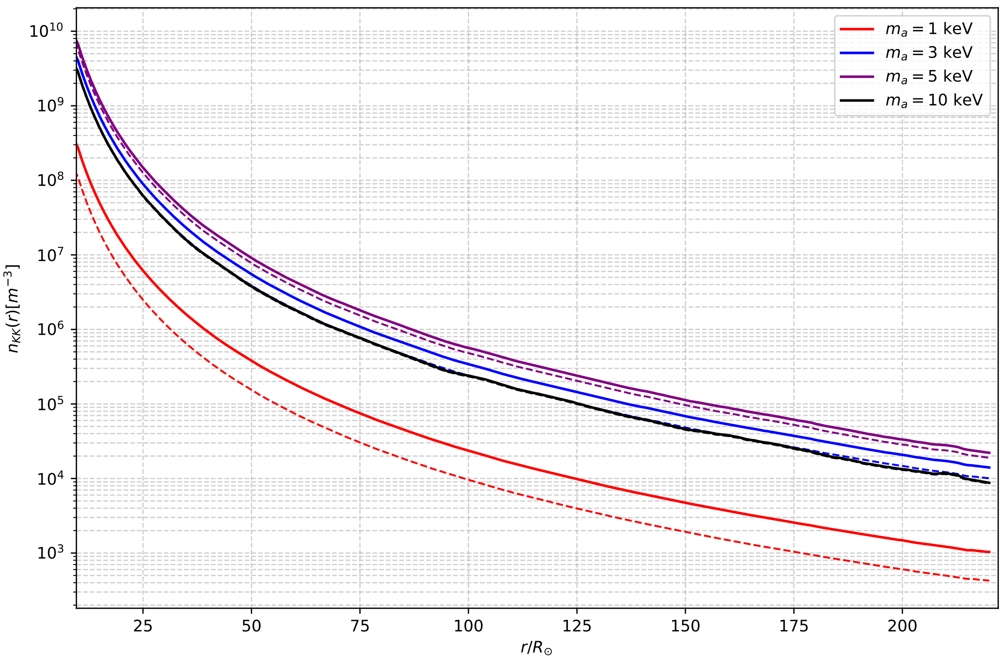

Figure 5.16: Number density associated with several axion masses. The dashed lines correspond to $\Gamma_{a\gamma\gamma} = \Gamma_0(m)\frac{m_a}{E}$, while the solid lines correspond to the full $\Gamma_{a\gamma\gamma}$ computed for photon coalescence. For light masses, the difference is of order $\mathcal{O}(1)$. For heavier masses, we encountered computational difficulties, but their contribution is suppressed by the Bose-Einstein distribution function. We observe that the number density scales as $\sim 1/r^4$, as found in Ref. [104].

At this point, we can apply this machinery either to a generic ALP with fixed mass or to Kaluza–Klein axions. As anticipated, the main difference is that, in the KK case, several modes can contribute, potentially enhancing the axion number density.

5.4.1 ALP

We first apply our computation to the simpler case of a generic ALP interacting with photons as in Eq. (5.4). Having obtained the axion number density for a fixed mass, we can compute the decay rate per unit volume for an axion of mass m at a distance r from the Sun simply as

$$R = n(r, m)\Gamma_0(m, g_{a\gamma\gamma}).$$

Here, $\Gamma_0(m, g_{a\gamma\gamma}) = g_{a\gamma\gamma}^2 m^3/64\pi$ is the standard decay rate. We show below R for some fixed values of the mass as a function of $g_{a\gamma\gamma}$, and for some fixed values of $g_{a\gamma\gamma}$ as a function of the mass.

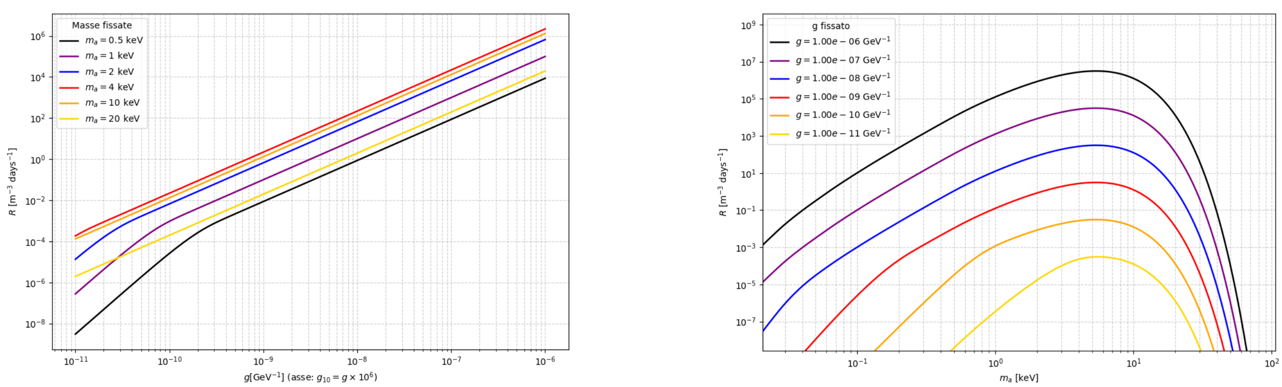

Figure 5.17: Rate as a function of the coupling $g_{a\gamma\gamma}$ for different axion masses (left) and as a function of the axion mass for some fixed coupling $g_{a\gamma\gamma}$.

We anticipate that the rate found in this case is smaller than the one obtained for KK axions. This is due to the fact that KK axions have a degeneracy of modes for fixed mass, which introduces an overall factor that increases the rate. Furthermore, the integral over the masses introduces a factor $m^{\delta-1}$, which shifts, for $\delta > 1$, the peak of the rate toward higher masses. Finally, we present the heat-map of R in the $(m, g_{a\gamma\gamma})$ plane.

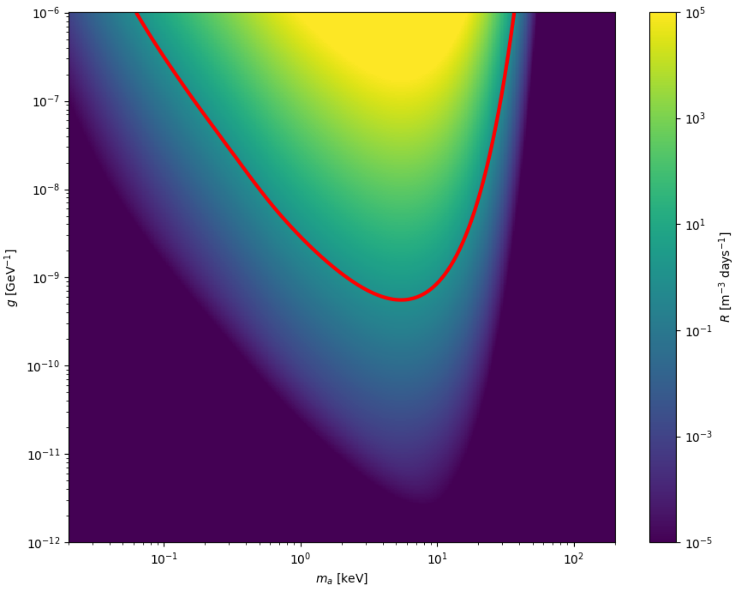

Figure 5.18: The maximum rate occurs for masses around 5 keV. An experiment probing the region within the red line would provide a direct upper bound of approximately 5×10^{-10} GeV^{-1}.

The red line represents the region corresponding to one event per m^3 per day. Increasing, for example, the detector volume or the exposure time would allow one to reach sensitivities of about 10^{-12} GeV^{-1}. Below, we show the same plot within the current bounds in the $(m, g_{a\gamma\gamma})$ plane.

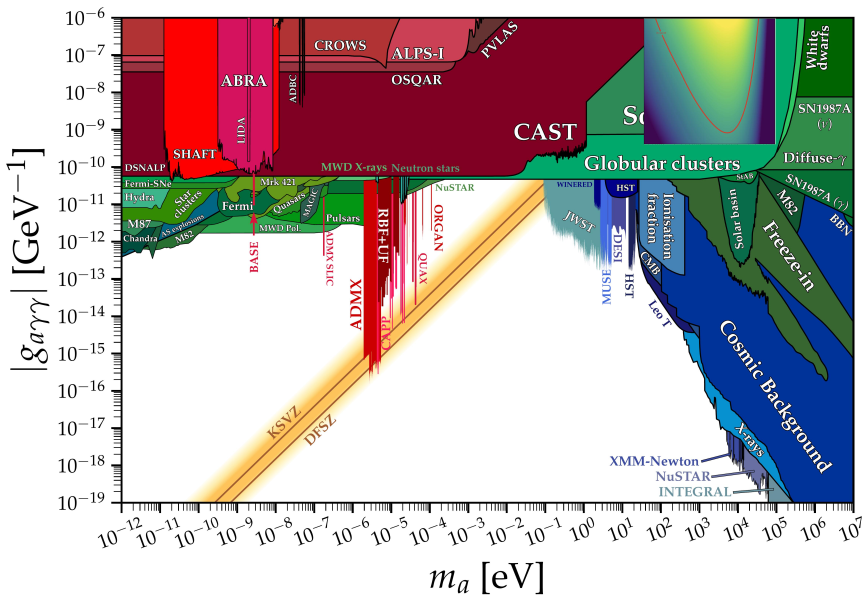

Figure 5.19: ALP parameter space with the region of interest where our considerations apply. The red line represents the region corresponding to one event per m^3 per day.

This region lies within the bound derived from the solar neutrino flux, as discussed in Chapter (3). If this region of parameter space can be probed directly, this bound should be regarded complementary to the other indirect bounds applied.

5.4.2 Kaluza–Klein Axions

We now apply our formalism to KK axions. As anticipated several times, we need to integrate over all masses in order to take into account the contribution of every KK mode. We use the expression

$$n_{\mathrm{KK}}(r) = \frac{2\pi^{\delta/2}}{\Gamma\left(\frac{\delta}{2}\right)} R^{\delta} \int_0^{\infty} m^{\delta-1} \, n_{\mathrm{KK}}(r, m) \, \mathrm{d}m, \tag{5.59}$$

where in practice we integrate only over masses larger than $2\omega_p$.

Here, δ is the number of extra dimensions in which the axion modes propagate, while R is the compactification radius. We present below the results for different values of (δ, g, R), chosen to match those used in Ref. [191].

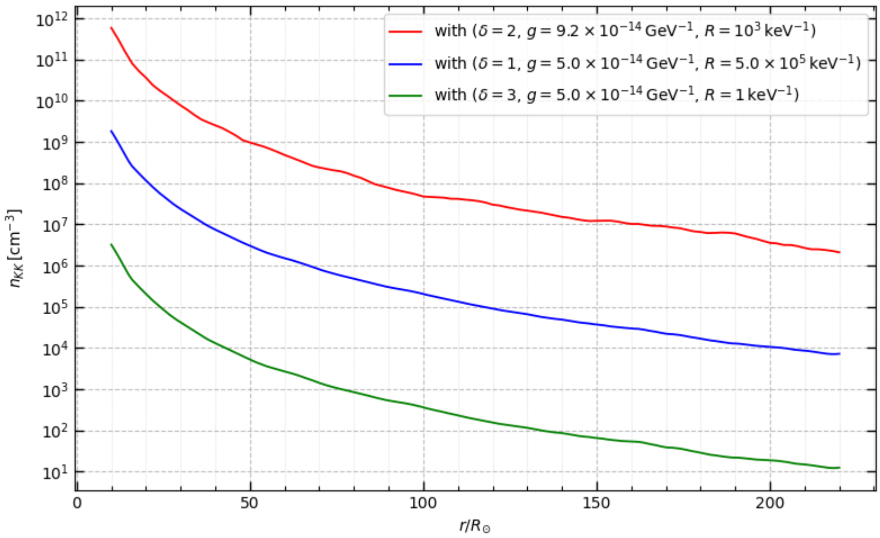

Figure 5.20: Total axion number density as a function of the distance from the Sun, normalized to the solar radius. As before, $n_{KK} \sim 1/r^4$. The integral required a significant computational effort and led to some numerical noise, so we slightly smoothed the curve with a moving average, although some bumps are still visible in the plot. For computational reasons, we omit the approximation used in Ref [191]. The difference between our results and those in Ref. [191] is consistent with an overall factor of about 3.

For this choice of parameters, we obtain a number density of about 2.6×10^6 cm^{-3} at $r = 1$ A.U., which is about 3 times larger than the value found in Ref. [191] and 16 times smaller than the result in Ref. [104].

Finally, we can compute the differential decay rate of trapped KK axions into photons, which is given by

$$\frac{dR}{dm} = \Gamma_{a\gamma\gamma} \frac{dn_{KK}}{dm} = \left(6.51 \times 10^{-12} \mathrm{days}^{-1}\right) g_{10}^2 \left(\frac{m}{\mathrm{keV}}\right)^3 \frac{dn_{KK}}{dm}. \tag{5.60}$$

We present below, in Fig. (5.22), the differential rate for the same parameter values used in Fig. 7 of Ref. [191], evaluated at $r = 1$ A.U. $\simeq 200\, R_\odot$. We obtain the same shape for the differential rate as in Ref. [191], with the same overall factor of about 3 already found for the number density. We first present $\frac{dn_{KK}}{dm}$ in Fig. (5.21).

We now point out that the values used in Ref. [193], as already observed in Ref. [191], are no longer compatible with the observed solar photon flux in the keV range measured by SPHINX. Furthermore, the number density adopted in Ref. [193], namely $n_{KK_{NEWSG}} = 4 \times 10^7 \mathrm{cm}^{-3}$, is not compatible with

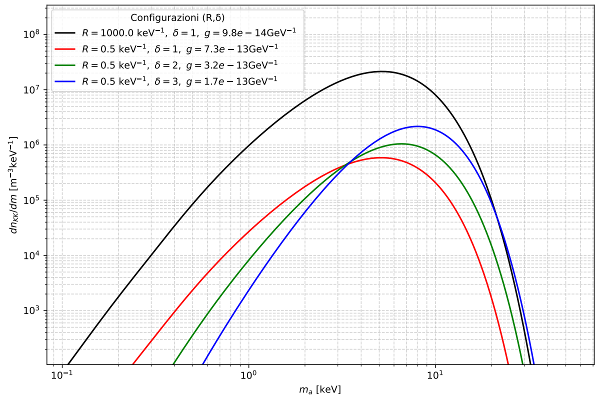

Figure 5.21: Differential number density as a function of the axion mass for the same parameter values (δ, R) used in Ref. [191].

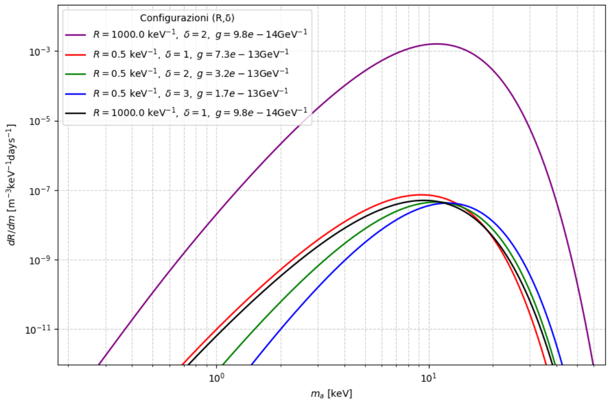

Figure 5.22: Differential rate as a function of the axion mass. The purple line represents the parameter values (δ, R) used in Ref. [104] and in Ref. [193]. The other lines represent the differential rate for parameter values currently compatible with the SPHINX data, as shown in Ref. [191]. The difference is remarkable.

our result, and is even less compatible with the value found in Ref. [191], differing by factors of 16 and 50, respectively.

The set of parameters and results used in Ref. [193] was first obtained in Ref. [104], as the parameter choice that, at the time, seemed most suitable to explain the coronal heating problem, while for the number density no explicit derivation was given. Moreover, as stated by the authors, that analysis was carried out without taking into account the effective photon mass inside the Sun. For this reason, and given that the results in Ref. [191] also differ significantly from those in Ref. [104], we believe that the results used in Ref. [193] should now be considered outdated.

We now present the event rate for different choices of the parameters of the Kaluza–Klein axion model as a function of the coupling $g_{a\gamma\gamma}$. We observe that the predicted event rate is very different between the benchmark parameters used in Ref. [193] and those compatible with the SPHINX observations.

Since the Kaluza–Klein axion model has at least two additional parameters, (δ, R), we present some more general results. In fact, at fixed number of extra dimensions, there is an upper bound on the compactification radius, as already discussed in Chapter (4). Since this parameter enters as a multiplicative factor $\sim R^\delta$ in Eq. (5.59), every result can in principle be rescaled accordingly. The edge of each shaded region corresponds to the upper bound $R_{\max}$ on the compactification radius for each specific value of δ. These limits are reported in the table below.

δ	$R_{\max}$ [keV^{-1}]	method
1	2.2×10^5	Pendulum
2	4.9×10^3	SN 1987A
3	6.2×10^{-1}	Colliders
4	3.1×10^{-3}	Colliders

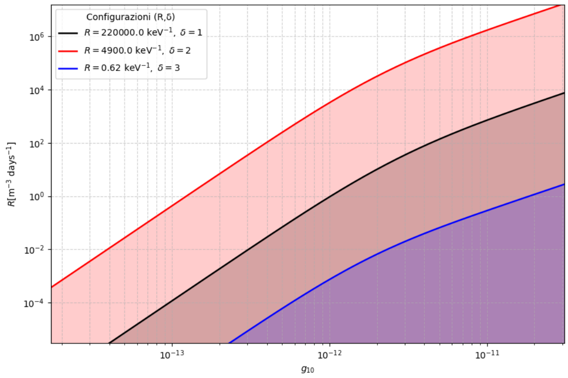

Figure 5.23: Event rate for the upper bound of R for each value of the number of extra dimensions.

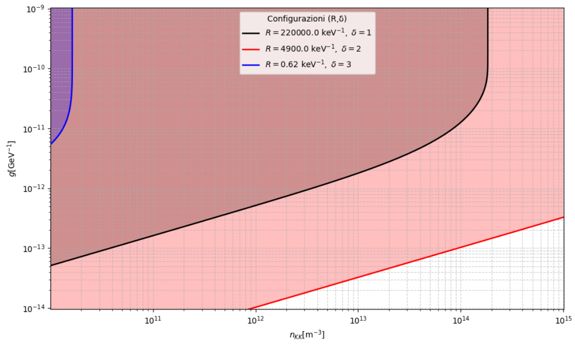

Figure 5.24: Number density of trapped KK axions for some choices of the parameters (δ, R).

For each value of δ, we present the most stringent bound. We exclude the bounds derived from neutron stars, since they do not apply if one assumes the existence of KK axions. These bounds can be found in Refs. [194, 179, 195]. The bounds always imply smaller values of $R_{\max}$ as δ increases. Since the overall factor scales as R^{δ}, we restrict our analysis to $\delta \leq 3$ in order to obtain phenomenologically relevant predictions.

Finally, in analogy with Fig. 9 of Ref. [193], we derive the number density as a function of the coupling $g_{a\gamma\gamma}$ for the maximum allowed compactification radius.

Chapter 6

Direct Detection of Solar Axions using a Gaseous Time Projection Chamber

In this final chapter, we use the results of the previous computations to derive sensitivity projections for the ALP parameter space using LIME, presented in Ref. [196], a detector operated at the INFN Gran Sasso National Laboratories (LNGS) within the CYGNO collaboration, see also Ref. [197]. We first describe the experimental setup, focusing on the detector geometry and gas properties relevant to the present analysis. We then discuss the data-analysis strategy and finally derive a sensitivity projection for the ALP parameter space.

6.1 LIME

The Long Imaging ModulE (LIME) is the detector considered in this analysis. LIME is a prototype of the larger CYGNO experiment, whose long-term goal is the construction of a large optically readout gaseous Time Projection Chamber with an active volume of about 100 m^3, to perform directional dark matter direct-detection searches. The LIME prototype was designed to validate Monte Carlo simulations, study background sources, and test the optical readout technology under realistic underground conditions.

Schematically, LIME consists of a box-shaped Time Projection Chamber (TPC) with a drift length of 50 cm and a transverse area of 33×33 cm^2. The active volume is filled with a gas mixture of He (60%) and CF_4 (40%), operated at a small overpressure with respect to atmospheric pressure. The field cage is made of 34 copper rings and is terminated by a 0.5 mm-thick copper cathode. Each ring is 10 mm thick, with a 4 mm spacing between adjacent rings.

At the end of the drift volume, the ionization electrons are amplified by a triple stack of Gas Electron Multipliers (GEMs). A GEM consists of a thin Kapton foil, approximately 50 μm thick, coated on both sides with thin copper layers approximately 5 μm thick. The foil is chemically perforated with a high density of microscopic holes. When a voltage difference is applied between the two copper electrodes, a strong electric field is generated inside the holes. The primary electrons produced in the gas drift towards the GEMs, enter the holes, and are multiplied through an avalanche process.

Gas Mixture

A TPC is a detector in which ionization charge drifts in a gas or liquid under an electric field, allowing three-dimensional track reconstruction. When a particle crosses the active volume, it may deposit part of its energy in the medium, producing ionization electrons. Under the action of the electric field, these electrons drift towards the amplification region, where they are multiplied through an avalanche process. The resulting charge and light signals can then be used to reconstruct the event topology and deposited energy.

LIME uses a gas mixture composed of 60% helium and 40% tetrafluoromethane, CF_4, in molar fraction. This mixture was chosen by the CYGNO collaboration as a compromise between detector stability, charge transport properties, and light yield. Some relevant transport properties of the gas mixture are shown in Fig. (6.1) and Fig. (6.2). In particular, Fig. (6.1) shows the transverse and longitudinal diffusion coefficients as a function of the drift field, while Fig. (6.2) shows the electron drift velocity as a function of the electric drift field. These quantities were evaluated with Garfield in Ref. [198], with an estimated accuracy of about 10%.

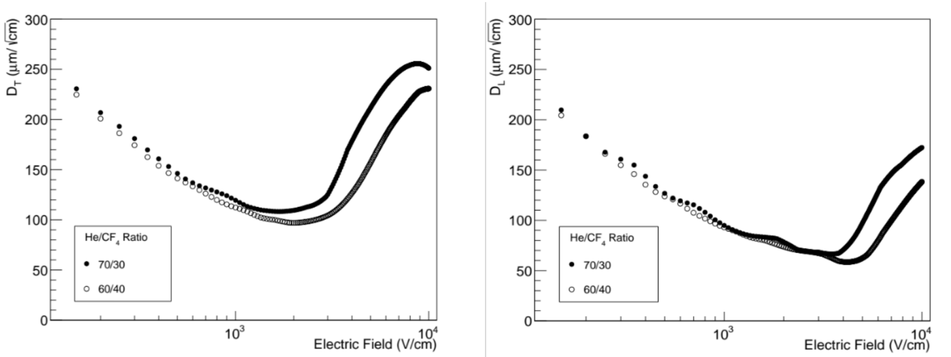

Figure 6.1: Transverse and longitudinal diffusion coefficients as a function of the drift field, computed with Garfield. The estimated accuracy of the simulation is about 10%. Figure adapted from Ref. [198].

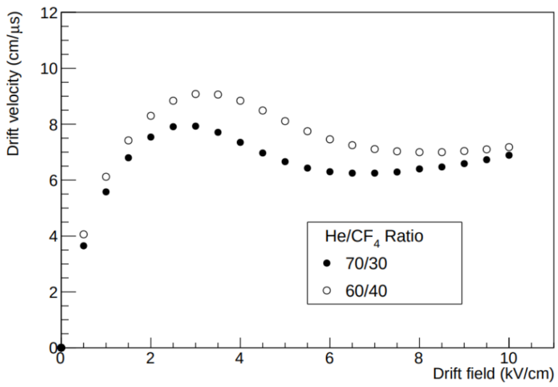

Figure 6.2: Electron drift velocity as a function of the drift field, computed with Garfield. The estimated accuracy of the simulation is about 10%. Figure adapted from Ref. [198].

For the present analysis, an important quantity is the photon absorption coefficient of the gas mixture, μ_{mix}, or equivalently the absorption length

$$\lambda(E) = \frac{1}{\mu_{\text{mix}}(E)}. \tag{6.1}$$

The absorption coefficient of the mixture is computed from the absorption coefficient of each individual component through:

$$\frac{\mu_{\text{mix}}}{\rho_{\text{mix}}} = \sum_i w_i \left(\frac{\mu}{\rho}\right)_i, \tag{6.2}$$

where w_i and ρ_i are respectively the mass fraction and the density of the i-th component. The mass fractions are obtained from the molar fractions according to

$$w_i = \frac{x_i M_i}{\sum_j x_j M_j}, \qquad x_i = \frac{p_i}{p}, \tag{6.3}$$

where M_i is the molar mass of the species, x_i is its molar fraction, p_i is its partial pressure, and p is the total pressure of the gas mixture.

The absorption length obtained with this procedure for the gas mixture used in LIME is shown in Fig. (6.3). This quantity determines the probability that a photon produced inside the active volume is absorbed before leaving the detector.
Given the type of event in which we are interested in, that are axions decaying in two coincident, back to back photons with the same energy, the absorption length is therefore a central ingredient in the estimate of the expected number of detectable events.

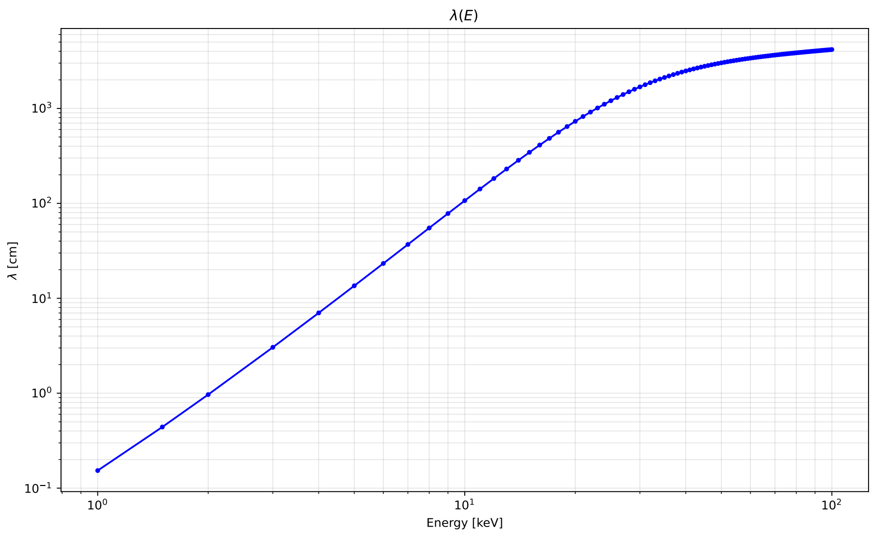

Figure 6.3: Photon attenuation length $\lambda(E)$ in the LIME gas mixture as a function of the photon energy.

Readout System

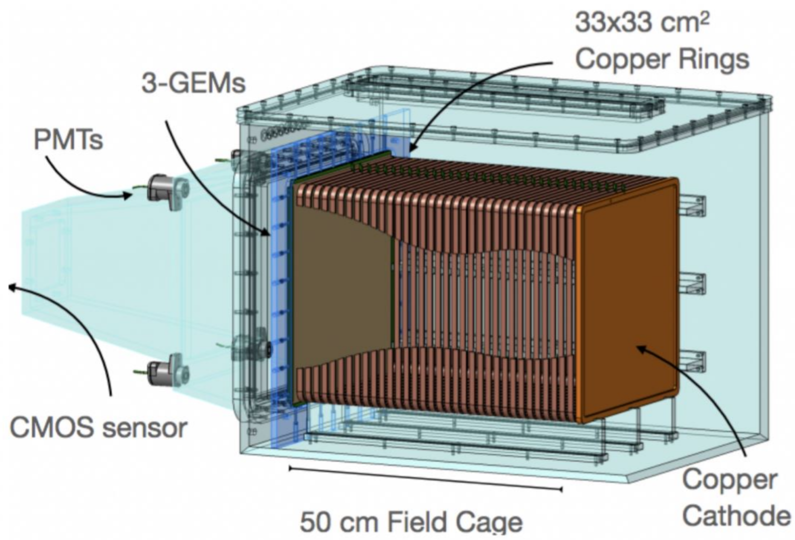

Figure 6.4: A schematic view of the LIME detector: we can see the field cage, the stack of the 3 GEMs, the PMTs and the CMOS camera.

The peculiarity of the CYGNO experimental approach is the so-called "optical readout" of the TPC. In fact, during the amplification process, scintillation light (visible and UV) is produced thanks to the CF4 component of the gas mixture. The capability to read-out the light by means of optical sensors instead of the charge makes the CYGNO detectors scalable to the large areas and volumes needed in the dark matter direct detection field. In LIME we have both optical read-out, consisting of a scientific CMOS camera (providing the 2D information of the track) with an exposure that can be set from 30 ms to 10 s and a set of four PhotoMultiplier Tubes (PMT, providing the z-coordinate information of the track) installed in a square around the CMOS sensor, at 19 cm from the last GEM, see Fig. (6.4). For our analysis only the PMT signals, providing the z evolution of the track, were considered. An example of the typical LIME dataset is shown in Fig. (6.5).

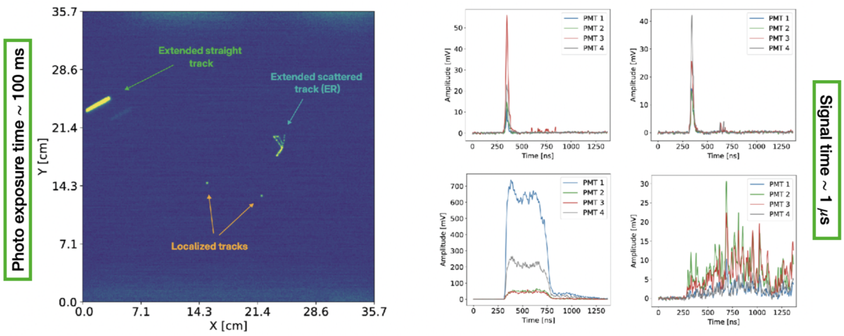

Figure 6.5: Example of a typical LIME dataset.

A Trigger and Data Acquisition system (TDAQ) has been designed for LIME to synchronize the readout of the camera with the signals recorded by the four PMTs, see Ref. [199]. In a typical running configuration, the camera is operated in continuous exposure mode with an exposure time of about 300 ms. The sensor rows are activated sequentially, starting from the top of the camera. The full activation of the sensor takes approximately 180 ms.
During the time interval in which all rows are simultaneously active, the camera is said to be in the Global Exposure (GE) condition. After this interval, the rows are progressively deactivated, again in a sequential way. Therefore, only events occurring during the GE window are guaranteed to be fully imaged by the camera.
During the full camera exposure window, the PMT signals are sent to a leading-edge discriminator. The trigger condition is defined by a MAJORITY-2 logic, requiring at least two PMTs to fire in coincidence, in particular it is also required that the camera is exposed and the digitizer it is ready to acquire. There are two possible boards for digitization: the first provided by CAEN V1742 with sampling rate of 750 MHz and 1024 samples stored, and the second provided by CAEN V1720E with a sampling rate of 250 MHz, and 4000 samples stored. We used the second board, since we are interested in studying signals within temporal windows up to 10 μs (which is, given the drift velocity, the length in time of the TPC).
The geometrical configuration of the PMTs with respect to the GEM plane is shown in Fig. (6.6).

We used PMTs to perform both temporal and energy analysis. Following Ref. [200], the light collected by the i-th PMT from a light-emitting event located at position (x_j, y_j) on the GEM plane can be described under the assumption of a perfectly diffuse source. If L_j denotes the light emitted per unit surface and per unit solid angle by the event j, the amount of light collected by the i-th PMT is

$$L_{ij} = \frac{L_j A_j A_i \cos^2\theta_{ij}}{R_{ij}^2}, \tag{6.4}$$

where A_j is the area of the emitting event, A_i is the effective area of the PMT, and R_{ij} is the distance between the source event and the PMT. The angle θ_{ij} is the angle between the normal to the GEM plane and the direction connecting the source event to the PMT. Since the PMTs are placed at a fixed distance $h = 19$ cm from the GEM plane, one has

$$\cos\theta_{ij} = \frac{h}{R_{ij}}. \tag{6.5}$$

With:

$$R_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2 + h^2} \tag{6.6}$$

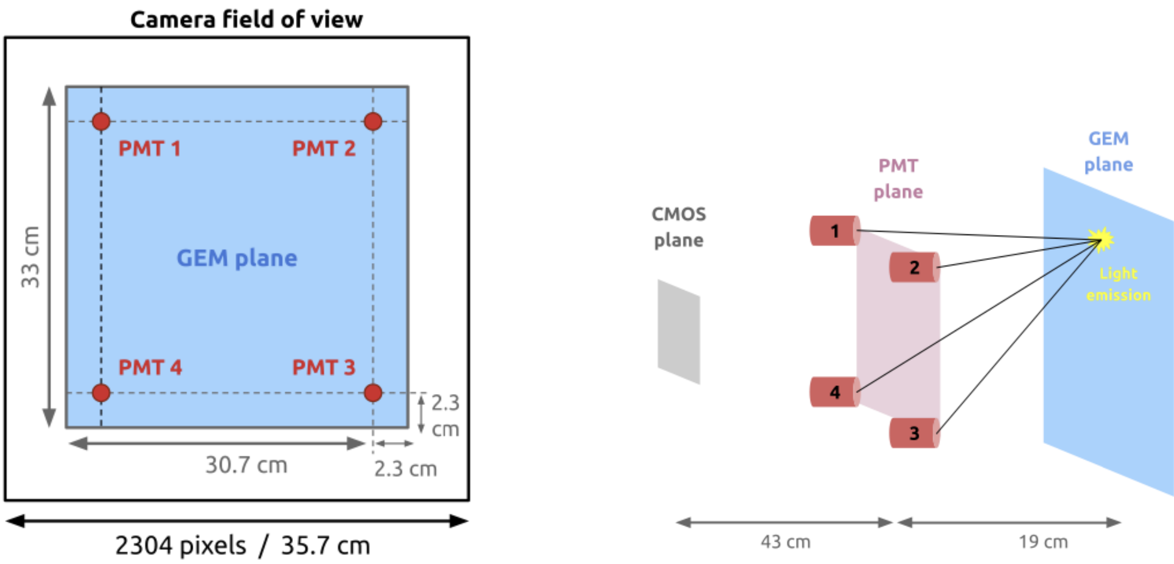

Figure 6.6: Relative disposition of the sensors with respect to the GEM plane, where light is emitted. Left: side view showing the field cage and the vertical distance between the PMTs and the GEMs. Right: front view showing the camera position at the center and the four PMTs at the corners. Figure taken from Ref. [200].

Therefore, up to geometrical and normalization factors, the collected light scales as

$$L_{ij} \propto \frac{L_j}{R_{ij}^4}. \tag{6.7}$$

The signal recorded by the i-th PMT is a voltage waveform $V_{ij}(t)$. The corresponding integrated charge is obtained as

$$Q_{ij} = \frac{1}{R} \int_{\Delta t} V_{ij}(t), dt \propto L_{ij}, \tag{6.8}$$

where $R = 50\ \Omega$ is the input impedance and Δt is the integration time window. The PMT response is affected by stochastic fluctuations, which become less relevant at high deposited energy, and by a systematic uncertainty of about 5%. The relative noise was estimated using a ^{55}Fe radioactive source, whose dominant K_α at 5.9 keV together with a weaker K_β X-ray lines produce an effectively monochromatic peak at 5.9 keV. At this energy, the relative noise was found to be approximately 15%.
In the following section, we will present the calibration performed using the same radioactive source.

6.2 Data Analysis

We analyzed the data taken by LIME during RUN5, looking for axion decays. As anticipated, the specific signature of axion decays we are looking for is the following: two coincident, back to back photons, with the same energy. In fact as shown in chapter (5) the axions that accumulated around the solar system are the ones that dominate on earth, therefore they are non-relativistic, meaning they decay in their rest frame. This is a very specific feature, that allows us to reach a good sensitivity even with a prototype as small as LIME. To support this expectation at a preliminary level, we present the probability that both photons are completely absorbed inside LIME as a function of photon energy.

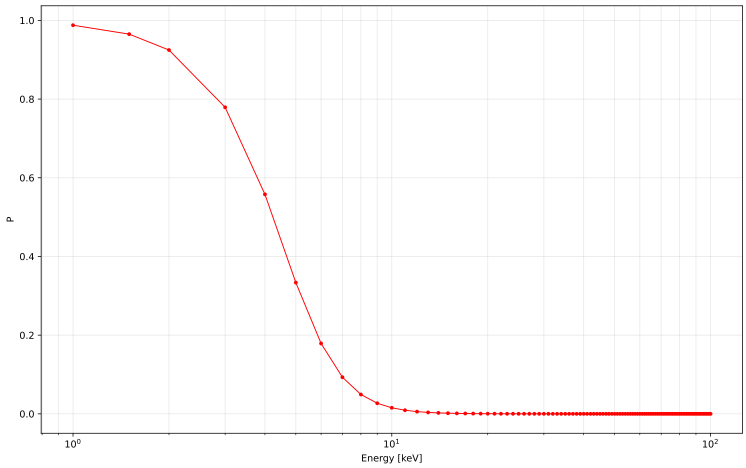

Figure 6.7: Probability for both photons, produced back to back and with the same energy to be absorbed inside LIME as a function of the photon energy. We see that we are sensitive to energies below the ~ 10 keV, corresponding to axions with masses below 20 keV.

The data taking lifetime of RUN5 is 1.363×10^7 s or about 157.8 days. Considering the active volume of LIME, this corresponds to an exposure of about 0.0235 m^3·years. To process all these data, we built a waveform–clustering algorithm and applied it to all RUN5 data. Given a 4 PMT waveform, the code first smooths each channel with a centered moving average to suppress high–frequency noise that could make the noise in a PMT signal oscillate around the threshold. We then apply a negative threshold to the smoothed signals (the waveform signals are negative): a sample is considered active for a given PMT when the smoothed value falls below a fixed threshold. A cluster starts as soon as at least two PMTs are simultaneously active (coincidence multiplicity ≥ 2); the start is shifted backward by a configurable number of pre–trigger samples to retain context. A cluster ends when no PMT remains active, i.e. when all PMTs return above the negative threshold. To avoid artificial fragmentation, consecutive clusters are merged whenever the start of the next cluster occurs within a minimum time gap t_{min} from the end of the previous one; this prevents splitting signals that belong to the same physical trace (e.g. an electron recoil ER) into multiple clusters see Fig. (6.8).

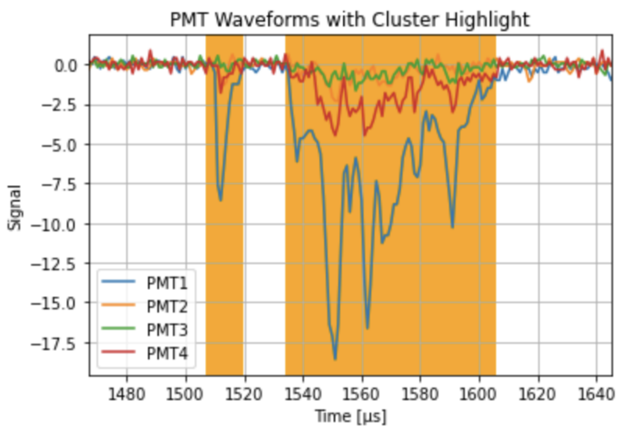

Figure 6.8: Example of ER, the shaded regions are the ones recognized as clusters by the code.

For each reconstructed cluster, the algorithm stores the start and end times, together with the corresponding temporal width. It also computes the integral of the waveform in each PMT over the cluster time window, as well as the total integral obtained by summing the four PMT contributions. In addition, for each PMT, the algorithm records the minimum waveform value within the cluster window, corresponding to the pulse amplitude for negative-going signals, and the sample index at

which this minimum occurs.
The total number of clusters identified in a given waveform is then simply given by the length of the reconstructed cluster list. This information is used in the following analysis to characterize the PMT activity associated with each recorded event and to select waveform structures compatible with the expected signal topology.
We now describe the selection cuts. They were designed to reject the vast majority of background events. These cuts are intended to select clusters compatible with the expected axion signature. In particular, below approximately 20 keV, photons are fully absorbed through the photoelectric effect, and the resulting photoelectron tracks are short compared to diffusion. Therefore, the signals associated with the absorption of these photons are expected to be short in time and to have a longitudinal width comparable to that of photoelectrons produced by the 5.9 keV X-rays from ^{55}Fe. Then we present the efficiency of these cuts on the signal, that is the percent of signal that survives after the employed cuts. Starting from 16016226 waveforms, the final selection yields 97 waveforms ($\approx 6.1 \times 10^{-4}\,\%$). A "good" cluster passes our cuts if:

- **Width window.** Its width lies between 20 and 33 time samples. We recall that we considered the data recorded by the V1720E Digitizer, for which 1 sample = 4 ns. This choice is motivated by the reference signal from the source ^{55}Fe placed at $z = 25$ cm above the GEMs plane, whose width distribution has mean 33.8 time samples and standard deviation 2.8 (see Fig. (6.9)). However we will see more quantivatively how this choice is reasonable when we will study the efficiency of this cut on the signal. This window effectively suppresses noise while retaining signal. For comparison, a looser window of 20 to 45 time samples would keep 1492 waveforms ($\approx 9.32 \times 10^{-3}\,\%$), but many of those do not match the expected two–photon topology. An example of waveform that is allowed if we take a width window $20 - 45$ is given in Fig. (6.11).

- **Total integral.** The sum of the four PMT integrals satisfies $\text{Integral}_{\text{tot}} \leq -200$ A.U. For the ^{55}Fe source in the same disposition we observe a mean of -310 A.U. with a standard deviation of 42 A.U. (see Fig. (6.10)). Again we will verify in the following the efficiency of this cut on the signal.

For each waveform, we also compute the number of not good clusters and the sum of their widths. At the event level, to select for events compatible with the axion two-photon decay signature, i.e. with two "good" clusters, we require:

- at least two good clusters;

- at most three bad clusters;

- sum of bad–cluster widths ≤ 1000 time samples.

At each iteration, we visually inspected the waveforms passing the cuts to verify qualitative improvement of the selection. Some examples of waveforms that passed all the cuts are presented in Fig. (6.12) and Fig. (6.13). Having so few remaining waveforms at this point to improve the analysis human inspection can remove clusters with not desired features.

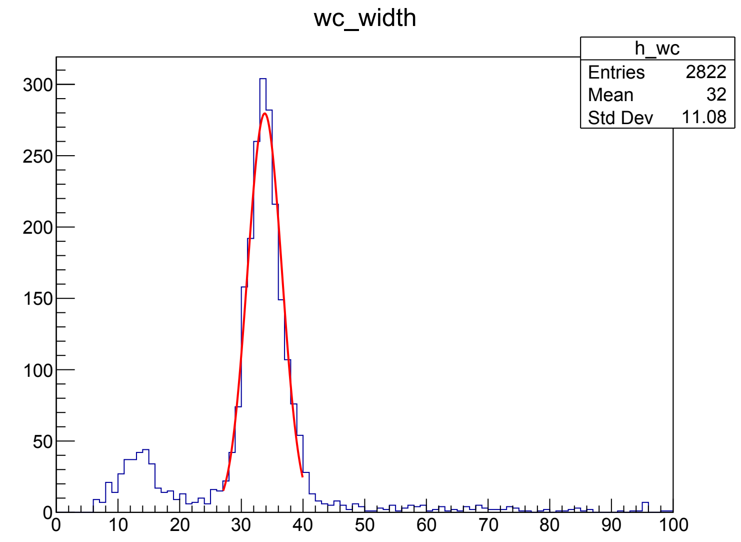

Figure 6.9: Width distribution for the ^{55}Fe reference sample.

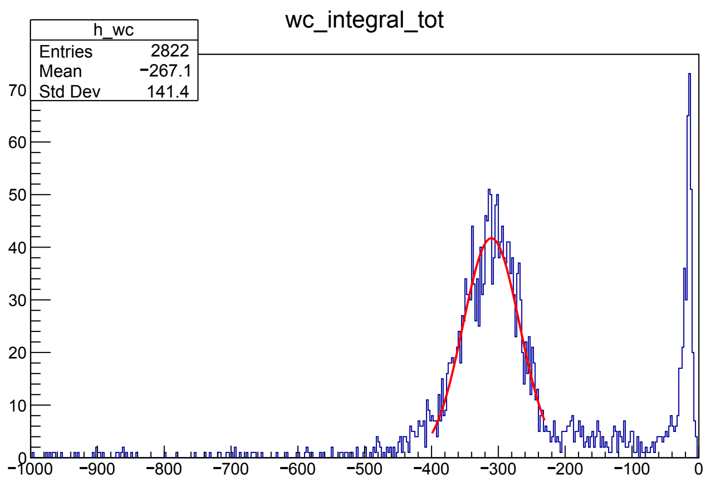

Figure 6.10: Total integral distribution for the ^{55}Fe reference sample.

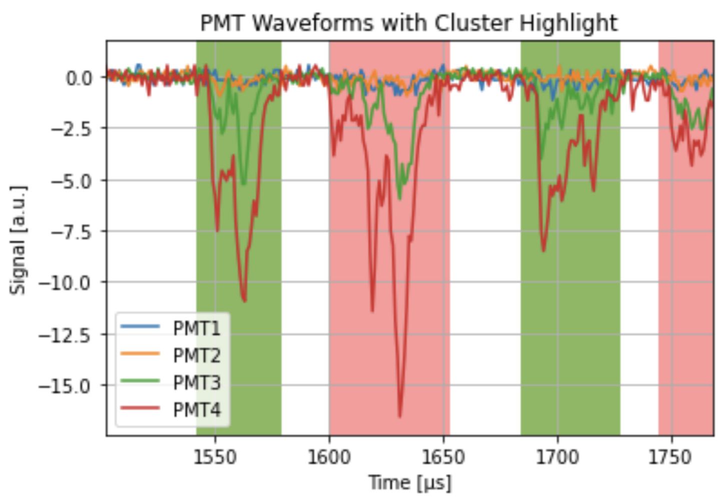

Figure 6.11: Example of a cluster type rejected by tightening the width window. Red and green shaded regions denote "bad" and "good" clusters, respectively.

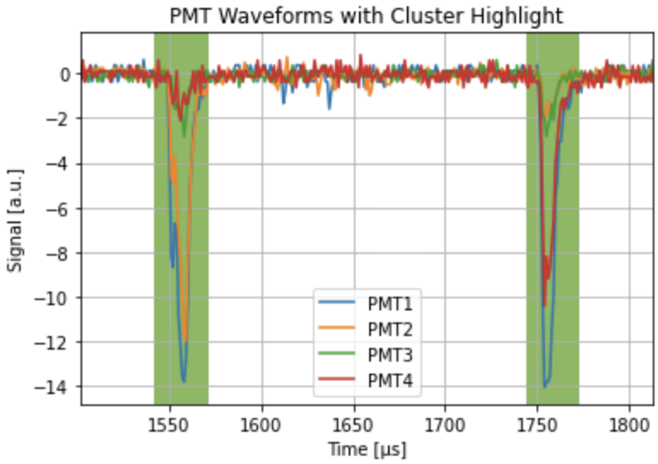

Figure 6.12: Example showing a non-monotonic behaviour.

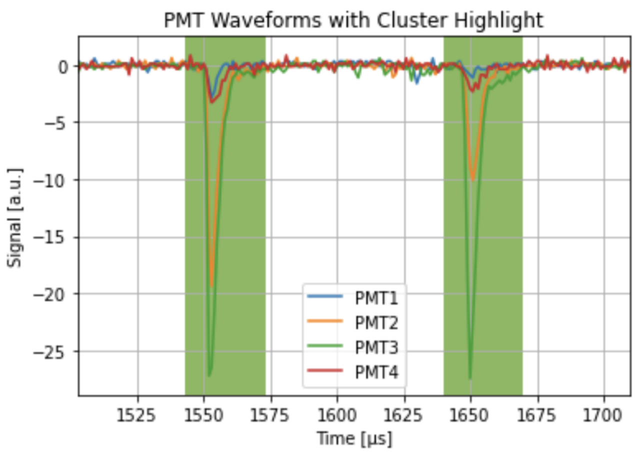

Figure 6.13: Example with monotonic cluster behaviour.

6.2.1 Efficiency

We now study quantitatively how the signal is affected by the cuts that we applied. We start from the cuts on the width of the clusters. We need first to understand how the width of a signal is determined. In fact, due to diffusion, we expect the electrons produced by primary ionization to drift in random directions, resulting in a diffuse signal. Following the general rule for a diffusive process the expectation is that the width, that we now indicate as σ follows the rule:

$$\sigma(z) = \sqrt{\sigma_0^2 + f_T^2 z} \tag{6.9}$$

To determine the parameters σ_0 and f_T practically one uses the same monoenergetic radioactive source (in our case ^{55}Fe) placing it at different positions z and studies how the signal spreads. Below the results of this procedure:

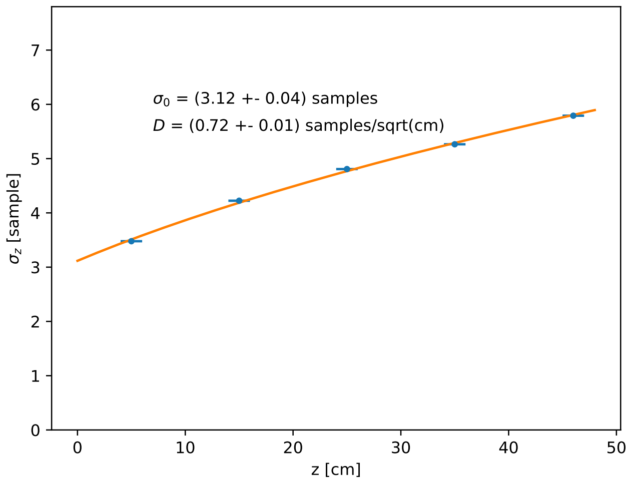

Figure 6.14: Width as a function of the source's distance from the GEMs plane.

To pass from samples to quantities expressed in physical units one needs one final ingredient. First of all we use that 1 sample = 4 μs, then we use the value of the drift velocity. With a drift field $E = 0.5$ kV/cm it is $v_{drift} = 4.4$ cm/μs and from that we get $\sigma_0 = 562 \pm 13$ μm and $f_T = 124.9 \pm 2.4$ μm/sqrt(cm). The parameter f_T (that is often denoted also as D) depends on the drift field too. For reference, Fig.(6.15) shows how drift velocity and longitudinal diffusion coefficient depend on the drift field according to Garfield simulation.

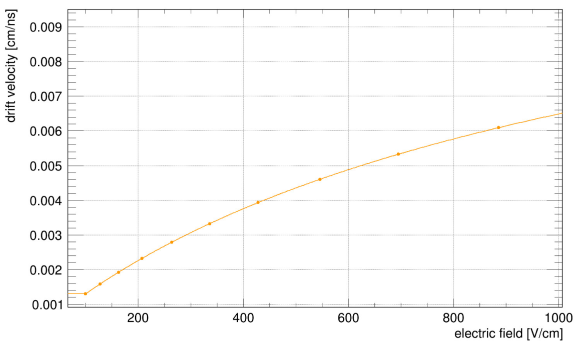

((a)) Drift velocity as a function of the drift field.

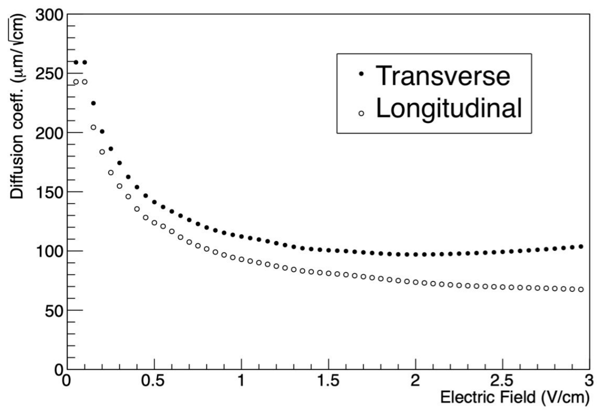

((b)) For our setup, $E = 0.5$ V/cm, we expect the diffusion coefficient to be $f_T \simeq 125$ μm/$\sqrt{\text{cm}}$.

Figure 6.15: Drift velocity and transverse diffusion coefficient as functions of the drift field.

We built a MC simulation of our signal and for each event (at fixed energy) where both photons are absorbed we studied how many pass the cut. In Fig. (6.16) we show an histogram of the width of the signal for different energies, while in Fig. (6.17) we show the ratio of events that pass our cuts over the number of events where both photons are absorbed in the detector. We just state without showing that, if the energy ratio were computed for each photon pair within the same event, all events would satisfy the cut.

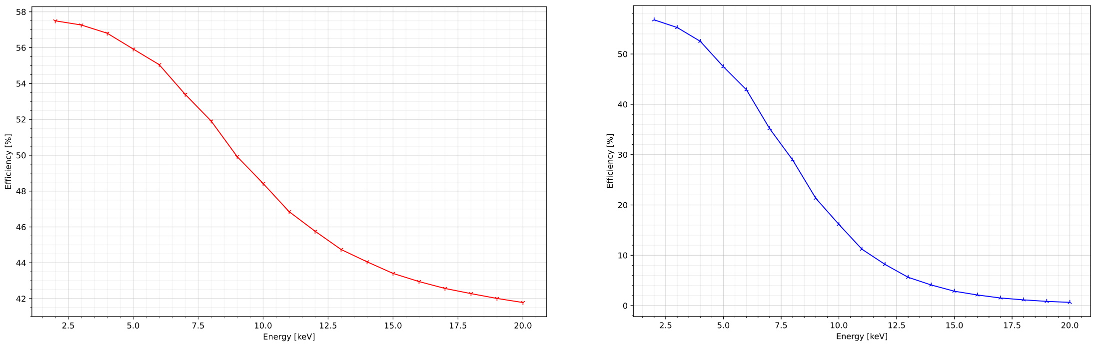

((a)) *efficiency over the events that are absorbed inside the detector*

((b)) *efficiency over all events.*

Figure 6.17: Efficiency of the cut on width as a function of the energy. As expected, no strong energy dependence is observed, since the profiles of the histograms in Fig. (6.16) are approximately the same. Moreover, a large fraction of signal events passes this cut for all the energies considered, while requiring also that both photons are absorbed again we see that we are sensitive to energies below 20 keV.

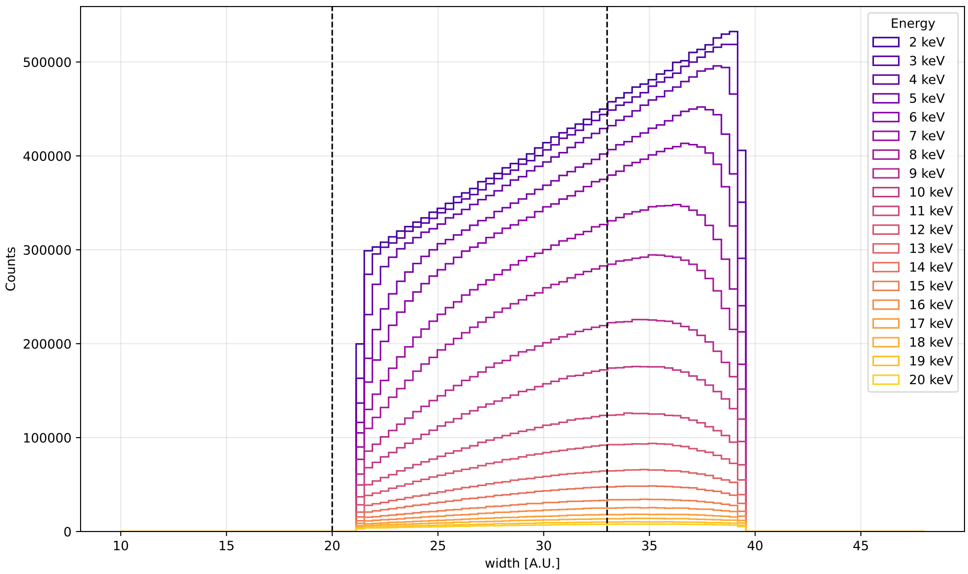

Figure 6.16: Histogram of the widths for the events simulated. For each energy we simulated 10^7 events, so that we can already have an idea of the efficiency at fixed number of events. The dotted lines are the cuts that we set.

We now consider the efficiency associated with the energy cut. We started determining C_i, the calibration constants introduced in Eq. (6.1), comparing the data of the ^{55}Fe of RUN5 at $z = 25$ cm against our simulation. To do so we first simulated the signal of the ^{55}Fe assuming an elliptical cone distribution for our source with vertex outside the chamber, at $O = (L_x/2,\ -11.5\text{ cm},\ L_z/2)$ and with angles on the xy plane fixed equal to 22° and on the yz plane equal to 12°, matching the source collimator aperture. In Fig. (6.18) the projection on every cartesian plane of the events simulated.

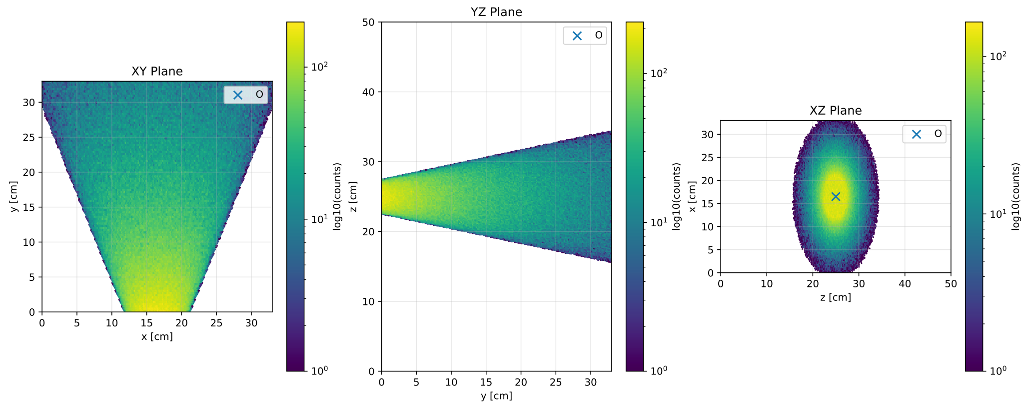

Figure 6.18: Number of events on each xyz plane. We expect the PMT with lower y coordinate, placed close to the ^{55}Fe source to have a larger signal than the upper PMTs.

We then assumed a fixed energy of 5.9 keV and used the corresponding attenuation length to estimate the energy measured by each PMT. A Gaussian smearing with $\sigma = 0.135$ is also applied to account for the finite energy resolution of the detector. This value is consistent with the final energy profile obtained from the total integral. The simulated distributions are shown below along with the real data in Figs (6.19). Note that given the geometrical disposition of the cone we have a pair PMT near the vertex of the cone and a pair that are instead more distant, for each pair of this type we expect the same behavior. After calibration we can compute the total energy that is just the sum of each PMT's contribution. Again, we present in Tab. (6.1) the calibration results.

Table 6.1: Calibration results for the four PMTs.

	μ [A.U.]	σ [A.U.]	Gain
PMT1	-19.0 ± 2.2	13.7 ± 2.4	710 ± 80
PMT2	-23.1 ± 1.4	14.1 ± 2.3	860 ± 60
PMT3	-34.4 ± 1.3	21.1 ± 1.8	1590 ± 60
PMT4	-28.5 ± 0.6	9.8 ± 0.8	1330 ± 30

Having found the calibration constants, we can compare the simulated total energy with the real data of our ^{55}Fe source. In Fig. (6.20) our results.

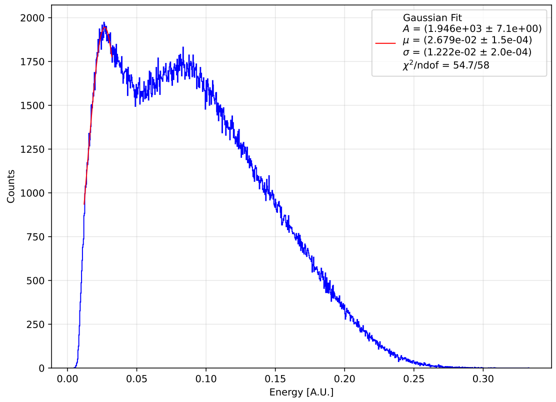

((a)) Simulation of the PMT closer to the ^{55}Fe source. Note that there is an implicit minus sign in the calibration constant. The red line corresponds to the Gaussian fit on the peaks that we identified to be present also in the figures below, and which allows us to find the calibration constants.

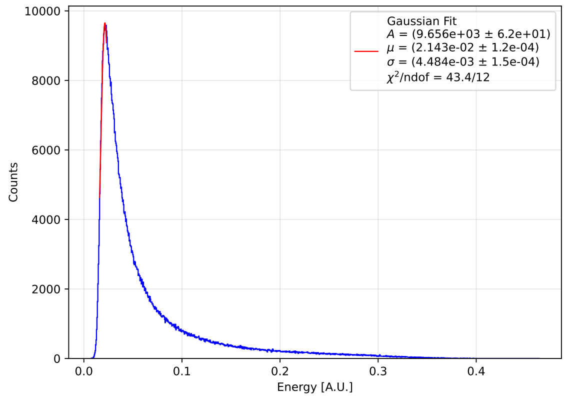

((b)) Simulation of the PMT more distant to the ^{55}Fe source. Note that there is an implicit minus sign in the calibration constant. The red line corresponds to the Gaussian fit on the peaks that we identified to be present also in the figures below, and which allows us to find the calibration constants.

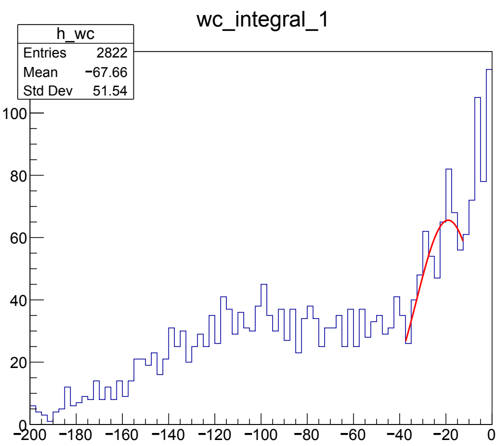

((c)) Real data of one of the closest PMTs, again the red line corresponds to the Gaussian fit that we identified above. $x-$axis is in arbitrary units, while $y-$axis in counts.

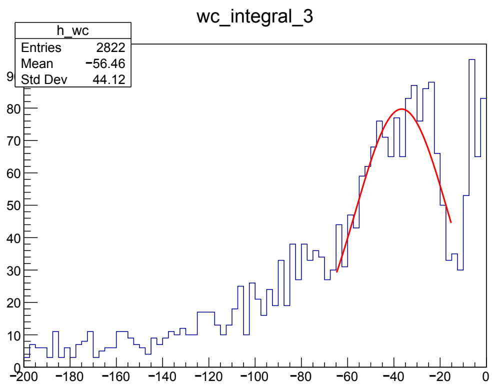

((d)) Real data of one of the more distant PMTs, again the red line corresponds to the Gaussian fit that we identified above. $x-$axis is in arbitrary units, while $y-$axis in counts.

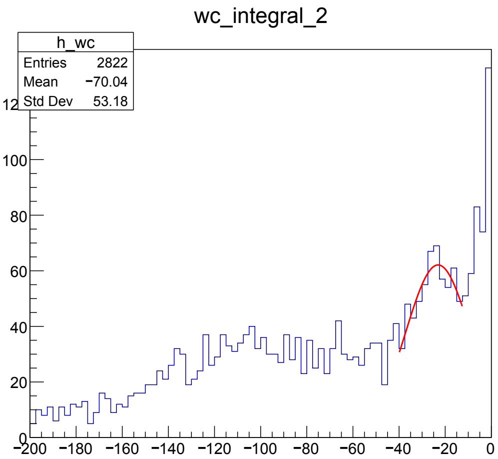

((e)) Real data of one of the closest PMTs, again the red line corresponds to the Gaussian fit that we identified above. $x-$axis is in arbitrary units, while $y-$axis in counts.

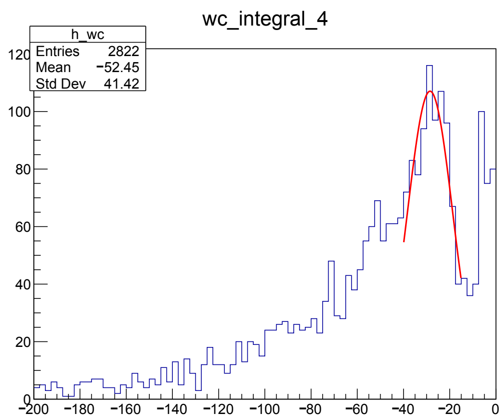

((f)) Real data of one of the more distant PMTs, again the red line corresponds to the Gaussian fit that we identified above. $x-$axis is in arbitrary units, while $y-$axis in counts.

Figure 6.19: Comparison between our simulations and real data of the ^{55}Fe source.

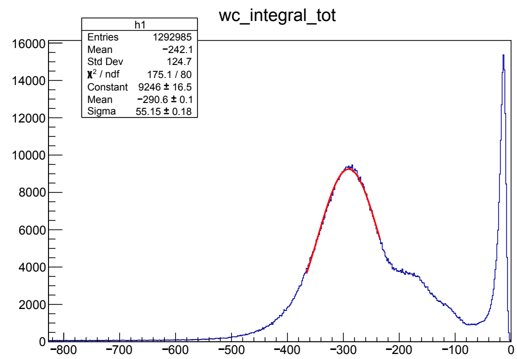

((a)) Real data of the ^{55}Fe source. Again in red the Gaussian fit, with mean at -290 A.U. and standard deviation $\sigma = 55$ A.U..

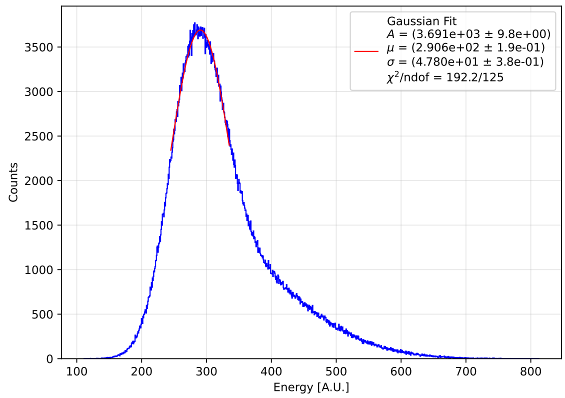

((b)) Simulation of the ^{55}Fe source. Again in red the Gaussian fit, with mean at 290 A.U. and standard deviation $\sigma = 58$ A.U..

Figure 6.20: Comparison between real data and simulation of ^{55}Fe. An important remark is that the units used in both the plot are indeed the same. The results of our simulation are compatible to the real data within our uncertanties.

. After having calibrated each PMT, and getting good compatibility between the simulated total energy and the real data, we move to the simulation of axion decays. As discussed above, the energy of the event placed at position (x_j, y_j, z_j) with real energy E inside the detector is given by:

$$E_{\text{simulated}} = \sum_i^4 \frac{C_i(E)}{\left((x_i - x_j)^2 + (y_i - y_j)^2 + h^2\right)^2} \tag{6.10}$$

Where the $i-$index refers to the $i-$th PMT, and for the calibration constant we have assumed linearity on the response of our detector:

$$C_i(E) = C_i(E^*)\frac{E}{E^*} \tag{6.11}$$

Here we used $E^* = 5.9$ keV, And $C_i(E^* = 5.9$ keV$)$ of course are the constants presented in Tab. (6.1). We then again added a smearing Gaussian factor, with its standard deviation σ depending on energy as seen in Ref. [201]:

$$\sigma(E) = 0.11 + 0.08\sqrt{\frac{E^*}{E}} \tag{6.12}$$

Again we used $E^* = 5.9$ keV. After having simulated our signal, we have computed the efficiency of the cut that we set as a function of the axion mass (or equivalently the photon's energy) we present in Fig. (6.21) our results, computed both over events where both photons are absorbed and over all events.

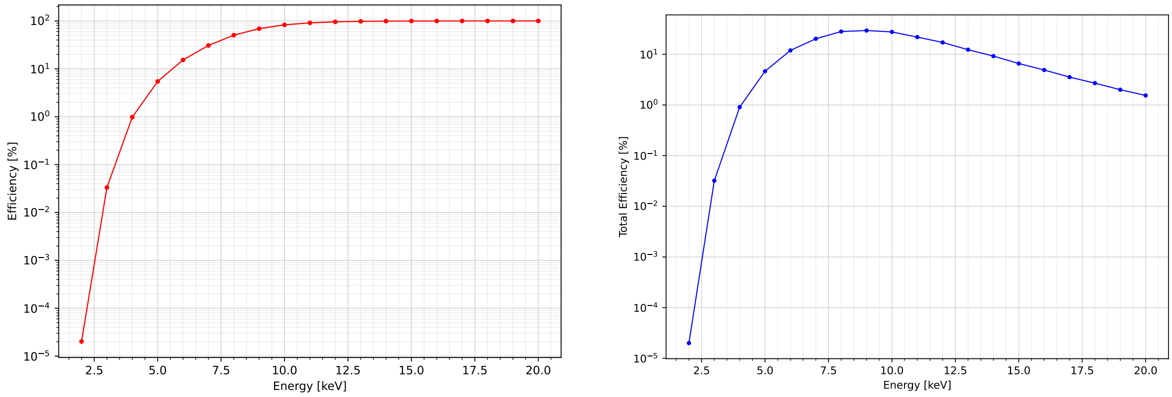

((c)) *efficiency over the events that are absorbed inside the detector*

((d)) *efficiency over all events.*

Figure 6.21: Efficiency of the cut on energy for different photon energies. We see that we are more sensitive at higher energies.

To get an idea of the type of signal that we expect, below an histogram of the energies of the events simulated for different axion masses.

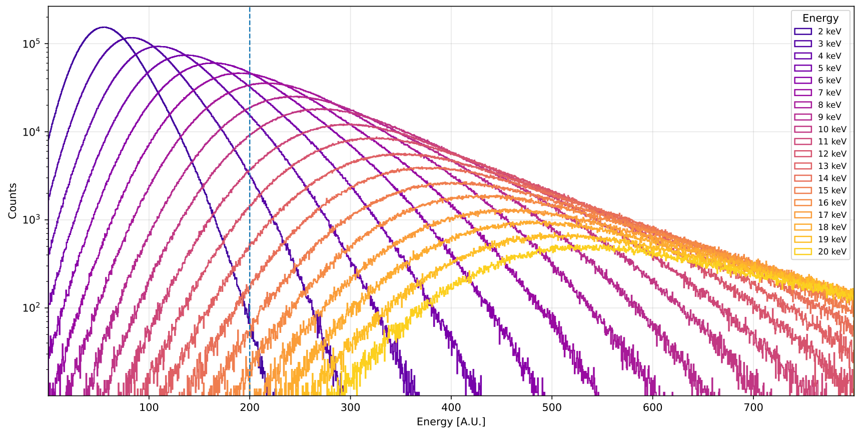

Figure 6.22: Histogram of the energies of the events simulated. For each energy we simulated 10^7 events, so that we can already have an idea of the efficiency at fixed number of events. The dotted line at $x = 200$ is the cut that we set.

6.2.2 Bayesian Sensitivity Projection for ALP Model and KK Axions

We are now ready to evaluate the total efficiency of the whole analysis. This quantity includes both the efficiency of the applied selection cuts and the probability that the two photons produced in the axion decay are absorbed in the detector. Therefore, the total efficiency, that represents the probability that an axion decay event is detected and survives the full selection procedure is given by:

$$\epsilon_{\text{tot}}(m_a) = \epsilon_{\text{cuts}}(m_a) \times P^{(2)}_{\text{abs}}(m_a), \tag{6.13}$$

where $\epsilon_{\text{cuts}}(m_a)$ is the combined efficiency of the analysis cuts and $P^{(2)}_{\text{abs}}(m_a)$ is the probability that both photons are absorbed in the detector.

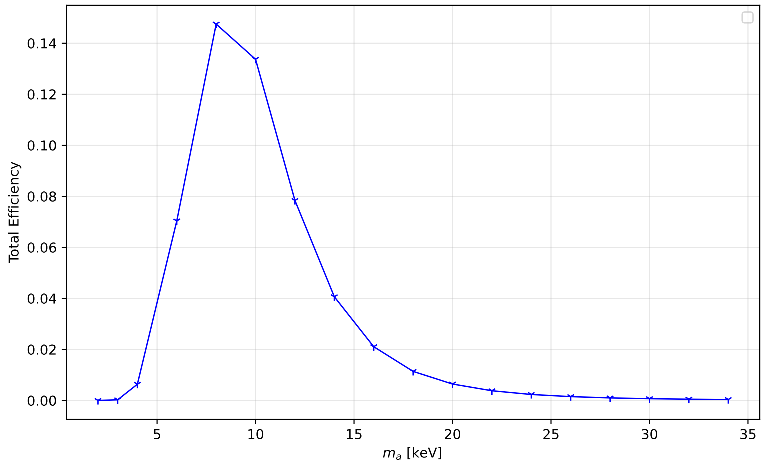

Figure 6.23: Total efficiency as a function of the axion mass. The efficiency includes the effect of the selection cuts and the probability that both photons are absorbed in the detector. The efficiency is highest approximately in the interval $m_a \sim 5$–15 keV.

For each axion mass, the expected number of selected signal events is obtained by multiplying the predicted signal yield before the selection cuts by the total efficiency:

$$N_S(g_{a\gamma\gamma}, m_a) = N_0(g_{a\gamma\gamma}, m_a)\, \epsilon_{\text{tot}}(m_a), \tag{6.14}$$

where $N_0(g_{a\gamma\gamma}, m_a)$ is the expected signal yield before event selection, evaluated as a function of the coupling $g_{a\gamma\gamma}$ and of the axion mass m_a.
The purpose of the following analysis is to obtain a sensitivity projection. Therefore, the result should not be interpreted as an observed upper limit derived from experimental data. Instead, we assume a background-only dataset, in which all the selected events are attributed to background processes.
Therefore the observed number of events and the nominal background expectation are fixed to be the same, and in particular $x_{\text{obs}} = N_{\text{bkg}} = 97$.
For a fixed axion mass m_a, the only free parameter in the model is the photon–axion coupling $g_{a\gamma\gamma}$. The expected total number of events is therefore

$$\mu(g_{a\gamma\gamma}; m_a) = N_{\text{bkg}} + N_S(g_{a\gamma\gamma}, m_a). \tag{6.15}$$

We use a Bayesian approach to infer the posterior probability density of the coupling. Given the observed data x_{obs}, the axion mass m_a, and a set of assumptions I, Bayes' theorem can be written as

$$p(g_{a\gamma\gamma} \mid x_{\text{obs}}, m_a, I) = \frac{\mathcal{L}(x_{\text{obs}} \mid g_{a\gamma\gamma}, m_a, I)\, \pi(g_{a\gamma\gamma} \mid I)}{\mathcal{Z}(x_{\text{obs}} \mid m_a, I)}. \tag{6.16}$$

Here, $p(g_{a\gamma\gamma} \mid x_{\text{obs}}, m_a, I)$ is the posterior probability density of the coupling, $\mathcal{L}(x_{\text{obs}} \mid g_{a\gamma\gamma}, m_a, I)$ is the likelihood, $\pi(g_{a\gamma\gamma} \mid I)$ is the prior probability density, and $\mathcal{Z}(x_{\text{obs}} \mid m_a, I)$ is the Bayesian evidence. Since the evidence does not depend on the coupling, it acts only as a normalization factor for the posterior distribution.
The likelihood is described by a Poisson distribution:

$$\mathcal{L}(x_{\text{obs}} \mid g_{a\gamma\gamma}, m_a) = \mathcal{P}\left(x_{\text{obs}} \mid \mu(g_{a\gamma\gamma}; m_a)\right), \tag{6.17}$$

where

$$\mathcal{P}(x \mid \mu) = \frac{\mu^x}{x!} e^{-\mu}. \tag{6.18}$$

The prior probability density is assumed to be uniform in the physical coupling within the interval covered by the signal-rate calculation $g_{a\gamma\gamma} \in [10^{-11}\ \text{GeV}^{-1}, 10^{-6}\ \text{GeV}^{-1}]$.

The posterior distribution is sampled using the `emcee` Python package, which implements an affine-invariant Markov Chain Monte Carlo sampler. For each fixed axion mass, the parameter space is one-dimensional and contains only the coupling $g_{a\gamma\gamma}$. We use an ensemble of 32 walkers, initialized at random positions within the allowed prior range. We present in Fig. (6.24) shows representative trajectories sampled every 15 points, for one axion mass.

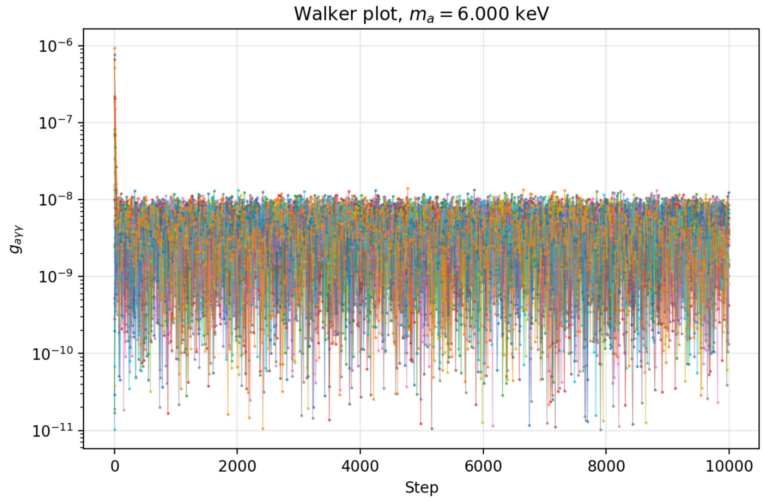

Figure 6.24: Example of walkers for $m_a = 6$ keV for 10^4 steps, we use a burn-in of 10^3 steps.

Each walker explores the parameter space according to the posterior probability density. After discarding an initial burn-in phase, the remaining samples represent draws from the posterior distribution

$$p(g_{a\gamma\gamma} \mid x_{\mathrm{obs}}, m_a, I). \tag{6.19}$$

The Bayesian credible upper limit at a credibility level q is defined as the value g_q for which the cumulative posterior probability is equal to q:

$$\int_{g_{\mathrm{min}}}^{g_q} p(g_{a\gamma\gamma} \mid x_{\mathrm{obs}}, m_a, I)\, dg_{a\gamma\gamma} = q. \tag{6.20}$$

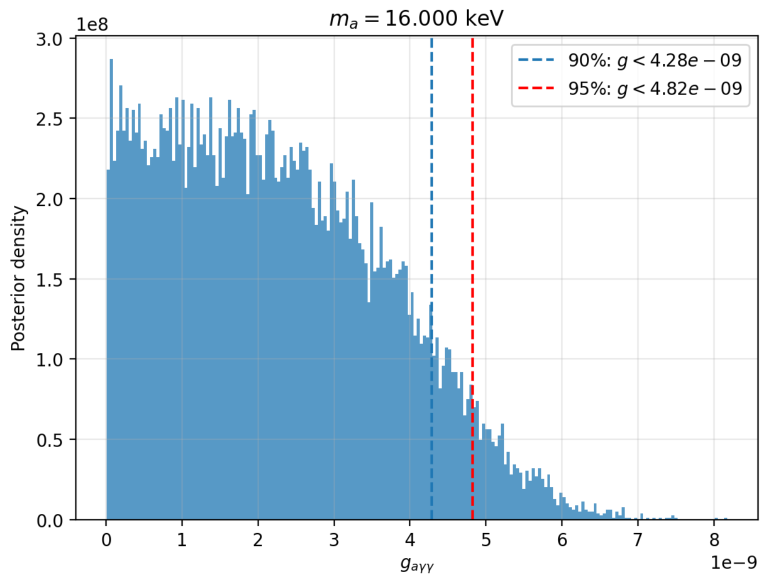

Figure 6.25: Posterior density distribution for $m_a = 16$ keV, the dashed line correspond to g_{90} and g_{95}.

For example, the 90% upper credible bound g_{90} is defined by

$$P\left(g_{a\gamma\gamma} \leq g_{90} \mid x_{\text{obs}}, m_a, I\right) = 0.90. \tag{6.21}$$

This means that, under the assumptions of the model, the likelihood, and the prior, 90% of the posterior probability lies below the value g_{90}. Equivalently, the posterior probability that the coupling is larger than g_{90} is 10%. We present in Fig. (6.25) an example of posterior that we got. Finally we plot g_{90} and g_{95} as a function of the axion mass, deriving a region of exclusion in the ALP parameter space.

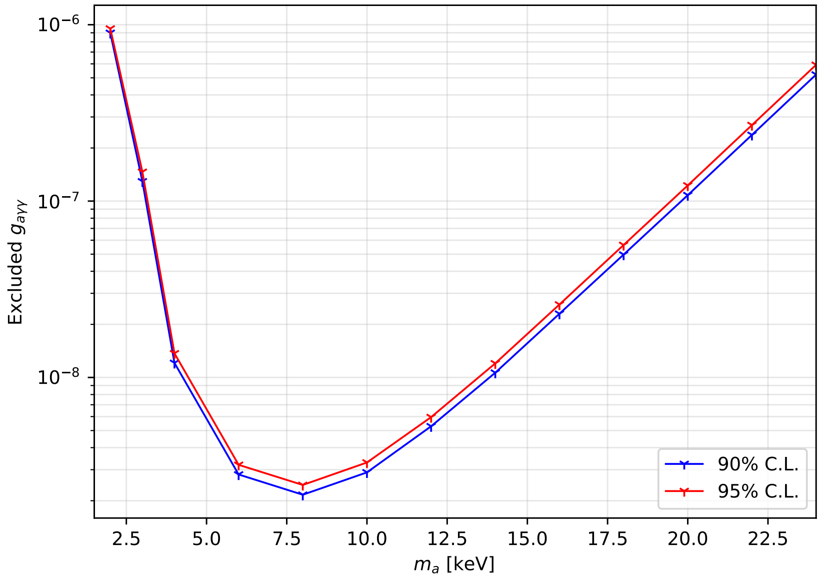

Figure 6.26: Exclusion region in the ALPs parameter space at 95 % (and 90 %) credibility level. We are in a region already excluded by indirect bounds.

Finally we present the same projection for CYGNO04. In this case the volume is eight times bigger, and therefore we have a double effect: the first is obviously that the exposure at even live time becomes 8 times bigger. The second effect instead is that also the probability for both photons to be absorbed in this detector increases, meaning that we are more sensitive at higher axion masses. We assumed the same efficiency for our cuts, and a signal 8 times the LIME's background signal. The results, along with LIME's results are presented below. As done at the end of chapter (5) we can now insert this

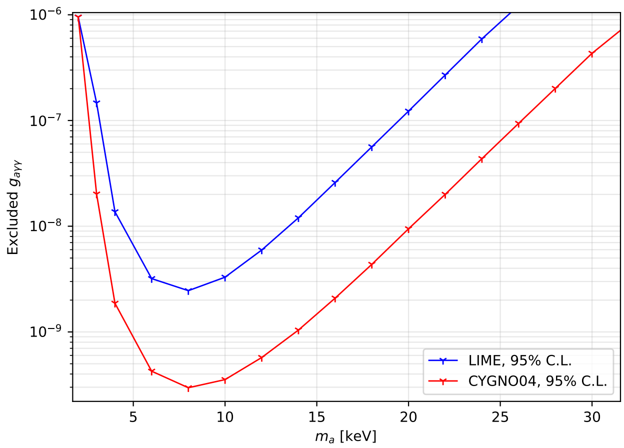

Figure 6.27: Comparison between LIME and CYGNO, we see that the difference is mainly at higher masses.

exclusion region in the ALP parameter space with the other bounds. We present below our results, as long as the CYGNO04 projection.

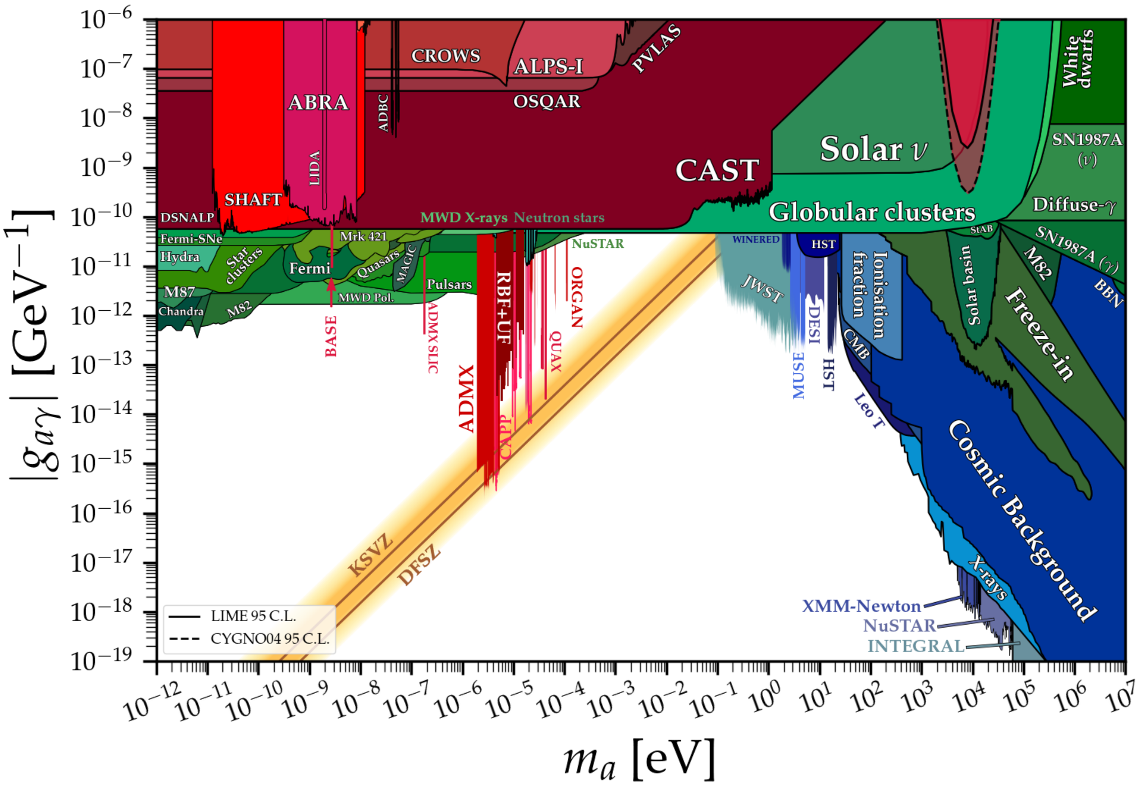

Figure 6.28: our projected exclusion region lies within the solar neutrinos argument

We see above that our sensitivity projection lies within an area that was already excluded, nevertheless it is the first case in which the bound is provided by an experiment. We must specify however, that our analysis was done under the assumption that all the recorded events were actually background events. In this sense our analysis should be regarded as a projection sensitivity on the ALP parameter space.

Finally, we present the constraints on KK axions. In this case, our results are less conclusive than in the ALP scenario. As shown, for instance, in Eq.(5.59), the compactification radius enters the event rate as an overall multiplicative factor. Consequently, large values of the compactification radius R can still be allowed provided that the axion–photon coupling $g_{a\gamma\gamma}$ is sufficiently small. Conversely, larger values of $g_{a\gamma\gamma}$ may be compatible with the data for sufficiently small values of R.
Starting from the differential rate for KK axions given in Fig.(5.22), which depends on the energy, or equivalently on the mass of the KK mode, we included the total detection efficiency and integrated over the full mass range to obtain the expected total signal rate:

$$N_{\rm S}(g_{a\gamma\gamma}, \delta, R = 1\ {\rm keV}^{-1}) = \int_0^\infty dm_a, \epsilon_{\rm tot}(m_a), N_0(m_a, g_{a\gamma\gamma}, \delta, R = 1\ {\rm keV}^{-1}). \tag{6.22}$$

To constrain the parameter space, as done in the ALP case, we adopt a Bayesian framework. We model the observed counts via a Poisson likelihood $\mathcal{L}(g_{a\gamma\gamma}, R) = {\rm Pois}(N_{\rm obs} \mid \mu)$, with an expected mean parameter

$$\mu(R, \delta, g_{a\gamma\gamma}) = N_{\rm bkg} + R^\delta N_{\rm S}(g_{a\gamma\gamma}, \delta, R = 1\ {\rm keV}^{-1}) = N_{\rm bkg} + N_{\rm S,R}(R, \delta, g_{a\gamma\gamma}). \tag{6.23}$$

This calculation was performed for each number of extra dimensions $\delta \in 1, 2, 3$ and for coupling values in the range

$$g_{a\gamma\gamma} \in [10^{-14}, 10^{-6}]\ {\rm GeV}^{-1},$$

while fixing the compactification radius to $R = 1\ \mathrm{keV}^{-1}$. Once the signal rate is known as a function of $g_{a\gamma\gamma}$, a statistical procedure analogous to that used in the ALP analysis can be applied. Following our previous assumptions, we set $N_{\mathrm{obs}} = N_{\mathrm{bkg}} = 97$. We assign a uniform prior to the physical coupling $g_{a\gamma\gamma}$ and to the compactification radius, for $R \in [10^{-4}, 10^{4}]\ \mathrm{keV}^{-1}$. After exploring the global joint posterior distribution via a Markov Chain Monte Carlo sampler deploying 40 walkers for 3×10^{4} steps (with a 2×10^{3} step burn-in), we compute the conditional posterior quantiles numerically.
Specifically, at each fixed value of the coupling $g_{a\gamma\gamma}$, we construct the 1D conditional posterior distribution for the radius, $p(R \mid g_{a\gamma\gamma}, \mathrm{data}) \propto \mathcal{L}(g_{a\gamma\gamma}, R)\, \pi(R)$. The limit boundary $R_{\mathrm{up}}(g_{a\gamma\gamma})$ at C.L. of $1-\alpha$ is then strictly defined as the upper quantile satisfying:

$$\int_{R_{\min}}^{R_{\mathrm{up}}(g_{a\gamma\gamma})} p(R \mid g_{a\gamma\gamma}, \mathrm{data})\, dR = 1 - \alpha, \tag{6.24}$$

with $1 - \alpha \in \{0.90, 0.95\}$. The parameter space lying strictly above the resulting $R_{\mathrm{up}}(g_{a\gamma\gamma})$ curve accounts for the remaining α of the conditional posterior probability mass and is therefore excluded. We show in Fig. (6.29) our results.

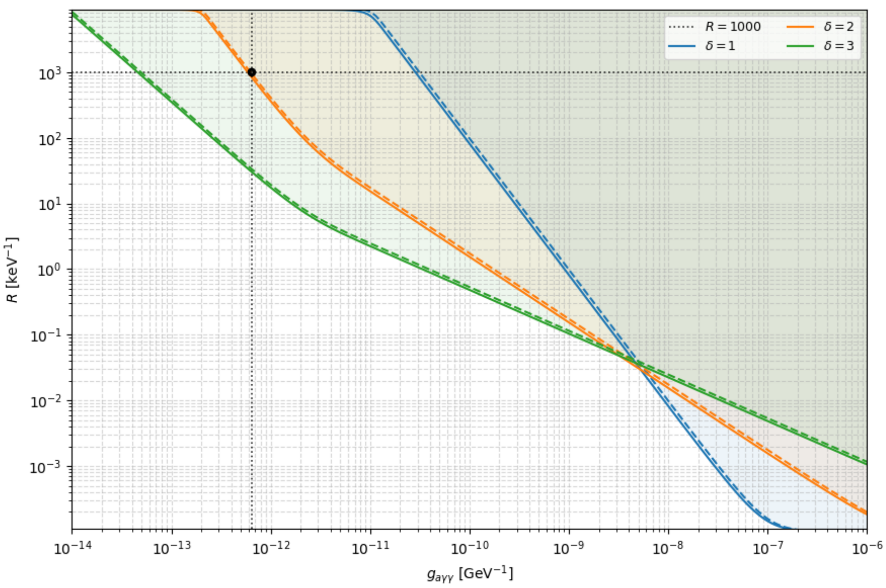

Figure 6.29: Bayesian credible exclusion regions in the $(g_{a\gamma\gamma}, R)$ plane for different numbers of extra dimensions δ. The region strictly above each curve is excluded. Solid lines correspond to a 90% Credible Level, while dashed lines correspond to a 95% Credible Level.

The horizontal dotted line in Fig. (6.29) corresponds to the benchmark value $R = 1000\ \mathrm{keV}^{-1}$. with $\delta = 2$, its intersection with the 95% credible-level exclusion curve yields

$$g_{a\gamma\gamma} < 6.44 \times 10^{-13}\ \mathrm{GeV}^{-1}.$$

This projected sensitivity, if confirmed by a dedicated LIME background study, would improve the current literature bound on the axion–photon coupling by approximately a factor of 1.4. In fact, this same benchmark was considered by the NEWS-G collaboration in Ref. [193], which obtained the limit

$$g_{a\gamma\gamma} < 8.99 \times 10^{-13}\ \mathrm{GeV}^{-1}.$$

It is important to stress, however, that the two analyses adopt substantially different KK-axion number densities. For R=1000 keV^{-1} and $\delta = 2$, Ref. [193] assumed a local KK-axion number density $n_{\mathrm{KK}} = 4.07 \times 10^{7}\ \mathrm{cm}^{-3}$, whereas our calculation, described in Sec. (5.4.2), gives $n_{\mathrm{KK}} = 2.6 \times 10^{6}\ \mathrm{cm}^{-3}$. The number density employed in our analysis is therefore about 16 times smaller. Since the expected signal rate scales linearly with the local axion number density, our bound is obtained with a predicted signal that is about sixteen times weaker than that assumed in Ref. [193]. We finally note that a subsequent calculation in Ref. [191] revised the expected KK-axion number density downward to $n_{\mathrm{KK}} = 0.79\times10^6\ \mathrm{cm}^{-3}$. To the best of our knowledge, the limit reported by the NEWS-G Collaboration has not yet been updated using this revised density estimate.

Conclusions

In this thesis, constraints on axion-like particles (ALPs) and Kaluza–Klein (KK) axions have been investigated through the study of the solar axion basin.

We started from the very beginning, first introducing the strong CP problem, and explaining how it provides a natural motivation for axions. We then reviewed the main axion models and presented the main arguments used to constrain their parameter space. Then we discussed extra-dimensional theories, together with the properties of KK axions.

The production of axions in the solar interior was studied by considering the main relevant production mechanisms, with particular focus on the axion-photon coupling. We derived the axion decay rate at Earth, as well as the energy spectra expected for KK axions.

The theoretical predictions were then compared with data collected by the LIME detector of the CYGNO collaboration at the Gran Sasso National Laboratories. A dedicated data analysis was performed, including the evaluation of the signal selection efficiency and the impact of the applied cuts. Under the assumption that all recorded events are background events, we obtained a sensitivity projection in the ALP and KK-axion parameter spaces.

For generic ALPs, our projected sensitivity, if confirmed by a dedicated LIME background study, would yield a direct-detection exclusion region for masses in the range $m_a \in [1, 30]\,\mathrm{keV}$, reaching couplings down to $g_{a\gamma\gamma} \leq 3 \times 10^{-8}\,\mathrm{GeV}^{-1}$. Although this parameter space is already excluded by indirect constraints, our result would provide a complementary direct-detection bound.

For KK axions, considering a specific benchmark choice of the model parameters, our projected sensitivity, if again confirmed by a dedicated LIME background study, would improve the current bound reported in the literature by approximately a factor of 1.4. This result relies on an updated calculation of the expected signal, which is approximately 16 times smaller than that adopted in previous state-of-the-art analysis.

Bibliography

[1] Hook. "TASI Lectures on the Strong CP Problem and Axions". In: *PoS* TASI2018 (2019), p. 004. DOI: 10.22323/1.333.0004. arXiv: 1812.02669 [hep-ph].

[2] Pendlebury, Afach, Ayres, Baker, Ban, Bison, Bodek, Burghoff, Geltenbort, Green, Griffith, Grinten, Grujic, Harris, Helaine, Iaydjiev, Ivanov, Kasprzak, Kermaidic, Kirch, Koch, Komposch, Kozela, Krempel, Lauss, Lefort, Lemiere, May, Musgrave, Naviliat-Cuncic, Piegsa, Pignol, Prashanth, Quemener, Rawlik, Rebreyend, Richardson, Ries, Roccia, Rozpedzik, Schnabel, Schmidt-Wellenburg, Severijns, Shiers, Thorne, Weis, Winston, Wursten, Zejma, Zsigmond. "A Revised Experimental Upper Limit on the Electric Dipole Moment of the Neutron". In: *Phys. Rev. D* 92.9 (2015), p. 092003. DOI: 10.1103/PhysRevD.92.092003. arXiv: 1509.04411 [hep-ex].

[3] Abel, Afach, Ayres, Baker, Ban, Bison, Bodek, Bondar, Burghoff, Chanel, Chowdhuri, Chiu, Clement, Crawford, Daum, Emmenegger, Ferraris-Bouchez, Fertl, Flaux, Franke, Fratangelo, Geltenbort, Green, Griffith, Grinten, Grujic, Harris, Hayen, Heil, Henneck, Helaine, Hild, Hodge, Horras, Iaydjiev, Ivanov, Kasprzak, Kermaidic, Kirch, Knecht, Knowles, Koch, Koss, Komposch, Kozela, Kraft, Krempel, Kuzniak, Lauss, Lefort, Lemiere, Leredde, Mohanmurthy, Mtchedlishvili, Musgrave, Naviliat-Cuncic, Pais, Piegsa, Pierre, Pignol, Plonka-Spehr, Prashanth, Quemener, Rawlik, Rebreyend, Rienacker, Ries, Roccia, Rogel, Rozpedzik, Schnabel, Schmidt-Wellenburg, Severijns, Shiers, Tavakoli Dinani, Thorne, Virot, Voigt, Weis, Wursten, Wyszynski, Zejma, Zenner, Zsigmond. "Measurement of the Permanent Electric Dipole Moment of the Neutron". In: *Phys. Rev. Lett.* 124.8 (2020), p. 081803. DOI: 10.1103/PhysRevLett.124.081803. arXiv: 2001.11966 [hep-ex].

[4] Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, 2014.

[5] Peskin, Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995.

[6] Yang, Mills. "Conservation of Isotopic Spin and Isotopic Gauge Invariance". In: *Physical Review* 96.1 (1954), pp. 191–195. DOI: 10.1103/PhysRev.96.191.

[7] Gell-Mann. *The Eightfold Way: A Theory of Strong Interaction Symmetry*. Tech. rep. CTSL-20. California Institute of Technology, Synchrotron Laboratory, 1961. DOI: 10.2172/4008239.

[8] Ne'eman. "Derivation of Strong Interactions from a Gauge Invariance". In: *Nuclear Physics* 26 (1961), pp. 222–229. DOI: 10.1016/0029-5582(61)90134-1.

[9] Gell-Mann. "A Schematic Model of Baryons and Mesons". In: *Physics Letters* 8.3 (1964), pp. 214–215. DOI: 10.1016/S0031-9163(64)92001-3.

[10] Zweig. *An SU(3) Model for Strong Interaction Symmetry and Its Breaking*. Tech. rep. CERN-TH-401. CERN, 1964. DOI: 10.17181/CERN-TH-401.

[11] Greenberg. "Spin and Unitary-Spin Independence in a Paraquark Model of Baryons and Mesons". In: *Physical Review Letters* 13 (1964), pp. 598–602. DOI: 10.1103/PhysRevLett.13.598.

[12] Han, Nambu. "Three-Triplet Model with Double SU(3) Symmetry". In: *Physical Review* 139 (1965), B1006–B1010. DOI: 10.1103/PhysRev.139.B1006.

[13] Bogoliubov, Struminsky, Tavkhelidze. *On Composite Models in the Theory of Elementary Particles*. Tech. rep. JINR D-1968. Dubna: Joint Institute for Nuclear Research, 1965.

[14] Khriplovich. "Green's Functions in Theories with Non-Abelian Gauge Group". In: *Soviet Journal of Nuclear Physics* 10 (1970). Translated from Yad. Fiz. 10 (1969) 409–420, pp. 235–242.

[15] Hooft. "Renormalization of Massless Yang-Mills Fields". In: *Nuclear Physics B* 33.1 (1971), pp. 173–199. DOI: 10.1016/0550-3213(71)90139-8.

[16] Gross, Wilczek. "Ultraviolet Behavior of Non-Abelian Gauge Theories". In: *Physical Review Letters* 30.26 (1973), pp. 1343–1346. DOI: 10.1103/PhysRevLett.30.1343.

[17] Politzer. "Reliable Perturbative Results for Strong Interactions?" In: *Physical Review Letters* 30.26 (1973), pp. 1346–1349. DOI: 10.1103/PhysRevLett.30.1346.

[18] Fritzsch, Gell-Mann, Leutwyler. "Advantages of the Color Octet Gluon Picture". In: *Physics Letters B* 47.4 (1973), pp. 365–368. DOI: 10.1016/0370-2693(73)90625-4.

[19] Nambu, Jona-Lasinio. "Dynamical Model of Elementary Particles Based on an Analogy with Superconductivity. II". In: *Physical Review* 124 (1961), pp. 246–254. DOI: 10.1103/PhysRev.124.246.

[20] Callan, Dashen, Gross. "The Structure of the Gauge Theory Vacuum". In: *Physics Letters B* 63.3 (1976), pp. 334–340. DOI: 10.1016/0370-2693(76)90277-X.

[21] Tong. "TASI Lectures on Solitons: Lecture 1: Instantons". In: (2005). TASI lectures. URL: https://www.damtp.cam.ac.uk/user/tong/tasi/instanton.pdf.

[22] Crewther. "Effects of Topological Charge in Gauge Theories". In: *Facts and Prospects of Gauge Theories*. 1978, pp. 47–153. DOI: 10.1007/978-3-7091-8538-4_3.

[23] Forkel. "A Primer on Instantons in QCD". In: (2000). arXiv: hep-ph/0009136 [hep-ph].

[24] Coleman. *Aspects of Symmetry: Selected Erice Lectures*. Cambridge: Cambridge University Press, 1985. ISBN: 9780521318273.

[25] Atiyah, Drinfeld, Hitchin, Manin. "Construction of Instantons". In: *Physics Letters A* 65.3 (1978), pp. 185–187. DOI: 10.1016/0375-9601(78)90141-X.

[26] Belavin, Polyakov, Schwartz, Tyupkin. "Pseudoparticle Solutions of the Yang-Mills Equations". In: *Physics Letters B* 59.1 (1975), pp. 85–87. DOI: 10.1016/0370-2693(75)90163-X.

[27] Forkel. "A Primer on Instantons in QCD". In: (2000). arXiv: hep-ph/0009136 [hep-ph].

[28] Hooft. "Symmetry Breaking Through Bell-Jackiw Anomalies". In: *Phys. Rev. Lett.* 37 (1976), pp. 8–11. DOI: 10.1103/PhysRevLett.37.8.

[29] 't Hooft. "Computation of the Quantum Effects Due to a Four-Dimensional Pseudoparticle". In: *Physical Review D* 14.12 (1976). Erratum: Phys. Rev. D 18, 2199 (1978), pp. 3432–3450. DOI: 10.1103/PhysRevD.14.3432.

[30] Hook. "TASI Lectures on the Strong CP Problem and Axions". In: (2018). arXiv: 1812.02669 [hep-ph].

[31] Navas. "Review of Particle Physics". In: *Phys. Rev. D* 110 (2024). Particle Data Group; see also 2025 update tables, p. 030001. DOI: 10.1103/PhysRevD.110.030001.

[32] Babu, Mohapatra. "A Solution to the Strong CP Problem Without an Axion". In: *Phys. Rev. D* 41 (1990), p. 1286. DOI: 10.1103/PhysRevD.41.1286.

[33] Barr, Chang, Senjanovic. "Strong CP Problem and Parity". In: *Phys. Rev. Lett.* 67 (1991), pp. 2765–2768. DOI: 10.1103/PhysRevLett.67.2765.

[34] Ellis, Gaillard. "Strong and Weak CP Violation". In: *Nucl. Phys. B* 150 (1979), pp. 141–162.

[35] Nelson. "Naturally Weak CP Violation". In: *Phys. Lett. B* 136 (1984), pp. 387–391. DOI: 10.1016/0370-2693(84)92025-2.

[36] Barr. "Solving the Strong CP Problem Without the Peccei-Quinn Symmetry". In: *Phys. Rev. Lett.* 53 (1984), pp. 329–332. DOI: 10.1103/PhysRevLett.53.329.

[37] Zeldovich, Kobzarev, Okun. "Cosmological Consequences of the Spontaneous Breakdown of Discrete Symmetry". In: *Soviet Physics JETP* 40 (1975), p. 1.

[38] Dine, Leigh, Kagan. "Supersymmetry and the Nelson-Barr mechanism". In: *Phys. Rev. D* 48 (1993), pp. 2214–2223. DOI: 10.1103/PhysRevD.48.2214. arXiv: hep-ph/9303296.

[39] Hiller, Schmaltz. "Solving the Strong CP Problem with Supersymmetry". In: *Phys. Lett. B* 514 (2001), pp. 263–268. DOI: 10.1016/S0370-2693(01)00814-0. arXiv: hep-ph/0105254.

[40] Barr. "Supersymmetric solutions to the strong CP problem". In: *Phys. Rev. D* 56 (1997), pp. 1475–1480. DOI: 10.1103/PhysRevD.56.1475. arXiv: hep-ph/9612396.

[41] Vecchi. "Spontaneous CP violation and the strong CP problem". In: *JHEP* 04 (2017), p. 149. DOI: 10.1007/JHEP04(2017)149. arXiv: 1412.3805 [hep-ph].

[42] Peccei, Quinn. "CP Conservation in the Presence of Pseudoparticles". In: *Phys. Rev. Lett.* 38 (1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440.

[43] Peccei, Quinn. "Constraints Imposed by CP Conservation in the Presence of Pseudoparticles". In: *Phys. Rev. D* 16 (1977), pp. 1791–1797. DOI: 10.1103/PhysRevD.16.1791.

[44] Vafa, Witten. "Parity Conservation in Quantum Chromodynamics". In: *Physical Review Letters* 53 (1984), p. 535. DOI: 10.1103/PhysRevLett.53.535.

[45] Barr, Seckel. "Planck-Scale Corrections to Axion Models". In: *Phys. Rev. D* 46 (1992), pp. 539–549. DOI: 10.1103/PhysRevD.46.539.

[46] Kamionkowski, March-Russell. "Planck-Scale Physics and the Peccei-Quinn Mechanism". In: *Phys. Lett. B* 282 (1992), pp. 137–141. DOI: 10.1016/0370-2693(92)90492-M. arXiv: hep-th/9202003.

[47] Randall. "Composite axion models and Planck scale physics". In: *Physics Letters B* 284 (1992), pp. 77–80. DOI: 10.1016/0370-2693(92)91928-3.

[48] Sikivie. "Invisible Axion Search Methods". In: *Reviews of Modern Physics* 93.1 (2021), p. 015004. DOI: 10.1103/RevModPhys.93.015004. arXiv: 2003.02206 [hep-ph].

[49] Asano. "Search for a Rare Decay Mode $K^+ \to \pi^+ \nu\bar{\nu}$ and Axion". In: *Physics Letters B* 107 (1981), pp. 159–162. DOI: 10.1016/0370-2693(81)91172-2.

[50] Dine, Fischler, Srednicki. "A Simple Solution to the Strong CP Problem with a Harmless Axion". In: *Phys. Lett. B* 104 (1981), p. 199. DOI: 10.1016/0370-2693(81)90590-6.

[51] Zhitnitsky. "On Possible Suppression of the Axion Hadron Interactions". In: *Sov. J. Nucl. Phys.* 31 (1980). [Yad. Fiz. 31, 497 (1980)], p. 260.

[52] Bardeen, Tye. "Current Algebra Applied to Properties of the Light Higgs Boson". In: *Physics Letters B* 74.3 (1978), pp. 229–232. DOI: 10.1016/0370-2693(78)90560-9.

[53] Kandaswamy, Salomonson, Schechter. "Mass of the Axion". In: *Physical Review D* 17 (1978), pp. 3051–3054. DOI: 10.1103/PhysRevD.17.3051.

[54] Srednicki. "Axion Couplings to Matter. 1. CP-Conserving Parts". In: *Nuclear Physics B* 260 (1985), pp. 689–700. DOI: 10.1016/0550-3213(85)90054-9.

[55] Georgi, Kaplan, Randall. "Manifesting the Invisible Axion at Low Energies". In: *Physics Letters B* 169 (1986), pp. 73–78. DOI: 10.1016/0370-2693(86)90688-X.

[56] Bardeen, Peccei, Yanagida. "Constraints on Variant Axion Models". In: *Nuclear Physics B* 279.3–4 (1987), pp. 401–428. DOI: 10.1016/0550-3213(87)90003-4.

[57] Kaplan. "Opening the Axion Window". In: *Nuclear Physics B* 260 (1985), pp. 215–226. DOI: 10.1016/0550-3213(85)90319-0.

[58] Mayle, Wilson, Ellis, Olive, Schramm, Steigman. "Constraints on Axions from SN 1987A". In: *Physics Letters B* 203 (1988), pp. 188–196. DOI: 10.1016/0370-2693(88)91595-X.

[59] Mayle, Wilson, Ellis, Olive, Schramm, Steigman. "Updated Constraints on Axions from SN1987A". In: *Physics Letters B* 219 (1989), pp. 515–520. DOI: 10.1016/0370-2693(89)91104-0.

[60] Kim. "Weak-Interaction Singlet and Strong CP Invariance". In: *Phys. Rev. Lett.* 43 (1979), p. 103. DOI: 10.1103/PhysRevLett.43.103.

[61] Shifman, Vainshtein, Zakharov. "Can Confinement Ensure Natural CP Invariance of Strong Interactions?" In: *Nucl. Phys. B* 166 (1980), p. 493. DOI: 10.1016/0550-3213(80)90209-6.

[62] Abbott, Sikivie. "A Cosmological Bound on the Invisible Axion". In: *Phys. Lett. B* 120 (1983), pp. 133–136.

[63] Dine, Fischler. "The Not So Harmless Axion". In: *Phys. Lett. B* 120 (1983), pp. 137–141.

[64] Preskill, Wise, Wilczek. "Cosmology of the Invisible Axion". In: *Phys. Lett. B* 120 (1983), pp. 127–132.

[65] O'Hare. *Axion Cosmology*. Lecture notes/review. NSW 2006, Australia, 2024.

[66] Gross, Pisarski, Yaffe. "QCD and Instantons at Finite Temperature". In: *Rev. Mod. Phys.* 53 (1981), p. 43.

[67] Wantz, Shellard. "Axion Cosmology Revisited". In: *Phys. Rev. D* 82 (2010), p. 123508. arXiv: 0910.1066 [astro-ph.CO].

[68] Borsanyi. "Calculation of the axion mass based on high-temperature lattice quantum chromodynamics". In: *Nature* 539 (2016), pp. 69–71. arXiv: 1606.07494 [hep-lat].

[69] Seckel, Turner. "Isothermal Density Perturbations in an Axion Dominated Inflationary Universe". In: *Phys. Rev. D* 32 (1985), p. 3178.

[70] Linde. "Generation of Isothermal Density Perturbations in the Inflationary Universe". In: *Phys. Lett. B* 158 (1985), p. 375.

[71] Bucher, Moodley, Turok. "The General Primordial Cosmic Perturbation". In: *Phys. Rev. D* 62 (2000), p. 083508. arXiv: astro-ph/9904231.

[72] Davis. "Cosmic Axions from Cosmic Strings". In: *Phys. Lett. B* 180 (1986), pp. 225–230. DOI: 10.1016/0370-2693(86)90300-X.

[73] Jia, Bian. "Cosmological simulation of axion-Higgs strings". In: *Phys. Rev. D* 111.6 (2025), p. 063552. DOI: 10.1103/PhysRevD.111.063552. arXiv: 2412.04218 [hep-ph].

[74] O'Hare, Pierobon, Redondo, Wong. "Simulations of axionlike particles in the postinflationary scenario". In: *Phys. Rev. D* 105.5 (2022), p. 055025. DOI: 10.1103/PhysRevD.105.055025. arXiv: 2112.05117 [hep-ph].

[75] Sikivie. "Of Axions, Domain Walls and the Early Universe". In: *Phys. Rev. Lett.* 48 (1982), pp. 1156–1159. DOI: 10.1103/PhysRevLett.48.1156.

[76] Córdova, Hong, Wang. "Axion domain walls, small instantons, and non-invertible symmetry breaking". In: *JHEP* 05 (2024), p. 325. DOI: 10.1007/JHEP05(2024)325. arXiv: 2309.05636 [hep-th].

[77] Turner. "Early-Universe Thermal Production of Not-So-Invisible Axions". In: *Phys. Rev. Lett.* 59 (1987), pp. 2489–2492. DOI: 10.1103/PhysRevLett.59.2489.

[78] Salvio, Strumia, Xue. "Thermal axion production". In: *JCAP* 01 (2014), p. 011. DOI: 10.1088/1475-7516/2014/01/011. arXiv: 1310.6982 [hep-ph].

[79] Hannestad, Mirizzi, Raffelt. "New cosmological mass limit on thermal relic axions". In: *JCAP* 07 (2005), p. 002. DOI: 10.1088/1475-7516/2005/07/002. arXiv: hep-ph/0504059.

[80] Archidiacono, Hannestad, Mirizzi, Raffelt, Wong. "Axion hot dark matter bounds after Planck". In: *JCAP* 10 (2013), p. 020. DOI: 10.1088/1475-7516/2013/10/020. arXiv: 1307.0615 [astro-ph.CO].

[81] Lesgourgues, Pastor. "Massive neutrinos and cosmology". In: *Phys. Rept.* 429 (2006), pp. 307–379. DOI: 10.1016/j.physrep.2006.04.001. arXiv: astro-ph/0603494.

[82] Baumann, Green, Meyers, Wallisch. "Phases of New Physics in the CMB". In: *JCAP* 01 (2016), p. 007. DOI: 10.1088/1475-7516/2016/01/007. arXiv: 1508.06342 [astro-ph.CO].

[83] Bennett, Buldgen, Salas, Drewes, Gariazzo, Pastor, Wong. "Towards a precision calculation of N_{eff} in the Standard Model II: neutrino decoupling in the presence of flavour oscillations and finite-temperature QED". In: *JCAP* 04 (2021), p. 073. DOI: 10.1088/1475-7516/2021/04/073. arXiv: 2012.02726 [hep-ph].

[84] Abazajian. "CMB-S4 Science Book, First Edition". In: (2016). arXiv: 1610.02743 [astro-ph.CO].

[85] Aghanim. "Planck 2018 results. VI. Cosmological parameters". In: *Astron. Astrophys.* 641 (2020), A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO].

[86] D'Eramo, Di Valentino, Giarè, Hajkarim, Melchiorri, Mena, Renzi, Yun. "Cosmological bound on the QCD axion mass, redux". In: *JCAP* 09 (2022), p. 022. DOI: 10.1088/1475-7516/2022/09/022. arXiv: 2205.07849 [astro-ph.CO].

[87] Caloni, Gerbino, Lattanzi, Visinelli. "Novel cosmological bounds on thermally-produced axion-like particles". In: *JCAP* 09 (2022), p. 021. DOI: 10.1088/1475-7516/2022/09/021. arXiv: 2205.01637 [astro-ph.CO].

[88] Giarè, Renzi, Melchiorri, Mena, Di Valentino. "Cosmological forecasts on thermal axions, relic neutrinos and light elements". In: *Mon. Not. Roy. Astron. Soc.* 511.1 (2022), pp. 1373–1382. DOI: 10.1093/mnras/stac086. arXiv: 2110.00340 [astro-ph.CO].

[89] D'Eramo, Hajkarim, Yun. "Thermal Axion Production at Low Temperatures: A Smooth Treatment of the QCD Phase Transition". In: *Phys. Rev. Lett.* 128.15 (2022), p. 152001. DOI: 10.1103/PhysRevLett.128.152001. arXiv: 2108.04259 [hep-ph].

[90] Abazajian. "CMB-S4 Science Case, Reference Design, and Project Plan". In: (2019). arXiv: 1907.04473 [astro-ph.IM].

[91] Hall, Jedamzik, March-Russell, West. "Freeze-In Production of FIMP Dark Matter". In: *JHEP* 03 (2010), p. 080. DOI: 10.1007/JHEP03(2010)080. arXiv: 0911.1120 [hep-ph].

[92] Carenza, Fore, Giannotti, Mirizzi, Reddy. "Thermal axions with multi-eV masses are possible in low-reheating scenarios". In: *Phys. Rev. Lett.* 126.7 (2021), p. 071102. DOI: 10.1103/PhysRevLett.126.071102. arXiv: 2010.02943 [hep-ph].

[93] Arias, Bernal, Osiński, Roszkowski, Venegas. "Revisiting signatures of thermal axions in nonstandard cosmologies". In: *Phys. Rev. D* 109.12 (2024), p. 123529. DOI: 10.1103/PhysRevD.109.123529. arXiv: 2308.01352 [hep-ph].

[94] Kulkarni, Ostriker. "What is the halo mass function in a fuzzy dark matter cosmology?" In: *Mon. Not. Roy. Astron. Soc.* 510 (2021), p. 1425. arXiv: 2011.02116.

[95] Dentler, Marsh, Hlozek, Laguë, Rogers, Grin. "Fuzzy dark matter and the Dark Energy Survey Year 1 data". In: *Mon. Not. Roy. Astron. Soc.* 515 (2022), p. 5646. arXiv: 2111.01199.

[96] Rogers, Peiris. "Strong Bound on Canonical Ultralight Axion Dark Matter from the Lyman-Alpha Forest". In: *Phys. Rev. Lett.* 126 (2021), p. 071302. arXiv: 2007.12705.

[97] Hložek, Marsh, Grin. "Using the full power of the cosmic microwave background to probe axion dark matter". In: *Mon. Not. Roy. Astron. Soc.* 476 (2018), p. 3063. arXiv: 1708.05681.

[98] Marsh, Pop. "Axion dark matter, solitons and the cusp-core problem". In: *Mon. Not. Roy. Astron. Soc.* 451 (2015), p. 2479. arXiv: 1502.03456.

[99] Matos, Ureña López, Lee. "Short review of the main achievements of the scalar field, fuzzy, ultralight, wave, BEC dark matter model". In: *Front. Astron. Space Sci.* 11 (2024), p. 1347518. arXiv: 2312.00254.

[100] Caputo, Raffelt. "Astrophysical Axion Bounds: The 2024 Edition". In: *PoS* COSMICWISPers (2024), p. 041. DOI: 10.22323/1.454.0041. arXiv: 2401.13728.

[101] Schlattl, Weiss, Raffelt. "Helioseismological constraint on solar axion emission". In: *Astropart. Phys.* 10 (1999), pp. 353–359. DOI: 10.1016/S0927-6505(98)00063-2. arXiv: hep-ph/9807476.

[102] Vinyoles, Serenelli, Villante, Basu, Redondo, Isern. "New axion and hidden photon constraints from a solar data global fit". In: *JCAP* 10 (2015), p. 015. DOI: 10.1088/1475-7516/2015/10/015. arXiv: 1501.01639 [astro-ph.SR].

[103] Gondolo, Raffelt. "Solar neutrino limit on axions and keV-mass bosons". In: *Phys. Rev. D* 79 (2009), p. 107301. DOI: 10.1103/PhysRevD.79.107301. arXiv: 0807.2926 [astro-ph].

[104] Di Lella, Pilaftsis, Raffelt, Zioutas. "Search for solar Kaluza-Klein axions in theories of low-scale quantum gravity". In: *Phys. Rev. D* 62 (2000), p. 125011. DOI: 10.1103/PhysRevD.62.125011. arXiv: hep-ph/0006327.

[105] Hannah, Hudson, Hurford, Lin. "Constraining the hard X-ray properties of the quiet Sun with new RHESSI observations". In: *Astrophys. J.* 724 (2010), pp. 487–492. DOI: 10.1088/0004-637X/724/1/487. arXiv: 1009.2918 [astro-ph.SR].

[106] Beaufort, Bastero-Gil, Luce, Santos. "New solar x-ray constraints on keV axionlike particles". In: *Phys. Rev. D* 108.8 (2023), p. L081302. arXiv: 2303.06968 [hep-ph].

[107] Gburek, Sylwester, Kowalinski, Siarkowski, Sylwester, Podgorski, Trzebinski, Kordylewski, Plocieniak, Kepa, Bakala, Kuzin, Reale, Phillips. "SphinX: The Solar Photometer in X-Rays". In: *Solar Physics* 283 (2013), pp. 631–649. DOI: 10.1007/s11207-012-0201-2.

[108] Redondo. "Solar axion flux from the axion-electron coupling". In: *JCAP* 12 (2013), p. 008. arXiv: 1310.0823 [hep-ph].

[109] Dent, Dutta, Newstead, Thompson. "Inverse Primakoff Scattering as a Probe of Solar Axions at Liquid Xenon Direct Detection Experiments". In: *Phys. Rev. Lett.* 125.13 (2020), p. 131805. arXiv: 2006.15118 [hep-ph].

[110] Di Lella, Zioutas. "Observational evidence for gravitationally trapped massive axion(-like) particles". In: *Astropart. Phys.* 19 (2003), pp. 145–170. arXiv: astro-ph/0207073 [astro-ph].

[111] Van Tilburg. "Stellar basins of gravitationally bound particles". In: *Phys. Rev. D* 104.2 (2021), p. 023019. arXiv: 2006.12431 [hep-ph].

[112] DeRocco, Wegsman, Grefenstette, Huang, Van Tilburg. "First Indirect Detection Constraints on Axions in the Solar Basin". In: *Phys. Rev. Lett.* 129.10 (2022), p. 101101. arXiv: 2205.05700 [hep-ph].

[113] Pshirkov, Popov. "Conversion of Dark Matter Axions to Photons in Magnetospheres of Neutron Stars". In: *J. Exp. Theor. Phys.* 108 (2009), pp. 384–388. arXiv: 0711.1264 [astro-ph].

[114] Huang, Kadota, Sekiguchi, Tashiro. "Radio telescope search for the resonant conversion of cold dark matter axions from the magnetized astrophysical sources". In: *Phys. Rev. D* 97.12 (2018), p. 123001. arXiv: 1803.08230 [hep-ph].

[115] Prabhu. "Axion production in pulsar magnetosphere gaps". In: *Phys. Rev. D* 104.5 (2021), p. 055038. arXiv: 2104.14569 [hep-ph].

[116] Straniero, Pallanca, Dalessandro, Domínguez, Ferraro, Giannotti, Mirizzi, Piersanti. "The RGB tip of galactic globular clusters and the revision of the axion-electron coupling bound". In: *Astron. Astrophys.* 644 (2020), A166. DOI: 10.1051/0004-6361/202038775. arXiv: 2010.03833 [astro-ph.SR].

[117] Straniero. *Overview of axion astrophysics and perspectives from future astronomical probes.* Talk at GGI Axion Institute. Work in progress, slides. 2023.

[118] Capozzi, Raffelt. "Axion and neutrino bounds improved with new calibrations of the tip of the red-giant branch using geometric distance determinations". In: *Phys. Rev. D* 102.8 (2020), p. 083007. DOI: 10.1103/PhysRevD.102.083007. arXiv: 2007.03694 [astro-ph.SR].

[119] Raffelt. "Astrophysical axion bounds diminished by screening effects". In: *Phys. Rev. D* 33 (1986), p. 897. DOI: 10.1103/PhysRevD.33.897.

[120] Straniero, Ayala, Giannotti, Mirizzi, Domínguez. "Axion-Photon Coupling: Astrophysical Constraints". In: *11th Patras Workshop on Axions, WIMPs and WISPs.* 2015, pp. 77–81. DOI: 10.3204/DESY-PROC-2015-02/straniero_oscar.

[121] Dolan, Hiskens, Volkas. "Advancing globular cluster constraints on the axion-photon coupling". In: *JCAP* 10 (2022), p. 096. DOI: 10.1088/1475-7516/2022/10/096. arXiv: 2207.03102 [hep-ph].

[122] Harris. "The white dwarf luminosity function from SDSS imaging data". In: *Astron. J.* 131 (2006), p. 571. DOI: 10.1086/497966. arXiv: astro-ph/0510820 [astro-ph].

[123] Rowell, Hambly. "White dwarfs in the SuperCOSMOS Sky Survey: the thin disc, thick disc and spheroid luminosity functions". In: *Mon. Not. Roy. Astron. Soc.* 417 (2011), pp. 93–113. DOI: 10.1111/j.1365-2966.2011.19029.x. arXiv: 1102.3193 [astro-ph.SR].

[124] Isern, Hernanz, Garcia-Berro. "Axion cooling of white dwarfs". In: *Astrophys. J. Lett.* 392 (1992), pp. L23–L25. DOI: 10.1086/186416.

[125] Leinson. "Impact of axions on the Cassiopea A neutron star cooling". In: *JCAP* 09 (2021), p. 001. DOI: 10.1088/1475-7516/2021/09/001. arXiv: 2105.14745 [hep-ph].

[126] Buschmann, Dessert, Foster, Long, Safdi. "Upper Limit on the QCD Axion Mass from Isolated Neutron Star Cooling". In: *Phys. Rev. Lett.* 128.9 (2022), p. 091102. DOI: 10.1103/PhysRevLett.128.091102. arXiv: 2111.09892 [hep-ph].

[127] Raffelt, Seckel. "Bounds on Exotic Particle Interactions from SN 1987A". In: *Phys. Rev. Lett.* 60 (1988), p. 1793. DOI: 10.1103/PhysRevLett.60.1793.

[128] Raffelt. "Astrophysical methods to constrain axions and other novel particle phenomena". In: *Phys. Rept.* 198 (1990), pp. 1–113. DOI: 10.1016/0370-1573(90)90054-6.

[129] Carenza, Fischer, Giannotti, Guo, Martínez-Pinedo, Mirizzi. "Improved axion emissivity from a supernova via nucleon-nucleon bremsstrahlung". In: *JCAP* 10 (2019). [Erratum: JCAP 05 (2020) E01], p. 016. DOI: 10.1088/1475-7516/2019/10/016. arXiv: 1906.11844 [hep-ph].

[130] Penrose, Floyd. "Extraction of rotational energy from a black hole". In: *Nature Phys. Sci.* 229 (1971), pp. 177–179. DOI: 10.1038/physci229177a0.

[131] Zouros, Eardley. "Instabilities of massive scalar perturbations of a rotating black hole". In: *Annals Phys.* 118 (1979), pp. 139–155. DOI: 10.1016/0003-4916(79)90237-9.

[132] Baryakhtar. "Dark Matter In Extreme Astrophysical Environments". In: *Snowmass 2021.* Mar. 2022. arXiv: 2203.07984 [hep-ph].

[133] Irastorza. "An introduction to axions and their detection". In: *SciPost Phys. Lect. Notes* 45 (2022). DOI: 10.21468/SciPostPhysLectNotes.45. arXiv: 2109.07376 [hep-ph].

[134] Ehret, Frede, Ghazaryan, Hildebrandt, Knabbe. "New ALPS Results on Hidden-Sector Lightweights". In: *Phys. Lett. B* 689 (2010), pp. 149–155. DOI: 10.1016/j.physletb.2010.04.066. arXiv: 1004.1313 [hep-ex].

[135] Ballou. "New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall". In: *Phys. Rev. D* 92.9 (2015), p. 092002. DOI: 10.1103/PhysRevD.92.092002. arXiv: 1506.08082 [hep-ex].

[136] Bähre. "Any light particle search II — Technical Design Report". In: *JINST* 8 (2013), T09001. DOI: 10.1088/1748-0221/8/09/T09001. arXiv: 1302.5647 [physics.ins-det].

[137] Hoogeveen. "Terrestrial axion production and detection using RF cavities". In: *Phys. Lett. B* 288 (1992), pp. 195–200. DOI: 10.1016/0370-2693(92)91977-H.

[138] Battesti, Fouche, Detlefs, Roth, Berceau, Duc, Frings, Rikken, Rizzo. "A photon regeneration experiment for axionlike particle search using X-rays". In: *Phys. Rev. Lett.* 105 (2010), p. 250405. DOI: 10.1103/PhysRevLett.105.250405. arXiv: 1008.2672 [physics.ins-det].

[139] Della Valle, Ejlli, Gastaldi, Messineo, Milotti, Pengo, Ruoso, Zavattini. "The PVLAS experiment: measuring vacuum magnetic birefringence and dichroism with a birefringent Fabry–Perot cavity". In: *Eur. Phys. J. C* 76.1 (2016), p. 24. DOI: 10.1140/epjc/s10052-015-3869-8. arXiv: 1510.08052 [physics.optics].

[140] Hartman, Rivere, Battesti, Rizzo. "Noise characterization for resonantly-enhanced polarimetric vacuum magnetic-birefringence experiments". In: *Rev. Sci. Instrum.* 88.12 (2017), p. 123114. DOI: 10.1063/1.4986871. arXiv: 1712.01278 [physics.ins-det].

[141] Asztalos. "A SQUID-based microwave cavity search for dark-matter axions". In: *Phys. Rev. Lett.* 104 (2010), p. 041301. DOI: 10.1103/PhysRevLett.104.041301. arXiv: 0910.5914 [astro-ph.CO].

[142] Al Kenany. "Design and operational experience of a microwave cavity axion detector for the 20–100 μeV range". In: *Nucl. Instrum. Meth. A* 854 (2017), pp. 11–24. DOI: 10.1016/j.nima.2017.02.012. arXiv: 1611.07123 [physics.ins-det].

[143] Chung. "CULTASK, Axion Experiment at CAPP in Korea". In: *13th Patras Workshop on Axions, WIMPs and WISPs*. 2018, pp. 97–101. DOI: 10.3204/DESY-PROC-2017-02/woohyun_chung.

[144] Alesini. "Search for invisible axion dark matter of mass $m_a = 43$ μeV with the QUAX–$a\gamma$ experiment". In: *Phys. Rev. D* 103.10 (2021), p. 102004. DOI: 10.1103/PhysRevD.103.102004. arXiv: 2012.09498 [hep-ex].

[145] Melcón. "Axion Searches with Microwave Filters: the RADES project". In: *JCAP* 05 (2018), p. 040. DOI: 10.1088/1475-7516/2018/05/040. arXiv: 1803.01243 [hep-ex].

[146] McAllister, Flower, Ivanov, Goryachev, Bourhill, Tobar. "The ORGAN Experiment: An axion haloscope above 15 GHz". In: *Phys. Dark Univ.* 18 (2017), pp. 67–72. DOI: 10.1016/j.dark.2017.09.010. arXiv: 1706.00209 [physics.ins-det].

[147] Salemi. "The search for low-mass axion dark matter with ABRACADABRA-10cm". In: *Phys. Rev. Lett.* 127.8 (2021), p. 081801. DOI: 10.1103/PhysRevLett.127.081801. arXiv: 2102.06722 [hep-ex].

[148] Gramolin, Aybas, Johnson, Adam, Sushkov. "Search for axion-like dark matter with ferromagnets". In: *Nature Phys.* 17 (2021), pp. 79–84. DOI: 10.1038/s41567-020-1006-6. arXiv: 2003.03348 [hep-ex].

[149] Crisosto, Sikivie, Sullivan, Tanner, Yang, Rybka. "ADMX SLIC: Results from a Superconducting LC Circuit Investigating Cold Axions". In: *Phys. Rev. Lett.* 124.24 (2020), p. 241101. DOI: 10.1103/PhysRevLett.124.241101. arXiv: 1911.05772 [astro-ph.CO].

[150] Devlin. "Constraints on the Coupling between Axionlike Dark Matter and Photons Using an Antiproton Superconducting Tuned Detection Circuit in a Cryogenic Penning Trap". In: *Phys. Rev. Lett.* 126.4 (2021), p. 041301. DOI: 10.1103/PhysRevLett.126.041301. arXiv: 2101.11290 [hep-ex].

[151] Li. "MADMAX: A Dielectric Haloscope Experiment". In: *40th International Conference on High Energy Physics*. Vol. ICHEP2020. PoS. 2021, p. 645. DOI: 10.22323/1.390.0645.

[152] CAST Collaboration. "New CAST Limit on the Axion-Photon Interaction". In: *Nature Phys.* 13 (2017), pp. 584–590. arXiv: 1705.02290 [hep-ex].

[153] Armengaud, Avignone, Betz, Brax, Brun. "Conceptual Design of the International Axion Observatory (IAXO)". In: *Journal of Instrumentation* 9 (2014), T05002. DOI: 10.1088/1748-0221/9/05/T05002. arXiv: 1401.3233 [physics.ins-det].

[154] Porras-Bedmar. "Novel bounds on decaying axionlike particle dark matter from cosmic background observations". In: *Physical Review D* 110 (2024), p. 103501. DOI: 10.1103/PhysRevD.110.103501. arXiv: 2407.10618.

[155] Brouwer. "Projected Sensitivity of DMRadio-m^3: A Search for the QCD Axion Below $1\,\mu$eV". In: *Phys. Rev. D* 106.10 (2022), p. 103008. DOI: 10.1103/PhysRevD.106.103008. arXiv: 2204.13781 [hep-ex].

[156] Berlin, D'Agnolo, Ellis, Nantista, Neilson, Schuster, Tantawi, Toro, Zhou. "Axion Dark Matter Detection by Superconducting Resonant Frequency Conversion". In: *JHEP* 07 (2020), p. 088. DOI: 10.1007/JHEP07(2020)088. arXiv: 1912.11048 [hep-ph].

[157] Nordström. "Über die Möglichkeit, das elektromagnetische Feld und das Gravitationsfeld zu vereinigen". In: *Physikalische Zeitschrift* 15 (1914), pp. 504–506.

[158] Kaluza. "Zum Unitätsproblem der Physik". In: *Sitzungsberichte der Preussischen Akademie der Wissenschaften zu Berlin, Mathematisch-Physikalische Klasse* (1921), pp. 966–972.

[159] Klein. "Quantentheorie und fünfdimensionale Relativitätstheorie". In: *Zeitschrift für Physik* 37.12 (1926), pp. 895–906. DOI: 10.1007/BF01397481.

[160] Antoniadis. "A Possible New Dimension at a Few TeV". In: *Physics Letters B* 246 (1990), pp. 377–384. DOI: 10.1016/0370-2693(90)90617-F.

[161] Horava, Witten. "Heterotic and Type I String Dynamics from Eleven Dimensions". In: *Nuclear Physics B* 460 (1996), pp. 506–524. DOI: 10.1016/0550-3213(95)00621-4. arXiv: hep-th/9510209.

[162] Horava, Witten. "Eleven-Dimensional Supergravity on a Manifold with Boundary". In: *Nuclear Physics B* 475 (1996), pp. 94–114. DOI: 10.1016/0550-3213(96)00308-2. arXiv: hep-th/9603142.

[163] Polchinski. "Dirichlet-Branes and Ramond-Ramond Charges". In: *Physical Review Letters* 75 (1995), pp. 4724–4727. DOI: 10.1103/PhysRevLett.75.4724. arXiv: hep-th/9510017.

[164] Arkani-Hamed, Dimopoulos, Dvali. "The Hierarchy Problem and New Dimensions at a Millimeter". In: *Physics Letters B* 429 (1998), pp. 263–272. DOI: 10.1016/S0370-2693(98)00466-3. arXiv: hep-ph/9803315.

[165] Arkani-Hamed, Dimopoulos, Dvali. "Phenomenology, Astrophysics and Cosmology of Theories with Sub-Millimeter Dimensions and TeV Scale Quantum Gravity". In: *Physical Review D* 59 (1999), p. 086004. DOI: 10.1103/PhysRevD.59.086004. arXiv: hep-ph/9807344.

[166] Lee, Adelberger, Cook, Fleischer, Heckel. "New Test of the Gravitational $1/r^2$ Law at Separations down to $52\,\mu$m". In: *Physical Review Letters* 124 (2020), p. 101101. DOI: 10.1103/PhysRevLett.124.101101. arXiv: 2002.11761 [hep-ex].

[167] Ponton, Poppitz. "Casimir Energy and Radius Stabilization in Five and Six Dimensional Orbifolds". In: *Journal of High Energy Physics* 2001.06 (2001), p. 019. DOI: 10.1088/1126-6708/2001/06/019. arXiv: hep-ph/0105021 [hep-ph].

[168] Sundrum. "Effective Field Theory for a Three-Brane Universe". In: *Phys. Rev. D* 59 (1999), p. 085009. DOI: 10.1103/PhysRevD.59.085009. arXiv: hep-ph/9805471.

[169] Csaki. "TASI Lectures on Extra Dimensions and Branes". In: (Apr. 2004). DOI: 10.48550/arXiv.hep-ph/0404096. arXiv: hep-ph/0404096.

[170] Giudice, Rattazzi, Wells. "Quantum gravity and extra dimensions at high-energy colliders". In: *Nucl. Phys. B* 544 (1999), pp. 3–38. DOI: 10.1016/S0550-3213(99)00104-9. arXiv: hep-ph/9811291.

[171] Mirabelli, Perelstein, Peskin. "Collider signatures of new large space dimensions". In: *Phys. Rev. Lett.* 82 (1999), pp. 2236–2239. DOI: 10.1103/PhysRevLett.82.2236. arXiv: hep-ph/9811337.

[172] Han, Lykken, Zhang. "On Kaluza-Klein states from large extra dimensions". In: *Phys. Rev. D* 59 (1999), p. 105006. DOI: 10.1103/PhysRevD.59.105006. arXiv: hep-ph/9811350.

[173] Hewett. “Indirect collider signals for extra dimensions”. In: *Phys. Rev. Lett.* 82 (1999), pp. 4765–4768. DOI: 10.1103/PhysRevLett.82.4765. arXiv: hep-ph/9811356.

[174] Cullen, Perelstein. “SN1987A constraints on large compact dimensions”. In: *Phys. Rev. Lett.* 83 (1999), pp. 268–271. DOI: 10.1103/PhysRevLett.83.268. arXiv: hep-ph/9903422.

[175] Hanhart, Phillips, Reddy, Savage. “Extra dimensions, SN1987A, and nucleon-nucleon scattering data”. In: *Phys. Lett. B* 509 (2001), pp. 1–9. DOI: 10.1016/S0370-2693(01)00312-0. arXiv: nucl-th/0102063.

[176] Hall, Smith. “Cosmological constraints on theories with large extra dimensions”. In: *Phys. Rev. D* 60 (1999), p. 085008. DOI: 10.1103/PhysRevD.60.085008. arXiv: hep-ph/9904267.

[177] Dimopoulos, Landsberg. “Black holes at the LHC”. In: *Phys. Rev. Lett.* 87 (2001), p. 161602. DOI: 10.1103/PhysRevLett.87.161602. arXiv: hep-ph/0106295.

[178] Giddings, Thomas. “High-energy colliders as black hole factories: The end of short distance physics”. In: *Phys. Rev. D* 65 (2002), p. 056010. DOI: 10.1103/PhysRevD.65.056010. arXiv: hep-ph/0106219.

[179] Kapner, Cook, Adelberger, Gundlach, Heckel, Hoyle, Swanson. “Tests of the Gravitational Inverse-Square Law below the Dark-Energy Length Scale”. In: *Phys. Rev. Lett.* 98 (2007), p. 021101. DOI: 10.1103/PhysRevLett.98.021101. arXiv: hep-ph/0611184.

[180] Sirunyan. “Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 97.9 (2018), p. 092005. DOI: 10.1103/PhysRevD.97.092005. arXiv: 1712.02345 [hep-ex].

[181] Hannestad, Raffelt. “Supernova and neutron-star limits on large extra dimensions reexamined”. In: *Phys. Rev. D* 67 (2003), p. 125008. DOI: 10.1103/PhysRevD.67.125008. arXiv: hep-ph/0304029.

[182] Dienes, Dudas, Gherghetta. “Invisible Axions and Large-Radius Compactifications”. In: *Phys. Rev. D* 62 (2000), p. 105023. DOI: 10.1103/PhysRevD.62.105023. arXiv: hep-ph/9912455.

[183] Flacke, Gripaios, March-Russell, Maybury. “Warped Axions”. In: *JHEP* 01 (2007), p. 061. DOI: 10.1088/1126-6708/2007/01/061. arXiv: hep-ph/0611278.

[184] Arkani-Hamed, Cohen, Georgi. “Anomalies on Orbifolds”. In: *Phys. Lett. B* 516 (2001), pp. 395–402. DOI: 10.1016/S0370-2693(01)00946-7. arXiv: hep-th/0103135.

[185] Scrucca, Serone, Silvestrini, Zwirner. “Anomalies in Orbifold Field Theories”. In: *Phys. Lett. B* 525 (2002), pp. 169–174. DOI: 10.1016/S0370-2693(01)01430-7. arXiv: hep-th/0110073.

[186] Harvey. “TASI 2003 Lectures on Anomalies”. In: (2005). arXiv: hep-th/0509097.

[187] Scrucca, Serone. “Anomalies in Field Theories with Extra Dimensions”. In: *Int. J. Mod. Phys. A* 19 (2004), pp. 2579–2642. arXiv: hep-th/0403163.

[188] Chang, Tazawa, Yamaguchi. “Axion model in extra dimensions with TeV scale gravity”. In: *Phys. Rev. D* 61 (2000), p. 084005. DOI: 10.1103/PhysRevD.61.084005. arXiv: hep-ph/9908515 [hep-ph].

[189] Vinyoles, Serenelli, Villante, Basu, Bergström, Gonzalez-Garcia, Maltoni, Peña-Garay, Song. “A New Generation of Standard Solar Models”. In: *The Astrophysical Journal* 835.2 (2017), p. 202. DOI: 10.3847/1538-4357/835/2/202. arXiv: 1611.09867 [astro-ph.SR].

[190] Raffelt. *Stars as laboratories for fundamental physics : the astrophysics of neutrinos, axions, and other weakly interacting particles.* 1996.

[191] Bastero-Gil, Beaufort, Santos. “Solar axions in large extra dimensions”. In: *Journal of Cosmology and Astroparticle Physics* 2021.10 (Oct. 2021), p. 048. DOI: 10.1088/1475-7516/2021/10/048. arXiv: 2107.13337 [hep-ph].

[192] Cara Giovanetti, Robert Lasenby, Ken Van Tilburg. “Orbital Dynamics of the Solar Basin”. In: *J. High Energ. Phys.* 2024.7 (2024). DOI: 10.1007/JHEP12(2024)007. arXiv: 2408.16041 [hep-ph].

[193] NEWS-G Collaboration. “Solar Kaluza-Klein axion search with NEWS-G”. In: *Physical Review D* 105.1 (Jan. 2022), p. 012002. DOI: 10.1103/PhysRevD.105.012002.

[194] Hannestad, Raffelt. “Supernova and neutron-star limits on large extra dimensions reexamined”. In: *Physical Review D* 67.12 (2003), p. 125008. DOI: 10.1103/PhysRevD.67.125008. arXiv: hep-ph/0304029 [hep-ph].

[195] CMS Collaboration. "Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13\,\mathrm{TeV}$". In: *Physical Review D* 97.9 (2018), p. 092005. DOI: 10.1103/PhysRevD.97.092005. arXiv: 1712.02345 [hep-ex].

[196] CYGNO Collaboration. "A 50 liter CYGNO prototype overground characterization". In: *European Physical Journal C* 83.10 (2023), p. 946. DOI: 10.1140/epjc/s10052-023-11988-9. arXiv: 2305.06168 [physics.ins-det].

[197] Amaro. "The CYGNO experiment". In: *Proceedings of Science*. Vol. ICHEP2022. 2022, p. 1036. DOI: 10.22323/1.414.1036.

[198] Veenhof. "Garfield, a drift chamber simulation program". In: *International Conference on Programming and Mathematical Methods for Solving Physical Problems*. Vol. C9306149. 1993, pp. 66–71.

[199] CYGNO Collaboration. *Technical Design Report – TDR CYGNO-04/INITIUM*. Tech. rep. INFN, 2023. URL: https://agenda.infn.it/event/29666/contributions/160562/attachments/95430/131165/CSN2-CYGNO-PM-230.pdf.

[200] Amaro, Antonietti, Baracchini. "Bayesian network 3D event reconstruction in the CYGNO optical TPC for dark matter direct detection". In: *Eur. Phys. J. C* 85.11 (2025), p. 1261. DOI: 10.1140/epjc/s10052-025-14965-6. arXiv: 2506.04973 [physics.ins-det].

[201] Torelli. "Feasibility of a directional solar neutrino measurement with the CYGNO/INITIUM experiment". In: (2024). DOI: 10.48550/arXiv.2408.03760. arXiv: 2408.03760 [physics.ins-det].